



Artificial intelligence and machine learning

Niklas Kühl¹ · Max Schemmer¹ · Marc Goutier² · Gerhard Satzger¹

Received: 16 May 2022 / Accepted: 23 September 2022 / Published online: 9 November 2022
© The Author(s) 2022

Abstract

Within the last decade, the application of “artificial intelligence” and “machine learning” has become popular across multiple disciplines, especially in information systems. The two terms are still used inconsistently in academia and industry—sometimes as synonyms, sometimes with different meanings. With this work, we try to clarify the relationship between these concepts. We review the relevant literature and develop a conceptual framework to specify the role of machine learning in building (artificial) intelligent agents. Additionally, we propose a consistent typology for AI-based information systems. We contribute to a deeper understanding of the nature of both concepts and to more terminological clarity and guidance—as a starting point for interdisciplinary discussions and future research.

Keywords Machine learning · Artificial intelligence · Terminology · Framework

JEL classification C0

Introduction

Artificial Intelligence (AI) has been named as one of the most recent, fundamental developments of the convergence in electronic markets (Alt, 2021) and has become an increasingly relevant topic for information systems (IS) research (Abdel-Karim et al., 2021; Alt, 2018). While a large body of literature is concerned with designing AI to mimic and

replace humans (Dunin-Barkowski, 2020; Fukuda et al., 2001), IS research in general, and decision support systems (DSS) research in particular, emphasize the support of humans with AI (Arnott & Pervan, 2005). Recent research in hybrid intelligence (HI) and human-AI collaboration offers a promising path in synthesizing AI research across different fields (Dellermann, 2019): The ultimate goal of HI is to leverage the individual advantages of both human and artificial intelligence to enable synergy effects (James & Paul, 2018) and to achieve complementarity (Hemmer et al., 2021).

However, in many cases in both research and practice, AI is simply equated with the concept of machine learning (ML)—negatively impacting terminological precision and effective communication. Ågerfalk (2020, p.2) emphasizes that differentiating between AI and ML is especially important for IS research: “Is it not our responsibility as IS scholars to bring clarity to the discourse rather than contributing to its decline? (...) It would mean to distinguish between different types of AI and not talk of AI as synonymous with ML, which in itself is far from a monolithic concept.”

The practical relevance of a clear understanding is underlined by observing confusion and misuse of the terms AI and ML: During Mark Zuckerberg’s U.S. senate hearing in April 2018, he stressed that Facebook had “AI tools to

Responsible Editor: Ioanna Constantiou

✉ Niklas Kühl
niklas.kuehl@kit.edu

Max Schemmer
max.schemmer@kit.edu


Marc Goutier
goutier@ise.tu-darmstadt.de

Gerhard Satzger
gerhard.satzger@kit.edu

¹ Karlsruhe Institute of Technology (KIT),
Kaiserstr. 89, 76133 Karlsruhe, Germany

² Technical University of Darmstadt (TU Darmstadt),
Hochschulstr. 1, 64289 Darmstadt, Germany

Fig. 1 Appearance of the terms “artificial intelligence” and “machine learning” and in AIS Senior Scholars’ Basket journals



identify hate speech” as well as “terrorist propaganda” (The Washington Post, 2018). Researchers, however, would usually describe tasks identifying specific social media platform instances as classification tasks in the field of (supervised) ML (Waseem & Hovy, 2016). The increasing popularity of AI (Fujii & Managi, 2018) has led to the term often being used interchangeably with ML. This does not only hold true for the statement of Facebook’s CEO above, but also across various theoretical and application-oriented contributions in recent literature (Brink, 2017; ICO, 2017; Nawrocki et al., 2018). Camerer (2017) even mentions that he still uses AI as a synonym for ML *despite* knowing it is inaccurate.

As the remainder of this paper shows, both concepts are not identical—although in many cases both terms will appear in the same context. Such ambiguity might lead to multiple imprecisions in both research and practice when conversing about the relevant concepts, methods, and results. This is especially important in IS research—being interdisciplinary by nature (D’Atri et al., 2008). Ultimately, misuse can either lead to fundamental misunderstandings (Carnap, 1955) or to research that ought to be undertaken not being conducted (Davey & Cope, 2008; Lange, 2008). After all, misunderstandings can potentially lead to low perceived trustworthiness of AI (Thiebes et al., 2021).

It seems surprising that despite the frequent use of the terms, there is hardly any helpful academic delineation—apart from the notion that ML is a (not well-defined) subset of AI (Campeato, 2020), comparable to other possible sub-disciplines of AI: Expert systems, robotics, natural language processing, machine vision, and speech recognition (Collins et al., 2021; Dejoux & Léon, 2018). Consequently, this paper aims to shed light on the relationship between the two concepts: We analyze the role of ML in AI and, more precisely, in intelligent agents, which are defined by their capability to sense and act in an environment (Schleiffer, 2005). We do so by taking an ML perspective on intelligent agents’ capabilities and their relevant implementation—with IS research in mind. To this end, we review the relevant literature for both terms and synthesize and conceptualize the results.

Our article’s contributions are twofold: First, we identify different contributions of ML to intelligent agents as

specific AI instantiations. We base this on an expansion of the existing AI framework by Russell and Norvig (2020)—explicitly breaking down intelligent agents’ capabilities into separate “execution” and “learning” capabilities. Second, we develop a typology to provide a common terminology for AI-based information systems, where we conceptualize which systems employ ML—and which do not. The result should provide guidance when designing and analyzing systems.

Next, in Section “Terminology”, we review relevant literature in the fields of AI and ML. In Section “The role of rational agents in information systems”, we then analyze the capabilities of intelligent agents in more depth and examine the role of ML in them. Section “Towards a typology for machine learning in AI systems” develops a framework and typology to differentiate the terms AI and ML and to explain their relationship. In Section “Conclusion”, we conclude with a summary.

Terminology

Over the last decade, both terms, artificial intelligence (AI) and machine learning (ML), have enjoyed increasing popularity in information systems (IS) research. An analysis of the “AIS Senior Scholars’ Basket”¹ journals since 2000,² illustrates how the occurrences of both terms increased in titles, abstracts, and keywords (Fig. 1). While over the last 21 years, we observe a small but constant number of publications covering AI-related topics, ML only gained relevance in the literature after 2017: The late reflection of ML—despite of the earlier adoption and spread in industry (Brynjolfsson & McAfee, 2017)—may raise questions about whether IS has picked up the topic early enough.

¹ As of March 2022, see <https://aisnet.org/page/SeniorScholarBasket>, last accessed 16.05.2022.

² We start with the year 2000, as it was the last point in time when a journal (JAIS) was added to AIS Senior Scholars’ Basket.

Table 1 AI research streams; based on Russell and Norvig (2020).

Depth \ Objective	Human	Rational
Thinking	Cognitive modeling	Laws of thought
Acting	Turing Test	Rational agent <i>(Perspective of this work)</i>

As the analysis demonstrates, the two terms do exist for quite some time, while their related subjects are highly and increasingly topical now. In this section, we will elaborate on the meaning of the terms.

Artificial intelligence

In 1956, a Dartmouth workshop, led by Minsky and McCarthy, coined the term “artificial intelligence” (McCarthy et al., 1956) —later taking in contributions from a variety of different research disciplines, such as computer science (K. He et al., 2016) and programming (Newell & Simon, 1961), neuroscience (Ullman, 2019), robotics (Brady, 1984), linguistics (Clark et al., 2010), philosophy (Witten et al., 2011), and futurology (Koza et al., 1996). While the terminology is not well defined across disciplines, even within the IS domain definitions do vary widely; Collins et al. (2021) provide a comprehensive overview. Recent AI definitions transfer the human intelligence concept to machines in its entirety as “the ability of a machine to perform cognitive functions that we associate with human minds, such as perceiving, reasoning, learning, interacting with the environment, problem solving, decision-making, and even demonstrating creativity” (Rai et al., 2019, p.1). Still, over the last decades various debates have been raging on the *depth* and *objectives* of AI. These two dimensions span the space for different AI research streams in computer science and IS that were categorized by Russell and Norvig (2020): On the one hand (depth dimension), it may target either the thought process or a concrete action (*thinking* vs. *acting*); on the other hand (objective dimension), it may try to either replicate human decision making or to provide an ideal, “most rational” decision (*human-like* vs. *rational decision*). The resulting research streams are depicted in Table 1.

According to the *cognitive modeling* (i.e., thinking humanly) stream, AI instantiations must be “machines with a mind” (Haugeland, 1989) that perform human thinking (Bellman, 1978). Not only should they arrive at the same output as a human when given the same input, but also apply the same reasoning steps leading to this conclusion (Newell & Simon, 1961). The *laws of thought* stream (i.e., thinking

rationally) requires AI instantiations to arrive at a rational decision despite what a human might come up with. AI must therefore adhere to the laws of thought by using logic-based computational models (McDermott & Charniak, 1985). The *Turing test* stream (i.e., acting humanly) implies that AI must act intelligently when interacting with humans. To accomplish such tasks, AI instantiations must perform human tasks at least as well as humans (Rich & Knight, 1991), which can be tested via the Turing test (Turing, 1950). Finally, the *rational agent* stream considers AI as a rational (Russell & Norvig, 2020) or intelligent (Poole et al., 1998) agent.³ This agent does not only act autonomously, but also with the objective of achieving the rationally ideal outcome.

Machine learning

Many researchers perceive ML as an (exclusive) part of AI (Collins et al., 2021; Copeland, 2016; Ongsulee, 2017). In general, learning is a key facet of human cognition (Neisser, 1967). Humans process a vast amount of information by utilizing abstract knowledge that helps them to better understand incoming input. Owing to their adaptive nature, ML models can mimic a human being’s cognitive abilities (Janiesch et al., 2021): ML describes a set of methods commonly used to solve a variety of real-world problems with the help of computer systems, which can learn to solve a problem instead of being explicitly programmed to do so (Koza et al., 1996). For instance, instead of explicitly telling a computer system which words within an tweet would indicate it to contain a customer need, the system (given a sufficient set of training samples) learns the typical patterns of words and their combination which results in a need classification (Kühl et al., 2020).

In general, we differentiate between *unsupervised*, *supervised*, and *reinforcement* ML. *Unsupervised* ML comprises methods that reveal previously unknown patterns in data.

³ In this case, the terms *rational* and *intelligent* are used interchangeably in related work (Gama et al., 2014; Koza et al., 1996; Russell & Norvig, 2020).

Consequently, unsupervised learning tasks do not necessarily have a “correct” solution, as there is no ground truth (Wang et al., 2009).

Supervised ML refers to methods that allow the building of knowledge about a given task from a series of examples representing “past experience” (Mitchell, 1997). In the learning process, no manual adjustment or programming of rules or strategies to solve a problem is required, i.e., the model is capable to learn “by itself”. In more detail, supervised ML methods always aim to build a model by applying an algorithm to a set of known data points to gain insight into an unknown set of data (Hastie et al., 2017): Known data points are semantically labeled to create a target for the ML model. So-called semi-supervised learning combines elements from supervised and unsupervised ML by jointly using labeled and unlabeled data (Zhu, 2005).

Reinforcement learning refers to methods that are concerned with teaching intelligent agents to take those kinds of actions that increase their cumulative reward (Kaelbling et al., 1996). It differs from supervised learning in that no correctly matched features and targets are required for training. Instead, rewards and penalties allow the model to continuously learn over time. The focus is on a trade-off between the exploration of the uncharted environment and the exploitation of the existing knowledge base.

The role of rational agents in information systems

To further elaborate on the role of ML within AI, we need to take a clear perspective on the different definitions of AI to be beneficial to IS research. IS traditionally utilizes ML in predictive analytics tasks within (intelligent) decision support systems (DSS) (Arnott & Pervan, 2005; Müller et al., 2016) where the goal is to generate the best possible outcome (Arnott, 2006; Hunke et al., 2022; Power et al., 2019). As Phillips-Wren et al. (2019, p.63) emphasize, DSS “should help the decision-maker think rationally”. The perspective of rationality is also endorsed by other researchers in the field (Bakos & Treacy, 1986; Dellermann, 2019; Kloör et al., 2018; Power et al., 2019; Schuetz & Venkatesh, 2020). Thus, in the following we will explore the relationship between ML and AI in IS from the lens of the *rational agent* stream as discussed above. Furthermore, we will focus on supervised ML as it is the most common type of ML (Jordan & Mitchell, 2015). In the remainder of this section, we will first distinguish different types of (rational) agents and then use the insights to differentiate between the necessary layers when designing them as part of information systems.

Types of rational agents

According to the selected research stream, intelligence manifests itself in how rational agents act. Five features characterize agents in general: they “operate autonomously, perceive their environment, persist over a prolonged time period, adapt to change, and create and pursue goals” (Russell & Norvig, 2020, p.4). An agent defines its action, not for itself, but within the environment it operates and interacts with. It recognizes the environment through its sensors, relies on an agent program to handle and digest input data, and performs actions via actuators. A *rational* agent targets to achieve the highest expected outcome according to one or multiple objective performance measures—which are based on current and past knowledge of the environment and possible actions. For example, a rational agent within a medical diagnosis system aims to maximize the health of a patient measured via blood pressure, heart rate, and blood oxygen (potentially while minimizing the financial costs of a treatment as a secondary condition) (Grosu, 2022).

The agent’s conceptualization and surroundings are summarized in the *agent-environment framework*. It consists of three components: an agent, an environment, and a goal. Intelligence is the measurement of the “agent’s ability to achieve goals in a wide range of environments” (Legg & Hutter, 2007, p. 12). The agent obtains input through perceptions that the environment generates. Observations of the environment are one type of perception, while others are reward signals that indicate how well the agents’ goals have been achieved. Based on these input signals, the agent decides to perform actions, which are subsequently communicated back as signals to the environment.

Rational agents in information system architectures

As we investigate the role of ML in AI for IS research, we also need — apart from the theoretical and definitory aspects of agents — to consider how the functionality of a rational agent is reflected in an IS architecture. The implementation of agents is a key step to embed their functionality into practical, real-world (intelligent) information systems in general or into DSS specifically (Gao & Xu, 2009; Zhai et al., 2020). Any rational agent needs to be capable of at two least two tasks: cognition (Lieto et al., 2018) and (inter)action with the environment (Russell & Norvig, 2020). If we map these capabilities to system design terms, then *acting* capabilities are the ones built into a *frontend*, while the *cognitive* capabilities are embedded in a *backend*.

The frontend as the interface to the environment may take various forms; it may be designed as a very abstract, machine-readable web interface (Kühl et al., 2020), a human-readable application (Engel et al., 2022; Hirt et al., 2019), or

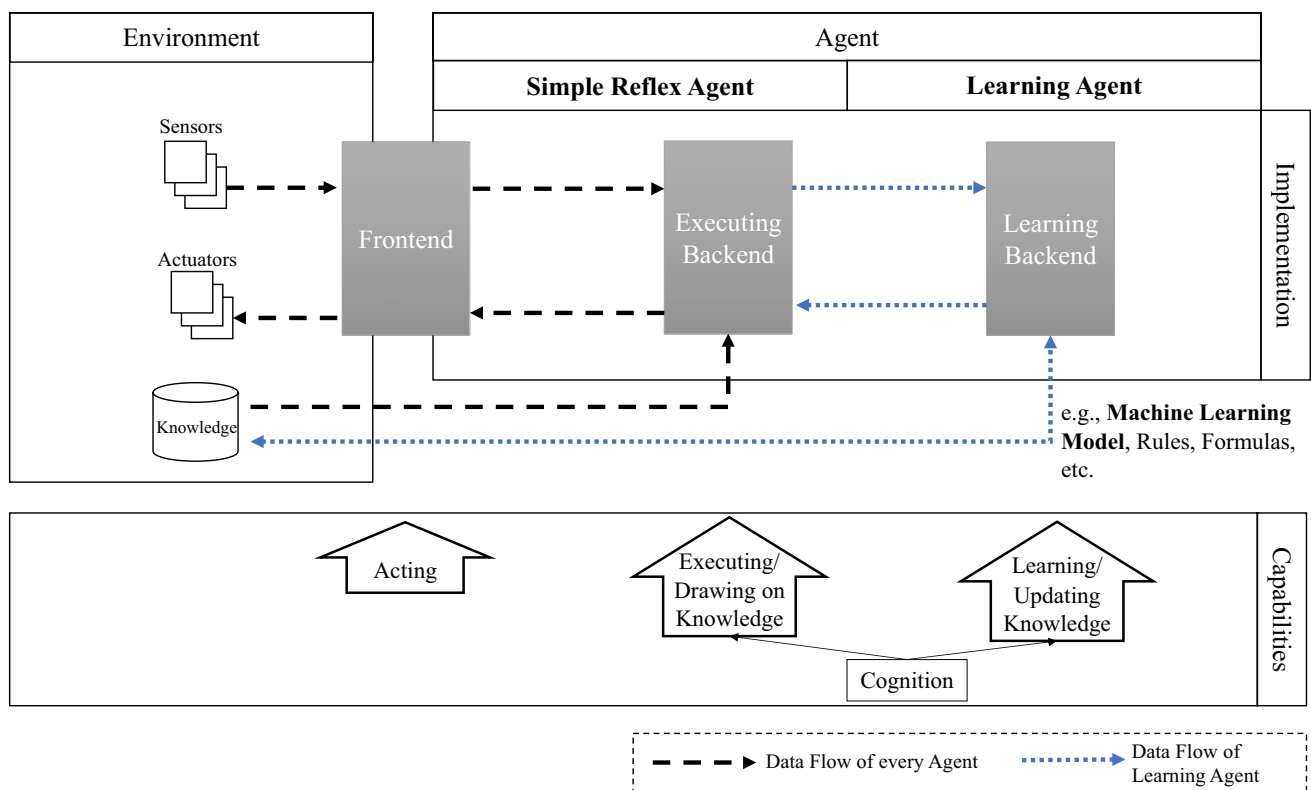


Fig. 2 Conceptual framework describing the general architecture for intelligent agents in AI-based information systems

even a humanoid template with elaborated expression capabilities (Guizzo, 2014). For the frontend to interact with the environment, two technical components are required: sensors and actuators. *Sensors* detect events or changes in the environment and forward the information via the frontend to the backend. They can, for instance, read the signals within an industrial process network (Hein et al., 2019), read visuals of an interaction with a human (Geller, 2014), but also perceive a keystroke input (Russell & Norvig, 2020). *Actuators*, on the other hand, are components responsible for moving, controlling, or displaying content. While sensors merely process information, actuators act, for instance, by automatically making bookings (Neuhofer et al., 2021) or changing a humanoid’s facial expressions (Berns & Hirth, 2006). One could argue that the Turing test (Turing, 1950) takes place at the environment’s interaction with the frontend, or, more precisely, when sensors and actuators are combined in a way to test the agent’s AI for *acting humanly*.

The backend provides the required functionalities to depict an intelligent agent’s *cognitive* capabilities. More precisely, this executing backend allows the agent to draw on its built-in knowledge. The backend translates signals from the frontend and transforms them into signals sent back to the frontend as

a response by executing actions. In some cases, there is an additional component modifying this response function over time, and thus modifying the execution part of the backend. We call this the learning part of the backend as depicted in Fig. 2. Within the next subsections, we will further elaborate this framework and its components.

The role of machine learning in rational agents

In terms of supervised ML, we need to further differentiate between the process task of *building* (training) adequate ML models (Witten et al., 2011) and the one of *executing* the deployed models (Chapman et al., 2000). To further understand ML’s role in intelligent agents, we partition the agent’s *cognition* layer into a *learning* sublayer (model building) and an *executing* sublayer (model execution).⁴ We, therefore, regard the implementation required by the learning sublayer as the *learning backend*, while the *executing backend* denotes the executing sublayer.

⁴ Russel and Norvig indicate a related relationship by differentiating between *learning elements* and *performance elements* (Russell & Norvig, 2020).

The learning backend first dictates *if* the intelligent agent is able to learn, and, second, *how* it does so—with respect to the algorithms it actually uses, the type of data processing it applies, and the handling of concept drift (Gama et al., 2014). Using the terminology of Russell and Norvig (2020), we distinguish two different types of intelligent agents: *simple-reflex agents* and *learning agents*. This differentiation holds explicitly in terms of a ML perspective on AI because it considers whether the underlying models in the *cognition* layer are trained just once and after that never touched (simple-reflex), or whether they are continuously updated to be adaptive (learning). Related work provides suitable examples of both. Kitts and Leblanc (2004) build a bidding agent for digital auctions as a simple-reflex agent: While building and testing the model for the agent may show convincing results, the system’s adaptive learning after deployment could be critical. Other examples of agents with models trained just once are common in different areas, for example, in terms pneumonia warning for hospitals (Oroszi & Ruhland, 2010), the (re)identification of pedestrians (Z. Zheng et al., 2017), and object annotation (Jorge et al., 2014). On the other hand, recent literature also provides examples of learning agents. Mitchell et al. (2015) present the concept of “never-ending learning” agents that strongly focus on continuously building and updating models in agents. Neuhofer et al. (2015) suggest an agent capable of personalization through a continuous learning processes of guest information for digital platforms, which an example of such an agent. Other examples include agents capable of making recommendations on music platforms (Liebman et al., 2015), regulating heat pump thermostats (Ruelens et al., 2015), acquiring collective knowledge across different tasks (Rostami et al., 2017), and learning the meanings of words (Yu et al., 2017). The choice of the learning type in agents (simple-reflex vs. learning agent) influences the agent’s general overall design and the contribution of ML.

As a result from the layers of agents and types of learning, our conceptual framework combining both is shown in Fig. 2. Regarding the previously mentioned ML methods, supervised ML *can* be the basis for either simple-reflex or learning agents, depending on whether the learning backend exists and on its feedback to the agent’s knowledge base. In terms of reinforcement learning, the agent, by definition, is a learning agent. However, there are also examples of where an agent functions *without* the utilization of ML—because the execution is based on rules (H. Wang et al., 2005), formulas (Billings et al., 2002) or other methods (Abasolo & Gomez, 2000). From this perspective, this means there can be AI without ML.

Towards a typology for machine learning in AI systems

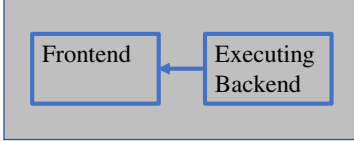
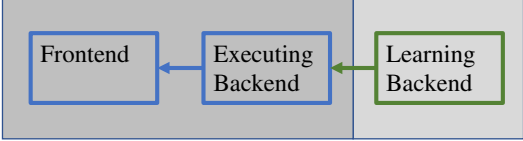
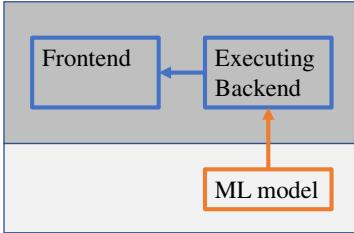
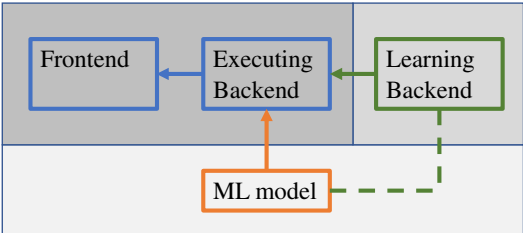
Based on the differentiation between simple-reflex and learning agents, we can now derive a typology for IS research. We refer to IS systems as *static AI-based systems* if they employ simple reflex agents that may be based on a model trained with ML. *Adaptive AI-based systems*, though, use learning agents, i.e., do have a learning backend—that may be based on ML, but alternatively also could be based, e.g., on rule-based knowledge representation. We, thus, propose the typology (as depicted in Table 2) for AI-based IS along the two dimensions: the existence of an ML-base for the executing backend and the existence of a learning backend.

We illustrate these findings in concrete IS research examples: *Static AI systems* are characterized by an executing backend which is based on algorithms not classified as ML and they lack a learning backend, i.e. they have a fixed response model (Chuang & Yadav, 1997). The executing backend of such systems is based on rules (like nested if-else statements), formulas (like mathematic equations describing a phenomena) or algorithms (like individual formal solution descriptions for specific problems). As an example for such systems, Hegazy et al. (2005) build a static AI system based on a self-developed algorithm and evaluate its performance within a cybersecurity context by simulating multiple attacks. Another example is provided by Ritchie (1990) who has developed an architecture and an instantiation of a static AI system for a traffic management platform.

In contrast, a *static ML-based AI system* has an executing backend which is based on ML. An example is provided in S. He et al. (2018). The authors develop an artifact to classify marketing on Twitter in either defensive or offensive marketing and show convincing prediction results. While their work did not aim at designing a productive artifact and is rather focused on showing the general feasibility of the approach, they choose a static ML-based AI system—which, however, might not be sufficient for permanent use: After the release of the article in 2018, Twitter changed its tweet size from 140 to 280 characters, thus changing the environment. It would be interesting to see how the developed model would need to adapt to this change. As another example, Samtani et al. (2017) build a model to identify harmful code snippets, typically utilized by hackers. They show how to design an artifact that can detect these *code assets* accurately for a proactive cyber threat intelligence. However, also in this case the environment and the assets of the hackers could and will change over time.

Adaptive AI systems, which are not based on ML, do comprise an executing backend with the flexibility to dynamically adapt the model to changing environments. This type

Table 2 Typology of AI-based information systems

Learning Backend Machine-Learning-based Executing Backend	No	Yes
No	Static AI system 	Adaptive AI system 
Yes	Static ML-based AI system 	Adaptive ML-based AI system 

of system is oftentimes enabled through the interaction between humans and AI systems. Most of the times, the system provides means and triggers for updates, while the human provides “manually encoded” knowledge updates. For example, Zhou et al. (2009) implement an adaptive AI system for pipeline leak detection which is based on a rule-based expert system and offers means to update the system online. In another example, Hatzilygeroudis and Prentzas (2004) develop an adaptive AI system to support the teaching process which has a specific component for knowledge updates. Both examples are inherently knowledge-based, but are explicitly designed to allow and force updates—although not on the basis of ML.

Finally, *adaptive ML-based AI-systems* implement learning in both sublayers of the cognition layer. For example, Q. Zheng et al. (2013) design a reinforcement-learning-based artifact to obtain information from hidden parts (“deep web”) of the internet. As their developed system perceives its current state and selects an action to submit to the environment (the deep web), the system continuously learns and builds up experience. In another example, Ghavamipoor and Hashemi Golpayegani (2020) build an adaptive ML-based AI system to predict the necessary service quality level and adapt an e-commerce system accordingly. As their system is continuously learning, their results show the total profits improve through effective cost reduction and revenue enhancement.

Conclusion

In this article, we clarify the relationship of machine learning (ML) in artificial intelligence (AI), particularly in intelligent agents, for the field of information systems research. Based on a rational agent view, we differentiate between AI agents capable of continuously improving as well as those who are static. Within these agents as instantiations of artificial intelligence, (supervised) ML can serve to support in different ways: either to contribute a once-trained model to define a static response pattern or to provide an adaptive model to realize dynamic behavior. As we point out, both could also be realized without the application of ML. Thus, “ML” and “AI” are not terms that should be used interchangeably—but as a conscious choice. Without question, ML is an important driver of AI, and the majority of modern AI cases will utilize ML. However, as we illustrate, there can be cases of AI without ML (e.g., based on rules or formulas).

This distinction enables our proposed framework to apply an intelligent agent’s perspective on AI-based information systems, enabling researchers to differentiate the existence and function of ML in them. Interestingly, as of today, many AI-based information systems remain static, i.e. employ once-trained ML models (Kühl et al., 2021). With increasing focus on deployment and life cycle management, we will see more adaptive AI-based systems that sense changes in the

environment and use ML to learn continuously (Baier et al., 2019). Our framework and the resulting typology should allow IS researchers and practitioners to be more precise when referring to ML and AI, as it highlights the importance of not using the terms interchangeably but clarifying the role ML plays in AI's system design.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abasolo, J. M., & Gomez, M. (2000). MELISA: An ontology-based agent for information retrieval in medicine. *Proceedings of the 1st international workshop on the semantic web (SemWeb2000)*, 73–82.
- Abdel-Karim, B. M., Pfeuffer, N., & Hinz, O. (2021). Machine learning in information systems - a bibliographic review and open research issues. *Electronic Markets*, 31(3), 643–670. <https://doi.org/10.1007/s12525-021-00459-2>
- Ågerfalk, P. J. (2020). Artificial intelligence as digital agency. *European Journal of Information Systems*, 29(1), 1–8. <https://doi.org/10.1080/0960085X.2020.1721947>
- Alt, R. (2018). Electronic markets and current general research. *Electronic Markets*, 28(2), 123–128. <https://doi.org/10.1007/s12525-018-0299-0>
- Alt, R. (2021). Electronic markets on the next convergence. *Electronic Markets*, 31(1), 1–9. <https://doi.org/10.1007/s12525-021-00471-6>
- Arnott, D. (2006). Cognitive biases and decision support systems development: a design science approach. *Information Systems Journal*, 16(1), 55–78. <https://doi.org/10.1111/j.1365-2575.2006.00208.x>
- Arnott, D., & Pervan, G. (2005). A critical analysis of decision support systems research. *Journal of Information Technology*, 20(2), 67–87. <https://doi.org/10.1057/palgrave.jit.2000035>
- Baier, L., Kühl, N., & Satzger, G. (2019). How to cope with change? Preserving validity of predictive services over time. *Hawaii International Conference on System Sciences (HICSS-52)*. <https://doi.org/10.5445/IR/1000085769>
- Bakos, J. Y., & Treacy, M. E. (1986). Information technology and corporate strategy: a research perspective. *MIS Quarterly*, 107–119. <https://doi.org/10.2307/249029>
- Bellman, R. (1978). In Boyd & Fraser. (Ed.), *An introduction to artificial intelligence: Can computers think?*
- Berns, K., & Hirth, J. (2006). Control of facial expressions of the humanoid robot head ROMAN. *IEEE International Conference on Intelligent Robots and Systems*, 3119–3124. <https://doi.org/10.1109/IROS.2006.282331>
- Billings, D., Davidson, A., Schaeffer, J., & Szafron, D. (2002). The challenge of poker. *Artificial Intelligence*, 134(1–2), 201–240. [https://doi.org/10.1016/S0004-3702\(01\)00130-8](https://doi.org/10.1016/S0004-3702(01)00130-8)
- Brady, M. (1984). Robotics and artificial intelligence. In M. Brady, L. A. Gerhardt, & H. F. Davidson (Eds.), *Artificial intelligence* (Vol. 26, Issue 1). Springer. <https://doi.org/10.1007/978-3-642-82153-0>
- Brink, J. A. (2017). Big data management, access, and protection. *Journal of the American College of Radiology*, 14(5), 579–580. <https://doi.org/10.1016/j.jacr.2017.03.024>
- Brynjolfsson, E., & McAfee, A. (2017). The business of artificial intelligence. *Harvard Business Review*, 1–20.
- Camerer, C. F. (2017). Artificial intelligence and behavioral economics. In *Economics of Artificial Intelligence*. University of Chicago Press.
- Campeato, O. (2020). *Artificial intelligence, machine learning, and deep learning*. Mercury Learning & Information.
- Carnap, R. (1955). Meaning and synonymy in natural languages. *Philosophical Studies*, 6(3), 33–47. <https://doi.org/10.1007/BF02330951>
- Chapman, P., Clinton, J., Kerber, R., Khabaza, T., Reinartz, T., Shearer, C., & Wirth, R. (2000). CRISP-DM 1.0. *CRISP-DM Consortium*, 76. <https://doi.org/10.1109/ICETET.2008.239>
- Chuang, T.-T., & Yadav, S. B. (1997). An agent-based architecture of an adaptive decision support system. *Americas Conference on Information Systems, Indianapolis, IN*.
- Clark, A., Fox, C., & Lappin, S. (2010). *The handbook of computational linguistics and natural language processing* (a. Clark, C. Fox, & S. Lappin (eds.)). Wiley-Blackwell. <https://doi.org/10.1002/9781444324044>
- Collins, C., Dennehy, D., Conboy, K., & Mikalef, P. (2021). Artificial intelligence in information systems research: A systematic literature review and research agenda. *International Journal of Information Management*, 60, 102383. <https://doi.org/10.1016/j.ijinfomgt.2021.102383>
- Copeland, M. (2016). What's the difference between artificial intelligence. *Machine learning, and deep learning*, 29. <https://blogs.nvidia.com/blog/2016/07/29/whats-difference-artificial-intelligence-machine-learning-deep-learning-ai/>. Accessed 3 May 2022.
- D'Atri, A., Marco, M., & Casalino, N. (2008). *Interdisciplinary aspects of information systems studies. The Italian association for information systems*. Physica Heidelberg. <https://link.springer.com/book/10.1007/978-3-7908-2010-2>
- Davey, B., & Cope, C. (2008). Requirements elicitation – What's missing? *Issues in Informing Science and Information Technology*, 5, 543–551. <https://doi.org/10.28945/I027>
- Dejoux, C., & Léon, E. (2018). *Métamorphose des managers à l'ère du numérique et de l'intelligence artificielle*. Pearson.
- Dellermann, D., Lipusch, N., Ebel, P., & Leimeister, J. M. (2019). Design principles for a hybrid intelligence decision support system for business model validation. *Electronic Markets*, 29(3), 423–441. <https://doi.org/10.1007/s12525-018-0309-2>
- Dunin-Barkowski, W. (2020). Editorial: Toward and beyond human-level AI. *Frontiers in Neurobotics*, 14. <https://doi.org/10.3389/fnbot.2020.617446>
- Engel, C., Ebel, P., & Leimeister, J. M. (2022). Cognitive automation. *Electronic Markets*, 32(1), 339–350. <https://doi.org/10.1007/s12525-021-00519-7>
- Fujii, H., & Managi, S. (2018). Trends and priority shifts in artificial intelligence technology invention: A global patent analysis. *Economic Analysis and Policy*, 58, 60–69. <https://doi.org/10.1016/j.eap.2017.12.006>
- Fukuda, T., Michelini, R., Potkonjak, V., Tzafestas, S., Valavanis, K., & Vukobratovic, M. (2001). How far away is "artificial man." *IEEE Robotics & Automation Magazine*, 8(1), 66–73. <https://doi.org/10.1109/100.924367>

- Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4), 1–37. <https://doi.org/10.1145/2523813>
- Gao, S., & Xu, D. (2009). Conceptual modeling and development of an intelligent agent-assisted decision support system for anti-money laundering. *Expert Systems with Applications*, 36(2), 1493–1504. <https://doi.org/10.1016/j.eswa.2007.11.059>
- Geller, T. (2014). How do you feel? Your computer knows. *Communications of the ACM*, 6(8), 24–26. [https://doi.org/10.1016/S1364-6613\(02\)01946-0](https://doi.org/10.1016/S1364-6613(02)01946-0)
- Ghavamiipoor, H., & Hashemi Golpayegani, S. A. (2020). A reinforcement learning based model for adaptive service quality management in E-commerce websites. *Business & Information Systems Engineering*, 62(2), 159–177. <https://doi.org/10.1007/s12599-019-00583-6>
- Grosu, R. (2022). Can artificial intelligence improve our health? In *Strategies for sustainability of the earth system* (pp. 273–281). Springer. https://doi.org/10.1007/978-3-030-74458-8_17
- Guizzo, E. (2014). *How Aldebaran robotics built its friendly humanoid robot, pepper*. IEEE Spectrum. <https://www.spectrum.ieee.org/how-aldebaran-robotics-built-its-friendly-humanoid-robot-pepper>
- Hastie, T., Tibshirani, R., & Friedman, J. (2017). *The elements of statistical learning: Data mining, inference and prediction* (Vol. 9). Springer.
- Hatzilygeroudis, I., & Prentzas, J. (2004). Using a hybrid rule-based approach in developing an intelligent tutoring system with knowledge acquisition and update capabilities. *Expert Systems with Applications*, 26(4), 477–492. <https://doi.org/10.1016/j.eswa.2003.10.007>
- Haugeland, J. (1989). *Artificial intelligence: The very idea*. MIT Press.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. <https://doi.org/10.1109/CVPR.2016.90>
- He, S., Rui, H., & Whinston, A. B. (2018). Social media strategies in product-harm crises. *Information Systems Research*, 29(2), 362–380. <https://doi.org/10.1287/isre.2017.0707>
- Hegazy, I. M., Faheem, H. M., Al-Arif, T., & Ahmed, T. (2005). Performance evaluation of agent-based IDS. *Proceedings of the 2nd international conference on intelligent computing and information systems (ICICIS 2005)* (pp. 314–319).
- Hein, A., Weking, J., Schrieck, M., Wiesche, M., Böhm, M., & Krcmar, H. (2019). Value co-creation practices in business-to-business platform ecosystems. *Electronic Markets*, 29(3), 503–518. <https://doi.org/10.1007/s12525-019-00337-y>
- Hemmer, P., Schemmer, M., Vössing, M., & Kühl, N. (2021). Human-AI complementarity in hybrid intelligence systems: A structured literature review. *PACIS 2021 Proceedings*.
- Hirt, R., Kühl, N., & Satzger, G. (2019). Cognitive computing for customer profiling: meta classification for gender prediction. *Electronic Markets*, 29(1), 93–106. <https://doi.org/10.1007/s12525-019-00336-z>
- Hunke, F., Heinz, D., & Satzger, G. (2022). Creating customer value from data: Foundations and archetypes of analytics-based services. *Electronic Markets*, 32(2), 1–19. <https://doi.org/10.1007/s12525-021-00506-y>
- ICO. (2017). *Big data, artificial intelligence, machine learning and data protection*. <https://www.ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf>
- James, H., & Paul, R. (2018). *Collaborative intelligence: Humans and AI are joining forces* (pp. 114–123). Harvard Business Review.
- Janiesch, C., Zschech, P., & Heinrich, K. (2021). Machine learning and deep learning. *Electronic Markets*, 31(3), 685–695. <https://doi.org/10.1007/s12525-021-00475-2>
- Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255–260. <https://doi.org/10.1126/science.aaa8415>
- Jorge, A. M., Leal, J. P., Anand, S. S., & Dias, H. (2014). A study of machine learning methods for detecting user interest during web sessions. *Proceedings of the 18th International Database Engineering & Applications Symposium on - IDEAS '14*, 149–157. <https://doi.org/10.1145/2628194.2628239>
- Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1996). Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4, 237–285. <https://doi.org/10.1613/jair.301>
- Kitts, B., & Leblanc, B. (2004). Optimal bidding on keyword auctions. *Electronic Markets*, 14(3), 186–201. <https://doi.org/10.1080/1019678042000245119>
- Kloör, B., Monhof, M., Beverungen, D., & Braäer, S. (2018). Design and evaluation of a model-driven decision support system for repurposing electric vehicle batteries. *European Journal of Information Systems*, 27(2), 171–188. <https://doi.org/10.1057/s41303-017-0044-3>
- Koza, J. R., Bennett, F. H., Andre, D., & Keane, M. A. (1996). Automated design of both the topology and sizing of analog electrical circuits using genetic programming. In J. S. Gero, & F. Sudweeks (Eds.), *Artificial Intelligence in Design '96*. Springer. https://doi.org/10.1007/978-94-009-0279-4_9
- Kühl, N., Hirt, R., Baier, L., Schmitz, B., & Satzger, G. (2021). How to conduct rigorous supervised machine learning in information systems research: The supervised machine learning report card. *Communications of the Association for Information Systems*, 48(1), 589–615. <https://doi.org/10.17705/1CAIS.04845>
- Kühl, N., Mühlthaler, M., & Goutier, M. (2020). Supporting customer-oriented marketing with artificial intelligence: Automatically quantifying customer needs from social media. *Electronic Markets*, 30(2), 351–367. <https://doi.org/10.1007/s12525-019-00351-0>
- Lange, P. G. (2008). Terminological obfuscation in online research. In *Handbook of Research on Computer Mediated Communication* (pp. 436–450). IGI Global. <https://doi.org/10.4018/978-1-59904-863-5.ch033>
- Legg, S., & Hutter, M. (2007). Universal intelligence: A definition of machine intelligence. *Minds and Machines*, 17(4), 391–444. <https://doi.org/10.1007/s11023-007-9079-x>
- Liebman, E., Saar-Tsechansky, M., & Stone, P. (2015). Dj-mc: A reinforcement-learning agent for music playlist recommendation. *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, 591–599.
- Lieto, A., Bhatt, M., Oltramari, A., & Vernon, D. (2018). The role of cognitive architectures in general artificial intelligence. *Cognitive Systems Research*, 48, 1–3. <https://doi.org/10.1016/j.cogsys.2017.08.003>
- McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (1956). A proposal for the Dartmouth summer research project on artificial intelligence. *Dartmouth Conference*. <https://doi.org/10.1609/aimag.v27i4.1904>
- McDermott, D., & Charniak, E. (1985). Introduction to artificial intelligence. *International Journal of Adaptive Control and Signal Processing*, 2(2), 148–149.
- Mitchell, T. M. (1997). Machine learning. *Annual Review Of Computer Science* (Issue 1). <https://doi.org/10.1145/242224.242229>
- Mitchell, T. M., Cohen, W., Hruschka, E., Talukdar, P., Betteridge, J., Carlson, A., Mishra, B. D., Gardner, M., Kisiel, B., Krishnamurthy, J., Lao, N., Mazaitis, K., Mohamed, T., Nakashole, N., Platanios, E. A., Ritter, A., Samadi, M., Settles, B., Wang, R., Wijaya, D., Gupta, A., Chen, X., Saparov, A., Greaves, M., & Welling, J. (2015). Never-ending learning. *AAAI Conference on Artificial Intelligence*, 2302–2310.

- Müller, O., Junglas, I., Brocke, J. V., & Debortoli, S. (2016). Utilizing big data analytics for information systems research: Challenges, promises and guidelines. *European Journal of Information Systems*, 25(4), 289–302. <https://doi.org/10.1057/ejis.2016.2>
- Nawrocki, T., Maldjian, P. D., Slasky, S. E., & Contractor, S. G. (2018). Artificial intelligence and radiology: Have rumors of the radiologist's demise been greatly exaggerated? *Academic Radiology*. <https://doi.org/10.1016/j.acra.2017.12.027>
- Neisser, U. (1967). Cognitive psychology. *Thinkingjudgement and Decision Makin*. <https://doi.org/10.1126/science.198.4319.816>
- Neuhofer, B., Buhalis, D., & Ladkin, A. (2015). Smart technologies for personalized experiences: a case study in the hospitality domain. *Electronic Markets*, 25(3), 243–254. <https://doi.org/10.1007/s12525-015-0182-1>
- Neuhofer, B., Magnus, B., & Celuch, K. (2021). The impact of artificial intelligence on event experiences: A scenario technique approach. *Electronic Markets*, 31(3), 601–617. <https://doi.org/10.1007/s12525-020-00433-4>
- Newell, A., & Simon, H. A. (1961). *GPS, a program that simulates human thought*. (Report of the Defense Technical Information Center). <https://www.apps.dtic.mil/sti/citations/AD0294731>
- Ongsulee, P. (2017). Artificial intelligence, machine learning and deep learning. *2017 15th International Conference on ICT and Knowledge Engineering (ICT&KE)*, 1–6. <https://doi.org/10.1109/ICTKE.2017.8259629>
- Oroszi, F., & Ruhland, J. (2010). An early warning system for hospital acquired. *18th European Conference on Information Systems (ECIS)*. <https://www.aisel.aisnet.org/ecis2010/93>
- Phillips-Wren, G., Power, D. J., & Mora, M. (2019). *Cognitive bias, decision styles, and risk attitudes in decision making and DSS*. Taylor & Francis. <https://doi.org/10.1080/12460125.2019.1646509>
- Poole, D. L., Mackworth, A., & Goebel, R. G. (1998). Computational intelligence and knowledge. *Computational Intelligence: A Logical Approach*, Ci, 1–22.
- Power, D. J., Cyphert, D., & Roth, R. M. (2019). Analytics, bias, and evidence: The quest for rational decision making. *Journal of Decision Systems*, 28(2), 120–137. <https://doi.org/10.1080/12460125.2019.1623534>
- Rai, A., Constantinides, P., & Sarker, S. (2019). Next generation digital platforms: Toward human-AI hybrids. *MIS Quarterly*, 43(1), iii–ix.
- Rich, E., & Knight, K. (1991). *Artificial intelligence*. McGraw-Hill.
- Ritchie, S. G. (1990). A knowledge-based decision support architecture for advanced traffic management. *Transportation Research Part A: General*, 24(1), 27–37. [https://doi.org/10.1016/0191-2607\(90\)90068-H](https://doi.org/10.1016/0191-2607(90)90068-H)
- Rostami, M., Kolouri, S., Kim, K., & Eaton, E. (2017). Multi-agent distributed lifelong learning for collective knowledge acquisition. *ArXiv preprint ArXiv:1709.05412*. <https://doi.org/10.48550/arXiv.1709.05412>
- Ruelens, F., Iacovella, S., Claessens, B. J., & Belmans, R. (2015). Learning agent for a heat-pump thermostat with a set-back strategy using model-free reinforcement learning. *Energies*, 8(8), 8300–8318. <https://doi.org/10.3390/en8088300>
- Russell, S. J., & Norvig, P. (2020). Artificial intelligence: A modern approach. In *Artificial Intelligence* (3rd ed.). <https://doi.org/10.1017/S0269888900007724>
- Samtani, S., Chinn, R., Chen, H., & Nunamaker Jr., J. F. (2017). Exploring emerging hacker assets and key hackers for proactive cyber threat intelligence. *Journal of Management Information Systems*, 34(4), 1023–1053. <https://doi.org/10.1080/07421222.2017.1394049>
- Schleiffer, R. (2005). An intelligent agent model. *European Journal of Operational Research*, 166(3), 666–693. <https://doi.org/10.1016/j.ejor.2004.03.039>
- Schuetz, S., & Venkatesh, V. (2020). Research perspectives: The rise of human machines: How cognitive computing systems challenge assumptions of user-system interaction. *Journal of the Association for Information Systems*, 21(2), 460–482. <https://doi.org/10.17705/1jais.00608>
- The Washington Post. (2018, April 10). *Transcript of Mark Zuckerberg's senate hearing*. <https://www.washingtonpost.com/news/the-switch/wp/2018/04/10/transcript-of-mark-zuckerbergs-senate-hearing/>
- Thiebes, S., Lins, S., & Sunyaev, A. (2021). Trustworthy artificial intelligence. *Electronic Markets*, 31(2), 1–18. <https://doi.org/10.1007/s12525-020-00441-4>
- Turing, A. M. (1950). Computing machine and intelligence. *MIND*, LIX(236), 433–460. <https://doi.org/10.1093/2Fmind/2FLIX.236.433>
- Ullman, S. (2019). Using neuroscience to develop artificial intelligence. *Science*, 363(6428), 692–693. <https://doi.org/10.1126/science.aau6595>
- Wang, H., Kwong, S., Jin, Y., Wei, W., & Man, K.-F. (2005). Agent-based evolutionary approach for interpretable rule-based knowledge extraction. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 35(2), 143–155. <https://doi.org/10.1109/TSMCC.2004.841910>
- Wang, K., Wang, B., & Peng, L. (2009). CVAP: Validation for cluster analyses. *Data Science Journal*, 904220071. <https://doi.org/10.2481/dsj.007-020>
- Waseem, Z., & Hovy, D. (2016). Hateful symbols or hateful people? Predictive features for hate speech detection on twitter. *Proceedings of the NAACL Student Research Workshop*, 88–93. <https://doi.org/10.18653/v1/N16-2013>
- Witten, I. H., Frank, E., & Hall, M. A. (2011). Data mining: Practical machine learning tools and techniques, third edition. In *Annals of physics* (Vol. 54, Issue 2). [https://doi.org/10.1002/1521-3773\(20010316\)40:6<9823::AID-ANIE9823>3.3.CO;2-C](https://doi.org/10.1002/1521-3773(20010316)40:6<9823::AID-ANIE9823>3.3.CO;2-C)
- Yu, Y., Eshghi, A., & Lemon, O. (2017). VOILA : An optimised dialogue system for interactively learning visually-grounded word meanings (demonstration system). *Proceedings of the SIGDIAL 2017 Conference*, 197–200.
- Zhai, Z., Martínez, J. F., Beltran, V., & Martínez, N. L. (2020). Decision support systems for agriculture 4.0: Survey and challenges. *Computers and Electronics in Agriculture*, 170, 105256. <https://doi.org/10.1016/j.compag.2020.105256>
- Zheng, Q., Wu, Z., Cheng, X., Jiang, L., & Liu, J. (2013). Learning to crawl deep web. *Information Systems*, 38(6), 801–819. <https://doi.org/10.1016/j.is.2013.02.001>
- Zheng, Z., Zheng, L., & Yang, Y. (2017). Pedestrian alignment network for large-scale person re-identification. *ArXiv Preprint ArXiv:1707.00408*.
- Zhou, Z.-J., Hu, C.-H., Yang, J.-B., Xu, D.-L., & Zhou, D.-H. (2009). Online updating belief rule based system for pipeline leak detection under expert intervention. *Expert Systems with Applications*, 36(4), 7700–7709. <https://doi.org/10.1016/j.eswa.2008.09.032>
- Zhu, X. J. (2005). *Semi-supervised learning literature survey*. University of Wisconsin-Madison, Department of Computer Sciences. <https://www.digital.library.wisc.edu/1793/60444>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.