https://doi.org/10.48009/4_iis_2020_135-142

Issues in Information Systems
Volume 21, Issue 4, pp. 135-142, 2020

AI, IoT, AND AIoT: DEFINITIONS AND IMPACTS ON THE ARTIFICIAL INTELLIGENCE CURRICULUM

John C. Stewart, Robert Morris University, stewartj@rmu.edu
Gary Alan Davis, Robert Morris University, davis@rmu.edu
Diane A. Igoche, Robert Morris University, igoche@rmu.edu

ABSTRACT

The term Artificial Intelligence (AI) has been used frequently in the media and in advertisements in recent years. Yet, few can offer a clear and concise definition of the term "AI." This lack of a definition complicates the communication and understanding among entities that focus on the collaboration, use, and application of AI. Further complicating the landscape is the fact that advances in technology have altered the perception of AI, and consequently, the ways in which they are defined. The goal of this paper is to derive a definition that capture the current collective scientific, rather than the societal or media-based definition of the discipline (and its subcategories), map this definition to emerging areas in AI such as AIoT, and determine how well this mapping translates to an applied curriculum at the highest ranked schools. Specifically, are we adequately defining AI within the confines of how it is currently applied, and how we aspire to apply it? Are college degree programs preparing students to make contributions to this field, in terms of the application and advancement of AI?

Keywords: IoT, AIoT, Artificial Intelligence, Machine Learning, AI Curriculum

INTRODUCTION

Confusion about what Artificial Intelligence (AI) actually is, or what it encompasses, is inherent in the collective inability to define it and determine a universally-agreeable definition. The field lacks a clearly delineated definition (Brahman, 2006; Nilsson, 2009; Bhatnagar et al., 2018; Monett and Lewis, 2018; Wang, 2019). From a purely research perspective, this might not be viewed as a particularly troublesome issue. Research will not wait in the absence of a universally accepted definition. However, problems arise when other viewpoints are considered. In AI project development, what criteria are used to evaluate systems without a framework to do so? Differing interpretations of AI leads to problems in the design and development of AI projects (Wang, 2019). Lastly, if we cannot adequately define AI, how can we determine what to teach in higher education AI curricula, and assess whether we are teaching the applicable things to produce professionals who will contribute to the field? In short, if we cannot define AI, then we cannot teach it.

DEFINING ARTIFICIAL INTELLIGENCE (AI)

Some of the difficulty in defining AI arises from the general perception that AI means robots and self-driving cars. While these are applications of AI, what are the specific components of AI that allow these types of applications? Making it even more difficult to define AI are the rapid technological improvements which render previous definitions obsolete. Algorithms that perform basic calculations are no longer perceived as AI, because they are now commonplace as functions of a computer system.

In the simplest (and broadest) terms, AI can be defined as the simulation or replication of human intelligence in machines (Franenfield, 2020). This definition makes the assumption that we are defining artificial intelligence based on what we determine to be human intelligence. The goal of AI is to capture the human characteristics of learning, reasoning, and perceiving. Are machines capable of learning and problem-solving at the level of humans? Have we reached that level of machine intelligence?

Machine Intelligence: The Internet of Things (IoT)

While it is not yet determined if machines have reached human problem-solving ability, machines have certainly extended the reach of human senses. Machines like cars, appliances, and intelligent personal assistants (e.g., Amazon

Echo) extend the reach of humans far beyond our physical capabilities. These commonplace machines make up what is known as the Internet of Things, or IoT. Simply stated, IoT is " . . . everyday objects that can be connected to the Internet and be recognized by other devices and contribute information to a database" (Marr, 2017, p. 1).

IoT has impacted our daily lives in almost every way imaginable. For example, intelligent personal assistants (e.g., Amazon Echo and Google Home) have made our homes smarter and safer by integrating with climate controls, entertainment systems, and security systems. Wearable technology, such as Apple Watch and FitBit, help us to stay healthy and connected while we are active and away from our hand-held devices. Finally, smart cars help us to navigate, schedule routine maintenance, and diagnose mechanical problems. As smart car technology evolves, our cars will inevitably become self-driving.

While IoT provides many benefits to consumers, it also collects and forwards massive amounts of data to companies. This endless flow of "streaming data" can help companies manage inventories, identify consumer demands, shorten purchasing cycles, and (most importantly) "learn" from the repositories of collected data (Marr, 2017). However, as prevalent and as useful as IoT is, is it artificial "intelligence?" Are these devices truly "thinking and reasoning" machines? Clearly, a specific definition of "intelligence" is needed.

Defining Intelligence

The difficulty in clearly defining artificial intelligence lies in defining "intelligence." As humans, we have yet to agree on what intelligence is. If we can distinguish that "artificial" means "non-human," then the difficulty remains to define "intelligence." If we are unable to accurately and distinctly define human intelligence, then basing the definition of AI on human intelligence becomes, by default, problematic.

If, however, we attempt to view AI in terms of its application and capability to solve problems, we overcome the requirement of having to include the expectation that it must mimic an intelligence that we cannot clearly define. So, if we remove the requirement that AI mimic or adhere to the constraints of what we perceive as human intelligence, this removal may allow for an expansion of, not only the definition, but the capability of AI. Stuart Russell and Peter Norvig point out in their definitive text, *Artificial Intelligence: A Modern Approach*, that "artificial flight" became possible when the Wright brothers and other aviation pioneers stopped imitating birds and began studying aerodynamics and using wind tunnels (Russell and Norvig, 2015).

Wang (1995) defined intelligence as "the capacity of an information-processing system to adapt to its environment, while operating with insufficient knowledge and resources" (p. 13). This approach eliminates the need to make the definition of artificial intelligence based upon its ability to mimic human intelligence and transfers the definition to the capability of an information system to comprehend inputs and reason solutions to problems. More precisely, the focus of the system is performing tasks and solving problems (Wang, 2019).

ARTIFICIAL INTELLIGENCE (AI) APPLICATIONS

AI applications have been divided into *weak* and *strong* intelligence (Rouse, 2020). A *weak* AI system is designed to do one specific thing. For example, an industrial robot that performs a repetitive task, or an intelligent agent that does one focused and specific task (e.g., shopping bot) are examples of weak AI.

A *strong* AI system will deal with much more complex interactions or situations. The complexities and complications in the decision-making with self-driving cars would be an example of strong AI. Strong AI would be the gathering and processing of massive amounts of external data on a constant basis to follow a specified route, and keep from colliding with a myriad of objects. A strong AI application has the ability to autonomously apply knowledge to new problems that are presented to it.

The Artificial Intelligence of Things (AIoT)

Are IoT and AI related? If so, how and where do they converge? The "strong" AI systems discussed previously are rapidly emerging into their own category: *Artificial Intelligence of Things*, or AIoT. If AI can be described as the intelligent "brain" of a system, then IoT should be thought of as the "digital nervous system" (Marr, 2019, p. 1). As in the human body, both systems must work in concert with one another. In information systems, the combination of these two systems creates powerful synergies. In a practical sense, combining AI with IoT allows the various connected devices to analyze the data that they collect, and then act on those data without human intervention.

How does AIoT allow "smart" devices to be even smarter? In retail sales, combining AI software with facial-recognition cameras allows demographic information to be collected from store traffic. AI can then use that demographic data to make decisions regarding product placement and marketing. To help ease traffic congestion, AI can be combined with drones to monitor traffic patterns in real-time. The data collected by the drones can then be combined with AI models to adjust speed limits and the timing of traffic lights; which manages traffic flow without human involvement. Finally, AIoT can make government buildings and workplaces smarter. Facial recognition and RFID sensors can be combined with AI models to identify and track visitors. Using AI to further analyze the collected data can trigger the automatic adjustment of lighting, temperature, and visitor access (Marr, 2019).

Mimicking Human Decision-Making

Another way to define AI systems is through the ability of the system to mimic some of the human characteristics that encompass processes of decision-making. But again, this definition determines that an AI system must be defined by its capability of being "human." An AI program or application is then defined by having these 3 cognitive capabilities (Rouse, 2020):

1. *Learning* - the ability to use data and rule-based algorithms to perform an analysis. Training and testing datasets allow for "learning" from historical data.

2. *Reasoning* – the capability of deciding on the best algorithm for the objective.

3. *Self-correcting* – the capability of the algorithm to constantly monitor and fine-tune the process to maximize the accuracy of the outcome.

Arend Hintze, Assistant Professor of Integrative Biology and Computer Science and Engineering at Michigan State University, defines AI by distinguishing the different types of AI based upon their application (Hintze, 2016):

The most basic type of AI is what he calls *Reactive Machines*. These applications only view the current situation from a direct perspective. That is the system is only reactive to the existing options. It has no ability to use memory or past experiences to make current decisions. An example is *Deep Blue*, the IBM supercomputer that played chess grandmaster Garry Kasparov. Deep Blue had the capability to choose the optimum move out of all possible moves. A Reactive Machine cannot remember what happened in the past; all it can do is perceive the current state.

Similarly, an application that determines credit worthiness for a particular loan can only consider the current application. The application has no concept of past applicants or a prior state of affairs. The application can only perceive the current inputs and base the decision on those.

Expert Systems would be an example of Reactive Machines. Expert Systems are designed to have a narrow focus to very specific tasks. In fact, designing a broader view would be a detriment and reduce the Expert System's capability in making the correct decision. In other words, each time the system encounters a specific situation, it will behave the same way.

The second type of AI, *Limited Memory*, has the ability to process prior data and incorporate it into the systems view of the world. As an example, a self-driving car cannot only look at the current situation, it must have the capability to identify objects in its surroundings and track them over time. The movements of all the other cars, the direction the road is taking, stop lights, etc. are all included in determining what the vehicle does and when.

However, all of the disparate bits of historical information about the car's surroundings are volatile and are not retained as part of a memory that the car can draw upon in a new situation. The car cannot compile experience the way a human can from years of driving and use that experience to develop a representation of the world. The difficulty then is developing AI systems that contain a complete representation, retain experiences, and have the capability of applying knowledge to new situations. The critical question then becomes, how can machines be allowed to create representations of their own?

The current capability of AI systems is represented by the two aforementioned types of AI. The next types to be discussed relate to what researchers aspire to build in the future (i.e., the "vision" of AI). The third type of AI is *Theory of the Mind*. This type of AI will not only develop representations in relation to the world, but also see and comprehend other systems or actors that exist and will determine how these objects impact the AI's behavior. Theory of the Mind Systems will be capable of comprehending and understanding the perceptions, requirements, and viewpoints of other systems. These systems will also be able alter its own behavior in order to interact with other systems on a "human" level.

The fourth type of AI is *Self-Awareness*. Self-Aware systems have the ability to be consciously aware of itself. This awareness is the key factor: awareness is an understanding of internal mechanisms that determine a given state, at a given time, and exhibit the ability to determine the motives and intentions of other actors that interact with the system. Wanting or intending to do something and knowing why you want or intend to do it, is the key difference. We generally know that a stranger knocking on our door is either selling something or attempting to enlighten us to their religious beliefs; therefore, we might consider not answering the door. Consequently, there are inferences that come from the comprehension of memory, and from the learning and functional ability to evaluate memory to make decisions. Currently, systems possess the ability of classification based solely on the current view of the present.

COMPONENTS OF AI

The current discipline of AI is actually a compilation of a number of research areas, each with its own objectives, methods, and applications (Wang, 2019). These "sub-disciplines" are collectively called AI as a matter of historical habit as opposed to any logical or theoretical similarity.

As previously discussed, we can classify AI on the basis of the components, or sub-disciplines that comprise the discipline of AI. Much of the discussion and reference to AI has been specifically focused on Machine Learning. Therefore, the topic of Machine Learning warrants further discussion.

Machine Learning

Machine learning uses a variety of different algorithms to make predictions in terms of classification or segmenting data into discernable categories. There are two basic types of Machine Learning algorithms:

1. *Supervised Learning*. In supervised learning historical data are used to "train" a model algorithm so that new classification predictions can be made. For example, based on the data input of a loan applicant, it can be predicted with a certain level of accuracy whether they will default on the loan or not. This prediction is based on prior data inputs on both credit worthy individuals and those that defaulted.

2. *Unsupervised Learning*. In Unsupervised learning there is no specific outcome variable (like loan default or not) and so, no specific classification. One example is clustering. In this case, records are clustered with similar records, based on the values of the variables, but separated from dissimilar records. Euclidean distance is used to collect records similar to each other and separate those dissimilar. For example, if we wanted to cluster major league pitchers based on their pitching statistics into 3 clusters, we would likely see a cluster containing the good pitchers, one containing the fair pitchers, and another with the poor pitchers.

Deep Learning is a sub-discipline of Machine Learning. In deep learning, we use complex artificial neural networks to solve problems in very large datasets that have different types of unstructured data. These datasets require intricate models with large numbers of hidden nodes to capture the complexity of the problem.

Language Processing (NLP)

NLP is the processing of actual language text (or audio) to predict outcomes like classification. For example, based on the text of an email, we may be able to determine if the message is "ham" (a desired email) or "spam." In addition, sometimes the sentiment (e.g., happy or sad) can be determined based on the words used in a message.

Automation

One of the functions of AI has been the concept of automating certain processes or repetitive tasks that humans do not want to do (or pay to have done), or are unwilling to do. Automation has allowed for the replacement of line workers with robots. Other processes can be automated as well, like loan applications and diagnosis of pap smears. In these cases, *Expert Systems* are developed in an attempt to capture the knowledge of experts, like loan officers or physicians. Expert Systems capture the knowledge of experts in "if-then" rules to make a determination. In a loan application, the knowledge of the loan officer is captured in a rule-based algorithm. This capturing of rules allows for the automation of the application process.

Machine Vision

Machine Vision uses image processing and pattern recognition to identify and compare images. One application would be the identification of signatures. Also, self-driving car applications will use image recognition and deep learning to alter the speed to keep the car away from obstacles and on its prescribed path.

Robotics

Robots are designed and produced to perform repetitive tasks that humans might have difficulty doing accurately, and on a consistent basis. For example, automobile production lines use robots for a variety of tasks.

METHODOLOGY: EVALUATING THE AI CURRICULUM

Now that we have a general idea of how we might define AI and its sub-disciplines, the question becomes, how do these definitions translate into what we are teaching in our degree programs? Our research question then is: Does what we have determined AI to be, and how we have defined it, in terms of AI applications in the real world, match what we are offering our students in the current structure of AI curricula?

In order to evaluate the appropriateness of curricula in higher education, we compared course titles and descriptions at highly rated AI schools to the AI sub-disciplines determined in the literature and outlined in this paper. This study looked at the course offerings in the curricula of the top 20 AI schools, as determined by *U.S. News and World Report*, a recognized source of college and university rankings. The number of courses offered at both the graduate and undergraduate levels were counted to determine the total number of courses in each sub-discipline of AI.

DISCUSSION

As seen in table 1, both introductory and advanced broadly-related AI courses are counted in the first column. Introductory and advanced courses would be those specific to areas like knowledge representation, planning, reinforcement learning, and automated reasoning. According to Russell and Norvig, NLP, computer vision, robotics, machine learning, automated reasoning, and knowledge representation comprise nearly all of AI (Russell and Norvig, 2015). When looking at the curricula in these top schools, introductory courses and advanced courses in specific areas, like knowledge representation and automated reasoning, were aggregated.

As seen in Table 1, these top AI schools offer courses that span the sub-disciplines of AI. Each school generally offers at least one course in each of the current areas that comprise AI. Institutions that are not top tier schools will, at a minimum, offer an Introductory AI course along with at least one course in each of the areas of NLP, Machine Learning, Computer Vision, and Deep Learning, in order to remain competitive. A course in Robotics would round out the curriculum, but it may be argued that the investment in resources may not be seen as justified by the administration. Emphasis on other aspects of AI in the curriculum may be more attractive to prospective students interested in AI beyond Robotics.

Table 1: AI Courses at Top-Rated Schools

AI School by Rank	Intro to AI	Machine Learning	Machine Vision	NLP	Deep Learning	Robotics	Total Courses
1	2	1	1	1	0	0	5
2	6	2	3	4	1	4	20
3	1	1	0	0	0	1	3
4	2	1	3	1	1	4	12
5	3	1	2	1	0	3	10
6	3	5	1	4	1	0	14
7	1	4	2	1	0	5	13
8	1	2	0	1	0	0	4
9	4	2	1	1	1	1	10
10	3	2	4	2	0	3	14
11	4	2	0	3	0	1	10
12	1	4	1	3	4	1	14
13	1	3	3	1	0	1	9
14	2	3	3	2	0	0	10
15	1	2	1	3	1	2	10
16	2	1	1	0	1	0	5
17	3	4	2	0	0	0	9
18	3	3	1	1	0	2	10
19	2	3	1	0	0	3	9
20	1	3	3	2	0	0	9
Total Courses	46	49	33	31	10	31	200

CONCLUSIONS

As discussed throughout this paper, an effective AI curriculum must be focused on pragmatic application, as opposed to theoretical concepts. Consequently, an effective curriculum would prepare graduates to develop practical applications of AI that are able to perform tasks and solve problems. We concur with Wang's definition of AI as a task oriented information processing system that can adapt to its environment, is constrained by limited knowledge and resources, but has the goal of solving problems (Wang, 1995, 2019). This allows for broad-based curriculum development without the limitation or restriction of conformance to any vague and elusive definition of human intelligence.

In addition to being applied and practical, an effective AI curriculum must prepare graduates to develop AI systems that are *autonomous*. As proposed in this study, the promise of fully autonomous AI relies on "strong" AI applications; applications that can *autonomously* apply the knowledge that they gather to solve new problems. The emerging field of AIoT will undoubtedly be a critical component to these strong AI applications. AIoT, as previously discussed, will allow everyday devices to collect data and make their own autonomous decisions. Advancements in AIoT takes the concept of autonomous devices further, and perhaps someday allowing everyday devices to be *self-aware*. It is imperative, therefore, that AI graduates learn the concepts and skills necessary to develop devices that can determine the motives and actions of other devices and react accordingly . . . all without human intervention.

Out of the top 20 AI schools examined in this paper, most schools include courses in nearly all of the AI sub-disciplines. We find, then, that, at least the top rated schools (according to US New and World Report), are coordinating their curricula to meet the AI research and application challenges. More precisely, most schools in the list offer courses in *Artificial Intelligence*, *Machine Learning*, and *Robotics*. However, as revealed by our analysis, some schools could consider adding the other available AI courses to make their degree programs more comprehensive. For example, *Deep Learning* is only offered at seven of the top-ranked schools. It should be noted that *some* Deep Learning topics may be included in an individual school's Machine Learning course, or courses. However, it is difficult to determine all possible course content from a course name and a limited course description.

Among the top-ranked schools, *Machine Vision* and *Natural Language Processing* (NLP) courses are notably more prevalent than Deep Learning courses. However, there are some schools in the list that do not offer *any* courses in Machine Vision or NLP. Specifically, three of the schools in the ranked list do not offer courses in Machine Vision and four of the schools do not offer courses in NLP. Considering all possible courses in the list, the most robust AI curriculum would, therefore, contain Deep Learning, Machine Vision, and NLP courses (in addition to *Robotics* and *Introduction/Advanced AI* courses).

Since future technological systems must rely on the practical applications of AI, the typical AI curriculum could be even more robust by including courses that specifically focus on IoT and AIoT. From the ranked listing of AI schools, seven of the schools do not offer courses in *Robotics*. Since IoT and AIoT both involve the development of robotic "things," a Robotics course would certainly be useful to the AI graduate. As AI research strives to develop the elusive self-aware applications, the highest goal of AI, it may be argued that AI degree programs might consider including the exploration of IoT and AIoT technologies as an important step in reaching that goal.

Future research might look at how AI systems are more precisely being applied in industry. What AI sub-disciplines are the basis for such systems? With this information, applied AI systems can then be mapped to specific courses and curricula, identifying adequate knowledge bases, capabilities, and limitations.

REFERENCES

Bhatnagar, S., Alexandrova, A., Avin, S., Cave, S., Cheke, L., Crosby, M., Feyereisl, J., Halina, M., Loe, B. S., Eigeartaigh, S. O., Martnez-Plumed, F., Price, H., Shevlin, H., Weller, A., Winfield, A., & Hern´andez-Orallo, J. (2018). Mapping intelligence: requirements and possibilities. *Philosophy and theory of artificial intelligence* (Muller, V.C., Ed.). Springer.

Franenfield, J. (2020, March 13). *Artificial Intelligence*. Investopedia, https://www.investopedia.com/terms/a/artificial-intelligence-ai.asp

Hintze, A. (2016, November 13). *Understanding the four types of AI, from reactive robots to self-aware beings.* The Conversation. https://theconversation.com/understanding-the-four-types-of-ai-from-reactive-robots-to-self-aware-beings-67616

Marr, B. (2017). What is the Internet of Things (IoT) and how will it change our world? *Forbes*. Retrieved May 28, 2020, from https://www.bernardmarr.com/default.asp?contentID=964

Marr, B. (2019). What is the artificial intelligence of things? When AI meets IoT. *Forbes*. Retrieved May 31, 2020, from https://www.forbes.com/sites/bernardmarr/2019/12/20/what-is-the-artificial-intelligence-of-things-when-ai-meets-iot/#1ae7e9e1b1fd

Monett, D., and Lewis, C.W. P. (2018). Getting clarity by defining artificial intelligence - a survey. *Philosophy and theory of artificial intelligence* (Muller, V.C., Ed.). Springer.

Nilsson, N. J. (2009). The quest for artificial intelligence: a history of ideas and achievements. Cambridge: Cambridge University Press.

Rouse, M. (2020 April). *Artificial intelligence*. TechTarget.com. Retrieved May 20, 2020, from https://searchenterpriseai.techtarget.com/definition/AI-Artificial-Intelligence

Russell, S., Norvig, P. (2015). Artificial intelligence: a modern approach (3rd Edition). Pearson Education.

Wang, P. (1995). *Non-axiomatic reasoning system: exploring the essence of intelligence.* [Doctoral Dissertation, Indiana University].

Wang, P. (2006). Symposium. *Artificial intelligence: what it is, and what it should be?*. (Lebiere, C. & Wray, R., Ed.). Between a Rock and a Hard Place: Cognitive Science Principles Meet AI-Hard Problems, 97–102.

Wang, P. (2008). What do you mean by “AI?”. (Wang, P., Goertzel, B., & Franklin, S., Ed.). *Proceedings of the First Conference on Artificial General Intelligence*, 362–373.

Wang, P. (2019). Defining artificial intelligence. *Journal of Artificial General Intelligence*, 10(2), 1–37. DOI: 10.2478/jagi-2019-0002