Individual assignment

Third iteration: The Imitation Game, The Language Game, The Learning Game and the Moving Game

1.

According to John McCarthy, who coined Artificial Intelligence (AI) in the mid-1950s (McCarthy et al., 1955), the term can be defined as "the science and engineering of making intelligent machines, especially intelligent computer programs" (McCarthy, 2007, p. 2). Cambridge Business English Dictionary (2018) presents two definitions of AI: "Computer technology that allows something to be done in a way that is similar to the way a human would do it" and "the study of how to produce machines that have some of the qualities that the human mind has, such as the ability to understand language, recognize pictures, solve problems, and learn". According to the American technology company Amazon, AI can be defined as "the field of computer science dedicated to solving cognitive problems commonly associated with human intelligence, such as learning, problem solving, and pattern recognition" (Amazon Web Services, 2018).

All of these definitions have similarities, but McCarthy's definition can be perceived as slightly more general than the others, since it doesn't say anything about what the term intelligence implies. The definitions retrieved from Amazon and Cambridge Business English Dictionary include this aspect; an intelligent machine has some of the qualities that the human mind has.

In this course, AI has been referred to as computer systems that can learn and improve on the basis of large data sources. AI systems has mainly been discussed as systems that are good at performing a single task. Such systems have for instance the ability to understand language or recognize pictures, but they only work within a very limited context. Noessel (2017), who divides AI into three parts, categorizes the intelligence of such systems as narrow. According to Noessel (2017), the other categories of AI are artificial general intelligence and artificial super intelligence. Since the definition retrieved from Amazon and Cambridge Business English Dictionary emphasizes the limitations of current AI systems, they can perhaps, to use Noessel's terminology, be perceived as more closely linked to artificial narrow intelligence, than artificial super intelligence. McCarthy's definition is formulated more general, and has a wider scope.

2.

Robotics as a field can be defined as "the study, design and use of robotic systems for manufacturing" (Deep et al., 2015). Another definition of the field, retrieved from Cambridge Business English Dictionary (2018), is as "the science of making and using robots". According to National Aeronautics and Space Administration (NASA), robotics can be defined as “the study of robots” (National Aeronautics and Space Administration, 2009).

These definitions have in common that they mention robotics in relation to the study of robots or robotic systems. A robot can be defined as “a reprogrammable, multifunctional manipulator designed to move material, parts, tools, or specialized devices through various programmed motions for the performance of a variety of tasks" (The Robot Institute of America in Ramon, 2014) . Deep et. al (2015) definition of robotics can be perceived as somewhat more specific than the ones retrieved from NASA (2009) and Cambridge Business English Dictionary (2018). The definition retrieved from NASA (2009) is possible the most general one as it only includes “the study” of robots, not the design, the making or the use of them.

3.

According to Patcha and Park (2007), Machine Learning (ML) can be defined as “the ability of a program and/or a system to learn and improve their performance on a certain task or group of tasks over time”. Mohri et al. (2012, p. 1) propose a definition of ML as “computational methods using experience to improve performance or to make accurate predictions”. As stated in a publication by the American media company Forbes, ML can also be defined as “a current application of AI based around the idea that we should really just be able to give machines access to data and let them learn for themselves” (Marr, 2016).

These definitions have in common that they emphasizes the importance of ML as systems/methods/programs that are able to learn and improve performance over time. The definition retrieved from Forbes implies AI as a broader concept of machines being able to carry out tasks in a “smart” way, and that ML is a current application of this field (Marr, 2016). Of these selected definitions, this is the only one emphasizing how ML and AI are related.

In this course, as shown in Figure 1, Machine Learning has been discussed as a current application of AI. Machine learning can perhaps be perceived as the statistical arm of AI, where the focus is on the study and constructions of algorithms that can learn from and make predictions on data with an emphasis on high statistical accuracy. It is important to notice that a system needs a lot of good quality data to be able to learn and to make accurate predictions. There are several Machine Learning techniques, both supervised and unsupervised. In this course, Deep Learning, a subset of Machine Learning that powers the most “human-like” AI, has been demonstrated. Deep learning is involving multiple hidden layers in an artificial neural network. In Deep Learning, algorithms are structured in layers to create an Artificial Neural Network that can learn and make intelligent decisions on its own. ML, as discussed in this course, have similarities with the collected ML definitions. As observed, ML is commonly referred to as a subsystem of AI that are able to learn and improve performance over time.

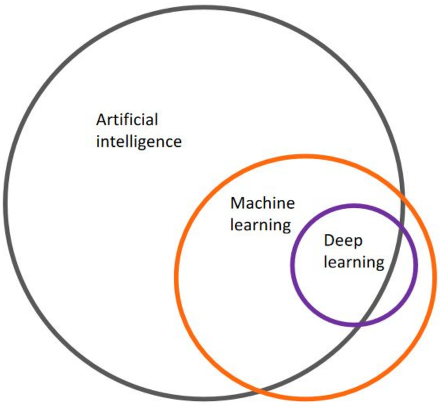

Figure 1: A Venn diagram, as shown in the course, showing the relationship between Artificial Intelligence, Machine Learning and Deep Learning.

4.

Although AI and Robotics are similar in many ways and can be combined in advanced systems, I understand these as two entirely separate technologies or fields of science. I understand AI as computer technology dedicated to solving cognitive problems in a way commonly associated with human intelligence, with or without interacting with the real-world. Unlike AI, I understand Robotics as the study, design and use of programmable machines, so-called robots, that requires an interaction with the real-world. Robots can carry out physical processes and may be controlled by a human operator and/or an AI system. As shown in Figure 2, I understand Artificial Intelligent Robots as robots controlled by an AI system.

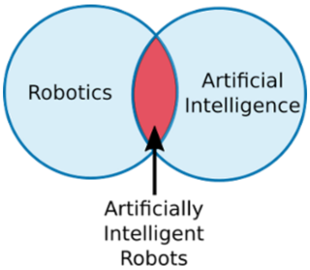

Figure 2: A Venn diagram showing the logical relationship between Robotics and Artificial Intelligence (Owen-Hill, 2017). The diagram is highlighting how the two fields intersect.

5.

I choose to define AI as a field of computer science dedicated to develop and produce intelligent system. Here, intelligence refers to having some of the same qualities as the human mind has, such as the ability to learn, understand language, recognize patterns, and solve problems. Machine Learning, a current application of AI, can be defined as the study and construction of algorithms that can learn from and make predictions on data.

When I formulated my definition of AI, I chose to draw upon the definitions formulated by McCarthy (1955), Amazon (2018) and Cambridge Business English Dictionary (2018). Inspired by the definitions formulated by Amazon and Cambridge Business English Dictionary, I chose to focus on what the term intelligence implies in this setting. I chose to define AI as a system that have some human-like qualities, for instance the ability to learn from data. The purpose of including this formulation in the definition, was to clarify the relation between AI and ML. In my definition of ML, I chose to include the importance of data, in addition to an attempt to convey how ML and AI are related. As I understand, ML is a current application of AI, as ML techniques can handle large amounts of data, or so-called Big Data, in an intelligent way. ML algorithms need a lot of data to learn and to able to make accurate predictions on data.

6.

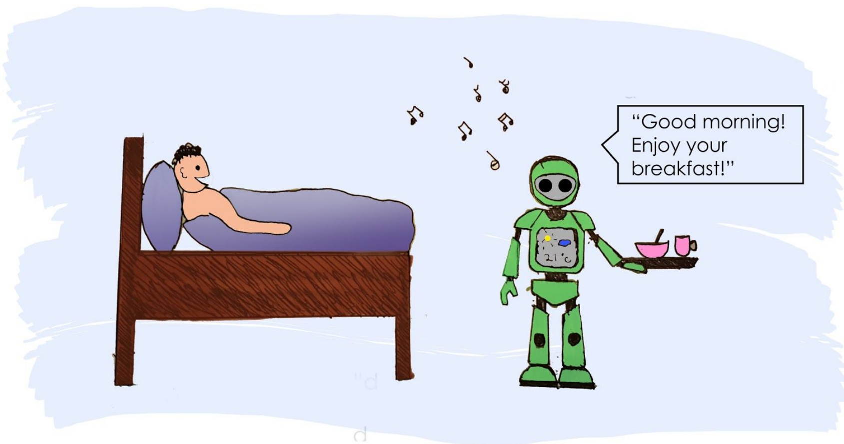

Figure 3: Interaction with an Artificial Intelligent Robot.

Figure 3 is showing an Artificial Intelligent Robot serving breakfast. I imagine waking up in the morning by the smell of freshly brewed coffee, and be served breakfast in bed by my personal assistant.

Interaction with AI, and designing for interactions with AI, currently concerns what Noessel (2017) calls Artificial Narrow Intelligence. In this course AI-based interactive systems has been referred to as interactive systems where important components are powered by Artificial Intelligence. AI-based interactive systems are typically set up for learning and improvement on the basis of large datasets and gathering of new data. This causes several challenges and opportunities for the field of Human-Computer Interaction (HCI). When designing for AI-based interactive systems, there are several factors that have to be taken into consideration.

As discussed in this course, there exist a set of tentative design principles for a user-centered design of an AI system. As the systems learns, it is important to design for change, and if possible, explain to the user the dynamic character of the system. Figure 4 is illustrating how a system can inform the user about its dynamic character. It is also important to convey the system's capabilities and be clear on its limitations. As mistakes are inevitable, the AI system has to be designed for uncertainty. The system has to be able to learn from mistakes, in addition to make recovery easy. As an AI system is fuelled by large data sets, it is also important to design for data capture. However, when collecting huge amounts of data, serious security and privacy issues may occur. There are for instance challenges in relation to wide scale electronic surveillance, profiling, and disclosure of private data. Hence, to adopt a *privacy by design* approach in the development and application of the AI system, is crucial. Data gathered through interaction should be used for system improvements, so users will benefit from the collected data.

In this course, there has also been proposed a set of principles of conversational interaction design. For the Natural Language User Interface (NLUI) to be successful, there is a need to understand conversation processes like speech acts and conversational implicature. If you for instance ask a Conversation Agent (CA) if it knows what time it is, you want it to give you the time, not a confirmation of its knowledge of time. In this course, four maxims of conversation have been discussed. It is important to be as informative as required (Maxim of Quantity), to speak what you believe is the truth (Maxim of Quality), to be relevant (Maxim of Relation) and to be clear and unambiguous (Maxim of Manner). The CA should also have a consistent persona, in addition to be able to present itself in a good way. It is also important that the CA has the ability for conversation repair, for instance by failing gracefully.

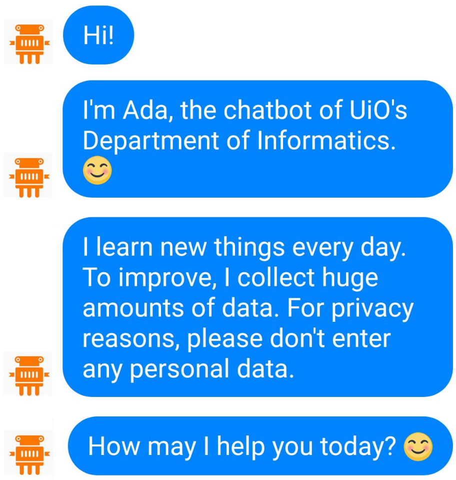

Figure 4: A sketch illustrating some of the characteristic of interaction design for AI-based systems. The chatbot is designed to explain its dynamic character. It is designed for data capture, but it also warns the user not to enter any personal data.

7. and 8.

In the article "On the Subject of Objects: Four Views on Object Perception and Tool Use", Tarja Susi and Tom Ziemke (2005) examine some of the theories on the relation between an agent and its environment. The article present a comparison of four different conceptions of how subjects perceive objects/artefacts/tools and their possible use: von Uexküll's *functional tone*, Heidegger's *equipment*, Gibson's *affordance*, and Kirsch's *entry point*. It is argued that the differences between these concepts deserve attention given that the relation between a subject and an object is important to understand human cognition and how humans interact with tools and technology. According to Susi and Ziemke, the subject-object, or agent-environment, relationship is also a major issue in regards to understanding how artificial objects, like for instance robots, meaningful can perceive and interact with their environment. In this task I have chosen to focus on von Uexküll's *functional tone*, and describe this perspective into more detail.

The German biologist Jakob von Uexküll's *functional tone* perspective is inspired by the Kantian insight that all knowledge is determined by the knower's subjective ways of perceiving and conceiving. Jakob von Uexküll based his concept on the idea that each subject ascribe meaning to the physical objects it is faced with, and construct its own subjective universe (Umwelt), a closed unit consisting of the subject's perceptual world and its effector world. Objects are initially neutral, but as subjects imprint meaning upon them, they are transformed into meaning-carries. As an object becomes a meaning-carrier, it assumes a certain functional tone. According to von Oxeküll, it's the subject's prevailing mood that determines which functional image will lend its tone to the perceptual image. If an object is used in different ways, it may also possess several effector images in the same subjective universe, which then lend different tones to the same perceptual image. Objects that are not transformed into meaning-carrieres by a subject, are totally neglected.

The German philosopher Martin Heidegger is the developer behind the concept of *equipment*, and his work has some overlaps with that of von Uexküll. Heidegger's perspective is focused on the individual's social and cultural embedding, and what it means for a being to exist. The concept of *equipment* is based on the idea that there is an interdependent relation between subjects and objects, and therefore, these cannot be considered as separate entities. Humans must be considered in their form of being-in-the-world, and the way individuals perceive objects depends on their ongoing activity/context.

The concept of *affordance* was developed by the American psychologist James J. Gibson. Gibson emphasises the reciprocal relationship between subject and object, and his concept is based on the idea that each subject lives in its own set of affordances which cut across the subject-object distinction. According to Gibson, objects can either be attached or detached. The latter group of objects includes tools, characterized by being graspable, portable, and manipulable. Affordances are, according to Gibson, objective properties of the environment working as entry point into the mutuality between a subject and its environment. Affordances, made available in the perceived patterns of light that are reflected from surfaces, are always in relation to the subject and its movements in the environment. Since the affordances a human perceive at a given moment are depending on the context of the subject's activity, perceived affordances are, according to Gibson, subjective.

The concept of *entry points* appears in the work of David Kirsh, a cognitive scientist influenced by Gibsonian psychology. According to Kirsh, human beings actively structure their surroundings by creating so-called entry points, which they use to scaffold their daily life. Entry points provide a structure that help people to improve their performance, and they may be objective or subjective. Kirsch's concept is based on the idea that emphasis lies on the co-adaptation of agent and environment. When discussing environments, Kirsh mainly refers to work contexts, such as offices.

9.

In the text “Is AI Riding a One-Trick Pony?”, Somers (2017) examine the current moment of AI. The text thoroughly describe the concept of deep learning, including an explanation of how the so-called backpropagation technique can train a deep neural net, i.e. a net with more than two layers. In 2012, Geoffrey Hinton, referred to as the “the father of deep learning” (Somers, 2017), and two of his students, published a paper where they showed that deep neural nets, trained using backpropagation, beat state-of-the-art systems in image recognition. This paper is, according to Somers (2017), the cause to the massive interest from the outside world in deep learning. For Hinton, who 26 years earlier had showed that backpropagation could train a deep neural net, this was, “a payoff long overdue” (Somers, 2017). Increasing computational power had finally made good of his discovery. However, despite recent progress in the field, we are still, according to Somers (2017), “largely in the dark about how deep-learning systems work, or whether they could ever add up to something as powerful as the human mind”.

10.

I choose to describe how interaction with AI is portrayed in the American documentary film “Lo and Behold, Reveries of the Connected World”, directed by Werner Herzog (2016). The film explores the beneficial opportunities the Internet, robotics, AI, and more have afforded humans, but also the dangers and ethical issues that arise due to the human drive for technological change. The film is divided into chapters, and in the chapter called “Artificial Intelligence”, Herzog visits Pittsburgh, and “Chimp”, a robot that can test its limbs of its own. Nevertheless, J Michael Vandeweghe, a robotics engineer at the Carnegie Mellon University, explains that we still are long ways away from a robot having a complete understanding of the world. In the film, Herzog asks several of his interviewees to reflect on the question “Does the Internet dream of itself?”. To answer this question, Vandeweghe presents a scenario that may be looked upon as a robot almost dreaming of itself. In the scenario, a robot is conceptualizing what is going to happen in the future, and “thinking” about different scenarios. When robots, via the Internet, start to exchange information with one another, we might have, according to Vandeweghe, a robot dreaming about places it hasn’t even been. One of the key things that spurred the research project behind “Chimp”, was realizing that it was too dangerous for humans to do certain operations, such as performing rescue missions in disaster zones.

Another one of Herzog’s interviewees, Elon Musk, is far more critical than Vandeweghe to the use of AI. Musk points out that even though AI can be used to help humanity, there are major risks of doing so. “The biggest risk”, he says, “isn’t that the AI will develop a will on its own, but rather that it will follow the will of people that establish its utility function” (Musk in Herzog & Maconick, 2016, 01:22:10). For instance, if an AI system is aiming to maximize the value of a hedge fund portfolio, a quite benign intent, it can decide that the best way to do so, is by starting a war. By this example, Musk raises awareness about the unforeseen dangers that may arise from the use of AI; it could have quite a bad outcome if it hasn’t been well thought out.

11.

I understand autonomy as the condition of being self-governing. A machine, like a robot, that have a high level of autonomy can act on its own accord for a long period of time. This doesn't mean that the robot has to be intelligent. A non-intelligent robot can for instance be programmed to pick up an object and place it elsewhere, and to continue this sequence of events until it is switched off. The robot is autonomous since it doesn't require any human input after it has been programmed, but it isn't able to carry out the task in a so-called intelligent way. I understand human autonomy as the quality of being self-governing, independent, and having a free will.

12.

The term “AI” was first coined in a workshop proposal submitted by John McCarthy (Dartmouth College), Marvin Minsky (Harvard University), Nathaniel Rochester (IBM), and Claude Shannon (Bell Telephone Laboratories) in August 1955 (Press, 2016; McCarthy et al., 1955). The workshop, which took place at Dartmouth College in 1956, is considered as the official birthdate of the field of AI (Press, 2016).

13.

In the article "What we talk about when we talk about context", Paul Dourish (2004) presents the approach of “embodied interaction”. Why is this model of context relevant for the field of HCI, according to Dourish?

14.

The article "Interactive robots as social partners and peers tutors for children: A field trial", written by Kanda et. al (2004), examine the proposition that robots and children can form relationships, and that children can learn from robots as they learn from their friends. Why does Kanda suggests that interactive robots should be designed to have something in common with their users?

15. a.

The article "Like Having a Really Bad PA", written by Luger and Sellen in 2016, addresses how CA systems fail to to bridge the gap between user expectation and system operation. The authors find user expectations to be “dramatically out of step with the operation of the systems” (Luger & Sellen, 2016, p. 5286). Luger and Sellen argues that the CA research of today, fail to truly understand how and why CA systems are used. By conducting 14 semi-structured interviews with users of different CA systems, Luger and Sellen seek to understand user experience of CA systems. The questions addressed in the article are; what factors currently motivates and limits the ongoing use of CAs in everyday life, and what should we consider in future design iterations? (Luger & Sellen, 2016, p. 5286-5287).

According to Luger and Sellen, users of CA systems have poor mental models of how their CA work, and they tend to have too high expectations regarding the system’s intelligence, capability and goals. Programmed trigger responses tend to give users unrealistic expectations, and lack of meaningful feedback, tend to reinforce the users’ poor mental models (Luger & Sellen, 2016). The authors found that the majority of their interviewees, used their CA on a daily basis. It was most common to use CAs for relative simple tasks, such as checking upcoming weather and setting alarms, particularly in situations where the users had their hands otherwise engaged. The majority of the participants had a reluctance to use their CA for complex or sensitive tasks, especially where they perceived a high social cost to failure. Factors that had negatively affected the interviewees use of their CA, was mainly that their CA had misunderstood their words or commands. In situations where the CA had responded to task requests by defaulting to on-screen web-search results, this was commonly perceived as a system failure. A majority of the participants did express their desire to have more natural conversational interactions with their CA, in addition to reporting issues regarding a lack of feedback and transparency. Concludingly, Luger and Sellen state that there is a need for humanlike cues and affordances relied upon by multimodal systems.

Luger and Sellen found that participants with technical skills were better able to “see beyond artificial humanlike qualities to devise their own mental models of interaction”. Less skilled users described greater levels of frustration, leading them to doubt the intelligence of their CA. These findings may indicate that user expectations should be “scaffold through more considered revelation of system intelligence through design" (Luger & Sellen, 2016, p. 5294). In my opinion, to bridge the gap between user expectation and system operation, AI-based systems should in general be designed to reveal their intelligence. During conversations, humans use a variety of cues to communicate intelligence, and some of these cues may be relevant to consider in the design of AI-based systems.

Luger and Sellen found that playful and humorous interactions had the effect of reinforcing anthropomorphic qualities, thus compounding users' expectations of system capability. Framing systems as anthropomorphic raises user expectations about the extent of capabilities, and may result in user dissatisfaction. However, the authors found that playful interactions with a CA system could act as affordances, in that they suggest the possibility of action. When designing AI systems, it may be important to reconsider the interactional promises made by humorous engagement and explore how such engagements could instead support user assessment of system intelligence.

The inability of users to assess the intelligence of the CA, was an overarching themes throughout Luger and Sellen's findings. According to Sellen and Luger, their findings indicate that whilst users applied a mental model of human communication, it was revised in light of their experiences with thier CA. In situations where users were not able to draw from a technical frame of reference, they tended to blame themselves, and often abandoned particular types of task requests. The article state that there is a need for more thoroughly investigations regarding how to convey system limitations and capabilities. In my opinion, these investigations may be relevant for the design of AI-based systems in general.

Luger and Sellen found that the majority of users engage with their CA system only up to the point that it ceases to provide utility. The principle use-case of a CA system is, according to Luger and Sellen, hands-free. This implies that it is an alternative primary task, rather than the conversation, that is the focus of attention. In their study, Luger and Sellen found that the primary user goal of their interviewees wasn't solely to use the CA, making the system a means to an end rather than an end in itself. If a CA reverted to screen-based response in situations where the user were engaged in activities that also required a level of visual attention, this was, from a user perspective, perceived as a system failure. According to Luger and Sellen, there is a need for more investigations related to the design goal of current CA system and how these might be rethought to deliver a more compelling user experience. Such investigation may also apply to the design of other AI-based systems.

16.

Due to technological evolution, the "possibilities for automating tasks of human operators have become more sophisticated, and as well as the possibilities to improve human-machine performance in complex systems" (Save & Feuerberg, 2013). According to Beer et al. (2014, p. 74), "developing fully autonomous robots has been a goal of roboticists and other visionaries since the emergence of the field, both in product development and science fiction”. However, automation isn’t “all or nothing” (Sheridan & Verplanck, 1978), i.e. a matter of either automating a task entirely or not automate at all, it’s rather to “decide on the extent of automating it” (Save & Feuerberg, 2013). In 1978, Sheridan and Verplanck proposed an initial 10-point scale of levels of automation, representing a continuum of levels between low automation (Level 1) and full automation (Level 10). At the lowest level, the computer offers no assistance, and the human must take all decisions and actions. At the highest level, the computer decides everything, acts autonomously, and totally ignores the human.

I understand levels of automation as the degree to which a task is automated, ranging from complete human control to complete computer control. As new technology evolves, it may be tempting to employ automatic means in systems and processes. However, when deciding on level of automation, several factors have to be taken into account. According to Save and Feuerberg (2013), “the choice of the ‘optimal’ level of automation in a specific task context is about matching the automation capabilities to a number of operational situations, while increasing the overall performance in efficient human-machine cooperation”. In my view, as tasks and contexts change, so does the optimal levels of automation. In my view, higher levels of automation can be given if the task is simple or straightforward. If the task for instance requires critical thinking, or involves consideration or discretion, the optimal level of automation may be lower. Deciding on level of automation involves finding the optimal balance between human and machine control. Lower levels of automation gives the human a higher level of control, while higher levels of automation gives the human a lower level of control.

An AI system can either have a static level of automation, or it could, depending on the task or context, pass control from automation to the human (adaptive automation). An AI system, like a self-driving car, can for instance, depending on the task or context, take all of the decisions, some of the decisions, or none of the decisions. Given that the self-driving car is equipped to “sense” the driver’s physiological or psychological state, this can be used for deciding on a suitable level of automation. AI components, like sensors or cameras, can probably detect certain dangers earlier than what humans are capable of. In situations where there, for instance, is a risk of the vehicle hitting an animal, the AI system can respectively override the human driver to prevent roadkill. This is an example of how adaptive automation can be used to increase safety. However, when implementing such systems, privacy and other issues have to be taken into account. As humans have a tendency to overestimate technology’s intelligence (Luger & Sellen, 2016), it’s important that an AI system is able to convey its level of automation, including its limitations and capabilities. There are tasks that AI systems can perform better than humans, and tasks that humans can perform better than AI systems. The optimal level of automation is, in my view, the level where humans and AI systems co-exist in a state of symbiosis, i.e. a mutually beneficial relationship that allows both parties to profit in tandem with the other.

17. References

Amazon Web Services. (2018). What is Artificial Intelligence? Retrieved September 6, 2018, from https://aws.amazon.com/machine-learning/what-is-ai/

Artificial Intelligence. (2018). *Cambridge Business English Dictionary.* Cambridge, United Kingdom: Cambridge University Press. Retrieved September 9, 2018, from https://dictionary.cambridge.org/dictionary/english/artificial-intelligence#dataset-cald4

Beer, J. M., Fisk, A. D. & Rogers, W. A. *Toward a framework for levels of robot autonomy in human-robot interaction. Journal of Human-Robot Interaction, 3(2)*, 74-99. doi:10.5898/JHRI.3.2.Beer

Deep, A., Singh, J., Narayan, Y., Chatterji, S. & Mathew, L. (2015). Robotic arm controlling using automated balancing platform. *Communication, Control and Intelligent Systems (CCIS), 2015*, 282-285. IEEE. Retrieved September 6, 2018, from https://www.researchgate.net/profile/Yogendra_Narayan3/publication/304412762_Robotic_arm_controlling_using_automated_balancing_platform/links/58cfd17392851c5009efa7da/Robotic-arm-controlling-using-automated-balancing-platform.pdf

Dourish, P. (2004). What we talk about when we talk about context. *Personal and ubiquitous computing 8 (1)*, 19–30. doi:10.1007/s00779-003-0253-8

Herzog, W. (Producer/Director) & Maconick, R. (Producer). (2016). *Lo and Behold, Reveries of the Connected World* [Motion Picture]. USA: NetScout Systems, Inc.

Kanda, T., Hirano, T., Eaton, D. & Ishiguro, H. (2004). Interactive robots as social partners and peers tutors for children: A field trial. *Human-Computer Interaction, 19(1)*, 61-84. doi:10.1207/s15327051hci1901&2_4

Luger, E. & Sellen, A. (2016). Like having a really bad PA: the gulf between user expectation and experience of conversational agents. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 5286-5297. ACM.

Marr, B. (2016). What Is The Difference Between Artificial Intelligence And Machine Learning? *Forbes*. Retrieved September 10, 2018, from https://www.forbes.com/sites/bernardmarr/2016/12/06/what-is-the-difference-between-artificial-intelligence-and-machine-learning/#2d274cf02742

McCarthy, J. (2007). What is Artificial Intelligence? Stanford, USA: Stanford University. Retrieved September 10, 2018, from http://www-formal.stanford.edu/jmc/whatisai.pdf

McCarthy, J., Minsky M., Rochester, N. & Shannon C. (1955). A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence. Retrieved September 10, 2018, from http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html

Mohri, M., Rostamizadeh, A. & Talwalkar, A. (2012). *Foundations of Machine Learning*. Cambridge, US: The MIT Press.

National Aeronautics and Space Administration. (2009). What is Robotics? Retrieved

September 10, 2018, from https://www.nasa.gov/audience/forstudents/k-4/stories/nasa-knows/what_is_robotics_k4.html

Noessel, C. (2017). Designing Agentive technology: AI that works for people. Rosenfeld Media.

Owen-Hill, A. (2017). What's the Difference Between Robotics and Artificial Intelligence? Retrieved September 11, 2018, from https://blog.robotiq.com/whats-the-difference-between-robotics-and-artificial-intelligence

Patcha, A. & Park, J. (2007). An overview of anomaly detection techniques: Existing solutions and latest technological trends. *Computer Networks 51(12)*, 3448-3470. 3448-3470. doi:10.1016/j.comnet.2007.02.001

Press, G. (2016). A Very Short History Of Artificial Intelligence (AI). *Forbes*. Retrieved September 10, 2018, from https://www.forbes.com/sites/gilpress/2016/12/30/a-very-short-history-of-artificial-intelligence-ai/#4043dae76fba

Ramon, M. C. (2014). Assembling and Controlling a Robotic Arm. *Intel Galileo and Intel Galileo Gen 2. Apress.* Berkeley, US. doi:10.1007/978-1-4302-6838-3_11

Robotics. (2018). In: *Cambridge Business English Dictionary.* Cambridge, United Kingdom: Cambridge University Press. Retrieved September 10, 2018, from https://dictionary.cambridge.org/dictionary/english/robotics

Save, L. & Feuerberg, B. (2012). Designing Human-Automation Interaction: A new level of Automation Taxonomy. In D. de Waard, K. Brookhuis, F. Dehais, C. Weikert, S. Röttger, D. Manzey, S. Biede, F. Reuzeau,and P. Terrier (Eds.), Human Factors: A view from an integrative perspective. Proceedings HFES Europe Chapter Conference Toulouse. Retrieved from https://www.hfes-europe.org/wp-content/uploads/2014/06/Save.pdf

Sheridan, T. B. & Verplank, W. L. (1978). Human and computer control of undersea teleoperators. Cambridge, USA: Massachusetts Institute of Technology, Man-Machine Systems Laboratory. Retrieved from www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA057655

Somers, J. (2017). Is AI Riding a One-Trick Pony? *MIT Technology Review*. Retrieved September 18, 2018, from https://www.technologyreview.com/s/608911/is-ai-riding-a-one-trick-pony/