IN5480 - Specialization in Research in Design of IT

Individual assignment fall 2020 - Iteration 3

Table of contents

1 Module one

This individual assignment is divided into four sections, where the first section defines the concept and history of AI. Then write about the word Robot and also discuss the relation between AI and Robots. Furthermore we will find and describe the definition of Universal Design and for the last section describe one of 18 guidelines from Microsoft.

1.1 Concepts, definition and history of AI and interaction with AI

1.1.1 How AI came about

The first use of the term AI, or Artificial intelligence, was first used in 1956 by an American mathematician and logician John McCarthy (Grudin, 2009:49). McCarthy brought together a group of leading researchers from different disciplines (mathematicians, psychologists and social scientists) to a workshop which was hosted by the Macy Foundation. The participant's optimistic forecasts attracted considerable attention.

1.1.2 Three definitions of AI

Definition 1:
"The theory and development for computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages." - (Oxford Reference, 2020)

This definition focuses on the technical ability to do human-like tasks, such as visual perception. The Oxford Dictionary defines AI as more a theoretical and practical framework, and due to the website it is hard to find out which date this definition was explained.

Definition 2:
"AI is a subfield of computer science aimed at specifying and making computer systems that mimic human intelligence or express rational behaviour, in the sense that the task would require intelligence if executed by a human." - (Bratteteig & Verne, 2018:1-2)

This definition focuses on how the AI mimics human intelligence, as well as providing a difference between human and machine intelligence. Bratteteig and Verne har two researchers from the DESIGN group at IFI, they focused on participatory design and user-centered processes. Since the article was from 2018, the definition was quite recent.

Definition 3:

"It is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable." - (McCarthy, 2007)

This definition was first used in a workshop held in 1956, by John McCarthy, which was also the first definition of Artificial Intelligence. This definition focuses more on how AI is a science to make intelligent machines to understand human intelligence, rather than acting like a human being.

1.1.3 Make one definition for AI

My definition:

"Artificial Intelligence (AI) are machines using computer science designed to learn from their experience, make decisions, adjust or perform human-like tasks."

To this definition I wanted to focus on how AI gives computers and machines the ability to "think", which gives them the ability to perform human-like tasks including cognitive abilities. I also wanted to pinpoint that AI were designed using computer science to be able to perform like they do. It has to be made by humans to be able to think like a human.

1.1.4 A company that works with AI

PwC

The company I chose for this task is PwC and how they present artificial intelligence on their websites. PwC presents AI as a definition that gaps the several tasks of the computer. In how great they may perform the tasks (weak AI), for example to identify the contents of a car, and a more general description of AI which is their ability to learn from their experience. One of the examples they use is the "Google DeepMind", which is a project that has a goal to program general AI. Moreover PwC explains how AI can affect you and your business. Here they are giving examples on which field the use of AI will be more effective than others. Within the transport field they use weak AI, for example Tesla, to gather information about the driving patterns of a car. Furthermore PwC also informs us about which tasks AI will replace and which gender this will affect the most (PwC, 2020).

1.1.5 The use and interaction with AI

Avengers: Age of Ultron

This movie is a sequel of The Avengers where the Avengers fight Ultron, an artificial intelligence obsessed with causing human extinction. As we learn about AI in this subject we can understand and notice the increased presence of AI and how science and engineering can have the potential of improving our lives - which can also destroy us. At least in this movie, the theme of *Avengers: Age of Utron* is an excellent example of how AI has gone really bad. One of the main characters Tony Stark creates a deadly autonomous robot getting ready for war, but Ultron's intelligence seeks to take over humans who created it. Ultron is built with a "normal" robot-like body with superhuman strength, speed and reflexes. The problem here is that iterations of Ultron consciousness were self-created. Meaning that for each time he was getting stronger og smarter, fulfilling two main desires, which was bringing peace and order to the universe. With this goal in mind Ultron's move was to eliminate all other intelligences. Throughout this movie we can see how wonderful AI can be, but also considering the possibility of how it also can go rogue.

1.2 Robots and AI systems

1.2.1 How the word Robot came about

The first use of the word *robot* was by a Czech playwright, novelist and journalist Karel Čapek introduced in one of his plays in 1920, Rossum's Universal Robots (R.U.R). The word robot is from an old Church Slavoni word, *robota*, which was translated to "forced labor" (Science Friday, 2020).

1.2.2 Two definitions of Robot

Definition 1:

"Any device for performing computational or physical tasks automatically. For example, in computing, a robot is a program that trawls the Web for information, indexing it, say, for a search…" (Oxford Reference, 2020)

This definition focuses on how robots are programmed to perform automation of machines. It also mentions that a robot could be any device, which I presume is a technical device.

Definition 2:

"A reprogrammable, multifunctional manipulator designed to move materials, parts, tools, or specialized devices through various programmed motions for the performance of a variety of tasks". (Thrun, 2004:11)

This definition focuses on all the different tasks a robot is programmed to perform. Here it is clearly that a robot is reprogrammable, which means that it may unlearn the tasks it has already learned earlier.

1.2.3 Make one definition for Robot

My definition

"A robot is a machine that is capable of working automatically and executing one or more specific tasks. It may replace human effort, but it does not resemble humans in appearance."

This definition focuses on how robots are able to perform automatic tasks that are humanlike functions. Robots are programmed to do one or more specified tasks with various programmed motions that can resemble a human, but it does not need to look like one.

1.2.4 The relation between AI and Robots

The idea of a robot is to do complex tasks, while an AI can simulate the human intelligence. A robot can most likely be hardware like a physical machine, and work as a tool. AI is quite the opposite, it can be a software that is able to understand, learn and reason. Robots are not able to learn something by themselves, but are programmed to do something specific.

1.2.5 A physical robot

ASIMO

ASIMO (Advanced Step in Innovation Mobility) is a humanoid robot developed by Honda with the goal to develop robots that are able to be useful to people. It is capable of moving and performing different humanlike taks without human operation. The coordination between the visual and sensors enables ASIMO to recognize faces and voices, and therefore it can recognize different people who are talking simultaneously. By its physical appearance, it is able to walk, run, run backwards, and hop on one or two legs. Because of its size, which is 130cm, it has become more flexible and is capable of adapting to changing external situations. ASIMO is also able to pour a drink or speak in sign language to communicate with people with hearing problems (Honda, 2020).

1.3 Universal Design and AI systems

1.3.1 Definition of Universal Design

Definition:
"Universal design is design that's usable by all people, to the greatest extent possible, without the need for adaptation or specialized design." (The Universal Design Project, 2020)

This definition focuses on how the design of something should be understandable for everyone, or even better. The design is making sure that it is functional for people with any type of impairment of the body or the mind. It should not give the impression of being specialized to someone with a disability. The core of universal design is to be able to adapt to changes (The Universal Design Project, 2020).

1.3.2 The potential of AI with respect to human

AI has a great potential in universal design. Seeing AI is an example on how we can include technology to people with visual impairment. This app has the ability to read printed text; object or obstacle recognition and labeling utilities. One of the reasons why Seeing AI has a lot of good press is because it combines the features mentioned in one app. As well as it is for free and anyone with Android can download it (VisionAware, 2020). I also believe AI has the potential of helping us in other areas like movement, emotions and reason, and an example of an AI-infused system helping us with emotions could be a mental health chatbot like Woebot. Other examples like how we may use the benefit of AI in larger industries where they use machine learning to make our jobs easier,

1.3.3 The potential of AI for people, including or excluding?

Face recognition can be extremely useful for people and society, therefore the technology needs to be developed thoughtfully and responsibly. For example, we now have smartphones with face recognition to open our phones instead of using a password. Face authentication gives us the feeling that it is made safer and more secure (Google AI, 2020). There are many concerns to worry about. Face recognition needs to be fair nevertheless, so it does not reinforce biases, but is this always the case? It seems like human bias can sneak into the AI systems. A study found that African American and Latino were discriminated against by the mortgage algorithm. It is said to be good at recognizing white people, but not black faces. Therefore AI should be developed well enough not to exclude people (Vox, 2020).

1.3.4 Do machines understand?

The concept of to “understand” and “understanding” something in technology depends on how you interpret what different “things” do. And by things I mean if it is a software like an AI, or something physical like a robot. For example, AI has now the ability to speech recognition, translation, object recognition, and face recognition, which is a lot of human-like tasks in general. As I mentioned earlier in this assignment AI can learn, reason and “think”, but it can not predict scenarios like humans can. As for machines and robots they can learn certain things and perform certain actions, but they can not perform beyond that or reason like we do. For the human brain, to understand is the ability to reflect your actions and change your behaviour dependent on the scenario or the other person. Therefore I will say that machines can understand, but to some degree, only if they have learned every scenario and contexts possible.

1.4 Guideline for Human-AI interaction

G5 - Match relevant social norms
Ensure the experience is delivered in a way that users would expect, given their social and cultural context.

Example: [Face Recognition] “Smartphones with face authentication should apply and recognise people of any skin color.”

For this section I chose Norman’s Seven Principles (HCI-06129), and already from the first principle there are similarities with the one of the 18 Guidelines from Microsoft. Norman’s principle tells us something about that we should use both knowledge in the world and knowledge in the head, which is quite similar to G13 to learn from user behaviour. These guidelines are used when we wish to assess the interaction between human and computers. These were the seven stages that could be used to transform difficult tasks. While the Guidelines for Human-AI are more focused on interaction with users and how to reply back to humans. They are not necessarily specified to difficult tasks, but more general tasks and performance.

2 Module two

This individual assignment is divided into three sections; the first section defines the characteristics of AI-infused systems and identifying one AI-infused system that exemplifies of the key characteristics, the second section we discuss about how the AI-infused system adheres to or deviates from two design guidelines, and at last the third section we discuss key challenges in the design of chatbots as well as how the guidelines adherence to these challenges.

2.1 Characteristics of AI-infused systems

2.1.1 Key characteristics of AI-infused systems

We differentiate between three types of depths when it comes to AI: *artificial super intelligence, artificial general intelligence* and *artificial narrow intelligence*. Usually, when we speak of interaction with AI, we refer to interaction with artificial narrow intelligence. In the lecture of module two in this course, Følstad identifies four key characteristics: *learning, improving, black bow* and *fuelled by large data sets*.

Learning

Learning refers to the system being dynamic and designed for changing while learning from the user's behaviour. Talking to a learning AI, by sending identical messages or interactions will not give identical responses. The reason is because the AI "*changes via learning over time"* (Amershi et al., 2019:2) through the interactions with people.

Improving

Improving refers to the AI's ability to become better over time by interacting with the users of the system. This usually comes from user input, feedback or making mistakes which, as a result of this, makes it a bit more intelligent for every time. With the fact that AI systems can improve themselves, indicates that they are not perfect - and can make mistakes too. AI's will gradually learn more and more from user feedback and become more accurate.

Black box

Black box refers to the view of AI's as *black boxes*, because of their way of not giving the users insight into the system behind the interaction. A black box refers to "*a system or process that uses information to produce a particular set of results, but that works in a way that is secret or difficult to understand*" (Cambridge Dictionary). The reason why we use this

term is because of our lack of understanding and insight of what happens between the input and output presented in front of us. If AI's are less of a black box design it will, according to Kocielnik et al. "*will lead to higher perceptions of understanding how the AI system works"* (Kocielnik, 2019:4).

Fuelled by large data sets

Fuelled by large data sets refers to the input of AI systems, which means, their knowledge and what powers it. In order for an AI to learn they need a huge amount of data, and this data is collected from user input, or input provided by the initial designer of the system. Through this type of data, AI systems can learn and improve over time to become better at its tasks.

2.1.2 Spotify as an AI-infused system

Spotify

The AI-infused system I chose was *Spotify* (which I use every single day, I am kind of addicted to music). Spotify is the largest on-demand music service application used for different devices, such as, smartphones or web pages. The firm is well known for their record of pushing the boundaries in technology using AI and machine learning to enhance the user experience through user and data insights. Every day tens of millions of users listen to music, and brands like Spotify accumulate a mountain of implicit customer data consisting of song preferences, keyword preferences, playlist data, geographic location of listeners and more.

One example is *"Discover Weekly",* where each Monday individual users are presented with a customised list of thirty songs. The recommended playlist consists of tracks that the user might have not heard before, but the recommendations generated are based on the user's search history pattern and music preference. The users are able to see these suggestions by Spotify and with that, find new and inspiring music to listen to without the need to search for songs. This application definitely embodies the characteristics of *learning* and *fuelled by large data sets*. Machine learning enables the recommendations to improve over time. Not only does it keep users returning, but it also enables exposure for artists who users may not search for originally. Spotify's machine learning is fueled both by user data and by external data, which is also why they have a deep understanding of their customer base and predictive recommendations (Outside Insight, 2020).

2.2 Human-AI interaction design

2.2.1 Main take-aways from the two papers

Guidelines for human-AI interaction

In this paper, Amersi et al. have gathered relevant research, evaluated, and validated it to propose 18 generally applicable guidelines for the Design of human-AI interaction. Their vision is that the *"guidelines can serve as a resource to practitioners working on the design of applications and features that harness AI technologies and to researchers interested in the further development of guidelines for human-AI interaction design"* (Amershi et al. 2019:1). Since the use and expansion of AI is increasing, they have the impression that clear guidelines are significant for this field, but it will also apply in better systems.

Will you Accept an Imperfect AI?

In this paper, Kocielnik et al. investigates the relations between the expectations of users, their acceptance of AI systems, and how they *"explore techniques for shaping end-users expectations of AI-powered technologies prior to use and study how that shaping impacts user acceptance of those technologies"* (Kocielnik et al., 2019:2). In this article, they offer three useful and technologies to mediate this issue; *accuracy indicator, example-based explanation* and *performance control*. The first one explicitly states the accuracy of the system. The second one, seeks to increase user understanding. The third one, allows the user to adjust the performance of the system directly.

2.2.2 How Spotify adheres to both G4 and G13

For this particular task I have chosen the two guidelines G4 and G13, which Spotify adheres to and I will briefly explain how these could inspire improvements.

Guideline 4 - Show contextually relevant information

Display information relevant to the user's current task and environment.

Every time I want to study or perhaps just relax, instead of searching, you can browse your way to a particular category and choose one of the many playlists in that category. In figure 1 I have added an image of the browse "*Focus"* as an example of what kind of playlist you may receive if you wish to study or listen to calm music. Spotify gives you all the music you'll ever need. If you search for an activity, for example "read", Spotify will suggest a playlist with

the searched genre, and when you choose to click on that suggested playlist you will get an overview over many other playlists with the same type of genre.

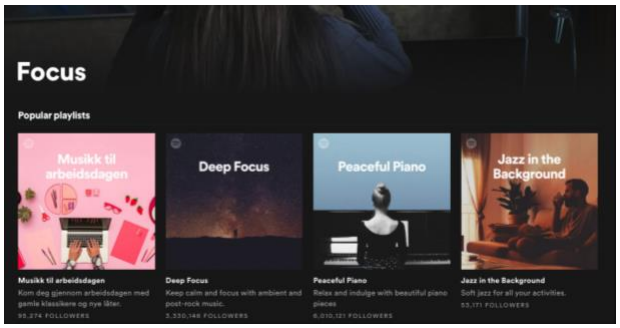

Figure 1: Browse your way to what interest you

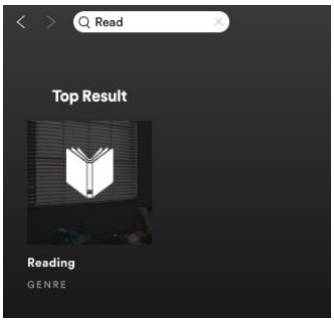

Figure 2: Search and receive a genre

G13 - Learn from user behavior

Personalize the user's experience by learning from their actions over time.

As already mentioned in this assignment, Spotify uses a set of powerful algorithms to use your music and your personal details to shape your entire listening experience. It will do everything to get you to listen to more music, and this algorithm will learn to recognize different aspects of the music that might be desirable. This is also how the "Daily Weekly" playlist works, but Spotify also has another folder "*Made For You"* that gives you an overview of songs and playlists you have listened to lately. It will learn your patterns and generate playlists according to how you "behave" on Spotify.

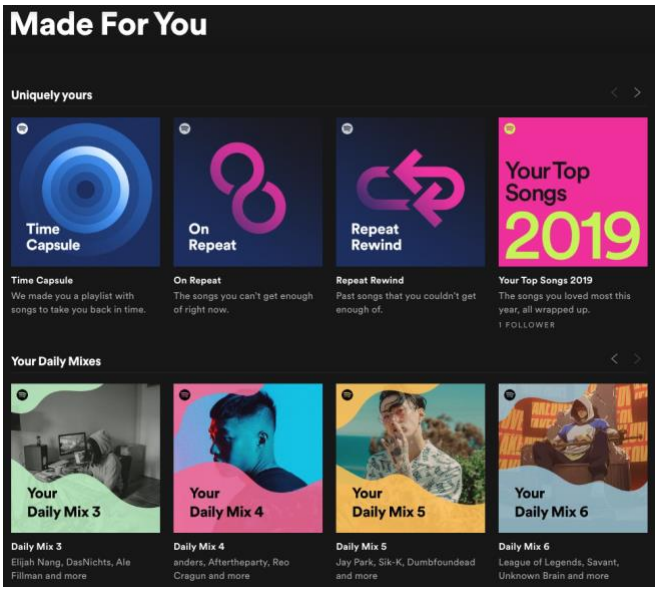

Figure 3: Made For You shows playlists according to your behaviour

2.3 Chatbots / conversational user interfaces

2.3.1 Key challenges in the design of chatbots

Chatbots and the New World of HCI

Some of the key challenges in the design of chatbots and conversational user interfaces, presented by Følstad and Brandtzæg, is the challenge of *conversations as the object of design*. If conversations are to take place *"the future era of chatbots and natural language user interfaces, the designer repertoire of graphical and interaction mechanisms will be greatly reduced"* (Følstad and Brandtzæg, 2019:41). The reason why a dialogue based on user interaction is more dependent on the user's input, is because the user interface *"is to a much greater degree a blank canvas where the content and features of the underlying service are mostly hidden from the user, and where the interaction is more dependent of the user's input"* (Følstad and Brandtzæg, 2019:41).

This represents a clearly challenge to the field of HCI, and as they mention in the paper *"we need to move from seeing design as [...] a task of explaining to the user which content and features are available and which steps to take to reach the desired goal, to [...] a task of understanding what the user needs and how she may be best served"* (Følstad and Brandtzæg, 2019:41).

Like Having a Really bad PA

Lugar and Sellen mention the challenges when the users expectations exceed the systems and issues with the conversational agents' feedback (CA). CA who participated in their study reported that, *"insufficient feedback or visibility of both the limits and the capabilities of the CA was often cited as a factor limiting users' ability to make the system work"* (Lugar and Sellen, 2016:5291). The participants had a difficult time figuring out what the system managed to do or not, which led them *"feeling overwhelmed by the unknown potential, or led them to assume that the tasks they could accomplish were highly limited"* (Lugar and Sellen, 2016:5291).

Their design of the CAs did not make it easy for the participants to know what they were capable of doing, and they also experienced that the agents were not capable of doing what they wanted them to do.

2.3.2 How adherence of G1 and G2 resolve some of the challenges

These guidelines are particularly relevant to managing user's expectations of the system, which will make less frustration and overall better user experience.

G1 - Make clear what the system can do
Help the user understand what the AI system is capable of doing.

G2 - Make clear how well the system can do what it can do
Help the user understand how often the AI system may make mistakes.

If the CAs had adhered these two guidelines, the majority of the participants would have had a much more pleasant experience. As a user it can be difficult to perceive, for example, how well a chatbot is able to answer - which can create bad user experiences if the expectations are the opposite. Therefore, the system must make it clear to the user what the system is capable of doing beforehand. Figuring out the functionalities of how a new technology works, can take a lot of time. As for the second guideline, the CAs would benefit a lot from using this, because many of the participants had troubles to follow up on the topics that they had talked about.

I believe that even if we have these improvements, nothing will ever be perfect. The interaction between us and a chatbot has its potentials and this interface represents a new world to be conquered, but we must remember that having issues within HCI is a huge challenge for us. As Følstad and Brandtzæg mention *"with the emergence of chatbots, even stronger attention on ethics and privacy is needed"* (Følstad and Brandtzæg, 2017:42).

3 Module three

This individual assignment of module three will only contain one section, where we will give an example of humans-robots collaboration. Furthermore, we will give some examples and reflect on the advantages and disadvantages if we decrease or increase their level of autonomy.

3.1 Human AI collaboration

3.1.1 Introduction

Firstly, Philips et al. (2016) reflects and explores why the analog of human-animal teams can be a useful mechanism for supporting human mental models of robots. They also describe functional benefits that human-animal teams do provide, such as capabilities that support teamwork. Being able to see robots as teammates rather than tools makes it a lot easier to gain trust, along with knowledge estimations and attitudes, which plays a significant role in mental model formation in relationships in general.

In this paper, they refer to mental models as structured, organized knowledge possessed by us humans that describes and predicts a system's purpose. The relationship between physical forms and mental models has been shown to influence our initial perceptions of our trust to robots, which is a major key drawn from the paper. Trust determines our interaction with robot teammates, and as they describe as a crucial template for fostering trusting relationships in human-robot teams.

3.1.2 BigDog

BigDog is a legged robot designed to function essentially as a pack if mule and traverse terrain not accessible by wheeled or tracked vehicles, and resembles a large dog (Philips et al., 2020, p.102). This robot was developed by Boston Dynamics, and its intent was to design and build a robot for the military force to carry cargo in order to reduce soldier load.

3.1.2.1 Advantage/disadvantage

The key function of this robot is to carry a large amount of weight through different and challenging terrain, which provide physical benefits that require advanced action and perception. The most beneficial skill the BigDog has is its innate ability to exceed current human abilities, and capable of navigation and balance by their controlling system. But it has

some disadvantages, such as having a loud engine, which gives off the military location to the enemy (The Guardian, 2020).

3.1.2.2 Level of autonomy

As they are building larger and stronger robots the idea of increasing BigDog's automation to the point where it could take actions by itself and unattended, would be a severe problem for soldiers working together. According to Shneiderman's (2020) paper this robot has a level of 4-5 when it comes to automation, where it executes suggestions if the human approves for example, traverse through different terrain or the robot suggests one alternative which is so walk besides the soldiers. Thus, the advantages we would gain from increasing the automation would maybe give the soldiers supplies and resources much quicker without needing a human to direct. On the contrary, decreasing would make the robot "less smarter" and would need more supervision and added tasks.

3.1.2.3 Explainability

BigDog's current explainability is the skill to give us feedback, the robot does respond to humans kickin it and will respond by trying to fall, but then get back up like a normal dog. It also gives us the kind of support and safety by being able to carry resources for us (Hagras, 2018, p. 29), promoting trust that it is built enough to handle the carry. Thus, it gives us some feedback the robot needs to respond more to our feedback. User wants the opportunity to provide feedback (Smith-Renner et al., 2020, p. 9), and then be able to respond to it not only the practical part, but understand tasks we wish the robot should do. Contrary, Smith-Renner et al. (2020, p. 1) argue that robots may also learn from experience, even without user feedback. So there is a possibility for BigDog to learn more tasks without feedback.

3.2.2 NAO robot

NAO is a robot used as an assistant by companies and healthcare centers to welcome, inform and entertain visitors. It is the first robot created by SoftBank Robotics and is famous around the world with its tremendous programming tool and he has become a standard in education and research. NAO has seven touch sensors, speech recognition, four microphones and speakers, as well as 25 degrees of freedom which enables him to move and adapt to his environment (SoftBank Robotics, 2020). His cameras allow him to recognize shapes, objects and humans, and he can also locate himself in space.

3.2.2.1 Advantage/disadvantage

As NAO is able to move around and interact with people he already has a very high degree of autonomy. A huge advantage with its abilities is how the robot is conducted for the purpose of using it to induce care-taking behaviours in children. Social robots like NAO can encourage care-giving behaviours, and are able to inform on an emotional level. It reinforces learning by teaching a paradigm, which allows children to gain deliberate practice with various skills by teaching specific skills to a robot (Phillips et al., 2016, p. 107). Although the disadvantages with the current state now is that it might still need human input to understand tasks, unless it provides machine learning and improves it skills through interactive feedback from humans.

3.2.2.2 Level of autonomy

NAO is able to move around by itself and interact with humans, he has according to Shneiderman's (2020) table a very high level of automation (around 9-10). The robot is able to decide all its movements and acts autonomously, but not ignoring the human. NAO can carry out many tasks without human intervention, although it still needs to be programmed to fit each specific user. A major issue that could occur, if its level of autonomy increased, the possibility of going rouge could be an option. The interaction with children would probably not be as safe as before and he could act on its own, but more automation could also mikt it more similar to interacting with a human. Thus, decreasing its autonomy would make the robot lose its purpose by teaching and learning by itself without supervision. But with less autonomy autonomy would make it less prone to errors and distraction.

3.2.2.3 Explainability

Since the high level of autonomy this robot is able to express some kind of safety feeling and fairness (Hagras, 2018, p. 29) towards the children by teaching, as well as being used in health care centers. NAO's current explainability promotes trust and acceptance, as Smith-Renner et al. (2020, p. 2) argues, a user's perceptions such as trust is shaped by the presence of the absence of feedback and explanations. Although it could encourage bias as Hagras (2018, p. 29) proclaims, how can we ensure that the AI system has not learned a biased view of the world based on data and objective function? How do we know that the children are being taught correctly or not? An explainability that NAO might need is their fairness, so the users can be comfortable knowing that their decision and their teaching is verified before used in educational settings. If the robot displays uncertainty and flaws, its activities will negatively impact users' perception (Smith-Renner, 2020, p. 9).

Appendix - Peer Review

Peer Review 1

In my peer review I received mostly good feedback, but I was encouraged to write a bit more about my own reflections and what I thought of the different definitions - which was really helpful. Therefore, I took the feedback into consideration when writing this assignment and I welcome the next person to review my assignment, as I wish to improve my writing skills more.

Peer Review 2

Mostly positive feedback to my assignment, although one of the tasks (2.3.1) were mainly quotes from the same article. My peer wished to see more discussion and my own reflections around this, which I took into consideration for my next iteration as well.

References

Amershi, Saleema, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, et al. “Guidelines for Human-AI Interaction.” In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, 3:1–3:13. CHI ’19. New York, NY, USA: ACM, 2019. https://doi.org/10.1145/3290605.3300233

Bratteteig, Tone and Guri Verne. 2018. “Does AI Make PD Obsolete? Exploring Challenges from Artificial Intelligence to Participatory Design.” n.d. Accessed September 3, 2020. https://dl.acm.org/citation.cfm?id=3210646.

Cambridge Dictionary. 2020. Make your words meaningful. *“Black Box”.* dictionary.cambridge.org/dictionary/english/black-box

Følstad, Asbjørn, and Petter Bae Brandtzæg. “Chatbots and the New World of HCI.” Interactions 24, no. 4 (June 2017): 38–42. https://doi.org/10.1145/3085558

Google AI. n.d. Accessed September 10, 2020. https://ai.google/responsibilities/facial-recognition/

Google HCI-060129. n.d. Accessed September 10, 2020. https://sites.google.com/a/nu.edu.pk/hci-060129/lectures-1/norman-s-7-principles

Grudin, Jonathan. 2009. “AI and HCI: Two Fields Divided by a Common Focus.” *AI Magazine* 30 (4): 48–48.

Hagras, H., Toward Human-Understandable, Explainable AI, Computer, 51, 9, 2018, 28- 36 https://ieeexplore.ieee.org/document/8481251

Honda. n.d. Accessed September 10, 2020. https://global.honda/innovation/robotics/ASIMO.html

Kocielnik, Rafal, Saleema Amershi, and Paul N. Bennett. “Will You Accept an Imperfect AI?: Exploring Designs for Adjusting End-User Expectations of AI Systems.” In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, 411:1–411:14. CHI ’19. New York, NY, USA: ACM, 2019. https://doi.org/10.1145/3290605.3300641

Luger, Ewa, and Abigail Sellen. "'Like Having a Really Bad PA': The Gulf Between User Expectation and Experience of Conversational Agents." In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, 5286–5297. CHI '16. New York, NY, USA: ACM, 2016. https://doi.org/10.1145/2858036.2858288

McCarthy, John. 2007. "What is Artificial Intelligence?". *Computer Science Department*. Accessed September 3, 2020. http://jmc.stanford.edu/articles/whatisai/whatisai.pdf

Outside Insight. 2020. *"How AI helps Spotify win in the music streaming world"*. outsideinsight.com/insights/how-ai-helps-spotify-win-in-the-music-streaming-world/

Oxford Reference. "Artificial Intelligence". "Robot" Accessed September 3, 2020. https://www.oxfordreference.com/view/10.1093/oi/authority.20110803095426960 https://www.oxfordreference.com/search?q=robot&searchBtn=Search&isQuickSearch=true

Phillips, E., Ososky, S., Swigert, B. and Jentsch, F. *Human-animal teams as an analog for future human-robot teams, Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol 56, Issue 1, (2016) pp. 1553 – 1557 DOI: https://doi.org/10.1177/1071181312561309

PwC. "Artificial Intelligence". Accessed September 3, 2020. https://www.pwc.no/no/teknologi-omstilling/digitalisering-pa-1-2-3/kunstig-intelligens.html https://www.pwc.no/no/pwc-aktuelt/ai-inntar-det-norske-markedet.html

Science Friday. The Origin Of The Word 'Robot'. Accessed September 10, 2020. https://www.sciencefriday.com/segments/the-origin-of-the-word-robot/

Shneiderman, B., *Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy*, arXiv.org (February 23, 2020). https://arxiv.org/abs/2002.04087v1

Smith-Renner, A., Fan, R., Birchfield, M., Wu, T., Boyd-Graber, J., Weld, D.S., and Findlater. L. 2020. No *Explainability without Accountability: An Empirical Study of Explanations and Feedback in Interactive ML*. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20). Association for Computing

Machinery, New York, NY, USA, 1–13. DOI: https://doi.org/10.1145/3313831.3376624

SoftBank Robotics. *"NAO".* Assessed 4. November 2020. https://www.softbankrobotics.com/emea/en/nao

The Guardian. *"US marines reject BigDog robotic pack horse because it's too noisy"*. Assessed 4. November 2020. https://www.theguardian.com/technology/2015/dec/30/us-marines-reject-bigdog-robot-boston-dynamics-ls3-too-noisy

The Universal Design Project. n.d. Accessed September 10, 2020. https://universaldesign.org/about

Thrun, Sebastian. 2004. "Toward a Framework for Human–Robot Interaction." https://doi.org/10.1080/07370024.2004.9667338

VisionAware. Seeing AI: Artificial Intelligence for Blind and Visually Impaired Users. Accessed September 10, 2020. https://visionaware.org/everyday-living/helpful-products/using-apps/seeing-ai-app/

Vox. n.d. Accessed September 10, 2020. https://www.vox.com/future-perfect/2019/4/19/18412674/ai-bias-facial-recognition-black-gay-transgender