Individual assignment in IN5480

Gwendolyn Borchsenius (gwendob)

MODULE 1

1. Concepts, definition and history of AI and interaction with AI

The history of Artificial Intelligence (AI) goes way back to the 20th century. The idea of robots had been represented by figures such as *The Tin Man* in the *Wizard of the Oz* and the robot in *Metropolis*, both movies from the first half of the 20th century. Alan Turing was one of the first persons that explored AI as a mathematical concept and raised the questions why computers could not solve problems and reason in the same way as humans can (Rockwell, 2017). Turing is often being considered as the father of AI. However, back then the technological possibilities were limited and computers could not store commands, they could only execute them. The term *artificial intelligence* was first used by mathematician John McCarthy in 1956 (Grudin, 2009). Even though *artificial intelligence* existed as a concept, computers were too expensive and technical underdeveloped to be able to execute these mathematical tasks. Some years later, the *Logic Theorist*, developed by Allen Newell, Cliff Shaw and Herbert Simon, was developed to imitate a human's way of problem solving. The Logic Theorist contributed immensely to the AI field, as it enabled problem solving from the technological perspective for the first time in history (Rockwell, 2017).

There are many explanations that seek to define AI. Turin's idea about AI, was a machine that could learn from its own experience by training a network of artificial neurons to perform specific tasks (Copeland). Interestingly, today's AI has not changed much from Alan Turin's concept of AI. However, the approach to AI has changed over the years. Whereas AI was mostly rule-based with algorithms and mathematical forms in the 20th century, the current approach is based on that the computer learns from its own experiences by processing large sets of data. Nowadays, techniques such as *Machine Learning (ML)* and *Deep Learning (DL)* are widely being used to recognise patterns in the data structures. Thus, the current AI software develops its intelligence from the data and the algorithms used for filtering the data. Nowadays algorithms, used for filtering the data, are more complex than in the 20th century and the system's behaviour can change over time as it learns more from the retrieved data (Bratteteig & Verne, 2018, p.2). Another possible, quite simple definition is that AI aims to mimic human intelligence and rational behaviour (Bratteteig & Verne, 2018, p.2). From my point of view, AI is a field within computer science that aims to imitate human decision making and behaviour. Machines are good at handling data and computing and they have become way more efficient in executing mathematical operations and calculations than humans. However there are certain tasks such as image recognition where humans still are superior. AI tries to imitate these human thinking structures in order to reach human reasoning.

One company that works with implementing AI in the automotive industry is BMW. On the website AI is presented as *fast, efficient and reliable* (BMW, 2019). Noticeably, all the words used to describe AI on the website, awake positive emotions. One of the main focuses on the website yields automated image recognition where AI is being used to compare an image with hundreds of other images in milliseconds. The website talks about AI as a promising service, easing the tasks of employees and customers.

iHuman, a movie about the use and interaction with AI, shows how far the development of AI has come, who stands behind the development and how AI might completely revolutionize our lives in the future. However the movie also reflects the uncertainty that exists around AI and raises serious concerns about privacy and power relations. Most people do not have insight in how AI is actually going to be used and for what purposes it will be used; it could be one of the biggest achievements of human history or it could also be the beginning of constant surveillance and tracking of human behaviour. The point is that whether we are going to benefit from AI or not depends on who is going to be in charge of decision making.

2. Robots and AI systems

Originally, the word robot comes from “rabota” which means work or labour (Deloitte Digital, 2018). The term robot was first used in the playwright “Rossum’s Universal Robots” (Deloitte Digital, 2018) in 1920. However it was not before the 1960s that the first actual robot came on the market, called “Shakey”, that could navigate itself through an environment (Deloitte Digital, 2018). As automation evolved over the years, robots became more important in areas where human work was not precise enough or where the work tasks were too dangerous to carry out for humans, such as in the automotive industry. Another interest of area concerns work where humans have failed to carry out tasks precisely due to stress, lack of sleep or distraction. This has especially raised interest in the automated and connected car driving field or high stress level jobs.

A very simple definition on a robot is provided by the Cambridge Dictionary: "*a machine controlled by a computer that is used to perform jobs automatically"* (Cambridge Dictionary). As the definition suggests, a robot is able to carry out certain tasks without direct human involvement, however controlled by a computer. The definition also implies that these specific tasks may be predetermined. This way, robots differ from AI as AI will seek to extend its neuronal network and learn from its behaviour. Another difference between robots and AI is that AI can be entirely software whilst robots normally are physical objects moving around, seeking to release human workload (Müller, 2020) . Another definition of robots describes robots as physical machines moving around, interacting with its surrounding, typically through sensors (Müller, 2020). However some robots may use AI software. Thus AI and robots as fields may also be combined. My idea of robots is that robots are programmed and produced in order to help human beings achieve bigger tasks, such as building a car or a house. The robot is not supposed to replace human work however it can support and release the workload for humans.

Singapore developed the robot dog "Spot" to encourage social distancing during the pandemic (Toh, 2020). The robot is currently being used in parks to help people keep distance from each other. When the dog locates people not holding the recommended minimum difference to each other, it politely reminds people to keep the distance for their own and other's safety. Spot is also outfitted with cameras to estimate the total number of people in the park at a certain time. The cameras do not track specific individuals. The dog's characteristics are similar to a real dog in the way he looks and moves around.

3. Universal Design and AI Systems

Universal design is about designing the environment in such a way that everyone regardless of their cognitive and physical abilities or impairments can be included in the society

(Digitaliseringdirektoratet, 2020). In Norway Universal Design is a legal requirement for both the public and the private sector. The way I understand this definition is that Universal Design is about making ICT available for everyone, creating a society that includes everyone and thus minimising the barriers for people with impairments.

From my point of view, AI has great potential to optimise tasks and improve effectiveness in the healthcare sector. Nurses, doctors and other healthcare workers are often exposed to immense levels of stress. As human capacity for perception is limited, stressful situations may cause an overflow of information input and as a result humans can perform poorly. AI is not exposed to stress hormones, raising temperature, lack of sleep and thus may perform better in certain stressful situations. AI may also be more accurate to execute certain tasks where human's motor skills are less inaccurate, e.g. surgery.

AI has also great potential to include people in our society in the future. In the educational sector for example AI may register where students are being challenged or struggling with the curriculum. AI then could provide individually customised tasks or learning activities for these students. AI may also generally engage students to learn through digital interactive learning lessons (Ashar & Cortesi, 2018). Thus AI might play an important role in including people, both young and old in our society in the future. However if not designed carefully, AI could also exclude certain user groups. For instance, facial recognition is not working equally well for people from different religions. The newer iPhones use facial recognition to unlock the phone and authorise transactions which causes problems for people from muslim countries that use burkas to cover their faces.

The third principle of the WCAG 2.1 guidelines deals with that web content must be understandable for the users. This includes that different users understand differently and use different amounts of time to receive and understand information. I am not certain about whether you could say that a machine can understand. The way I see it, machines can interpret input. For example a programme written in python can be programmed to interpret the user input as an int. However there are certain human actions such as emotions and gestures that are not being interpreted by a machine unless we program the machine to do so.

This leads me to my conclusion that machines do not understand as humans do, however they can be trained to understand human behaviour.

4. Guideline for Human-AI-interaction

Guideline number 13 is about learning from user behaviour. Different users react and behave differently in one and the same situation. The guideline's goal is to make AI be able to recognise these human behaviours and reactions and learn from them in order to improve. By doing so, the AI software is supposed to adapt to the user's personal behaviour.

The Interaction Design Foundation (Wong, 2020) illustrates ten guidelines for interfaces and their implementation. The framework is similar to the AI-guidelines in the way that they both want to make the user as satisfied as possible with the machine or technology. They are also similar in the way that the technology is illustrated as a tool to satisfy and help the user experience a better interaction. However they differ in adapting towards the user. Once an interface is released and on the market, it is hard to make changes. The AI guidelines in contrast seek the AI software to continuously adapt and learn from the user as the software generates more data about the user's personal behaviour.

MODULE 2

1. Characteristics of AI-infused systems

In the lecture (Følstad, 2020), four main characteristics of AI were discussed. These are learning, improving, black Box and fuelled by large data sets. Learning relates to the AI system's capability to adapt to and learn from the user's behaviour. This learning process is dynamic. Improving relates to the systems ability to improve accuracy and response over time by learning from user input, feedback and errors. Black box relates to the characteristics of hiding certain processes and actions that are going on in the background. In the litterature, this is also being referred to as *behind the scenes processes* (Amershi et. el, 2019*)*. Finally, AI-infused systems learn and improve from large data sets. The quality of the data sets is also essential for the accuracy of the AI-infused system (Karahasanovic, 2020). For instance, little representative datasets can cause the AI-infused system to behave racist as it was the case with Google's vision AI (Kayser-Bril, 2020).

The literature available through the course, discusses additional key characteristics of AI-infused systems. For instance, Amershi et al. (2019) describe three main key characteristics of AI-infused systems: inconsistency , uncertainty and behind the scenes personalisation. Inconsistency relates to the difficulty to predict behaviour and outcome. In addition to inconsistent user behaviour, AI is constantly changing due to deep learning processes that in general make it hard to predict behaviour on both sides. Lack of reliability can lead to poor user experience and in the worst cases lead to dangerous outcomes. Further, AI-infused systems can give false negative and false positive outcomes as they often perform under uncertainty. This uncertainty can again lead to errors. Finally, Amershi et al. (2019) discuss behind the scenes personalisation as a key characteristic of AI-infused systems relating to the for the user hidden activities going on in the background. These background activities often happen on behalf of the user. For instance, the authors discuss personalisation of content that can either match the user's preferences but when poorly aligned also hide for the user important content. Other characteristics of AI-infused systems are described in

Kocielnik et al. (2019). The authors argue that AI mechanisms such as natural language understanding and object recognition are probabilistic, however almost never completely accurate. Further, AI behaviour may pose transparency issues as most of the AI algorithms work below the surface, hidden for the user.

One example of an AI-infused system is the speech recognition system “Siri” of Apple. The slogan “Siri does more than ever. Even before you ask.” (Apple, 2020) gives the impression that AI is involved to recognise user intentions. The slogan also hints to the in the lecture omitted characteristics of *black box* where certain actions and processes are not shown to the user. Personally, I never use Siri so much except for when driving the car. However, I get the strong feeling that Siri constantly listens and works in the background even though I am not aware. Sometimes, Siri accidentally turns on even though I was not intending to interact with Siri, leaving me confused. This is an example for the *black box* or *behind the scenes* principle as Siri constantly works in the background without me being aware of what Siri is filtering and analysing. Further, this is an example of that AI-infused systems are probabilistic in the way that Siri works quite well most of the time but sometimes turns on even though I was not intending to interact with Siri.

2. Human-AI interaction design

Kocielnik et al. (2019) examine the impact of the user’s expectation towards the AI-infused system on user acceptance. Studies show that low expectations towards usability decrease user satisfaction and willingness to continue using the product. The authors argue that AI poses additional challenges impacting user satisfaction and acceptance as they almost never operate completely accurate. However, most users expect error free user experience, leading to a conflict between the capabilities of AI-infused systems and the user’s expectations. Therefore, Kocielnik et al. argue that the end user’s expectations should be shaped prior to use in order to minimise this gap. Further, the authors discuss how pre-use adjustment of user expectations can impact positively on transparency and improve trust. Most commonly, transparency techniques are used to explain why certain AI decisions have been made.

The main argument in Amershi et al. (2019) paper is that conventional guidelines and design principles are not applicable when designing AI-infused systems. This is due to the natural behaviour of AI mechanisms as described earlier. More specifically, unpredictability and inconsistency challenge use of conventional design principles. For instance, the design principle of error prevention cannot simply be applied to AI-infused systems as errors are common in AI algorithms due to unpredictable and inconsistent behaviour. Further, the authors argue that variability in AI designs due to different forms of interaction and capabilities, challenge intuitive and effective design of AI-infused systems. Thus, the common accepted design guidelines and principles cannot simply be transferred to AI-infused systems. However, shared understanding and standards for design are important in order to achieve reliable and consistent AI technologies. The main challenge is therefore to develop reusable guidelines and design principles that yield all different types of AI-infused systems in order to improve user experience and build trust among users when interacting with AI-infused systems.

Guideline G1 "Make clear what the system can do. Help the user understand what the AI system is capable of doing." (Amershi et al., 2019) relates to letting the user understand what the system he or is interacting with can or cannot do. The authors discuss displaying all metrics of an activity tracker app in order to help the user understand what the app measures and how. This is to help shaping user expectations and improve user experience and interaction. Drawing back on Siri as an example of an AI-infused system, I would say guideline G1 is met by viewing an overview of what the user can ask or do with help of Siri. The system also gives examples on how to interact with the system in order to fulfill a certain task. Noticeably, the overview does not show what Siri cannot do. This could be improved in order to also show the user what Siri is not capable of.

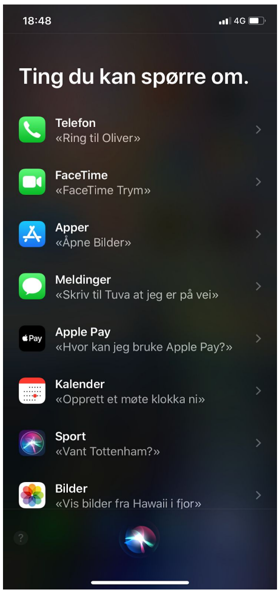

Guideline G5 “Match relevant social norms. Ensure the experience is delivered in a way that users would expect, given their social and cultural context.” (Amershi et al., 2019). The example being discussed by the authors relates to using semi formal voice in voice assistants that spells out “okay” rather than for example “k”. Drawing back to Siri as an example, guideline G5 is not as easy to identify as G1. I asked “What’s up?” to see how Siri handles casual, informal smalltalk. Interestingly, Siri answers in a neutral, quite formal tone while at the same time being funny and informal. This gives the impression that Siri even can handle slang and humour.

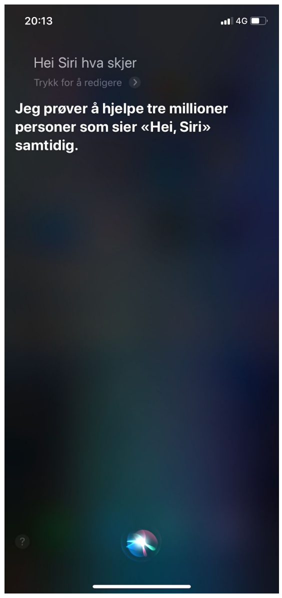

3. Chatbots / conversational user interfaces

The key challenges of AI-infused systems described in Amershi et al. (2019) and Kocielnik et al. (2019) also apply to Chatbots. Shaping expectations can be crucial in order to make the user understand what the chatbot is capable of and what not in order to make the user understand the chatbot's limitations. As Luger and Sellen (2016, p. 5292) state, Conversational Agents often fail to bridge the gap between user expectation and system operation. This can be due to poor mental models, giving the user unrealistic expectations of what the chatbot is capable of doing or not. Different contexts of use can challenge chatbots in addition (Luger & Sellen, 2016, p. 5289). Chatbots are used for different purposes such as in customer service, for assistance or for social matters. Different purposes pose different context of use and thus need for different types of chatbots. For instance, a chatbot assisting older people to send or receive messages is most likely to be designed differently than a chatbot for customer service. Thus, agreeing on standards and general accepted design principles can be challenging as the users and use contexts differ. Luger and Sellen (2016) also discuss the challenge of supporting the ongoing user engagement. Ideally, a conversation should result in a “binding hypnotic effect” (Luger & Sellen, 2016, p. 5295) that keeps the user wanting to continue to interact with the system. Currently, AI-infused conversational agents are far from this goal. Personally, I often stop interacting with chatbots before achieving my goal because I experience the chatbot to be too cumbersome.

Drawing back on the guidelines G1 and G2;

> G1: "*Make clear what the system can do. Help the user understand what the AI system is capable of doing.*" (Amershi et al., 2019)
>
> 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*." (Amershi et al., 2019)

I think both guidelines should generally be followed no matter what the context of use. It is always important for the user to know what the limitations of the system is. Revisiting Kocielvik's argument, it seems legit to know of the AI-infused system's limitations prior to use. Guideline number one can naturally also be applied to chatbots in order to shape expectations and improve user experience. The second guideline relates to not only making clear what the system can do but also making clear how well the system can carry out the task. The user should be informed about potential errors, preferably prior to use, in order to align the user's expectation towards the chatbot.

MODULE 3

According to Phillips et al. (2016) robots have become more interactive team members, aiming to help humans achieve their goals rather than tools used by humans. However, robots still lack humanlike teams and social competences. Human-animal teams work as analog for human-robot interaction and can therefore help to refine mental models of users, making it easier to understand the robots limitations and capabilities. Through their article the authors provide a taxonomy for human-robots collaboration where the authors examine three different types of human robot interactions: physical teams, emotional teams and cognitive benefits.

Physical teams relate to the human-animal relationships that aim to provide physical benefits by either replacing, multiplying or augmenting a human's physical needs. For example, horses, camels, mules and elephants have traditionally been used for cargo. This area is now extensively investigated within the robots field. The military robot Big Dog for instance is designed to carry cargo and relieve a soldier's load. Big Dog is also designed for different or uncertain terrain.

Emotional teams relate to the social component of human-animal relationships and the ability to express and perceive emotions or build social relationships of trust. One of the most common forms of human-animal relationships is that of companionship, typically provided by domesticated pets. Pets can decrease stress, anxiety and lower blood pressure. Recently, researchers investigate how robots can provide companionship and comfort in the same way as animals do. For instance, Paro, a seal looking robot, has been designed to provide companionship for older people and alleviate depression.

Cognitive benefits relate to animals providing humans with additional sensory information that can improve decision making, especially in extreme situations where human cognition is limited. Animals, due to their sensitive sensory organs, can serve to replace cognitive capabilities by detecting risks to humans. Until the late 1980s, canaries were used to detect poisonous gases in the coal mines. In the area of robots, nano UAV robots were used by soldiers in Afghanistan to replace human cognitive capabilities. The sensors of the robots could provide additional sensory information to the soldiers, important for scouting and reconnaissance and thus improve the decision making process.

Shneiderman (2020) examines levels of automation and criticises the conventional view that a system either is fully automated or fully controlled by humans. Instead, the author argues that the goal should be to both ensure high levels of human control and high levels of automation, ideally complementing each other. Drawing back on the idea of robots acting as interdependent teammates, rather than independent from humans, Shneiderman (2020, p.5) argues that human performance can be dramatically increased. This is presented through the two dimensional human-centered artificial

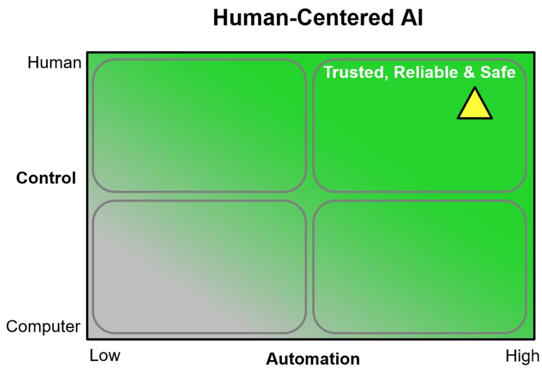

Figure 1: Human Centered AI (Shneidermann 2020, p. 7)

intelligence framework (see figure 1), suggesting that designers are not forced to choose between either automation or human control but can combine high levels of both.

Shneidermann argues that the desired goal is often in the upper right quadrant where both automation and human control is high. These systems usually require complex decision making and involve creative, intelligent thinking. The lower right quadrant, with high automation and low human control, describes highly automated systems that require rapid action. Examples for such high-automated systems are anti-lock brakes, airbag systems and defensive weapon systems. The upper left corner describes tasks with high human control and low automation. For instance, playing instruments and cooking requires mostly human control. Even though certain systems can assist humans in fulfilling the task (e.g. timer for cooking noodles), it is still up to the person to carry out the task. The lower left quadrant yields simple automated tasks such as clocks or land mines.

All of the earlier described examples for human-robot teams fit into the upper right quadrant as the robots all are highly automated but are interdependent on humans. Big Dog, Paro and the UAV robots all act as teammates, aiming to help humans achieve certain goals. UAV robots such as the Black Hornet Nano require high automation as they need to receive GPS data and navigate themselves in the air as well as send relevant signals to the operator's terminal. At the same time, reading and analysing the received data requires human control. I would therefore argue that the Black Hornet Nano is in the middle of the upper right corner as the system is completely responsible for analysing the surroundings and sending relevant signals to the operating soldiers, however final decisions are still depending on the soldiers. Increasing the automation level of the detector, could save the operator from having to carry around a display for receiving data from the device. However, this would require full trust of the operator as any errors could lead the entire troop into an ambush. Paro is another example of a highly automated system that also requires human control. The seal looking robot is not independent of humans, acting itself but will react when humans interact with it. The output of the robot however, is out of human control. There are no switches or buttons to manipulate the seal. Increasing automation in this case could give a more realistic look and feel of the robot. Currently, engineers investigate how robot dolphins could replace real dolphins being held captive in zoos and theme parks. Finally, Big Dog is also located in the upper right corner as the robot can walk itself as well as investigate the surrounding terrain. However, I suppose the robot does not itself know where to march without human input. Increasing automation in this case could improve the robot's independence.

Hagras (2020) discusses the concept of explainable AI (XAI) to enable responsible and trusted AI and in order to meet user acceptance and increase trust towards the system. XAI implies that actions and processes, enabled through AI should be easy to understand by humans. Especially critical decisions should not be hidden for the user and be transparent. Furthermore, XAI systems should be able to justify their actions and follow logical decision making, including predicting their future actions. As Big Dog operates in extremely critical situations where wrong decisions could harm the entire troop or even cause death, trust towards the system becomes increasingly important. This requires that Big Dog does not make decisions such as walking away from the troop with the entire load on it's own. Instead, the actions should be transparent and follow rational decision making. For instance, a causal, logical decision would be if the robot would change it's direction due to difficult terrain. This should also be transparent to the troop in order to increase trust towards the system.

REFERENCES

Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., ... & Teevan, J. (2019). *Guidelines for human-AI interaction*. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (paper no. 3). ACM. Retrieved from https://www.microsoft.com/en-us/research/uploads/prod/2019/01/Guidelines-for-Human-AI-Interaction-camera-ready.pdf

Apple (2020). *Siri does more than ever. Even before you ask*. Retrieved from https://www.apple.com/siri/

Ashar, A. & Cortesi, S. (April 26, 2018). *Why inclusion matters for the future of AI*. Retrieved from https://digital.hbs.edu/artificial-intelligence-machine-learning/inclusion-matters-future-artificial-intelligence/

BMW Groups (July 15, 2019. *Fast, efficient, reliable: Artificial intelligence in BMW Group Production.* Retrieved from https://www.press.bmwgroup.com/global/article/detail/T0298650EN/fast-efficient-reliable:-artificial-intelligence-in-bmw-group-production?language=en

Bratteteig, T. & Verne, G. *Does AI make PD obsolete? Exploring challenges from Artificial Intelligence to Participatory Design.* Retrieved from https://dl.acm.org/doi/pdf/10.1145/3210604.3210646

Copeland, B.J. *Alan Turing And The Beginning Of AI.* Retrieved from https://www.britannica.com/technology/artificial-intelligence/Evolutionary-computing

Deloitte Digital (August 14, 2018). *A brief history of robots*. Retrieved from https://www.deloittedigital.com.mt/en/news/2018/8/a-brief-history-of-robots-450

Digitaliseringsdirektoratet (2020). *Kva er universell utforming?*. Retrieved from https://uu.difi.no/kva-er-universell-utforming

Følstad. A. (September 22, 2020). *Interaction with AI. Module 2* [pdf]. retrieved from https://www.uio.no/studier/emner/matnat/ifi/IN5480/h20/Undervisningsmateriale/interacting-with-ai-2020---module-2---session-1---handouts.pdf

Grudin, J. (2009). *AI and HCI: Two Fields Divided by a Common Focus*. Retrieved from https://aaai.org/ojs/index.php/aimagazine/article/view/2271

Hagras, H. (2020), *Toward Human-Understandable, Explainable AI*, Computer, 51, 9, 2018, 28- 36. retrieved from https://ieeexplore.ieee.org/document/8481251

Kayser-Bril. N (April 7, 2020). *Google apologizes after its Vision AI produced racist results.* retrieved from https://algorithmwatch.org/en/story/google-vision-racism/

Kocielnik, R., Amershi, S., & Bennett, P. N. (2019). *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 (paper no. 411). ACM. Retrieved from https://www.microsoft.com/en-us/research/uploads/prod/2019/01/chi19_kocielnik_et_al.pdf

Luger, E., & Sellen, A. (2016). *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 (pp. 5286-5297). ACM. Retrieved from https://www.microsoft.com/en-us/research/wp-content/uploads/2016/08/p5286-luger.pdf

Müller, V. (April 30, 2020). *Ethics of Artificial Intelligence and Robotics.* Retrieved from https://plato.stanford.edu/entries/ethics-ai/#AIRobo

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. Retrieved from https://doi.org/10.1177/1071181312561309

Rockwell, A. (August 28, 2017). *The History of Artificial Intelligence*. Retrieved from http://sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence/

Toh, M. (May 8, 2020). *Singapore deploys robot 'dog' to encourage social distancing.* Retrieved from https://edition.cnn.com/2020/05/08/tech/singapore-coronavirus-social-distancing-robot-intl-hk/index.html

Wong, E. (August 10, 2020). *User Interface Design Guidelines: 10 Rules of Thumb.* Retrieved from https://www.interaction-design.org/literature/article/user-interface-design-guidelines-10-rules-of-thumb

APPENDIX

The feedback that I received on the different iterations was mostly positive. However, there were some mistakes in writing that I corrected. I have also been advised to describe the guidelines in the second module more specifically and come up with more examples.