IN5480 Individual assignment fall 2019

,"I do not see why [the computer] should not enter any one of the fields normally covered by the human intellect, and eventually compete on equal terms. I do not think you can even draw the line about sonnets, though the comparison is perhaps a little bit unfair because a sonnet written by a machine will be better appreciated by another machine."-Alan Turing 1949 (Grudin, J 2009).

Innhold

Module 1

Introduction

In module 1 am going to explore three different topics within the vast field of AI:

-Concepts, definition and history of interaction with AI.

-Human Robot Interaction.

-Universal Design and Interaction with AI

These three topics will be in three sections. In the first section I will look at different concepts, definitions and take a deep dive into the history of interaction with AI. In the second section I will look at robots and AI systems where I will discuss the relation between the two. Is there a difference between "a robot" and "an AI"? And for the last section I will look at the topic of universal design in relation to AI.

Concepts, definition and history of interaction with AI

The history of AI

The history of AI and how it came to be had a lot of different contributors. At the very beginning of this text I added a citation from a paper Alan Turing wrote in the London Times in 1949 where he talks about the possibility of computers entering the fields normally covered by the human intellect and eventually competing on equal terms. One year later Turing's "Computing Machinery and Intelligence," were published, just as Claude Shannon's "Programming a computer for Playing Chess" and Isaac Asimov's *I, Robot* (Grudin, J 2009). In the late 1940s and the early 1950s the Macy Foundation sponsored a series of conferences where they brought together leading mathematicians, psychologists, and social scientists where they discussed topics such as cybernetics (Grudin, J 2009). As we can see there were a lot of different papers written and conferences held that helped shape the concept of AI. The actual term *artificial intelligence* didn't appear before in 1956 as a call for participation in a workshop written by John McCarthy. The field has ever since attracted more and more attention, and evolved to the point it is as of today.

Definitions of AI

In AI Magazine in 1987 Roger C. Shank wrote: “What AI is depends heavily on the goals of the researchers involved, and any definition of AI is dependent upon the methods that are being employed in building AI models. Last, of course, it is a question of results.”.

Later in the article Shank gives his own definition of AI: “AI is the science of endowing programs with the ability to change themselves for the better as a result of their own experiences” (Roger C. Shank 1987).

If we look at another example closer to our time in an article written by Internet Society they state that traditional AI refers to “an artificial creation of human-like intelligence that can learn, reason, plan, perceive, or process natural language.” (Internet Society 2017). They also write that Artificial intelligence is further defined as “narrow AI” or “general AI”. In which they say Narrow AI is designed to preform specific tasks within a domain where as General AI is hypothetical and not domain specific but can learn and perform tasks anywhere.

Both definitions come from communities with very technical backgrounds. To get another perspective I looked into the humanitarian perspective regarding the topic of AI. I read a text written by *Psychology Today* where they defined AI as “the ability of computers to perform human-like feats of cognition including learning, problem-solving, perception, decision-making, and speech and language.” (Psychology Today 2019). From reading this I could see that the definition of AI were similar across the different fields computer science and psychology.

Looking at these different definitions I can see some similarities and some differences between them all. Based on my own understanding of the topic my definition of AI is as follows: “Artificial Intelligence refers to a computer being able to simulate understanding based on cognition”. To explain my definition, I must separate it into two parts, the first being the “simulating understanding” and the second being “based on cognition”.

So, what do I mean by “simulating understanding”? Well if we look at how the computer works, it takes in input in the form of data. It does not have the same understanding(cognition) of its surroundings as us humans. But if it is able to simulate understanding to a certain level as to where we humans perceive the computer as intelligent or “thinking as a human” that is when it has successfully created an artificial intelligence. If it had real understanding, we would not have had the need to create the term Artificial Intelligence because it would just be Intelligence.

To explain the second part “based on cognition” we need to look at the word cognition. Cognition is defined by Lexico (Oxford dictionary 2019) as *“The mental action or process of acquiring knowledge and understanding through thought, experience, and the senses.”*. So, what I mean by “based on cognition” is that for a computer to simulate understanding it must be able to process the information it takes in as input and learn from it. It will not do it the same way as us humans but for it is to be perceived as an “Intelligent being” it also must simulate cognitive thinking to some degree. If it is able to do this it will directly affect the perception of understanding.

Apple

According to thomasnet.com apple is the biggest artificial intelligence software company in the world. I did some more research into apple and found “Apple machine learning journal” where they have documentation on their projects relating to AI. If we take Siri as an example, from the documentation on their page about Siri (links to the page will be in the section “Further reading” down below) we can see that Siri is presented as a service to help users do different tasks while using Apple’s products such as the Iphone. This goes for other “AI software” they have as well. They do different services to support the user in using their products, this can be for example text recognition or facial recognition.

Her

For a movie to describe I have chosen the movie “Her” directed by Spike Jonze. The way in which the human interaction with AI is portrayed in this work is through a fictional world where humans have created artificial intelligent virtual assistants. The movie shows a relationship between a lonely writer and his virtual assistant. What is interesting about this movie is how their relationship develops and how real feelings get involved when it gets hard for the main character to differentiate a computer from reality. The movie plays with the concept of the “Turing test” (Stanford Encyclopedia of Philosophy 2016) which I find really interesting.

Robots and AI systems

"I can't define a robot, but I know one when I see one."

~ Joseph Engelberger

The word “robot”

The very first appearance of the word “robot” was in 1921 in Karel Capek’s play R.U.R., or Rossum’s Universal Robots. “Robot” comes from the Czech for “forces labor” (Matt Simon 2018).

Definitions of Robot

There were a few different definitions of “robot” that I looked at and I ended up choosing Hans Peter Moravec’s definition:

“Robot, any automatically operated machine that replaces human effort, though it may not resemble human beings in appearance or perform functions in a humanlike manner.” (2019)

What he means by this is that a robot is a machine that will work independently from humans with the help of automation doing tasks previously done by humans. Even though it replaces human effort it does not mean it needs to resemble human appearance or that it is able to perform functions in similar way to humans.

I also chose the definition Leo Center for Service Robotics uses in their article “DEFINING ROBOTS AND ROBOTICS”:

> "A robot is an actuated mechanism programmable in two or more axes with a degree of autonomy, moving within its environment, to perform intended tasks. Autonomy in this context means the ability to perform intended tasks based on current state and sensing, without human intervention."

But later in the article they write:

> This is the most frequently used definition of a robot. You may sense that this definition does leave room for interpretation.
>
> -a degree of autonomy - What amount is a degree?
>
> -Intended tasks - intended by whom? The programmer, the end user, the robot?
>
> -moving within its environment - What qualifies as 'environment'? Does a virtual agent moving through its software environment qualify?
>
> -without human intervention - If a human is needed to 'intervene' does that disqualify an object as a robot? (Leo Center for Service Robotics 2019)

What I found interesting about this definition is that they picked it apart and criticized it themselves. This definition is a very general definition that as Leo Center for Service Robotics says, “leaves room for interpretation”. What the definition says is that a robot is some sort of mechanism that can be programmed so that it will work autonomously (being able to do tasks without human intervention).

With my own definition I will try to simplify it: “A robot is a machine that is programmed to be able to perform tasks without the help of humans if the right conditions are met”. What I mean by this definition is that a robot is programmed to operate/perform a task independently, but under some conditions robots cannot operate without any intervention from humans. An example of this would be to compare a robot in a factory vs a lawn mower robot. In the factory the perfect conditions are met so the robot is able to work autonomously without any human intervention. In the case of the lawn mower robot the robot would need a human to intervene if for example a tree in the garden fell unto the lawn. The robot is blocked by the tree and no longer able to perform its task before a human removes the tree. In e perfect world every robot would be able to do the task they are programmed to do, but as of today that is not the case in many scenarios.

Robots and AI relations

I will state that robots and AI has similarities but is not the same. They are similar in the way that they want to be able to perform tasks on a "human level" independently from human interaction, but robots do not necessarily need to inhabit signs of artificial intelligence to perform autonomous tasks. The same goes for the AI, it does not have to be autonomous to be conceived as an artificial intelligence. Both parts have limits that can be mitigated to some degree by combining both entities, if we create a robot with artificial intelligence. If we look at my definition of a robot "A robot is a machine that is programmed to be able to perform tasks without the help of humans if the right conditions are met". This is achievable without any AI being a part of it. The same goes for AI that can function as a software not needing the autonomous machinery that would classify it as a robot.

The Mavic 2

I wanted to present a commercial robot the Mavic 2. This robot is a drone and it moves in aerial-space by the help of propellers. It has sensors to help it dodge obstacles. The interaction with the Mavic 2 goes through their own app software. You control it with your phone connected to the Mavic's controller. You can either remote control it or you can set it to trace mode. In this mode the user will decide on an object by drawing a square around an object on the phone within the Mavic's "vision". The Mavic will then follow this object as it moves keeping a set distance.

Universal design and AI systems

Definition of universal design

"Universal Design is the design and composition of an environment so that it can be accessed, understood and used to the greatest extent possible by all people regardless of their age, size, ability or disability." This is a definition by the National Disability Authority on universal design (2019). What it means is that universal design is about the creation of design that includes everyone regardless of factors that would differentiate them from the "norm". My understanding of universal design is that it is a set of guidelines for inclusivity. Everything we create should be available to every person and by incorporating universal design we get one step closer to an inclusive society.

Potential of AI with respect to human perception, human movement and human cognition/emotions

With all the work that is being done within the field of AI there is a lot of potential with respect to these different topics and there are a lot of examples. Affectiva has presented Human Perception AI. The possibilities here lie in the AI abilities to detect emotions that can be used in social robots, or you can detect attention that can save people from road accidents. With movement you now have artificial intelligent assistants with voice recognition helping users do tasks by only speech, making them able to do tasks that previously required movement. When it comes to cognition/emotions I want to bring back the example of social robots but you also have therapeutical robots like PARO the seal that help people with dementia.

The potential of AI for including and excluding people.

AI has a big potential when it comes to including users. Tasks that users were not able to do excluding them from using certain software/products have been made easier by the help of AI. For example, by making it possible for a user to do commands by voice instead of using "hands on" interaction. But AI also carries the potential of excluding people. One big example of this is some of the facial recognition software that could not detect the faces of African-Americans or other dark skin people. To mitigate the exclusion in AI we have to create AI with inclusion in mind. The principles of universal design should also be in the center when we create AI-products.

Module 2

Introduction

In module 2 I am going to look at three topics in the same way as I did in module 1. These topics are listed below:

Characteristics of AI-infused systems.

Human-AI interaction design.

Chatbots / conversational user interfaces

Characteristics of AI-infused systems.

"Systems that have features harnessing AI capabilities that are directly exposed to the end user."

(Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., ... & Teevan, J. 2019)

Below I will try to identify and describe some key characteristics of AI-infused systems.

Learning- The AI should always evolve and adapt to the user by learning. This can be different habits, patterns and so on. This will make the AI dynamic and more efficient for the user.

Improving- The Ai should always improve. Mistakes are inevitable but it is important for the AI to learn from the mistakes and not repeat the same ones.

Black Box- The AI cakes in input but what happens between the input and output is very opaque. AI behaves independently and without the need of being able to show everything that is going on "behind the scenes".

Fueled by large data sets- The AI learns by acquiring knowledge and act upon this acquired knowledge. This knowledge is data either given to the AI or collected by the AI stored as large data sets.

Unpredictable- The AI will change over time by learning and improving, but this can also be cause for confusion for the end user. The AI can suddenly change the output it gives form certain input overtime changing its behavior. It can be hard to predict for the user witch direction the AI changes.

Bixby, Samsung's virtual assistant.

For an AI infused system I wanted to write a little about my Samsung smartphone. My Samsung phone has a virtual assistant called Bixby. Bixby can help you do different tasks like setting alarms, making a phone call, answer questions and so on. When we look at the characteristics, I identified in the section above we can see that a lot of them relates to Bixby.

Learning- Bixby learns from your usage of the smartphone so it can for example make better customized news suggestions. You can also train Bixby so that it will better understand your voice. This also goes into the improving characteristic.

Fueled by large data sets – Bixby collects data about the phone's user that will fuel its future behavior.

Human-AI interaction design.

My main takeaway from Kocielnik et al. (2019) was that the interaction with AI will differentiate based on the expectations the users have of the AI in question, if we can change the expectations the user has it will impact the way in witch the user interact with the system. I also did not expect the results they got from the study "that focus on high precision rather than high recall of a

system performing at the same level of accuracy can lead to much lower perception of accuracy and decreased acceptance"(Amershi, 2019).

My main takeaway from Amershi et al. (2019) was that with all the new AI-infused systems we need to rethink the way in witch we design these systems. We need to incorporate new ways of designing that before was not necessary. In this article they presented 18 design guidelines for human-AI interaction.

I have selected to write about G7 -support efficient invocation, and G8 -support efficient dismissal.

Bixby really adheres to G7. You can use the command "hey Bixby" to wake it up and then follow up with a desired task. This is mostly used in all virtual assistants in smartphones.

In the way Samsung created the phone: "Samsung galaxy s8" it deviates from G8. It is hard to dismiss the bixby services. They created a button on the phone that can only be used to activate bixby. It is no possible to change it to do anything else, the only thing you can do is turn it completely off so that when you press the button nothing happens.

I think the G7 guideline already works pretty well within the Bixby system, but Samsung can learn a lot from the G8 guideline. They can use the G8 guideline to improve their system by making their system more customizable for the user and "free". They really push their own software and makes it hard for the users that owns a Samsung smartphone to dismiss their own services, Bixby is one of the examples of this, they don't make it easy to turn off the Bixby services on their phones if the user would wish to do so.

Chatbots / conversational user interfaces.

When it comes to conversational user interfaces it has its technological challenges. The systems are not yet perfect and probably won't be for many years. You also have the challenges of speech recognition, but this is also technical. What I got from reading the paper by Luger & Sellen (2016). Was that you need to design these systems with the expectations of the user in mind. You must adapt to the technological limitations of the conversational user interfaces by changing the expectations of the users, and this is not an easy task. The challenge is to teach the user of the limitations of the system, and how they can use it. Often the systems get a negative feedback because users have the expectations not matching the capabilities of the system.

The article by Følstad & Brandtzaeg (2017) mentioned many challenges when it comes to the design of chatbots. First, we must change the usual role of the designer when we talk about "designing an interface" since chatbots is a whole new way of interacting with a system. As they say in the article "we need to move from seeing design as an explanatory task… …to an

interpretational task" Følstad & Brandtzaeg (2017). We need to move from the design of interfaces to a service-oriented design. Another challenge is to look away from the usual one user, one interface design and look at it as a network when designing. The designers need to look at the multi-agent aspect of interaction. New technology also can create new digital divides and biases that the designer needs to take into consideration and try to deviate from. And the last challenge I want to mention is to understand conversational processes. For chatbots to work efficiently as a natural language interface we need to design the actual language and conversations that will be used by the systems.

One of the big issues with current chatbots and conversational user interfaces is that they don't meet the expectations of the users. They have limitations. This can cause the service they provide to break down rather quickly. Most of us have experienced some sort of chatbot or conversational interface to get stuck because they can't complete a task you have asked of the system. By adhering to Guidelines G1-Make clear what the system can do, and G2 -Make clear how well the sytem can do what it can you can help shape the user's expectations of the system. These guidelines will help to make the users understand what the system is capable of doing and its limitations. By doing this the user will keep h*s commands or tasks within the systems "capability barriers", making the interaction between Human and AI as efficient as possible to the extent of what the system is capable of.

Module 3

Introduction

In this module I will look at some taxonomy and examples of human-robot collaboration in the article "Human-Animal Teams as an Analog for Future Human-Robot Teams: Influencing Design and Fostering Trust". (Philips, 2016). While I am looking at these examples, I will describe their level of autonomy as described by Endsley, 2004. I will then go into and reflect upon the advantages and disadvantages if we were to decrease or increase their current level of autonomy.

Augment/Extend Physical Capabilities.

The first example I have chosen is the example of the ARMAROB's FRIEND robot that goes under the taxonomy of augment/extend physical capabilities. The robot FRIEND (Functional, Robot arm with user-frIENdly interface for Disabled people) is a robot developed by ARMAROB research consortium (Philips, 2016). The goal of the robot is to provide enhanced

dexterous manipulation for people suffering from paralysis or other skeletal-muscular disorders (Philips, 2016)

The robot itself works together with its user in a what Philips calls a team. The human is still in control of the robot, but the robot helps extend his/her capabilities as categorized by the type of taxonomy. So, if we look at the levels of autonomy as described by Endsly we see that it is not close to the level of full automation. I would argue that this robot has a level of automation that goes within what Endlsy call the 'Action support/tele-operation' described as "Computer aids in doing each action as instructed". The taxonomy goes under 'extension' of physical capabilities and not fully replacement of physical capabilities so there is still a level of human interaction in the process. I would not say it has a level of Batch processing the robot does no fully operate by itself but aids in different tasks, but it still physically aids in tasks so i would say it is a higher level then SA support that's is why I ended up at Action support/tele-operation.

So, what if we were to increase or decrease the level of automation of the FRIEND robot? If we look at decreasing the level of autonomy, I only see disadvantages. The reason why I think this is, is because of the user group of the robot. The users of the robot are people with physical limitations or disabilities. If we were to decrease the level of autonomy the robot would no longer assist the user in the tasks needed to be done. I think the level of autonomy the robot has is the minimum requirement to be able to be useful for the different tasks the user group of the robot are facing.

Then what if we look at increasing the level of autonomy? I can see some advantages and disadvantages when it come to the increasing in level of autonomy. I can see the possibilities where by increasing the level of autonomy the robot would be able to assist users with even greater physical disabilities. But if we were to increase the level of autonomy to a level where the robot starts making the decision it goes into the cognitive abilities and out of the taxonomy of extending only physical capabilities. I see the advantage of raising the level to Batch process where the robot can completely carry out tasks for the user and not only aid in doing physical tasks. My reasoning for this is that it will then reach out to a greater amount of users with even more severe physical disabilities, but it will also stay in within the taxonomy in question. I believe that if we were to raise the level of autonomy to an even greater level there may arise ethical issues where the users themselves get the feeling of losing control of their own lives.

Multiply Cognitive Capabilities

For the next example I have chosen to talk about drones. In the image example used by Philips she used an image of a bird and a Nano robot looking like a small drone. Within the taxonomy of multiply cognitive capabilities she talks about the humans using animals to multiply their cognitive capabilities by for example detecting poisonous gas.

A drone can reach places humans cannot and give them a different cognitive understanding of a situation, multiplying their cognitive capabilities. Drones are usually on a level of automation of Batch processing. They completely carry out tasks but are still controlled by a human operator. Drones can differentiate in the levels of autonomy, they can increase to a level of supervisory control depending on the type of drone and the task set by the human operator of the drone.

So, with the drone example the level of autonomy can differentiate greatly. Now if we look at the decreasing and increasing of autonomy levels there are some advantages and some disadvantages. If we first look at the decreasing of autonomy the drone the disadvantage of this would be that we lose a lot of functions that supports the operator in controlling the drone. For example, usually drones have sensors that prevents them from crashing into objects forcing them to stay at a certain distance as to not crash into other objects. If we were to decrease the level of autonomy to Manual control the drone would lose the capability to a certain degree override the operator's control. This will make it easier for the operator to crash the drone. Also if we lose the level of SA support the drone will be hard to operate as the operator will not get the information needed from the drone to be able to operate it at a distance. The drone would then act more as a "remote controlled toy" that you can by in the toy store. If the level of autonomy goes below the level of SA support, I do not believe it would stay in the taxonomy of multiplying cognitive capabilities.

Now if we look at increasing the level of autonomy, I would like to use military drones as an example. The advantage of increasing the level of autonomy Would be that the drone could do tasks independently, for example spying and gathering information on an enemy. By increasing the levels of automation the military would be able to delegate less people to these tasks as they are now done by drones and focus their efforts and manpower elsewhere. A big disadvantage when it comes to increasing the level of autonomy is the ethical question when it comes to military drones. When you have robots that multiply our cognitive capabilities we need to know if we can trust the data presented to us by these robots. What factors does a drone bring in when deciding on who is a possible target when spying on the enemy? And if we look at drones that carries out bomb missions it would be problematic to increase the level of autonomy to full automation or even a level above shared control. It would no longer only multiply our cognitive capabilities, Then it arises the question if we are going to give the robots power to decide upon someone's life.

Appendix 1

Changes after feedback from iteration 1

I ended up with not making any changes after I got the feedback from another group member. I got positive feedback on the two first subjects of the assignment. The only critique I got was that

the part on universal design, the third topic, were a little short and could have had a little more content. I were already on the maximum number of pages and since I got such a good feedback on the two other topics I did not want to change them. Or shorten them down to add more to the third topic. I was limited by the boundaries of the assignment.

Appendix 2

Changes after feedback form iteration 2

In the feedback on iteration 2 he wished that I would fix the structure of the text by adding table of contents, so I added this for the third and final iteration.

References

Module 1

Apple Machine Learning Journal. (n.d.). Retrieved September 17, 2019, from Apple Machine Learning Journal website: https://machinelearning.apple.com/

Artificial Intelligence. (n.d.). Retrieved September 17, 2019, from Psychology Today website: https://www.psychologytoday.com/basics/artificial-intelligence

Artificial Intelligence & Machine Learning: Policy Paper. (n.d.). Retrieved September 17, 2019, from Internet Society website: https://www.internetsociety.org/resources/doc/2017/artificial-intelligence-and-machine-learning-policy-paper/

Bostrom, N., & Yudkowsky, E. (2014). The ethics of artificial intelligence. In K. Frankish & W. M. Ramsey (Eds.), The Cambridge Handbook of Artificial Intelligence (pp. 316–334). https://doi.org/10.1017/CBO9781139046855.020

Cognition | Definition of Cognition by Lexico. (n.d.). Retrieved September 17, 2019, from Lexico Dictionaries | English website: https://www.lexico.com/en/definition/cognition

Defining robots and robotics. (2015, November 24). Retrieved September 18, 2019, from LEO Center for Service Robotics website: http://www.leorobotics.nl/definition-robots-and-robotics

Design, I. (2019, August 27). How to Recognize Exclusion in AI. Retrieved September 25, 2019, from Medium website: https://medium.com/microsoft-design/how-to-recognize-exclusion-in-ai-ec2d6d89f850

Estacio, A. (2018, September 10). Buyers Guide – DJI Mavic 2 Pro In-depth Review. Retrieved September 25, 2019, from DJI Guides website: https://store.dji.com/guides/mavic-2-pro-review/

Everything You Ever Wanted To Know About Robots. (n.d.). Wired. Retrieved from https://www.wired.com/story/wired-guide-to-robots/

Grudin, J. (2009). AI and HCI: Two Fields Divided by a Common Focus. AI Magazine, 30(4), 48–48. https://doi.org/10.1609/aimag.v30i4.2271

Intelligence. (n.d.). Retrieved September 17, 2019, from Psychology Today website: https://www.psychologytoday.com/basics/intelligence

Oppy, G., & Dowe, D. (2003). The Turing Test. Retrieved from https://seop.illc.uva.nl/entries/turing-test/

Robot | technology. (n.d.). Retrieved September 18, 2019, from Encyclopedia Britannica website: https://www.britannica.com/technology/robot-technology

Siri. (n.d.). Retrieved September 17, 2019, from Apple website: https://www.apple.com/siri/

Top Artificial Intelligence (AI) Software Companies in the USA and Internationally. (n.d.). Retrieved September 17, 2019, from https://www.thomasnet.com/articles/top-suppliers/ai-software-companies

What is Universal Design | Centre for Excellence in Universal Design. (n.d.). Retrieved September 25, 2019, from http://universaldesign.ie/What-is-Universal-Design/

Zijderveld, G. (n.d.). Our Evolution from Emotion AI to Human Perception AI. Retrieved September 25, 2019, from https://blog.affectiva.com/our-evolution-from-emotion-ai-to-human-perception-ai

Module 2

Amershi, S., Inkpen, K., Teevan, J., Kikin-Gil, R., Horvitz, E., Weld, D., … Bennett, P. N. (2019). Guidelines for Human-AI Interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, 1–13. https://doi.org/10.1145/3290605.3300233

Følstad, A., & Brandtzæg, P. B. (2017). Chatbots and the New World of HCI. *Interactions*, *24*(4), 38–42. https://doi.org/10.1145/3085558

Kocielnik, R., Amershi, S., & Bennett, P. N. (2019). Will You Accept an Imperfect AI?: Exploring Designs for Adjusting End-user Expectations of AI Systems. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, 1–14. https://doi.org/10.1145/3290605.3300641

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 - CHI '16*, 5286–5297. https://doi.org/10.1145/2858036.2858288

Module 3

Endsley, Mica R.. Designing for Situation Awareness: An Approach to User-Centered Design, Second Edition CRC Press. 2011 (chapters 2 and 10)

Phillips, E., Schaefer, K. E., Billings, D. R., Jentsch, F., & Hancock, P. A. (2016). Human-animal Teams As an Analog for Future Human-robot Teams: Influencing Design and Fostering Trust. *J. Hum.-Robot Interact.*, *5*(1), 100–125. https://doi.org/10.5898/JHRI.5.1.Phillips