IN5620- Interaction with AI and autonomous systems

Individual assignment

Rameen Khan (rameenk)

Fall 2022

Content

Module 1 ..... **1**

Module 2 ..... **7**

References **11**

Appendix **13**

Module 1

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

History of AI

Some may argue that the earliest idea of AI originates from myths about robots told by the ancient Greeks and that modern AI stems from "classical philosophers' attempts to describe human thinking as a symbolic system" (Lewis, 2014). However, we usually talk about Artificial Intelligence as a term first used after the 1956 Dartmouth Summer Research Project on Artificial Intelligence hosted by John McCarthy – an American mathematician and logician. Just before this, in 1950 Alan Turing presented an article on the notion of machines being able to think in a way of simulating human beings (Turing, 1950). This paper opened for other mathematicians, psychologists, and social scientists to discuss the areas around artificial intelligence before the workshop in 1956 which is said to be the conference that started it all. After that, AI flourished as a major research field filled with visions of AI becoming "autonomous and intelligent, competing with people on equal terms" (Grudin, 2009). In the 60s/70s these visions became a promise for many people struggling to see the light during the uncertainties of their time with the cold war and nuclear wars hanging over their heads (Grudin, 2009).

Definitions

Bratteteig & Verne (2018) – "*AI is a subfield of computer science aimed at specifying and making computer systems that mimic human intelligence or express rational behavior, in the sense that the task would require intelligence if executed by a human"*. As this article writes about how AI might make participatory design obsolete, the definition focuses more on the human aspects and the imitation of human behavior (often in more social settings).

McCarthy (2007) – "*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 ".* This definition focuses on the science of the technology in comparison to the definition stated above.

IBM (2020) – *"Artificial intelligence leverages computers and machines to mimic the problem-solving and decision-making capabilities of the human mind."*. I find this definition more straightforward and understandable on a ground level thought it lacks some technological aspects to cover both the social and technical parts of AI.

My definition – Developing an AI system includes the process of combining large data sets with processing algorithms. Accordingly, I find it important to include the notion of a *process* in the definition of AI as well as an emphasis on both the technological aspects and human influences. A mix of the three definitions stated above might form a suitable definition for AI: *"Artificial intelligence is the process of utilizing science and engineering to have computer systems independently mimic or simulate human behaviour and intelligence, including rational problem-solving and decision-making capabilities".*

Human autonomy vs machine autonomy

Both human and machine autonomy refers to the level of self-governing in various important tasks and choices in different contexts. This implies that "taking on significant choices increases autonomy while offloading significant choices to others decreases autonomy." (Formosa, 2021). Human autonomy can in addition be comprehended as a moral principle, meaning beliefs or guidelines about what is considered right or wrong. As "moral agents", we humans need to address these principles and respect the human autonomy by consenting – something that's not part of respecting machine autonomy. This also insinuate that only human autonomy may be subject to nudging, manipulation, and surveillance (Formosa, 2021).

The issue of respecting human autonomy doesn't go without the concern that a growing machine autonomy (i.e., humans transfer decisions and decision-making powers to machines and AI) may lead to loss of some human autonomy. Machine autonomy refers to a systems capability "of achieving its goal within a defined scope without human interventions while adapting to operational and environmental conditions." (Ezenkwu & Starkey, 2019). As the situation is today neither humans, nor machines, can be considered fully autonomous as they

are not capable of performing all possible tasks in all possible circumstances (Ezenkwu & Starkey, 2019). Still, humans have the upper hand as developing autonomous systems is dependent on our understanding of how we achieve such a high degree of human autonomy.

Automatic and autonomous

Automatic refers to systems developed by a set of predefined parameters and operates without human or other external interventions. Their actions are based on predefined heuristics. Autonomous on the other hand is more like a "driverless car" rather than the automatic "self-driving" one. This means that autonomous systems are self-governing as described earlier, indicating that they learn based on data and as a result adapts to their dynamically changing environments.

Brief article review

The article "Does AI make PD obsolete? Exploring challenges from Artificial Intelligence to Participatory Design" written by Tone Bratteteig and Guri Verne, writes about how AI and the continuous growth of AI over time challenges participatory design (PD). PD attempts to understand the potential technology has to change current use situations by including current and future users to participate in decision-making regarding the technologies' design and use. As AI here is described as an opportunity for personalized service, "knowing" and "understanding" you, to make your life easier, the article questions what the role of PD would be as it many similar ways emphasize design for users' needs. AI's unpredictability may be seen as a threat to PD as the data it learns and builds on often can be insufficient or biased and thus evolve to perform both better or worse (Bratteteig & Verne, 2018). This can cause users to lose the control over their activities, something PD on the other hand works to maintain. Though many of the challenges presented can somewhat be solved with a balance, AI will, as stated in the article, because of technical characteristics also pose complex challenges that reach beyond the scope of any PD project.

Contemporary company working with AI

Facebook (Meta) is a well-known social networking platform and there is no doubt that they use AI to meet and improve user needs. For instance, they use machine learning algorithms to automatically translate posts on the users' feeds. Their most known AI algorithms are still probably the ones that analyzes likes/dislikes, photos, videos, and articles that the user has interacted with to recommend new or similar content to the user. By then filtering and giving feedback to the system, the AI is further developed and becomes more accurate. Though many of the users still are unaware of exactly how the AI works, Facebook (or Meta) presents many of their tools and research on their webpage such as how they are "creating personalized and meaningful interactions" with Conversational AI, or how they are working with integrity by "building AI solutions to keep people safe on social platforms" (Meta AI, 2022). They talk about AI in a reassuring manner making it seem like just good: *"Meta creates*

breakthrough technologies and advances AI to connect people to what matters and to help keep communities safe." (Tech at Meta, 2022).

Human interaction with AI in fiction

Black mirror is maybe one of the most popular Netflix series when it comes to dystopian futures and AI. They have several episodes usually portraying how an AI is used for a beneficial purpose in the beginning of the show and then ends with how it might have been overestimated. For instance, in season 4, episode 4 – called "Hang the DJ", a dating app measures the love compatibility of people and give them an exact number in time of how long their "relationship" will last. The idea of a partner with guaranteed compatibility is a dream for many, especially when you don't have to do anything to get the compatibility rating. However, the show takes a turn when the two protagonists discovers that they are meant for each other despite the app telling them otherwise. This episode raises questions around AI as a replacement for personal decision-making, and whether you should trust them blindly or make your own choices as well.

1.2 Robots and AI systems

History

The word robot stems from the Czech word 'robota', meaning 'forced labor'. It was first used in a play by Karel Capek in 1921. Here, the robots did all the work the humans didn't want to do and were portrayed as human bodies with no feelings or intellectual life (Wallén, 2008).

Definitions

Robots.ieee (2020) – "*A robot is an autonomous machine capable of sensing its environment, carrying out computations to make decisions, and performing actions in the real world.". This definition focuses on the autonomy of robots, emphasizing the lack of human intervention needed to function. The robot here has to sense and interact with the physical world around it.*

ISO (2012) – *"actuated mechanism programmable in two or more axes with a degree of autonomy, moving within its environment, to perform intended tasks".* Similar to the definition above, this definition focuses on the robot's physical interaction with its surrounding environment but specifies a robot as something moving which may restrict some examples of robots one may have, such as smart home systems.

My definition – I find it relevant to include the robot's interaction with its physical surrounding as it was mentioned in both the definitions above, as well as adding its impact on said environment: *"A robot is a physical device that is aware of and can interact with its surrounding environment to perform actions and thus make changes in the real world.".*

Relation between AI and robots

Looking at the definitions stated above and earlier, AI and robots are similar as they both should be able to interact in a certain environment and therefor function somewhat autonomously without human interventions. They differ, however, on their tangibility-level as the definitions of robots all point to robots being a physically existing machine, while AIs don't necessarily have to have a specific form at all – they are usually software. Another difference is the emphasis the AI definitions include on mimicking human behaviour and intellect, whereas the definitions on robots don't express the need for such level of rational problem-solving and intellect.

Example: contemporary robot

As described on their homepage, the robot Spot is "an agile mobile robot that navigates terrain with unprecedented mobility, allowing you to automate routine inspection tasks and data capture safely, accurately, and frequently." (Boston Dynamics, 2022). This "robot dog" is equipped with cameras, microphones, and various sensors to measure and inspect machines and tools used to analyze potential malfunctions or threats. Spot became part of the Norwegian company Aker BP in 2020 and is working on an oil rig in the Norwegian sea to make offshore operations safer, more efficient, and more sustainable (Aftenposten, 2020).

1.3 Universal Design and AI systems

Definition

"Universal design is design that's usable by all people, to the greatest extent possible, without the need for adaption or specialized design." (The UD Project, 2022).

This definition of universal design consists of three equally important parts; the first being universal design is design that's usable by all people, meaning that whether you have a disability or not shouldn't matter, this environment, product, space, software, etc. should be designed for everyone to use.

The second part; to the greatest extent possible, implies that there is no "standard human" - everyone is different with different user needs that has to be accounted for. This includes for instance upper/lower body movement, strength, and/or sensation, vision, cognition and memory, speech and/or communication, etc.

The last part: without the need for adaption or specialized design refers to inclusiveness across generations and/or health conditions. No one should feel like they are left out or discriminated because of their situation whatever it may be.

The potential of AI

There are several specially developed AIs to help aid towards inclusion with respect to human perception, human movement, and human cognition/emotions. However, many of them are so embedded in our daily life that many people might not even realize that the AI is there. For example, is use of texting on videos a good way to include people who are hearing impaired. This includes videos we watch on YouTube, the ones we encounter on our Facebook feed, the

series, and movies we watch on Netflix, and even our lectures. Similarly, use of speech-to-text in shape of ex. Voice-assistants such as Siri or various home systems such as Google Home or Amazons Alexa might work well to include people with reduced possibility to use their hands to get the same information or perform the same tasks as someone who easily can type or press buttons.

This way, AI has a great potential to include more people, and as technology and software is getting more available by the day, people have easier access to the content. Chatbots is an example of AI that's being used by more companies on their websites to allow increased inclusion for example for people who find it difficult or stressful to talk to people in-person. AI also have the trait of continuous learning which means that the more the user utilizes the AI the more it can adapt and customize to the users needs.

However, it is important to remember that all good unfortunately can be used for bad. The same goes for AI. A very criticized example of how AI has the potential to exclude is the use of facial recognition. Firstly, there are many incidents where biased data-input in AI have resulted in the AI failing to recognize people with darker skin tone. To take it to the next level, facial recognition is used in countries like China to target specific minorities by recognizing certain traits that's "typical" for these people (Harwell & Dou, 2020). Not only is this exclusion, but also a violation of human rights.

"Understand" and "Understanding"

Even thought it may seem like its little to no difference between the two terms, there is a small gap where they can differ in meaning. I think of "understand" as a term used when one is given some information, which then gets processed and further results in an action based on the processed information. I would hence claim that machines can understand as they usually can receive input which they in turn process based on defined rules and algorithms before they give us an output. However, I find the term of "understanding" more fluid as we humans have a complex way of making sense of the information we receive. The context in which we are situated are not taken into consideration or understood by the rules and algorithms of the machines and thus restrict the machines "understanding".

1.4 Guideline for Human-AI interaction

Human-AI Interaction Design Guidelines

Human-AI Interaction Design Guideline 4: "Show contextually relevant information".
This guideline refers to the importance of giving the user relevant information at the right moments in the task or environment they are situated in as this can impact expectations and/or needs. Netflix has a lot of functionalities that fulfills this guideline, for instance the "resume where you left of" functionality as well as getting a recommended movie or series as a "next up" after finishing a movie or series.

HCI design guidelines

Don Normans seven design principles are meant to help designers communicate with the users by making products and services more efficient, effective, and delightful to use (Medium, 2020). This differs a bit from the Human-AI Interaction Design Guidelines as they don't focus as much specifically on the interaction with users but include a broader societal perspective as well as specific system requirements. Norman's guidelines are however similar to the Human-AI Interaction Design Guidelines on several aspects, firstly with their weightage on feedback – giving users continuous responses to their interaction with the system. The principle of mapping is also illustrated in both guidelines as the user's understanding of what the system can do and what effect it will lead to can result in more intuitive use.

Module 2

2.1 Characteristics of AI-infused systems

Key characteristics

As described in the assignment introduction, AI-infused systems are "systems that have features harnessing AI capabilities that are directly exposed to the user" (Amershi et al., 2019). These systems are characterized by a range of functionalities as well as challenges. Amershi et al. identifies, based on AIs constant evolving on interaction with the users, characteristics such as uncertainty, implying that these systems may cause more errors (both what they refer to as "false positives" and "false negatives") than systems with a deterministic nature, inconsistency, where changes in the system may look small to the users but can have great consequences for its output, and "behind the scenes personalization" (or automated personalization), which might hide some information from users. This is an important topic for discussion among designers as such systems are difficult to predict and hence design for. In addition to the designers' struggles to predict these systems outcomes, Kocielnik et al. (2019) and Yang et al. (2020) describe how another key characteristic of AI-infused systems is the inaccuracy in expectations with the users towards the system. This may have root in the fact that the users aren't properly informed about what the system actually is capable of doing and so it's a higher chance for unintended results.

The characteristics discussed by Amershi et al. is quite similar to what is described in both the first lecture of module 2 and in Kocielnik et al., especially the learning-aspect. Because these AI-infused systems are learning they are also constantly improving. However, this doesn't happen without uncertainties as the algorithms used to develop these systems aren't precise and therefore inconsistent as Amershi et al. pointed to, as well as the data the systems are "taught" from are external and develops over time (Yang et al., 2020). Følstad (2022), also refer to the characteristics of the "black box", meaning that much of the processing the AI

systems do happens in a non-visible way (Kocielnik et al., 2019), as well as they're driven by big sets of data which is further driven by the users' interactions. This may however cause the systems to learn stuff that results decisions not intended by the developers (Yang et al., 2020).

Netflix as an AI-infused system

Netflix is an AI-infused system that I am quite familiar with. Its recommendation system uses AI to learn about patterns in your watched content on the platform to present movies and series that may suit your preferences better. This fits the learning characteristic mentioned in the numerous articles above. Functionalities like "liking" or "disliking" content adds to this learning process and is also part of the data sets in the large data sets needed to improve and hence make more accurate suggestions over time, presenting personalized content to the users. All this obviously happens in the "black box" and is not visible to the users. However, when users have the possibility to help in the process by liking or disliking they are more aware of what the system is capable of.

2.2 Human-AI interaction design

Amershi et al. (2019) takeaways

In this article, Amershi et al. introduces 18 guidelines for human-AI interaction and the design of AI systems. These are based on previous studies of humans and AI, as well as other relevant guidelines, to create better human-AI interactions and also maintain them to ultimately avoid negative outcomes that may occur when interacting with an AI technology. To better address the amount of challenges related to human-AI interaction the guidelines are categorized into four categories based on the moment in the user-interaction-journey where they are relevant – you have the *initial* phase, *during interaction*, *when wrong*, and lastly *over time*. These guidelines were also validated through some iterations of evaluation but has potential to be further developed and refined in future work as the "current technology landscape is shifting towards the increasing inclusion of AI in computing applications" (Amershi et al., 2019).

Kocielnik et al. (2019) takeaways

Unlike Amershi et al. who focuses on presenting the guidelines to better, more human-centric AI-infused systems, Kocielnik et al. aim at discussing and steering the users' expectations to these systems. Because of how the systems can be presented it's very common for users' expectations to exceed what the AI is capable of doing and hence why they often have a hard time accepting an imperfect AI-system. To counteract these expectations it is proposed, as in Amershi et al., to be consistent in providing the user with necessary and relevant information related to the systems capabilities as this can reduce errors and improve user experience. The authors also designed three "expectation adjustment techniques" with the purpose of

managing expectations and help users recognize AI systems not being perfect or to their ideal. These are *accuracy indicators*, *example-based explanations*, and lastly *control*. They used these on two versions of an "email scheduling assistant" to explore whether they can improve user satisfaction and acceptance of imperfect AIs.

Netflix and design guidelines

G12 – Remember recent interactions

When finishing a movie or a series Netflix remembers what these are and use them to provide related suggestions. In addition, the content you haven't finished watching yet will be located under a "continue to watch" area and will show where you last stopped watching. This could be beneficial for easier accessibility to continue watching something instead of having to search up whatever it was that you were watching and finding the part where you left of.

G13 – Learn from user behavior

The recommendation system present personalized content to users based on their actions, including what they usually watch in addition to the feedback they give when hitting the "like" or "dislike" button. Now they also have added to this functionality by giving the user the option to "double-like" with two thumbs instead of one for content you like even better as well as the regular "one-thumb" option and "thumb-down" dislike button. This makes it easier to distinguish between what the user really like and what they can enjoy watching and hence give better recommendation results. You will often find such suggestions beneath a heading that reads something like "Top picks for [your name]".

Bender et al. (2021) on language models

This article is a critical discussion on possible risks with large language models and whether they are necessary or not while considering various costs associated with these models, such as environmental and financial costs, and opportunity costs. The authors define risks leading to potential harm like abusive or foul language leading to stereotypes, denigration, and an increase in extremist ideologies. This is however not something that can be "fixed" with quantity to even out for more diverse content as the authors point to possible overrepresentation of, for instance, various user groups. These and other factors like misinterpretations, controversial imaging, and biases gets picked up and used as learning data in language models and becomes of risk to make substantial harm. The article propose future NLP research to focus on considering these risks and costs, and being cautious of the data sets that go into the learning data to reduce the occurrence of such risks and their potential harms.

2.3 Chatbots/ conversational user interfaces

Challenges with designing chatbots

From both the lectures and the articles of Bender et al. (2021) and Koecelnik et al. (2019), natural language understanding is pointed out to be one of the key challenges with chatbots as it is a core component in these systems. Natural language is considered a challenge because of the abusiveness and misinterpretations that can happen because of, for instance, language variations or different contexts in which a message is expressed that chatbots aren't capable of deciphering. This is also reflected in Følstad & Brandtzaeg (2017) as they discuss how our (the users') desired interaction with chatbots is that of which resembles human-human interactions, or human conversation. This makes the users' expectations way higher than what is the reality, which discussed earlier also gets enhanced by the lack of communication of the systems' capabilities. Building on this, Luger & Sellen (2016) describes how users with somewhat technical skills are more capable to "see beyond artificial humanlike qualities", as opposed to the users with less knowledge who rather got frustrated "[…] leading them to question the 'intelligence' of the system, indicating that user expectations of CAs should be scaffold through more considered revelation of system intelligence through design" (Luger & Sellen, 2016, p.5294). Følstad & Brandtzaeg also point to the importance of designers' understanding of the users' motivation for using chatbots to solve their challenges more effectively (Følstad & Brandtzaeg, 2017). Designing a chatbot, where the understanding and conversation is crucial, means that we must change our focus towards designing services instead of designing objects.

Adherence to guidelines

Guideline G1 ("Make clear what the system can do")

Because many users aren't aware of what questions the chatbots are capable of answering, they tend to ask questions they won't get a desired answer to. If chatbots were to adhere to this guideline it would raise the users' awareness and provide input more suited to what the chatbot can answer. This is also part of expectation maintenance to give the users realistic expectations of the system, including what it can't do and relevant feedback when an error or misunderstanding occurs.

Guideline G2 ("Make clear how well the system can do what it can do")

Adhering to this guideline would better manage the expectations of the user and hence lessen the frustrations that may occur when interacting with a system you expect too much from. As chatbots often are designed to mimic human conversation, it's no wonder why some would treat it as such. By informing the users about how well the system can respond, one could reckon them being more forgiving of its flaws.

References

Amershi, S., Vorvoreanu, M. & Horvitz, E. (2019). Guidelines for human-AI interaction design. (Microsoft research blog). Hentet fra https://www.microsoft.com/en-us/research/blog/guidelines-for-human-ai-interaction-design/

Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collison, P., Suh, J., Iqbal, S., Bennett, P.N., Inkpen, K., Teevan, J., Kikin-Gil, R. & Horvitz, E. (2019, May). *Guidelines for human-AI interaction design.* In CHI 2019, Glasgow, Scotland.

Bender, E.M., Gebru, T., McMillan-Major, A. & Schmitchell, S. (2021, March) On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21*), New York, NY, 610-623. https://doi-org.ezproxy.uio.no/10.1145/3442188.3445922

Boston Dynamics (2022). Spot. Hentet fra https://www.bostondynamics.com/products/spot

Bratteteig, T. & Verne, G. (2018). Does AI make PD obsolete? exploring challenges from artificial intelligence to participatory design. *Proceedings of the 15th Participatory Design Conference: Short Papers, Situated Actions, Workshops and Tutorial, 2(8),* 1-5. https://doi.org/10.1145/3210604.3210646

Ezenkwu, C.P., Starkey, A. (2019). Machine Autonomy: Definition, Approaches, Challenges and Research Gaps. In: Arai, K., Bhatia, R., Kapoor, S. (eds) Intelligent Computing. CompCom 2019. Advances in Intelligent Systems and Computing, vol 997. Springer, Cham. https://doi.org/10.1007/978-3-030-22871-2_24

Formosa, P. Robot Autonomy vs. Human Autonomy: Social Robots, Artificial Intelligence (AI), and the Nature of Autonomy. *Minds & Machines* 31, 595–616 (2021). https://doi.org/10.1007/s11023-021-09579-2

Følstad, A. & Brandtzaeg, P.B. (2017). Chatbots and the new world of HCI. *Interactions, 24*(4), 38-42. https://doi-org.ezproxy.uio.no/10.1145/3085558

Grudin, J. (2009). AI and HCI: Two fields divided by a common focus. *The AI magazine, 30(4),* 48-57. 10.1609/aimag.v30i4.2271

Harwell, D. & Dou, E. (2020, 08.12). Huawei tested AI software that could recognize Uighur minorities and alert police, report says. *The Washington Post*. Hentet fra https://www.washingtonpost.com/technology/2020/12/08/huawei-tested-ai-software-that-could-recognize-uighu r-minorities-alert-police-report-says/

IBM (2020, 03.06). Artificial Intelligence (AI). Hentet fra https://www.ibm.com/cloud/learn/what-is-artificial-intelligence

ISO (2012). Robots and robotic devices – Vocabulary. Hentet 13.09.22 fra https://www.iso.org/obp/ui/#iso:std:iso:8373:ed-2:v1:en

Johansen, P.A. (2020, 16.02). -Her er robothunden Spot. Nå har den fått fast jobb i Norge. *Aftenposten.* Hentet fra https://www.aftenposten.no/norge/i/vQP3rw/her-er-robothunden-spot-naa-har-den-faatt-fast-jobb-i-norge

Kocielnik, R., Amershi, S. & Bennett, P.N. (2019, May). *Will You Accept an Imperfect AI? Exploring Designs for Adjusting End-user Expectations of AI Systems.* In CHI 2019, Glasgow, Scotland.

Lewis, T. (2014). A Brief History of Artificial Intelligence. Hentet fra https://www.livescience.com/49007-history-of-artificial-intelligence.html

Luger, E. & Sellen, A. (2016, May). *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 (CHI '16). New York, NY, USA, 5286-5297. https://doi-org.ezproxy.uio.no/10.1145/2858036.2858288

McCarthy, J. (2007). What is Artificial Intelligence? Hentet 13.09.22 fra http://jmc.stanford.edu/articles/whatisai/whatisai.pdf

Meta AI (2022). Research. Hentet 13.09.22 fra https://ai.facebook.com/research/?utm_source=hp#research-areas

Robots.ieee (2020). What is a Robot? Hentet 13.09.22 fra https://robots.ieee.org/learn/what-is-a-robot/

Tech at Meta (2022). Artificial Intelligence. Hentet 13.09.22 fra https://tech.fb.com/artificial-intelligence/

The UD Project (2022). What is Universal Design? Hentet fra https://universaldesign.org/definition

Turing, A.M. (1950). Computing Machinery and Intelligence. *Mind, 59(236),* 433-460. http://www.jstor.org/stable/2251299

Wallén, J. (2008). The history of the industrial robot. *Technical report from Automatic Control* Hentet fra https://www.diva-portal.org/smash/get/diva2:316930/FULLTEXT01.pdf

Yang, Q., Steinfeld, A., Rosé, C. & Zimmerman, J. (2020, April). *Re-examining Whether, Why, and How Human-AI Interaction Is Uniquely Difficult to Design.* In CHI 2020, Honolulu, HI.

Appendix

As I didn't get any feedback as to how I could approve module 1 I didn't make any changes to that part. However, to module 2 I was told that I could add more to the challenges of designing chatbots, maybe by adding another source. Based on this I chose to add aspects of the article written by Luger & Sellen (2016) which aims at understanding interactional factors affecting everyday use of people using Cas (Conversational Agents). This content was easy to integrate in the already written text and worked as great support for the points already made.