My Schultheiss - mysc

IN5620 Individual assignment 1 - Second iteration of Module 1

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

The history of AI

The term "AI" or "Artificial intelligence" was according to the Council of Europe (2022) first used by John McCarthy from the Institute of Technology of Massachusetts and Marvin Minsky from Carnegie-Mellon University. They defined it as "The construction of computer programs that engage in tasks that are currently more satisfactorily performed by human beings because they require high-level mental processes such as: perceptual learning, memory organization, and critical reasoning.".

Still, the technology behind AI is known to have been developed some years before being defined in a term, with Alan Turing and the universal machine. Turing took this further towards AI by raising the question of the possible intelligence of a machine and the Turing imitation test, where humans are to guess whether they are talking to a human or a machine (Council of Europe, 2022).

Definitions of AI

In the encyclopedia Britannica (2022), Copeland defines AI as "The ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings."

In order to understand the definition of AI, it may be necessary to define what intelligence is. For example, Copeland in the encyclopedia Britannica explains that almost all human behavior is called intelligent, which does not apply to other animals, even in situations where they act with complicated behavior. This is because we define our intelligence as more than instinctive, while other animals act on instinct and based on one ability. In contrast, humans operate on a combination of many diverse skills (Copeland, 2022). In my understanding, this impliest hat digital computers or computer-controlled robots act in such a way that we associate them with ourselves by regarding their actions as mirrored and learned from us.

Another way of defining AI was done by Simmons and Chappel in 1988; "The term artificial intelligence denotes behavior of a machine which, if a human behaves in the same way, is considered intelligent." The definition has its origin in machine engineering in the late 80s. It can, in my opinion, give more clarity as to whether the machine is intelligent or merely acts in a way that reminds us of human behavior. In addition, it can be used to debate whether we should consider/be careful when discussing AI and possibly personifying computers.

Bratteteig and Verne (s. 1-2, 2018) describe AI as "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." This definition has its origin in the field of Participatory design and once more creates a clear divide between actual human intelligence and the seeming imitation of it.

My definition would be that AI is a broad term for describing machines that can perform tasks similarly to how a human using its intellect would do it.

Similarities and differences between human and machine autonomy
Some aspects of human autonomy and behavior are transferred to machines. The design of AI attributes is based on human techniques for problem-solving, resulting in machines calculating based on a framework for implementing a known but inexact method of solving a problem (Simmons & Chappell, 1988, p. 15).

Simmons and Chappell write that the term human intelligence covers the ability to solve specific groups of problems and to come up with solutions for new issues. Their article is from 1988, but I believe that their arguments broadly cover today's systems. The systems cannot learn or create new concepts through experience or rearrange pieces of knowledge in novel ways (Simmons & Chappel, 1988, p. 15). This shows that AI computers have less autonomy, as they have restricted possibilities for solving problems compared to humans.

Similarities between the concept of automatic and the concept of autonomous
Both terms cover machines acting without human assistance. In a digital matter, they both have pre-set parameters programmed that work as a starting point for their work (Matteson, 2019). The word “autonomy” means that an individual/thing makes descicions on their own, and therefore the concept of autonomous may often be used in situations where one is actually talking about something automatic.

Review of the article “The Ethics of AI Ethics: An Evaluation of Guidelines” by Thilo Hagendorff
I find this article very interesting but also scary. It alarms me to read the results Hagendorff presents from his research of current AI Ethics frameworks and actual implementation. As the evaluation is given, it gets clearer and more apparent that there are severe problems with AI use throughout all (or almost all) parts of our society and that it is critical to enforce more covering rules and regulations as soon as possible.

Another aspect is how AI is used and abused in politics and governing people in different countries. For example, not only is there a “scoring system” in China to ensure inhabitants’ “good living”, but automation of propaganda, bots, fake news, election fraud, etc. is hardly regulated or guided (Hagendorff, 2020, p. 105). A sentence in his conclusion shows the

bigger picture of AI ethics in our society: “[...] The purposes for which AI systems are developed and applied are not in accordance with societal values or fundamental rights such as beneficence, non-maleficence, justice, and explicability” (Hagendorff, 2020, p. 114). To me, the author provides a clear and concise text that seems hard to misunderstand. It highlights weaknesses in AI ethics guidelines and the lack of addressing these.

Google’s AI presentation

Google talks about AI as a service and an idea they apply to products and new domains and want to share with the world. They also research to develop the field and accessibility of AI further. The webpage is divided into six categories; About, Responsibilities, Research, Education, Tools, and Blog.

They introduce their principles of AI with the claim that they “aspire to create technologies that solve important problems and help people in their daily lives” (Google, 2022). I think this goes to show their definition of AI as not only a service but also a future-oriented idea.

“Kimi,” the voice assistant

“Kimi” is a movie about (amongst other things) a woman working to improve a voice assistant (similar to Alexa or Siri) by going through sound clips of frustrated people trying to tell their “Kimi” what to do without getting a satisfying result. She uses these to alter algorithms programmed in the devices so that they can meet user demands to a larger degree. After hearing a sound clip of a woman getting raped, the movie plot escalates as the rapist is the company's founder, and she (the main person) is stalked by men trying to kill her. Many attempts later, they find themselves in a battle in her apartment, where the main person uses her “Kimi” device to fight the predators and win.

Human interaction with AI is shown as a breach of user privacy and is crucial for saving a life and a rape victim. I think there is an underlying attempt to show some of the company utilization of user data collected and stored from voice assistants that daily users potentially do not think about - such as using voice clips to improve their devices.

1.2 Robots and AI systems

Origin of the word “Robot”

According to Matt Simon (2020), the word Robot comes from Czech “forced labor” and was first used in 1921 by Karel Capek. He had created a play called “R.U.R” (Rossum’s Universal Robots), where human-looking robots wound up going on a killing spree, creating the foundation of our general skepticism towards these types of machines (Simon, 2020).

Definitions of the word “Robot”

Sebastian Thrun introduces three different definitions of the word robots. One of them is from 1977 when a robot was defined as “a reprogrammable, multifunctional manipulator

designed to move materials, parts, tools, or specialized devices through various programmed motions for the performance of a variety of tasks" (Thrun, 2004, p. 11). In my understanding, this can be described as defining a machinery robot, whereas more modern robots that are used for f.ex. social interactions are left out. I would explain it to suggest that robots are machines that, to a larger extent, automatize industrial processes.

The following definition goes: "[a robot is] an automatic device that performs functions normally ascribed to humans or a machine in the form of a human." (Thrun, 2004, p. 11). This seems to be a more modern term, which shows that robots can replace humans in several different processes - not only in industrial work but also in healthcare work, surveillance, etc.

My definition of the word "robot" would be: "a physical computer that is programmed to act automatically". In this way, I think I manage to separate machines and robots, something I feel is necessary because machines, in my opinion, act automatically but not autonomically. Machines can also be manual, whereas robots have to be electric. After some research though, I am no longer sure if any machines or AI can be autonomous, as the term implies descicion making that would extend the commands that are programmed into the different things by humans.

AI and Robots

In line with my definitions of "AI" and "Robot", I think one of the main differences between the two is that Robots are physical machines, while AI's are systems that can be programmed and added to f.ex. robots or other systems and digital machines. There are similarities, such as their acting autonomically, but the functionalities/data material/abilities of AI continuously grows as they receive and process new data input. If a robot doesn't have AI, it will keep its original functionalities.

Kaspar, the minimally expressive robot

Kaspar is a robot used in research projects in the area of robot-assisted play (Dautenhahn et al. 2009, p. 380). An example for this is education and therapy for children with autism. Human interaction with Kaspar is meant to be natural, creating space for the observer/interaction partner to communicate without potentially being overwhelmed by social cues. Therefore he is designed with minimal facial expressions (Dautenhahn et al. 2009, p. 375).

Kaspar's movements are restricted to its upper body, as it does not have functions in its legs. It is designed with facial/head and gestural expressions that allow social interaction and communicative and collaborative games. In one specific situation, it is utilized to teach a child with autism to teach turn-taking skills (Dautenhahn et al., 2009, p. 384). This is illustrated in the figure underneath, borrowed from figure 16, p. 384 in the report.



1.3 Universal Design and AI systems

Universal design

The Irish National Disability Authority describes universal design as “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 age, size, ability or disability.” (2022). I want to use the “GAP model” to explain the definition. In my understanding, the model illustrates that the gap between user abilities and demands from society should be as narrow as possible. In other words, we need to work on lowering requirements from the environment and, at the same time, strengthen user abilities in situations where they may meet challenges with participating and doing what they want or need to in their environment. This can be done through Universal Design - a term that aims to minimize the gap between users' abilities and environmental/societal demands.

Inclusive and exclusive AI

START (Symptom Tracker And Resources for Treatment) is a screening tool created for healthcare workers to track mental health status and connect users to mental healthcare resources. The platform draws questions from AI algorithms that efficiently capture the most relevant clinical symptom data (Lee et al., 2020). This example shows how AI can support mental health and human emotions. It also provides a solution where technology can be used to f.ex. fill roles of health care professionals in cases where there is a lack of access to such people and potential clients struggle with affording or finding time to prioritize such services. I think it shows the potential to fill gaps in access to mental health professionals by reducing the time professionals need to spend on mapping of potential treatment, and instead connecting the professionals and patients at later stages. These aspects could create space to help an increasing amount of people. On the other hand, I think there are concerning potentials of replacing humans in critical roles where human-AI interaction is insufficient.

AI can also include other groups of people in situations described in the KASPAR section. Not only can it be utilized for working with autistic children and their social skills, but it also has the potential to help in broader therapy settings and schools (Dautenhahn et al., 2009, p. 392). The criticism and concern described above also go for this solution.

An example of AI potentially excluding people is fingerprint scanning used in f.ex. stores to confirm age when buying products with age limits, in phones for access, or in airports for identity confirmation. :) It excludes people who do not have fingerprints. Fortunately, there are examples where one can choose between solutions, such as face scanning or numeral codes to unlock phones.

WCAG 2.1 and the concept of "understanding"
To me, "understand" and "understanding" means that one perceives, processes, and makes sense of information in the way it is intended to be.

It is problematic to say that machines "understand" when our definition of the term is regarding human understanding. This is because machines "understanding" is in the form of receiving, processing, and acting in line with the information they have received. Machines cannot develop deeper learning and individual thinking in the way that humans can through increased understanding and interpreting and, therefore, cannot be described as understanding.

1.4 Guideline for Human-AI interaction

Microsoft AI Guidelines point 05.; Match relevant social norms
Cultural and social contexts vary across the world. As the internet is global and AI systems are accessible across many of these contexts, it is essential to address several ways of use, understanding, and interpretations of interaction with AI.

Jakob Nielsen's 10 Usability Heuristics for User Interface Design and Microsoft's Ai Guidelines
The HCI and AI guidelines are similar in many ways. They address the same issues and introduce similar suggestions for design. However, Nielsen's Guidelines address the user's whole experience to a large extent, whereas Microsoft focuses solely on the interaction. For example, Nielsen recommends aesthetic and minimalist design of interfaces in his 8. heuristic. This difference is very natural regarding what the guidelines are meant to address.

Literature

Bratteteig, T & Verne, G. (2018). Does AI make PD obsolete? Exploring challenges from Artificial intelligence to Participatory design. https://doi.org/10.1145/3210604.3210646

Copeland, B. J. (2022). Artificial intelligence. Britannica. https://www.britannica.com/technology/artificial-intelligence/Is-strong-AI-possible

Dautenhahn, K. et al. (2009). KASPAR – A Minimally Expressive Humanoid Robot for Human–Robot Interaction Research. Applied Bionics and Biomechanics. 6. 369-397. https://doi.org/10.1155/2009/708594

Fuglerud, Kristin. (2014). Inclusive design of ICT: The challenge of diversity. 10.13140/2.1.4471.5844.

Hagendorff, T. The Ethics of AI Ethics: An Evaluation of Guidelines. *Minds & Machines* 30, 99–120 (2020). https://doi.org/10.1007/s11023-020-09517-8

Lee, F., Liston, C. & Forese, L. L. (2020). This AI tool helps healthcare workers look after their mental health. https://www.weforum.org/agenda/2020/09/ai-tool-mental-health-frontline-workers

Matteson, S. (2019). Autonomous versus automated: What each means and why it matters. https://www.techrepublic.com/article/autonomous-versus-automated-what-each-means-and-why-it-matters/

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

National Disability Authority. What is Universal Design. https://universaldesign.ie/what-is-universal-design/

Nielsen, J. (2020). 10 Usability Heuristics for User Interface Design. https://www.nngroup.com/articles/ten-usability-heuristics/

Simmons, B. & Chappell, S. G. (1988). Artificial intelligence - Definition and Practice. The journal of engineering, vol 13, no 2.

Simon, M (2020). The WIRED Guide to Robots. https://www.wired.com/story/wired-guide-to-robots/

The Council of Europe (2022). History of Artificial Intelligence. https://www.coe.int/en/web/artificial-intelligence/history-of-ai

Thrun, S. (2004). Toward a Framework for Human-Robot Interaction, Human–Computer Interaction, 19:1-2, 9-24, DOI: 10.1080/07370024.2004.9667338

IN5620 Individual assignment 2 - First iteration of Module 2

My Schultheiss - mysc

Characteristics of AI-infused systems

Draw on the first lecture of Module 2 and three of the mandatory articles (Amershi et al. (2019), Kocielnik et al. (2019), Yang et al., (2020)).

Identify and describe key characteristics of AI-infused systems.

Yang et al. (2020) write that AI-infused systems is a poorly-defined term for artificial intelligence (AI) (p. 4). Amershi et. al. (2019) however, refer to the term as a definition of "systems that have features harnessing AI capabilities that are directly exposed to the end user". In this assignment, I will use the latter definition as a standing point for what AI-infused systems are.

In the first lecture of module 2, Asbjørn Følstad defines four characteristics to be key for AI-infused systems. These are learning, improving, black boxes and being fuelled by large data sets. One aspect of AI-infused systems as defined by Følstad (2022, p. 39), is that they function similarly to black boxes in airplanes. This means that it can be hard for users (and developers) to understand why the system does what it does.

The characteristic of evolving by learning is made possible through data input. Because data input has traditionally been small and narrow, Ai-infused systems have been inconsistent. The nuances of tasks and settings have been a starting point for poorly understood, probabilistic behavior that leads to low abilities to recognize certain conditions that end users do not see as distinct, but that shape the outcome of how the systems respond. They change via learning over time and with such input therefore give different outputs to different users, something that can create confusion for the users and more drastically lead to abandonment of the technology (Amershi et. al., 2019).

Amershi et. al. (2019) writes that AI-infused systems lack attributes to make them intuitive and effective. Improving is described as another characteristic of AI-infused systems. This happens through evolving over time, and therefore such attributes are key. The evolvement includes learning-, reasoning- and self-correcting processes that rely on large data sets to work with, and therefore need the right algorithms to reach the desired outcome (Burns et. al., 2022). As written in the lecture by blabla (2022, p. 43), the systems are programmed to learn over time from user behavior through their interactions. Therefore larger data sets accommodate the learning- and self-correcting processes.

Identify one AI-infused system which you know well, that exemplifies some of the above key characteristics.

An AI-infused system that is well-known to me is google's AI-toolkit to find child abuse content. This is used to scan the content of every user of their cloud-saving solutions, in order to identify photographs that may be evidence of child abuse. The background and reason for using this tool-kit is to find and save children from pedophile abusers (Diaz, K., 2022). Therefore it is programmed with algorithms that recognize both content that potentially is child abuse, and whether the content is new or not.

The older images have been used to create a learning process in the system. This way the AI can distinguish images that have been spread around the internet for a longer period of time (typically many years), from newer images. The reason for such a distinction is that older content does not provide opportunities of potentially being able to save the child that is exposed, to the same extent that newer images do.

Discuss the implications of these characteristics for the example system, in particular how users are affected by these characteristics

Users are affected in different ways by this system. Some of them do not know that they are part of it before they are contacted by authorities that have started investigations on their actions, and some will be included in the algorithms of the system without consenting or knowing about it, as the AI-infused system works as a black box where users cannot see any of its actions. The latter unknowing "users" are children that may be featured in abusive content, because the system learns and evolves by being fed such content.

There are also users that are in direct contact with the system, as they have to go through all content marked by the system as abusive content and check whether the system is correct or not. These users know how it works and their part of it, and therefore are affected in very different ways than the other users; they work on improving the AI through sorting out content that is wrongfully marked as abusive, and therefore have insight into what is going on. They also choose what data to feed the system with to create the learning- and improving process.

Summaries of main takeaways in Amershi et al. (2019) and Kocielnik et al. (2019) Guidelines for human-AI interaction

Amershi et al. (2019) discuss human-AI interaction and how AI-infused systems can become more human-centric. They propose that by following certain guidelines when shaping such systems, one can achieve a human-centric design that will better human-AI interaction.

The authors have researched guidance proposed over the last 20 years and created 18 AI usability guidelines based on their findings (Amershi et al., 2019). They write that they "[...] anticipate the guidelines will be useful to evaluate existing products and emerging design ideas." The guidelines are further described to be possible heuristic evaluation points for usability testing or inspecting. As the guidelines have not been utilized at the time of their conclusions, one takeaway is that further research can examine and evaluate the uses and values of the guidelines (Amershi et al., 2019).

Will You Accept an Imperfect AI? Exploring Designs for Adjusting End-user Expectations of AI Systems

Kocielnik et al. (2019) study whether a High Recall- or a High Precision version of a Scheduling Assistant lead to higher user satisfaction and acceptance of the system. Both types of systems have a 50% level of accuracy, but the High Recall version is designed with a focus on avoiding false negatives, and the High Precision version on avoiding false positives. This means that High Recall systems are to improve by learning from false negative results, such as predicting that an email containing the sentence "how about lunch? Maybe 1:30?" is in fact a meeting request. High Precision systems on the other hand, learn from emails containing sentences such as "Yesterday's meeting was good". The study showed that a focus on High Precision rather than High Recall systems can lead to lower perceptions of accuracy and decreased acceptance (Amershi et al., 2019).

Select two of the design guidelines in Amershi et al. (2019). Discuss how the AI-infused system you used as example in the previous task adheres to, or deviates from, these two design guidelines.

For this task I have chosen the design guidelines G1: "Make clear what the system can do." and G13: "Learn from user behavior". Google's AI-infused system can be argued to deviate from G1 for users that are unaware of their participation in the system, while adhering to it for administrative users. This is because, as already explained, the users that administrate the system have a responsibility to make sure it improves and find the content it is supposed to, and therefore knows exactly what the system can do.

On the other hand we have people who store their images and videos in Google Photos. As it is this cloud-saving service the AI-infused system scans for content, I would describe them as unknowing users. They do after my understanding not get insight into the system or its capabilities. I have not been able to find any information on the system in Google's websites or other reports other than news articles on an incident where an unknowing user was in fact (wrongfully) tagged as a possible pedophile (Diaz, K., 2022). Because the system learns through humans manually separating abusive and non-abusive content I would say that G13 also is a guideline that it deviates from.

Briefly discuss whether/how these two design guidelines could inspire improvements in the example system.
The example system I chose is hard to discuss when it comes to improvements through the design guidelines. If it was clear to the users of Google Photos that the system scans their files, some or all users may stop using them and make the AI-system less useful or useless. The system is already becoming more known through different mediums. On the other hand, the G1-guideline could inspire improvements for the sorting-workers by displaying metrics that facilitate more efficient alterations of the algorithms. I am not sure how G13 could inspire improvements.

Summarize the argument by Bender et al. 2021 concerning problematic aspects of textual content and solutions based on large langue models.
Bender et al. (2021) argue that one risk of language models is that they are shaped by their training data. Because the training data is often to a large extent provided by privileged people, and such data includes human hegemonic worldviews and language, the result is systems with encoded biases. Another argument is that bad actors can take advantage of the language model's abilities to produce seemingly coherent texts on specific topics. They can for example generate text that provides content by non-existent users (Bender et al., 2021, p. 617).

Extremely large amounts of parameters in the systems shape the output they provide. An example of this is that users can gain access to the training data of the language models and shape it. Also, a risk that is argued to concern the language models is machine translation which can cause readers to misunderstand the original authors' intentions (Bender et al., 2021, p. 618).

Careful planning before beginning to build datasets or systems trained by datasets can mitigate the risks of large language models. This includes careful data collection, adopting frameworks to describe what uses their models are meant for and evaluations of the systems, and re-aligning research goals (Bender et al., 2021, p. 618). Bender et al. (2021) also argue that value sensitive design is a methodology that can help in mitigating risks (p. 619).

Discuss key challenges in the design of chatbots / conversational user interfaces.
Yang et al. (2020) introduce potentially problematic areas that need to be addressed in the design process of AI systems. These areas have been reported by human-centered and user-experience researchers and practitioners in their work with such systems (2020, p. 2). During the presentation Asbjørn Fjølstad held about interaction with AI and autonomous systems, he talked about these challenges and how they can be solved, as illustrated by Yang et al. (2020, p. 2). The illustration shows a human-centered design process through the

double diamond model, and includes elements that can affect the resulting AI system throughout the whole process. Some key challenges that Fjølstad pointed out in his presentation are that it can be hard to foresee potential effects of AI, explain AI behaviors to users, and anticipate/mitigate unpredictable AI behaviors (2022, p. 13).

One challenge as mentioned above is that AI systems can make unpredictable errors that affect user experience (Yang et al., 2020, p. 1). In HCI design, practitioners use rapid iterations with prototyping and assessing human consequences of the design in order to improve it, and because they cannot do this with AI systems, user experience and societal consequences of these is hard to foresee (Yang et al., 2020, p. 3). The other key challenges can come in to play in a loop where lack of assessment and improvement lead to challenges with explaining the AI behaviors to users and foreseeing its potential effects.

Revisit Guidelines G1 and G2 in Amershi et al. (2019). Discuss how adherence to these could possibly resolve some of the challenges in current chatbots / conversational user interfaces.

Adherence to Guideline G1 “make clear what the system can do” could resolve challenges with explaining AI behaviors to users. If there was a transparency in system functionalities and restraints, such explanations may come continuously and create a bigger understanding. It could also help with resolving the other key challenges I have mentioned, and in that way create a better HCI design process in total. Including G2 “Make clear how well the system can do what it can do” could also create a possibility of improving the process to a larger extent. If users understood what a system could do to a larger extent it could also be easier to work with including users in the process of improving the system and assessing potential future effects and unpredictable behaviours it may show.

Literature

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. (https://www.microsoft.com/enus/research/uploads/prod/2019/01/Guidelines-for-Human-AI-Interaction-camera-ready.pdf)

Bender, E. M., Gebru, T., McMillan-Major, A., & Mitchell, M. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (pp. 610-623). ACM. https://dl.acm.org/doi/pdf/10.1145/3442188.3445922

Diaz, K. (2022). Google's AI labels a parent a pedophile after taking photos intended for the pediatrician: when you ignore privacy and it goes wrong. https://thenewstrace.com/googles-ai-labels-a-parent-a-pedophile-after-taking-photos-intended-for-the-pediatrician-when-you-ignore-privacy-and-it-goes-wrong/246095/

Følstad, A. (2022). INTERACTION WITH AI AND AUTONOMOUS SYSTEMS – MODULE 2. https://www.uio.no/studier/emner/matnat/ifi/IN5620/h22/interacting-with-ai-2022---module-2---session-1---handout.pdf

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. https://www.microsoft.com/enus/research/uploads/prod/2019/01/chi19_kocielnik_et_al.pdf

Yang, Q., Steinfeld, A., Rosé, C., & Zimmerman, J. (2020). Re-examining Whether, Why, and How Human-AI Interaction Is Uniquely Difficult to Design. In Proceedings of the 2020 CHI conference on human factors in computing systems (Paper no. 164). https://dl.acm.org/doi/abs/10.1145/3313831.3376301