

Second Iteration

University of Oslo | Department of Informatics

Course: IN5620
Student: ratungla
Lecturer: Jo Herstad

Delivery: 09.11.2022

Pages: 12

Innholdsfortegnelse

1.1 Concepts, definitions and history of AI **4**

1.1.1 History of AI *4*

1.1.2 Definitions of AI *4*

1.1.3 Human vs. machine autonomy *5*

1.1.4 Article I Holmquist (2007) – Intelligence on Tap *5*

1.1.5 Contemporary company working with AI *6*

1.1.6 AI in popular culture *6*

1.2 Robots and AI systems **6**

1.2.1 How the term robot came about *6*

1.2.2 Different definitions of the word robot *6*

1.2.3 The relationship between AI and robot *7*

1.2.4 Describe the interaction between human and a contemporary robot *7*

1.3 Universal design and AI systems **7**

1.3.1 Universal design defined *7*

1.3.2 AI potential with respect to human perception, human movement and human cognition/emotions *8*

1.3.3 AI - Exclusion or inclusion? *8*

1.3.4 Do machines understand? *8*

1.4 Guidelines for Human-AI interaction **9**

1.4.1 Description of a Microsoft Human-AI interaction guideline *9*

1.4.2 HCI design guidelines vs. Human-AI interaction guidelines *9*

2.1 Characteristics of AI-infused systems **10**

2.1.1 Key characteristics of AI-infused systems *10*

2.1.2 Key characteristics of an AI-infused system in practice I Apple Music *11*

2.2 Human-AI interaction design **11**

2.2.1 Amershi et al. (2019) and Kocielnik et al. (2019) I Summary *11*

2.2.2 How Apple Music adheres to design guidelines *12*

2.2.3 Bender et al. (2021) | Critical discussion regarding AI systems – Language *13*

2.3 Chatbots | Conversational user interfaces **14**

2.3.1 Key challenges regarding designing chatbots *14*

2.3.2 Could adherence to G1 and G2 resolve challenges with chatbots? *14*

Literature **16**

1.1 Concepts, definitions and history of AI

1.1.1 History of AI

Although the concept of what we now define as artificial intelligence (AI) has been explored through the world of science fiction, popular media and through pioneers like Alan Turing during the first half of the 20th century, the term AI was first introduced at the Dartmouth Summer Research Project on Artificial Intelligence (DSRPAI) in 1956, by John McCarthy (Anyoha, 2017). The conference was held in order to further explore and discuss the possibilities of AI as a research-field (Anyoha, 2017).

1.1.2 Definitions of AI

Hearing the term AI, thoughts about autonomous robot armies attacking humanity or Siri never being able to understand what you try to say, might arise. As the scope and possibilities of AI and information technology continuously expands, finding one definitive all-encompassing definition could pose a challenge. Let's look at three different definitions found in a 2021 Medium.com article (Ribeiro, 2021):

> *"AI is the science and engineering of making intelligent machines"*
> -John McCarthy
>
> *"AI is getting computers to do tasks that would normally require human intelligence"*
> -Deloitte
>
> *"Artificial intelligence is a type of technology that mimics the human thought, empowering machines to act on their own and to perform functions similar to human intelligence, such as the ability to perceive, learn, reason, and act."*
> -Jair Ribeiro

The abovementioned definitions have origins from three different domains. Whereas McCarthy's definition has a scientific basis, Deloitte is a private company with independent goals for growth and profit. Lastly, Jair Ribeiro (2021) represent a media-/news-outlet.

In all three definitions, intelligence is a common denominator. They differ when describing which systems are encompassed, ranging from machines to computers. They also define different scopes, types of systems and what those systems do. While McCarthy's definition describes machines just being intelligent, Deloitte describe computers doing tasks that normally would require human levels of intelligence. Ribeiro (2021) expands upon this, describing systems mimicking human thought and problem-solving patterns through perception, learning, reasoning and acting.

Based on these definitions, I've made an AI-definition that encapsules both AI as a scientific field and a type of technology. The new definition also expands upon AI and its ability to imitate human behavior, not just explicitly intelligence:

AI encompasses the science, engineering and technology of systems exhibiting or imitating human-like behavior and intelligence, in order to perceive, learn, reason and act.

1.1.3 Human vs. machine autonomy

Autonomy can be considered as a condition of thinking, and as a distinct characteristic of intelligence (Martinez, 2019). Humans and animals alike can, based on this, be characterized as autonomous. They're free to think and act upon these thoughts. As machines mostly just do what they're told, autonomy for machines can be seen as an ability of decision-making without any outside input (Martinez, 2019). This will still be by design and its original programming, but with the ability to learn and provide some output that was not originally programmed. In the same way as humans are "programmed" through genomes and environmental factors, autonomous machines have a less complex programming giving it wanted attributes and abilities to somewhat act outside of just direct control.

1.1.4 Article I Holmquist (2007) – Intelligence on Tap

Do you remember all those CAPTCHA-tasks you completed to prove you were not a robot? Well now, thanks to *your* impeccable capabilities in identifying traffic lights and fire hydrants, the robots have learned enough and are ready to replace you. But only at *certain, limited* tasks. In the article "Intelligence on Tap: Artificial Intelligence as a New Design Material", Holmquist (2017) highlights recent advances in machine learning through neural networks, due to increased processing power and huge amounts of available data (which may or may not have been ethically collected). Like training a dog with treats, neural networks learn through reinforcing/weighting nodes in its network when they succeed at pre-defined tasks.

This can make AI capable of solving complex problems and the recent advances is already making AI-technology and machine learning through neural networks more accessible without being an expert in the field. As a designer, "intelligence on tap" can be used in the future, but not without limitations. AIs can still not form an actual understanding of the world. Designers or whoever uses AI-generated solutions as design basis/material therefore needs to tread lightly, manage expectations and understand the current limitations of the technology for the technology to be useful (Holmquist, 2017).

1.1.5 Contemporary company working with AI

I've had a look at Levity.ai, a company delivering what they describe as an "out-of-the-box AI technology" for classifying documents based on set parameters (Kleinings, 2022). Clients can train neural networks with their own datasets and be provided with the necessary service in order to do said tasks. Both the service and the infrastructure is hosted by Levity for customers to subscribe to for a monthly fee. They therefore describe their services as AIaaS (AI as a service).

1.1.6 AI in popular culture

The 2009 science fiction movie Moon is set in space, where an astronaut is living and working in complete isolation, only accompanied by the AI-robot, GERTY. GERTY is portrayed as being able to understand and communicate natural language, while also being able to recognize and show emotion (using emojis). The emotional response from GERTY seems to be limited to reflecting the current mood of the astronaut through sympathizing. GERTY has its own pre-defined (programmed) goals, that differs from the astronaut.

Although GERTY is used to stifle the astronaut's jarring loneliness, the longing for actual human-to-human interaction is evident through the depiction of deteriorating mental health. The movie highlights the importance of human-to-human/sentient interaction and illustrates how AI companions might not suffice when trying to cope with loneliness.

1.2 Robots and AI systems

1.2.1 How the term robot came about

Looking at the etymology of the word "Robot" we can find the first English citation of the word back in 1839, with a meaning derived from the Czech word *robota*, meaning forced labour/slavery (Interesting Literature, 2022). The current understanding of the word robot however, was first seen in the Czeck play *Rossum's Universal Robots,* featuring a factory making "artificial people" and having them serve and work for humans (Interesting Literature, 2022). Nobody enjoys being slaved around and the robots uproared, as you can imagine.

1.2.2 Different definitions of the word robot

In the same way AI has a multitude of different definitions, there can be found a plethora of different definitions of the word "Robot". Let's have a closer look at two such definitions:

"Actuated mechanism programmable in two or more axes with a degree of autonomy, moving within its environment, to perform intended tasks." (ISO, 2012)

"Any automatically operated machine that replaces human effort, though it may not resemble human beings in appearance or perform functions in a humanlike manner." (Moravec, 2022)

The ISO 8373:2012-definition (ISO, 2012) defines robots as a programmable and movable *mechanism*, with the intent of solving tasks. Encyclopedia Britannia (Moravec, 2022) has a slightly different definition, defining robots as automatically *machines* that replaces human efforts in order to solve tasks, although not necessarily with human-like features and characteristics. The latter definition does not specify movement as an ability. I've created my own definition that might be a bit vague, trying to encompass all types of programmable and physically movable mechanisms that are able to perform taskes based on input:

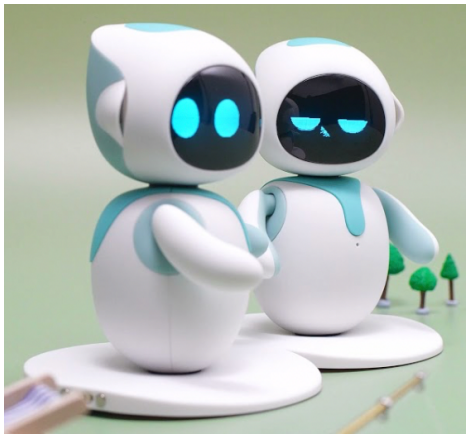
Robots are programmable, sensory and movable physical mechanisms, designed to perform intended tasks within a given environment.

1.2.3 The relationship between AI and robot

AI technology is defined as using human-like intelligence in order to solve tasks, while robots are defined as mechanisms performing programmed tasks using "human effort". Robots act through direct programming and sometimes with human-like movement, instead of autonomous actions or intelligence. A robot *can* have AI, but it doesn't *have* to in order to still be classified as a robot. Based on previously mentioned popular culture depicting humanities doom due to robot uprisings, huAIs with legs might not be the greatest idea?

1.2.4 Describe the interaction between human and a contemporary robot

Meet Eilik, described as "Tech with heart" by it's creators (Energize Lab, n.d.). This robot is minted for the consumer market as a social companion, responding to touch, movement, sound and its own position within a physical space. It mimics human emotion through digital "emoji-like" facial expressions, and through it's movement. It can spin 360 degrees and move it's head and arms.



Eilik seems to have interchanging moods, inviting users to interact by using a language that are recognizable to us all, visible emotion. If it's sad you can pick it up, pat it on the head or similar and it will respond to the sensory input and output a response.

1.3 Universal design and AI systems

1.3.1 Universal design defined

The Norwegian Digitalisation Agency (Digdir) defines universal design as:

"Designing, or accommodating, the main solution with regards to physical conditions, so that the solution may be used by as many people as possible, regardless of disability" (UUtilsynet, n.d.)

Universal design as per Digdirs definition (n.d.) is about designing for inclusion and availability regardless of individual uniqueness, variability or physical/cognitive abilities. The goal is to make systems that are accessible and usable for as many individuals as possible and not hindering use, which in sum should benefit all users, not just the minority.

1.3.2 AI potential with respect to human perception, human movement and human cognition/emotions

Every single human being is unique regarding capabilities, interests, culture, emotion, wants and needs. The current “one size fits all”-paradigm of software development might become a thing of the past, due to the future of AI-capabilities. Designing systems for every single unique user, with tailored user interfaces and functions is not feasible or possible today. There is however huge potential for a new genre of “generic AI-software”, adapting to you, your capabilities, emotions and needs, with dynamically changing UIs and functions.

For students, maybe the “Focus”-mode of current cell phones/PCs could identify patterns of what triggers them to be distracted and then remove/hide these UI-elements/functions when a deadline is closing in, making it easier to work efficiently. Or maybe it could even understand how someone learn most efficiently and tailor/present content in a way that works for them. Maybe a boring text could be made richer by automatically generating or adding relevant images, videos or similar, in the same way social media algorithms are designed to keep you continuously engaged.

1.3.3 AI - Exclusion or inclusion?

As Holmquist (2017) described the future of AI as a design material (seen in task 1.4), he also highlighted some of the risks associated with neural networks and machine learning. These networks learn from existing user data, which will most likely also contain current biases, posing a risk of possibly reinforcing existing exclusionary tendencies, bias and prejudice in society.

There’re already multiple examples of inherited AI-bias in real-world. US hospitals used an algorithm on more than 200 million patients, trying to identify which patients would most likely need extra medical care. In 2019 it was identified that the algorithm heavily favored white patients over black patients, not because of parameters separating by skin color, but due to the groups healthcare cost history (Shin, 2020). To counter inherited bias and increase inclusiveness, AI-systems and neural networks could also be designed to identify imbalance or possible bias and flag these in datasets.

1.3.4 Do machines understand?

In the context of the WCAG 2.1 principles (W3C, 2018) and the guidelines for Human-AI-interaction, I would argue that the words “understand” and “understanding”

represent the ability to not just perceive or know something (like an output from an interface), but rather being able to categorize and associate whatever is perceived as part a context, connected to a larger system of knowledge and as part of a complex subjective understanding of how the world works.

Based on this I'll argue that machines and AI in their current state (and in their walled-off garden) does not currently have the capability to understand something, but rather holds the ability to sense and perceive and know something about what was perceived. I do believe this is subject to change in the future.

1.4 Guidelines for Human-AI interaction

1.4.1 Description of a Microsoft Human-AI interaction guideline

Guideline 16: Convey the consequences of user actions (Amershi et al., 2019a).

Make sure to convey the user how and if an interaction with the AI-system might impact system behavior in the future. If the AI-system detects behavior or data that may have been unintentional or not representative of the target it learns from or interacts with, an option should be given to delete or invalidate the data in order to not cause unintended unwanted behavior.

1.4.2 HCI design guidelines vs. Human-AI interaction guidelines

When taking a look at Jakob Nielsen and Rolf Molich's Ten User Interface Guidelines (building upon Ben Shneiderman's "eight golden rules"), one can see that these heuristics does have some similarities to the Human-AI interaction guidelines from Microsoft (Wong, 2020).

While HCI design guidelines are more generalized and applicable to the design of every type of interface, the Human-AI guidelines are specialized for this specific type of technology. This is specialization is highlighted by the Human-AI guidelines being divided into different categories or "stages", such as "initially", "during interaction", "when wrong" or more the more general "over time".

One of the ten user interface guidelines from Nielsen & Molich, "Consistency & Standards", recommends designs that implement a recognizable and established design language, similar to other contemporary interfaces (Wong, 2020). To me, this translates to following current norms within design and interaction. The human-AI interaction guideline "Match relevant social norms", is similar in that it recommends an experience that is delivered in a way users would expect, given their context and societal norms (Amershi et al., 2019a).

2.1 Characteristics of AI-infused systems

2.1.1 Key characteristics of AI-infused systems

Amershi et al. (2019b) defines AI-infused systems as “systems that have features harnessing AI capabilities that are directly exposed to the end user”. Such systems are widely available and adopted in modern society, encountered on many bedside tables in the form of Amazon Alexa-enabled smart speakers or whenever you attempt to persuade a chatbot into connecting you to a human representative.

Even though AI-infused systems can have varying capabilities and take on different shapes and forms, some key characteristics are defined in current AI research-literature. They’re often categorized as somewhat inaccurate, where a certain “margin of error” is to be expected. In this regard, Kocielnik et al. (2019) highlights their probabilistic nature and proneness to inaccuracy. Amershi et al. (2019b, p. 1) further describes AI-infused systems as prone to making errors, often producing false positives and negatives and generally “performing under uncertainty”.

Established usability guidelines calls for consistency through user interfaces and how systems respond to interaction (Amershi et al., 2019b). AI-infused systems, however, tends to be sensitive to minor variations in context or user input, making them inconsistent in how they respond and react (Amershi et al., 2019b). Kocielnik et al. (2019) also highlights the large impact even minor user action can have on how an AI-infused system respond, as a key characteristic of such systems.

A third characteristic with AI-infused systems is their “black boxed” inner workings. Curious to see what caused the system to respond the way it did? Unfortunately, the “box” containing all the “AI magic” tends to be opaque and often challenging to understand. Yang et al. (2020, p. 1) discuss these limiting characteristics as design challenges within human-AI interaction, highlighting the “uncertainty surrounding AI’s capabilities” and the “AI’s output complexity” as design challenges. Amershi et al. (2019b).

Identify one AI-infused system which you know well, that exemplifies some of the above key characteristics. Discuss the implications of these characteristics for the example system, in particular how users are affected by these characteristics.

2.1.2 Key characteristics of an AI-infused system in practice | Apple Music

I often listen to music, currently using Apple Music as my music streaming platform of choice. Apple Music offers, similarly to Spotify, AI-curated playlists based on the listeners listening history. Personally, I often use this functionality as an attempt to find new music that I might enjoy. More often than not, Apple Music keeps on recommending songs from this one specific artist I don't really enjoy. Maybe I stumbled upon and played a song by that artist by accident ages ago, but I can't remember having listened to any of their songs. This "inaccuracy", causing "false positives" in my predicted music could be the result of me mistakenly clicking on the wrong song at some point, but I have no way of knowing *why* the recommendation are as they are, as Apple Music does not inform me why specific music or artists are recommended.

2.2 Human-AI interaction design

2.2.1 Amershi et al. (2019) and Kocielnik et al. (2019) | Summary

Both articles mentioned in the title explores how AI-infused systems can be designed in to increase acceptance and usability for end-users. Kocielnik et al. (2019) focuses on how unrealistic expectations regarding the capabilities of AI-infused systems can lead to negative perceptions of the technology for users, whenever those expectations are not met.

Popular media and enthusiastic tech companies often portray AI as something more than what it currently is, creating expectations that might cause disappointment whenever a user encounters the probabilistic and unpredictable nature of AI-infused systems. A takeaway from this article can be the importance of designing AI-infused systems which doesn't try to pretend like it's something it's not and set realistic expectations for users regarding what the system can and cannot do, in order to increase acceptance.

Similartly to the worries of Kocielnik et al. (2019), Amershi et al. (2019a, p. 3) are also worried about how the “inconsistent and unpredictable behaviour” of AI-infused systems can erode users’ confidence and “lead to abandonment of AI technology”, stating that “AI-infused systems can violate established usability guidelines of traditional user interface design” (Amershi et al., 2019b, p. 1). Due to the unpredictable nature of these systems, they suggest 18 AI design guidelines that can be used in applied in order to increase usability and user acceptance.

2.2.2 How Apple Music adheres to design guidelines

Chosen AI-design guidelines

G9 – Support efficient correction
G11 – Make clear why the system did what it did



The previously mentioned AI-infused music streaming service, Apple Music, does adhere to Amershi et al. (2019b) AI design guideline G9, “Support efficient correction”. I previously indicated some annoyance over songs and artist I don’t necessarily enjoy being featured repeatedly in personally curated playlists. I haven’t corrected the algorithm though, but you can easily give feedback per song, stating if you like or dislike it. I think one of the reasons for me not correcting the algorithm is because I don’t know if it might somehow affect the system’s ability to suggest songs I actually enjoy listening to.

The system does not adhere to G11 – Making clear why the system did what it did. Song and playlists are created and presented with no information regarding *why* these songs were chosen. It is probably based on historical listening patterns, but it could be interesting if users were presented with a statement like “This song was recommended to you based on your interest in Röyksopp and Susanne Sundfør”.

2.2.3 Bender et al. (2021) | Critical discussion regarding AI systems – Language

What are the consequences of our ever-growing hunger for more data and bigger datasets? This continuous hunt for better machine learning could benefit us as a society in the long term, but Bender et al. (2021) recommend taking a step back in order to assess the risks and potential harms associated with data gathering for improved NLP (Natural language processing) through large language models (LM).

The article highlights the substantial environmental costs related to utilizing such language models, as the energy consumption required to train such models (which might not be generated from renewable sources) can equal that of “a trans-American flight” (Bender et al., 2021, p. 612). While this is problematic by itself, the distribution of potential environmental harm versus the achieved benefits might be unevenly distributed. Is it justifiable to train large English-based LMs that benefits the western world, while 800.000 people in Sudan are affected by drastic floods due to global waming? Bender et al. (2021) argues that the societies benefiting from LMs are usually not the same as those who are potentially harmed by their existence.

In addition, there’s a risk of encoding bias into the LMs using data gathered from the internet. Even though the internet can be seen as diverse, encompassing different viewpoints, cultures and demographics, “internet access itself is not evenly distributed” (Bender et al., 2021, p. 613). In this regard, language and opinions from groups overrepresented on the internet, such as younger demographics and those from developed countries are likely to become more prevalent in the LM. Hence, there’s a risk of an uneven representation. Due to the large quantities of data it’s also difficult to filter out discrimination, hate speech etc.

In order to mitigate the risk of encoded biases, Bender et al. (2021, p. 610) recommends “budgeting for curation and documentation”, meaning that there should only be collected data in quantities that are possible to analyze and curate, in order to avoid “documentation debt”.

2.3 Chatbots I Conversational user interfaces

2.3.1 Key challenges regarding designing chatbots

Having to interact with chatbots are often described as frustrating by users attempting to find information or when their original intention was to speak to an actual human representative. Conversational interactions between humans are usually boundless and conditioned on a mutual understanding of natural language. As users are presented with an AI-based conversational interface for the first time, it's possible that they'll assume that you can interact with them through natural conversation, just as you would with another human. These expectations are likely to be reinforced by chatbots taking a on a "human shape", being named "Kari" and perhaps having a profile picture looking like an actual person. As illustrated through chapter 2.2.1 and through Kocielnik et al. (2019), not managing users expectation for the capabilities of AI-infused systems could cause set the user up for a disappointing experience.

Simulating human conversational interaction using chatbots is difficult, as it's near impossible to account for every possible question or statement. As Yang et al. (2020) states, it's challenging for designers to envision all possible interactions, especially as AI-infused chatbots will learn, adapt and evolve over time based on interactions and it's given context. In order to properly train a chatbots NLP (natural language processing), there might be a need for large LMs, which as illustrated in chapter 2.2.3, and by Bender et al. (2021) could pose ethical challenges.

2.3.2 Could adherence to G1 and G2 resolve challenges with chatbots?

As stated in the previous chapter, setting realistic expectations for the usability of AI-infused systems is critical in order to increase user acceptance and to better the perception of such systems (Kocielnik et al., 2019). There's often a clear deviation between what chatbots are presented as (human-like, but completely lacking human ability to understand natural language), and what they're actually capable of doing. To help mitigate this, adhering to the AI-design guidelines from Amershi et al. (2019b) could help increase how chatbots are perceived.

In the regard, setting realistic expectations is the main intent of guideline G1 (Make clear what the system can do) and G2 (Make clear how well the system can do what it can do) (Amershi et al., 2019b). Instead of trying to exhibit chatbots as human, with human-like capabilities for understanding natural language, clearly stating what it can and cannot do and its level of potential inaccuracy, could make interaction with chatbots easier. This way, users won't have to "trip and fall" unnecessarily on their way to achieve their goals with the interaction.

Literature

Amershi, S., Vorvoreanu, M., & Horvitz, E. (2019a). *Guidelines for human-AI interaction design*. Microsoft. Retrieved September 13, 2022 from https://www.microsoft.com/en-us/research/blog/guidelines-for-human-ai-interaction-design/

Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019b). Guidelines for Human-AI Interaction. 1-13. https://doi.org/10.1145/3290605.3300233

Anyoha, R. (2017). *The History of Artificial Intelligence*. Harvard University. Retrieved September 12, 2022 from https://sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence/

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021, 2021-03-03). On the dangers of stochastic parrots: Can language models be too big?

Energize Lab. (n.d.). *Eilik*. Engerize Lab. Retrieved September 13, 2022 from https://energizelab.com/consumerview/eilik

Holmquist, L. (2017). Intelligence on tap: artificial intelligence as a new design material. *Interactions (New York, N.Y.)*, *24*(4), 28-33. https://doi.org/10.1145/3085571

Interesting Literature. (2022). *The Curious Origin of the Word 'Robot'*. Interesting Literature. Retrieved September 13, 2022 from https://interestingliterature.com/2016/03/the-curious-origin-of-the-word-robot/

ISO. (2012). *ISO 8373:2012(en)*. ISO. Retrieved September 13, 2022 from https://www.iso.org/obp/ui/#iso:std:iso:8373:ed-2:v1:en

Kleinings, H. (2022). *What is AIaaS? Your Guide to AI as a Service*. Levity. Retrieved September 12, 2022 from https://levity.ai/blog/aiaas-guide

Kocielnik, R., Amershi, S., & Bennett, P. (2019, 2019-05-02). Will You Accept an Imperfect AI?: Exploring Designs for Adjusting End-user Expectations of AI Systems.

Martinez, R. (2019). Artificial intelligence: Distinguishing between types & definitions. *Nevada Law Journal*, *19*(3), 9. https://scholars.law.unlv.edu/cgi/viewcontent.cgi?article=1799&context=nlj

Moravec, H. P. (2022). robot. In *Encyclopedia Britannica*.

Ribeiro, J. (2021). *These Are the Best Definitions of Artificial Intelligence You Can Read Today.* Medium.com. Retrieved September 12, 2022 from https://medium.com/swlh/these-are-the-best-definitions-of-artificial-intelligence-you-can-read-today-7c53c0e38584

Shin, T. (2020). *Real-life Examples of Discriminating Artificial Intelligence*. Towards Data Science. Retrieved September 13, 2022 from https://towardsdatascience.com/real-life-examples-of-discriminating-artificial-intelligence-cae395a90070

UUtilsynet. (n.d.). *Regulations*. UUtilsynet. Retrieved September 13, 2022 from https://www.uutilsynet.no/english/regulations/278

W3C. (2018). *Web Content Accessibility Guidelines (WCAG) 2.1*. W3C. Retrieved September 13, 2022 from https://www.w3.org/TR/WCAG21/

Wong, E. (2020). *User Interface Design Guidelines: 10 Rules of Thumb*. Interaction Design Foundation. Retrieved September 13, 2022 from https://www.interaction-design.org/literature/article/user-interface-design-guidelines-10-rules-of-thumb

Yang, Q., Steinfeld, A., Rosé, C., & Zimmerman, J. (2020, 2020-04-21). Re-examining Whether, Why, and How Human-AI Interaction Is Uniquely Difficult to Design.