IN5620 Individual assignment fall 2022 – Module 2

Characteristics of AI-infused systems

AI-infused systems are systems that directly exposes its AI capabilities to the end user. To identify these systems, you can look at certain characteristics:

One of the most typical characteristics to an AI-infused system is how inconsistent and unpredictable the system is. Errors are common and makes the system less reliable due to potentially dangerous and costly actions (Amershi et al, 2019). This applies to both false positives and negative positive errors, which makes it difficult to create human-verification systems. These errors can occur due to the sensitivity of user-input. Captcha (a human-verification system) is a well-known example to how unpredictable and inconsistent an AI-infused system can be.

Content impacted by user actions is another characteristic of an AI-infused system that is being more and more seen in today's new technologies (Kocielnik et al, 2019). This characteristic involves both user-generated content and behind the scenes personalization. The system can for instance use machine learning algorithms to show specific content to the user based on human-system interactions such as personalized advertisements and images created by user input. Usually, these technologies are probabilistic. Which implies that such systems are capable of almost always operate but can be less than perfect in terms of accuracy (Amershi et al, 2019).

AI-infused systems makes decisions when completing tasks, however, in some cases they do not provide transparency to the end-user in how the AI model get to a decision (Kocielnik et al, 2019). These systems are often referred to "black-box" systems and is another typical characteristic of certain AI-infused systems. Example of black-box systems are personality-tests in job applications. How do they work and how does the AI sort the applicants?

Identify one AI-infused system and its characteristics:

Google Lens is one of many fascinating and useful AI tools for the day-to-day life. This one give you the ability to use your smartphone camera to do many different tasks such as looking up products, translating text and even look up dish photos from a restaurant menu without photos. However, even though it proves itself useful in many scenarios there are some implications that may occur from using Google Lens as the AI system is probabilistic. Most of

the times it shows the user the wanted information and content, but it is not always correct. For instance, when you look up food photos for each menu item in a restaurant, it could give you an unrealistic expectation to what the restaurant offers. Another probabilistic example is when you want to translate text. It is not always perfect. It could be errors occurring by the text-recognizing AI, or even the translator AI itself. This can affect how the user perceive certain text in foreign countries for instance. When pointing the camera at a certain object such as a perfume, the user impact the content that is displayed next which goes into the "impacted by user actions" characteristics.

Human-AI interaction design

The paper by Kocielnik et al (2019) points out that AI-systems such as a scheduling assistant will always be imperfect, however it is possible to optimize and visualize it in a way it meets end-users expectations by designing expectation adjustment techniques. The second paper by Amershi et al (2019) investigates Microsoft's 18 Human-AI interaction guidelines and how they can contribute to making AI-infused systems more human-centric.

Select two design guidelines and how Google Lens deviate from them

Google Lens is not an independent application as it is embedded into the standard "Google" app which is just the browser search engine put into an application with extra features such as music recognition, text translation and Google Lens. Google Lens is only seen as a camera-icon next to the input-field and is not explained any further which makes it a "hidden" feature that is not obvious to the user, which deviates from Microsoft's AI design guideline 1. To solve this problem, Google could display an introduction showing the application's AI capabilities during the first interactions with the application. Doing this will ensure that the users know what the application is capable of, resulting into optimized og intended use of the application. Another guideline Google Lens deviates from is G10. If Google can't find what you were looking for or does not recognize the object it will not give you suggestions unless Google knows what kind of object it is. To be fair, there is not much you can do to improve this other than asking the user for more input to base its search on.

Critical discussion of AI-infused systems based on large language models

Bender et al. (2021) discusses the risks of expanding already large language models in their arguments such as an environmental factor and bias in terms of harmful language. The language model size scales with the environmental impact. Bigger the data, the more consumption of

resources. The same goes for training data. Bigger the language model is the more potential and damaging bias can influence the natural language processing systems as it becomes more and more difficult to moderate. Tay (Microsoft's AI Twitter account) is a great example to the bias factor, where the bot it taught to use racial slurs and degrading language when interacting with human twitter users.

Chatbots / conversational user interfaces

As time goes on the importance of AI in HCI and UX design increases. With this increase follows many challenges to almost all stages of a typical design process (Yang et al, 2020). One of which is understanding the capabilities and limits of AI systems. It is difficult for designers to understand what an AI system can do and cannot do, which makes brainstorming, sketching and prototyping processes a lot more challenging. In addition, creating thoughtful interactions is difficult as designers struggle to set user expectations appropriately due to AI systems' lack of consistency and its unpredictable behaviors (Yang et al, 2020), especially when chatbots are wired to communicate with specific topics in mind. Another challenge is the user-experience design of the chatbot. Many chatbots today do not clarify their functionality to the end-user, so many of their core functionality is hidden from the users. Which means the user must explore the chatbot AI system themselves with no meaningful intentions to find out what the system can do or cannot do. Looking at you DNB:)

Revisiting guidelines 1 and 2 – How can they be beneficial for today's chatbots?

The first guideline "Make clear what the system can do" makes sure that the user is told what the system is capable of. Looking at current chatbots, this guideline will benefit both the users and the AI system. If the user knows the capabilities of the chatbot during the first interactions, it will result in optimal and efficient use and better user-experience. It is also beneficial for the company that uses support chatbots for instance. Is it clear what it is capable of, they might not need to contact a human for simple tasks. The same thing goes for the second guideline "Make clear how well the system can do what it can do". Knowing how well the chatbot can do certain tasks important for making the end-user expectations more realistic. This can be achieved by designing user-expectation adjustment techniques.

References:

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

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/en-us/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)