Ergonomics in Design, Vol. 47, 2022, 288–298
https://doi.org/10.54941/ahfe1001948

Making Human-AI Interactions Sustainable: 7 Key Questions for an Ergonomics Perspective on Artificial Intelligence

Myriam Frejus, Dounia Lahoual, and Marion Gras-Gentiletti

EDF Research and Development, Palaiseau, France

ABSTRACT

Recent advances in artificial intelligence (AI) have led to new forms of human-AI interaction and confirmed the need for human-centered AI design. But what are the human factors that need to be addressed for successful AI design? This paper looks at seven key questions critical to designing human-AI interactions in a sustainable way. It also examines recent and emerging factors in relation to challenges posed by earlier forms of automation and AI (such as expert systems). The aim of our research is to propose a framework for multidisciplinary efforts essential to human-centered intelligent system design; it identifies potential activity centered ergonomics contributions and the issues that need to be addressed through situated studies of sociotechnical systems and human activities.

Keywords: Human – AI interaction, Activity centered design, Activity centered ergonomics

INTRODUCTION

Artificial intelligence (AI) has been a subject of academic and industrial study for over fifty years. But today it has undergone a revival, owing specifically to advances in machine learning and big data processing, but also to the massive media attention it has received from leading digital players in recent years. However, although artificial and human processes have been indissociable from the outset, with AI paradigms being shaped by theories of human cognition and psychology all along, the question of how to define human-AI interactions remains hotly disputed. Starting with Allan Turing's earliest works, and the philosophical divide between Engelbart and McCarthy as to whether AI should be used to enhance or to replace human capacities, AI has been approached as a competitor to human cognition, a question that is still a matter of debate. Likewise, there is no consensus among media and science experts on where AI will take us, and this too is nothing new. This is why the automation of systems and the first two waves of AI (and expert systems, in particular) can no doubt shed light on potential third-wave AI scenarios (Xu, 2019). Nevertheless, many experts, both in AI and Human-Computer Interaction (HCI), have argued that it is no longer a question of defining who is in control, but rather of establishing a sound complementarity between artificial and human systems (Licklider, 1960; Roth, Bennett & Woods, 1987;

Dekker & Woods, 2002). We are headed towards a form of mutual reliance to achieve maximal performance. (Dellermann et al., 2019) envision hybrid intelligence systems "combining human and artificial intelligence to behave more intelligently than each of the two could be in separation". Within this context, numerous experts are calling for human-centered AI design (Amershi et al., 2019; Inkpen et al., 2019; Schneiderman, 2020). In their paper presented at CHI 2020, Yang et al. identify the challenges facing research communities working on issues of user-centered and lean startup AI design, who "do not know how to bring a human centered view to AI." Unfortunately, these issues are rarely approached from the perspective of human behavior, and the purpose and potential role of ergonomists have not been clearly defined. In short, human-centered design still considers the human user to be a cog in the design process machine (for example, human input is integrated into the machine learning pipeline), neglecting the fact that all human activities are carried out within an overarching sociotechnical system (Inkpen et al., 2019). Activity centered and cognitive ergonomists use situated analysis to understand current and future human activities in natural contexts and to acquire an overview of the entire sociotechnical system. But what are the challenges specific to human-centered AI, and how can ergonomists help resolve them?

This paper presents a synthetic overview of seven key challenges and identifies design aspects needing to be addressed in terms of ergonomics contributions. Its objective is to provide a framework for multidisciplinary efforts essential to human-centered AI design involving, more specifically, human factors experts. Our work is informed by relevant research in HCI, CSCW, and in the human and social sciences, as well as by our own involvement in industrial projects for the deployment of AI-capable chatbots, VA interfaces for consumer use, and home automation systems with integrated AI-IoT (Fréjus & Martini, 2016; Lahoual & Fréjus, 2019a, 2019b; Gras-Gentiletti et al., 2021).

7 KEY QUESTIONS FOR DEVELOPING HUMAN-AI INTERACTION

How can AI be Put to Good Use? Clearly Identifying the Needs AI may Potentially Meet

Recalling this may seem trivial, but AI is only a means to an end, and as with any other system, its usefulness is commensurate with the services it provides. The first of the "23 principles" outlined at the Asilomar Conference on Beneficial AI in 2017 states that "the goal of AI research should be to create not undirected intelligence, but beneficial intelligence." And yet the beneficial nature of AI is by no means a given. In an article published in 1983, Bainbridge condemns the shortcomings of automation, liable to spread rather than eliminate the problems of human-machine interactions. Woods (1996) denounces the cognitive overload caused by expert systems. In other words, the value proposition of AI is not always evident. Take, for example, interaction chatbots that simply replicate interaction functionalities that are already available (and more efficient). Therefore, identifying the circumstances of real user activity as a social practice can make it possible to accurately define proactive technology behaviors, to "*make visible the mundane, [the] 'seen but unnoticed*'" (Hyland et al., 2018). Conversely, an incorrect description of human factors may call into question user appropriation of systems. (Yang & Newman, 2012) show for example that the predictive features of the NEST "smart" thermostat ignore the ways in which temperature is collaboratively negotiated as a social practice. The technology is therefore either abandoned or misunderstood. Seniors, for example, may be resistant to using home automation technologies, despite the obvious assistance such AI-driven devices could provide around the house. Many elderly people are concerned that such systems will make them less autonomous by encouraging a lazy lifestyle and thus lead to greater dependency, the very thing we fear most when it comes to ageing (Portet et al., 2013). The study (Lahoual & Fréjus, 2019a, 2019b) on how vocal assistants are used shows that functional usefulness (in this example "dictating a grocery list" to a VA) can only be evaluated when real use situations are taken into account ("going grocery shopping"); usefulness is challenged when real uses are neglected : "*I tested the grocery list function. I used it a lot at first, but then, I went back to paper lists (…) You just say 'Google, add this to my grocery list.' But once you got to the supermarket, you had to get out your phone, read the list, push the shopping cart, use the thing to scan articles… It's too complicated.*" (Woman, age 50). Furthermore, human activity tends to be approached from the angle of single user-machine interactions, ignoring the activity's collaborative dimension, whereas, given the ubiquitous possibilities of AI – the multitude of places of use, of devices, user profiles, and temporalities – cooperation is a defining factor. Innovation and AI expectations have seen a growing number of technology-driven projects to the detriment of projects directed solely by the needs of future users, revealing an innovation-centered rather than user-centered vision of design (Gras-Gentiletti et al., 2021). Current uncertainties as to what intelligent systems will ultimately be able to do and how they will be able to do it (what Yang et al. (2020) call the *capability uncertainty*) renew questions of how future human-machine interactions can be anticipated, and sustained over time. Which brings us to the second of our key challenges.

How can a Sustainable Relationship be Developed by Anticipating and Guiding Changing Human-AI Interactions?

AI technology has opened the way for multimodal, ubiquitous, evolutive, context-aware, and natural human-machine interactions. One of its first mature technologies, the voice-commanded virtual assistant or voice assistant (VA) has emerged as the prototype of natural interaction systems. However, several studies have outlined an array of difficulties that users encounter (Velkovska & Zouinar, 2018; Lahoual & Fréjus, 2019b). The illusion of fluid and natural interactions tends to create a gap between user expectations and the real capabilities of VA devices, which are experienced as a source of substandard use and frustration. The filmed observations and interviews of Lahoual & Fréjus (2019b) reveal activities of supervision, verification, diagnosis, and problem-solving. These are not only time-consuming; they interrupt the flow of user routine. Instead of facilitating the additional task as expected, the voice recognition technology is an impediment that places the user in the role of the assistant's assistant! The risk, of course, is seeing a drop in use and even the abandoning of a technology altogether. However, other users may accept these failures and forgive system errors more readily when they expect to see future improvements (Lahoual & Fréjus, 2019b). Although hoped-for functionalities may not be offered or may for the time being fall short of user expectations, the system's evolutive capacity constitutes one of its greatest strengths and can potentially determine system acceptance.

This is a distinctive feature of artificial intelligence: the interactions and services offered by AI can evolve and adapt over time. Adaptability therefore is no longer seen from the user and designer perspective alone. While this constitutes a potential advantage in terms of gaining user acceptance, it also risks seeing system interactions and uses being constantly called into question. The challenge therefore is to design human-machine interactions that can be sustained over time (Fréjus & Martini, 2016). How to design lasting interactions is a question that applies to any interactive system, for that matter, the aim of which is to create successful human-system couplings (Woods, 1996) that make it possible for systems to at once assist human operators while making it possible for them to evolve (such as by acquiring new skills), even as the individual user's needs and abilities continue to change. But this notion of a sustainable relationship cannot be reduced to challenges of ethical, social, and legal matters alone. It also signifies that the tightly coupled interactions in question need to be designed to last over time, namely by integrating all aspects of design, including utility, usability, desirability, and so forth. However, what is new today is the capacity for both the human practitioner and the system to respectively evolve, offering opportunities for development but also risks of loss (of skills, in particular), as well as breakdowns and adverse consequences. On the design side, this means that systems evolutions and user-system couplings need to be both considered and, more importantly, anticipated. Today's system designers need to therefore look at how the joint "intelligences" (machine and human) inherent to interactions can be deployed, according to the logic of plasticity. The development of a system and its coupling with human users is directly tied to the study of human attributes, thus reinforcing the need for human factors contributions in areas such as activity anticipation/simulation, behavior and behavior evolution modeling, longitudinal human-AI coupling analysis, shared environments and mutual intelligibility, design principles and criteria for interaction sustainability, and so forth.

What Form(S) Should AI Take?

Designing digital systems with "intelligence" capabilities raises the question of what form they should take. Should AI technologies have distinctive personalities? Should they be humanlike in design? If interactions are meant to be as natural as possible and the system is capable of adapting to specific contexts, human users will naturally employ the same interaction codes used in human-human conversation: people who use voice assistants often use polite language like "thank you" and "good-bye" (Lopatovska & Williams, 2018).

But the same users will also have expectations about AI information delivery, as well. For it is generally agreed today that customer-facing "autonomous" systems are expected to have personalities. It is not uncommon in industrial environments to encounter cobots humanized by operators, with a hat and glasses, for example, and seen as a full-fledged colleague as a result. Amazon's wake-word, the name "Alexa," and voice assistants with distinctive voices are consistent with this line of thinking. Anthropomorphism can therefore be a way for creating human-machine affinities that can contribute to service improvements. Animals can be attributed with human characteristics as well, like the stuffed baby seal, Paro, shown successfully to treat pathologies related to ageing (Jøranson et al., 2017). Humanization therefore is not only a question of outward appearance. It is about the attribution of personality and affect. Humanization can also be a source of embarrassment stemming from our response to the presence of an artificial agent or to agent personification choices (giving orders to a female, for example, can be indicative of a sexist behavior (Lahoual & Fréjus, 2019a). Therefore, our acceptance of AI interfaces with humanlike design and our sense of affinity for them come with strings attached. First, the systems must behave in a consistently "human" way, in other words, not to interact like a human at certain times and like a machine at others. This is the case of chatbots, whose conversations at first seem humanlike but may later, over the course of the exchange, become mechanical and even ineffective (Budiu, 2018). Users must also be able to decide on levels of humanization through personalization options: to choose virtual assistant wake-words, for example, or the type of language to be used ("coded" or natural). And indeed, humor, playfulness, and emotional response are key aspects of interactions with AI technologies, and with machines in general. Let us not forget that cognitive realism and perceptive realism are two very different things. The question of humanized AI necessarily involves exploring the interpersonal aspects at play in human-human relations, by integrating notions of interaction pleasure. This area of exploration could constitute a real step forward by shifting the focus from human-machine interaction to human-anthropomorphized interaction, although we must bear in mind that an efficient human-machine relationship may be the better choice. Thus, human-centered AI design involves setting goals according to functionalities, context of use, and technological maturity.

How will Human-Machine Interactions be Modified by "Autonomous" Systems" (and What Role will the User Play)? Defining the Capacity of Sociotechnical Systems for Action

The potential autonomy of AI-driven systems powered by self-learning and adaptative capabilities could make AI a gamechanger in terms of human-machine interaction. The notion of system autonomy has been widely criticized (Schneiderman. 2020). The automation of system results in various changes with regards to the role and the status of the user. In the first place, we are no longer just users but trainers who teach the machine to do a good job. The other notable evolution that the advent of autonomous systems involves is the modification of assisted activities. In other words, once an element of an activity is delegated to a system, the user becomes a supervisor and controller. (Gras-Gentiletti et al., 2021) show how the work of jurists is altered by the introduction of a legal chatbot, because jurists now receive demands for validating chatbot answers. The user participates, too, by personalizing the system's role. Users need to be involved in both courses of action and problem-resolution decisions made by machines at critical moments (Costanza et al., 2014). It can thus be maintained that by assigning the required behavior to the system, the users build its intelligence based on their own understanding. User involvement can also mitigate a lack of available data. Because the fewer the AI-building events, the more the norms and tenets for determining functionalities are used by designers (staff recommendations, security/safety rules, traffic regulations, etc.); and yet, these models fail to account for the variations, contingencies, and adjustments to which humans respond and which make the system work. Real activity modeling studies must therefore be made available early on, to identify the unpredictable behaviors and forms of human adjustments that shape situations, and to involve users in both the defining and the functioning of the system (Nilsson et al., 2018). Interaction is the means for seeing autonomous systems "grow", whether human or artificial, and respond appropriately in each environment. Negotiation models could be one source of inspiration for enabling systems to provide additional information in preparation of future user returns and to provide alternate forms of justification in response to user dissatisfaction (Pollack, Hirschberg & Webber, 1982).

None of these questions are particularly new. Since the emergence of the first automata and expert systems, research has been exploring issues of transforming the system-as-prosthesis-centered vision into a system-as-assistant-centered vision, with a particular focus on questions of human-machine relations (Woods, Johannesen & Potter, 1991). The question of design thus becomes one of how to configure sociotechnical assemblages that make possible the construction of mutual intelligibility and take into account the asymmetric relations and differences between human agents and artefacts. The system therefore cannot be defined without human actors.

What Kind of Technology and Situation Intelligibility do Human Actors Expect? Moving Beyond Explainability

Appropriating a system, interacting with it, supervising it, being able to take control of it, understanding its possibilities, trusting it, and anticipating what it will do – these are among the many expectations users have when interacting with an intelligent system or when immersed in an intelligent environment. But actions like these depend on the prerequisite that the system and system behavior be intelligible to the user (Bellotti & Edwards, 2001). However, it is important not to equate intelligibility with the explainability of AI systems. Although intelligibility can at times be derived from explainability, it cannot be reduced to explainability alone. Explainability is primarily "oriented" toward algorithmic verification and validation. As such, it is mainly a question for AI makers. Intelligibility, on the other hand, takes into account the human view of the system. It needs to be equal parts computational and social. Computational intelligibility is the system's capacity to account for own behavior so as to better support human-machine interactions (accountability); social intelligibility is the consideration of social contexts to determine what people do in particular social circumstances. Therefore, a system's behavior can be understood, but if its behavior is not appropriate or socially acceptable, it will remain unintelligible. System intelligibility takes into account the system's observable behavior in relation to a set of particular circumstances, not just behavior on its own. When working on making a system appropriable, it is not a question of explaining the internal workings of the system (the how) but of justifying its behavior (the why) (Woods, Johannesen & Potter, 1991; Woods, 1996). A qualitative approach must therefore be taken to how recommendations are justified. Necessary user explanations and clarifications need to be determined from the perspective of action usefulness, referring to situations where human activity can be aided by intelligent systems. As such, ergonomists have a twofold part to play: to determine the needs for justifications and for information about how the system works, and to design the system's formal aspects and content in terms of its qualitative value.

How can Systems Respond to What is Implied and Intended? Contextual Relevance is the Key to Intelligence

The presumed intelligence of next generation machines presupposes an interest in their capacity for contextually relevant responses to specific situations. However, since context recognized by the machine and the actual context of the user's inquiry are two different things, there is a fundamental asymmetry to their two situations (Suchman, 1999). Next generation artificial intelligence artifacts need to be able to manage the inevitable discrepancy between the context of the user (influenced by background, elements of situational relevance, intentions, preferences) and that of the machine (based on material and environmental factors, and learning). Implicit to the claims that the improvement of the systems' technical performances raises expectations of new capabilities, such as natural interaction, is the underlying question of context-appropriate response. A study by Velkovska & Zouinar (2018) shows how VA users employ expressions that rely on situation for significance, including terms like "here," "there," and "now," and such words as "the" and "that," in reference to earlier ideas. AI users also expect autonomous systems to be able to transfer context and therefore make assumptions from one interaction to the next. For example, Budiu (2018) shows that the user of a banking chatbot who happens to have two credit cards is forced to clarify which account the query refers to each time he asks a question. Context awareness is therefore multidimensional but is also a question of providing the right service at the right time, without obstructing user action (Nilsson et al., 2018). Such questions of coherent in-context interpretations are major barriers in artificial intelligence development, both for algorithm designers and cognitive science communities exploring human behavior to define appropriate response. On one side, there is the issue of teaching algorithms to recognize common sense, and on the other, that of developing a model of human action in order to qualify situations of use. The successful implementation of future systems requires the integration of ergonomics findings (for example, semiological studies of courses-of-action (Theureau, 2003) or the CSCW analyses of sociotechnical systems) to direct next generation context- and situation-sensitive learning, to associate contexts and concrete services, and, lastly, to design meaningful interactions that are contextually relevant.

How to Guarantee a Legal, Ethical, and Political Framework for AI?

Addressing human factors in AI implies dealing with ethical issues related to both the use and the design of such technologies. This point is no doubt the most widely recognized and debated topic among issues related to the social and human impacts of artificial intelligence. For that matter, it has been around since the 1960s and various national and international commissions have since issued recommendations and warnings on the subject. So how can we guarantee the design of ethical AI systems? The think tank Doteveryone recommends not just the training of computer scientists in questions of ethics, but a greater involvement of human factors specialists in AI design and decision-making. One obvious problem is that the notion of "ethical" is a highly personal matter. Vaughan (1996) calls this the "normalization of deviance," a social phenomenon that blinds teams to past errors once an error has become accepted. She maintains that it is important to develop a culture of challenging unjustified orders and consensus – a culture more in line with that of the human and social sciences, another reason why their practitioners need to be included in design processes ensuring the development of responsible technology. Privacy and data protection have given rise to a set of specific design principles (Danezis et al., 2014). Many privacy and confidentiality concerns can be resolved through design interaction possibilities and system intelligibility enabling users to control their data and how it is used. User-centered design makes it possible to develop an intelligible, usable, and reversible consent path guaranteeing user control over what happens to their data. The legal aspects of AI are still wide open to debate, and system automation has lent renewed urgency to one question in particular: responsibility. AI explainablility and the possibility of demystifying "black box" techniques are key, as tied in with the earlier-mentioned questions of intelligibility, justification, and trust.

Therefore, ethical and legal questions reconfirm the need for AI projects to be informed by social, human, and legal science contributions. Karwowski (2018) suggests rounding out the team with the creation of "*human and artificial factors*" (AI-ergonomics or cyber ergonomics) as a "*subdiscipline concerned with the understanding of interrelations between humans, intelligent agents, and other artificial cognitive agents in any social context (…) and the development of human-centered principles for the design and integration of artificial intelligent systems beneficial to humankind.*"

CONCLUSION

The aim of this paper is to look at the questions that human-centered AI system design inevitably raise and that constitute significant scientific barriers. Recent technological developments, and namely the experimental features they make possible, raise specific new questions or reformulate others that have emerged with automation or earlier waves of AI. These questions demonstrate the myriad dimensions that need to be considered to avoid reproducing the disasters of earlier AI systems but also to avert what amounts to a technology-driven design aimed solely at acquiring social and technological acceptance. It is agreed today that human and artificial agent relations need to be approached in terms of complementarity and, as such, can only be defined by a human-centered design approach. We have insisted on the importance of establishing human-AI interactions from the perspective of sustainability, beyond social, ethical, and legal criteria, and capable of responding to the key questions we have enumerated here. In keeping with Suchman and Weber (2015), we are interested in moving beyond the question of interaction-based design for a conceptualization of configurations involving human and artificial agents. To meet these ambitions, activity centered ergonomists have conducted situated analyses of sociotechnical systems and of human activities, which make it possible to approach design from the perspective of in-situ actors and ultimately to guarantee the technology's appropriation. Their findings allow designers to better understand what users expect from so-called intelligent systems in both functional and interactive terms. Furthermore, approaches like these shed new light on such regularly mentioned categories as privacy, fairness, and explainability, which by pointing to missing problematics (collective activity and multiple user profiles, for example) or by challenging certain definitions (notions of wrong system behavior, for example), can supplement existing guidelines for the design of human-AI interaction (Gras-Gentiletti et al., 2021).

REFERENCES

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., and Horvitz, E. (2019). Guidelines for human-AI interaction. In Proceedings CHI 2019. ACM.

Asilomar principles (2020). Future of Life Institute.

Bainbridge, L. (1983). Ironies of Automation. Automatica, Vol. 19 (6), pp. 775–77.

Bellotti, V., and Edwards, K. (2001). Intelligibility and accountability: human considerations in context-aware systems. Human Computer Interaction, 16, 2.

Budiu, R. (2018). The User Experience of Chatbots. Rapport Nielsen Norman Group.

Costanza, E., Fischer, J., Colley, J., Rodden, T., Ramchurn, J., and Jennings, N. (2014). Doing the laundry with agents: a field trial of a future smart energy system in the home. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 813–822.

Danezis, G., Domingo-Ferrer, J., Hansen, M., Hoepman, J., Metayer, D., Tirtea, R., and Schiffner. S. (2014). Privacy and data protection by design-from policy to engineering. ENISA.

Dekker, S. W. A. and Woods D. D. (2002). MABA-MABA or abracadabra? Progress on human-automation coordination. Cognition, Technology, and Work, 4, 240–244.

Dellermann, D., Calma, A., Lipusch, N., Weber, T., Weigel, S. and Ebel. P. (2019). The Future of Human-AI Collaboration: A Taxonomy of Design Knowledge for Hybrid Intelligence Systems. In HICSS. Hawaii, USA.

Fréjus M. and Martini. D. (2016). Why Energy Consumption Feedback Is not (Only) a Display Issue. In A. Marcus (Ed.): DUXU 2016, Part II, pp. 461–471, Springer International Publishing Switzerland.

Gras-Gentiletti, M., Fréjus, M., Bourmaud, G. and Decortis, F. (2021). Interpersonal mediations and learning: what are the links? An example of the chatbot appropriation. IEA 21st triennial congress.

Hyland, L., Crabtree, A., Fischer, J., Colley, J. and Fuentes C. (2018). What do you want for dinner? Need anticipation and the design of proactive technologies for the home. Computer Supported Cooperative Work (CSCW), 27:917–946.

Inkpen, K., Chancellor, S., De Choudhury, M., Veale, M.l and Baumer. E. P.S. (2019). Where is the Human? Bridging the Gap Between AI and HCI. CHI'19 Extended Abstracts, May 4–9, Glasgow, Scotland UK.

Jøranson, N., Pedersen, I., Mork Rokstad, A-M. and Ihlebaek. C. (2017). Change in quality of life in older people with dementia participating in Paro-activity: a cluster-randomized controlled trial. J Adv Nurs. Jun;73(6):1531.

Karwowski. W. (2018). The human use of artificial intelligence. 20th Congress of the International Ergonomics Association (IEA 2018).

Lahoual, D. and Fréjus, M. (2019a). Développement et utilisation d'écosystèmes d'objets connectés: concevoir des dispositifs et des services grand public appropriables et durables pour une gestion connectée de l'habitat. Technical Report EDF Research & Development.

Lahoual, D. and Fréjus., M. (2019b). When users assist the voice assistants: from supervision to failure resolution. In Proceedings CHI'19.

Licklider, J.C.R. (1960). Man-Computer Symbiosis. IRE Transactions on Human Factors in Electronics HFE-1, 1: 4–11.

Lopatovska, I. and Williams. H. (2018). Personification of the Amazon Alexa: BFF or a Mindless Companion? In Proceedings of ACM CHIIR conference.

Nilsson, T., Crabtree, A., Fische J. and Koleva. B. (2018). Breaching the Future: Understanding Human Challenges of Autonomous Systems for the Home, July 2018.

Pollack, M. E., Hirschberg, J. and Webber, B. (1982). User participation in the reasoning processes of experts systems. AAAI-82 Proceedings.

Portet, F., Vacher, M., Golanski, C., Roux C. and Meillon B. (2013). Design and evaluation of a smart home voice interface for the elderly: acceptability and objection aspects. Pers Ubiquit Comput (2013) 17:127–144.

Roth, E. M., Bennett, K. B. and Woods D. D. (1987). Human interaction with an "intelligent" machine. Int. J. Man-Machine Studies 27, 479-525.

Schneiderman, B. (2020). Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy, International Journal of HCI, 36:6, 495-504.

Suchman. L. (1999). Human/Machine Reconsidered. Conference Sociality/Materiality, Brunel University, UK, September 9-11, 1999.

Suchman, L. and Weber J. (2015). Human-Machine Autonomies. In N. Bhuta, S. Beck, R. Geiβ, HY. Liu, and C. Kreβ (Eds.), Autonomous Weapons Systems: Law, Ethics, Policy (pp. 75–10). Cambridge, UK: Cambridge University Press.

Theureau, J. (2003). "Course-of-action analysis and course-of-action centered design", Handbook of cognitive task design, pp. 55–81.

Vaughan. D. (1996). The Challenger Launch Decision. Risk Technology, Culture and Deviance at NASA, Chicago, The University of Chicago Press.

Velkovska J. and Zouinar. M. (2018). The illusion of natural conversation: interacting with smart assistants in home settings. CHI'18 Extended Abstracts, Canada.

Woods, D.D. (1996). Decomposing Automation: Apparent Simplicity, Real Complexity, In R. Parasuraman and M. Mouloula, editors, Automation Technology and Human Performance, Erlbaum.

Woods, D. D., Johannesen, L. and Potter, S. (1991). Human Interaction with Intelligent Systems: An Overview and Bibliography. SIGART Bulletin, Vol. 2, No. 5.

Xu, W. (2019). Toward Human-Centered AI: A Perspective from Human- Computer Interaction. Interactions, july-august, ACM.

Yang, R. and Newman, M. W. (2012). Living with an intelligent thermostat: advanced control for heating and cooling systems. Proceedings of the 2012 ACM Conference on Ubiquitous Computing, 09/2012.

Yang, Q., Steinfeld, A., Rosé, C. and Zimmerman, J. (2020). Re-examining Whether, Why, and How Human-AI Interaction is Uniquely Difficult to Design. CHI 2020, April 25–30, 2020, Honolulu, HI, USA.