Human-Centered Artificial Intelligence: Beyond a Two-Dimensional Framework

Matthew Pacailler(✉), Sarah Yahoodik, Tetsuya Sato, Jeremiah G. Ammons, and Jeremiah Still

Old Dominion University, Norfolk, VA 23508, USA
Mpaca001@odu.edu

Abstract. Shneiderman's Human-Centered Artificial Intelligence (HCAI) framework suggests that high human control of automation is necessary to create reliable, safe, and trustworthy systems. The HCAI framework demonstrates that there is no need to sacrifice human control when incorporating higher levels of automation. We propose that Shneiderman's two-dimensional framework is static and unable to incorporate contextual factors such as the decision for a human-in-the-loop system, cognitive limitations of the user, and user characteristics. The HCAI framework, while an essential foundation, ought to reflect the flexibility of AI systems, while meeting individual differences and situational requirements.

Keywords: HCAI · Artificial Intelligence · Automation

1 Introduction

Artificial Intelligence (AI) is a system that performs a specific task, drawing on a single human ability such as visual perception, reasoning, and understanding context [1]. AI is of the utmost importance and continues to improve. AI's history is one of hope and fantasy [2]. In recent years, basic research on AI has been focusing on robots and pattern recognition [3]. Major companies have begun incorporating AI into their systems. For example, Microsoft announced real-time translation robots and image recognition products, along with Facebook. Amazon has incorporated autonomous robots into its delivery system. In addition, many universities are helping develop AI, leading to the creation of robot cars, cleaning robots, and four-foot walking robots [3]. These recent advancements often blur the boundaries between autonomy and automation systems.

Automation varies across levels and stages [6]. The most widely used framework describing the levels of automation is Sheridan and Verplank's Framework, which bases the levels of automation on a continuum that ranges from low (Level 1) to high (Level 10) [4, 7]. As the autonomy of automation increases, so does the level of automation. Thus, autonomy in the automation is highest at the highest level of automation. A trade-off between human control and autonomy is implicit in this framework; if the amount of human input needed increases, the level of automation lowers. Thus, high human control is at the lowest level of automation. In addition to levels of automation, there are 4 different stages of automation [5, 6]. Stage 1 of automation consists of acquiring

J. Y. C. Chen et al. (Eds.): HCII 2022, LNCS 13518, pp. 471–482, 2022.
https://doi.org/10.1007/978-3-031-21707-4_33

information from the environment. Stage 2 consists of using the information and analyzing it. Stage 3 occurs when the automation chooses or decides a course of action based on the previous analysis. Stage 4 carries out the chosen action. It is important to note that levels of automation can vary within each stage of automation. Thus, there are two dimensions, where a higher level and later stage equate to more automation [6].

Shneiderman introduces a topic referred to as Human-Centered Artificial Intelligence (HCAI). HCAI explores the interaction between HCI with AI. Shneiderman's definition of AI suggests that AI systems can perceive, think, decide, and act. Such systems can analyze emotions, adapt to a changing environment, and have equal status to a human being [8–10]. Classically, the goal of HCAI is to maintain a human-centered view, creating a future where technology is built around human control to create a reliable, safe, and trustworthy (RST) environment. Doing so will keep humans in power by creating systems that allow high levels of human control and high levels of automation. Thus, devices should be made to amplify human ability, empower people, and ensure human control [8–10]. As a result, technology should not be looked at as divine beings, but as a tool or appliance, that allows humans to enhance their abilities. Shneiderman provides this framework to overcome the stigma around more automation leading to less human control. Shneiderman refers to HCAI as the Second Copernican Revolution [10]. Similar to how many believed earth was the center of the solar system before Copernicus developed a sun-centered model. Shneiderman wants researchers and designers to move away from the mindset of AI being the focal point, and humans revolving around AI, and instead lean towards a human-centered model. Thus, instead of focusing on how to improve AI and machine autonomy, we should focus on improving the user experience. Shneiderman believes that an HCAI approach will eliminate the fear associated with a future of autonomous robots taking over the world, or on a much smaller scale, taking over jobs. He noted that just because humans are of focus, does not mean that designers should build products that emulate the appearance or behavior of a human. Such products tend to lead to fear.

While designing AI around the needs of the user is an important consideration for RST environments, our motivation is to expand the conversation beyond a two-dimensional framework. This paper contributes to the expansion as a translation of human factors literature from automation to the HCAI framework. We believe that it is important to highlight the process underlying the high human and automation control quadrant to have a better understanding of designing automation around the human. We do not seek to replace the framework but instead enrich the conversation with human factors considerations.

We would like to address that Shneiderman's definition of AI is broad. This definition includes AI with machine learning algorithms and adaptive systems, but also automation in general [8]. Although the argument can be made that the boundary between AI and automation is fuzzy, we believe there should be a distinction between automation that uses sensors and AI capable of learning and making decisions.

Furthermore, human interaction with AI will be different depending on how the AI functions. As van Berkel and colleagues suggested, there are three paradigms of human interaction with AI: intermittent, continuous, and proactive [34]. Intermittent interaction is described as a conversation where a user inputs a cue, the AI responds then the user will react. Continuous interaction is like commentary where user input is now continuously monitored and given suggestions by AI. Proactive interaction has the AI monitoring the environment with sensors and can complete make decisions and act with or without human input. High human control is desirable for intermittent and continuous human-AI interaction, but perhaps less critical with proactive human-AI interaction in cases that are not time sensitive and safety critical. Perhaps what we seek with proactive AI interaction is not control but coordination by having AI that are transparent. Although Shneiderman believes in a teammate fallacy when designing AI [10], there is evidence suggesting that user perception and expectation of teamwork exist for users [35].

Although we believe there is a distinction between AI and automation there is still a strong connection that should not be ignored. Indeed, The HCAI framework is a beneficial framework for designing automation and AI for RST systems. However, we suggest that this framework can be enriched with human factors considerations. We that a two-dimensional approach to AI does not consider contextual factors such as the decision for a human-in-the-loop system, cognitive limitations of the user, and user characteristics. We highlight that human control over automation is a continuous process where the operator observes the feedback from the current system and adjusts the automation when necessary. As with all new frameworks, there are assumptions that must be carefully considered to grow their utility.

2 Context of Human-in-the-Loop

Wickens and colleagues suggest four purposes of automation: performing functions that humans cannot perform due to inherent limitations, alleviating high workload of tasks that humans can perform, augmenting or assisting in human performance, and economic reasons [11]. The HCAI framework accounts for alleviating humans of workload and augmenting human performance but disregards the other purposes of automation. It may be advantageous to consider the purpose of the automation before human involvement.

When referring to a system that involves human-in-the-loop, the human controls the AI and monitors the situation [12]. Human-in-the-loop is most beneficial in dynamic environments where a failure of the AI can lead to disastrous consequences in time-sensitive or safety-critical situations. To operate in these environments, AI needs open-world algorithms. These algorithms enable the AI to update their database about unknown objects and are capable of decision-making using this information. Example situations where AI benefits from human-in-the-loop are Urban Search and Rescue (USAR) scenarios and adversarial tampering.

USAR scenarios consist of a human-robot team where humans understand the physical layout of a building but are physically removed from the environment. The robot is sent in with established goals, perceives environmental stimuli, and communicates acquired information. Humans choose to update the initial goal based on new information [13]. This scenario fits Shneiderman's HCAI ideal, where high human and automation control are desirable. Due to unknown elements, there are situations where the AI misinterprets a signal or incorrectly changes the goal of the mission. When humans monitor the AI, ideally all false alarms will be filtered out, and correct detections will be dealt with appropriately.

Adversarial tampering is another situation where it is beneficial for human-in-the-loop with AI. The AI needs state-of-the-art out-of-distribution detectors (OOD) to accomplish open-world algorithms. This is incorporated with a type of machine learning that is used to filter out unwanted or ambiguous inputs in dynamic environments. However, even with state-of-the-art OODs, Sehwag and colleagues found these OODs can be evaded or manipulated with relative ease [14]. Therefore, it is beneficial to have human-in-the-loop when malicious intent is known because humans can monitor the situation for AI failures.

Although there are situations where high human control is beneficial, there will be times where it is not beneficial or even necessary. For scenarios such as economic purposes, enabling humans to perform tasks they are unable to accomplish, or decision-making made by users who do not fully understand AI, it would be beneficial to exclude humans from the system. Although transparency of the AI may keep human involvement in some of these scenarios, we assume that there is a complex interaction of the user with AI preventing ease of implementing transparency, such as a proactive interaction paradigm described earlier [34]. Human-out-of-the-loop systems have no human in physical control nor monitoring the situation or have humans in physical control but not monitoring the situation [12]. These systems thrive in static environments with predictable conditions. This is due to most AI algorithms running on a closed-world assumption [15]. This logic indicates that unknown objects are not important, which in return cannot be processed by the algorithm under any circumstance. Situations where time is unlimited, and decisions made by AI will not result in endangering life could benefit from removing the human from the system. Detection of malware is a situation where keeping humans out of decision-making could increase cyber security. Older versions of anti-virus programs require users to make decisions about every virus encountered. When a user is not an expert in cyber security, this may lead to incorrect decisions. Fortunately, modern software can automatically block or quarantine infected files. Machine learning techniques can be incorporated to detect new and advanced malware [16]. In this situation, high levels of human control could lead to less cyber security when the user has a minimum understanding of malware.

At home service robots demonstrate a situation where a dynamic environment can incorporate AI and humans-out-of-the-loop. This situation can provide certain humans with a service that they cannot perform. For example, if a person cannot walk in their home without assistance, the service robot can deliver a requested item. The user can give a command to the robot requesting an unknown item or retrieving an item when the location is unknown [15]. The user has control over the robot but is unable to monitor the robot's decision making. For example, when trying to retrieve a water bottle from a different room the robot will have to decide where to look. When the robot is incorrect, it can update its databank and try again until the correct decision is made. These low consequence situations enable a human-out-of-the-loop system.

3 Cognitive Limitations

To maintain a high level of human control and automation control, Shneiderman provided suggestions for redesigning various products or services based on the Prometheus

Principles [8]. Shneiderman's principle emphasized providing informative feedback on the visual interface allowing users to understand and control the automated system. For example, patient-controlled analgesia devices can be designed to allow sensory feedback while the patients control the pain medication, creating RST systems. Indeed, providing informative feedback can potentially achieve a high level of human control and automation control. However, in a multitasking environment, maintaining high levels of human control can be challenging due to the user's cognitive limitations. Notably, implementing AI can potentially direct the user's attention away from the primary task since users are tasked to monitor the AI [17], degrading primary task performance. The high level of automation could constrain the operators from performing multiple tasks in various professional environments, including air traffic control [18] and aircraft cockpit [19–22]. It is critical to consider the attentional limitations of highly automated systems within the context of HCAI framework.

Air traffic controllers typically monitor the aircraft and navigate the aircraft to the right path. Implementing automated alert systems to air traffic control can direct the air traffic controllers to critical events on the visual screen. However, the automated alert system could disrupt the air traffic controller's primary task, directing attention away from the primary task and degrading primary task performance [18]. Alternatively, pilots operating the aircraft could overlook the automated alert system's notification due to the reliability of the automation [19, 20] and the attentional demand imposed by the primary task [21, 22]. Several incident reports indicated that implementing high levels of automation control allowed operators to behave counterproductively. For example, the National Transportation Safety Board [23] reported that Asiana Airlines Flight 214 collided at San Francisco International Airport, resulting from the pilot's misuse of the autothrottle system. Specifically, the pilot failed to recognize that the autothrottle system did not control the airspeed while approaching the runway. The pilot's misuse of the autothrottle system is attributed by the pilot's overreliance on the automation [20, 24]. In aviation, AI has been widely used to optimize various tasks such as a pilot's flight operation. However, the reliability of AI and the attentional demand of the concurrent task could potentially constrain the system from establishing high levels of human control and automation control. This challenge can be best described by referring to the theoretical framework of attention allocation.

Theoretical models of attention allocation could potentially explain the degraded performance in a multitasking environment involving high levels of human control and automation control. Particularly, the unitary resource model of attention [11, 25] could indicate possible limitations of the HCAI framework. The unitary resource model of attention indicates that users have limited attentional resources to allocate to a particular task [11, 25]. Attentional resources refer to a unitary group of mental energy that is allocated to different information processing stages, supporting the user's mental processing [25]. Within the unitary resource model of attention, task performance depends on the correspondence between the attentional demand imposed by the task and the number of attentional resources. Particularly, users' task performance can degrade when the attentional resources supplied does not suffice the attentional demand. Alternatively, users can establish successful task performance when the attentional resources supplied suffice the attention demand. Although the attentional limitation for using highly automated AI systems may not be apparent in a single task environment, implementing in a multitasking environment could potentially degrade the user's task performance due to the high attentional demand for monitoring the AI. Based on the unitary resource model of attention, system designers are challenged to alleviate the attentional demand imposed by the automated task in a multitasking environment. Thus, it is critical to consider alternative approaches to address the cognitive limitation for implementing RST systems in multitasking environments. One consideration for maintaining high levels of human control and high levels of automation is to reduce the complexity of the automation. The complexity of the automation is a critical factor that increases the attentional demand of the automated task [17]. Designing simpler automated systems could potentially reduce the attentional demand, allowing RST systems to maintain high levels of human control and automation control.

4 User Characteristics

Research on human-automation interaction has centered on professional users in tightly controlled (and regulated) safety-critical fields such as aviation, air traffic control, nuclear power, patient care, and military technology [26]. Embedded in most of this research is an implicit or explicit expectation that operators of automated systems are highly qualified, knowledgeable, and invested in avoiding adverse outcomes. The concept of having high levels of automation and high levels of human control in this context is feasible; experts by the very nature of their experience can leverage complicated, interdependent automated systems to fit their needs and goals. However, more and more, we are seeing AI and automation applications seeping into everyday life, fundamentally changing who uses these systems. The ubiquity of AI means that people with a variety of attitudes, experiences, and characteristics will be interacting with these systems; some users may be less able or interested in modifying or controlling the automated systems. Unlike professional automation operators, everyday automation users will most likely not tolerate extensive, mandatory training on automation capabilities or AI functionality. Everyday users' potential lack of investment in the AI systems they use or how they operate does not mean that we should exclude them from the automation control narrative. Indeed, incorporating (and anticipating) casual users' characteristics, attitudes and capabilities can help us design more inclusive and personalized AI systems that minimize the possibility of user misuse and societal backlash.

Although most consumer products that leverage AI are seemingly low stakes (inaccurate autocorrect may be annoying, but rarely results in injury or death), automated vehicles are expected to use AI to integrate information gained from vehicle and infrastructure sensors to constantly update existing road environment maps and allow vehicles to make real-time routing decisions. Advanced driver assistance systems (ADAS) demonstrate how consumers use and approach emerging technologies and can provide lessons for the future deployment of highly automated vehicles and AI applications in general.

First, there is the challenge, mentioned above, of ADAS becoming ubiquitous in new cars. For example, all new Toyota vehicles come with their proprietary suite of ADAS functions (i.e., Toyota Safety Sense) standards, even at the most basic trim levels. Such a wide implementation of ADAS features like lane-keeping assist, forward collision warning, and blind-spot detection is expected to save lives and prevent injuries [27], but against this, we also must weigh how people with little to no interest in advanced technologies will use these features. In a survey, most drivers (83%) could not predict how adaptive cruise control would function in a particular situation [28]. A full 40% of respondents reported that features in their vehicle had acted in a way that they did not anticipate, with most respondents reporting that they did not engage in any additional information seeking behavior about their vehicle's advanced features. Such disengagement may signal casual users will exert less control over automated systems, either because they do not possess the knowledge to do so or because they simply are not interested in doing so.

Second, there is the potential for individual users to reject or discount automated systems when they fail to meet their performance expectations. It should be a goal of AI system designers to instill the appropriate amount of trust in automated features; too much trust may result in overreliance, too little would result in complete disuse. Repeated exposure to the emerging technology may be the best way to ensure proper calibration of trust and use. In an 18-month longitudinal study, drivers of a vehicle with ADAS features gradually adopted most of the advanced vehicle features while at the same time acknowledging the limitations of the features [29].

Finally, we must acknowledge and anticipate broader societal backlash to AI, especially once the technology becomes ubiquitous. Generally, AI has limited transparency, especially to the passive user [30]. This lack of transparency may lead to concerns about privacy and prioritization of technology over people which may translate into negative attitudes towards AI. Establishing unbiased organizations to evaluate the ethics of AI is one route to prevent this wider backlash [9]. Shneiderman's framework argues that automation should serve the user, not the other way around, but it is essential to consider how AI serves an individual user may not necessarily serve the wider public.

The notion of individual differences in terms of attitudes, experience, and characteristics in the context of AI has not been widely considered. Incorporating these factors into the design and deployment of AI can help us increase the personalization and flexibility of AI systems to ensure optimal adoption while still ensuring that humans are fully in control.

5 Future Directions: Dynamic Automation and Human Control

As AI becomes sophisticated and users and contexts become varied, a static framework illustrating human-AI interaction may become outdated. Modern frameworks need to focus on human-in-the-loop components, operator limitations, and individual characteristics (i.e., attitudes, experiences, characteristics).

A user's attitudes, experiences, and characteristics will influence how they deploy automation (Fig. 1). If they opt to deploy a high level of automation, they can still retain a high level of control over the task by choosing to remain in-the-loop. The decision for the user to retain control or completely hand over control to automation will be influenced by contextual factors such as whether the environment is dynamic or static, if the user has sufficient automation experience, if the automation is being deployed in a safety-critical setting, and the transparency of the system (Fig. 2). However, this decision does not have to be static. Depending on the demands of the task, and the goals of the user, the operator has the flexibility to continually adjust the level of automation that they use. By continuously monitoring and analyzing automation performance, the user can adjust what level of automation they deploy and their level of control, resulting in a feedback loop. This feedback loop will be greatly impacted by the attentional resources available to the operator, allowing them to acquire information about the automation performance and choose appropriate levels of automation. However, concurrent tasks may take attentional resources away from monitoring and adjusting the automation level and control (Fig. 3).

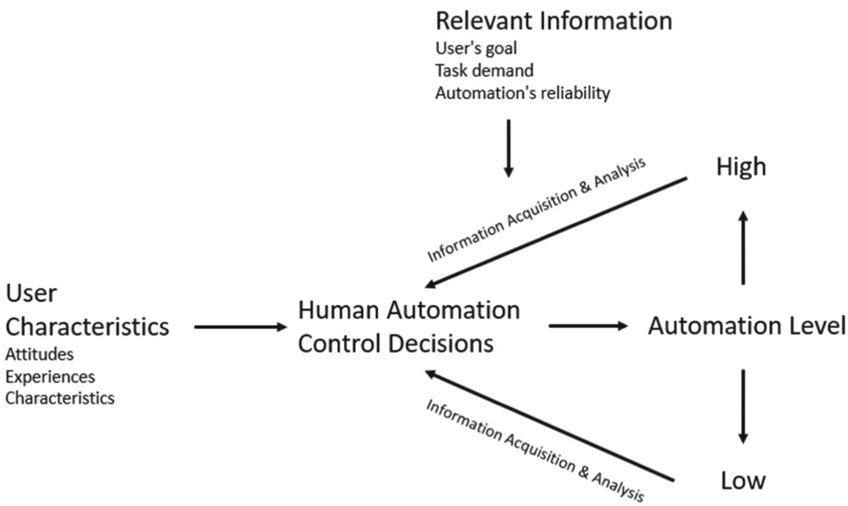

Fig. 1. Dynamic automation and human control based on user characteristics

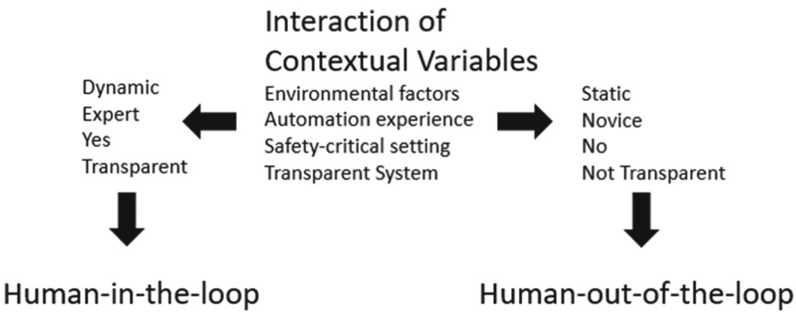

Fig. 2. Context that influences human-in-the-loop decision-making

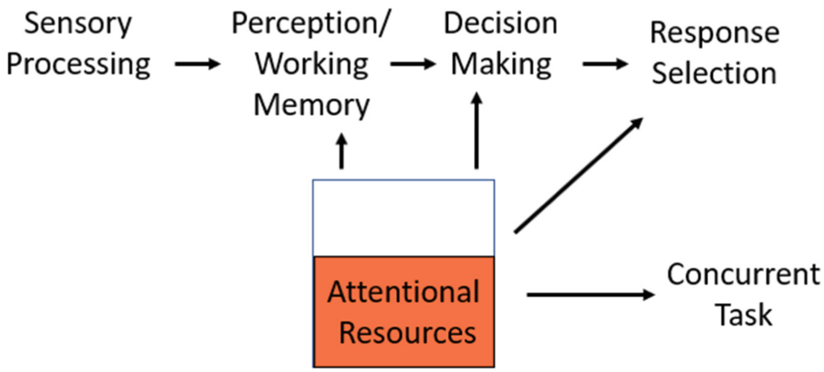

Fig. 3. Attentional resources portrayed with a simple human information processing model

Shneiderman [8] provides a scenario where automated vehicles can achieve a RST environment with high levels of human and computer control by designing a system that leads to an air traffic control equivalent for managing vehicles on the roadway. This scenario suggests that humans working from a remote station would manage traffic flow by changing speed limits in response to congestion and weather conditions. The HCAI framework encourages roadway safety engineers to transform conventional speed limits without considering human attentional limitations. Driving safely requires years of experience. During which time, operators form expectations to manage information overload. Making speed limits variable will demand continuous attentional resources. The operator of the vehicle will have to search for and verify ever-changing signage. When making critical decisions, dynamic changes in roadway conditions can negatively affect situation awareness [31]. In the worst-case scenario, changing conditions might cause an accident as drivers entering a different road will not be aware of the changes [32, 33]. They might enter a blind corner and accelerate as usual only to find traffic traveling significantly slower than expected. The road is wet, and their truck slides out of control while braking, resulting in a collision. As we approach this scenario with our dynamic approach in mind, we understand that human controlled traffic systems need to communicate with drivers like air traffic controllers do with pilots. Now drivers will need to monitor tower communications instead of talking on the phone or listening to music. Maintaining situational awareness will require additional training and, in some cases, a heavier cognitive load [17, 18].

A dynamic approach acknowledges that a two-dimensional framework between the level of automation and human control may not be sufficient in capturing the complexities of how people use AI systems. AI is open and dynamic; it is reasonable to assume that users' deployment and approach to AI will be equally as flexible. There are fundamental assumptions that needed to be made for this dynamic framework on human control and automation level. First, there is limited research on human-AI interaction, especially in regard to the level of automation. Second, understanding the impact of contextual variables, such as the environment, on users' automation control decisions needs to be further explored. Design recommendations from the trust and automation literature may be a good place to understand how users interact with automation. For example, Hoff and Bashir have identified multiple factors that influence trust such as transparency of feedback and ease of use [36]. Finally, the exact delineation of "high" and "low" automation levels may vary according to the context of AI. Future clarification may lead to the evolution of the HCAI framework.

6 Conclusions

Shneiderman's goal was to focus on designing AI that serves as tools/appliances to improve the user experience. HCAI framework argues that AI users do not need to sacrifice control even at high levels of automation; however, this framework does not incorporate situational, personal, or attentional context. As we stated previously, there are times when user preference or the environment determines when it is more beneficial to have humans in-the-loop or out-of-the-loop. Therefore, future work needs to expand on the HCAI framework by allowing the user to continually consider their context and attentional demand, allowing for more flexible use of AI.

References

1. Hengstler, M., Enkel, E., Duelli, S.: Applied artificial intelligence and trust—the case of autonomous vehicles and medical assistance devices. Technol. Forecast. Soc. Chang. **105**, 105–120 (2016)
2. Buchanan, B.G.: A (very) brief history of artificial intelligence. AI Mag. **26**(4), 53 (2005)
3. Lu, H., Li, Y., Chen, M., Kim, H., Serikawa, S.: Brain intelligence: go beyond artificial intelligence. Mob. Netw. Appl. **23**(2), 368–375 (2018)
4. Vagia, M., Transeth, A.A., Fjerdingen, S.A.: A literature review on the levels of automation during the years. What are the different taxonomies that have been proposed? Appl. Erg. **53**, 190–202 (2016)
5. Parasuraman, R., Sheridan, T., Wickens, C.D.: A model of types and levels of human interaction with automation. IEEE Trans. Syst. Man Cybern. Part A **30**, 286–297 (2000)
6. Onnasch, L., Wickens, C.D., Li, H., Manzey, D.: Human performance consequences of stages and levels of automation. Hum. Factors **56**(3), 476–488 (2014)
7. Sheridan, T.B., Verplank. W.L.: Human and computer control of undersea teleoperators. Massachussetts Institute of Technology, Cambridge, Massachusetts (1978)
8. Shneiderman, B.: Human-centered artificial intelligence: reliable, safe & trustworthy. Int. J. Hum. Comput. Interact. **36**(6), 495–504 (2020). https://doi.org/10.1080/10447318.2020.1741118
9. Shneiderman, B.: Bridging the gap between ethics and practice: guidelines for reliable, safe, and trustworthy Human-Centered AI systems. ACM Trans. Interact. Intell. Syst. **10**(4), 1–31 (2020). https://doi.org/10.1145/3419764
10. Shneiderman, B.: Human-centered AI: a second Copernican revolution. AIS Trans. Hum. Comput. Interact. **12**(3), 109–124 (2020)
11. Wickens, C.D., Hollands, J.G., Banbury, S., Parasuraman, R.: Engineering Psychology and Human Performance. Taylor & Francis, London (2015)
12. Merat, N., et al.: The "Out-of-the-Loop" concept in automated driving: proposed definition, measures and implications. Cogn. Technol. Work **21**(1), 87–98 (2018). https://doi.org/10.1007/s10111-018-0525-8

13. Talamadupula, K., Benton, J., Schermerhorn, P., Kambhampati, S., Scheutz, M.: Integrating a closed world planner with an open world robot: a case study. In: Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2010), Atlanta, GA, pp. 1561–1566 (2010)
14. Sehwag, V., et al.: Analyzing the robustness of open-world machine learning. In: 12th ACM Workshop on Artificial Intelligence and Security, London, UK, pp. 105–116 (2019)
15. Jiang, Y., Walker, N., Hart, J., Stone, P.: Open-world reasoning for service robots. In: Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling (ICAPS 2019), Berkeley, CA, pp. 725–733 (2019)
16. Anderson, B., Storlie, C., Lane, T.: Improving malware classification: bridging the static/dynamic gap, pp. 3–14. Association for Computing Machinery, October 2012
17. Bainbridge, L.: Ironies of automation. In: Analysis, Design and Evaluation of Man-Machine Systems, pp. 129–135 January1983. https://doi.org/10.1016/B978-0-08-029348-6.50026-9
18. Imbert, J.-P., Hodgetts, H.M., Parise, R., Vacho, F., Dehais, F., Tremblay, S.: Attentional costs and failures in air traffic control notifications. Ergonomics **57**(12), 1817–1832 (2014). https://doi.org/10.1016/B978-0-08-029348-6.50026-9
19. Bailey, N.R., Scerbo, M.W.: Automation-induced complacency for monitoring highly reliable systems: the role of task complexity, system experience, and operator trust. Theor. Issues Ergon. Sci. **8**(4), 321–348 (2007). https://doi.org/10.1080/14639220500535301
20. Parasuraman, R., Molloy, R., Singh, I.L.: Performance consequences of automation-induced 'complacency.' Int. J. Aviat. Psychol. **3**(1), 1–23 (1993). https://doi.org/10.1207/s15327108ijap0301_1
21. Karpinsky, N.D., Chancey, E.T., Palmer, D.B., Yamani, Y.: Automation trust and attention allocation in multitasking workspace. App. Ergon. **70**, 194–201 (2018). https://doi.org/10.1016/j.apergo.2018.03.008
22. Sato, T., Yamani, Y., Liechty, M., Chancey, E.T.: Automation trust increases under high-workload multitasking scenarios involving risk. Cogn. Technol. Work **22**(2), 399–407 (2019). https://doi.org/10.1007/s10111-019-00580-5
23. National Transportation Safety Board: Decent Below Visual Glidpath and Impact with Seawall Asiana Airlines Flight 214 Boeing 777-200ER, HL7742, San Francisco, California, 6 July 2013 (Rep. NTSB-AAR1401). Author, Washington, DC (2013)
24. Parasuraman, R., Victor, R.: Humans and automation: use, misuse, disuse, abuse. Hum. Factors **39**(2), 230–253 (1997). https://doi.org/10.1518/001872097778543886
25. Kahneman, D.: Attention and Effort. Prentice-Hall, Englewood Cliffs (1973)
26. Janssen, C.P., Donker, S.F., Brumby, D.P., Kun, A.L.: History and future of human-automation interaction. Int. J. Hum. Comput. Stud. **131**, 99–107 (2019). https://doi.org/10.1016/j.ijhcs.2019.05.006
27. Benson, A., Tefft, B., Svancara, A., Horrey, W.: Potential Reductions in Crashes, Injuries, and Deaths from Large-Scale Deployment of Advanced Driver Assistance Systems (Research Brief). AAA Foundation for Traffic Safety, Washington, D.C. (2018)
28. McDonald, A.B., McGehee, D.V., Chrysler, S.T., Askelson, N.M., Angell, L.S., Seppelt, B.D.: National survey identifying gaps in consumer knowledge of advanced vehicle safety systems. Transp. Res. Rec. **2559**(5), 1–6 (2016). https://doi.org/10.3141/2559-01
29. Lindgren, T., Fors, V., Pink, S., Osz, K.: Anticipatory experience in everyday autonomous driving. Pers. Ubiquit. Comput. **24**(6), 747–762 (2020). https://doi.org/10.1007/s00779-020-01410-6
30. Fröhlich, P., Baldauf, M., Meneweger, T., Tscheligi, M., de Ruyter, B., Paternó, F.: Everyday automation experience: a research agenda. Pers. Ubiquit. Comput. **24**(6), 725–734 (2020). https://doi.org/10.1007/s00779-020-01450-y
31. Endsley, M.R.: Towards a theory of situation awareness in dynamic systems. Hum. Factors **37**(1), 32–64 (1995)

32. Borowsky, A., Shinar, D., Parmet, Y.: Sign location, sign recognition, and driver expectancies. Transp. Res. Part F **11**, 459–465 (2008)
33. Thompson, C., Sabik, M.: Allocation of attention in familiar and unfamiliar traffic scenarios. Transp. Res. Part F **55**, 188–198 (2018)
34. van Berkel, N., Skov, M.B., Kjeldskov, J.: Human-AI interaction: intermittent, continuous, and proactive. Interactions **28**(6), 67–71 (2021). https://doi.org/10.1145/3486941
35. Zhang, R., McNeese, N.J., Freeman, G., Musick, G.: "An Ideal Human": expectations of AI teammates in human-AI teaming. In: Proceedings of the ACM on Human-Computer Interaction, vol. 4, no. 246, pp. 1–25 (2020).https://doi.org/10.1145/3432945
36. Hoff, K.A., Bashir, M.: Trust in automation: integrating empirical evidence on factors that influence trust. Hum. Factors **57**(3), 407–434 (2015). https://doi.org/10.1177/0018720814547570