Designing for Continuous Interaction with Artificial Intelligence Systems

Philipp Wintersberger
TU Wien
Vienna, Austria
philipp.wintersberger@tuwien.ac.at

Niels van Berkel
Aalborg University
Aalborg, Denmark
nielsvanberkel@cs.aau.dk

Nadia Fereydooni
Georgia Institute of Technology
Atlanta, United States
nadia.fereydooni@gatech.edu

Benjamin Tag
University of Melbourne
Melbourne, Australia
benjamin.tag@unimelb.edu.au

Elena L. Glassman
Harvard University
Cambridge, United States
glassman@seas.harvard.edu

Daniel Buschek
University of Bayreuth
Bayreuth, Germany
daniel.buschek@uni-bayreuth.de

Ann Blandford
University College London
London, United Kingdom
a.blandford@ucl.ac.uk

Florian Michahelles
TU Wien
Vienna, Austria
florian.michahelles@tuwien.ac.at

ABSTRACT

The increasing capabilities of Artificial Intelligence enable the support of users in a continuously growing number of applications. Current systems typically dictate that interaction between user input and AI output unfolds in discrete steps, as is the case with, for example, conversational agents. Novel scenarios require AI systems to adapt and respond to continuous user input, e.g., image-guided surgery and AI-supported text entry. In and across these applications, AI systems need to support more varied and dynamic interactions in which users and AI interact continuously and in parallel. Current methods and guidelines are often inadequate and sometimes even detrimental to user needs when considering continuous usage scenarios. Realizing a continuous interaction between users and AI requires a substantial change in perspective when designing Human-AI systems. In this SIG, we support the exchange of cutting-edge research contributing to a better understanding and improved methods and tools to design continuous Human-AI interaction.

CCS CONCEPTS

• **Human-centered computing** → **Human computer interaction (HCI)**; **Interaction paradigms**; **HCI theory, concepts and models**.

KEYWORDS

Human-AI interaction, design guidelines, continuous interaction, AI, ML, explainability

CHI '22 Extended Abstracts, April 29-May 5, 2022, New Orleans, LA, USA

ACM ISBN 978-1-4503-9156-6/22/04.
https://doi.org/10.1145/3491101.3516409

ACM Reference Format:
Philipp Wintersberger, Niels van Berkel, Nadia Fereydooni, Benjamin Tag, Elena L. Glassman, Daniel Buschek, Ann Blandford, and Florian Michahelles. 2022. Designing for Continuous Interaction with Artificial Intelligence Systems. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI '22 Extended Abstracts), April 29-May 5, 2022, New Orleans, LA, USA.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3491101.3516409

1 BACKGROUND

The number of touch-points between humans and computing systems steadily increases, and automation is gradually entering many workplaces and even safety critical domains such as (automated) driving or computer-aided surgery. In our private lives, we are continuously interacting with an ecosystem of systems that significantly impact our behaviour, e.g., by recommending and deciding which music we listen to, which places we visit, or which news we read. These developments must be seen under consideration of the 'always-on-mentality' (i.e., users simultaneously engaged in multiple conversations, consuming media content from various sources, etc.) that can impair productivity and well-being [3, 17, 22].

Additionally, the rise of Artificial Intelligence (AI) contributes to systems' and agents' capabilities of carrying out work more autonomously and with less defined interactions with their users. This requires to go beyond the classical interaction paradigms of the past and focus research and design efforts on what we call 'Continuous Interaction' (CI). CI focuses on a seamless synergy between users and intelligent systems and aims at a better integration of technology into our lives – for example, by putting more emphasis on the temporal dimension of interactions. In the following, we discuss arguments highlighting the need for such an approach, present core research themes, ideas, and concepts, and propose to organize a SIG meeting on CI at CHI'22 to explore and discuss different views, and to shape a community of researchers who work on related topics across different domains.

1.1 From Discrete to Continuous Interaction

The vast majority of current systems dictate that interaction between user input and machine output unfolds in discrete steps. This interaction style can, for example, be seen in the interplay between humans and conversational agents, in which the user and the AI system operate in alternating order: the user provides a command which is subsequently interpreted by the system to then output a reply to the user. Indeed, many interactive intelligent systems imply such a dialogue structure [9]. However, in contrast to these 'intermittent' interaction patterns, future systems will have to support more varied and dynamic usage scenarios in which, for example, user input and machine output can occur continuously and in parallel [25]. Examples of this continuous interaction can be found in the interaction with, and generation of, multimedia content [21], image-guided surgery [24], AI-supported text entry [5], or interaction with automated vehicles [26].

These examples highlight the potential of AI to support, accompany, and enhance activities across a wide range of contexts in everyday life, as long as we can successfully support end-user interaction without distracting the users, interrupting their focus activities, or increasing the cognitive demand on them. Realising a continuous and effortless information exchange between the user and such(often AI-driven) systems requires a substantial change in perspective.

At a first glance, this may appear trivial, but a closer look reveals that current methods and guidelines, originally meant to foster the interaction with AI systems, are often inadequate and sometimes even detrimental to user needs when considering interaction in continuous usage scenarios. For example, experiments frequently assess compliance and reliance rates to assess if users comply with AI or automated systems [13, 18, 20]. However, this method that is only useful when the interaction can be modelled as a series of single, subsequent trials (allowing users to track back AI decisions to dedicated events) but not when interaction is continuous and in parallel. Consequently, we see a strong need to develop new guidelines and evaluation methods suitable for describing and assessing CI.

On top of these problems resides the issue that 'we' as users do not use just one system, we are using a great variety of systems, and we frequently multitask concurrently but also sequentially. However, a large part of HCI research primarily focuses on isolated interaction and has only to a limited degree considered the context of continuous usage scenarios, in which interruptions to users' ongoing tasks can have severe consequences. Although various papers have addressed multitasking scenarios and the adverse effects of interruptions and task switches in different domains [3, 15, 16, 27], we see a lack of a holistic contemplation of the topics of Human-AI task parallelism, handover, interruption, and delegation, as it is present in real-world scenarios. Also, applications designed for continuous usage should adapt their interfaces to fit the situation users find themselves in, but apart from some particular use-cases (such as Android Auto or Apple CarPlay in the driving domain), interfaces do not really adapt to users' context – instead, the user is required to adapt to the technology to get things done. This results in strained interactions, in which the user and system interact iteratively rather than in a continuous manner. The need for a focus on such topics is frequently articulated. For example, Bradshaw et al. [4] have argued that *"there is need for kinds of breakthroughs in human-machine teamwork that would enable systems to not merely do things for people, but with people and other systems"*.

1.2 Definitions and Core Themes

Discussing about appropriate definitions and themes relevant for continuous interaction will become a central element of the proposed SIG-meeting. However, to provide some guidance as a starting point, **we define the concept of 'Continuous Interaction' as research addressing "the extended and seamless synergy between intelligent systems and human users over extended periods of time, with the aim to maximise benefit and minimise effort for the users"**. In particular, we suggest to focus future research endeavours on investigating and exploiting the capabilities of AI to better support:

- **Distribution of Tasks between Human and AI**: On a high level, CI asks which tasks should be automated or performed by humans. The origins of this question date back to the early times of automation research [7] but have gained additional momentum due to the recent progress in AI [6]. Instead of automating everything that is technically feasible, a human-centered application requires AI to be used mainly in situations beneficial to human users (which changes depending on context). This relates not only to (important) themes such as fairness or the *"moral crumple zone"* (i.e., users being left alone with edge cases or tasks impossible to automate [10]), it also requires AI to not remove potentially meaningful activities and interactions from users (*"amputation"*) [12, 19]. In this regard, Schmidt and Herrmann [23] have argued we should not ask what humans can do better than machines, but *"how a team of humans and machines collaborating is better than either human and technology alone?"*
- **Information Flow between Applications and Services**: Today, most work tasks require interactions with a particular tool or application. CI aims at optimizing the process of human users repeatedly interacting with various user interfaces over a longer time period to achieve certain goals, as well as the information exchange between users and involved systems/applications. It aims at providing AI-controlled attention, content, information, task, and collaboration management, meaning to automatically extract and provide the information suitable to progress towards users' overall goals, status summaries to allow for easy continuation and re-uptake of activities, integration into existing services, and proactive but non-intrusive hints for appropriate activities in different situations and across both real and virtual environments [11]. We argue the emergence of continuous human-AI scenarios, augmented and virtual reality, or conversational interfaces should yield to new interaction paradigms instead of again reinforcing existing concepts such as the Desktop Metaphor [28] or isolated applications (both analogies that put management responsibilities entirely into the user's hands).
- **Switching between Activities and Timing of Communication**: Despite the demonstrated negative effects, humans increasingly multitask. The aim of CI is neither to

eliminate nor to reduce multitasking, but to adapt system behavior in a way to intelligently deliver information and interrupt users appropriately. Various studies have shown that interruptions and their timing have significant impact on performance, stress, and wellbeing [3, 27], and multiple research projects have addressed *"attention management systems"* [2, 8]. Systems managing user attention are not limited to task interruptions but information delivery in general. A major issue here is communication timing, which has (at least from our point of view) not received sufficient attention, given that most computerized systems deliver information immediately and without an overarching strategy across multiple devices and applications. Twenty years after Horvitz's initial mixed-initiative principles [14], guidance on timing in particular seems to remain abstract (e.g., Microsoft's recent AI design guidelines suggest to *"time services based on context"* [1]) and current products often still avoid the question (e.g., word suggestions on modern smartphone keyboards are simply always on or off). Consequently, one important core question of CI is: *What is the best timing for interruptions and information delivery under consideration of the context and the whole ecosystem a user finds themself in?*

- **Continuous Task Work through Adaptive User Interfaces**: Considerations emerging from the the issues discussed above may liberate us from some issues (such as if and when human-AI cooperation should occur, if humans should better alternate or control a given process together, or when a system should communicate to users), but CI also requires to focus on designing appropriate interaction patterns. Many groups address such issues already for dedicated use cases, such as in the UbiComp, MobileHCI, the AutomotiveUI, or the CUI communities. However, what is needed are forms of Continuous Interaction: supporting the task of writing a document like this one, started on a text processor, enhanced by thoughts and ideas collected during a drive or hike, and reviewed during an encounter with a colleague. Instead of the individual apps and specific user interfaces tailored by app developers today, the goal and the function should be in the foreground. We call for researching the potential of AI in supporting this process, by identifying interaction possibilities latent in the environment, by supporting the mental models of tasks and goals in different contexts rather than particular app designs, by connecting to existing functions implemented in apps, to finally provide a stream of interactions weaved into the fabric of situations and context.

1.3 SIG Goals

This SIG invites viewpoints and discussions on designing, developing and evaluating for Continuous Interaction with AI systems in order to make interaction more seamless in the future. In our current discussions, we as authors have diverging viewpoints on the topic ourselves. Consequently, we want to present our standpoints to a larger audience and invite conference attendees to bring additional views to the table that allow us to better shape common interests and elaborate on potential definitions, involved domains and applications, and diverse approaches regarding the concept of CI. The SIG aims at establishing connections between researchers with different backgrounds to identify issues relevant for continuous interaction, and to holistically discuss societal, individual, but also technological issues. We aim for building up a community (such as collaborative project proposals or experiments) and, ultimately, generate a concrete research agenda for this timely topic. We believe elaborating on the differences between classical, discrete and continuous interaction can lead to cutting-edge research, new approaches for technology design, and a better implementation of AI driven systems in the future.

REFERENCES

[1] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N. Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, and Eric Horvitz. 2019. Guidelines for Human-AI Interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3290605.3300233

[2] Christoph Anderson, Isabel Hübener, Ann-Kathrin Seipp, Sandra Ohly, Klaus David, and Veljko Pejovic. 2018. A survey of attention management systems in ubiquitous computing environments. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 2 (2018), 1–27.

[3] Brian P. Bailey and Joseph A. Konstan. 2006. On the need for attention-aware systems: Measuring effects of interruption on task performance, error rate, and affective state. *Computers in Human Behavior* 22, 4 (2006), 685 – 708. https://doi.org/10.1016/j.chb.2005.12.009 Attention aware systems.

[4] Jeffrey M Bradshaw, Robert R Hoffman, David D Woods, and Matthew Johnson. 2013. The seven deadly myths of "autonomous systems". *IEEE Intelligent Systems* 28, 3 (2013), 54–61.

[5] Daniel Buschek, Martin Zürn, and Malin Eiband. 2021. The Impact of Multiple Parallel Phrase Suggestions on Email Input and Composition Behaviour of Native and Non-Native English Writers. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 732, 13 pages. https://doi.org/10.1145/3411764.3445372

[6] Mary Missy Cummings. 2014. Man versus machine or man+ machine? *IEEE Intelligent Systems* 29, 5 (2014), 62–69.

[7] Joost CF de Winter and Dimitra Dodou. 2014. Why the Fitts list has persisted throughout the history of function allocation. *Cognition, Technology & Work* 16, 1 (2014), 1–11.

[8] Tilman Dingler, Benjamin Tag, Sabrina Lehrer, and Albrecht Schmidt. 2018. Reading Scheduler: Proactive Recommendations to Help Users Cope with Their Daily Reading Volume. In *Proceedings of the 17th International Conference on Mobile and Ubiquitous Multimedia (MUM 2018)*. ACM, New York, NY, USA, 239–244. https://doi.org/10.1145/3282894.3282917

[9] Malin Eiband, Daniel Buschek, and Heinrich Hussmann. 2021. How to Support Users in Understanding Intelligent Systems? Structuring the Discussion. In *26th International Conference on Intelligent User Interfaces*. Association for Computing Machinery, New York, NY, USA, 120–132. https://doi.org/10.1145/3397481.3450694

[10] Madeleine Clare Elish. 2019. Moral crumple zones: Cautionary tales in human-robot interaction. *Engaging Science, Technology, and Society* 5 (2019), 40–60.

[11] Nadia Fereydooni and Bruce N. Walker. 2020. Virtual Reality as a Remote Workspace Platform: Opportunities and Challenges. (August 2020). https://www.microsoft.com/en-us/research/publication/virtual-reality-as-a-remote-workspace-platform-opportunities-and-challenges/

[12] Anna-Katharina Frison, Philipp Wintersberger, Andreas Riener, and Clemens Schartmüller. 2017. Driving Hotzenplotz: A Hybrid Interface for Vehicle Control Aiming to Maximize Pleasure in Highway Driving. In *Proceedings of the 9th International Conference on Automotive User Interfaces and Interactive Vehicular Applications* (Oldenburg, Germany) *(AutomotiveUI '17)*. ACM, New York, NY, USA, 236–244.

[13] Brittany E. Holthausen, Philipp Wintersberger, Bruce N. Walker, and Andreas Riener. 2020. Situational Trust Scale for Automated Driving (STS-AD): Development and Initial Validation. In *12th International Conference on Automotive User Interfaces and Interactive Vehicular Applications* (Virtual Event, DC, USA) *(AutomotiveUI '20)*. Association for Computing Machinery, New York, NY, USA, 40–47. https://doi.org/10.1145/3409120.3410637

[14] Eric Horvitz. 1999. Principles of Mixed-Initiative User Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Pittsburgh, Pennsylvania, USA) *(CHI '99)*. Association for Computing Machinery, New York, NY, USA, 159–166. https://doi.org/10.1145/302979.303030

[15] Christian P Janssen, Sandy JJ Gould, Simon YW Li, Duncan P Brumby, and Anna L Cox. 2015. Integrating knowledge of multitasking and interruptions across different perspectives and research methods.

[16] Christian P. Janssen, Andrew L. Kun, Stephen Brewster, Linda Ng Boyle, Duncan P. Brumby, and Lewis L. Chuang. 2019. Exploring the Concept of the (Future) Mobile Office. In *Proceedings of the 11th International Conference on Automotive User Interfaces and Interactive Vehicular Applications: Adjunct Proceedings* (Utrecht, Netherlands) *(AutomotiveUI '19)*. ACM, New York, NY, USA, 465–467. https://doi.org/10.1145/3349263.3349600

[17] Anita C Keller, Laurenz L Meier, Achim Elfering, and Norbert K Semmer. 2020. Please wait until I am done! Longitudinal effects of work interruptions on employee well-being. *Work & Stress* 34, 2 (2020), 148–167.

[18] David Keller and Stephen Rice. 2009. System-wide versus component-specific trust using multiple aids. *The Journal of General Psychology: Experimental, Psychological, and Comparative Psychology* 137, 1 (2009), 114–128.

[19] Holger Klapperich and Marc Hassenzahl. 2016. Hotzenplotz: Reconciling Automation with Experience. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction* (Gothenburg, Sweden) *(NordiCHI '16)*. Association for Computing Machinery, New York, NY, USA, Article 39, 10 pages. https://doi.org/10.1145/2971485.2971532

[20] Zhuoran Lu and Ming Yin. 2021. Human Reliance on Machine Learning Models When Performance Feedback is Limited: Heuristics and Risks. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 78, 16 pages. https://doi.org/10.1145/3411764.3445562

[21] Jon McCormack, Toby Gifford, Patrick Hutchings, Maria Teresa Llano Rodriguez, Matthew Yee-King, and Mark d'Inverno. 2019. In a Silent Way: Communication Between AI and Improvising Musicians Beyond Sound. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–11. https://doi.org/10.1145/3290605.3300268

[22] Brandon T McDaniel and Michelle Drouin. 2019. Daily technology interruptions and emotional and relational well-being. *Computers in Human Behavior* 99 (2019), 1–8.

[23] Albrecht Schmidt and Thomas Herrmann. 2017. Intervention user interfaces: a new interaction paradigm for automated systems. *interactions* 24, 5 (2017), 40–45.

[24] Niels van Berkel, Omer F. Ahmad, Danail Stoyanov, Laurence Lovat, and Ann Blandford. 2021. Designing Visual Markers for Continuous Artificial Intelligence Support: A Colonoscopy Case Study. *ACM Trans. Comput. Healthcare* 2, 1, Article 7 (dec 2021), 24 pages. https://doi.org/10.1145/3422156

[25] Niels van Berkel, Mikael B. Skov, and Jesper Kjeldskov. 2021. Human-AI Interaction: Intermittent, Continuous, and Proactive. *Interactions* 28, 6 (nov 2021), 67–71. https://doi.org/10.1145/3486941

[26] Philipp Wintersberger, Hannah Nicklas, Thomas Martlbauer, Stephan Hammer, and Andreas Riener. 2020. Explainable Automation: Personalized and Adaptive UIs to Foster Trust and Understanding of Driving Automation Systems. In *12th International Conference on Automotive User Interfaces and Interactive Vehicular Applications* (Virtual Event, DC, USA) *(AutomotiveUI '20)*. Association for Computing Machinery, New York, NY, USA, 252–261. https://doi.org/10.1145/3409120.3410659

[27] Philipp Wintersberger, Andreas Riener, Clemens Schartmüller, Anna-Katharina Frison, and Klemens Weigl. 2018. Let Me Finish before I Take Over: Towards Attention Aware Device Integration in Highly Automated Vehicles. In *Proceedings of the 10th International Conference on Automotive User Interfaces and Interactive Vehicular Applications* (Toronto, ON, Canada) *(AutomotiveUI '18)*. Association for Computing Machinery, New York, NY, USA, 53–65. https://doi.org/10.1145/3239060.3239085

[28] Colton Zuvich and Jordan Silbert. 2019. Workspaces: A Speculative Re-Imagining of Organizing Windows for Desktop Interfaces. (2019).