

Augmenting Human Cognition With Generative AI: Lessons From AI-Assisted Decision-Making

Zelun Tony Zhang

zhang@fortiss.org

fortiss GmbH, Research Institute of the Free State of Bavaria

Munich, Germany

LMU Munich

Munich, Germany

Leon Reicherts

leon.reicherts@microsoft.com

Microsoft Research

Cambridge, United Kingdom

Abstract

How can we use generative AI to design tools that augment rather than replace human cognition? In this position paper, we review our own research on AI-assisted decision-making for lessons to learn. We observe that in both AI-assisted decision-making and generative AI, a popular approach is to suggest AI-generated end-to-end solutions to users, which users can then accept, reject, or edit. Alternatively, AI tools could offer more incremental support to help users solve tasks themselves, which we call process-oriented support. We describe findings on the challenges of end-to-end solutions, and how process-oriented support can address them. We also discuss the applicability of these findings to generative AI based on a recent study in which we compared both approaches to assist users in a complex decision-making task with LLMs.

CCS Concepts

• **Human-centered computing** → **Interaction paradigms; HCI theory, concepts and models**; • **Information systems** → *Decision support systems*.

Keywords

Generative AI, tools for thought, AI-assisted decision-making, process-oriented support, forward reasoning

ACM Reference Format:

Zelun Tony Zhang and Leon Reicherts. 2025. Augmenting Human Cognition With Generative AI: Lessons From AI-Assisted Decision-Making. In *CHI '25 T4T: Workshop on Tools for Thought—Research and Design for Understanding, Protecting, and Augmenting Human Cognition with Generative AI, April 26, 2025, Yokohama, Japan*. ACM, New York, NY, USA, 5 pages.

1 Introduction

In recent years, the perspective that AI should not replace, but augment and enhance human cognition has become increasingly popular [1, 3, 13]. Now, with the rapid progress in generative AI (GenAI), this perspective becomes even more relevant, as AI can generate satisfactory, or sometimes even impressive solutions to more and more tasks with little human effort. But how can we leverage these capabilities without sidelining human reasoning? Or more ambitiously, how can we combine artificial with human intelligence to produce better results than either could on its own? These questions have been extensively researched in the field of *AI-assisted*

decision-making as one of the primary applications of human-AI collaboration prior to the popularization of GenAI. Therefore, to understand how to design GenAI systems to protect and augment human thinking, it seems natural to turn to AI-assisted decision-making for lessons to learn. In this position paper, we review our own work in AI-assisted decision-making for how it can inform the design of GenAI tools for thought. We start by outlining several challenges of using AI to propose end-to-end solutions (Section 2), continue by presenting alternative approaches and how they can overcome these challenges (Section 3), and end by discussing the applicability of these findings to GenAI, based on our recent study on supporting complex decisions with LLMs (Section 4).

2 Challenges of End-to-End Solutions

There is a tendency in AI-assisted decision-making to design AI in a recommendation-centric way: The AI provides an end-to-end recommendation, i.e., a recommendation for the final decision straight from the input data, without involving the user. The user can then decide whether to accept or overrule this recommendation. A similar tendency can be observed in GenAI as well, where the AI often assists by generating an end-to-end solution to the problem, and the user is only involved afterwards as an editor. In our research, we identified several challenges that come with this approach.

Increasing overreliance over time. In [22], we measured how participants' overreliance developed over repeated interactions with a recommendation-centric AI, and found that overreliance increased over time. This was to be expected for long, tiring study sessions that are routinely used in AI-assisted decision-making research. However, overreliance even increased when participants only had short interactions with the AI, which were spread over a longer period. The results indicate that providing end-to-end recommendations makes it difficult for users to remain meaningfully engaged with the decision-making task, similar to the difficulty of supervisory control in automation [2].

Overreliance in difficult tasks. We conducted two studies using different methods and decision tasks to study how users' reliance on recommendation-centric AI depends on decision difficulty [17, 19], as the vision is often for AI to augment human reasoning in *difficult* tasks. The results of both studies were consistent in that participants could often rely appropriately on AI recommendations in easy decisions, but became much more over-reliant in difficult decisions. The problem is that for decisions where users are highly uncertain about their answer, recommendation-centric support does little to improve their understanding of the

decision. Consequently, users cannot meaningfully contribute and have little choice but to rely on the AI recommendation.

Helpful when correct, but a burden when wrong. A common assertion in the discourse about using AI in high-stakes domains is that AI models are black boxes that need to be opened. However, when we interviewed professional pilots about their thoughts on deploying AI into cockpits [20], transparency was not their main concern. Instead, they were much more concerned that AI might not be intelligent enough to respond appropriately to complex situations and become a burden when workload is already high—a phenomenon known as *clumsy automation* in automation research [5]. This finding points to a seemingly trivial, but often neglected requirement for the design of AI systems: that imperfect AI outputs should have minimal negative consequences for users, which Gu et al. [7] call *AI-resilient*. In a follow-up study [21], we found why recommendation-centric support tends to violate this requirement. Pilots stated that they would not blindly trust AI, but they also rejected the idea that they would have to review every AI recommendation for its correctness. This seeming contradiction highlights that instead of augmenting users' reasoning, end-to-end recommendations burden users with an effortful review task that is not integrated into their existing reasoning processes—and one that becomes more effortful the more complex the decision is.

3 Benefits of Process-Oriented Support

Providing end-to-end solutions is often the most obvious approach, but it is not the only one. Rather than solving tasks *for* users, we can also aim to help users solve the task *themselves* [6] by identifying and focusing on their primary challenges when solving a task. Rather than end-to-end solutions, the AI would provide more incremental support targeted at these primary challenges. The goal would be to help users reason *forward* through their reasoning processes, rather than pushing them to reason *backward* from an end-to-end solution. We call this alternative paradigm *process-oriented support* (Figure 1).

We empirically compared recommendation-centric to process-oriented support in the use case of diversions in commercial aviation [18]. When making diversion decisions, pilots' primary challenge is to gather and integrate the required information from multiple sources. Our process-oriented support concept addressed this challenge by pulling relevant information together in a continuously updated table view. AI continuously evaluated the information and highlighted possible limitations at the surrounding airports within the table. We found that this continuous provision of local hints was better accepted by pilots, led to less overreliance, and resulted in similar decision times as recommendation-centric support. Interestingly, only the combination of recommendations and continuous provision of local hints led to faster decisions, as in that case, recommendations were better integrated into pilots' reasoning, which removed the need for the time-consuming review of the recommendation. This also shows that process-oriented support does not preclude using AI to present recommendations or complete solutions. However, process-oriented support emphasizes that *if* recommendations are given, they have to be embedded into the process such that users are still encouraged to reason forward—either through recommendations that are easy to verify [17, 19],

or through additional incremental support leading up to the final solution [18].

Process-oriented support can take very different forms depending on the task. In another study, we investigated stock investment decision-making [11]. One of the key challenges of this task is to stay consistent with one's strategy and not to act too emotionally in reaction to market events. To address this challenge, we designed a chatbot with the aim to scaffold the decision-making process of experienced stock investors by asking reflective questions. We found that these probing questions from the chatbot got investors to consider aspects of their decision which they previously had not adequately taken into account, giving them the opportunity to reflect on certain aspects of an intended investment before proceeding with it. Other examples of process-oriented decision support are often found in healthcare, e.g. Lindvall et al. [9] (key challenge addressed with AI: finding small structures in huge images) or Zhang et al. [16] (key challenge: uncertainty about patients' future state).

Overall, these studies demonstrate that by targeting AI support at the core challenges in users' existing working processes, we can design AI tools that engage and integrate much more with users' own reasoning than end-to-end solutions. As a result, process-oriented support—if properly designed—can address all of the challenges of end-to-end solutions laid out in Section 2. For one, as users' contributions remain key to producing a result, they are less likely to disengage from the task and become increasingly over-reliant on AI over time. Process-oriented AI tools further aim to enhance users' own understanding of the problem, which in turn should also help with making sense of and contextualizing AI support even in difficult decisions, leading to more appropriate reliance, as we could observe in our diversion use case [18]. Lastly, since process-oriented support is integrated into users' reasoning, reviewing AI outputs becomes less effortful, as exemplified by the faster decisions when recommendations were combined with continuous provision of local hints [18]. This reduces the burden when AI makes mistakes, provided subpar AI outputs can be easily dismissed or corrected.

4 Applicability to Generative AI

As the findings above are not unique to specific AI technologies, they likely also apply to many of the tasks GenAI tools are designed and used for. In fact, the challenges described in Section 2 may become even more pronounced as GenAI becomes increasingly capable of solving more and more complex tasks end-to-end. At the same time, GenAI offers novel possibilities to design AI tools in a process-oriented manner. For instance, while the investment chatbot mentioned in Section 3 was effective at triggering reflections in investors, its rule-based question prompts were sometimes not sufficiently relevant for what they were thinking or intending to do at the given moment [11]. Through their ability to process unstructured human reasoning, LLMs may be able to address this challenge and enable the design of AI tools that are more tightly integrated with users' thought processes.

We explored how lessons from AI-assisted decision-making with "traditional" AI apply to GenAI by comparing an end-to-end approach with a more process-oriented approach for a GenAI tool for a complex decision-making task [12]. Participants had to invest in a selection of ETFs (exchange-traded funds) to match a given investor

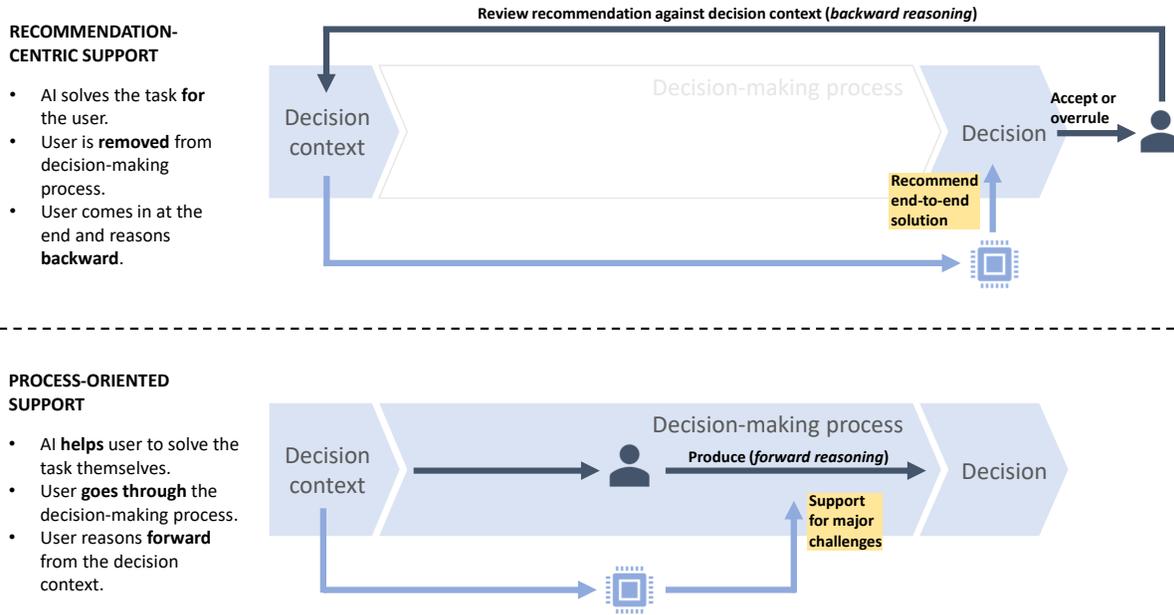


Figure 1: Recommendation-centric vs. process-oriented support. Adapted from [18].

profile. They interacted with two types of GenAI support: *RecommendAI* directly recommended a set of ETFs to invest in, along with a short explanation for each recommendation. *ExtendAI* first asked participants to describe their own investment plan and their rationale for the intended trades. The AI would then extend the rationale by embedding feedback into it, with the aim of highlighting additional factors to consider and uncovering blind spots.

Our findings show that *ExtendAI* is a promising way to conceptualize human-GenAI interactions in complex (decision-making) tasks. With *ExtendAI*, participants achieved slightly better diversified portfolios with fewer trades than with *RecommendAI*. This indicates that *ExtendAI* helped participants better understand how to improve their plan in a targeted way (forward reasoning), while participants tended to review *RecommendAI*'s recommendations individually with a less holistic understanding of their impact on their portfolio (backward reasoning). *ExtendAI* was further perceived as better integrated with participants' reasoning. As a result, while with *RecommendAI*, participants faced the question of whether to trust the recommendations or not, this was less of a concern with *ExtendAI*. Participants could more easily make sense of its feedback as it was directly tied to their own thoughts. Lastly, with *ExtendAI*, participants appeared to perceive more ownership of their decisions, while with *RecommendAI*, they tended to blame the AI for unsatisfactory outcomes.

However, we also found *RecommendAI* to have its advantages. In particular, it was perceived as more actionable and insightful, and to give more fresh ideas as it was not directly tied to participants' own thoughts. Overall, preference between the two types of AI support was evenly split among participants, highlighting trade-offs along various dimensions when moving along the spectrum between end-to-end and process-oriented tools for thought. Some of the dimensions which we found in our research are:

- **Timing of the support:** AI tools can provide assistance earlier or later in the user's reasoning process. Support in the early stages (e.g., via recommendations) can help the user get started and explore fresh perspectives, but may create a strong anchor that inhibits independent human reasoning. Support at later stages (such as with *ExtendAI*) can be embedded into the user's own reasoning, but when it comes "too late", it can be difficult for the user to accommodate suggestions from the AI that do not align with their own thinking/plans, as they have already "settled" on what they are going to do [12]. A good sweet spot seems to be support while the user's thoughts are "in the making," for example when the user has a rough idea of what to do but is still in the process of identifying the pros and cons for doing so.
- **Type of support:** AI tools can offer various types of support, from *providing*, to *curating*, to *interpreting* data. The type of support needs to be chosen based on the application context (e.g., type of task, user expertise, AI performance). While AI developers are often biased toward data interpretation, this form of support requires especially careful consideration, as it can easily diminish user autonomy when not properly embedded into the user's thinking process [12, 15, 16, 18]. In our diversion use case [18], we found that a combination of support types—continuous provision of local hints (curation) and recommendations (interpretation)—can be effective.
- **Degree of externalization required from the user:** To integrate AI support with the user's reasoning, some degree of cognitive externalization is often required from the user, such as with *ExtendAI*, where users had to describe their rationale first before the AI can then extend it. In that case, the externalization itself was also helpful to some participants to better think through their decisions (similar to

so-called “rubber-duck debugging” and the self-explanation effect [4]). However, externalizing one’s thoughts can also be challenging and effortful. It is thus important to consider if the amount of effort required to externalize one’s thinking, and the form in which this is done, is adequate with respect to what users might gain from it. For example, our participants suggested providing a direct manipulation interface to specify intended trades and relying only on natural language input to describe the rationale for each trade [12].

- **Degree of friction introduced into the process:** One possible goal of AI tools is to help the user correct their own mistakes or notice their own oversights. To that end, it may be necessary to introduce a certain degree of friction into the user’s reasoning process, as with *ExtendAI* [12] or the probing questions from the chatbot in our stock investment study [11]. However, it is important to find the right balance that helps the user to reflect on their own thinking without becoming too annoying or disruptive to the user.

All in all, the results of our ETF investment study are in line with findings from our previous work—that process-oriented support tools can integrate more deeply with human reasoning and even lead to better task outcomes than the end-to-end approach—while also surfacing nuanced considerations in the design of such tools. The study confirms the applicability of these findings to GenAI, at least within the context of decision-making tasks. But even moving beyond to other knowledge work tasks for which GenAI is commonly used, we are confident that the strategy of process-orientation can serve as a valuable starting point in cases where protecting and augmenting human cognition is desired. For instance, results similar to ours have been found in tasks such as ideation [10], coding [14], or writing [8].

5 Conclusion

In summary, based on our previous work on AI-assisted decision-making, we suggest that to preserve and enhance human reasoning in GenAI applications, the common approach of solving complex reasoning tasks end-to-end may often be inappropriate. In many applications, it may be easier for users to engage with the task and contribute their own reasoning with incremental, process-oriented support that facilitates forward reasoning, which might even lead to better outcomes than backward reasoning from end-to-end solutions. In a first exploratory study, we investigated both approaches applied to LLMs in the context of complex investment decision-making, confirming the applicability of these findings to GenAI, while also highlighting nuances on how to integrate process-oriented GenAI tools with human reasoning. Going forward, we suggest that process-oriented support can be an effective framework to inform the design of both “traditional” AI-assisted decision-making tools but also GenAI-based tools for thought.

References

- [1] Zeynep Akata, Dan Balliet, Maarten de Rijke, Frank Dignum, Virginia Dignum, Gusztai Eiben, Antske Fokkens, Davide Grossi, Koen Hindriks, Holger Hoos, Hayley Hung, Catholijn Jonker, Christof Monz, Mark Neerincx, Frans Oliehoek, Henry Prakken, Stefan Schlobach, Linda van der Gaag, Frank van Harmelen, Herke van Hoof, Birna van Riemsdijk, Aimee van Wynsberghe, Rineke Verbrugge, Bart Verheij, Piek Vossen, and Max Welling. 2020. A research agenda for hybrid intelligence: augmenting human intellect with collaborative, adaptive, responsible, and explainable artificial intelligence. *Computer* 53, 8 (Aug. 2020), 18–28. doi:10.1109/MC.2020.2996587
- [2] Lisanne Bainbridge. 1983. Ironies of automation. *Automatica* 19, 6 (Nov. 1983), 775–779. doi:10.1016/0005-1098(83)90046-8
- [3] Shan Carter and Michael Nielsen. 2017. Using artificial intelligence to augment human intelligence. *Distill* 2, 12 (Dec. 2017), e9. doi:10.23915/distill.00009
- [4] Michelene T.H. Chi, Nicholas De Leeuw, Mei-Hung Chiu, and Christian Lavancher. 1994. Eliciting self-explanations improves understanding. *Cognitive Science* 18, 3 (1994), 439–477. doi:10.1207/s15516709cog1803_3
- [5] Richard I. Cook, David D. Woods, Elizabeth Mccolligan, and Michael B. Howie. 1991. Cognitive consequences of clumsy automation on high workload, high consequence human performance. In *Fourth Annual Workshop on Space Operations Applications and Research (SOAR 90)*. NASA, Albuquerque, NM, USA, 543–546. https://ntrs.nasa.gov/citations/19910011398
- [6] Alan Dix. 2007. Designing for appropriation. In *Proceedings of the 21st British HCI Group Annual Conference on People and Computers (BCS-HCI '07, Vol. 2)*. BCS Learning & Development Ltd., Lancaster, UK, 27–30. doi:10.14236/ewic/HCI2007.53
- [7] Ziwei Gu, Ian Arawjo, Kenneth Li, Jonathan K. Kummerfeld, and Elena L. Glassman. 2024. An AI-resilient text rendering technique for reading and skimming documents. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. ACM, Honolulu, HI, USA, 898:1–898:22. doi:10.1145/3613904.3642699
- [8] Zhuoyan Li, Chen Liang, Jing Peng, and Ming Yin. 2024. The value, benefits, and concerns of generative AI-powered assistance in writing. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, Honolulu, HI, USA, 1048:1–1048:25. doi:10.1145/3613904.3642625
- [9] Martin Lindvall, Claes Lundström, and Jonas Löwgren. 2021. Rapid assisted visual search: supporting digital pathologists with imperfect AI. In *Proceedings of the 26th International Conference on Intelligent User Interfaces (IUI '21)*. ACM, College Station, TX, USA, 504–513. doi:10.1145/3397481.3450681
- [10] Peinuan Qin, Chi-Lan Yang, Jingshu Li, Jing Wen, and Yi-Chieh Lee. 2025. Timing matters: How using LLMs at different timings influences writers’ perceptions and ideation outcomes in AI-assisted ideation. doi:10.48550/arXiv.2502.06197 arXiv:2502.06197 [cs].
- [11] Leon Reicherts, Gun Woo Park, and Yvonne Rogers. 2022. Extending chatbots to probe users: enhancing complex decision-making through probing conversations. In *Proceedings of the 4th Conference on Conversational User Interfaces (CUI '22)*. ACM, Glasgow, UK, 2:1–2:10. doi:10.1145/3543829.3543832
- [12] Leon Reicherts, Zelun Tony Zhang, Elisabeth von Oswald, Yuanting Liu, Yvonne Rogers, and Mariam Hassib. 2025. AI, help me think—but for myself: Assisting people in complex decision-making by providing different kinds of cognitive support. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. ACM, Yokohama, Japan, 1–19. doi:10.1145/3706598.3713295
- [13] Ben Shneiderman. 2020. Human-centered artificial intelligence: three fresh ideas. *AIS Transactions on Human-Computer Interaction* 12, 3 (2020), 109–124. doi:10.17705/1thci.00131
- [14] Priyan Vaithilingam, Tianyi Zhang, and Elena L. Glassman. 2022. Expectation vs. experience: Evaluating the usability of code generation tools powered by large language models. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems (CHI EA '22)*. Association for Computing Machinery, New York, NY, USA, 1–7. doi:10.1145/3491101.3519665
- [15] Nur Yildirim, Hannah Richardson, Maria Teodora Wetscherek, Junaid Bajwa, Joseph Jacob, Mark Ames Pinnock, Stephen Harris, Daniel Coelho De Castro, Shruthi Bannur, Stephanie Hyland, Pratik Ghosh, Mercy Ranjit, Kenza Bouzid, Anton Schwaighofer, Fernando Pérez-García, Harshita Sharma, Ozan Oktay, Matthew Lungren, Javier Alvarez-Valle, Aditya Nori, and Anja Thieme. 2024. Multimodal healthcare AI: Identifying and designing clinically relevant vision-language applications for radiology. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. ACM, Honolulu, HI, USA, 444:1–444:22. doi:10.1145/3613904.3642013
- [16] Shao Zhang, Jianing Yu, Xuhai Xu, Changchang Yin, Yuxuan Lu, Bingsheng Yao, Melanie Tory, Lacey M. Padilla, Jeffrey Caterino, Ping Zhang, and Dakuo Wang. 2024. Rethinking human-AI collaboration in complex medical decision making: a case study in sepsis diagnosis. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. ACM, Honolulu, HI, USA, 445:1–445:18. doi:10.1145/3613904.3642343
- [17] Zelun Tony Zhang, Felicitas Buchner, Yuanting Liu, and Andreas Butz. 2024. You can only verify when you know the answer: Feature-based explanations reduce overreliance on AI for easy decisions, but not for hard ones. In *Proceedings of Mensch und Computer 2024 (MuC '24)*. ACM, Karlsruhe, Germany, 156–170. doi:10.1145/3670653.3670660
- [18] Zelun Tony Zhang, Sebastian S. Feger, Lucas Dullenkopf, Rulu Liao, Lou Süßlin, Yuanting Liu, and Andreas Butz. 2024. Beyond recommendations: From backward to forward AI support of pilots’ decision-making process. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW2 (Nov. 2024), 485:1–485:32. doi:10.1145/

3687024

- [19] Zelun Tony Zhang, Seniha Ketenci Argin, Mustafa Baha Bilen, Doğan Urgan, Sencer Melih Deniz, Yuanting Liu, and Mariam Hassib. 2024. Measuring the effect of mental workload and explanations on appropriate AI reliance using EEG. *Behaviour & Information Technology* (Nov. 2024), 1–19. doi:10.1080/0144929X.2024.2431055 Advance online publication.
- [20] Zelun Tony Zhang, Yuanting Liu, and Heinrich Husmann. 2021. Pilot attitudes toward AI in the cockpit: implications for design. In *2021 IEEE 2nd International Conference on Human-Machine Systems (ICHMS)*. IEEE, Magdeburg, Germany, 33:1–33:6. doi:10.1109/ICHMS53169.2021.9582448
- [21] Zelun Tony Zhang, Cara Storath, Yuanting Liu, and Andreas Butz. 2023. Resilience through appropriation: Pilots' view on complex decision support. In *Proceedings of the 28th International Conference on Intelligent User Interfaces (IUI '23)*. ACM, Sydney, NSW, Australia, 397–409. doi:10.1145/3581641.3584056
- [22] Zelun Tony Zhang, Sven Tong, Yuanting Liu, and Andreas Butz. 2023. Is over-reliance on AI provoked by study design?. In *Human-Computer Interaction – INTERACT 2023*, José Abdelnour Nocera, Marta Kristin Lárusdóttir, Helen Petrie, Antonio Piccinno, and Marco Winckler (Eds.). Springer, York, UK, 49–58. doi:10.1007/978-3-031-42286-7_3