Journal of Artificial Intelligence Research 80 (2024) 919-929 Submitted 01/2024; published 07/2024

Viewpoint: Hybrid Intelligence Supports Application Development for Diabetes Lifestyle Management

Bernd J. W. Dudzik B.J.W.DUDZIK@TUDELFT.NL
Delft Univeristy of Technology,
2628 XE Delft, The Netherlands

Jasper S. van der Waa JASPER.VANDERWAA@TNO.NL
TNO
3769 DE Soesterberg, The Netherlands

Pei-Yu Chen P.Y.CHEN@TUDELFT.NL
Roel Dobbe R.I.J.DOBBE@TUDELFT.NL
Íñigo M.D.R. de Troya I.M.D.R.DETROYA@TUDELFT.NL
Delft Univeristy of Technology,
2628 XE Delft, The Netherlands

Roos M. Bakker ROOS.BAKKER@TNO.NL
TNO
3769 DE Soesterberg, The Netherlands
&
Leiden University
2311 BE Leiden, The Netherlands

Maaike H. T. de Boer MAAIKE.DEBOER@TNO.NL
Quirine T. S. Smit QUIRINE.SMIT@TNO.NL
TNO
3769 DE Soesterberg, The Netherlands

Davide Dell'Anna D.DELLANNA@UU.NL
Emre Erdogan E.ERDOGAN1@UU.NL
Pınar Yolum P.YOLUM@UU.NL
Shihan Wang S.WANG2@UU.NL
Utrecht University
3584 CS Utrecht, The Netherlands

Selene Báez Santamaría S.BAEZSANTAMARIA@VU.NL
Lea Krause L.KRAUSE@VU.NL
Vrije Universiteit Amsterdam,
1081 HV Amsterdam, The Netherlands

Bart A. Kamphorst BART.KAMPHORST@WUR.NL
Wageningen University & Research,
6706 KN Wageningen, The Netherlands

Abstract

Type II diabetes is a complex health condition requiring patients to closely and continuously collaborate with healthcare professionals and other caretakers on lifestyle changes. While intelligent products have tremendous potential to support such *Diabetes Lifestyle Management (DLM)*, existing products are typically conceived from a technology-centered perspective that insufficiently acknowledges the degree to which collaboration and inclusion of stakeholders is required. In this article, we argue that the emergent design philosophy of *Hybrid Intelligence (HI)* forms a suitable alternative lens for research and development. In particular, we (1) highlight a series of pragmatic challenges for effective AI-based DLM support based on results from an expert focus group, and (2) argue for HI's potential to address these by outlining relevant research trajectories.

1. Introduction

One of the fastest growing lifestyle-related diseases is Type II Diabetes, with 536.6 million people diagnosed in 2022 (10.5% of the world population) and 783.2 million people predicted to be diagnosed by 2045 (12.2%) (Sun et al., 2022). Diabetes puts an enormous strain on healthcare, as it requires regular monitoring and can cause complications requiring prolonged care (Okemah et al., 2018). Since type II Diabetes can be prevented or mitigated through lifestyle changes (Magkos et al., 2020), much research focuses on AI-based support products to assist patients in lifestyle management (Contreras and Vehi, 2018). Thus far, however, widespread adoption of such products has been hindered by a variety of issues (cf. Lie et al. (2017)) that can be traced to a prevalent technology-centered design philosophy underlying development (Ziefle et al., 2010; D'Haeseleer et al., 2021).

Hybrid Intelligence (HI) is an emerging design philosophy aimed at augmenting human capabilities with AI-based capabilities and vice versa (Wiethof and Bittner, 2021; Akata et al., 2020). The central idea behind HI is that agents – artificial or human – use their complementary capabilities to strengthen and support the socio-technical system as a whole, resulting in improved performance overall (Dellermann et al., 2019). HI adopts a holistic systems perspective, considering constituting agents' different capacities in research and design. To achieve HI, Akata et al. (2020) prescribe that the socio-technical system should support *Collaboration*, *Adaptivity*, *Explainability* and *Responsibility* in behavior among agents.

This article argues that embracing HI as a perspective is especially beneficial for developing effective support products in *Diabetes Lifestyle Management (DLM)*, a domain where collaboration and inclusion of many different actors is required for success. Concretely, we (1) discuss a series of *Pragmatic Challenges* for intelligent support of DLM based on an expert focus group, and then (2) propose a set of *Research Trajectories* from within a HI perspective to address these challenges.

2. Pragmatic Challenges for Effective AI-based DLM Support

In the following, we outline a series of pragmatic challenges (identified as *C1-C4* below) for the successful design, development, or adoption of AI-based support products for DLM. These are derived from the results of three participatory design sessions (Spinuzzi, 2005) of four hours each, respectively focusing on (1) problem identification, (2) an envisioned HI system design, and (3) the further iteration of main support functions. Sessions were organized in accordance with the institutional policy of the *Nederlandse organisatie voor Toegepast-Natuurwetenschappelijk Onderzoek*[1] *(TNO)* that ethical approval was not required as no personal information would be collected, and there was no foreseeable risk of harm to the participants. Participants included two general practitioners, two lifestyle coaches, three diabetes researchers, four IT healthcare company representatives, two interaction designers, and four human-AI interaction researchers guiding the sessions. Sessions focused on the initial design phase of ideation (Goel, 1995), with the first one revolving around stakeholder identification (Vos and Achterkamp, 2006), and the others iterating on support function ideas through storytelling (Hunsucker and Siegel, 2015) and story-boarding activities (Mitchell and Nørgaard, 2011). Participants were kept engaged and critical through Socratic questioning. A major limitation was the absence of patient representatives. A report was written to summarize the identified scenarios and support functions. Finally, the authors reflected on these outcomes in a separate workshop, formulating the challenges below.

C1: Continuous Involvement of Diverse Stakeholders The insights obtained from participants reinforce existing findings about the meaningfulness of the continuous involvement of stakeholders throughout a product's life cycle – including the design, development, maintenance, and revision phases. Without such involvement, support products are unlikely to operate safely and effectively. Involvement also facilitates integration in organisations (Makarius et al., 2020) and contributes to successful deployment (Rafner et al., 2022). Participants mentioned, however, that different stakeholder groups have preconceptions of each other that may hinder communication and understanding. For example, technologists (e.g., AI experts) may falsely believe that domain experts (e.g., doctors) are overly technology-averse. Such gaps in understanding a problem, its context, and the role of technology are well known in cognitive science and system engineering (Rasmussen, 1987; Leveson, 2012). Importantly, the literature points out that keeping stakeholders involved requires participation to be approached critically (Bødker and Kyng, 2018; Sloane et al., 2022). A mere focus group to derive specifications at the start of a project does not constitute participation but merely consultation. Finally, all participants agreed that, unfortunately, patients and their families are often poorly represented. This is a serious problem, as these groups are both critically relevant for and at risk in DLM support.

C2: Long-term Patient Engagement One challenge participants highlighted was the difficulty of keeping patients engaged with a support product. This observation is in line with findings suggesting that engagement wanes after an initial novelty effect (van Olmen, 2022; Schmidt-Kraepelin et al., 2019; Mustafa et al., 2022). A particularly strong influence limiting patients' engagement is that many products assume high levels of digital and health literacy, which does not reflect the target population (AshaRani et al., 2021). Notably, participants observed a frequent overestimation of patients' capability to interact with support technology. An awareness of such limitations might be provided through adequate stakeholder involvement, but systems will still be required to overcome the challenges of retaining patient engagement over time.

1. English: Netherlands Organisation for Applied Scientific Research

C3: Managing Evolving Domain Knowledge The next challenge stems from the fact that DLM is a heavily researched topic, with roughly 15.000 new publications in 2022 on PubMed (McEntyre and Lipman, 2001). This information overload makes it difficult for developers to incorporate the latest insights into their products. Participants point out that relying on healthcare professionals to keep up with the evolving domain knowledge is not a solution due to their high workload.

C4: Accounting for Interdependence Finally, healthcare systems are complex institutions comprising many different actors that need to collaborate in the treatment of diabetes, e.g., a doctor requiring both patient input and insurance coverage information. In particular, DLM support involves a significant degree of *interdependence* (Gerpott et al., 2018) within a complex web of stakeholders. That is, different parties beyond the patient (e.g., healthcare professionals or family members) have relevant outcomes associated with the process, and crucially, these may strongly depend on one another. Importantly, stakeholders may also pursue conflicting interests or be subject to different constraints, e.g., healthcare providers might prioritize cost minimization, and healthcare professionals may optimize time budgets. While support products are envisioned to partake in this overarching collaborative environment, how to make these appropriately responsive remains largely unclear.

3. Hybrid Intelligence-based Research Trajectories

In the following, we outline four HI-based research trajectories (identified as *T1-T3* below), each touching on one or more of the above challenges.

T1: Developing Methods for Adaptive Stakeholder Involvement As discussed, it is a major challenge to meaningfully involve a significant number of heterogeneous stakeholder groups throughout a product's life cycle ($\rightarrow$*C1*), especially in high interdependence settings ($\rightarrow$*C4*). Consequently, continuous involvement requires support products with the right capabilities to integrate dynamically into the web of ongoing collaborations.

HI's socio-technical perspective can effectively guide these developments by emphasizing collaboration among agents and inspiring solutions that actively leverage interdependencies for better outcomes. Existing design tools, such as the *Coactive Design Method* (Johnson et al., 2014), could form a meaningful starting point to design for this purpose. Ultimately, however, we contend that novel design methodologies and tools must be created to ensure adaptive involvement over time and organizational embedding Sherson et al. (2023); Dell'Anna et al. (2024). Additionally, the continuous identification, elimination, and mitigation of risks emerging from this embedding may require the development of adequate control structures (Rasmussen, 1997).

T2: Developing Inclusive Interaction Capabilities We contend that long-term patient engagement ($\rightarrow$*C2*) requires support products that can explicitly account for differences between patient capabilities and adapt their functionality dynamically in interaction. Importantly, this should also encompass differences in beliefs and goals and how they might relate to relevant outcomes for patients and other stakeholders in the DLM process ($\rightarrow$*C4*). We believe HI provides a suitable conceptual framework for guiding this enterprise because it considers not only adaptivity but also responsibility, explainability, and collaboration as central (Akata et al., 2020). Building on this, we identify the following three key capabilities:

- **Dynamic Conversational Expression** Supporting DLM requires dynamic, ongoing interactions to keep patients engaged in the process and cater to their specific informational needs and capabilities (Kamphorst and Kalis, 2015; Kamphorst, 2017). Conversational agents are the dominant paradigm to facilitate these in healthcare support products (Laranjo et al., 2018; Milne-Ives et al., 2020; Kramer et al., 2020; Parmar et al., 2022). Recent advancements in the development of large language models (LLMs) allow for natural language content to be created dynamically (Dai et al., 2019; Brown et al., 2020), opening new avenues for linguistic variation in the interactions (Sallam, 2023). Integrating reinforcement learning techniques could improve long-term engagement further (Gao et al., 2018; Zou et al., 2019) by introducing continual optimization in the provided support. However, using such technologies is not without risk: LLMs are known to have shortcomings – such as hallucinating factually wrong content, data biases, and a lack of interpretability – that have the potential for disastrous consequences in healthcare contexts (Yang et al., 2023). Therefore, research must be accompanied by additional efforts to understand, mitigate, and explain these risks.

- **Dynamic Patient Modeling** To effectively support patients, understanding their drives and barriers to lifestyle changes is required. To do so, prior research suggested computational representations in terms of values, norms and similar broad motivational construct (Tielman et al., 2018; Kließ et al., 2019; Cranefield et al., 2017; Ajmeri et al., 2017). However, these can be too static to account for dynamic changes in situational influences over time.

  Consequently, an adaptive patient modeling approach where a fine-grained model can be updated by observing patients and their context within face-to-face interactions could be needed Specifically, because many existing collaborations between a patient and a healthcare professional play out in *face-to-face interactions*, understanding perceptions of interdependence (Gerpott et al., 2018) in these settings seems crucial for effective AI-based support (e.g., based on automatic analysis of nonverbal behavior (Dudzik et al., 2021)). However, it is still unclear to what extent relevant high-level concepts for patient models can be inferred solely from observing actors' behavior (Armstrong and Mindermann, 2018). As such, research on (pro-)active interaction capabilities for building understanding (e.g., through dialog) rather than purely passive monitoring is crucial.

- **Effective Theory of Mind Reasoning** Apart from modeling patients individually, supporting effective collaborative interactions also requires understanding their differences in beliefs and goals with other stakeholders – e.g., in medical consultation meetings. For this reason, we propose that intelligent solutions must possess a computational Theory of Mind (ToM) of involved parties (Erdogan et al., 2022). Research must strive to develop solutions for identifying conflicts in beliefs within the data-intensive and privacy-sensitive setting spanned by DLM. Moreover, a pragmatic challenge for implementing effective ToM reasoning in a support product is computational efficiency (De Weerd et al., 2013, 2022; Baker et al., 2011).

T3: Facilitating Collaborative Support for Knowledge Management Finally, we believe that HI-based solutions are well-equipped for addressing the rapidly changing knowledge relevant to DLM (→*C3*). In particular, HI's focus on collaboration and its understanding of technological applications within larger socio-technical systems facilitate identifying where and how functionality could be provided. In essence, it encourages not only addressing but also capitalizing on the setting's interdependencies (→*C4*). For example, HI-based solutions might focus on continuously providing healthcare professionals with timely information about specific patients instead of directly attempting to advise patients. Overall, this approach could effectively synergize with existing human capabilities, resulting in better outcomes (e.g., more effective consultations or reduced workloads).

Essential for such *collaborative support* are computational knowledge structures that can be continuously updated. Knowledge graphs are a popular tool for this purpose (El-Sappagh and Elmogy, 2020; Cote and Robboy, 1980), as they are both intuitive to create and interpretable for humans (Tiddi and Schlobach, 2022). While creating them has traditionally been labor-intensive and prone to bias, new data-driven approaches can facilitate or even automate this process (Asim et al., 2018; Cimiano and Völker, 2005). Additionally, it will be worthwhile examining how a product can teach such knowledge to a domain expert (e.g, akin to TrueLearn (Bulathwela et al., 2020)). Research into knowledge management is vital to further facilitate development, especially in efforts to combine symbolic knowledge representations with data-driven techniques.

4. Conclusion

In this viewpoint, we have argued that Hybrid Intelligence (HI) provides a suitable lens for developing AI-based Diabetes Lifestyle Management (DLM) support products. Considering products as part of a larger, dynamic, and interdependent socio-technical system will help address many issues currently preventing effective solutions. Based on participatory design sessions, we have presented a series of pragmatic challenges to be addressed (identified as *C1-C4* in *Section 2* above) and outlined relevant research trajectories for doing so (identified as *T1-T3* in *Section 3* above). We plan to actively pursue these trajectories in future research, working towards an effective HI system for DLM support.

Acknowledgements

Bernd Dudzik and Jasper van der Waa contributed equally to this article. We thank all participants in the participatory design sessions and expert workshop. This research was conducted as part of a collaboration between the research programs *Hybrid Intelligence* (024.004.022) and *Ethics of Socially Disruptive Technologies* (024.004.031) – both funded through the Gravitation program of the Dutch Ministry of Education Culture, and Science and the Netherlands Organization for Scientific Research –, and TNO's Appl.AI early research program.

References

Ajmeri, N., Murukannaiah, P. K., Guo, H., and Singh, M. P. (2017). Arnor: Modeling social intelligence via norms to engineer privacy-aware personal agents. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 230–238.

Akata, Z., Balliet, D., De Rijke, M., Dignum, F., Dignum, V., Eiben, G., Fokkens, A., Grossi, D., Hindriks, K., Hoos, H., et al. (2020). A research agenda for hybrid intelligence: augmenting human intellect with collaborative, adaptive, responsible, and explainable artificial intelligence. *Computer*, 53(08):18–28.

Armstrong, S. and Mindermann, S. (2018). Occam's razor is insufficient to infer the preferences of irrational agents. *Advances in neural information processing systems*, 31.

AshaRani, P., Jue Hua, L., Roystonn, K., Siva Kumar, F. D., Peizhi, W., Ying Jie, S., Shafie, S., Chang, S., Jeyagurunathan, A., Boon Yiang, C., et al. (2021). Readiness and acceptance of ehealth services for diabetes care in the general population: Cross-sectional study. *Journal of medical Internet research*, 23(9):e26881.

Asim, M. N., Wasim, M., Khan, M. U. G., Mahmood, W., and Abbasi, H. M. (2018). A survey of ontology learning techniques and applications. *Database*, 2018.

Baker, C., Saxe, R., and Tenenbaum, J. (2011). Bayesian theory of mind: Modeling joint belief-desire attribution. In *Proceedings of the annual meeting of the cognitive science society*, volume 33.

Bødker, S. and Kyng, M. (2018). Participatory design that matters—facing the big issues. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 25(1):1–31.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Bulathwela, S., Perez-Ortiz, M., Yilmaz, E., and Shawe-Taylor, J. (2020). Truelearn: A family of bayesian algorithms to match lifelong learners to open educational resources. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 565–573.

Cimiano, P. and Völker, J. (2005). A framework for ontology learning and data-driven change discovery. In *Proceedings of the 10th international conference on applications of natural language to information systems (NLDB)*, volume 3513, pages 227–238. Springer.

Contreras, I. and Vehi, J. (2018). Artificial intelligence for diabetes management and decision support: literature review. *Journal of medical Internet research*, 20(5):e10775.

Cote, R. A. and Robboy, S. (1980). Progress in medical information management: Systematized nomenclature of medicine (snomed). *Jama*, 243(8):756–762.

Cranefield, S., Winikoff, M., Dignum, V., and Dignum, F. (2017). No pizza for you: Value-based plan selection in bdi agents. In *IJCAI*, pages 178–184.

Dai, Z., Yang, Z., Yang, Y., Carbonell, J., Le, Q. V., and Salakhutdinov, R. (2019). Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*.

De Weerd, H., Verbrugge, R., and Verheij, B. (2013). How much does it help to know what she knows you know? an agent-based simulation study. *Artificial Intelligence*, 199:67–92.

De Weerd, H., Verbrugge, R., and Verheij, B. (2022). Higher-order theory of mind is especially useful in unpredictable negotiations. *Autonomous Agents and Multi-Agent Systems*, 36(2):30.

Dell'Anna, D., Murukannaiah, P. K., Dudzik, B., Grossi, D., Jonker, C. M., Oertel, C., and Yolum, P. (2024). Toward a quality model for hybrid intelligence teams. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, pages 434–443.

Dellermann, D., Ebel, P., Söllner, M., and Leimeister, J. M. (2019). Hybrid intelligence. *Business & Information Systems Engineering*, 61:637–643.

D'Haeseleer, I., Gielis, K., and Vanden Abeele, V. (2021). Human-centred design of self-management health systems with and for older adults: Challenges and practical guidelines. In *International Conference on Information and Communication Technologies for Ageing Well and e-Health (ICT4AWE)*, volume 7, pages 90–102.

Dudzik, B., Columbus, S., Hrkalovic, T. M., Balliet, D., and Hung, H. (2021). Recognizing Perceived Interdependence in Face-to-Face Negotiations through Multimodal Analysis of Nonverbal Behavior. In *Proceedings of the 2021 International Conference on Multimodal Interaction*, number 1, pages 121–130, New York, NY, USA. ACM.

El-Sappagh, S. and Elmogy, M. M. (2020). Medical case based reasoning frameworks: Current developments and future directions. *Virtual and Mobile Healthcare: Breakthroughs in Research and Practice*, pages 516–552.

Erdogan, E., Dignum, F., Verbrugge, R., and Yolum, P. (2022). Abstracting minds: Computational theory of mind for human-agent collaboration. *HHAI2022: Augmenting Human Intellect*, 354:199–211.

Gao, J., Galley, M., and Li, L. (2018). Neural approaches to conversational ai. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 1371–1374.

Gerpott, F. H., Balliet, D., Columbus, S., Molho, C., and de Vries, R. E. (2018). How do people think about interdependence? a multidimensional model of subjective outcome interdependence. *Journal of Personality and Social Psychology*, 115:716–742.

Goel, V. (1995). *Sketches of thought*. MIT press.

Hunsucker, A. J. and Siegel, M. A. (2015). Once upon a time: Storytelling in the design process. In *Proceedings of the 3rd International Conference for Design Education Researchers*, volume 1, pages 443–454.

Johnson, M., Bradshaw, J. M., Feltovich, P. J., Jonker, C. M., Van Riemsdijk, M. B., and Sierhuis, M. (2014). Coactive design: Designing support for interdependence in joint activity. *Journal of Human-Robot Interaction*, 3(1):43–69.

Kamphorst, B. A. (2017). E-coaching systems: What they are, and what they aren't. *Personal and Ubiquitous Computing*, 21(4):625–632.

Kamphorst, B. A. and Kalis, A. (2015). Why option generation matters for the design of autonomous e-coaching systems. *AI & society*, 30:77–88.

Kließ, M. S., Stoelinga, M., and van Riemsdijk, M. B. (2019). From good intentions to behaviour change: Probabilistic feature diagrams for behaviour support agents. In *PRIMA 2019: Principles and Practice of Multi-Agent Systems: 22nd International Conference, Turin, Italy, October 28–31, 2019, Proceedings 22*, pages 354–369. Springer.

Kramer, L. L., Ter Stal, S., Mulder, B. C., de Vet, E., and van Velsen, L. (2020). Developing embodied conversational agents for coaching people in a healthy lifestyle: scoping review. *Journal of medical Internet research*, 22(2):e14058.

Laranjo, L., Dunn, A. G., Tong, H. L., Kocaballi, A. B., Chen, J., Bashir, R., Surian, D., Gallego, B., Magrabi, F., Lau, A. Y. S., and Coiera, E. (2018). Conversational agents in healthcare: a systematic review. *Journal of the American Medical Informatics Association*, 25(9):1248–1258.

Leveson, N. G. (2012). *Engineering a Safer World: Systems Thinking Applied to Safety*. MIT Press, Cambridge, MA, USA.

Lie, S. S., Karlsen, B., Oord, E. R., Graue, M., and Oftedal, B. (2017). Dropout from an ehealth intervention for adults with type 2 diabetes: a qualitative study. *Journal of medical Internet research*, 19(5):e187.

Magkos, F., Hjorth, M. F., and Astrup, A. (2020). Diet and exercise in the prevention and treatment of type 2 diabetes mellitus. *Nature Reviews Endocrinology*, 16(10):545–555.

Makarius, E. E., Mukherjee, D., Fox, J. D., and Fox, A. K. (2020). Rising with the machines: A sociotechnical framework for bringing artificial intelligence into the organization. *Journal of Business Research*, 120:262–273.

McEntyre, J. and Lipman, D. (2001). Pubmed: bridging the information gap. *Cmaj*, 164(9):1317–1319.

Milne-Ives, M., de Cock, C., Lim, E., Shehadeh, M. H., de Pennington, N., Mole, G., Normando, E., and Meinert, E. (2020). The effectiveness of artificial intelligence conversational agents in health care: systematic review. *Journal of medical Internet research*, 22(10):e20346.

Mitchell, R. and Nørgaard, M. (2011). Using diy cartoon storyboards, live sketching and co-sketching to involve young and older users in participatory design. In *Proceedings of IASDR2011, the 4th World Conference on Design Research*, volume 31.

Mustafa, A. S., Ali, N., Dhillon, J. S., Alkawsi, G., and Baashar, Y. (2022). User engagement and abandonment of mhealth: a cross-sectional survey. In *Healthcare*, volume 10, page 221. MDPI.

Okemah, J., Peng, J., and Quiñones, M. (2018). Addressing clinical inertia in type 2 diabetes mellitus: a review. *Advances in therapy*, 35(11):1735–1745.

Parmar, P., Ryu, J., Pandya, S., Sedoc, J., and Agarwal, S. (2022). Health-focused conversational agents in person-centered care: a review of apps. *NPJ digital medicine*, 5(1):21.

Rafner, J., Bantle, C., Dellermann, D., Söllner, M., Zaggl, M. A., and Sherson, J. (2022). Towards hybrid intelligence workflows: Integrating interface design and scalable deployment. In *HHAI2022: Augmenting Human Intellect*, pages 310–313. IOS Press.

Rasmussen, J. (1987). Mental models and the control of actions in complex environments. Technical Report 2656, Risø National Laboratory, Risø.

Rasmussen, J. (1997). Risk management in a dynamic society: a modelling problem. *Safety science*, 27(2-3):183–213. Publisher: Elsevier.

Sallam, M. (2023). Chatgpt utility in healthcare education, research, and practice: Systematic review on the promising perspectives and valid concerns. In *Healthcare*, volume 11, page 887. MDPI.

Schmidt-Kraepelin, M., Thiebes, S., Stepanovic, S., Mettler, T., and Sunyaev, A. (2019). Gamification in health behavior change support systems-a synthesis of unintended side effects. In *Proceedings of the 14th International Conference on Wirtschaftsinformatik*, pages 1032–1046.

Sherson, J. F., Rabecq, B., Dellermann, D., and Rafner, J. (2023). A multi-dimensional development and deployment framework for hybrid intelligence. In *HHAI*, pages 429–432.

Sloane, M., Moss, E., Awomolo, O., and Forlano, L. (2022). Participation is not a design fix for machine learning. In *Equity and Access in Algorithms, Mechanisms, and Optimization*, pages 1–6.

Spinuzzi, C. (2005). The methodology of participatory design. *Technical communication*, 52(2):163–174.

Sun, H., Saeedi, P., Karuranga, S., Pinkepank, M., Ogurtsova, K., Duncan, B. B., Stein, C., Basit, A., Chan, J. C., Mbanya, J. C., et al. (2022). Idf diabetes atlas: Global, regional and country-level diabetes prevalence estimates for 2021 and projections for 2045. *Diabetes research and clinical practice*, 183:109119.

Tiddi, I. and Schlobach, S. (2022). Knowledge graphs as tools for explainable machine learning: A survey. *Artificial Intelligence*, 302:103627.

Tielman, M. L., Jonker, C. M., and van Riemsdijk, M. B. (2018). What should i do? deriving norms from actions, values and context. In *MRC@ IJCAI*, pages 35–40.

van Olmen, J. (2022). The promise of digital self-management: a reflection about the effects of patient-targeted e-health tools on self-management and wellbeing. *International Journal of Environmental Research and Public Health*, 19(3):1360.

Vos, J. F. and Achterkamp, M. C. (2006). Stakeholder identification in innovation projects: Going beyond classification. *European Journal of Innovation Management*, 9(2):161–178.

Wiethof, C. and Bittner, E. (2021). Hybrid intelligence-combining the human in the loop with the computer in the loop: a systematic literature review. In *Forty-Second International Conference on Information Systems, Austin*, pages 1–17.

Yang, R., Tan, T. F., Lu, W., Thirunavukarasu, A. J., Ting, D. S. W., and Liu, N. (2023). Large language models in health care: Development, applications, and challenges. *Health Care Science*, 2(4):255–263.

Ziefle, M. et al. (2010). *Human-centered design of e-health technologies: concepts, methods and applications: concepts, methods and applications.* IGI Global.

Zou, L., Xia, L., Ding, Z., Song, J., Liu, W., and Yin, D. (2019). Reinforcement learning to optimize long-term user engagement in recommender systems. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2810–2818.