

Document Version

Final published version

Licence

CC BY-NC

Citation (APA)

Van Zoelen, E., Mioch, T., Tajaddini, M., Fleiner, C., Tsaneva, S., Camin, P., Gouvêa, T. S., Baraka, K., De Boer, M. H. T., & Neerincx, M. A. (2023). Developing Team Design Patterns for Hybrid Intelligence Systems. In P. Lukowicz, S. Mayer, J. Koch, J. Shawe-Taylor, & I. Tiddi (Eds.), *HHAi 2023: Augmenting Human Intellect - Proceedings of the 2nd International Conference on Hybrid Human-Artificial Intelligence* (pp. 3-16). (Frontiers in Artificial Intelligence and Applications; Vol. 368). IOS Press. <https://doi.org/10.3233/FAIA230071>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Developing Team Design Patterns for Hybrid Intelligence Systems

Emma VAN ZOELEN^{a,b}, Tina MIOCH^{a,c,1}, Mani TAJADDINI^a,
Christian FLEINER^d, Stefani TSANEVA^{e,f}, Pietro CAMIN^g, Thiago S. GOUVÊA^h,
Kim BARAKAⁱ, Maaïke H. T. DE BOER^b and Mark A. NEERINCX^{a,b}

^a*Delft University of Technology, the Netherlands*

^b*TNO, the Netherlands*

^c*University of Applied Sciences Utrecht, the Netherlands*

^d*Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany*

^e*Vienna University of Economics and Business, Austria*

^f*TU Wien, Austria*

^g*University of Twente, the Netherlands*

^h*German Research Center for Artificial Intelligence (DFKI), Germany*

ⁱ*VU Amsterdam, the Netherlands*

ORCID ID: Emma van Zoelen <https://orcid.org/0000-0002-8793-9242>, Tina Mioch

<https://orcid.org/0000-0003-2457-5872>, Mani Tajaddini

<https://orcid.org/0000-0001-9038-3524>, Christian Fleiner

<https://orcid.org/0000-0003-4280-670X>, Stefani Tsaneva

<https://orcid.org/0000-0002-0895-6379>, Pietro Camin

<https://orcid.org/0000-0002-0943-5575>, Thiago S. Gouvêa

<https://orcid.org/0000-0002-0727-5838>, Kim Baraka

<https://orcid.org/0000-0003-4381-4234>, Maaïke H. T. de Boer

<https://orcid.org/0000-0002-2775-8351>, Mark A. Neerincx

<https://orcid.org/0000-0002-8161-5722>

Abstract. With [artificial intelligence \(AI\)](#) systems entering our working and leisure environments with increasing adaptation and learning capabilities, new opportunities arise for developing [hybrid \(human-AI\) intelligence \(HI\)](#) systems, comprising new ways of collaboration. However, there is not yet a structured way of specifying design solutions of collaboration for [hybrid intelligence \(HI\)](#) systems and there is a lack of best practices shared across application domains. We address this gap by investigating the generalization of specific design solutions into design patterns that can be shared and applied in different contexts. We present a human-centered bottom-up approach for the specification of design solutions and their abstraction into team design patterns. We apply the proposed approach for 4 concrete [HI](#) use cases and show the successful extraction of team design patterns that are generalizable, providing re-usable design components across various domains. This work advances previous research on team design patterns and designing applications of [HI](#) systems.

Keywords. Hybrid Intelligence, Team Design Patterns, Use-case based research, Human-centered AI, Co-evolution, Interdependence

¹Corresponding Author: Tina Mioch, tina.mioch@hu.nl

1. Introduction

The integration of AI in various aspects of our lives is becoming increasingly prevalent, resulting in a growing frequency of human-AI interaction. The advancements in AI research enable new opportunities for technology to perform part of the work autonomously, e.g., in the medical, financial, legal, and military fields [1]. This leads to new research challenges regarding the cooperation between the AI-technology and humans, which receives attention from both business perspectives, e.g., Google [2], IBM [3], and research perspectives (e.g., [4,5]). Several research & development agendas have emphasized the need to address these challenges, in particular on creating synergy between human and AI systems through cooperation on the individual and teaming levels [6,7].

This synergy can be achieved by HI, which can be characterized as a human-AI system that is able to accomplish complex goals by combining human and AI, collectively working on a shared objective, being dependent on each other, and by co-evolving through mutual adaptation and learning, both implicitly as well as explicitly. The general rationale is (1) that humans and AI have complementary capabilities that, when combined, augment each other [8], and (2) that humans and AI will co-learn over time, adapting to each other and to the dynamic environment in which they operate [9]. The design of such a situated interdependent development process involves the establishment (and possibly growth) of a symbiotic relationship between humans and AI that benefits humans.

Research and development of AI are often focused on designing specific (technical) solutions for demarcated tasks in certain use cases. Hybrid Intelligence involves a more holistic approach to advance human-AI work processes, addressing the relevant dyadic, team, organizational and societal aspects of these processes [5]. The main research challenge for HI has been identified as “how to build adaptive intelligent systems that augment rather than replace human intelligence, leverage our strengths, and compensate for our weaknesses while taking into account ethical, legal, and societal considerations” [6], with four subchallenges: collaborative HI, adaptive HI, responsible HI, and explainable HI. This requires the creation of reusable and generic solutions to hybrid intelligence problems, as well as frameworks with which solutions can be designed and implemented [10,11]. One of these generic solutions is the specification of design patterns, focusing on the human-AI interactions and collaborations in a human-AI team. For that reason, a group of HI and domain experts met at the workshop *Human-Centered Design of Symbiotic Hybrid Intelligence (HCDSHI 2022)*², held in June 2022 in Amsterdam, to explore the design and evaluation of HI systems and the application of design patterns for the iterative design of such systems.

We investigated the following research question:

Can specific design solutions of hybrid intelligence systems be abstracted into team design patterns in a way that they can be shared and applied in different contexts?

In this paper, we present an approach for extracting generalizable team design patterns from specific use cases and show that these design patterns are applicable across domains. We applied a bottom-up approach, starting with use cases in which hybrid in-

²<https://ii.tudelft.nl/humancenteredsymbioticHI/>

telligence systems are expected to improve performance, and abstracting general design patterns for these HI systems from the use cases. We begin by describing the background of hybrid intelligence and applications of design patterns for the development of HI systems in Section 2, continue with an overview of the applied approach (Section 3) and then present concrete use cases, discussed during the workshop (Section 4). We apply pattern engineering to identify HI patterns (Section 5) and discuss accompanying research challenges (Section 6). Lastly, we conclude in Section 7.

2. Background

Hybrid Intelligence In the last years, more and more research has been done on hybrid intelligence. The influence of AI technology on our daily lives increases, leading to the necessity of AI systems to synergistically work with humans [5]. We adopt the definition of hybrid intelligence as “the ability to achieve complex goals by combining human and AI, thereby reaching superior results to those each of them could have accomplished separately, and continuously improve by learning from each other.” [8]. This contrasts hybrid intelligence with, for example, human-in-the-loop learning, in which a human interacts with the AI system during the learning phase to improve the system through human involvement [12]. In HI, the human’s and AI’s capabilities are augmented, leveraging the strengths of the individual actors, and compensating for their weaknesses. Another essential aspect is that the human-AI system improves, with both agents (human and artificial) learning from each other and their surrounding environment.

In the last years, first steps have been made towards developing design principles for hybrid intelligence. For example, Dellerman et al. [13] specify prescriptive knowledge about form and function (i.e., design principles) as well as principles of implementation (i.e., a specific instantiation of these design principles for a HI system, a decision support for business model validation). Additionally, in [12], a systematic literature review was conducted on hybrid intelligence learning processes in which three patterns and three sub-patterns of collaborative learning processes were derived, namely exploration, integration and decision support (with corresponding sub-patterns assimilation, exploitation and explanation). Although these are good first steps, there is not yet a structured way of specifying design solutions for hybrid intelligence systems and best practices that can be shared across the different application domains. A possible way of specifying these design solutions is through design patterns.

Design Patterns. Design patterns offer a framework for organizing and sharing design knowledge within a particular field [14,15]. They consist of elements that can guide the design process and provide a general understanding of how to approach a design problem [16]. Design patterns are not meant to be rigid templates that must be followed exactly; instead, they offer a general idea of how to address a recurring problem [17] and can be adapted to the design problem at hand [18]. Each instantiation of a design pattern is dependent on the specific situation in which it is being used. As the demand for designing systems that involve both humans and AI grows, it becomes increasingly important to document and share successful approaches and experiences in the field. Design patterns have the potential to reduce development time, avoid repeating work, easily disseminate design knowledge, and create stable and cohesive systems.

There has been some research done on design patterns for hybrid systems. Van Bekkum et al. [19] suggested design patterns for neuro-symbolic (AI) systems, showing that a large number of systems could be composed by a small set of elementary patterns as building blocks. Van Diggelen et al. [11] introduced the concept of team design patterns as design patterns extended to meet varying teaming needs and teaming contexts and proposed a graphical language for describing the design choices, which has been applied to, e.g., decision-making in human-AI teams for first response [20] and moral decision-making [10]. However, there has not yet been shown that these team design patterns, extracted from design solutions of hybrid intelligence systems, can actually be shared and applied in different contexts and systems.

Human-Centered Design. As described above, the area of hybrid intelligence focuses specifically on combining the strengths of humans and AI systems by designing and developing the hybrid system as a whole. A related development is that in recent years, it has become more common within the broader AI research community to take a human-centered approach, popularizing the term *human-centered AI* [21]. This development was influenced by the research area of human-centered design. According to ISO standard [22], “human-centered design is an approach to interactive systems development that aims to make systems usable and useful by focusing on the users, their needs and requirements [...]”. Several people have emphasized the importance of taking a human-centered approach to designing AI as a way of creating AI systems that benefit people and contribute positively to society [23,24]. This fits well with the hybrid intelligence mindset, in which all agents, and specifically humans, benefit and grow from their collective actions.

Practically speaking, taking a human-centered approach often means working with tools like personas to thoroughly investigate the humans interacting with your designed system before deciding what the system should do technically. It often involves engaging with users through interviews or ethnography and studying the essential values in the context of a use case. Within our work, we have taken elements from this approach by looking at the users and stakeholders within a use case first, rather than at the functionalities of the technology. We have used methods similar to personas and storyboards (see [25,26] for examples of personas and storyboards for human-centered AI) in the process of designing hybrid human-AI systems, to ensure a holistic view on the socio-technical system in which an AI system operates, rather than to focus solely on technical challenges.

3. Approach

With the goal of answering the research question, we organized the HCDSHI’22 workshop co-located with the *first International Conference on Hybrid Human-Artificial Intelligence (HHAI’22)*³ to facilitate the bottom-up process of analyzing concrete use cases to move to drafts of design patterns. The main body of the workshop consisted of a work session in which participants were divided into four groups of about six people. Each group was assigned one of the following use cases to work on: *First Response*; *Au-*

³<https://www.hhai-conference.org/hhai2022/>

tonomous monitoring of animal wildlife; Personalized (emotional) care; Assembly and maintenance process.


The groups were guided through 5 phases with storyboarding material (paper, markers, post-its, printed illustrations of stakeholders and robots, etc.; see Figure 1) and a series of templates that posed questions to the group about design-pattern-related aspects of the use cases in a human-centered manner, based on the Socio-Cognitive Engineering method [27]. The templates focused on user story, functions of the AI system(s), effects of the system and how such effects could be measured. The complete templates can be found at [28]. The structure of the activities was as follows:

- (a) Use case pitches: the use cases mentioned above were based on work submitted by participants. The authors of the use case explained the use case in a three-minute pitch.
- (b) Defining stakeholders: the groups defined all human users and other stakeholders involved in the use case. The template centered around the ‘user story’ and contained questions about the motivations, values, and goals of these people, motivating participants to define the stakeholders as personas (providing name, age, and hobby).
- (c) Mapping the hybrid system: the groups mapped out the hybrid human-AI system relevant to their use case using the storyboarding material. The focus was on the interactions between the different humans and AI systems, how each might learn from such an interaction, and the information flows within the system (see Fig. 1 for an example).
- (d) Defining functions: the groups derived what the different functionalities of the AI systems were or should be, guided by questions about objectives, inputs, outputs, and interactions with users.
- (e) Visual presentations: the groups presented their results and the participants provided feedback based on the visual maps created and the explanations.
- (f) Iterating and defining effects: as a final step, groups iterated on their visual hybrid human-AI system maps and defined what possible (positive and negative) effects the implementation of the hybrid human-AI system might have on society, task performance, or stakeholders. If they had time left, they worked on the fourth template, which hypothesizes how such effects could be measured in an implementation or experiment.


The output of these sessions were visual maps such as the one shown in Figure 1, as well as filled-out templates. After the workshop, we organized a series of digital group meetings with participants to continue working on the use cases. We made digital versions of each use case map and translated them into more generalizable team design patterns through brainstorming sessions and discussions.

4. Use cases

In this section, the four use cases presented and discussed during the HCDSHI’22 workshop are described. For each use case, we provide a short problem description, introducing the domain, followed by a specific design solution of a possible hybrid system, including actors, tasks, and goals of the situation.



(a)



(b)

Figure 1. (a) An example of the visual mapping of a hybrid human-AI system for the *First response* use case as created during the workshop, using storyboarding materials; (b) An overview of the workshop activities.

4.1. Use case ‘First response’

First responder (FR) organizations work in environments that are highly complex, dynamic, and unpredictable, such as buildings on fire or sites affected by natural disasters. To support them, increasingly advanced AI technology is introduced, which in turn leads to new challenges, such as cognitive overload of **FRs** due to having to monitor and make decisions based on the incoming data. By introducing human-AI teams, the workload of the human can be reduced and the incident can be handled more effectively.

In the present use case we concentrate on the collaboration between unmanned vehicles and **FRs**, as visualized in Figure 2. A factory is on fire, while people might still be in the building. The human-AI team arriving at the location consists of **FRs** and several **unmanned ground vehicles (UGVs)** and **unmanned aerial vehicles (UAVs)** that can (semi-)autonomously explore the environment. The **UGV** interprets the environment, detecting objects and persons. When the **UGV** detects a victim, it analyses the status of the victim (e.g., through infrared cameras, object detection) and determines whether it can handle the situation by itself or needs (human) help. If the **UGV** decides it cannot handle the situation, it notifies a **FR** with the available information who takes over the direct communication with the victim and sets next steps. The **UAVs** also explore the environment, gathering sensor and video data. An AI system analyses this data and interprets it (e.g., object or fire detection). The interpreted information is sent to a **FR** who can act on this information by validating or discarding it, allowing the AI system to learn from such feedback.




Figure 2. Overview of the hybrid human-AI teams and their roles in the context of a first response use case.

4.2. Use case ‘Assembly and maintenance processes’




Figure 3. A hybrid human-AI maintenance team and the roles of the agents.

Industry workers must adapt to the new environment caused by the trend of mass customization in the context of Industry 4.0 [29]. Small lot sizes and wide design varieties require assembly workers to perform ad-hoc assembly and maintenance steps without prior detailed training which requires constant concentration and puts high cognitive workload on the worker. Human-AI teams pose a potential solution to reduce the workers’ cognitive workload.

In a fictional use case, illustrated in Figure 3, an inexperienced service technician works in collaboration with an AI agent to repair a mobile phone. The AI agent analyzes the broken mobile phone with object detection and queries the company’s knowledge base to provide the service technician with relevant information while the service technician improves the analysis quality by validating the results and giving feedback on the results to the AI agent.

Another design solution could be that after the AI agent identifies the problem, the AI agent provides an interactive step-by-step guidance to the service technician where the service technician can edit step descriptions by communicating with the AI agent.

4.3. Use case ‘Autonomous monitoring of animal wildlife’

Biodiversity is declining faster than at any time in human history [30], negatively impacting human well-being and economic prosperity all over the planet. Reversing biodi-

versity loss and promoting its sustainable use require evidence-based management practices. For that purpose, it has been agreed that the best available data and technology should be made accessible to experts, decision makers, and the public [31, Target 21].

In this use case, the hybrid team aims at monitoring animal wildlife in a given territory for an extended, multi-year period. A visual representation is given in Figure 4. The AI has access to on-site sensors (e.g., camera traps, microphones from static arrays [32] and aboard autonomous vehicles). The human in the team is an expert who can identify *essential biodiversity variables (EBVs)* [33], such as species- or community-level abundance, from the acquired sounds and images. The AI aims at autonomously reporting biodiversity status in a manner compatible with the language of human experts. Thus, it interacts with the expert and learns how to map sensor signals to an ontology based on *EBVs*. In this process, the AI uses self-supervised methods [34] to learn data structure, and graphically communicates to the expert what it has learned. The expert experiences this interaction as AI-aided data exploration, and provides knowledge to the AI both explicitly through data annotation as well as implicitly through browsing decisions. The AI learns about ecological phenomena and puts its beliefs to test by presenting hypotheses to the expert. The expert owns final judgment about factual truth and about how satisfactory the AI model is.




Figure 4. A hybrid human-AI team and their roles in the context of monitoring animal wildlife.

4.4. Use case 'Personalized (emotional) care'

Another concrete scenario for a human-AI team is located in the medical domain. In particular, we look at the use case presented in [35], where a social robot is employed in an emotional support youth program. Younger patients might experience stress when undergoing mental therapy, therefore, a social robot could help to overcome the communication barrier by providing a safe environment for the patient preventing them from directly communicating with an adult [36].

In the described scenario from [35], the child has one-to-one interactions with the robot where they can share personal experiences or discuss specific issues, as visualized in Figure 5. The shared verbal insights from the child are transformed into a personalized *knowledge graphs (KGs)* and specific events are mapped to an abstract emotional support *KGs*. Based on the two *KGs* the AI can select an appropriate response to provide emotional care (e.g., encouragement) to the child. The expert has a supervision role and inspects the additions to the personalized *KGs* to verify the correctness of the abstract *KGs* mappings. In case an issue is found, the expert can make improvements such as

removing, adding, or correcting elements of the **KGs** and its mappings. The AI learns from each **KGs** adjustment and improves over time.




Figure 5. A hybrid human-AI team aiming at providing personalized (emotional) care to young patients.

5. Team Design Patterns

In this section, we describe the **team design patterns (TDPs)** that were identified based on the design solutions described in the use cases.

Examining the four outlined use cases, we can extract three primary design patterns from the domain-specific hybrid intelligence solutions, namely *AI Advisor and Human Performer TDP*, *AI Performer and Human Assistant TDP*, and *AI Performer and Human Validator TDP*. We extracted the **TDPs** by iterating on the derived **TDPs** on different levels of abstraction. We targeted to extract **TDPs** at a level of abstraction for which we could show that the **TDPs** are applicable for not only the use case it has been derived from, but also for at least one of the other use cases.

In the following, we describe the derived **TDPs** in detail.

AI Advisor and Human Performer. To reduce cognitive workload for the human participant, the task of analyzing plausible choices is outsourced to the AI agent. While the AI analyses the resources and provides recommendations, the human agent has final decision on the choice of alternatives, as shown in Table 1. Examples of how this pattern can be implemented are seen in the maintenance (the AI recommends steps to the technician; Sect. 4.2) and wildlife monitoring (the AI suggests hypotheses to the expert; Sect. 4.3) use cases. This design pattern also allows *co-learning* among the team members. On one side, feedback from the human participant is used to improve the AI's recommendations over time. On the other hand, the suggestions provided to the human actor can expand their knowledge and give them new perspectives.

AI Performer and Human Assistant. The AI and the human agent have their own capabilities, which are inherently limited. While the AI is good at performing a particular task, it can happen that human assistance is required at some point in time. In such cases, the AI proactively request assistance from a human agent. An instance of this pattern is seen in the first response use case (Sect. 4.1), where a **UGV** interprets the environment and decides if it can handle the situation itself or a fire fighter's assistance is needed (e.g., when a conscious human victim is encountered). This team design pattern is represented in Table 2, where we also see that the pattern can lead to reduced cognitive workload for the human team member. This pattern relates to the idea of symbiotic autonomy where limited agents ask for help from external agents to be able to carry out their tasks [37].

Table 1. Team Design Pattern: AI Advisor and Human Performer



Description	The AI agent performs an analysis (0) and provides a set of possible options in the form of recommendations (e.g., a possible course of action) to the human participant (1), who validates the options and takes the final decision on the optimal solution (2). Lastly, the human actor responds with improvement items concerning the recommendations (3).
Visualization	
Expected effect	- The human agent saves cognitive load by outsourcing the options set creation to the AI. Thus, the quality of the human’s primary work increases. - Inexperienced human agents improve their solution deduction capabilities. - The quality of the AI’s recommendations improves over time.
Use when	- The human must retain high concentration for their primary work. - The human lacks experience to identify a wide range of options, but is experienced enough to choose the optimal solution from the given choices.


Table 2. Team Design Pattern: AI Performer and Human Assistant

Description	The AI performs a certain task (0) and monitors its own state to determine whether there are aspects that need human involvement (1). Once the AI recognizes that human intervention is necessary, it proactively requests assistance (2) from the human by communicating relevant information. The human interprets the situation and makes a decision and subsequently subsequently hands back the task/responsibility to the AI (3).
Visualization	
Expected effect	The human gets involved only when their capabilities are needed. This frees up limited (cognitive) resources of the human.
Use when	- There are limited resources and the team needs to be efficient. - The human should make some decisions, but the AI can work on other tasks autonomously.

AI Performer and Human Validator. While autonomous AI agents reduce human’s cognitive workload and let them concentrate on other tasks, human (ethical) supervision is often needed. This design pattern is represented in Table 3 and its implementations are found in all four use cases. In the personalized care example (Sect. 4.4) a domain expert supervises the AI interactions with the patient to ensure a stable and safe environment. In the first response use case (Sect. 4.1) a fire fighter validates that the AI correctly interprets the situation, while in the maintenance scenario (Sect. 4.2) the technician validates

the result of the suggested repairs. Lastly, in the wildlife monitoring scenario (Sect. 4.3) an expert validates the visual information provided by the AI and extends it by providing further annotations. In all four examples, the AI uses the feedback received from the human actor to improve over time.

Table 3. Team Design Pattern: AI Performer and Human Validator

<i>Description</i>	The AI Agent autonomously performs a certain task (0). It provides an overview of essential information without any recommended course of action to the human actor (1). The human agent has a supervision role, verifying that the information is correct (2). If they believe an adjustment is needed, they responds with improvement items (3).
<i>Visualization</i>	
<i>Expected effect</i>	- The human actor saves time and concentration power by outsourcing the information retrieval process to the AI. - The quality of the human’s primary work increases by retaining a higher level of concentration. - The quality of the AI’s information retrieval improves over time.
<i>Use when</i>	- The available information cannot be processed by the human in reasonable time or in satisfying quality. -The human must retain high concentration for their primary work. -The AI agent can autonomously perform a task, but ethical supervision is needed.

It becomes clear that a specified hybrid intelligence solution can be generalized into **TDPs** that are applicable to different domains. Thus, solutions can be shared across domains and specific hybrid intelligence implementations by means of **TDPs**, both by means of abstracting the comparable design solutions towards one **TDP**, but also by using specified **TDPs** as inspiration for specifying design solutions in other use cases.

6. Discussion

We have presented a workshop method for extracting **TDPs** from specific design solutions (based on use cases) that can serve as a reference for the newly forming research community of Hybrid Intelligence. We showed that the method was successful in deriving generalizable **TDPs** that were applicable for different domains. Participants of the workshop were **HI** experts or domain experts, with different technical backgrounds. The method and **TDPs** were new for many workshop participants; however, the participants were able to specify **HI** design solutions with the workshop method and specify **TDPs**, which shows that the method is practical and intuitive and not overly complicated.

The group size in general was perceived by the participants as appropriate; however, a limitation was the number of involved domain experts. Only a few domain experts (e.g., a firefighter for the *first response use case*) attended the workshop. As domain experts

provide valuable information to characterize the human actor's interaction strategy with an AI agent, we recommend having at least one domain expert per group and in general to increase the amount of domain experts involved.

Regarding the specification of TDPs, it is noteworthy that TDPs can have different levels of abstraction. In this paper, we have presented only relatively high-level TDPs, as our goal was to show the generalizability and applicability for different domains. This is also possible for lower level TDPs, but not necessarily for the given set of use cases presented and developed during the workshop. In fact, the workshop's groups concluded with TDPs of different abstraction levels which were later harmonized by the authors to simplify TDP profiles (which were presented in this paper). For instance, a lower level TDP included the applied technology of a use case (e.g., computer vision with stereo camera) to specify the context and effects of the TDP in more detail.

A limitation of the current work is that due to the small representation of domain experts, real-world constraints might have been overlooked in the specification of the use cases and TDPs, which might have led to oversimplification of the TDPs or that the TDPs do not correspond to practical considerations. In addition, the defined TDPs have not yet been evaluated in real, implemented HI systems.

7. Conclusion

With the progress of AI capabilities, research and industry struggle to harness and scope synergy effects in scenarios where both human and AI agents collaborate to complete a main goal (defined as HI). TDPs have been identified as a useful instrument to scope human-AI collaborative work, to highlight the requirements and the resulting effects of such collaboration, and to derive and share reusable design solutions. In this paper, we reported the procedure and the outcome of an academic workshop aiming to abstract TDPs from four defined use cases. Each use case addressed a different domain, respectively emergency response, manufacturing industry, wildlife observation, and child education. The TDPs were abstracted under the consideration of the collaboration between the miscellaneous stakeholders, including domain and AI experts. One of the three final TDPs could be applied to the context of all use cases indicating the impact on design decisions of HI systems TDPs might have, and their potential to find and share generic HI designs with stakeholders from different disciplines. Given the positive outcome of the workshop, future work will focus on evaluating the TDPs in real-world scenarios and improving the TDP abstraction and standardization process.

Acknowledgments

We would like to thank all participants of the workshop *Human-Centered Design of Symbiotic Hybrid Intelligence (HCDSHI 2022)*, held in June 2022 in Amsterdam, for their valuable input and contribution. Furthermore, we would like to emphasize that this paper is the result of a group effort, and that all co-authors contributed greatly to the results presented in this paper as well as the writing.

References

- [1] Rahwan I, Cebrian M, Obradovich N, Bongard J, Bonnefon JF, Breazeal C, et al. Machine behaviour. *Nature*. 2019 Apr;568(7753):477-86. Available from: <http://www.nature.com/articles/s41586-019-1138-y>.
- [2] Knight W. Your Best Teammate Might Someday Be an Algorithm; 2017. Publication Title: MIT Technology Review. Available from: <https://www.technologyreview.com/2017/07/10/4384/your-best-teammate-might-someday-be-an-algorithm/>.
- [3] Geyer W, Weisz J, Santos Pinhanez C, Daly E. What is human-centered AI?; 2021. Publication Title: IBM Research Blog. Available from: <https://research.ibm.com/blog/what-is-human-centered-ai>.
- [4] Johnson M, Vera AH. No Ai is an island: The case for teaming intelligence. *AI Magazine*. 2019;40(1):16-28.
- [5] Peeters MMM, van Diggelen J, van den Bosch K, Bronkhorst A, Neerincx MA, Schraagen JM, et al. Hybrid collective intelligence in a human-AI society. *AI & SOCIETY*. 2021 Mar;36(1):217-38. Available from: <https://doi.org/10.1007/s00146-020-01005-y>.
- [6] Akata Z, Balliet D, de Rijke M, Dignum F, Dignum V, Eiben G, et al. A Research Agenda for Hybrid Intelligence: Augmenting Human Intellect With Collaborative, Adaptive, Responsible, and Explainable Artificial Intelligence. *Computer*. 2020 Aug;53(8):18-28. Available from: <https://ieeexplore.ieee.org/document/9153877/>.
- [7] Seeber I, Bittner E, Briggs RO, de Vreede T, de Vreede GJ, Elkins A, et al. Machines as teammates: A research agenda on AI in team collaboration. *Information and Management*. 2020;57(2):103174. Available from: <https://doi.org/10.1016/j.im.2019.103174>.
- [8] Dellermann D, Ebel P, Söllner M, Leimeister JM. Hybrid Intelligence. *Business & Information Systems Engineering*. 2019 Oct;61(5):637-43. Available from: <http://link.springer.com/10.1007/s12599-019-00595-2>.
- [9] Van Zoelen EM, Van Den Bosch K, Neerincx M. Becoming team members: Identifying interaction patterns of mutual adaptation for human-robot co-learning. *Frontiers in Robotics and AI*. 2021;8:692811.
- [10] van der Waa J, Verdult S, van den Bosch K, van Diggelen J, Haije T, van der Stigchel B, et al. Moral Decision Making in Human-Agent Teams: Human Control and the Role of Explanations. *Frontiers in Robotics and AI*. 2021 May;8:640647. Available from: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8190710/>.
- [11] Van Diggelen J, Johnson M. Team design patterns. *HAI 2019 - Proceedings of the 7th International Conference on Human-Agent Interaction*. 2019:118-26.
- [12] Wiethof C, Bittner E. Hybrid Intelligence - Combining the Human in the Loop with the Computer in the Loop: A Systematic Literature Review; 2021.
- [13] Dellermann D, Lipusch N, Ebel P, Leimeister JM. Design principles for a hybrid intelligence decision support system for business model validation. *Electronic Markets*. 2019 Sep;29(3):423-41. Available from: <https://doi.org/10.1007/s12525-018-0309-2>.
- [14] Alexander C. A pattern language: towns, buildings, construction. Oxford university press; 1977.
- [15] Chung ES, Hong JI, Lin J, Prabaker MK, Landay JA, Liu AL. Development and evaluation of emerging design patterns for ubiquitous computing. In: *Proceedings of the 5th conference on Designing interactive systems: processes, practices, methods, and techniques*; 2004. p. 233-42.
- [16] Baltzer MCA, López D, Flemisch F. Towards an interaction pattern language for human machine cooperation and cooperative movement. *Cognition, Technology and Work*. 2019;21(4):593-606. Available from: <https://doi.org/10.1007/s10111-019-00561-8>.
- [17] Schulte A, Donath D, Lange DS. Design patterns for human-cognitive agent teaming. In: *International Conference on Engineering Psychology and Cognitive Ergonomics*. Springer; 2016. p. 231-43.
- [18] Löwgren J. Inspirational Patterns for Embodied Interaction. *Knowledge, Technology & Policy*. 2007;20(3):165-77.
- [19] van Bekkum M, de Boer M, van Harmelen F, Meyer-Vitali A, Teije At. Modular design patterns for hybrid learning and reasoning systems: a taxonomy, patterns and use cases. *Applied Intelligence*. 2021;51(9):6528-46.
- [20] Mioch T, Sterkenburg R, Beuker T, Neerincx MA. Actionable Situation Awareness: Supporting Team Decisions in Hazardous Situations. In: *Proceedings of the International ISCRAM Conference 18th International Conference on Information Systems for Crisis Response and Management, ISCRAM 2021 Blacksburg 23 May 2021 through 26 May 2021*; 2021. p. 62.

- [21] Auernhammer J. Human-centered AI: The role of Human-centered Design Research in the development of AI. In: Synergy. Held online; 2020. Available from: <https://dl.designresearchsociety.org/drs-conference-papers/drs2020/researchpapers/89>.
- [22] ISO 9241-210:2019(en), Ergonomics of human-system interaction — Part 210: Human-centred design for interactive systems;. Available from: <https://www.iso.org/obp/ui/#iso:std:iso:9241:-210:ed-2:v1:en>.
- [23] Shneiderman B. Human-Centered AI. Issues in Science and Technology. 2021;37(2):56-61. Publisher: Arizona State University. Available from: <https://www.jstor.org/stable/27092030>.
- [24] Xu W. Toward human-centered AI: a perspective from human-computer interaction. Interactions. 2019 Jun;26(4):42-6. Available from: <https://dl.acm.org/doi/10.1145/3328485>.
- [25] Ghunaim H, Weekes TR, Eskridge TC. Designing an AI Assistant for Student Telehealth: A Case Study Using Human-Centered Design. In: 2022 IEEE 3rd International Conference on Human-Machine Systems (ICHMS); 2022. p. 1-6.
- [26] Holzinger A, Kargl M, Kipperer B, Regitnig P, Plass M, Müller H. Personas for Artificial Intelligence (AI) an Open Source Toolbox. IEEE Access. 2022;10:23732-47. Conference Name: IEEE Access.
- [27] Neerincx MA, van Vught W, Blanson Henkemans O, Oleari E, Broekens J, Peters R, et al. Socio-Cognitive Engineering of a Robotic Partner for Child's Diabetes Self-Management. Frontiers in Robotics and AI. 2019;6. Available from: <https://www.frontiersin.org/articles/10.3389/frobt.2019.00118>.
- [28] Human-centered Symbiotic HI Workshop - XWiki;. Available from: <https://xwiki.ewi.tudelft.nl/xwiki/bin/view/Human-centered%20Symbiotic%20HI%20Workshop/>.
- [29] Büttner S, Mucha H, Funk M, Kosch T, Aehnelt M, Robert S, et al. The design space of augmented and virtual reality applications for assistive environments in manufacturing: a visual approach. In: Proceedings of the 10th International Conference on Pervasive Technologies Related to Assistive Environments; 2017. p. 433-40.
- [30] IPBES. Global assessment report on biodiversity and ecosystem services of the Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services. Zenodo; 2019. Version Number: 1. Available from: <https://zenodo.org/record/3831673>.
- [31] Conference of the Parties to the Convention on Biological Diversity. Kunming-Montreal Global biodiversity framework (CBD/COP/15/L.25); 2022. Available from: <https://www.cbd.int/doc/c/e6d3/cd1d/daf663719a03902a9b116c34/cop-15-1-25-en.pdf>.
- [32] Sugai LSM, Silva TSF, Ribeiro Jr JW, Llusia D. Terrestrial passive acoustic monitoring: review and perspectives. BioScience. 2019;69(1):15-25.
- [33] Pereira HM, Ferrier S, Walters M, Geller GN, Jongman R, Scholes RJ, et al. Essential biodiversity variables. Science. 2013;339(6117):277-8.
- [34] LeCun Y, Misra I. Self-supervised learning: The dark matter of intelligence; 2021. Available from: <https://ai.facebook.com/blog/self-supervised-learning-the-dark-matter-of-intelligence/>.
- [35] Saccardi I, Islakoglu DS, Neerincx A, Vinella FL. Symbiotic Child Emotional Support with Social Robots and Temporal Knowledge Graphs. In: Workshop on Human-Centered Design of Symbiotic Hybrid Intelligence, collocated with the first International Conference on Hybrid Human Artificial Intelligence (HHAI2022); 2022. <https://doi.org/10.48550/arXiv.2205.13229>.
- [36] Neerincx A, Rodenburg DL, de Graaf MM, Masthoff JF. Social Robots to Support Child and Family Care: A Dutch Use Case. In: Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction; 2021. p. 367-71.
- [37] Rosenthal S, Biswas J, Veloso MM. An effective personal mobile robot agent through symbiotic human-robot interaction. In: AAMAS. vol. 10; 2010. p. 915-22.