Teaming.AI: Enabling Human-AI Teaming Intelligence in Manufacturing

Thomas Hoch[1], Bernhard Heinzl[1], Gerald Czech[2], Maqbool Khan[1,3], Philipp Waibel[4], Stefan Bachhofner[4], Elmar Kiesling[4] and Bernhard Moser[1]

[1] *Software Competence Center Hagenberg, Softwarepark 32a, Hagenberg, 4232, Austria.*
[2] *Upper Austrian Fire Brigade Association, Hauptstraße 1–5, Linz, 4041, Austria.*
[3] *Sino-Pak Center for Artificial Intelligence, PAF-IAST, Khanpur Road, Pakistan.*
[4] *Vienna University of Economics and Business, Welthandelspl. 1, Vienna, 1020, Austria.*

Abstract

Teaming.AI aims to overcome the lack of flexibility as a limiting factor of human-centered AI collaboration by envisioning a teaming framework that integrates the strengths of both, namely the flexibility of human intelligence and the scaling and processing capabilities of machine intelligence. In Teaming.AI, this will be achieved by employing a teaming model that structures the interactions between humans and AI systems, and a knowledge graph that dynamically supports the teaming model to cope with process, regulatory and context knowledge. We expect that the developed Teaming.AI platform provides the human team members with a better understanding and control of automated services and decision support within the manufacturing environment, leading to a more trustful collaboration between the human and AI.

Keywords

Fault detection and diagnosis, decision-making and cognitive processes, human-centred automation, knowledge modeling, knowledge based systems

1. Introduction

The Teaming.AI project aims to address the open problem of the “missing middle” (see [1]) in scenarios where humans and AI systems collaborate towards a common goal. This missing middle is defined along a spectrum between human-only to machine-only activities. Human only activities include leading, empathizing, creating, and judging; machine-only activities include transacting, iterating, predicting and adapting. The “missing middle” lies in between these extremes - i.e., human and machine hybrid activities. These can be broken down into teaming activities where (i) “humans complement machines” (i.e., train, explain, sustain) and (ii) “AI gives humans superpowers” (amplify, interact, embody). Such hybrid activities are neglected in the state of the art and deserve more recognition, especially given the observation that human intelligence outperforms current AI systems in a wide field of applications, particularly in terms of flexibility and taking context into account.

The envisioned Teaming.AI approach aims to support the systematic development and evolution of AI systems in manufacturing in order to address current limitations of today’s narrow AI systems. Such systems typically lack self-adaptive capabilities and the ability to assimilate and interpret new information outside of its predefined programmed parameters. They are typically tailored to solve specific tasks in a specific predefined setting; changes in this underlying setting typically requires system adaptations, ranging from fine-grained parameter adaptations to fully-fledged re-design and

Proceedings of the Workshop of I-ESA’22, March 23–24, 2022, Valencia, Spain
EMAIL: thomas.hoch@scch.at (T. Hoch); bernhard.heinzl@scch.at (B. Heinzl); gerald.czech@ooelfv.at (G. Czech); maqbool.khan@scch.at (M. Khan); philipp.waibel@wu.ac.at (P. Waibel); stefan.bachhofner@wu.ac.at (S. Bachhofner); elmar.kiesling@wu.ac.at (E. Kiesling); bernhard.moser@scch.at (B. Moser)
ORCID: 0000-0003-0074-0052 (T. Hoch); 0000-0001-8297-7533 (B. Heinzl); 0000-0001-7656-0184 (M. Khan); 0000-0002-5562-4430 (P. Waibel); 0000-0001-7785-2090 (S. Bachhofner); 0000-0002-7856-2113 (E. Kiesling); 0000-0001-8373-7523 (B. Moser)

re-development of AI systems. In order to tackle this challenge, Teaming.AI has to provide a flexible framework to specify mechanisms for collaborative self-adaptation of the overall system that may involve both human actors and AI agents. Employing a teaming model provides a flexible way for performing adaptations on multiple levels, taking inspiration from conceptual models of self-adaptive systems developed in the software engineering literature (cf. [2]).

2. Related work

The teaming intelligence of humans has been studied and practiced several decades ago in the research community to increase the productivity and lessen task completion duration. The advancement in robot technologies pushed the teaming concept to a new era of human collaboration with machines, agents and AI systems. Within recent years, the advancement in technology has created robots and AI systems being able to perform a variety of tasks in manufacturing, space, agriculture, healthcare, autonomous vehicles and in other real-life scenarios [3]. Human-robot teaming is studied from multiple perspectives, such as: concepts and design components [4], perspective on analysis and implementation issues [5], human-robot interaction theory [6], Human-robot cross-training [7], and mutual trust between human and robot in decision making [8].

[9] studied human teamwork and identified five core components for effective teaming (see Figure 1A), considering not only whether the team performed well (e.g., completed the team task) but also how the team interacted (i.e., team processes, teamwork) to achieve the team outcome. They argued, that team effectiveness can be improved by a well-designed coordination mechanisms to ensure that the “Big Five” are consistently updated and that relevant information is distributed throughout the team. Recently, research has made advancements towards achieving common goals by human and autonomous systems using their unique capabilities for specific portions of a task as a team. The current collaborative teaming concepts like human-agent teaming [10] and human-autonomy teaming [11] are the motivation of our proposed Teaming.AI platform towards a novel approach for human-AI teaming.

3. High-level perspective on teaming

Although the study of [9] focuses purely on human teaming and not human-AI teaming, we believe this theory builds a solid foundation for the digitalization of human-AI teaming interaction for two reasons. First, the clear segregation of teamwork and coordination mechanisms supports separation of concerns in digitalization. Second, we believe that team effectiveness as a goal instead of team performance keeps human team members more in the focus than AI, because team performance only incorporates the outcome of the work, while team effectiveness also takes the interactions among team members into account. To be an effective team member, the AI must take part in the coordination activities of the team, and it needs to know what information to share or when to ask for assistance. Being capable of observing one another’s state, sharing information, or requesting assistance is regarded by [12] as Teaming Intelligence. [13] captured human-AI teaming requirements beyond traditional task-based approaches towards human-autonomy teaming (i.e. human-AI teaming preserving human autonomy). We believe that this fits well to achieve team effectiveness as defined by [9]. Human autonomy teaming requires understanding interdependency. [13] defines an Interdependence Analysis tool to understand how people and automation can effectively team by providing insight into the interdependence relationships used to support one another throughout an activity.

In Teaming.AI, we follow these design principles and analyze the interdependence relationships along the four dimensions of the 4S framework as described by [12]. Starting from the analysis of team and task structure, the skills of the team actors are identified and linked to the different teaming activities. Different to [13] and their concept of jointness, we expect that, at the most granular level, an activity is either performed by a human team member or an automated AI service. However, the performer of this activity can be supported either by a human or the AI by providing additional insights the performer can rely on. We introduce abstract activities as a mechanism to model this performer/supporter pattern. We envision the supporter role as a more passive role that monitors the current state of the production process and interacts if needed, similar as described by [14].

4. Teaming model

[12] defined Teaming Intelligence as intelligently managing the interdependencies of coordination work. Teaming.AI offers a method to manage these interdependencies and interactions by modelling them in a structured manner and linking these models to relevant activities, resources, and constraints (policies). To this end, the teaming model is comprised of multiple sub-models, in particular:

Teaming Process Model: The teaming process model defines the individual teaming processes and tasks, describing the state, structure, skills, and strategy of teaming interaction between human and the Teaming.AI platform according to the 4S framework. The teaming process model is instantiated and executed by the teaming engine within the Teaming.AI platform.

Activity Model: To achieve a high interchangeability, the information that is required for the concrete/abstract activities is separated from the teaming process model. The activity model is responsible for storing and querying the activity information. The activity model enriches the activities in the teaming process model with additional information required to execute the processes, such as necessary inputs, preconditions, and generated outputs.

Event Model: The teaming model is used by the teaming engine (see Figure 1B) to orchestrate the teaming aspects of the process execution and to act in case specific events are detected. If an event is detected, the teaming engine uses the teaming process model to decide on the next tasks that must be performed, together with the information who is performing the task, by considering policies and other aspects (e.g., human skills or organizational roles).

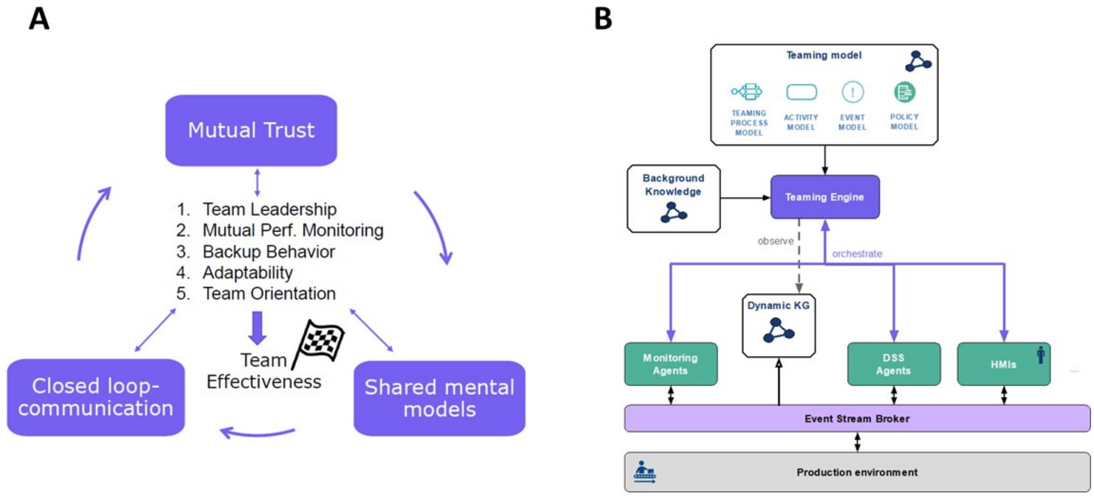

Figure 1: The Big Five of teamwork and their coordinating mechanisms (left). Overview of the Teaming.AI architecture (right).

Policy Model: The policy model enriches the overall teaming model with additional information regarding rules that control the teaming process in order to achieve effective teaming interaction and fulfill the team's goals. In particular, this encompasses external policies adhering to legal and ethical requirements or company regulations, as well as internal policies that are rules driven by the teaming process and provide a mechanism to increase flexibility and make the teaming process more adaptive at runtime.

These teaming model elements are formalized and stored in a knowledge graph, which makes it possible to associate and ground them in application-specific background knowledge – i.e., a concrete description of organizational roles and responsibilities, the production system, its resources and its environment, as well as industrial products and production processes. The teaming model should provide means to model effective teaming interaction according to the "Big Five" framework as well as enabling coordinating mechanisms that form a trust-enhancing communication cycle.

5. Teaming.AI platform overview

The Teaming.AI platform supports the development and execution of a flexible model for dynamic teaming of human stakeholders and AI systems in order to improve learning and knowledge transfer. A key goal is to enable better coordination of work sharing across teams of human agents and AI components. The central coordination element in the Teaming.AI platform is the Teaming Engine, which monitors the execution environment, tracks the dynamic context of the enacted teaming process in the production environment and applies policies to orchestrate teaming processes. This includes making decisions based on specified policies, e.g., who executes a specific task, when roles between task performer and task supporter need to be switched etc.

Figure 1B depicts the architectural components of the Teaming.AI platform. The interaction and communication are based on events, which are handled by a central event stream broker. Events can be enriched either automatically or manually with specific process knowledge (e.g., machine data or error descriptions). The knowledge graph runtime is responsible for filtering and aggregating these events into meaningful so-called complex events. These complex events are stored in a dynamic data knowledge graph and analyzed further in order to identify higher level correlations that can be used for decision making (e.g., to automate quality inspection of work pieces). With the use of a knowledge graph [15], we strive for solutions that allow for the generation of ML models that are easier to interpret and can make the derived information semantically explicit.

The knowledge in Teaming.AI has both static and dynamic parts. As static knowledge, we consider all knowledge that only changes at low frequencies (e.g., less than daily), for example product data, organizational structures, and policies. Dynamic knowledge on the other hand changes at higher frequencies, which may include data streams (e.g., state of machines or work pieces). These updates are retrieved from the event stream broker and need to be incorporated into the knowledge graph, e.g., by means of stream reasoning (see [16]) or online machine learning (see [17]).

Most current knowledge graph solutions have comparatively low update rates and would be considered static in the above frequency-of-change based definition. Hence, novel techniques are required that refine the current state of the art in knowledge graph processing. In Teaming.AI, we follow a modular approach that facilitates purpose-driven, agile construction of reusable knowledge graphs across multiple layers of abstraction and perspectives. This means e.g., that every layer of the knowledge graph represents a partial view on the real-world system that links relevant aspects for a given perspective (e.g. business / operational).

6. Conclusion

A key element for successful human-AI teamwork is a careful design and implementation of the coordinating mechanisms involved. Mutual trust increases if the appropriate amount of information is shared through a closed-loop communication between humans and AI components. The envisioned Teaming.AI platform has the goal to orchestrate the information exchange and to organize the collected information within a layered knowledge graph, reduce the information to its key aspects and semantically enrich this knowledge with context information. Transparent storage and processing of information is the foundation for a decision support system that can be understood and further analyzed by human team members.

7. Acknowledgements

The project Teaming.AI has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No. 957402.

8. References

[1] P. R. Daugherty, H. J. Wilson, Human + machine: Reimagining work in the age of AI, Harvard Business Press, Harvard, 2018.

[2] D. Weyns, Software engineering of self-adaptive systems: an organised tour and future challenges, 2017. URL: https://people.cs.kuleuven.be/~danny.weyns/papers/2017HSE.pdf
[3] A. Bauer, D. Wollherr, M. Buss, Human–robot collaboration: A survey, International Journal of Humanoid Robotics 05 (2008) 47–66. doi: 10.1142/S0219843608001303.
[4] L. M. Ma, T. Fong, M. Micire, Y. Kim, K. M. Feigh, Human-robot teaming: Concepts and components for design, 2017. URL: http://www.fsr.ethz.ch/papers/FSR_2017_paper_57.pdf
[5] A. Chella, F. Lanza, A. Pipitone, V. Seidita, Human-robot teaming: Perspective on analysis and implementation issues, in: A. Finzi, A Farinelli, S. Anzalone, F. Mastrogiovanni, CEUR Workshop Proceedings, vol. 2352. URL: http://ceur-ws.org/Vol-2352/short3.pdf
[6] N. C. Krämer, A. M. R. von der Pütten, S. C. Eimler, Human-agent and human-robot interaction theory: Similarities to and differences from human-human interaction, in: M. Zacarias, J. Valente de Oliveira (Eds.), Human-Computer Interaction: The Agency Perspective, Springer, Berlin, 2012, pp. 215-240. doi: 10.1007/978-3-642-25691-2_9.
[7] S. Nikolaidis, J. A. Shah, Human-robot cross-training: Computational formulation, modeling and evaluation of a human team training strategy, in: 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI), IEEE, New York, 2013, pp. 33-40. doi: 10.1109/HRI.2013.6483499.
[8] M. Chen, S. Nikolaidis, H. Soh, D. Hsu, S. Srinivasa, Planning with trust for human-robot collaboration, in: Proceedings of the 2018 ACM/IEEE International Conference on Human Robot Interaction, IEEE, New York, 2018, pp. 307–315. doi: 10.1145/3171221.3171264.
[9] E. Salas, D. E. Sims, C. S. Burke, Is there a "big five" in teamwork?, Small group research 36 (2005) 555-599.
[10] J. Y. C. Chen, M. J. Barnes, Human–agent teaming for multirobot control: A review of human factors issues, IEEE Transactions on Human-Machine Systems 44 (2014) 13-29. doi: 10.1109/THMS.2013.2293535.
[11] G. J. Lematta, C. J. Johnson, E. Holder, L. Huang, S. Bhatti, N. J. Cooke, Team interaction strategies for human–autonomy teaming in next generation combat vehicles, Proceedings of the Human Factors and Ergonomics Society Annual Meeting 64 (2020) 77-81. doi: 10.1177/1071181320641022.
[12] M. Johnson, A. Vera, No AI is an island: the case for teaming intelligence, AI Magazine 40 (2019) 16-28. doi: 10.1609/aimag.v40i1.2842
[13] M. Johnson, M. Vignatti, D. Duran, Understanding human-machine teaming through interdependence analysis, in: M. D. McNeese, E. Salas, M. R. Endsley (Eds.), Contemporary Research, CRC Press, Boca Raton, 2020, pp. 209–233.
[14] D. Şahinel, C. Akpolat, O. C. Görür, F. Sivrikaya, S. Albayrak, Human modeling and interaction in cyber-physical systems: A reference framework, Journal of Manufacturing Systems 59 (2021) 367–385. doi: 10.1016/j.jmsy.2021.03.002
[15] A. Hogan, E. Blomqvist, M. Cochez, C. D'Amato, G. De Melo, C. Gutierrez, S. Kirrane … A. Zimmermann, Knowledge graphs, ACM Computing Surveys 54 (2022) 1-37. doi: 10.1145/3447772.
[16] D. Dell'Aglio, E. Della Valle, F. van Harmelen, A. Bernstein, Stream reasoning: A survey and outlook, Data Science 1 (2017) 59-83. doi: 10.3233/DS-170006.
[17] S. C. Hoi, D. Sahoo, J. Lu, P. Zhao, Online learning: A comprehensive survey, Neurocomputing 459 (2021) 249-289. doi: 10.1016/j.neucom.2021.04.112.