

Empowering Large Language Models in Hybrid Intelligence Systems through Data-Centric Process Models

Carsten Maletzki¹, Eric Rietzke^{1,2}, Ralph Bergmann^{1,3}

¹German Research Center for Artificial Intelligence (DFKI),
Branch University of Trier, Behringstraße 21, 54296 Trier, Germany

²LiveReader GmbH, Zur Imweiler Wies 3, 66649 Oberthal, Germany

³Artificial Intelligence and Intelligent Information Systems,
University of Trier, Behringstraße 21, 54296 Trier, Germany
Carsten.Maletzki@dfki.de

Abstract

Hybrid intelligence systems aim to leverage synergies in closely collaborating teams of humans and artificial intelligence (AI). To guide the realization of such teams, recent research proposed design patterns that capture role-based knowledge on human-AI collaborations. Building on these patterns requires hybrid intelligence systems to provide mechanisms that orchestrate human and AI contributions accordingly. So far, it is unclear if such mechanisms can be provided based on shared representations of the required knowledge. In this regard, we expect ontology-based data-centric process modeling to be a promising direction for hybrid intelligence systems that aim to support knowledge-intensive processes (KiPs). We illustrate this through exemplary process models (realized with our ontology- and data-driven business process model – ODD-BP) that reflect the team design patterns for hybrid intelligence systems. We point out that relying on such process models enables multiple actors to fulfill roles jointly and allows them to address individual shortcomings. This is examined by discussing integrating large language models (LLMs) into the process models and describing how complementary AI actors could help to empower LLMs to fulfill their role in human-AI collaboration more comprehensively. Future work will extend the provided concepts while their evaluation initially focuses on the KiP of medical emergency call handling.

Introduction

Advances in artificial intelligence (AI) recently sparked much discussion about how its capabilities can be integrated into our day-to-day lives. Hybrid intelligence approaches this challenge by aiming to join humans and AI as teammates that leverage synergies, improve from mutual learning, reach sophisticated goals, and act responsibly concerning ethical, legal, and social implications (Akata et al. 2020; Dellermann et al. 2019a). To guide the realization of this vision, recent research proposed team design patterns that could find various applications in cognitive work (van Diggelen and Johnson 2019; van Zoelen et al. 2023). In the area of business process management, such work is an integral part of so-called knowledge-intensive processes (KiPs), which are typically executed by workers who primarily fo-

cus on creating, sharing, and applying knowledge (Davenport 2005; Di Ciccio, Marrella, and Russo 2015; Vaculín et al. 2011). Such knowledge workers are estimated to make up a significant proportion of modern workforces, with vivid examples to be found in the roles of managers, researchers, or doctors (Davenport 2005; Di Ciccio, Marrella, and Russo 2015).

To effectively leverage synergies in human-AI teams, it seems crucial to ground collaborations on clearly defined responsibilities. The team design patterns for hybrid intelligence systems approach this by capturing role-based knowledge on each teammate’s obligations when taking on tasks collaboratively (van Zoelen et al. 2023). Ensuring compliance with these roles can be a fundamental requirement for hybrid intelligence systems. In this context, the orchestration of individual contributions is typically delegated to a collaboration mechanism (Hemmer et al. 2021). So far, limited research has addressed how such mechanisms can be provided to realize the team design pattern for hybrid intelligence systems.

We argue that data-centric process modeling (Rietzke et al. 2021; Vaculín et al. 2011) could be a promising approach to realize role-based collaboration mechanisms to support KiPs. In this context, required knowledge is expressed through a network of data flows between tasks with assigned actors responsible for execution. Since roles in KiPs can have extensive requirements that a single actor might not be able to meet on their own, multiple actors should be allowed to fulfill roles jointly. To this end, we expect that an actor’s shortcomings in their role could be addressed by incorporating complementary contributions by other actors into the process model. Throughout this work, we elaborate on the mentioned aspects by providing exemplary process models based on our ontology- and data-driven business process model (ODD-BP). When discussing how multiple actors can jointly fulfill roles, we will use the example of large language models (LLMs). Our work contributes to current research on hybrid intelligence systems, especially towards applying the team design patterns in KiPs while focusing on LLMs.

In the following sections, we will first elaborate on the fundamentals of hybrid intelligence systems, KiPs, and data-centric process modeling. Afterward, we will briefly introduce the team design patterns for hybrid intelligence sys-

tems and ODD-BP. Then, we will provide exemplary process models to implement the team design patterns with ODD-BP and use the example of LLMs to illustrate how multiple AI actors can be combined to fulfill individual roles in human-AI collaboration more comprehensively. We will close this paper with a summary and an outlook on our future work to expand, implement, and evaluate our approach.

Foundations

Over the recent years, providing human-centric tools that harness AI's increasingly powerful capabilities has become a central point of discussion. Hybrid intelligence systems focus on harnessing AI to augment human intellect and capabilities while avoiding a general substitution (Akata et al. 2020). The overarching rationale in this context is that humans and AI feature complementary skills whose combination inhibits promising synergies (van der Aalst 2021; Akata et al. 2020; Dellermann et al. 2019b). The emerging level of competencies then yields what is witnessed as hybrid intelligence: The ability of humans and AI to achieve sophisticated goals that are out of reach for either humans or AI alone (Akata et al. 2020; Dellermann et al. 2019a). To this end, humans and AI are envisioned to form closely collaborating teams to which they contribute individual skills, learn from shared experiences, and improve over time (Akata et al. 2020; Dellermann et al. 2019a). While humans contribute skills like creativity, empathy, flexibility, and common sense, AI, in contrast, provides its fast, scalable, and consistent analytical capabilities (van der Aalst 2021; Dellermann et al. 2019b). To orchestrate the confluence of these capabilities, hybrid intelligence systems typically rely on collaboration mechanisms that enable task-oriented teamwork (Hemmer et al. 2021). To some extent, the requirements for collaboration mechanisms can be derived from a disparate body of work, among which the team design patterns from van Zoelen et al. can be found (van Zoelen et al. 2023). A pattern generally describes a generic solution for a recurring problem that can be adapted and combined to solve specific issues (Alexander et al. 1977). In this context, the team design patterns for hybrid intelligence systems address the issue of orchestrating human-AI teams and propose a set of roles that impose actors with specific responsibilities when collaborating (van Zoelen et al. 2023). To this end, the knowledge captured by the team design patterns about collaboration is considered common sense (van Diggelen and Johnson 2019). Consequently, collaboration mechanisms must empower human and AI actors with common sense knowledge that allows them to take on defined roles and contribute accordingly. To our knowledge, the team design patterns for hybrid intelligence systems have only selectively been considered so far (e.g., (Gouvêa et al. 2023)), while building on shared representations to orchestrate human-AI collaborations has not yet been discussed.

The general approach of describing human-AI collaboration based on patterns can be traced back to van Diggelen and Johnson, who addressed collaborations in physical and cognitive work (van Diggelen and Johnson 2019). In this context, KiPs provide application scenarios that predominantly focus on cognitive work. KiPs are data-centric

business processes that typically show a strong dependency on the knowledge provided by participants to perform interconnected decision-making tasks (Vaculín et al. 2011). For example, the KiP of medical emergency call handling is essentially described as an iterative procedure where humans derive decisions from mental pictures that originate from knowledge-based assessments of available information (Møller et al. 2021). KiPs can be identified in manifold domains while they typically tend towards unpredictable and emergent executions, which requires flexible support (Dav- enport 2005; Di Ciccio, Marrella, and Russo 2015).

Data-centric process modeling can be used to achieve flexible process-oriented support for KiPs (Rietzke et al. 2021). Compared to widely known control-flow oriented approaches, like BPMN¹, tasks in data-centric process models are described based on their data instead of sequential requirements. Data requirements are typically represented through input and output relations between tasks and data elements. Various approaches consider such data elements as part of artifacts that are generated, processed, and possibly archived throughout their lifetime (e.g., (Cohn and Hull 2009)). Since tasks in data-centric process models get executable when their input data elements are available, they enable flexible process executions driven by known data instead of a specific order of tasks (Rietzke, Bergmann, and Kuhn 2018). Data-centric process models can also be modeled based on semantics defined by ontologies, which aims to reduce semantic inaccuracies (Thomas and Fellmann 2009; Rietzke et al. 2021). In that case, semantically modeled processes can also be structurally adapted through automatic process planning (Heinrich et al. 2008). Regarding KiPs, this allows them to cope with evolving process executions more extensively.

Team Design Patterns for Hybrid Intelligence Systems

This section introduces the team design patterns for hybrid intelligence systems (van Zoelen et al. 2023). They consist of three patterns for human-AI collaboration that describe distinct ways of division of labor towards shared goals and role-based communication obligations. The patterns were derived from workshops between technical and domain experts that discussed human-AI collaboration in emergency response, autonomous animal wildlife monitoring, assembly/maintenance processes, and personalized care.

1. AI Advisor and Human Performer

The *AI Advisor and Human Performer* pattern, depicted in figure 1, directs the leading authority to humans. At the same time, the AI actor aims to recommend appropriate options to guide human actions. Therefore, AI has to utilize its analytical capabilities (0) to provide a human actor with appropriate options (1). The human actor then has to assess the options and choose the most suitable (2). After the decision, the human corresponds with feedback on the helpfulness of provided options (3), which, in turn, the AI uses to improve its performance over time.

¹<https://www.omg.org/spec/BPMN>




Figure 1: AI Advisor and Human Performer Pattern. Based on (van Zoelen et al. 2023)

2. AI Performer and Human Assistant

The *AI Performer and Human Assistant* pattern (figure 2) gives an AI actor full autonomy when performing a task and making decisions (0) and only if the AI actor encounters a limitation (1) requires it to interact with humans by requesting assistance (2). With their diverse skills, humans interpret the situation to perform the task for assistance (3) and then return responsibilities (4).




Figure 2: AI Performer and Human Assistant Pattern. Based on (van Zoelen et al. 2023)

3. AI Performer and Human Validator

The *AI Performer and Human Validator* pattern (figure 3) also allows an AI actor to perform a task and make decisions autonomously (0). Still, the AI actor must stay within the guardrails of human supervision. To this end, AI must communicate relevant information about its actions to human supervisors (1) who perform a validation (2). If humans identify the need for an intervention, they respond with corresponding feedback (3).




Figure 3: AI Performer and Human Validator Pattern. Based on (van Zoelen et al. 2023)

ODD-BP – Ontology and Data-Driven Business Process Model

ODD-BP is an approach to data-centric process modeling designed to address the needs of KiPs in terms of flexibility and was further motivated by fostering the division of labor between humans and AI (Rietzke, Bergmann, and

Kuhn 2019). Applying ODD-BP leads to a unified semantic knowledge base that enables an organization to manage its processes, data, and actors holistically. To this end, the ODD-BP metamodel (figure 4) fundamentally divides tasks into *user tasks* and *system tasks*. While user tasks have to be performed manually by human actors, system tasks are executed automatically by calling the respective AI actor. The data-centric perspective on process models in ODD-BP results from linking these tasks to data elements that either mark their prerequisites for execution or the resulting output (input: *required_by*; output: *delivers*). Data elements in this context can be represented through *dataobjects* and *attributes*. Dataobjects describe entities whose attributes are processed by tasks. To add precise semantics to these data elements that can be understood by humans and AI equally, they are declared as instances of domain-specific classes in an ontology. By following this approach, unambiguous descriptions are possible, for example, describing that a person (represented by a dataobject) has a name (stored as a value of the respective attribute).

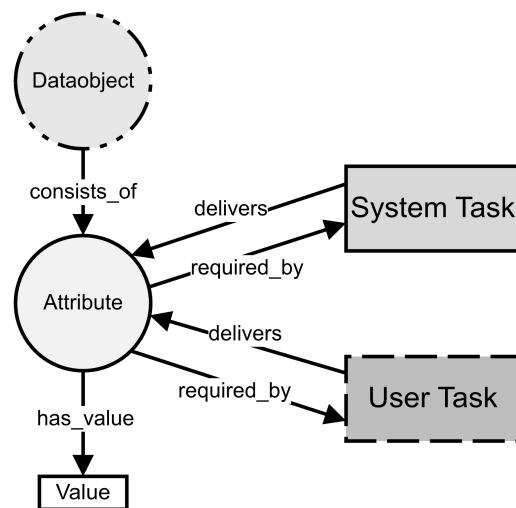


Figure 4: Excerpt of the ODD-BP Metamodel. Based on (Rietzke et al. 2021)

Modeling Team Design Patterns with ODD-BP

This section provides exemplary process models to realize the team design patterns for hybrid intelligence systems with ODD-BP. It also discusses integrating LLMs into the process models and points out that LLMs elicit shortcomings that could prevent them from solely fulfilling roles in human-AI collaboration for KiPs. However, these shortcomings could be addressed by integrating complementary AI actors into the process models.

AI Advisor and Human Performer

The team design pattern *AI Advisor and Human Performer* empowers an AI actor to provide decision-making options to a human actor. Conversely, it enables a human to decide on this basis and provide feedback on the options' quality. The

ODD-BP process model shown in figure 5 realizes this pattern. To benefit the clarity of provided visualizations, all figures in this section only depict tasks and attributes – the visualization of corresponding dataobjects and values is omitted. Further, labels of connecting arrows are omitted, as the direction of the arrow already implies the type used. The exemplary process model shown in figure 5 utilizes a system task to integrate an AI actor responsible for analyzing available process data (represented by a single symbolic attribute) and deriving a set of decision-making options. While the process data is modeled as input of this task, resulting decision options are its output. To enable a human actor to consider identified options during decision-making, they are further set as input for a corresponding user task. In this context, a human actor is also granted access to the process data that led to the decision options. This access might also be extended to a broader range of process data to consider during decision-making. This should enable human actors to identify responsible actions by carefully weighing available alternatives.

Suppose the human actor identifies a given option suitable for the current situation. In that case, this option is returned as a decision output of the user task. Further output of this user task can be human feedback on the options' quality, which lays the foundation to improve the AI actor. As a foundation for learning from feedback, the process model contains a system task that initiates a training procedure that revises the model used by the AI actor to generate decision options. To this end, provided options, feedback, and the underlying process data are set as input for this task.

LLMs could provide versatile decision support for KiPs

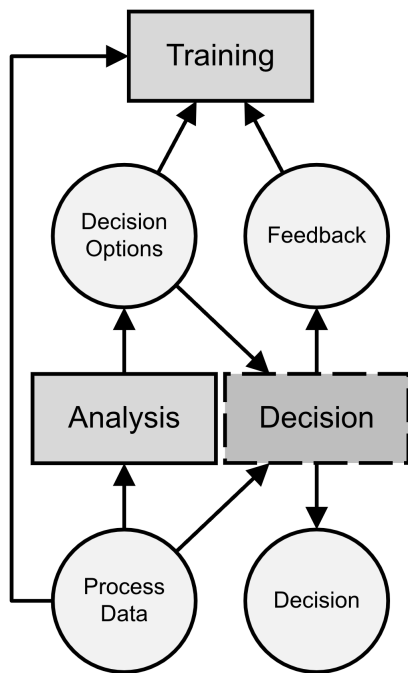


Figure 5: Exemplary Process Model to Realize the AI Advisor and Human Performer Pattern

in typical downstream tasks that involve decision-making based on natural language text (e.g., sentiment analysis). To this end, figure 6 shows an integration of an LLM into the exemplary process model from figure 5. To learn from feedback, an LLM could either be retrained entirely, fine-tuned, or approached with an altered prompt, which is called in-context learning (Brown et al. 2020). One way to effectively realize in-context learning is few-shot prompting (Brown et al. 2020). Few-shot prompting enables LLMs to learn with significantly lower training examples than completely retraining or fine-tuning a model (Brown et al. 2020). To this end, few-shot prompting only requires adding a set of training examples to a prompt that illustrates how to solve a given type of task. At the same time, various example selection strategies (e.g., random or similarity-based) can be applied to influence the LLMs performance (Brown et al. 2020; Liu et al. 2022). A prerequisite to enabling few-shot approaches in ODD-BP process models is that feedback from human actors in different situations can be composed into assessable training sets. Since ODD-BP manages processes in a uniform semantic knowledge base, every decision option and feedback can be queried with a query language like SPARQL² (Rietzke et al. 2021). To this end, figure 6 depicts a system task that performs the required query to compose the training set from which it selects few-shot examples. In

²<https://www.w3.org/TR/sparql11-query/>




Figure 6: Integration of an LLM and Few-shot Prompting

this context, the system task implements the appropriate selection strategy concerning the type of analysis task to solve. If the selection strategy is similarity-based, this task requires access to the process data to be analyzed by the LLM and for which suitable examples are to be found. Note that provided decision options and feedback are not defined as input of the example selection task, as they are not available at the time of example selection and, therefore, would prevent an execution. Identified few-shot examples are subsequently passed to the system task that integrates the LLM to analyze the process data. In this context, the few-shot examples are added to the prompt before calling the LLM.

AI Performer and Human Assistant

In ODD-BP process models, AI actors can be integrated to support KiPs by analyzing input data and deriving required outputs. In this respect, AI actors may fail to provide the required outputs. Possible reasons for this can be that AI actors lack skills or that the quality of available data does not suffice to derive outputs. The *AI Performer and Human Assistant* pattern addresses such limitations by referring to humans and asking them to close such gaps. In doing so, an AI actor transfers its authority to decide to a human actor. Figure 7 shows an exemplary process model to realize the *AI Performer and Human Assistant* pattern. In this context, it orchestrates the human-AI collaboration and the hand-over of responsibilities by incorporating a system task to analyze process data and output a decision or identified limitation. A limitation triggers a user task that asks a human actor to assist by assessing available process data and derive a suitable decision instead. To this end, the assistance user task might take different process data as input than the original system task. This can be reasonable if, for example, the model of the AI actor was built to assess specific data that suffices for decision-making in most cases. In con-

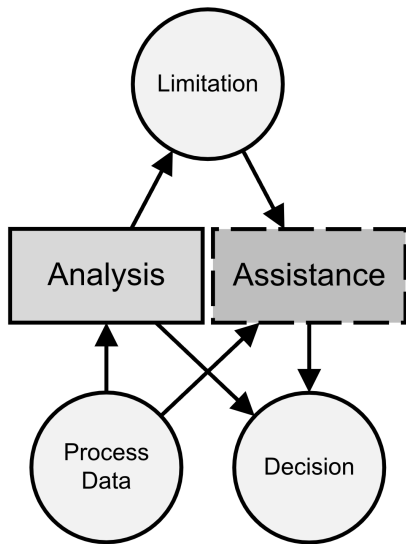


Figure 7: Exemplary Process Model to Realize the AI Performer and Human Assistant Pattern

trast, isolated cases might be more complex and require extensive individual treatment. This can be portrayed by considering integrating an LLM into the process model, which we primarily expect to support KiPs in processing natural language text. If decisions cannot be derived from available text, it may be because the text lacks adequate information or the LLM lacks the required skills. Either way, if the LLM identifies a limit, the exemplary process model to realize the *AI Performer and Human Assistant* pattern helps to activate a human actor for assistance. This human actor might then consider further data for which the LLM’s underlying method might not be ideal. Although a human actor might be able to assist the LLM, it would be preferable if the system would learn from observing human problem-solving. In this context, case-based reasoning (CBR, e.g., (Aamodt and Plaza 1994; Bergmann et al. 2021)) could help to solve cases based on occasional human assistance while further being able to consider a broad spectrum of data types. Fundamental to CBR is a collection of recorded cases and their respective solutions (so-called case base). To identify a solution for a current situation, CBR initially determines the extent to which a current situation is similar to recorded cases. Afterward, a solution to the current situation is derived from a known solution to a similar situation. Figure 8 shows how CBR could augment LLMs by incorporating a system task triggered by a limitation report. Similar to the previous exemplary process model for implementing few-shot prompting, this system task relies on a query to collect process data. In this context, the query retrieves recent cases

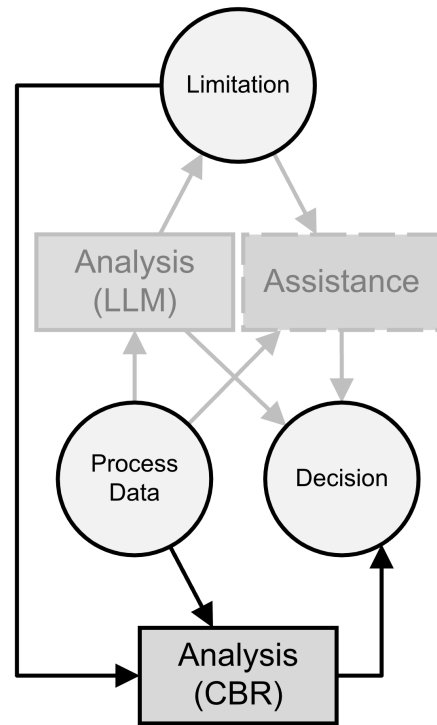


Figure 8: Integration of an LLM Augmented with Case-based Reasoning

and their respective solutions. Based on a similarity calculation to currently available process data, a decision is derived and returned as an output of the system task. Consequently, this can substitute human interventions in situations similar to those already solved.

AI Performer and Human Validator

The *AI Performer and Human Validator* pattern puts AI contributions under human supervision while considering human feedback to improve over time. Figure 9 illustrates how an ODD-BP process model can represent this pattern. Similar to the previously provided process models, it incorporates a system task integrating an AI actor to analyze relevant process data as a foundation for autonomous decision-making. The decision made in this context is the input of a user task requiring a human actor to validate the AI's decision. In this context, the human is granted access to process data relevant to evaluate the decision's correctness. The output of this user task represents the validation result, possibly correcting the AI actor's prior result. A system task is integrated for initiating a model training to learn from this feedback, similar to the exemplary process model to realize the *AI Advisor and Human Performer* pattern. As a foundation for this training procedure, this task is provided with relevant process data and the validation results.

To facilitate human validations, the *AI Performer and Human Validator* pattern specifies that AI should provide humans with appropriate information (van Zoelen et al. 2023). Suppose an LLM is integrated into the process model. In that case, decisions may be affected by so-called hallucinations. These can be plausible but not necessarily correct results that

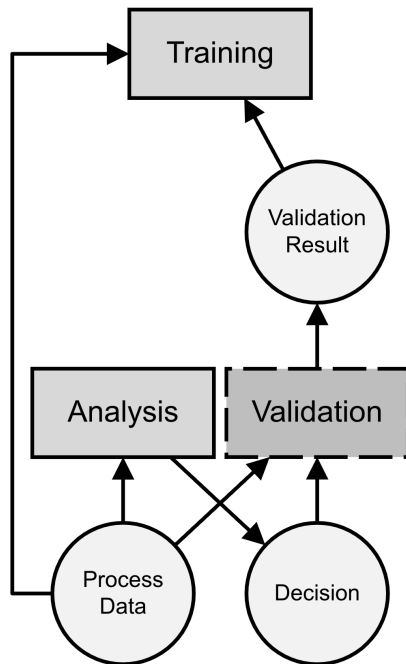


Figure 9: Exemplary Process Model to Realize the AI Performer and Human Validator Pattern

can deviate from established world knowledge (Zhang et al. 2023; Lenat and Marcus 2023). To facilitate human validations, it would be helpful if questionable results were recognized automatically in advance and marked accordingly. In this context, Lenat and Marcus consider combining logic-based AI with LLMs to identify LLM-generated results that have no logical underpinning (Lenat and Marcus 2023). Figure 10 provides the fundament for such approaches by adding a system task that integrates a pre-validation of an LLM's decision in the context of currently available process data. Since the validation user task takes the output of this task as input, a human can incorporate this additional information during validation.

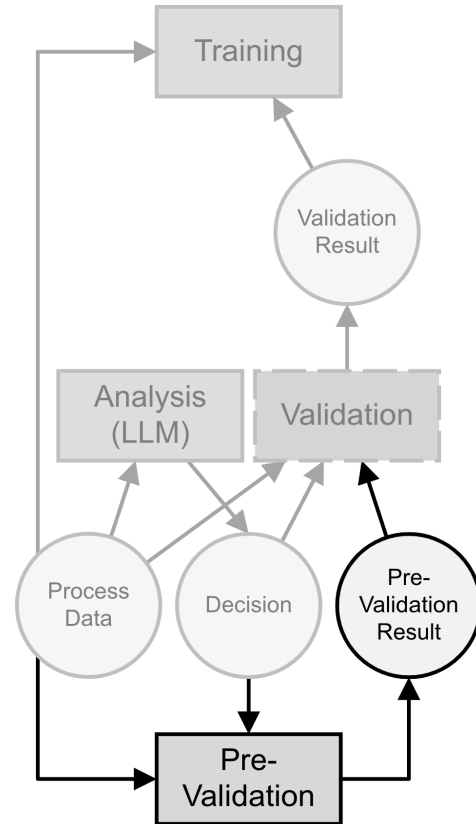


Figure 10: Integration of an LLM and a Logic-based Pre-Validation

Conclusion

This paper has shown that the team design patterns for hybrid intelligence systems can be realized with ODD-BP, resulting in shared representations of the required knowledge to orchestrate human-AI collaborations. Through exemplary integrations of LLMs in such data-centric process models, we have illustrated that the weaknesses of an AI actor regarding its role in human-AI collaboration could be addressed by integrating further complementary AI actors. Here, integrated AI actors might help LLMs adapt to feedback effectively, overcome limitations, and provide additional information to facilitate human supervision. To this

end, integrated AI actors provide domain and common-sense knowledge that empowers the LLM to more extensively fulfill the roles defined by the team design patterns.

Limitations and Future Work

This work aimed to indicate whether ontology-based data-centric process modeling (using the example of ODD-BP) can serve as a basis to realize human-AI collaborations in KiPs as envisioned by the team design patterns for hybrid intelligence. Since the proposed process models were not implemented and evaluated, the results are limited in their significance. Future work will address this limitation by implementing exemplary process models for various use cases. For this purpose, we plan to build on our tool Notitia³, whereby, as a first application scenario, we will focus on emergency call handling and progress as recently described (Maletzki, Elsenbast, and Reuter-Oppermann 2024). Since this work solely regarded exemplary integrations of LLMs, it is still unclear to what extent augmentations of other methods can be realized and whether this could be used to expand the team design patterns. Future work will address this gap and investigate the extent to which overarching patterns that are suitable for reuse can be identified.

Acknowledgements

This work is funded by the Federal Ministry for Economic Affairs and Climate Action under grant No. M01MK21005A.

The authors used GPT-4 to refine their writing. To this end, the model was prompted with original sentences and asked for alternative formulations. These alternatives were taken as an inspiration to improve original expressions.

References

- Aamodt, A.; and Plaza, E. 1994. Case-Based Reasoning: Foundational Issues, Methodological Variations, and System Approaches. *AI Communications*, 7(1): 39–59.
- Akata, Z.; Balliet, D.; de Rijke, M.; Dignum, F.; Dignum, V.; Eiben, G.; Fokkens, A.; Grossi, D.; Hindriks, K.; Hoos, H.; Hung, H.; Jonker, C.; Monz, C.; Neerincx, M.; Oliehoek, F.; Prakken, H.; Schlobach, S.; van der Gaag, L.; van Harmelen, F.; van Hoof, H.; van Riemsdijk, B.; van Wylsberghe, A.; Verbrugge, R.; Verheij, B.; Vossen, P.; and Welling, M. 2020. A Research Agenda for Hybrid Intelligence: Augmenting Human Intellect With Collaborative, Adaptive, Responsible, and Explainable Artificial Intelligence. *Computer*, 53(8): 18–28.
- Alexander, C.; Ishikawa, S.; Silverstein, M.; Jacobson, M.; Fiksdahl-King, I.; and Angel, S. 1977. *A Pattern Language: Towns, Buildings, Construction*. Oxford University Press.
- Bergmann, R.; Minor, M.; Bach, K.; Althoff, K.-D.; and Muñoz-Avila, H. 2021. Fallbasiertes Schließen. In Görz, G.; Schmid, U.; and Braun, T., eds., *Handbuch der Künstlichen Intelligenz*, volume 6, 343–394. Berlin, Boston: De Gruyter.
- Brown, T. B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D. M.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; and Amodei, D. 2020. Language Models Are Few-Shot Learners. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M. F.; Lin, H., ed., *Proceedings of the 34th International Conference on Neural Information Processing Systems*, volume 3 of *NIPS’20*, 1877–1901. Curran Associates Inc.
- Cohn, D.; and Hull, R. 2009. Business Artifacts: A Data-centric Approach to Modeling Business Operations and Processes. *IEEE Data Engineering Bulletin*, 32(3): 3–9.
- Davenport, T. H. 2005. *Thinking for a Living: How to Get Better Performance and Results from Knowledge Workers*. Harvard Business School Press.
- Dellermann, D.; Calma, A.; Lipusch, N.; Weber, T.; Weigel, S.; and Ebel, P. 2019a. The Future of Human-AI Collaboration: A Taxonomy of Design Knowledge for Hybrid Intelligence Systems. In Bui, T. X., ed., *Proceedings of the 52nd Hawaii International Conference on System Sciences*, 274–283. ScholarSpace.
- Dellermann, D.; Ebel, P.; Söllner, M.; and Leimeister, J. M. 2019b. Hybrid Intelligence. *Business & Information Systems Engineering*, 61: 637–643.
- Di Ciccio, C.; Marrella, A.; and Russo, A. 2015. Knowledge-Intensive Processes: Characteristics, Requirements and Analysis of Contemporary Approaches. *Journal on Data Semantics*, 4: 29–57.
- Gouvêa, T. S.; Kath, H.; Troshani, I.; Lüers, B.; Serafini, P. P.; Campos, I. B.; Afonso, A. S.; Leandro, S. M. F. M.; Swanepoel, L.; Theron, N.; Swemmer, A. M.; and Sonntag, D. 2023. Interactive Machine Learning Solutions for Acoustic Monitoring of Animal Wildlife in Biosphere Reserves. In Elkind, E., ed., *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, 6405–6413. International Joint Conferences on Artificial Intelligence Organization.
- Heinrich, B.; Bewernik, M.-A.; Henneberger, M.; Krammer, A.; and Lautenbacher, F. 2008. SEMPA – A Semantic Business Process Management Approach for the Planning of Process Models. *Wirtschaftsinformatik*, 50: 445–460.
- Hemmer, P.; Schemmer, M.; Vössing, M.; and Kühl, N. 2021. Human-AI Complementarity in Hybrid Intelligence Systems: A Structured Literature Review. In Vogel, D.; Shen, K. N.; Ling, P. S.; Ravishankar, M. N.; and Zhang, X. J., eds., *25th Pacific Asia Conference on Information Systems, PACIS 2021*.
- Lenat, D.; and Marcus, G. 2023. Getting from Generative AI to Trustworthy AI: What LLMs might learn from Cyc. arXiv:2308.04445.
- Liu, J.; Shen, D.; Zhang, Y.; Dolan, B.; Carin, L.; and Chen, W. 2022. What Makes Good In-Context Examples for GPT-3? In Agirre, E.; Apidianaki, M.; and Vulić, I., eds., *Proceedings of Deep Learning Inside Out (DeeLIO 2022)*:

³<https://notitia.world/>

- The 3rd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, 100–114. Association for Computational Linguistics.
- Maletzki, C.; Elsenbast, C.; and Reuter-Oppermann, M. 2024. Towards Human-AI Interaction in Medical Emergency Call Handling. In Bui, T. X., ed., *57th Hawaii International Conference on System Sciences*, 3374–3383. ScholarSpace.
- Møller, T. P.; Jensen, H. G.; Viereck, S.; Lippert, F.; and Østergaard, D. 2021. Medical dispatchers' perception of the interaction with the caller during emergency calls - a qualitative study. *Scandinavian Journal of Trauma, Resuscitation and Emergency Medicine*, 29: 45.
- Rietzke, E.; Bergmann, R.; and Kuhn, N. 2018. Semantically-Oriented Business Process Visualization for a Data and Constraint-Based Workflow Approach. In Teniente, E.; and Weidlich, M., eds., *Business Process Management Workshops*, 142–150. Springer International Publishing.
- Rietzke, E.; Bergmann, R.; and Kuhn, N. 2019. ODD-BP - an Ontology- and Data-Driven Business Process Model. In Jäschke, R.; and Weidlich, M., eds., *Proceedings of the Conference on "Lernen, Wissen, Daten, Analysen"*, volume 2454 of *CEUR Workshop Proceedings*, 310–321. CEUR-WS.org.
- Rietzke, E.; Maletzki, C.; Bergmann, R.; and Kuhn, N. 2021. Execution of Knowledge-Intensive Processes by Utilizing Ontology-Based Reasoning: ODD-BP: An Ontology- and Data-Driven Business Process Model. *Journal on Data Semantics*, 10: 3–18.
- Thomas, O.; and Fellmann, M. 2009. Semantic Process Modeling – Design and Implementation of an Ontology-based Representation of Business Processes. *Business & Information Systems Engineering*, 1: 438–451.
- Vaculín, R.; Hull, R.; Heath, T.; Cochran, C.; Nigam, A.; and Sukaviriya, P. 2011. Declarative business artifact centric modeling of decision and knowledge intensive business processes. In *IEEE 15th International Enterprise Distributed Object Computing Conference*, 151–160. IEEE Computer Society Conference Publishing Service.
- van der Aalst, W. M. P. 2021. Hybrid Intelligence: to automate or not to automate, that is the question. *International Journal of Information Systems and Project Management*, 9(2): 5–20.
- van Diggelen, J.; and Johnson, M. 2019. Team Design Patterns. In Oka, N.; Koda, T.; Obaid, M.; Nakanishi, H.; Mubin, O.; and Tanaka, K., eds., *Proceedings of the 7th International Conference on Human-Agent Interaction, HAI '19*, 118–126. Association for Computing Machinery.
- van Zoelen, E.; Mioch, T.; Tajaddini, M.; Fleiner, C.; Tsaneva, S.; Camin, P.; Gouvêa, T. S.; Baraka, K.; de Boer, M. H. T.; and Neerincx, M. A. 2023. Developing Team Design Patterns for Hybrid Intelligence Systems. In Lukowicz, P.; Mayer, S.; Koch, J.; Shawe-Taylor, J.; and Tiddi, I., eds., *Proceedings of the Second International Conference on Hybrid Human-Artificial Intelligence*, 3–16. IOS Press.
- Zhang, Y.; Li, Y.; Cui, L.; Cai, D.; Liu, L.; Fu, T.; Huang, X.; Zhao, E.; Zhang, Y.; Chen, Y.; Wang, L.; Luu, A. T.; Bi, W.; Shi, F.; and Shi, S. 2023. Siren's Song in the AI Ocean: A Survey on Hallucination in Large Language Models. *arXiv:2309.01219*.