Qeios

Research Article

Collaborative Intelligence: A scoping review of current applications

Emma Schleiger[1], Claire Mason[1], Claire Naughtin[1], Andrew Reeson[1], Cecile Paris[1]

1. The Commonwealth Scientific and Industrial Research Organisation, Canberra, Australia

Humans and artificial intelligence (AI) systems have complementary strengths. This complementarity creates the potential to improve performance by combining inputs from human and AI on a common task or goal. A systematic review of academic and grey literature was conducted to investigate whether real-world examples of 'collaborative intelligence' could be identified. Applications utilising collaborative intelligence had to have (1) complementarity (i.e., the collaboration improves performance beyond that which could be achieved by the human or the AI alone), (2) a shared objective and outcome, and (3) sustained, task-related interaction between human and AI. The literature search yielded 1250 documents but only 16 applications meeting these criteria were identified. Most collaborative applications were at the prototyping stage but they could perform a variety of types of work (creative, industrial, healthcare, emergency response and knowledge work) and delivered a range of benefits (efficiency, creativity and safety). In most applications the AI had a virtual presence but there were examples where the AI also had a cyber-physical form. These early applications reveal the potential for collaborative intelligence to be applied in a range of domains and formats.

Corresponding author: Emma Schleiger, emma.schleiger@data61.csiro.au

Introduction

Industry 4.0 is underway with organizations adopting automated systems, building their internet of things and using big data, smart systems, and cyber-physical systems to expand their capabilities and improve productivity (Mason, Ayre, and Burns 2022, Schuh et al. 2020, Szász et al. 2021, Veile et al. 2020). The next wave of innovation, known as Industry 5.0, focuses on human-centric technology development and adoption (European Commission et al. 2021, Maddikunta et al. 2022). Rather than

simply automating tasks that were previously performed by humans, there is greater focus on using collaboration between humans and smart machines to augment the capability of humans (Sindhwani et al. 2022).

The term collaborative intelligence refers to collaborative human-AI systems that leverage the different attributes and strengths of each agent to achieve further improvements in work outcomes (Daugherty and Wilson 2018, Seeber et al. 2020, Billman et al. 2006, Jarrahi 2018). Organizations investing in AI-human collaboration are expected to boost revenues by 38% within five years (Shook and Knickrehm 2018). While the literature provides a strong rationale for utilising collaborative intelligence (Traumer, Oeste-Reiß, and Leimeister 2017, Dellermann et al. 2019, Madni and Madni 2018), the construct has not been tested through systematic empirical research. A systematic review of AI applications is needed to determine whether any collaborative intelligence applications have been developed and, if so, how they are being used and what benefits they provide. In this paper, we develop a set of criteria for identifying applications that utilize collaborative intelligence and use them to systematically analyse potential collaborative intelligence applications reported in academic and grey literature. Our goal is to test whether collaborative intelligence exists as more than a theoretical construct and, if so, to describe the types of collaborative intelligence that are currently technologically and economically feasible.

Why Collaborative Intelligence?

Examples of more collaborative human-AI interactions (Poser and Bittner 2020, Seeber et al. 2020, Epstein 2015, Kolbeinsson, Lagerstedt, and Lindblom 2019) are described using a range of terms including human-robot interaction (Cesta, Orlandini, and Umbrico 2018), human/machine in the loop (Ostheimer, Chowdhury, and Iqbal 2021), hybrid intelligence (Ostheimer, Chowdhury, and Iqbal 2021) and collective intelligence (Dellermann et al. 2019). There is currently little consensus across the literature regarding how to differentiate the range of ways in which humans and AI can collaborate and when such interactions represent "collaborative intelligence". In this study we use descriptions of collaborative intelligence suggested by other researchers to draw out three defining characteristics of collaborative intelligence. First, the collaboration involves a sequence of shared actions between human and AI agents towards a shared objective (Kolbeinsson, Lagerstedt, and Lindblom 2019, Cienki 2015, Wang et al. 2020). Second, to enable this level of interaction the AI agents must have the ability to share and respond to information about the task and adapt to changes in the state of the human

agent and the task (Kolbeinsson, Lagerstedt, and Lindblom 2019, Wang et al. 2020). Finally, the collaboration between the human and AI agents improves the performance, novelty, productivity, or quality of work above what could be done individually (Dellermann et al. 2019, Madni and Madni 2018). Together, these descriptions provide three criteria for identifying AI systems that enable collaborative intelligence:

- Complementarity: The goal of the interaction between human and AI agents is to leverage their unique strengths to achieve improved outcomes that could not be achieved by either agent individually This excludes human-AI interactions that use the human to teach the AI so that in the long run the AI can perform the task independently. It also excludes applications that are designed to probe or test the dynamics of collaboration or teamwork rather than to complete a task (Gao et al. 2021).

- Shared objective: The human and AI agents are focused on the same objective and the activities of the human and AI agents are integrated and indivisible in the final output that is produced (Dellermann et al. 2019, Dubey et al. 2020, Johnson et al. 2014). The workflow must go beyond a simple division of labour or a transactional relationship.

- Sustained interaction: Interaction between the human and AI agents must extend beyond a static interaction such as a single question/answer dynamic (Traumer, Oeste-Reiß, and Leimeister 2017). Reciprocal communication which enables each agent to understand changes in the state of the objective or the other agent and respond adaptively is critical for all collaborations and is a key feature of collaborative intelligence (Madni and Madni 2018, McDermott et al. 2018).

The motivation for developing collaborative intelligence applications (as opposed to AI applications with collaborative capability) has two sources. First, whilst the capability of AI has been improving rapidly, there are still many tasks that AI cannot perform despite these being simple tasks for a human. The strength of AI lies in its computational power and its ability to process very large amounts of data, recognize patterns and evaluate alternative decision options (Jarrahi 2018, Agrawal, Gans, and Goldfarb 2019). However, AI struggles to understand common-sense situations (Jarrahi 2018), make intuitive decisions or judgements based on indescribable factors (Jarrahi 2018, Agrawal, Gans, and Goldfarb 2019, Goldfarb and Lindsay 2020) and respond to novel situations (Jarrahi 2018) – tasks that a human can perform well on. These complementary strengths suggest that there are likely to be many fields in which performance can be optimized by using a combination of human intelligence and AI (De Luca 2021).

The strength of collaborative intelligence is demonstrated in the evolving use of human-computer chess teams that combine human intuition and computational power (Kasparov 2010). When IBM's Deep Blue program defeated world chess champion Gary Kasparov in 1997, the ability to process hundreds of millions of moves per second outperformed human creativity and imagination (Kasparov 2010). Less than a decade later, in 2005, in a chess tournament of human-computer teams, two amateur players won against teams of grandmasters and the most powerful computer programs. The amateurs developed a superior process to leverage the most value from their computers demonstrating "weak human + machine + better process was superior to a strong computer alone and...superior to a strong human + machine + inferior process" (Kasparov 2010, Thompson 2013).

The second argument for utilising collaborative intelligence lies in the potential to improve the quality and scope of work for humans. Some types of work are inherently more motivating than others (Hackman and Oldham 1975). By automating functions that are less motivating for humans but allowing the human to add value by performing more rewarding tasks we can improve the quality of work. While it is not within the scope of this paper to discuss the optimal design of collaborative work involving both humans and AI, this important topic is already being explored (Parker and Grote 2019). Current research suggests that the experience of using AI can affect a human worker's experience of predictability, controllability, meaningfulness, and fairness (Parker and Grote 2019, Battina 2018, Langer and Landers 2021, Oh et al. 2018). In addition, by allowing human capability to be augmented by AI capability, there is potential to address skills gaps within the existing workforce or increase the pool of workers who can perform a given role.

Research objectives

Our first objective was to determine whether there are current applications that met our definition of collaborative intelligence. If examples of collaborative intelligence applications were identified, the second objective was to characterize the way in which collaborative intelligence applications are developing, describing (1) what types of tasks they perform, (2) the roles of the human and the AI agents, (3) how the human and AI agents interact and (4) the benefits achieved through the collaboration.

Methodology

A systematic methodology was adopted for the literature search, using the Preferred Reporting Items for Systematic Reviews and Meta-Analyses guidelines (Liberati et al. 2009, Arksey and O'Malley 2005). Scopus, ProQuest, Web of Science, and IEEE Xplore databases were used for the literature search because they contain a mix of academic and grey literature sources across a broad range of topics. Additional directed internet searches were carried out using the Google search engine in private browser mode to avoid the impact of cookies and previous searches on the returned results. In addition, where secondary sources of potential applications were revealed during the review of a full-text article, the secondary reference was also reviewed.

Keywords were identified and selected through an iterative approach including pilot searches and reviewing key documents. These keywords were "human" and "artificial intelligence collaboration"; "hybrid intelligence" and "artificial intelligence"; "collective intelligence" and "artificial intelligence"; "human computer collaboration" and "artificial intelligence"; "hybrid teamwork" and "artificial intelligence"; "cobot" and "artificial intelligence"; "human machine collaboration" and "artificial intelligence"; "work" and "artificial intelligence". Proximity and wildcard syntax were used where databases allowed, along with commonly used acronyms and synonyms for artificial intelligence (e.g., AI, machine learning, ML). The results of the searches were limited to English language documents published between January 1 2012 and 31 December 2021.

The keywords were simplified for the Google search protocol to accommodate the broader range of potentially relevant documents and to avoid being overly restrictive or returning results that replicated those done in the database search. The following key terms were used: human machine collaboration artificial intelligence machine learning; hybrid intelligence; collective intelligence artificial intelligence machine learning; human computer collaboration; hybrid teamwork artificial intelligence machine learning; cobot artificial intelligence machine learning; human artificial intelligence machine learning collaboration. Restrictive limiters such as proximity searches and quotation marks were also removed. The Google relevancy ranking was relied on to identify the most relevant documents, and the first 30 items returned of each of the seven targeted google searches were analysed for applications of collaborative intelligence.

Criteria

Literature Review

The database search returned 1,250 documents. The titles and executive summaries/abstracts (where available) of these documents were initially screened to identify those that were eligible for inclusion for full-text review. To be included there needed to be an indication or reference to an AI application that involved tasks completed by both AI and human. Of the 1,250 documents, 335 were eliminated as duplicate documents and a further 445 were excluded after the initial screening process. The remaining 470 documents were assessed for inclusion in the systematic review through a detailed full-text review. Each document was assessed against the following eligibility requirements:

- Complementarity: The collaboration between the human and AI agents improves the performance, novelty, productivity, or quality of work above what could be done individually. The outcome of the collaborative task in a collaborative intelligence application achieves a better result than either the human or the AI would alone.
- Shared objective and output: The human and AI agents are focused on the same objective and the final output represents an integration of their individual contributions.
- Sustained period of interaction: Interaction between the human and artificial intelligence must occur over time, rather than via a singular or static interaction.

From an analysis of the full texts, additional records were excluded because either (1) the full-text article could not be sourced after extensive searching (n=3), or (2) the document did not meet one or more of the inclusion criteria (n=448). A total of 451 documents were eliminated from further analysis based on these full-text exclusion criteria, with the remaining 19 documents detailing 16 collaborative intelligence applications (3 collaborative intelligence applications were detailed in two unique documents) were included in the systematic review for further analysis (see Figure 1).
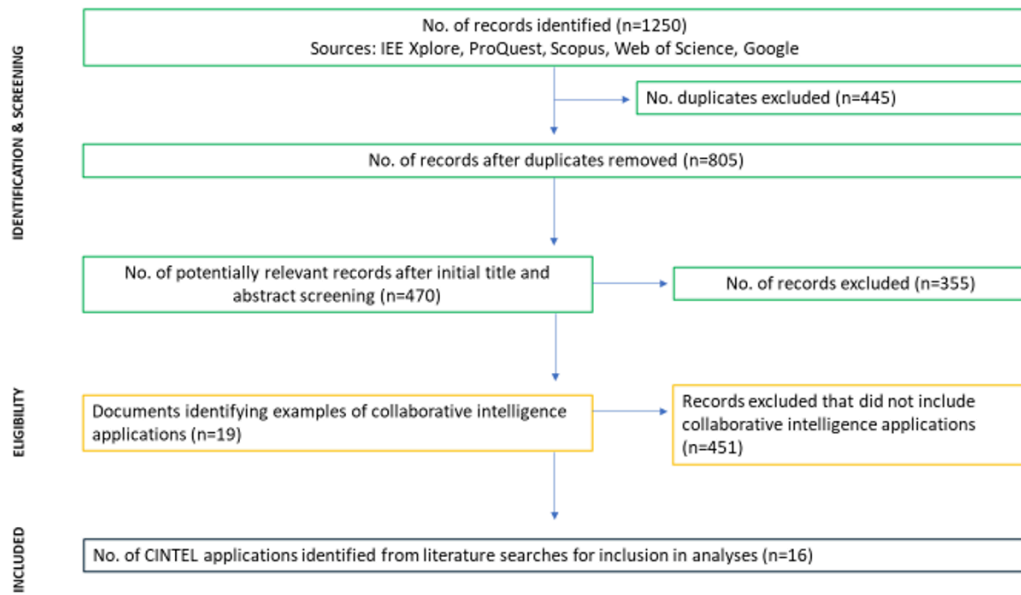
Figure 1. Results of the literature search.

Results

An extensive search (encompassing 1,250 documents from the academic and grey literature) revealed 16 examples of AI applications that met the criteria for collaborative intelligence (see Table 1). Analysis of these applications revealed that they were designed to perform various types of work. We identified five types of collaborative intelligence applications: creative agents, industrial agents, healthcare agents, emergency services agents and knowledge work agents.

Type	Description	Collaborative process	Details
Creative agents	**Evolver** (Feldman 2017) is designed to collaboratively produce generative graphic design artifacts based on constraints controlled by a human graphic designer in an iterative design process. Evolver enhances human creativity during the production of artworks and graphic designs.	The human agent provides Evolver with a design brief. Based on the seed material and design constraints, Evolver begins generating design outputs. The design outputs can be modified and edited by the human agent and fed back to Evolver to create alternative designs based on updated and specified features.	**Interface:** GUI **Stage of Development:** Prototype **Country of Development:** Canada **Sector:** Academia/research
	Flow Machines Professional (Avdeeff 2019) is an AI-assisted music composing system. The goal of Flow Machines is to enhance the creativity of humans in the production and development of music.	The human agent provides Flow-Machines with a set of scores that they want to take inspiration from. In response, Flow machines suggests melodies, accompaniments, and instrumentation. The process is iterative, where human agents can provide feedback to Flow Machines and provide additional desired parameters until they reach a satisfactory result.	**Interface:** GUI **Stage of Development:** In commercial use **Country of Development:** Japan **Sector:** Private organization (Sony)
	StoryDrawer (Zhang et al. 2021) is an intelligent agent that supports children aged 5-10 years old in creating oral stories through collaborative drawing. This application aims to enhance the creativity of the fictional stories and provide an educational experience for children learning to	StoryDrawer produces a random story element in the form of a doodle on the screen, which the child can add. Based on the doodle, the child begins orally telling a story. StoryDrawer analyses the oral description and extracts keywords to generate additional illustrations in real-time. If new material is desired by the child during storytelling, they can create a drawing and StoryDrawer will generate a sketch based on their drawings, or the child can generate a random doodle for	**Interface:** GUI **Stage of Development:** Prototype **Country of Development:** China **Sector:** Academia/research

Type	Description	Collaborative process	Details
	develop creative writing skills.	inspiration. The child can provide feedback to StoryDrawer; if the feedback is negative, StoryDrawer will provide another 'idea' in the form of a doodle.	
	An unnamed robot dancer application performs classical dancing patterns in harmony with the artistic moods expressed by a human pianist (Saffiotti et al. 2020, Thörn, Knudsen, and Saffiotti 2020). The purpose of the collaboration between the robot and human dancers is to spark inspiration and creativity in the human co-performers and enhance the audience's experience.	The robotic AI system analyses the artistic expression of the human dancer and dynamically responds with dancing/gestures to match this expression. The AI system may be creative and proactive in setting parameters. In response to how the robot performs, the human dancer may adjust their moves, resulting in a harmonious joint performance between the human and the robotic agent.	**Interface:** Cyber-physical system **Stage of Development:** Prototype **Country of Development:** Sweden **Sector:** Academia/research
	Shelley (Yanardag, Cebrian, and Rahwan 2021, O'Brien 2017) is an AI bot that collaboratively creates short horror stories with human users on Twitter. The application aims to spark creativity and introduce novel ideas for fictional horror story writing. Shelley also provides a form of entertainment for the	Shelley generates and posts new stories on Twitter and invites followers to participate in collaborative story generation. Shelley tracks the responses sent by the human users and automatically generates and posts new continuations. Shelley also responds to stories initiated by Twitter users.	**Interface:** GUI **Stage of Development:** In use **Country of Development:** Turkey and Germany **Sector:** Academia/research

Type	Description	Collaborative process	Details
	humans interacting with it on Twitter.		
	Calliope (Urban Davis et al. 2021) is a collaborative virtual reality system that enables designers to explore and manipulate generative design solutions to a given design problem. The objective of Calliope is to provide innovative and novel designs to enhance human creativity.	Calliope facilitates an iterative design process, while encouraging a collaborative dialogue between the machine and designer. The human user generates or selects an object/form they would like to create from a range of available options. The AI then generates various designs based on specifications and edits that the human makes to the object. The AI can also integrate elements of design from various samples.	**Interface:** Virtual reality headset **Stage of Development:** Prototype **Country of Development:** USA, Canada **Sector:** Academia/research and private organization (Autodesk)
	The **Cobbie** (Lin et al. 2020) robot is designed to support the human design process by providing inspirational stimuli at the ideation stage. The application sketches on top of the human's sketches when invited by the human designer. The objective of the collaboration is to inject new inspiration and creativity into the in early-stage design of consumer goods in a human-led manner.	Cobbie and the human have shared pens for sketching. The human chooses which pen (stroke type) to use at each stage. The human can move Cobbie to select what section of the sketch it uses for inspiration. Cobbie has buttons that allow the human to direct it draw or pause sketching. The human can provide feedback to the Cobbie on the quality of its design.	**Interface:** Cyber-physical system **Stage of Development:** Prototype **Country of Development:** China **Sector:** Academia/research

Type	Description	Collaborative process	Details
	Wordcraft (Coenen et al. 2021) is an AI-assisted editor for story writing in which a human writer and AI system collaborate to write a story. The collaborative writing assistant supports the writing process from planning through to writing and editing.	Wordcraft comprises a traditional text editor and an AI assistant. The human writer can ask the AI to complete specific tasks, like expanding upon a plot or story that the human agent has started, elaborating on specific elements of a story and filling in elements of the story to spark new ideas. A side-by-side word-processor and dashboard are used to enable communication between the human and the AI agent.	**Interface:** GUI **Stage of Development:** Prototype **Country of Development:** USA **Sector:** Private organization (Google)
Industrial agents	The **COMPLEMANT** system (COllaborative robot aMPLifying and Extending huMAN capabiliTies) (Bettoni et al. 2020) comprises a cobot and smart decision-maker for industrial and manual processes. The collaboration is designed to increase productivity and reduce the fatigue and stress of human industrial workers.	Wearable devices collect the human agent's mental and physical workload and physiological responses during industrial tasks. A vision system monitors the environment and work process. The COMPLEMANT system uses inputs from these sensors to compute and suggest optimal working configurations for human workers and cobots to complete the industrial process. The human worker can ask for reconfigurations or choose to reject the suggestions made by COMPLEMANT.	**Interface:** Cyber-physical system **Stage of Development:** Prototype **Country of Development:** Switzerland **Sector:** Academia/research
	ARMAR-6 (Asfour et al. 2018) is a humanoid robot able to assist industrial workers. It can carry out various tasks, including grasping, mobile manipulation, integrated perception and natural	ARMAR-6 can safely and intuitively collaborate with humans in various industrial and maintenance tasks by combining advanced sensorimotor skills with learning and reasoning abilities. ARMAR-6 can infer when a human needs help, and proactively offer the most appropriate assistance. ARMAR-6 can recognize human activities and	**Interface:** Cyber-physical system **Stage of Development:** Prototype **Country of Development:**

Type	Description	Collaborative process	Details
	language understanding. It is equipped with a suite of sensors and the cognitive abilities to facilitate natural, safe collaboration with humans. The objective of the collaboration is to provide support so that humans can concentrate on the 'skilled' part of a job whilst the robot takes on heavy lifting and support roles.	intentions, reason about situations, and interact with humans in a natural way. ARMAR-6 can also grasp and manipulate objects bimanually to accurately and safely use tools such as power drills and hammers.	Germany **Sector:** Academia/research
	Bionicworkplace (Festo 2018) is an intelligent collaborative workstation that works with humans in customized, short manufacturing runs. The objective of the system is to support workers to carry out manufacturing tasks more efficiently, relieving them of tiring or hazardous tasks.	The Bionicworkplace workstation is equipped with a bionic arm, various assistance systems, peripheral devices, and tools. It uses sensors and camera systems to register the positions of the human operator, components, and tools. It uses this information to derive the optimal program sequence of operations to be carried out between the human and the robot to manufacture the target item. The system continuously learns from each action. Humans can intuitively control the Bionicworkplace by means of gestures, touch, or speech.	**Interface:** Cyber-physical system **Stage of Development:** In commercial use **Country of Development:** Germany **Sector:** Private organization (Festo)
	SHERLOCK solution (Dimitropoulos et al. 2020, Dimitropoulos et al. 2021) provides robotic assistance for human workers on a variety of manufacturing and industrial tasks by	The SHERLOCK solution system is a general-purpose robotic system that can be adopted for use across a range of manufacturing and industrial tasks. It is comprised of three modules that can capture the human operator, environment, and process status of a task. It can identify the tasks that are being executed	**Interface:** Cyber-physical system **Stage of Development:** Prototype **Country of Development:**

Type	Description	Collaborative process	Details
	identifying what actions need to be completed and then supporting the human to complete the task by moving and holding objects based on the human operator's preferences and ergonomic factors. The system is designed to improve productivity, quality and safety and enable customisation of manufacturing and assembly processes.	by the human operator using vision-based machine learning and provide customized support to reduce physical strain and promote ergonomically correct posture. It also responds to inputs from the human operator around their preferences for how they would like to carry out the industrial task.	Greece **Sector:** Academia/research
Healthcare agents	**HALS (Human-Augmenting Labelling System)** (Diao, Chen, and Kvedar 2021, van der Wal et al. 2021) assists humans with cellular annotation on tissue images to identify and label pathological cell types. It is designed to decrease the human annotator's workload and increase the effectiveness of the annotated data, enabling annotation of data sets that were previously cost prohibitive.	A human annotator labels cells within a small region of the tissue slide. An untrained classifier AI system is trained based on the human's annotations and then begins making suggestions to the human, which the human can accept or reject. HALS learns from these corrections and progressively improves its predictions progressively over time. The AI also identifies patches of the slide to annotate next, guiding the annotator around the image. The machine has three deep learning models– to (1) learn the labels provided by an annotator (2) provide recommendations to that annotator designed to increase their speed, and (3) determine the next best data to label to increase the overall quality of annotations while minimising total labelling burden.	**Interface:** GUI **Stage of Development:** Prototype **Country of Development:** USA, South Korea **Sector:** Academia/research and private organization (Salesforce)

Type	Description	Collaborative process	Details
	The **SAGE** patient management system (Goldberg, Belyaev, and Sluchak 2021) supports accurate diagnosis and treatment of pathologic syndromes. The AI agent does not propose a diagnosis, but it analyses the health provider's diagnosis and treatment plan in relation to data collected from the patient to identify where there may be logical inconsistencies and contradictions in the diagnosis and ongoing management of the patient.	A physician determines input parameters (vital signs, labs, imaging) for a patient and makes a preliminary diagnosis. SAGE checks the data for possible errors and develops a diagnosis. If SAGE's prognosis matches the physician's opinion, SAGE does not manifest itself in any way. If they do not match, SAGE displays the parameters which do not match the physician's diagnosis and prompts the physician with an additional question. SAGE determines the most useful question that would challenge the physician's opinion. If an answer to this question matches SAGE hypothesis, but the physician's original assessment of the situation remains unaltered, the SAGE asks a new question (maximum two additional questions). SAGE displays the trends of the overall severity of the patient's condition, the severity of leading syndromes relative to the dynamics of type and intensity of treatment. SAGE offers trends of parameters that need attention when defining treatment. The system detects and analyses contradictions in the provider's decision-making process, as it compares the provider's prognosis with the outcomes which are reflected in improvement or deterioration of the patient's condition. It brings to attention various contradictions in the management of the patient.	**Interface:** GUI **Stage of Development:** Prototype **Country of Development:** USA **Sector:** Academia/research
Emergency Response agents	**DroneResponse** (Agrawal, Cleland-Huang, and	Teams of humans and semi-autonomous UAVs are deployed to find people in need of rescue and assist in their safe recovery.	**Interface:** Cyber-physical

Type	Description	Collaborative process	Details
	Steghöfer 2020) is a cyber-physical system (CPS) involving the collaboration of humans and smart, semi-autonomous unmanned aerial vehicles (UAVs) to carry out rescue operations and emergency response missions. The purpose of the collaboration is to support faster, more successful, and safer emergency responses.	Directed by the human personnel, the UAVs scans an area to locate the person in need of rescue and request feedback (e.g., to confirm a potential rescuee sighting). Once the rescuee is located, the AI agent can request additional support from other UAVs (e.g., one equipped with a flotation device that can be dropped to the rescuee) while streaming imagery of the rescuee until human rescuers arrive. Humans can leverage this information to make more informed decisions and directly intervene in the UAV's behaviour (e.g., directing the UAV to collect imagery from a different angle or assigning additional UAVs to provide support to the rescue effort).	**Stage of Development:** Prototype **Country of Development:** USA **Sector:** Academia/research
Knowledge work agents	**ForSense** (Rachatasumrit et al. 2021) is a web browser extension that supports and improves online research work. The AI agent accelerates and extends the human's online research activity by identifying additional material and suggesting ways of classifying the material.	The human begins identifying relevant material from webpages and organizing it into groups. The AI uses that material to identify other website material relevant to the research and suggest groups the new material may belong to. The human can accept or reject the new material and refine the original groupings based on their own sensemaking or suggestions provided by the AI. If the human decides to regroup material the AI recomputes the word embeddings so that it can begin to identify website material relevant to the new groupings.	**Interface:** GUI **Stage of Development:** Prototype **Country of Development:** USA **Sector:** Academia/research and private organization (Microsoft)

Table 1. Collaborative intelligence applications identified from the review.

Characteristics of collaborative intelligence applications

The second objective of this review was to profile the current state of collaborative intelligence applications by describing their characteristics.

Stage of development

Although our analysis covered documents published between 2011 and 2021, all of the 16 collaborative intelligence applications included in the analysis were documented after 2017, indicating the emerging nature of this field. Three of the collaborative intelligence applications had been released for use publicly, namely, Flow Machines Professional (Avdeeff 2019), Shelley (Yanardag, Cebrian, and Rahwan 2021, O'Brien 2017) and Bionicworkplace (Festo 2018, Kärcher et al. 2017) (see Table 1). Bionicworkplace was the only identified collaborative intelligence application that was commercially available. Bionicworkplace, developed by Festo, is a complex and adaptable cyber-physical system comprised of various sensors, tools, and abilities to provide a flexible, collaborative workstation. The adaptability of the cobot is especially beneficial for increasing productivity in the development of short-runs of customized items (Kärcher et al. 2017).

The other two publicly available collaborative intelligence applications were offered as free digital products. Flow Machines, developed by Sony (Pachet, Roy, and Carré 2021), enables collaboration between humans and AI to enhance and inspire creativity during music production (Pachet, Roy, and Carré 2021). Shelley, a TwitterBot developed by a group of researchers to probe the success of human–AI collaborative fictional horror story development, has been deployed for public use on the Twitter platform, resulting in over 500 collaborative narratives. Shelley was designed to enhance the emotional impact of human stories, introducing novel and surprising directions to the narrative (Yanardag, Cebrian, and Rahwan 2021). The collaborative stories between Shelley and humans were found to induce greater negative effect and state anxiety than those created by Shelley alone, indicating the success of collaborative creation (Yanardag, Cebrian, and Rahwan 2021). There is no indication that it has been used to generate any financial profit for the developers (Yanardag, Cebrian, and Rahwan 2021). The remaining 13 collaborative intelligence applications were defined as prototypes and proof of concept designs, developed in universities and research organisations. Some of these applications appeared to be undergoing further development towards bringing the application to market for public use, for example DroneResponse, AMAR-6 and Forsense (see Table 1).

Collaboration channels

The applications that we identified were evenly divided in terms of whether the collaboration between the human and AI occurred in a virtual or physical environment. The virtual collaborations occurred through a graphical user interface (GUI) (n=7) and a virtual reality headset (n=1). Forsense (see Table 1) is an example of collaborative intelligence that occurs in a virtual environment. It was developed for people carrying out exploratory research online (e.g. collating, organizing and making sense of information) (Rachatasumrit et al. 2021). The system provides collaborative support to human users allowing them to accelerate, improve and coordinate their search tasks. The collaboration is enabled through a web browser extension. The GUIs used in the identified collaborative applications have an interface that enables the human and the AI to provide feedback and respond to one another, although the human is the arbiter in decision-making (with the exception of Shelley (Yanardag, Cebrian, and Rahwan 2021, O'Brien 2017)).

The collaborations involving cyber-physical systems used robots (or cobots) or drones. The human and the cobot collaborated through a virtual communication channel using a GUI, sensors or wearable devices. DroneResponse (Agrawal, Cleland-Huang, and Steghöfer 2020) is an example of a cyber-physical system where communication occurs through a GUI. The DroneResponse prototype was developed in 2020 by a group of academics from the United States using semi-autonomous UAVs that collaborate with human agents through GUIs to provide faster, more successful and safer emergency responses than human or UAVs could alone (Agrawal, Cleland-Huang, and Steghöfer 2020). In this application, the GUI enabled bi-directional communication around mission plans and situational changes between the human and UAV rescue teams. ARMAR-6 (Asfour et al. 2018) was another example of a GUI-based cyber-physical cobot, which collaborated with human technicians on maintenance tasks. Communication with the human agent occurred through gestures, voice commands and various sensors used by the cobot to detect changes in the state of a task or their human collaborator.

In terms of the way that information is shared between the human and AI agents during collaboration, the SAGE application (Goldberg, Belyaev, and Sluchak 2021) presents a unique design choice. It is a patient management system designed to collaborate with human medical practitioners to improve diagnostics and patient treatment and care. The designers of SAGE sought to develop a system that communicates with human agents in an explainable way to build trust and ultimately better collaborative outcomes. Non-collaborative patient management systems will simply provide a

diagnosis, but SAGE collaborates with the practitioner by probing data collected from the patient during the course of treatment to identify and communicate any indicators that are not consistent with the practitioners' original diagnosis or prognosis. SAGE uses a high-level and intuitive visual interface to communicate any issues of concern but the practitioners can interrogate SAGE to understand the patient indicators that underpin the issues that SAGE identifies. In this way, SAGE can reduce the impact of the human practitioners' biases or limited attention and improve decision confidence (Goldberg, Belyaev, and Sluchak 2021). The explainability and transparency of the system's decision making was designed to address socio-technological barriers such as trust between humans and machines that affect the quality of collaborative decision making (Goldberg, Belyaev, and Sluchak 2021).

Types of benefits/outcomes sought from the collaboration

Enhanced productivity or accuracy was a common objective for the collaborative intelligence applications that we identified. Improving human worker satisfaction and safety was another motivation for collaborative intelligence applications. The third type of outcome sought was creativity.

An example of an application that was designed to increase productivity is HALS (Diao, Chen, and Kvedar 2021, van der Wal et al. 2021). HALS is a collaborative human–AI labelling workflow that assists human pathologists with cellular annotation of pathological cell types in tissue samples. The system is designed to enable accurate annotation of large data sets that were previously cost prohibitive. The collaborative use of HALS by a group of seven expert pathologists found a manual work reduction of 91% along with a boost in data quality of 4.34% (van der Wal et al. 2021).

There were several collaborative intelligence applications that were designed to improve human safety, either for a human client or for a human worker. Bionicworkspace (a cobot used in collaborative manufacturing and production tasks) is designed to reduce physical strain on human workers as well as improving their productivity (Kärcher et al. 2017). On the other hand, SAGE (Goldberg, Belyaev, and Sluchak 2021) and DroneResponse (Agrawal, Cleland-Huang, and Steghöfer 2020) improve safety by supporting accurate patient diagnosis and management (SAGE) and providing faster search and response for people in need of rescue (DroneResponse).

Evolver (Feldman 2017) is example of a collaborative intelligence application designed to enhance human creativity during the production of artworks and graphic designs. Evolver collaboratively

produces generative graphic design artifacts based on constraints controlled by a human graphic designer in an iterative design process enabled through a software program. A group of ten designers went on to test the application and how it impacted their outputs and the creative process. It was found that the collaboration provided accessibility to alternative design solutions, helping designers to step out of their current frame of reference to assist with the ideation process (Feldman 2017). Although the designers largely reported positive experiences associated with collaborating with Evolver, the issue of authorship was raised by several participants. This is likely to be an issue that arises with a number of co-creative collaborative agents and could be a challenge to the adoption of these technologies into professional creative practices.

Discussion

Our systematic review of the literature yielded 16 examples of collaborative intelligence applications in which human and AI agents work collaboratively to achieve a shared outcome that achieves more than either agent could achieve on their own. Of these examples, almost all were in early/prototype stages of development rather than commercially available products. There was a mix of embodied cyber-physical and virtual software systems and the channels and format of the communication between the human and the AI also took several forms, ranging from sensors and computer vision to natural language processing and GUI key commands. The applications were designed for a variety of fields including healthcare, manufacturing, graphic design, emergency services and creative writing.

Technological feasibility of collaborative intelligence

The range of applications that we identified illustrate that collaboration between humans and AI is technologically feasible across a range of domains and towards multiple ends. Furthermore, the applications delivered a range of benefits. Working with creative collaborative intelligence applications improved both efficiency and creativity. Manufacturing and assembly collaborative intelligence applications improved efficiency and health and safety. Knowledge work applications improved the accuracy and coverage of the decisions and classifications that were produced. We infer that it is technologically feasible to combine a variety of human and AI capabilities and thereby achieve benefits in terms of efficiency, quality, creativity, safety, and human enjoyment.

Current constraints and future directions for collaborative intelligence

The collaborative intelligence applications identified in this review were restricted in terms of the range of environments in which they could operate and the variety of tasks that they could perform. Whereas human to human collaborators often make frequent switches between collaborating in the real world (e.g., to scope requirements in an initial meeting) and collaborating virtually (when they write a document together or share information and advice via online channels), the current collaborative intelligence applications do not have this agility. Collaborative intelligence applications can work with humans in a virtual environment (e.g., generating digital content) or in the real world where they take a cyber-physical form. However, while the cyber-physical forms of collaborative intelligence often communicate with their human collaborator via online channels, they are generally not capable of collaborating with them in this environment. Whereas human workers will often collaborate over multiple stages of a project or task, the collaborative intelligence applications were limited to discrete tasks and stages of production or decision-making. This uniquely human ability, the capacity to transfer capability in one domain to a related domain, will be a key challenge in the further development of future collaborative intelligence applications.

The review also revealed hundreds of applications that met some but not all the criteria for collaborative intelligence. These applications suggests that future applications of collaborative intelligence will emerge in finance, investment, and insurance; defence and security operations; and within scientific and research.

Limitations

One limitation of this study is that it is based on published academic and grey literature written in English. There are likely to be more collaborative intelligence applications under development not captured in the literature because they are not yet ready for commercialisation, and represent valuable intellectual property. Patent databases may offer another fruitful dataset for identifying additional examples of collaborative intelligence.

A second issue that affected the review was the lack of detailed information on all the potential collaborative intelligence applications that were identified. Several AI applications potentially met the criteria for collaborative intelligence but were not described in sufficient detail to be assessed fully. Thus, while this paper provides the important foundations for documenting current applications of collaborative intelligence, it may not capture the full range of applications that currently exist.

Conclusion

The dominance (up to now) of AI applications that automate work can be attributed to the fact that collaborative AI requires advanced capabilities such as the ability to model the human view of the world and engage in dialogue with a human collaborator. The latest wave of AI is developing these capabilities (Stowers et al. 2021). The complementarities of human intelligence and AI and the range of benefits already associated with collaborative intelligence applications suggests that collaborations between human intelligence and AI can catalyze a new wave of innovation that enables more efficient, safer, sustainable and enjoyable work and lives (Nahavandi 2019).

Conflict of interest

The authors report there are no conflicts of interest to disclose.

References

- Agrawal, A., J. Cleland-Huang, and J. P. Steghöfer. 2020. Model-Driven Requirements for Humans-on-the-Loop Multi-UAV Missions. 2020 IEEE Tenth International Model-Driven Requirements Engineering (MoDRE), 31-31 Aug. 2020.

- Agrawal, A., J. S. Gans, and A. Goldfarb. 2019. Exploring the impact of artificial intelligence: Prediction versus judgment. *Information Economics and Policy* 47:1-6. doi: https://doi.org/10.1016/j.infoecopol.2019.05.001.

- Arksey, H. and L. O'Malley. 2005. Scoping studies: Towards a methodological framework. *International Journal of Social Research Methodology* 8 (1):19-32. doi: 10.1080/1364557032000119616.

- Asfour, T., L. Kaul, M. Wächter, S. Ottenhaus, P. Weiner, S. Rader, R. Grimm, Y. Zhou, M. Grotz, F. Paus, D. Shingarey, and H. Haubert. 2018. ARMAR-6: A Collaborative Humanoid Robot for Industrial Environments. 2018 IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids), 6-9 Nov. 2018.

- Avdeeff, M. 2019. Artificial intelligence & popular music: Skygge, flow machines, and the audio uncanny valley. *Arts* 8 (4). doi: 10.3390/arts8040130.

- Battina, D. S. 2018. The future of artificial intelligence at work: A review on effects of decision automation and augmentation on workers targeted by algorithms and third-party observers. *International Journal of Innovations in Engineering Research and Technology* 5 (7).

- Bettoni, A., E. Montini, M. Righi, V. Villani, R. Tsvetanov, S. Borgia, C. Secchi, and E. Carpanzano. 2020. Mutualistic and adaptive human–machine collaboration based on machine learning in an injection moulding manufacturing line. *Procedia CIRP* 93. doi: 10.1016/j.procir.2020.04.119

- Billman, D., G. Convertino, J. Shrager, P. Pirolli, and J. Massar. 2006. Collaborative intelligence analysis with CACHE and its effects on information gathering and cognitive bias. Human Computer Interaction Consortium Workshop.

- Cesta, A., A. Orlandini, and A. Umbrico. 2018. Fostering Robust Human-Robot Collaboration through AI Task Planning. 51st CIRP Conference on Manufacturing Systems (CIRP CMS), Stockholm, SWEDEN, May 16-18.

- Cienki, A. 2015. Insights into coordination, collaboration, and cooperation from the behavioral and cognitive sciences: A commentary. *Interaction Studies* 16 (3):553-560.

- Coenen, A., L. Davis, D. Ippolito, E. Reif, and A. Yuan. 2021. Wordcraft: A Human-AI Collaborative Editor for Story Writing. Ithaca: Cornell University Library, arXiv.org.

- Daugherty, P. R. and H. J. Wilson. 2018. *Human + machine: Reimagining work in the age of AI.* Brighton, USA: Harvard Business Review Press.

- De Luca, G. 2021. The development of machine intelligence in a computational universe. *Technology in Society* 65:101553. doi: https://doi.org/10.1016/j.techsoc.2021.101553.

- Dellermann, D., A. Calma, N. Lipusch, T. Weber, S. Weigel, and P. Ebel. 2019. The future of human-AI collaboration: A taxonomy of design knowledge for hybrid intelligence systems. *Hawaii International Conference on System Sciences.*

- Diao, J. A., R. J. Chen, and J. C. Kvedar. 2021. Efficient cellular annotation of histopathology slides with real-time AI augmentation. *NPJ Digital Medicine* 4 (1). doi: 10.1038/s41746-021-00534-0.

- Dimitropoulos, N., T. Togias, G. Michalos, and S. Makris. 2020. Operator support in human-robot collaborative environments using AI enhanced wearable devices.

- Dimitropoulos, N., T. Togias, N. Zacharaki, and G. Michalos. 2021. Seamless human–robot collaborative assembly using artificial intelligence and wearable devices. *Applied Sciences* 11 (12):5699. doi: http://dx.doi.org/10.3390/app11125699.

- Dubey, A., K. Abhinav, S. Jain, V. Arora, and A. Puttaveerana. 2020. HACO: A framework for developing human-AI teaming. *Proceedings of the 13th Innovations in Software Engineering Conference on Formerly known as India Software Engineering Conference.*

- Epstein, S. L. 2015. Wanted: Collaborative intelligence. *Artificial Intelligence* 221:36-45. doi: https://doi.org/10.1016/j.artint.2014.12.006.

- European Commission, Directorate-General for Research Innovation, M. Breque, L. De Nul, and A. Petridis. 2021. *Industry 5.0: Towards a sustainable, human-centric and resilient European industry*: Publications Office.

- Feldman, S. 2017. Co-creation: human and AI collaboration in creative expression. Proceedings of the conference on Electronic Visualisation and the Arts, London, United Kingdom.

- Festo. 2018. Bionicworkplace: Human-robot collaboration with artificial intelligence. Proceedings of 2018 International Conference on Hydraulics and Pneumatics, Romania.

- Gao, X., L. Yan, G. Wang, and C. Gerada. 2021. Hybrid recurrent neural network architecture-based intention recognition for human-robot collaboration. *IEEE Transactions on Cybernetics*:1-9. doi: 10.1109/TCYB.2021.3106543.

- Goldberg, S., S. Belyaev, and V. Sluchak. 2021. Dr. Watson type artificial intellect (AI) systems. Ithaca: Cornell University Library, arXiv.org.

- Goldfarb, A. and J. Lindsay. 2020. Artificial intelligence in war: Human judgment as an organizational strength and a strategic liability. Washington D.C., USA.

- Hackman, J. R. and G. R. Oldham. 1975. Development of the job diagnostic survey. *Journal of Applied Psychology* 60 (2):159-170. doi: 10.1037/h0076546.

- Jarrahi, M. H. 2018. Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Business Horizons* 61 (4):577-586. doi: https://doi.org/10.1016/j.bushor.2018.03.007.

- Johnson, M., J. M. Bradshaw, P. J. Feltovich, C. M. Jonker, M. B. van Riemsdijk, and M. Sierhuis. 2014. Coactive design: Designing support for interdependence in joint activity. *Journal of Human-Robot Interaction* 3 (1):43–69. doi: 10.5898/JHRI.3.1.Johnson.

- Kärcher, N., M. Moerdijk, S. Schrof, C. Trapp, M. Purucker, M. Baltes, N. Peltzer, I. Lamich, A. Hildebrandt, and R. Neumann. 2017. BionicCobot: Sensitive helper for human-robot collaboration. Germany.

- Kasparov, G. 2010. The Chess Master and the Computer. USA.

- Kolbeinsson, A., E. Lagerstedt, and J. Lindblom. 2019. Foundation for a classification of collaboration levels for human-robot cooperation in manufacturing. *Production & Manufacturing Research* 7 (1):448-471. doi: 10.1080/21693277.2019.1645628.

- Langer, M. and R. N. Landers. 2021. The future of artificial intelligence at work: A review on effects of decision automation and augmentation on workers targeted by algorithms and third-party observers. *Computers in Human Behavior* 123:106878. doi: https://doi.org/10.1016/j.chb.2021.106878.

- Liberati, A., D. G. Altman, J. Tetzlaff, C. Mulrow, P. C. Gøtzsche, J. P. A. Ioannidis, M. Clarke, P. J. Devereaux, J. Kleijnen, and D. Moher. 2009. The PRISMA statement for reporting systematic reviews and meta-analyses of studies that evaluate healthcare interventions: Explanation and elaboration. *BMJ* 339:b2700. doi: 10.1136/bmj.b2700.

- Lin, Y. Y., J. H. Guo, Y. Chen, C. Yao, and F. T. Ying. 2020. It Is Your Turn: Collaborative Ideation with a Co-Creative Robot through Sketch. CHI Conference on Human Factors in Computing Systems (CHI), Electr Network, Apr 25-30.

- Maddikunta, P. K. R., Q. V. Pham, P. B, N. Deepa, K. Dev, T. R. Gadekallu, R. Ruby, and M. Liyanage. 2022. Industry 5.0: A survey on enabling technologies and potential applications. *Journal of Industrial Information Integration* 26:100257. doi: https://doi.org/10.1016/j.jii.2021.100257.

- Madni, A. M. and C. C. Madni. 2018. Architectural Framework for Exploring Adaptive Human-Machine Teaming Options in Simulated Dynamic Environments. *Systems* 6 (4):44.

- Mason, C. M., M. Ayre, and S. M. Burns. 2022. Implementing industry 4.0 in Australia: Insights from advanced Australian manufacturers. *Journal of Open Innovation: Technology, Market, and Complexity* 8 (1):53.

- McDermott, P., C. Dominguez, N. Kasdaglis, M. Ryan, I. Trahan, and A. Nelson. 2018. Human-Machine Teaming Systems Engineering Guide. USA.

- Nahavandi, S. 2019. Industry 5.0 - A Human-Centric Solution. *Sustainability* 11 (16):4371.

- O'Brien, M. 2017. "Nightmare machine writes bone-chilling tales." *Chronicle - Herald*, 10/31/2017 Oct 31. https://www.proquest.com/newspapers/nightmare-machine-writes-bone-chilling-tales/docview/1958636667/se-2?accountid=26957

- Oh, C., J. Song, J. Choi, S. Kim, S. Lee, and B. Suh. 2018. I Lead, You Help but Only with Enough Details: Understanding User Experience of Co-Creation with Artificial Intelligence. Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, Montreal QC, Canada.

- Ostheimer, J., S. Chowdhury, and S. Iqbal. 2021. An alliance of humans and machines for machine learning: Hybrid intelligent systems and their design principles. *Technology in Society* 66:101647. doi: https://doi.org/10.1016/j.techsoc.2021.101647.

- Pachet, F., P. Roy, and B. Carré. 2021. Assisted Music Creation with Flow Machines: Towards New Categories of New. In *Handbook of Artificial Intelligence for Music: Foundations, Advanced Approaches, and Developments for Creativity*, edited by Eduardo Reck Miranda, 485-520. Cham: Springer International Publishing.

- Parker, S. K. and G. Grote. 2019. Automation, algorithms, and beyond: Why work design matters more than ever in a digital world. *Applied Psychology*. doi: https://doi.org/10.1111/apps.12241.

- Poser, M., and E. Bittner. 2020. *Hybrid teamwork: Consideration of teamwork concepts to reach naturalistic interaction between humans and conversational agents.* Wirtschaftsinformatik (Zentrale Tracks).

- Rachatasumrit, N., G. Ramos, J. Suh, R. Ng, and C. Meek. 2021. ForSense: Accelerating online research through sensemaking integration and machine research support. *Transactions on Interactive Intelligent Systems* 12 (4). doi: 10.1145/3532853.

- Saffiotti, A., P. Fogel, P. Knudsen, L. de Miranda, and O. Thörn. 2020. On human-AI collaboration in artistic performance. https://www.semanticscholar.org/paper/On-human-AI-collaboration-in-artistic-performance-Saffiotti-F%C3%B6gel/b94d654e88d9f5fe2f1c371f8d0e80c44830ceef

- Schuh, G., R. Anderl, R. Dumitrescu, A. Krüger, and M. ten Hompel. 2020. Industrie 4.0 Maturity Index: Managing the Digital Transformation of Companies - Update 2020. Germany.

- Seeber, I., E. Bittner, R. O. Briggs, T. de Vreede, G. J. de Vreede, A. Elkins, R. Maier, A. B. Merz, S. Oeste-Reiß, N. Randrup, et al. 2020. Machines as teammates: A research agenda on AI in team collaboration. *Information & Management* 57 (2):103174. doi: https://doi.org/10.1016/j.im.2019.103174.

- Shook, E. and M. Knickrehm. 2018. Reworking the revolution. Accenture.

- Sindhwani, R., S. Afridi, A. Kumar, A. Banaitis, S. Luthra, and P. L. Singh. 2022. Can industry 5.0 revolutionize the wave of resilience and social value creation? A multi-criteria framework to analyze enablers. *Technology in Society* 68:101887. doi: https://doi.org/10.1016/j.techsoc.2022.101887.

- Stowers, K., L. L. Brady, C. MacLellan, R. Wohleber, and E. Salas. 2021. Improving teamwork competencies in human-machine teams: Perspectives from team science. *Frontiers in Psychology* 12. doi: 10.3389/fpsyg.2021.590290.

- Szász, L., K. Demeter, B. G. Rácz, and D. Losonci. 2021. Industry 4.0: A review and analysis of contingency and performance effects. *Journal of Manufacturing Technology Management* 32 (3):667-694. doi: 10.1108/JMTM-10-2019-0371.

- Thompson, C. 2013. *Smarter Than You Think: How Technology Is Changing Our Minds for the Better.* East Rutherford: Penguin Publishing Group.

- Thörn, O., P. Knudsen, and A. Saffiotti. 2020. Human-Robot Artistic Co-Creation: A Study in Improvised Robot Dance. 2020 29th IEEE International Conference on Robot and Human

Interactive Communication (RO-MAN), 31 Aug-4 Sept. 2020.

- Traumer, F., S. Oeste-Reiß, and J. M. Leimeister. 2017. Towards a Future Reallocation of Work between Humans and Machines – Taxonomy of Tasks and Interaction Types in the Context of Machine Learning. Thirty Eighth International Conference on Information Systems, South Korea.

- Urban Davis, J., F. Anderson, M. Stroetzel, T. Grossman, and G. Fitzmaurice. 2021. Designing co-creative AI for virtual environments. *Creativity and Cognition.* doi: 10.1145/3450741.3465260.

- van der Wal, D., I. Jhun, I. Laklouk, J. Nirschl, L. Richer, R. Rojansky, T. Theparee, J. Wheeler, J. Sander, F. Feng, et al. 2021. Biological data annotation via a human-augmenting AI-based labeling system. *NPJ Digital Medicine* 4 (1):145. doi: 10.1038/s41746-021-00520-6.

- Veile, J. W., D. Kiel, J. M. Müller, and Kai-Ingo Voigt. 2020. Lessons learned from industry 4.0 implementation in the German manufacturing industry. *Journal of Manufacturing Technology Management* 31 (5):977-997. doi: 10.1108/JMTM-08-2018-0270.

- Wang, D., E. Churchill, P. Maes, X. Fan, B. Shneiderman, Y. Shi, and Q. Wang. 2020. From Human-Human Collaboration to Human-AI Collaboration: Designing AI Systems That Can Work Together with People. Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems, Honolulu, HI, USA.

- Yanardag, P., M. Cebrian, and I. Rahwan. 2021. Shelley: A crowd-sourced collaborative horror writer. *Creativity and Cognition.* doi: 10.1145/3450741.3465251.

- Zhang, C., C. Yao, J. Liu, Z. Zhou, W. Zhang, L. Liu, F. Ying, Y. Zhao, and G. Wang. 2021. StoryDrawer: A co-creative agent supporting children's storytelling through collaborative drawing. *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems.*

Declarations

Potential competing interests: No potential competing interests to declare.