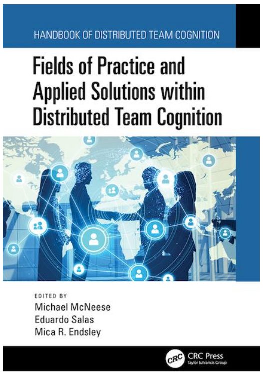

Fields of Practice and Applied Solutions within Distributed Team Cognition

Michael D. McNeese, Eduardo Salas, Mica R. Endsley

Distributed Cognition and Human-Co-Robot Manufacturing Teams

Publication details
https://test.routledgehandbooks.com/doi/10.1201/9780429459542-7
Lora A. Cavuoto, Ann M. Bisantz

7 Distributed Cognition and Human-Co-Robot Manufacturing Teams

Issues in Design and Implementation

Lora A. Cavuoto and Ann M. Bisantz

CONTENTS

INTRODUCTION

Collaborative robots (co-robots) present a potentially transformative technology for small, medium, and large manufacturers because they enable manufacturing's desire for rapid, flexible automation capable of semi-manual/automated tasks. Manufacturing co-robots exploit the flexibility and creative aspects of (human) manual work with the efficiency and productivity of automation—if they are designed and integrated into manufacturing operations effectively.

Co-robots have been deployed across a range of manufacturing environments, with the largest implementations existing in the automotive, electronics, and medtech industries. In automotive manufacturing, for example, manufacturers have installed co-robots on assembly lines for high-precision tasks such as screw-driving, sealant dispensing, and plug insertion. Other applications include electronics inspection and carton packaging. A recent study by BMW and MIT showed that human-co-robot teams performing cooperative work reduced human idle time by about 80% and were more productive than human- or robot-only teams (Unhelkar, Lasota et al., 2018). Most current applications involve payloads in the range of 1 to 5 kg. Co-robot adaptability and affordability, with costs on the order of $50,000 to $100,000 depending on the type and payload requirements, mean that small and medium-sized enterprises (SMEs) have the ability to implement these systems. Scale-up of co-robot installations in manufacturing will continue, due to faster, lower cost deployment than traditional automation, the smaller required footprint of the technology, the potential reduction in ergonomic risk to operators, and an expected labor shortage of skilled manufacturing workers. However, effective scale-up requires a fundamental understanding of manufacturing-specific human-co-robot interaction. Failure to do so can be detrimental in terms of efficiency, productivity, flexibility, and quality, as has been the case with over half of the early introductions of advanced manufacturing technologies and cellular manufacturing (Chung, 1996; Charalambous, Fletcher, & Webb, 2015).

Defining Human-Co-Robot Teams

In designing for effective use, it is important to define the terminology related to co-robots and what distinguishes them from traditional industrial robots. A robot includes the robot arm and its means of control (both software and electronics control system) and a robot system includes the robot along with any accompanying end-effector and parts that are manipulated by the robot arm. Operators include all personnel that interact with the system, including production personnel as well as any maintenance and programming personnel.

In a traditional manufacturing environment with industrial automation, industrial robots are typically separated from the human operator by means of a cage or other physical barrier. When the robot is active, the operator is prevented from entering the blocked area, minimizing the risk of interference between the human and robot. For these systems, the human and robot coexist, operating at different times in separate workspaces. In a second type of robot operation, termed cooperation, a similar industrial robot is used in the same workspace as a human, however all human tasks are done when the robot is inactive, and robot functioning starts once the human is at a safe distance away. An example of cooperation includes tasks where the human operator is responsible for loading and unloading parts. In contrast, with collaborative operation, the human and co-robot share the same space and are active at the same time. As defined by the Robotic Industries Association (RIA; ANSI/RIA Technical Report R15.606–2016) and the International Organization of Standardization (ISO) Technical Specification (ISO/TS 15066, 2016; International Organization for Standardization, 2016; Robotics Industries Association, 2016), a collaborative robot is "a robot that can be used in a collaborative operation within a collaborative workspace" where collaborative operation is a "state in which a purposely designed robot system and an operator work within a collaborative workspace" and collaborative workspace is a "space within the operating space where the robot (including the workpiece) and a human can perform tasks concurrently during production operation." A collaborative robot is defined by the task the robot is performing, the timing of that task with respect to the operator's task, and the space in which the task is being performed, not by the robot itself. Note that collaborative robots, by these definitions, involve collaboration between people and robots occupying the same physical space (and typically, working together to manufacture the same physical artifact), contemporaneously. That is, the notion of collaborative robots, as defined and discussed in this chapter, does not refer to asynchronous collaboration or collaboration at a distance.

An important distinction between co-robots and traditional industrial robots is that co-robots are typically power and force limited, minimizing the risk of injury to the operator if contact occurs. Contact can be intended (through hand guidance), incidental (unintended impact between the operator and robot), or resulting from a failure of the co-robot. The co-robot has sensors and safety functions that detect the contact and stop. Allowable levels of contact force are based on the level of risk and likelihood of pain for the operator (Muttray, Melia et al., 2014; International Organization for Standardization, 2016). Christiernin (2017) provides a relevant classification framework for characterizing the level of interaction with co-robots, ranging from low levels with no collaboration (traditional robots that perform work in a separated, gated space from humans) and co-location only when the robot is idle or moving slowly under human command, to higher-level collaborative situations. These include, for example, humans and robots working in co-located spaces, but where the co-robots' sensors detect human movements and avoid contact, or human-guided robot movement. In most cases, co-robot movement is preplanned based on task requirements and knowledge of expected human movement paths, but adaptive learning is being integrated to support adaptability to unexpected situations. Adaptive learning is discussed later in this chapter. At the highest "collaboration" level, the humans and robots are "working together in the same physical space, solving problems together" and robots adapt their behaviors based on human activities, and may initiate actions.

The primary focus of much of the initial research and development work regarding the integration of co-robots into manufacturing environments has focused on the safety aspects of eliminating machine guarding with co-robot installation. Thus, important questions regarding the roles that co-robots make take on in human-co-robot teams and how both co-robots and teamwork will need to be designed to support successful human-co-robot teams still remain. In response, this chapter identifies theoretical frameworks drawn from human factors and distributed cognition which can inform the study and design of human-co-robot manufacturing teams, including situation awareness, communication, trust/reliability, and function allocation. The impacts of the teams on operator safety, production quality, and efficiency are also explored. Implications for system (both robot and work environment) design and operator training are presented.

THEORETICAL FRAMEWORKS FOR CONSIDERING HUMAN-CO-ROBOT MANUFACTURING TEAMS

Situation Awareness

A critical theoretical lens with which to view distributed human-co-robot teams engaged in the highest level of collaboration (Christiernin, 2017) is that of situation awareness. Endsley's (1995) three-level model of situation awareness can be applied: Level 1, awareness of dynamic system state variables; Level 2, understanding of the meaning of those variable values in context; and Level 3, prediction of the future state of those variables. For any manufacturing situation, these levels would apply to human awareness of relevant manufacturing process variables that make up the work task. Additionally, however, the introduction of a co-robot team member requires additional layers of awareness. Currently, the three most common applications of manufacturing co-robots are in machine tending (i.e. loading raw material or unloading finished parts from a machine like a press), pick and place (i.e. moving parts from one location to another, often with precision), and dispensing (e.g., depositing a bead of fluid adhesive). Consider a joint assembly task in which a co-robot retrieves a part from a 3-D printer and places a part in a fixture, the human operator visually inspects the part before manually installing a clip, and then advances the fixture to a second co-robot, which applies a bead of adhesive. From the human operator perspective, relative to the co-robot, levels of situation awareness in this case would translate to (1) knowing the position and directional velocity of the co-robot arms and actions being taken by the end-effectors for both co-robot arms, (2) understanding the task that is being performed by the robot (e.g., for the first co-robot, the tasks of moving to placing the part, releasing the part, or moving to retrieve the next part; for the second co-robot, tasks of preparing to dispense, dispensing along a path, or resetting to the start of the path), and (3) predicting the future actions and states of the arms (trajectory and end position) and effectors and the next tasks to be performed. Considering the co-robot to be a team member (rather than a system the human is controlling) requires a similar sense of awareness on the part of the co-robot itself. That is, to truly support collaboration as a distributed team, the co-robot should itself have awareness of the work task variables, particularly, in this example, whether or not the human operator has released the previous part (and is therefore ready to receive a new part) and the precise location and orientation of the part, and also be able to (1) sense the position and directional velocity of the human team-member, (2) understand (infer) the task that the person is engaged in, and (3) predict the next movements or actions of the person.

In addition to these components (human and co-robot awareness of the work task; mutual situation awareness of the other's state, activity, and future actions), processes of *team* situation awareness are also critical (Salas, Prince, Baker, & Shrestha, 1995). Team situation awareness implies a shared understanding of the situation and relevant elements in the task environment (e.g., in this case, progress through the manufacturing process, quality-related variables) developed through individual and team processes (e.g., individual information seeking, communication of information among team members; Salas et al., 1995). Team situation awareness is related to the more general concepts of transactive memory systems, shared mental models, and team cognition (Kozlowski, Grand, Baard, & Pearce, 2015; Morrow & Fiore, 2013). A transactive memory system includes shared knowledge of a team including both individuals' knowledge and the understanding of which team member knows what information (Wickens, Hollands, Banbury, & Parasuraman, 2013). Shared mental models include a shared understanding of the task or problem structure along with the knowledge and skills of other team members, and allow team members to work flexibly to adapt to changing work demands and to communicate more effectively (Morrow & Fiore, 2013).

Importantly, then, both team members should know what the other is aware of. That is, the human must understand what the robot knows about the work environment and the humans' movements and activities, and perhaps more critically, understand the limits of that awareness. This understanding is fundamentally tied to another theoretical lens—that of trust—discussed later in the chapter, and can be supported through explicit and transparent displays of not only what the co-robot is doing, but of the model that the co-robot has of the human partner. Team awareness is also reliant on successful communication, also discussed later in the chapter.

To date, although situation awareness in human-robot teams has been considered in a number of domains (Riley & Endsley, 2004; Riley & Endsley, 2005; Riley, Strater, Chappell, Connors, & Endsley, 2012, there has has been limited research specifically with regard to manufacturing teams involving co-robots. Gombolay, Bair, Huang, and Shah (2017) present one example, designed to address "one major gap in the robotics and human factors literature . . . the study of situational awareness wherein humans plan and execute a sequence of actions collaboratively within a human-robot team" (p. 600). They conducted research to understand how the situation awareness of the human team member changed across varying levels of robot autonomy in assigning tasks to human vs. robotic team members, in a situation where two humans and one robot collaborated on a physical assembly task, in part to understand how these systems should be designed to take into account the preferences of the human team member regarding task scheduling and workflow. They found that participant situation awareness of which team members were assigned and performed a task decreased when task assignment was done autonomously vs. by the human participant. However, this study did not address questions of situation awareness more broadly, in terms of mutual awareness of the other's actions in close proximity or shared awareness among the human and robot team member over aspects of the physical task. Additionally, it was not clear how (or if) the robot contribution to team situation awareness was assessed.

An important area of related research includes that of intent inferencing through the use of motion capture and kinematics. A current limitation of most co-robots that restricts their integration into a work cell is an inability to track human movement, estimate the future position, and react accordingly (Ivaldi, Fritzsche et al., 2017). Implementation of the co-robot requires knowledge of all tasks to be performed and detailed programming of the motion paths to be followed. True collaboration requires the robot to understand the human partner and to predict what the human will do next (Hayes & Scassellati, 2013; Ivaldi, Fritzsche et al., 2017). In order to achieve the first piece of this cycle, there is a need to track and model human dynamics. This includes both the body posture and the forces applied at the contact location, the motion and the effort. Most systems that monitor human motion near the robot use the information for safety purposes, allowing the robot speed to adjust when the operator is within predefined areas. Due to sensor costs and reliability concerns, this is traditionally limited to detecting the separation between the robot and the human and having the robot slow down when the distance is below a predefined threshold. Models of worker kinematics remain primarily in the research domain. Recent studies have focused on basic or isolated assembly tasks assigned between the human and co-robot with each agent working separately (Johannsmeier & Haddadin, 2017). For a screw-placing and sealing task, one study of human-aware motion planning showed ~5% faster task performance by the human, ~20% more concurrent motion, and reductions in idle time for both the human and robot (Lasota & Shah, 2015). An important result from this study was that the human-aware motion planning resulted in improved subjective evaluation of robot understanding, robot interference, and robot safety compared to the non-aware motion (Lasota & Shah, 2015). A human-co-robot movement chain is needed to model the shared task assignment for the human-co-robot team. Understanding this combined movement chain will support the goals of minimizing the interference between the human and co-robot and maximizing task performance. As indicated by Hayes and Scassellati (2013), anticipation of object placement and movement during task performance itself directly affects the possible actions and the available space for the collaborator (whether human or co-robot) on the team. Work cell design and task assignment can be affected by the incorporation of monitoring.

However, despite these advances, and in addition to solving the technical challenges related to sensing and motion capture, several important research questions exist when considering the problem of maintaining shared situation awareness in human-co-robot teams. One is the need to provide easy to understand displays or other signals that allow the human operator to develop and maintain an awareness over the co-robot actions. Current displays tend to be single indicators (e.g., red/green lights) that do not provide more detailed state information. Further discussion of human-co-robot communication is provided in the next section. More challenging, but critical to the creation of true team situation awareness, is the need to develop robust methods by which the co-robot can sense and appropriately understand human movement, infer activity, and predict future actions—without requiring explicit and continued input on the part of the human team member.

Communication

In a distributed cognitive system, development of a shared understanding—of the task situation and other team members' goals, activities, and level of shared awareness—depends on shared knowledge (including based on past experiences working together), shared awareness of the world in which the team is functioning, and successful communication about information which is required but not mutually known. For instance, human team members in a traditional manufacturing cell would, through past experience, understand task elements that are particularly challenging and also understand which team members have more or less expertise with certain tasks (and thus may require assistance). Because they are working together, at the same location, they share sensory input regarding, for instance, the operating state of machinery or characteristics of incoming parts. These shared experiences and shared understanding of the work environment contribute to the development of a "common ground," or a set of mutually shared knowledge, beliefs, and assumptions among communicative partners that support understanding and dialog (Clark & Brennan, 1991). Common ground—which encompasses background or fixed knowledge as well as aspects of the current situation and environment—is critical in facilitating, or in some cases reducing the need for, overt communication.

The need to explicitly communicate information and maintain shared situation awareness adds to the workload for distributed teams—team members must recognize the need for communication, act to communicate, check understanding, and effect repair of any communication failures. In human teams, maintenance of common ground and communication can be accomplished in subtle, sometimes implicit, ways that relieve some of this burden. For instance, facial expressions, gestures, or conversational tone can be used to direct attention, communicate confusion, or indicate the need for repair. Previous research on communication in complex work systems has documented the use of other, more implicit communication mechanisms. Co-workers may overhear conversations or observe others' actions and use that information as a cue to begin or inform their own activity (e.g., overhearing a phone conversation to initiate a station announcement in a mass transit control room; Heath & Luff, 1991; using an overheard confirmation of a ship's bearing in creating a positional plot; Hutchins, 1995) rather than requiring an explicit instruction or question. However, relying on auditory interfaces may be infeasible in manufacturing environments due to the ambient noise from machine operation.

Suchman's (1987) seminal work on situated action and human-machine communication revealed how the lack of common ground between human and machine team members could result in communicative error and task failure, even in a highly proceduralized task (making a photocopy). In her research, people were using a newly designed copier with advanced (for the time) capabilities, that included providing step by step directions to the human user and then inferring human intent based on limited sensing (e.g., whether a document had been placed on, or removed from, the glass). Inability to sense error and subsequent mischaracterization of the human user intent (i.e., what process step was being performed) led to an inability for the human-machine team to complete the task. This example, while outside of the manufacturing environment, is nonetheless telling: even for relatively straightforward, highly proceduralized tasks (e.g., as many manufacturing tasks are typically characterized), effective human-co-robot communication, including maintenance of common ground and detection and repair of communication failure, will be essential.

One thread of research related to human-co-robot communication has addressed the use of robot "facial expressions" or anthropomorphized display elements to convey information to human team members. For instance, a number of researchers are considering how robots might convey meaning or direct a teammate's attention through eye gaze behaviors in non-manufacturing environments (Mutlu, Yamaika, Janda, Ishiguro, & Hagita, 2009; Mwangi, Barakova, Diaz-Boladeras, Mollofre, & Roauterbert, 2018). Sauppe and Mutlu (2015) took an ethnographic approach to understanding human workers' social understanding and interaction with co-robots (specifically, the Baxter robot by Rethink Robotics) in three manufacturing work environments. They found that human co-workers were able to use signals from the robot (displayed "eye" movement and expressions) to infer the co-robot's status and intended next actions, and used other cues (e.g., operating sounds or lack of sounds) to identify or diagnose work problems. Operators also expressed a desire for additional communication capabilities, through speech and through strategic use of the robot's display screen (currently used to show the "eyes").

Other research has focused on the use of human gestures to communicate with or control co-robot actions. For instance, Tsarouchi, Matthaiakis, Makris, and Chryssolouris (2017) describe a human-co-robot hybrid assembly cell in which the human team member used gestures (sensed and recognized by the robots) to start, stop, or provide directional guidance to the robots. Importantly, however, the gestures were predetermined and part of a gesture "library" that could be understood by the co-robots, and thus represent an explicit means of communication, rather than an example of recognition of meaning in gestures which occur naturally during conversation or interpretation of human activity (ie. implicit communication). Related research regarding robot-produced gestures has studied how robots might use movement to indicate emotions (Li & Chignell, 2011; Novikova & Watts, 2015) and the interpretability of robot gestures by humans. For instance, Busch, Grizou, Lopes, and Stulp (2017) studied how a co-robot might "learn" to adapt their movements so they are more interpretable by people (specifically, allowing a person to interpret which button an industrial co-robot was going to press).

Important research and design questions related to human-co-robot communication therefore relate to the need to develop methods which support implicit as well as explicit communication, reduce workload related to explicit communication, support "natural" methods of communication and maintenance of common ground, and support effective detection and resolution of communicative problems. For instance, it should be possible for human team members to signal confusion or directives through voice and tone, using natural language. Co-robots should understand human movements and use that understanding to manage activity, without the need for explicit direction, and co-robots should use movements and gestures that are interpretable by their human teammates. Similar to recommendations regarding situation awareness, the co-robot's understanding of the work task, and its understanding of the intentions of the human team member, should be immediately and transparently apparent to the human team member.

Trust

Integral to the concept of high levels of human-co-robot collaboration is the notion of trust. Research on human trust in automated systems has a long history (Muir, 1994; Muri & Moray, 1996; Lee & Moray, 1994) and has leveraged models of trust developed in social science (e.g., Barber, 1983; Rempel, Holmes, & Zanna, 1985). A working definition of trust in the context of human-machine systems is provided by Lee and See (2004, p. 54): "the attitude that an agent will help achieve an individual's goals in a situation characterized by uncertainty and vulnerability."

Relevant findings from the body of research on trust in human-machine systems include the fact that there is at least some evidence which suggests that human concepts of trust are similar across trust in other people vs. trust in automated systems (Jian, Bisantz, & Drury, 2000), and that trust in systems grows through experience with successful operation, but also can be lost when systems do not operate as expected. Development and maintenance of trust is related to system reliability, and importantly, human expectation of and understanding related to the circumstances when automated systems should be expected to perform competently (e.g., concepts of robustness and understandability; Sheridan, 1988). Research in the context of automated decision aids suggests that human assessments of trust are higher when explicit information about the validity and reliability of system recommendations are provided (Seong & Bisantz, 2008).

Lee and See (2004) present a conceptual feedback model of factors which moderate behaviors of the intention to use, and actual reliance on, automated systems, including individual factors (i.e., perceived risk, self-confidence, predispositions to trust) and situation or system-related factors (e.g., workload, time constraints, effort to use). Feedback—that is, the experience of individuals relying on automation, and information gained about the automation through observation of its performance—impacts the development of trust as well as the "calibration" of trust. Calibration refers to appropriately understanding the contexts (temporal and functional) in which the automation should be trusted, or not. Importantly, in terms of design of co-robot displays or support for human-co-robot communication, Lee and See note that the "observability" of automation performance affects the development of trust. Related concepts include misuse stemming from overtrust (relying on a system even though it should not be trusted), disuse (abandoning a system in circumstances when it actually could aid performance Parasuraman & Riley, 1997), and complacency (failing to adequately monitor performance and therefore intervene when required, due to factors such as overtrust or workload; Parasuraman & Manzey, 2010). Also related, Meyer (2004) defines compliance as acting as directed by the system, while reliance is refraining from acting (e.g., exerting control) unless told explicitly otherwise. In a human-co-robot team, these behaviors would manifest (from the human perspective) as the human taking direction from the co-robot (compliance) or allowing the co-robot to work autonomously (reliance).

It is reasonable to assume that within a human-co-robot team, that the human team member will develop some attitude of trust related to the co-robot by observing its performance in the joint work environment, allowing the person to both rely on the co-robot to contribute to shared work tasks effectively, but also (and perhaps more importantly) to keep the person physically safe. Expectations related to co-robot competency can result from experience as well as directed training regarding expectations. Komatsu, Kurosawa, and Yamada (2012) describe an "adaptation gap" which arises when human expectations of a robot's behaviors (influenced by, for instance, anthropomorphic features) are not consistent with those behaviors. If a person's initial expectation of the robot exceeds the actual performance, attitudes towards use and acceptance would be negative. Similarly, unexpected behaviors due to co-robot faults (through mechanical or software failure) or programming changes (for instance, due to software updates) may negatively impact trust. In the case of human-co-robot teams, then, expectations set by training or factors such as the form or appearance of the co-robot regarding safety and competence should correspond to co-robot performance. For instance, Sauppe and Mutlu (2015) discuss how the co-robot's design features led to feelings of safety and comfort on the part of the human team members, but that such "increased sociality has the potential to create false expectations that may risk worker safety," requiring designers to match perceived and actual safety. Research in other complex systems has identified errors when interacting with automated systems in situations where otherwise highly competent systems unexpectedly lack specific capabilities: the human team member simply does not expect sophisticated automation to be unable to perform tasks that (from the human's perspective) are straightforward. For example, pilots were often "surprised" when sophisticated flight control systems did not protect them against inadvertently breaking a common altitude rule (Sarter & Woods, 1997).

There has also been some study of trust within the framework of human-robot interaction or human-co-robot teams. Sadrfaridpour, Saeidi, Burke, Madathil, and Wang (2016) created a mathematical, time-series-based measure of trust that was based primarily on the difference in task performance speeds between the human and co-robot team members. The work was premised on the assumption that humans would trust a co-robotic team member more if the co-robot adjusted its speeds based on human performance variability (e.g., slowing down due to muscle fatigue)—that is, if the co-robot adapted to the human team member's performance. Yagoda and Gillan (2012) developed a scale to measure trust in human-robot interactions which addressed factors beyond the robotic system to include team processes (e.g., communication, coordination, situation awareness), team configuration (other non-robotic team members), and aspects of task and the situational context. Charalambous, Fletcher, and Webb (2016) developed a scale to assess trust between humans and robots specifically in industrial settings (not specific to co-robots), finding that trust included factors related to assessments of the reliability of the robot gripper, the speed and fluidity with which the robot moved, safety in interactions, and comfort with features such as size.

Research challenges related to trust include helping the human team member appropriately calibrate trust regarding the co-robot's ability to perform the work task, but particularly with respect to safety—so that the human team member clearly understands the performance boundaries of the co-robot. The human team member must trust the co-robot in order to effectively collaborate in close physical proximity (a goal of these technologies). Co-robots are intended to reliably sense human motion or touch and avoid movements which release enough energy to be harmful to the human team member. Co-robots also could proactively warn human team members if the human is encroaching in an unsafe space (either related to the co-robot or to other aspects of the manufacturing environment). However, it is unlikely that the co-robot would behave in a way to "save" the human team member from some other unsafe act (e.g., deliberately disabling safety guarding) or unexpected danger (e.g., an out of control fork truck), in the way, for instance, that a human co-worker might act to warn or intervene.

Coordination of Action and Task Allocation in Distributed Teams

The coordination of action is fundamental to the functioning of even the most basic teams. Hutchins (1995) writes that "all divisions of labor, whether that labor is physical or cognitive, require distributed cognition to coordinate the activities of the participants" (p. 176) and uses as an example the coordination needs of two laborers working together to drive in a rail tie spike. At a micro-level, close coordination of action in both time and space requires team members to understand both the requirements for, and effects of, their own actions, as well as those of others. Team members must communicate (implicitly or explicitly), have knowledge about the task goals and results of actions, and share awareness of the current state of the work product and each others' actions. At a macro-level, negotiating the coordination of work tasks involves the allocation and often the dynamic re-allocation of tasks among team members.

Models of function allocation, particularly as related to level of automation, may provide some guidance. A number of frameworks for describing levels of automation have been proposed within the human factors literature, and include levels ranging from fully automated, to partial levels in which the human operator may or may not be involved in the choice or action and/or informed about action execution, to fully manual (Parasuraman, Sheridan, & Wickens, 2000). The framework proposed by Parasuraman et al. (2000) further breaks down the description of levels of automation across four task components—information acquisition, information analysis, decision selection, and action implementation—thus allowing a more differentiated model where task elements are automated to different degrees. These models can assist in characterizing the level of autonomy taken by the co-robot team member in selecting and completing tasks, or in contributing to a shared work task. For example, high levels of human-co-robot collaboration would generally require the co-robots to acquire information (through sensors), make inferences (regarding the state of the work task and human team member), and implement actions. However, the degree to which the robot determines the choice and timing of actions (vs. looking to the human team member for direction) could vary based on the work task, human preference, or the joint experience of the human and co-robot working together.

Research in task allocation across human and co-robot team members has been conducted from the perspective of optimizing task scheduling. For example, Tsarouchi et al. (2017) described the use of algorithmic, automated methods for assigning sequential tasks to humans or co-robots in hybrid assembly cells, based on parameters such as capability, task execution times, waiting time, and resource availability in order to maximize throughput and utilization rates. In the example cases, the optimized task allocation reduced completion time by 78% from the manual task time and brought down human resource utilization to one task. As suggested by the researchers, this would allow the human operator to work on other processes or across multiple cells. Gombolay et al. (2017) cite progress in computational methods that support task allocation and sequencing within human-robot teams, but still recognize that coordinating these team activities safely and efficiently is "a challenging computational problem."

Task coordination, however, is more than the assignment of tasks to team members. Gombolay et al. (2017) note that computationally-oriented research has not generally considered how to incorporate support for situation awareness of the human team member (regarding task assignments), if work and task assignments are made automatically, and specifically studied that issue (as described previously). Task coordination also requires adapting performance to the capabilities and needs of team members, which can change within a work shift, as team members change (e.g., interacting with individuals who might be left vs. right handed), or even within the performance of a single task. The research by Sadrfaridpour et al. (2016) considered how co-robots might need to adapt performance on one parameter (speed) in response to changes in the human team member (e.g., due to fatigue) and investigated various levels of autonomy in making that change, including in response to a request from the human team member, automatically by the co-robot, or a mixed approach, where either team member could request/initiate the speed change. Other research in co-robotics is considering coordination in tasks such as object handover between humans and co-robots, through methods which address the initial positioning of the robot and development of acceptable movement trajectories by taking into account the human team member's visual field and ease of reach (Sisbot, Marin-Urias, Broquere, Sidobre, & Alami, 2010; Broquere et al., 2014).

Importantly, however, these models and methods of function or task allocation do not seem poised to characterize some of the complex coordination which might be envisioned for the highest levels of human-co-robot collaboration. For instance, consider the actions of two human team members in a different, complex domain. Hutchins (1995) describes a set of highly skilled and coordinated actions among two naval personnel as they jointly held, slid, and rotated a plotting tool on a navigational chart, without any verbal exchange or explicit division of labor—"they simply created this coordination in the doing of the task" (p. 220). Likewise, two people might coordinate action in a joint assembly task (say, where one person is holding or positioning two parts, while a co-worker is using a tool to join the parts together). The first worker may subtly shift the part position to make it easier for the second to reach, or may change the position and force of his grip to counteract or assist in the actions of the co-worker using the tool. This implicit coordination, made possible through experience in performing the task, and in predicting trouble, might be supplemented with explicit requests for action and responses (e.g., "can you move that a bit so I can get a better view"). How (or even whether) co-robotic team members can be trained to achieve the levels of shared awareness, communication, and self-awareness of their own role in such task coordination is an open question.

INTEGRATING CO-ROBOTS ON TEAMS: KEY MANUFACTURING CHALLENGES

The theoretical frameworks described in this chapter reveal open questions regarding the nature of how people and co-robots can communicate and coordinate activities as members of a distributed manufacturing team. As cited previously, researchers in the broader area of human-robot interaction, particularly those focused in the area of social robotics, have taken on challenges related to creating more natural interactions between people and robots in non-manufacturing environments, including communication through natural language, use and interpretation of nonverbal communication, classification of human movements, developing awareness through sensing and interpretation of human movement, and developing environmental and situational models that support common communication ground. For instance, in creating the robot bartender JAMES, researchers are developing methods to use human positions, movements, gestures, eye gaze, and facial expressions to allow JAMES to interpret customer states (e.g., attending to a menu, interacting with other customers, attending to the bartender) to correctly determine which customers are waiting for service and manage interaction across multiple customers (Foster, Gashler, & Giuliani, 2017). There is significant interest and research in the area of robots that provide social connections, health monitoring, and task assistance for older adults (Robinson, MacDonald, & Broadbent, 2014), which for some tasks may require close physical coordination (e.g., for bathing or dressing assistance). Other research is addressing the common ground problem in "everyday" communication with a robot by developing methods for the robot to sense and model objects in the shared physical environment combined with what the human can see and might gesture to (Lemaignan et al., 2012). However, implementing human-co-robot teams in manufacturing environments requires additional, systems-level considerations, including impacts on safety, workplace, and task design.

Impacts on Operator Safety and Risk Assessment

One of the first, and continuing, concerns with the implementation of co-robots in manufacturing environments is the safety of the operator. Industrial robotics has traditionally involved robots moving at high speed, but separated from the operator by a physical barrier, such as a cage and with emergency stop buttons. With humans and co-robots sharing a workspace, these methods are no longer sufficient for protecting the operator from contact with the robot. The payloads and end-effector designs (which can include sharp, hot, or otherwise hazardous tips) required for manufacturing tasks pose additional safety risks beyond those of home or everyday social environments. Current planning for collaboration requires a detailed assessment of the tasks being performed and the risk presented by those tasks, regardless of whether one is designing a new work cell or modifying an existing cell into a collaborative work cell (Marvel, Falco et al., 2015). As part of this, research on adaptive automation for general human-robot interaction has considered the potential for human operator complacence and shown benefits to regular adjustment of expected operator interaction in encouraging maintained engagement (Parasuraman, Mouloua, & Hilburn, 1999; Parasuraman & Hancock, 2008; Scerbo, 2018).

Risk assessment for human-co-robot operation should be proactive and involve identifying hazards, mitigating risks, and documenting the process. Marvel et al. (2015) and the National Institute of Standards and Technology (NIST) have laid out a methodology for defining risks for human-co-robot collaborative tasks. While a thorough risk assessment process should still consider the severity and likelihood of the hazard, as well as the possibility of avoidance, a new, but equally critical, aspect for collaborative operation is separating the safety of the robot itself from the safety of the robot system (Marvel, Falco et al., 2015). As described earlier in the chapter, the robot system includes the robot and the accompanying end-effector and part attached at the end of the robot. The advancements in communication, situation awareness, and coordination between the human operator and the robot are necessary components for a robot to be safe. However, even if a specific robot has been deemed safe for one application, the addition of a cutting tool as the end-effector can change the hazards and risk level for the robot system, requiring further safeguards to be put in place. The end-effector and part are dependent on the task being performed and, therefore, the risk assessment should be performed at the task element level and consider the processes, tools, and environment involved (Marvel, Falco et al., 2015).

Addressing the identified risk for human-co-robot teams and tasks remains based on the hierarchy of safety controls (Barnett & Brickman, 1986). At the top, as the most effective control, is elimination of the risk through removing the hazard. This includes implementation of inherently safe design measures, such as choosing appropriate robot system components for the task. The next level is engineering control, preventing a hazard through programming and other means to separate the operator from the dangerous component. The third level is administrative control, changing the exposure of the operator to the hazard through work instructions, administrative restrictions, and other policy changes. Along with, and within, these levels, Michalos et al. (2015) described three main strategies for different safety types that are unique to human-co-robot interaction. Crash safety focuses on limiting the force exerted on the human operator in the case of a collision. Active safety employs a series of proximity and contact sensors to detect a possible collision, and stops the co-robot operation. Adaptive safety identifies necessary corrective actions in the planned co-robot motion path to avoid a collision and not stop the operation. It will be important for human team members to understand these safety strategies—both through training and during operation—to maintain appropriate levels of trust and coordinate activities with the co-robot.

The Robotic Industries Association (RIA) and the International Organization for Standardization (ISO) Technical Specifications provide guidance for system designers on hazard identification and risk assessment methods and necessary risk mitigation approaches for the design of safe collaborative workspaces (ISO/TS 15066, 2016). Of note, these recommendations and the proposed standard began with the starting point of existing standards for industrial robots, and thus focus more heavily on safety and risk assessment. Separate standards, including ISO 13482:2014, focus on safety requirements for personal care robots and robots used as medical devices (Harper & Virk, 2010). The safety standards provide a starting point for integrating collaborative robots, however they fail to consider other guidelines for safety and effective design, such as the use of gestures and social cues to support shared situational awareness and trust between the operator and the co-robot. Standards should be used in project planning, system requirements specification, and requirements verification (Harper & Virk, 2010), but with awareness of the limitations as a minimum level required to support operator safety.

Giuliani et al. (2010) described three design principles to support safe human-robot interaction that apply to standard industrial robots and to co-robots. These principles are robustness, fast reaction time, and context awareness (Giuliani, Lenz et al., 2010). The system must be robust to errors, compensating for and correcting incorrect inputs. It must also be able to respond quickly to sensor information gathered on the environment and the operator. Frustration and distrust of the co-robot can result from the perceived non-responsiveness of a slow co-robot. Advanced communication between the operator and co-robot can support the human's perception of the reaction speed. According to Giuliani et al. (2010), the robot system should be designed based on the specific context so that the robot is able to interpret ambiguous input information. These design principles affect the means of speech and vision processing used to support safe interaction. Templates and patterns for speech communication and gestures that the robot recognizes may be needed to allow for faster reaction from the co-robot.

Impacts on Work Cell Design and Ergonomics

Work cells include the space in which the human-co-robot team operates, along with the co-robot, the application software, machine vision system for detecting the operator, the tool and end-effectors, and the power source. Current methods for designing a safe work cell begin with performing a risk assessment, as described previously, and defining the collaborative workspace. The motion paths are then designed to be within the barriers and space constraints. Then the safety requirements for the robot are determined and the appropriate robot is selected. Robot selection depends on the tasks that need to be performed and scaled to the payload. Michalos et al. (2015) identified the following categories for requirements in work cell design: robot type (single vs. dual arm), robot payload and power generating ability, part characteristics (both geometry and weight), and the assembly and manufacturing processes that will be used. Typical co-robot payloads are 1 to 5 kg to support small parts assembly. Inappropriate selection, such as choosing a co-robot with larger capabilities, increases the difficulty of working collaboratively and safely, since the forces required to stop the co-robot increase (Zanchettin, Ceriani et al., 2016). Robot selection needs to include tests of the application in response to collision and error. Error can be caused by the human operator, a sensor malfunction, or an external source. In order to maximize productivity, the human and co-robot need to be able to work at similar speeds. Human operator task performance typically ranges from 500 mm/s to 1.5 m/s (Kamali, Moodie et al., 1982). Co-robots also differ in terms of their reach, repeatability, and accuracy of parts pick up and placement. The positioning requirements, accuracy requirements, and space and weight requirements need to be clearly defined prior to implementation. These requirements depend on a detailed task analysis and assessment of function allocation between the human and co-robot. Task factors and end effect selection do not only affect task performance. As described earlier, Charalambous et al. (2016) showed that trust of the co-robot by the operator was dependent on end-effector reliability, speed and smoothness of motion, and comfort of the operator with the size of the robot and end-effector.

Additional work cell ergonomics to consider include the means of communication between the operator and machines, the clarity of controls, and the intuitiveness of the interface available to the operator, in order to support shared awareness and task coordination, as described previously. In addition to understanding the general means of operation, the operator must know the contact needed to operate or stop the robot, or it should be built into the mode of collaborative operation. Many systems are also designed under ideal conditions with peak operator performance, but considerations of operator stress and fatigue should be factored into the work cell design.

As the introduction of a collaborative work cell involves a substantial investment, researchers emphasize the use of simulation to evaluate safe and effective design of the robot system. One digital human modeling tool proposed by Maurice et al. (2016) allows for comparison of different co-robots based on ergonomic outcomes for varying tasks. Robot systems can be compared on a series of 28 ergonomic indicators, from operator balance to awkward postures. Peternel, Tsagarakis, Caldwell, and Ajoudani (2018) have proposed a muscle fatigue model where the robot takes over the task once the human reaches a pre-defined fatigue threshold. As an advanced collaborative robotics research effort, the robot learns the task from the cooperation with the human so that it is prepared to take over once fatigue develops. Pearce et al. (2018) demonstrated an optimization approach for task scheduling of human-co-robot teams that minimizes both ergonomic risk and task duration. They showed that for tasks with poor hand/wrist posture (an indicator of ergonomic risk) and those that were repetitive, the method performed well for minimizing the idle time for both human and co-robot. However, tasks with high forces or that required high precision were not well addressed with the method. This approach applied the Strain Index as the means of identifying ergonomic risk, limiting implementation to only ergonomic risk to the distal upper extremity. Further work is needed on methods modeling the joint objectives of maximizing performance while minimizing negative human operator outcomes (whether they are physical or cognitive) for a broader range of tasks and conditions, and incorporating the workspace in the determination. A challenge is that these optimization models are typically run pre-implementation, and thus cannot dynamically adjust task allocation. Research is also needed on the integration of safety and human factors considerations of human-co-robot teams into design tools for manufacturing environments (such as Siemens Jack or CATIA), so that the collaborative work environment and the impact on the manufacturing process can be simulated (Michalos et al., 2015). It is also critical to define metrics to adequately assess human-co-robot team performance and determine what outcomes can be used to measure good interaction (Marvel et al., 2015).

AREAS FOR FUTURE RESEARCH

Co-robots have increased the flexibility of manufacturing process design, but, as described throughout this chapter, new challenges and research questions for human-co-robot teaming have been introduced. Here we highlight a few directions for future research into human-co-robot teaming in manufacturing environments. In this chapter, and in most current research and implementations, the human-co-robot team being considered has focused on one human and one co-robot working in a common workspace at the same time. Important questions remain regarding the challenges of asynchronous and remote collaboration, along with interaction as multiple co-robots and humans engage in teamwork. In larger teams there are concerns for the roles of each member and how individual and team situation awareness can be developed and maintained. As mentioned previously, critical to the creation of true team situation awareness is the need to develop robust methods by which the co-robot can sense and appropriately understand human movement, infer activity, and predict future actions—without requiring explicit and continued input on the part of the human team member. Accompanying such methods is the need to explore learning and adaptability of the co-robot. For example, many industrial co-robots are hand-guided during the training process to program initial movement plans based on expected task and worker positioning. This limits the flexibility of implementation unless new algorithms are programmed. With co-robots that could sense and predict human team member action, the system could learn from process data to minimize the likelihood of interference and potentially prevent errors from occurring.

Accompanying these technological questions, a number of important research questions exist related to human-co-robot communication and trust. One is the need to provide easy-to-understand displays that allow the human operator to maintain awareness over the co-robot actions. Even with the current typically pre-planned co-robot movements the displays tend to be basic indicators. As the teams become more complex or co-robot movement less predictable, investigating the best means of communicating detailed state information will become more critical. This includes the need for methods of implicit communication and the development of co-robot movements and gestures that can be interpreted by the human operator. Communication of co-robot intent and movement will also impact the calibration of trust for the co-robot's ability and safety.

CONCLUSION

The introduction of collaborative robots into manufacturing environments raises a number of important issues regarding the design of robots and the work environment. Clearly, as with any new manufacturing technology, manufacturing system designers are considering aspects of safety, quality, productivity, and ergonomic risk related to the use of co-robots. To the extent that co-robots are intended to take on true team member roles, however, critical design challenges remain. By examining these challenges, humans will be able to develop appropriately calibrated trust in their co-robot teammates and both human and co-robot team members will be able to develop a shared awareness of each other and the manufacturing task, mutually interpret each other's actions and intentions, and seamlessly and safely coordinate action.

REFERENCES

Barber, B. (1983). *The logic and limits of trust*. New Brunswick, NJ: Rutgers University Press.

Barnett, R. L., & Brickman, D. B. (1986). Safety hierarchy. *Journal of Safety Research*, *17*, 49–55.

Broquere, X., Finzi, A., Mainprice, J., Rossi, S., Sidobre, D., & Staffa, M. (2014). An attentional approach to human-robot interactive manipulation. *International Journal of Social Robotics*, *6*, 533–553.

Busch, B., Grizou, J., Lopes, M., & Stulp, F. (2017). Learning legible motion from human-robot interactions. *International Journal of Social Robotics*, *9*, 765–779.

Charalambous, G., Fletcher, S., & Webb, P. (2015). Identifying the key organisational human factors for introducing human-robot collaboration in industry: an exploratory study. *The International Journal of Advanced Manufacturing Technology*, *81*, 2143–2155.

Charalambous, G., Fletcher, S., & Webb, P. (2016). The development of a scale to evaluate trust in industrial human-robot collaboration. *International Journal of Social Robotics*, *8*, 193–209.

Christiernin, L. (2017). How to describe interaction with a collaborative robot. In *HRI '17 companion, March 2017*. Vienna, Austria: ACM.

Chung, C. A. (1996). Human issues influencing the successful implementation of advanced manufacturing technology. *Journal of Engineering and Technology Management*, *13*, 283–299.

Clark, H. H., & Brennan, S. E. (1991). Grounding in communication. In L. B. Resnick, J. M. Levine & J. S. D. Teasley (Eds.), *Perspectives on socially shared cognition*. Washington, DC: American Psychological Association.

Endsley, M. (1995). Towards a theory of situation awareness in dynamic systems. *Human Factors*, *37*, 32–64.

Foster, M. E., Gashler, A., & Giuliani, M. (2017). Automatically classifying user engagement for dynamic multi-party human-robot interaction. *International Journal of Social Robotics*, *9*, 659–674.

Giuliani, M., Lenz, C., Müller, T., Rickert, M., & Knoll, A. (2010). Design principles for safety in human-robot interaction. *International Journal of Social Robotics*, *2*(3), 253–274.

Gombolay, M., Bair, A., Huang, C., & Shah, J. (2017). Computational design of mixed-initiative human-robot teaming that considers human factors: situational awareness, workload, and workflow preferences. *International Journal of Robotics Research*, *36*, 597–617.

Harper, C., & Virk, G. (2010). Towards the development of international safety standards for human robot interaction. *International Journal of Social Robotics*, *2*, 229–234.

Hayes, B., & Scassellati, B. (2013). Challenges in shared-environment human-robot collaboration. *Learning*, *8*(9).

Heath, C., & Luff, P. (1991). Collaborative activity and technology design: Task coordination in London underground control rooms. In *Proceedings of the 2nd European conference on computer-support cooperative work, September, 1991*. Springer, Dordrecht.

Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.

International Organization for Standardization. (2016). *ISO/TS 15066:2016*. Retrieved from https://www.iso.org/standard/62996.html

Ivaldi, S., Fritzsche, L., Babič, J., Stulp, F., Damsgaard, M., Graimann, B., Bellusci, G., & Nori, F. (2017). *Anticipatory models of human movements and dynamics: The roadmap of the AnDy project*. Digital Human Models (DHM).

Jian, J. Y., Bisantz, A. M., & Drury, C. G. (2000). Foundations for an empirically determined scale of trust in automated systems. *International Journal of Cognitive Ergonomics*, *4*(1), 53–71.

Johannsmeier, L., & Haddadin, S. (2017). A hierarchical human-robot interaction-planning framework for task allocation in collaborative industrial assembly processes. *IEEE Robotics and Automation Letters*, *2*, 41–48.

Kamali, J., Moodie, C. L., & Salvendy, G. (1982). A framework for integrated assembly systems: Humans, automation and robots. *International Journal of Production Research*, *20*, 431–448.

Komatsu, T., Kurosawa, R., & Yamada, S. (2012). How does the difference between users' expectations and perceptions about a robotic agent affect their behavior? *International Journal of Social Robotics*, *3*, 109–116.

Kozlowski, S. W. J., Grand, J. A., Baard, S. K., & Pearce, M. (2015). Teams, teamwork, and team effectiveness: Implications for human systems integration. In D. Boehm-Davis, F. Durso, & J. D. Lee (Eds.), *APA handbook of human systems integration* (pp. 555–571). Washington, DC: American Psychological Association.

Lasota, P. A., & Shah, J. A. (2015). Analyzing the effects of human-aware motion planning on close-proximity human—robot collaboration. *Human Factors*, *57*, 21–33.

Lee, J. D., & Moray, N. (1994). Trust, self-confidence, and operators' adaptation automation. *International Journal of Human Computer Studies*, *40*, 153–184.

Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, *46*(1), 50–80.

Lemaingnan, S., Ros, R., Sisbot, E., Alami, R., & Beetz, M. (2012). Grounding the interaction: Anchoring situated discourse in everyday human-robot interaction. *International Journal of Social Robotics*, *4*, 181–199.

Li, J., & Chignell, M. (2011). Communication of emotion in social robots through simple head and arm movements. *International Journal of Social Robotics*, *3*(2), 125–142.

Marvel, J. A., Falco, J., & Marstio, I. (2015). Characterizing task-based human—robot collaboration safety in manufacturing. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, *45*, 260–275.

Maurice, P., Padois, V., Measson, Y., & Bidaud, P.(2016, June). A digital human tool for guiding the ergonomic design of collaborative robots. In *4th International Digital Human Modeling Symposium (DHM2016)*. Montreal, Canada.

Meyer, J. (2004). Conceptual issues in the study of dynamic hazard warnings. *Human Factors*, *46*(2), 196–204.

Michalos, G., Makris, S., Tsarouchi, P., Guasch, T., Kontovrakis, D., & Chryssolouris, G. (2015). Design considerations for safe human-robot collaborative workplaces. *Procedia CIrP*, *37*, 248–253.

Morrow, P., & Fiore, S. (2013). Team cognition: Coordination across individuals and machines. In J. D. Lee & A. Kirlik (Eds.), *The Oxford handbook of cognitive engineering* (pp. 200–215). New York: Oxford University Press.

Muir, B. M. (1994). Trust in automation: Part I. Theoretical issues in the study of trust and human intervention in automated systems. *Ergonomics*, *37*, 1905–1922.

Muri, B. M., & Moray, N. (1996). Trust in automation: Part II. Experimental studies of trust and human intervention in a process control simulation. *Ergonomics*, *39*(3), 429–460.

Mutlu, B., Yamaika, Y., Janda, T., Ishiguro, H., & Hagita, N. (2009). Nonverbal leakage in robots: Communication of intentions through seemingly unintentional behavior. In *Proceedings of HRI '09, March 11–13*. La Jolla, CA: ACM.

Muttray, A. et al. (2014). *Collaborative robots—Investigation of pain sensibility at the man-machine interface*. Institute for Occupational, Social and Environmental Medicine for the Johannes Gutenberg University of Mainz.

Mwangi, E., Barakova, E., Diaz-Boladeras, M., Mollofre, A., & Roauterbert, M. (2018). Directing attention through gaze hints improves task solving in human-humaniod interaction. *International Journal of Social Robotics*, *10*, 343–355.

Novikova, J., & Watts, L. (2015). Towards artificial emotions to assist social coordination in HRI. *International Journal of Social Robotics*, *7*, 77–88.

Parasuraman, R., & Hancock, P. A. (2008). Mitigating the adverse effects of workload, stress, and fatigue with adaptive automation. In P. A. Hancock & J. L. Szalma (Eds.), *Performance under stress* (pp. 61–74). Burlington, VT: Ashgate Publishing.

Parasuraman, R., & Manzey, D. (2010). Complacency and bias in human use of automation: An attentional integration. *Human Factors*, *52*, 381–410.

Parasuraman, R., Mouloua, M., & Hilburn, B. (1999). Adaptive aiding and adaptive task allocation enhance human-machine interaction. In *Automation technology and human performance: Current research and trends* (pp. 119–123). Mahwah, NJ: Lawrence Erlbaum Associates.

Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, *39*(2), 230–253.

Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, *30*(3), 286–297.

Pearce, M., et al. (2018). Optimizing makespan and ergonomics in integrating collaborative robots into manufacturing processes. *IEEE Transactions on Automation Science and Engineering*, *15*(4), 1772–1784.

Peternel, L., Tsagarakis, N., Caldwell, D., & Ajoudani, A. (2018). Robot adaptation to human physical fatigue in human—robot co-manipulation. *Autonomous Robots*, 1–11.

Rempel, J. K., Holmes, J. G., & Zanna, M. P. (1985). Trust in close relationships. *Journal of Personality and Social Psychology*, *49*(1), 95–112.

Riley, J. M., & Endsley, M. R. (2004). The hunt for situation awareness: Human-robot interaction in search and rescue. In *Proceedings of the Human Factors and Ergonomics 2004 annual meeting*. Human Factors and Ergonomics Society. Los Angeles, CA: SAGE Publications.

Riley, J. M., & Endsley, M. R. (2005). Situation awareness in HRI with collaborating remotely piloted vehicles. In *Proceedings of the Human Factors and Ergonomics 2005 annual meeting*. Human Factors and Ergonomics Society. Los Angeles, CA: SAGE Publications.

Riley, J. M., Strater, L. D., Chappell, S. L., Connors, E. S., & Endsley, M. R. (2012). Situation awareness in human robot interaction: Challenges and user interface requirements. In F. Jenstsch & M. Barnes (Eds.), *Human-robot interactions in future military operations*. New York, NY: Routledge.

Robinson, H., MacDonald, B., & Broadbent, E. (2014). The role of healthcare robots for older people at home: A review. *International Journal of Social Robotics*, *6*, 575–591.

Robotics Industries Association. (2016). *RIA TR R15.606–2016 for robots & robotic devices—Collaborative robots*. Retrieved from https://www.robotics.org/robotics/technical-report-ria-tr-r15-606-2016-for-robots-and-robotic-devices-collaborative-robots

Sadrfaridpour, B., Saeidi, H., Burke, J., Madathil, K., & Wang, Y. (2016). Modeling and control of trust in human-robot collaborative manufacturing. In *Robust intelligence and trust in autonomous systems* (pp. 115–141). Boston, MA: Springer.

Salas, E., Prince, C., Baker, D. P., & Shrestha, L. (1995). Situation awareness in team performance: Implications for measurement and training. *Human Factors*, *37*(1), 123–126.

Sarter, N., & Woods, D. D. (1997). Team play with a powerful and independent agent: Operational experiences and automation surprises on the airbus A-320. *Human Factors*, *39*, 553–569.

Sauppe, A., & Mutlu, B. (2015). The social impact of a robot co-worker in industrial settings. In *Proceedings of CHI 2015* (pp. 3613–3622). Seoul, Korea: ACM.

Scerbo, M. W. (2018). Theoretical perspectives on adaptive automation. In *Automation and human performance* (pp. 57–84). Abingdon: Routledge.

Seong, Y., & Bisantz, A. M. (2008). The impact of cognitive feedback on judgment performance and trust with decision aids. *International Journal of Industrial Ergonomics*, 608–625.

Sheridan, T. B. (1988). Trustworthiness of command and control systems. In *IFAC Man-machine systems: Selected papers from the Third IFAC/IFIP/IEA/IFORS conference* (pp. 427–431). Oulu, Finland, 14–16 June 1988 .

Sisbot, E. A., Marin-Urias, L. M., Broquere, X., Sidobre, D., & Alami, R. (2010). Synthesizing robot motions adapted to human presence. *International Journal of Social Robotics*, *2*, 239–343.

Suchman, L. (1987). *Plans and situation action*. Cambridge: Cambridge University Press.

Tsarouchi, P., Matthaiakis, A., Makris, S., & Chryssolouris, G. (2017). On a human-robot collaboration in an assembly cell. *International Journal of Computer Integrated Manufacturing*, *30*, 580–589.

Unhelkar, V. V., et al. (2018). Human-aware robotic assistant for collaborative assembly: Integrating human motion prediction with planning in time. *IEEE Robotics and Automation Letters*, *3*, 2394–2401.

Wickens, C. D., Hollands, J. G., Banbury, S., & Parasuraman, R. (2013). *Engineering psychology and human performance* (4th ed.). New York: Routledge Publishing.

Yagoda, R., & Gillan, D. J. (2012). You want me to trust a ROBOT? The development of a human-robot interaction trust scale. *International Journal of Social Robotics*, *4*, 235–248.

Zanchettin, A. M., et al. (2016). Safety in human-robot collaborative manufacturing environments: Metrics and control. *IEEE Transactions on Automation Science and Engineering*, *13*, 882–893.