

EXAMINING THE EFFECTS OF COGNITIVE ASSISTIVE AGENTS ON TEAM COORDINATION IN MANUFACTURING TEAMS

Begerowski, S.R., Waldherr, F., Biddlecom, J., Traylor, A., Krugh, M., Mears, L., Shuffler, M.L.
Clemson University

This article details the motivation and design of an experiment to investigate the effects of artificially intelligent cognitive assistive agents on coordination efforts in manufacturing teams. As automation solutions become more accessible and products rapidly grow in complexity, there are significant calls to leverage abilities of both artificial agents and human workers to maximize team functioning and product output. As such, we propose an experimental design where we introduce a cognitive agent with two levels of autonomy (low, and high) into a team of participants during an assembly task. We hypothesized that cognitive assistive technologies would enhance coordination within assembly teams, leading to higher productivity and reduced errors, with initial data suggesting trends in support of these hypotheses. We seek to demonstrate the value of cognitive agents in augmenting human workers, allowing manufacturers to see the benefit of increased productivity while retaining value and relevance of human labor in the face of technological development.

INTRODUCTION

The field of manufacturing has seen fairly consistent advancement throughout history, a trend punctuated by particularly disruptive developments. The mechanization of production was followed by electrification and, in turn, by automation. At present, manufacturing facilities benefit from the large-scale use of data in connecting systems to improve their capability and their responsiveness to changes throughout the production process. Throughout all of these changes, however, the presence of a human workforce in manufacturing has persisted. Assembly tasks continue to see a large degree of human involvement, as high levels of variability and complexity often render automation efforts prohibitively expensive and highlight limitations of their capability. This has driven a persistent interest in improving the abilities of human workers. Consequently, developing systems where human workers work collaboratively with artificially intelligent agents is often regarded to hold critical importance in further developing manufacturing.

Although the mechanisms by which human teams form and operate have been the subject of a large body of research, how they relate to human-autonomy teams (HATs) is less studied. Accordingly, this research investigates the effect of an artificially intelligent cognitive assistive agent with varying levels of capability on a team's coordination when performing a manufacturing task, and what effect that has on the team's productivity.

The purpose of this paper is twofold. First, we provide a summary of relevant concepts. This topic is, by nature, interdisciplinary. As such, it bears inclusion of an overview of concepts critical to understanding human-autonomy teams. Second, we detail an experiment in which teams of two humans and one agent must engage in an assembly task, varying the level of agent autonomy (i.e., high, moderate, low).

BACKGROUND

Artificial Intelligence

The phrase "Artificial Intelligence" (AI) originated with the 1956 Dartmouth Summer Research Program, to

describe a machine capable of thought: that is, a machine which is able to make decisions without having received specific instruction regarding those decisions. Since 1956, work in AI has been essentially continuous, though advancement in the field has seen a variety of upsurges and downturns. At present, AI research is undergoing an "explosion" driven by improvements in information technology and big data (Wang et al., 2021), which is in turn driving efforts to support the practical implementation of artificially intelligent systems in multiple domains. To date, excepting a relatively small number of applications including predictive maintenance and quality control (Demlehner et al., 2021), this implementation has not been achieved on an appreciable scale in manufacturing processes. This is attributed partially to difficulties finding usable data (Peres et al., 2020) and a lack of collaboration (Trakadas et al., 2020).

The implementation of artificially intelligent systems in manufacturing settings can take several different forms: namely, as a physical or cognitive agent, differentiated by their interactions with human workers. Physical agents interact with their human counterparts through physical means (e.g., a collaborative robot positioning parts or supplying humans with necessary components). By contrast, a cognitive agent facilitates communication and the sharing of information (e.g., offering adaptive instructions in response to humans' actions). In either case, their use contributes to a more collaborative environment than does traditional automation, a point of significant interest in industry (Trakadas et al., 2020).

Beyond the form of an artificial agent, consideration of its capabilities is important. A simple cognitive agent may, for example, be capable of detecting and marking inconsistencies in images recorded for quality control. Another may detect specific types of errors during the manufacturing process and communicate the need to correct them to a human operator. O'Neill and colleagues (2020) define a scale characterizing these levels of autonomy. On the low end, a system only responds to decisions made and input by a human operator. A case with partial agent autonomy generates alternatives and presents these to the human to make the decision, and a high autonomy agent makes the decision itself, perhaps not even involving a human in the process.

Behavioral Processes in HATs

The advancement of artificial intelligence has sparked interest in the transition of an agent used as a teammate rather than a tool, particularly in manufacturing environments, in order to meet increasingly complex production needs. Teams are often conceptualized as two or more members working interdependently towards a common goal (Salas et al., 1992). In recent years, this notion of teaming has expanded to include at least one artificial agent, evolving into the distinct subfield of Human-Autonomy teaming. Researchers are exploring what constitutes an artificial teammate and how this affects affective, behavioral, and cognitive states of teams (see O'Neill et al., 2020 for a review). While affective states have garnered much attention (Glikson & Woolley, 2020; Hancock et al., 2011), behavioral processes (e.g., coordination) and cognitive states (e.g., team cognition) remain understudied in HATs (O'Neill et al., 2020).

Behavioral processes in teams are vital to team performance (Salas et al., 2015). One key process in teams is coordination: “the enactment of behavioral mechanisms necessary to perform a task and transform resources into outcomes” (Sims & Salas, 2007). Coordination is a main driver of team effectiveness (Rousseau et al., 2006; Stewart, 2006), and manifests in various forms (e.g., implicit, explicit). Explicit coordination reflects intentional efforts of team members to address interdependencies between roles and tasks, using mechanisms such as direct communication. Conversely, implicit coordination refers to the management of interdependencies through a common ground that the team achieves by adjusting behaviors without instruction (Rico et al., 2008). HATs tend to rely on explicit coordination due to increased workloads and need for role clarity (Fiore et al., 2005). However, there are consistent calls to better understand implicit coordination in HATs (McNeese et al., 2018).

Implicit Coordination

Metacognitive activities of team members enable implicit coordination. Transactive memory systems (TMSs) are a key part of metacognitive activities in teams. Namely, TMSs enable teams to leverage a shared understanding of what should be done, when, and by whom, forthcoming needs and behaviors can be anticipated and acted upon (Fiore et al., 2005; Klimoski & Mohammed, 1994). TMSs are multidimensional in nature, comprised of three dimensions: specialization (e.g., team members hold various specialties), credibility (e.g., team members can rely on others for said specialties), and general coordination processes (e.g., team members can leverage others' specialties and credibility to enhance intra-team behavioral processes; Lewis, 2003).

This underscores the two critical factors of implicit coordination: anticipation and dynamic adjustment. When coordinating implicitly, a team must anticipate actions and needs of teammates and task demands as well as dynamically adjust behavior to meet these needs without communicating or planning directly (Cannon-Bowers et al., 1993; Espinosa et al., 2004).

Implicit coordination notably leads to improved team performance beyond explicit coordination, particularly when

teams are actively engaging in tasks (Rico et al., 2008). Despite this, current use of autonomous agents often depends on explicit coordination. In a manufacturing context, agents often send parts to human teammates after a request is made, representing explicit coordination in the team. Conversely, there are industry calls to explore applications of highly autonomous agents. As such, the goal of this experiment is to test how varying levels of autonomy affect the performance of manufacturing teams.

Given that high autonomy agents can engage in behaviors consistent with implicit coordination (e.g., anticipation of team needs), and that implicit coordination is related to higher performance, we hypothesize that:

H1: Teams in the high autonomy condition will perform better than those in the low condition such that these teams will take less time to accurately complete the task.

Moreover, highly autonomous agents exhibit characteristics across the three dimensions of TMSs. Such agents can demonstrate mastery of a function (i.e., specialization), perform with high accuracy (i.e., credibility), and demonstrate both physical and cognitive behaviors to aid in progression towards team goals (i.e., general coordination processes). That is, highly autonomous agents can exhibit factors consistent with human teammate functioning, resulting in our hypothesis that:

H2: Teams in the high autonomy condition will report higher perceived levels of implicit coordination than the low condition.

METHODOLOGY

Participants

Participants are undergraduate students enrolled in psychology courses at a public university in the Southeastern United States. They are recruited through Sona Systems and offered participatory research credits in line with course requirements upon completion of the study. Each team includes two participants in addition to one artificial agent.

Design

Teams of two humans and one assistive agent will complete an assembly task of building three carts using Funphix building toys. One cart is divided into three sub-assemblies: upper frame, lower frame, and body. Teams are instructed to build three carts sequentially, so they must complete one full build before moving onto the next cart.

The role of the agent is to send kits to participants from the warehouse area to the build area. The warehouse area is dedicated storage space for all parts and houses the agent whereas the build area is designated space for participants to complete the assembly task – resembling a physical setup commonly found in manufacturing plants (Figure 1). Level of autonomy is manipulated across two levels (i.e., low, high).




Figure 1: Experiment Set-Up. Image on the top shows the warehouse area that houses the agent and Funphix materials. Image on the bottom shows the end of the conveyor in the build area, where participants can retrieve parts sent by the agent.

Procedure

Teams in both conditions engage in an assembly task: building three carts. After completing informed consent protocols, the participants are trained to familiarize themselves with Funphix materials and agent interaction. During training, participants practice creating a frame using Funphix tubes as well as become familiar with agents used in the study during training. All kits are sent down a conveyor, with a physical agent (e.g., robotic arm) picking and placing the kit onto the line. This process is demonstrated to participants during training so that they are aware of the contributions made by the assistive agent. Participants are then shown the cameras that function as the visual system for the cognitive agent, which are used for the systems monitoring functionality of the cognitive agent.

Once training is complete, participants receive instructions for their condition as well as images of the final product to aid building. When the team finishes their builds, the task is complete, and the participants engage in a post-survey.

Manipulations and Measures

Levels of Autonomy. Teams are randomly assigned to one of two conditions: low and high autonomy. The low autonomy

condition consists of human participants and an agent with only physical capabilities. The high autonomy condition is manipulated using a “Wizard of Oz” approach in which a trained experimenter performs system monitoring behaviors from another room. In the high autonomy condition, the “agent” monitors progress. As the team concludes building the current kit, the agent sends the next kit and only informs the team after it was sent.

Team Roles. Assembly teams are typically composed of individuals with general core competencies as well as specialized roles. To model this phenomenon, we imposed team roles for the human participants: the blue role and the red role. While all participants can engage with yellow, green, and black parts, only the blue role can use blue parts while only the red role can use red parts. The agent specializes in transportation of parts, as well as selection of parts in the high autonomy condition. In addition to modeling specialization in assembly teams, these roles encourage collaboration during building as the kits will come in assorted colors.

Transactive Memory Systems. TMSs are assessed using a 5-point Likert scale that contains items across the three dimensions (e.g., specialization: “Different team members are responsible for expertise in different areas.”; credibility: “I was comfortable accepting procedural suggestions from other team members”; general coordination: “Our team worked together in a well-coordinated fashion”; Lewis, 2003).

Implicit Coordination. Implicit coordination is assessed through a self-report measure by human participants. Participants are asked about implicit coordination efforts across four items using a 7-point Likert scale (e.g., “my team proactively helped individual members when they needed assistance” Fisher et al., 2012).

Team Efficiency. Coordination is frequently indirectly measured through time-based metrics (Harbers et al., 2011; O’Neill et al., 2020). As such, we also measure team efficiency by looking at the time taken to accurately complete the task in addition to the number of errors a team makes.

Materials

Funphix. Funphix building toys are used to build the carts and frames, consisting of tubes (long and short), connectors (three-way, four-way, T-style), wheels, panels, screws, and screw keys. Parts can be black, blue, green, yellow, or red. The parts provided are separated into three kits: upper frame, lower frame, and body. Each kit contains only the relevant parts for that portion of the build. All kits are created prior to the experiment to ensure the same kits are provided across sessions.

Assistive agent. To model assembly processes and restrict participants’ view of the human experimenter selecting kits, a roller conveyor is used to send kits to participants from a warehouse area. Once the correct kit is selected and placed in a staging area by the experimenter, a Universal Robots UR16 collaborative robot manipulates the kit, placing it on the conveyor.

RESULTS

Though full data collection is underway, initial results provide insight on emerging trends. The following discusses preliminary trends of the data in a descriptive manner, consisting of data from seven teams ($N_{\text{High}} = 4$; $N_{\text{Low}} = 3$).

Consistent with H1, teams in the high autonomy condition appear to perform better (Figure 2). That is, teams exhibit faster performance time (i.e., minutes to complete task) and better accuracy (i.e., make less errors).




Figure 2. Average Performance Indicators by Autonomy Condition

Initial trends also appear to be consistent with H2. That is, teams in the high autonomy condition are reporting higher levels of perceived implicit coordination (Figure 3).




Figure 3. Average Self-Reported Implicit Coordination by Autonomy Condition

To better understand these emerging trends, we also assessed the three facets of TMSs by autonomy condition (Figure 4). Interestingly, it appears that perceived general coordination and credibility may not differ by condition; however, specialization might.




Figure 4. Average Transactive Memory System Facets by Autonomy Condition

DISCUSSION

Initial results suggest that, like in human-human teams, implicit coordination may in fact improve performance of human-autonomy teams. The “agent” with high autonomy sends the parts as needed, reducing the need for participants to multitask between building, performance monitoring, and part selection. We propose this directly reduces the overall time to complete the building task, while also allowing for further focus on building in order to improve accuracy. This is of particular importance for industries such as manufacturing, in which both build time and accuracy function as key performance indicators to build a competitive edge.

While there is a difference in self-reported perceived implicit coordination, with teams in the high autonomy condition reporting higher averages, the difference is small. Given this, we conducted exploratory descriptive analyses to better understand factors at play. Though credibility and coordination seem to remain consistent, specialization is notably higher for teams in the high autonomy condition. Perhaps, it is the notion of more teamwork behaviors (e.g., systems monitoring, backup behaviors) that results in higher perceptions of specialization. That is, the agent may appear to hold unique capabilities that other teammates do not, in this case part selection and transportation, resulting in a specialized contribution to the team. Future research should focus on disentangling these perceptions across the various dimensions in addition to relationships with performance, to address both academic and industry calls for expanding beyond explicit coordination in HATs.

IMPLICATIONS

The results of this experiment promise to hold relevance in a number of different contexts: both in a purely academic sense, and as relating to continuing efforts to develop AI and manufacturing systems well suited for use with one another. Taking a strict academic view, this work contributes to building a body of knowledge regarding human teaming with artificial agents: specifically, how the sort of coordination commonly found in effective human teams develops in the presence of an artificially intelligent agent. Knowledge about the formation of teams can also inform how their creation is

facilitated, both in industrial settings and in continued experimentation. Furthermore, this information can influence the direction pursued in the development of real artificial agents (in contrast to the simulated agent used in this setting) by clarifying which approaches provide a benefit to manufacturing processes.

REFERENCES

- Cannon-Bowers, J., Salas, E., & Converse, S. (1993). Shared Mental Models in Expert Team Decision Making. In N. J. Castellan (Ed.), *Individual and Group Decision Making* (1st ed., p. 26). Psychology Press.
- Demlehner, Q., Schoemer, D., & Laumer, S. (2021). How can artificial intelligence enhance car manufacturing? A Delphi study-based identification and assessment of general use cases. *International Journal of Information Management*, 58, 102317. <https://doi.org/10.1016/j.ijinfomgt.2021.102317>
- Espinosa, J. A., Lerch, F. J., & Kraut, R. E. (2004). Explicit versus implicit coordination mechanisms and task dependencies: One size does not fit all. In *Team cognition: Understanding the factors that drive process and performance*. (pp. 107–129). American Psychological Association. <https://doi.org/10.1037/10690-006>
- Fiore, S. M., Jentsch, F., Becerra-Fernandez, I., Salas, E., & Finkelstein, N. (2005). Integrating Field Data with Laboratory Training Research to Improve the Understanding of Expert Human-Agent Teamwork. *Proceedings of the 38th Annual Hawaii International Conference on System Sciences*. <https://doi.org/10.1109/HICSS.2005.327>
- Fisher, D. M., Bell, S. T., Dierdorff, E. C., & Belohlav, J. A. (2012). Facet personality and surface-level diversity as team mental model antecedents: Implications for implicit coordination. *Journal of Applied Psychology*, 97(4), 825–841. <https://doi.org/10.1037/a0027851>
- Glikson, E., & Woolley, A. W. (2020). Human Trust in Artificial Intelligence: Review of Empirical Research. *Academy of Management Annals*, 14(2). <https://doi.org/10.5465/annals.2018.0057>
- Hancock, P. A., Billings, D. R., Schaefer, K. E., Chen, J. Y. C., de Visser, E. J., & Parasuraman, R. (2011). A Meta-Analysis of Factors Affecting Trust in Human-Robot Interaction. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 53(5), 517–527. <https://doi.org/10.1177/0018720811417254>
- Harbers, M., Bradshaw, J. M., Johnson, M., Feltovich, P., van den Bosch, K., & Meyer, J.-J. (2011). Explanation and Coordination in Human-Agent Teams: A Study in the BW4T Testbed. *2011 IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology*, 17–20. <https://doi.org/10.1109/WI-IAT.2011.83>
- Klimoski, R., & Mohammed, S. (1994). Team mental model: construct or metaphor? *Journal of Management*, 20(2). [https://doi.org/10.1016/0149-2063\(94\)90021-3](https://doi.org/10.1016/0149-2063(94)90021-3)
- Lewis, K. (2003). Measuring transactive memory systems in the field: Scale development and validation. *Journal of Applied Psychology*, 88(4). <https://doi.org/10.1037/0021-9010.88.4.587>
- McNeese, N. J., Demir, M., Cooke, N. J., & Myers, C. (2018). Teaming With a Synthetic Teammate: Insights into Human-Autonomy Teaming. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 60(2), 262–273. <https://doi.org/10.1177/0018720817743223>
- O'Neill, T., McNeese, N., Barron, A., & Schelble, B. (2020). Human–Autonomy Teaming: A Review and Analysis of the Empirical Literature. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 001872082096086. <https://doi.org/10.1177/0018720820960865>
- Peres, R. S., Jia, X., Lee, J., Sun, K., Colombo, A. W., & Barata, J. (2020). Industrial Artificial Intelligence in Industry 4.0 - Systematic Review, Challenges and Outlook. *IEEE Access*, 8, 220121–220139. <https://doi.org/10.1109/ACCESS.2020.3042874>
- Rico, R., Sánchez-Manzanares, M., Gil, F., & Gibson, C. (2008). Team Implicit Coordination Processes: A Team Knowledge–Based Approach. *Academy of Management Review*, 33(1). <https://doi.org/10.5465/amr.2008.27751276>
- Rousseau, V., Aubé, C., & Savoie, A. (2006). Teamwork Behaviors. *Small Group Research*, 37(5). <https://doi.org/10.1177/1046496406293125>
- Salas, E., Dickinson, T. L., Converse, S. A., & Tannenbaum, S. I. (1992). Toward an understanding of team performance and training. In R. W. Swezey & E. Salas (Eds.), *Teams: Their training and performance* (pp. 3–29). Ablex Publishing.
- Salas, E., Shuffler, M. L., Thayer, A. L., Bedwell, W. L., & Lazzara, E. H. (2015). Understanding and Improving Teamwork in Organizations: A Scientifically Based Practical Guide. *Human Resource Management*, 54(4). <https://doi.org/10.1002/hrm.21628>
- Sims, D. E., & Salas, E. (2007). When teams fail in organizations: What creates teamwork breakdowns? In J. Langan-Fox, C. L. Cooper, & R. J. Klimoski (Eds.), *Research companion to the dysfunctional workplace: Management challenges and symptoms* (pp. 302–318). Edward Elgar.
- Stewart, G. L. (2006). A Meta-Analytic Review of Relationships Between Team Design Features and Team Performance. *Journal of Management*, 32(1). <https://doi.org/10.1177/0149206305277792>
- Trakadas, P., Simoens, P., Gkonis, P., Sarakis, L., et al. (2020). An Artificial Intelligence-Based Collaboration Approach in Industrial IoT Manufacturing: Key Concepts, Architectural Extensions and Potential Applications. *Sensors*, 20(19), 5480. <https://doi.org/10.3390/s20195480>
- Wang, L., Liu, Z., Liu, A., & Tao, F. (2021). Artificial intelligence in product lifecycle management. *The International Journal of Advanced Manufacturing Technology*, 114(3–4), 771–796. <https://doi.org/10.1007/s00170-021-06882-1>