

# Toward a Taxonomy of Teamwork Support to Help Inform Collaborative AI Efforts

Matthew Johnson and Jeffrey M. Bradshaw

The Florida Institute for Human and Machine Cognition  
mjohnson@ihmc.org and jbradshaw@ihmc.org

## Abstract

There are two primary impediments to progress in work on collaborative AI. The first is definitional; the second is implementational. We will discuss both challenges—hoping that it will help advance the conversation among practitioners about how we might move together in a more coherent fashion toward the goals of collaborative AI.

## Introduction

There are two primary impediments to progress in work on collaborative AI. The first is definitional; the second is implementational. We will discuss both challenges—hoping that it will help advance the conversation among practitioners about how we might move together in a more coherent fashion toward the goals of collaborative AI. Here, we will not attempt to justify the need for collaborative AI but will assume that readers already believe that to achieve AI’s ambitious goals—at least for the foreseeable future—collaboration among hybrid groups of people and various software and hardware entities is essential if we are to achieve consequential ends in a safe and effective manner.

## The Definitional Challenge: What Is Teamwork?

We are convinced that work to understand and measure the effectiveness of various forms of collaborative AI will benefit from agreement among researchers and practitioners on a common set of goals and definitions. However, current attempts to characterize the domain of collaborative AI are complex and unwieldy. Multiple disciplines throw different concepts and constructs at the problem, varied domains and techniques confound the conversation, and there is no agreed-to framework to interpret different results and assemble these interrelated moving parts into a coherent system. We propose *interdependence* as the common factor that unifies and explains these moving parts and undergirds the

different terms people use to talk about them. We argue that teamwork and collaboration are simply managing interdependence within joint activity. Elsewhere, we have described in detail a sound and practical theoretical framework based on interdependence that enables researchers to predict and explain experimental results in terms of interlocking relationships among well-defined operational principles (Johnson & Bradshaw, 2021).

Our central interest in this paper is *teamwork*, a challenging and productive form of collaboration that we distinguish from *independent activity* or other kinds of *joint activity* as follows (see also figure 1):

- *Independent activity*. Clearly, some work is independent. Those who engage in independent work have no intention of working with others. Examples include going on a hike by oneself or writing a paper alone.
- *Competitive activity*. The adversaries involved in competitive sports activities, though playing against each other, are—clearly interdependent. For this reason, it is appropriate to call their work “joint activity.” However, the competitors do not intend to work together cooperatively but rather are consciously trying to work against each other to produce a win-lose outcome selfishly favoring only their own team.
- *Social activity*. Although members of a running club intend to practice their sport at the same time and place for social reasons, the work itself is typically done more or less independently by each runner. While club members may provide support, tips, and encouragement for one another, the skill of running ultimately cannot be improved vicariously for someone else.
- *Norm-governed cooperation*. Although drivers on a busy road are in a state of interdependence and may cooperate with each other in ways that are mutually beneficial to other drivers, in most cases they do not share a common destination. In ordinary situations, they rely for the most part on traffic laws and norms (e.g., driving courteously) to manage their interdependence with other drivers.

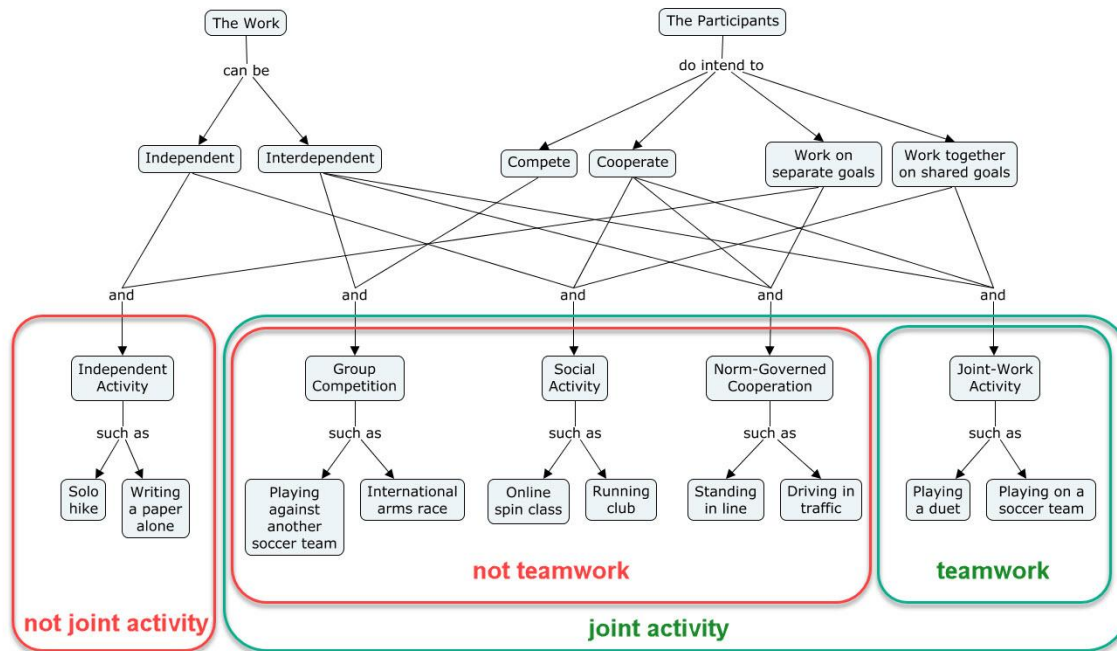


Figure 1. Some examples of activity that are not teamwork and some that are.

- *Teamwork.* Musicians playing a duet and members of the same soccer team consciously intend to work together and are clearly in an interdependent state as they jointly pursue common rather than competing goals. This sort of joint-work activity can be rightfully called “teamwork” in the sense that other kinds of joint activity cannot and is more in line with the goals of collaborative AI—including collaboration among hybrid teams made of multiple humans and machines.

Figure 1 does not exhaust the possibility for additional nuanced forms of joint and independent activity. There are no doubt interesting examples we have not covered. And some applications will require very simple forms of collaboration. However, this framework provides a common basis for contrasting approaches based on the critical differentiating collaborative aspects of interdependence and intention. We hope the brief discussion in this section will suffice to give readers enough of an intuition about what constitutes “teamwork” to proceed with the next phase of our discussion.

### The Implementational Challenge: Principles and Mechanisms to Enable Effective Teamwork

Clarity on the definition of what teamwork is provides a sound starting point, but operationalizing that definition into a domain specific solution still presents many challenges. Implementation requires understanding the kinds of support

needed to facilitate the satisfaction of joint requirements in teamwork, thus assuring successful human-machine team performance.

From the definition and its examples depicted by figure 1, it should be clear that interdependence is the defining characteristic of teamwork (Johnson et al., 2011). That interdependence comes from two primary sources: the work and the participants. In both teamwork and collaboration, how effectively and efficiently the group manages these interdependencies is the main basis for assessing the quality of the teamwork or collaboration (Johnson et al., 2014).

In our view, non-trivial teamwork is facilitated to the degree that joint activity is accomplished by some combination of human effort and technology to support three things: observability, predictability, and directability. These three interdependence relationships are consistent with long-standing principles in human-centered design (Johnson et al., 2014; Johnson & Bradshaw, 2021):

1. **Observability** refers to how clearly pertinent aspects of one’s status — as well as one’s knowledge of the team, task, and environment — are observable to others. This is commonly referred to as “transparency,” but we prefer the term “observability.” The complement to making one’s own status observable to others is being able to observe status.
2. **Predictability** refers to how clearly one’s intentions can be discerned by others and used to predict future actions, states, or situations. The complement to sharing

one's own intentions with others is being able to predict the actions of others. It requires being able to receive and understand information about the intentions of others, to be able to predict future states, and to take those future states into account when making decisions.

3. **Directability** refers to one's ability to be directed and influenced by others. The complement is to be able to direct or influence the behavior of others.

A given set of joint requirements may not necessitate that the three primary forms of support for interdependent work (observability, predictability and directability) be available in equal measure. To satisfy a given joint requirement, support for only one or two of the three may be sufficient—and sometimes that may be all that is possible. However, our observations suggest each provides a unique form of facilitation that can become almost essential in a given setting.

Of important note, is that these abilities are not related to functions or taskwork. They are *relational*. They exist between two parties. As such, they are complementary in nature (Johnson et al., 2014). Just as it is helpful to be able to observe, predict, and direct teammates, so it is important to be observable, predictable, and directable oneself. Though it is true that one is sometimes required to observe without being observable, supportive mutual observability is usually a plus. Given the inherent differences in humans and machines, the kind and degree of observability, predictability and directability that is possible for a machine will rarely be identical and symmetrical to what is possible for a human, but that is to be expected. Indeed, the same thing is true with any two people. To achieve robust and resilient teamwork,

as much support for observability, predictability and directability should be provided as can be leveraged effectively. Additionally, richer degrees of support often enable increased collaborative capabilities.

These three forms of teamwork support can be implemented by any number of specific mechanisms. For example, team members who want to make their actions predictable by announcing their intentions could equally well communicate them over a radio channel or using a hand signal — or a combination of both. But, of course, the choice of an appropriate signaling mechanism depends not only on the capabilities of the team member sending a message but also on the capabilities of the one receiving it. A signal cannot be successfully “observed” unless the receivers are able to attend, perceive and interpret them, thus they must be complementary.

Another key implementational challenge is that it is difficult to separate teamwork activity from taskwork activity. A simplistic view is that taskwork is what would be done if one did the work alone and teamwork is the additional stuff necessary if done as a team. However, this simple view can mislead developers to believe the two are cleanly separable. Teamwork permeates all activity and alters how that activity is performed. Sensing, interpretation, option development and selection, planning and execution are all modified by additional considerations and context when done as a team interdependently rather than as an individual independently. We propose that the mechanisms that facilitate such an adaptation are observability, predictability and directability and help to characterize the nature and quality of collaboration.

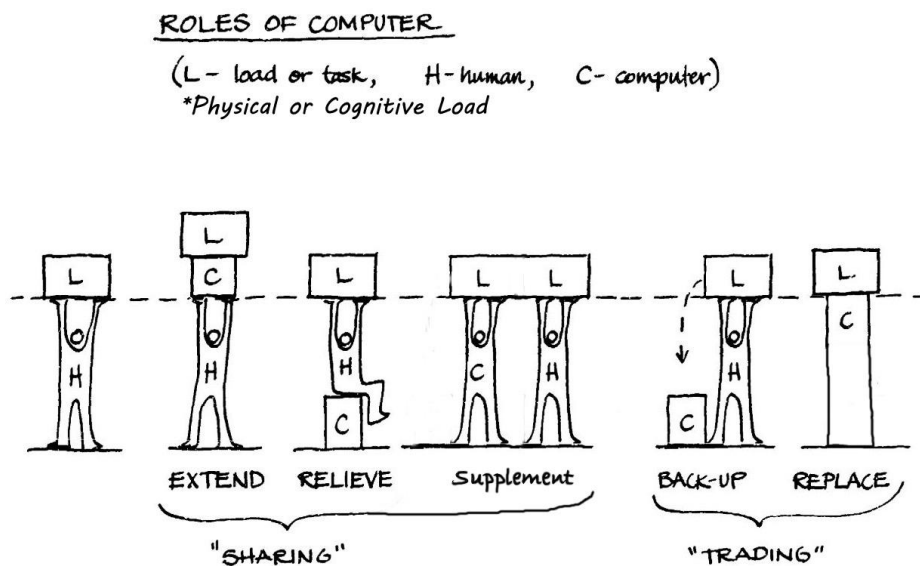


Figure 2. Potential roles of machine, modified from Sheridan & Verplank (Sheridan & Verplank, 1978).

## Discussion

There are many ways to view teamwork and collaboration. A common approach is to consider what people are good at and what machines are good at and try to allocate work accordingly. This approach is probably the oldest and was introduced in the 1950's by Fitts as the HABA-MABA (humans-are-better-at/machines-are-better-at) approach (Fitts, 1951). However, this function allocation perspective ignores the fact that allocation alone is insufficient. As noted by Drucker long ago "When it comes to the job itself, however, the problem is not to dissect it into parts or motions, but to put together an integrated whole" (Drucker, 1954). The observability, predictability and directability (OPD) framework can help provide more nuanced guidance. For example, when allocating, what options are available to shape how the allocated task might be performed (directability)? Once allocated, what behavior should I expect, how long will it take and what is its likelihood of success (predictability)? During execution, how can I tell if things are going well or if they are going poorly and performer needs help (observability)? These qualities will influence the fluency and quality of teamwork, as well as what type of teaming or collaboration is even possible.

Another common approach is to consider roles. Is the system a supervisor, a coach, a teacher, a coordinator or a partner? Roles provide good intuition with respect to the responsibilities of the team members, but are not specific enough to be sufficient design guidance. Then can help shape design questions based on the OPD framework. For example, what kind of observability is needed by a supervisor as opposed to a coach? How much predictability is needed by a coordinator versus a teacher? What kind of directability makes a teacher effective and what does a partner need?

Another view on roles comes from Sheridan and Verplank (Sheridan & Verplank, 1978). There is a pervasive misconception that AI or automation is about replacement (Bradshaw et al., 2013; Johnson & Vera, 2019). Sheridan and Verplank did not consider replacement the only or even the best option. They highlighted additional roles, such as backup, supplementation and relief. We have additionally added extending human capability to the list as a desirable role. Again, these roles help provide intuition about the relationship between people and machines, but the OPD framework can ground that intuition into specific implementation requirements.

## Conclusion

We have attempted to address the two primary impediments to progress in work on collaborative AI. We have laid out a simple and clear definition of teamwork and collaboration as managing interdependence in joint work. We have argued that interdependence in joint activity derives mainly

from two sources: the work and the participants. To address the implementational challenge, we provided the concepts of observability, predictability and directability as critical design guidance. These are consistent with long standing principles of human-centered computing. Our goal was to provide guidance that can be used by researchers and practitioners of collaborative AI to both shape design requirements and rate how well a given software or hardware implementation fares in meeting these requirements. With further discussion and refinement, we hope that such an approach will be helpful in moving the field of collaborative AI forward in a disciplined and coherent fashion, allowing approaches to be compared and evaluated with respect to a common set of criteria.

## Acknowledgments

This material is based upon work sponsored by the Office of Naval Research under grant Agreement number N00014-23-1-2886.

## References

- Bradshaw, J. M., Hoffman, R. R., Woods, D. D., & Johnson, M. (2013). The seven deadly myths of "autonomous systems." *IEEE Intelligent Systems*, 28(3). <https://doi.org/10.1109/MIS.2013.70>
- Drucker, P. F. (1954). *The Practice of Management*. Harper Business. <http://www.amazon.com/The-Practice-Management-Peter-Drucker/dp/0060878975>
- Fitts, P. M. (1951). *Human engineering for an effective air-navigation and traffic-control system*.
- Johnson, M., Bradshaw, J., Feltovich, P., Jonker, C., van Riemsdijk, B., & Sierhuis, M. (2011). The Fundamental Principle of Coactive Design: Interdependence Must Shape Autonomy. In M. De Vos, N. Fornara, J. Pitt, & G. Vouros (Eds.), *Coordination, Organizations, Institutions, and Norms in Agent Systems VI* (Vol. 6541, pp. 172–191). Springer Berlin / Heidelberg. [https://doi.org/10.1007/978-3-642-21268-0\\_10](https://doi.org/10.1007/978-3-642-21268-0_10)
- Johnson, M., & Bradshaw, J. M. (2021). How Interdependence Explains the World of Teamwork. In *Engineering Artificially Intelligent Systems* (pp. 122–146). Springer. [https://doi.org/10.1007/978-3-030-89385-9\\_8](https://doi.org/10.1007/978-3-030-89385-9_8)
- Johnson, M., Bradshaw, J. M., Feltovich, P. J., Jonker, C. M., van Riemsdijk, B. M., & Sierhuis, M. (2014). Coactive Design: Designing Support for Interdependence in Joint Activity. *Journal of Human-Robot Interaction*, 3(1), 43–69.
- Johnson, M., & Vera, Alonso. H. (2019). No AI is an Island: The Case for Teaming Intelligence. *AI Magazine, Spring*, 16–28.
- Sheridan, T. B., & Verplank, W. (1978). *Human and Computer Control of Undersea Teleoperators*. Man-Machine Systems Laboratory, Department of Mechanical Engineering, MIT.