

How Interdependence Explains the World of Teamwork

Matthew Johnson and Jeffrey M. Bradshaw

Florida Institute for Human and Machine Cognition, Pensacola FL32502, USA
mjohnson@ihmc.org and jbradshaw@ihmc.org

Abstract. Current attempts to understand human-machine systems are complex and unwieldy. Multiple disciplines throw different concepts and constructs at the problem, but there is no agreed-to framework to 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. In this chapter, we will describe 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. Our exposition is not intended to be exhaustive, but instead aims to describe the basic principles in a way that allows the gist to be grasped by a broad cross-disciplinary audience through simple illustrations.

Keywords: Interdependence, Teamwork, Joint Activity, Situation Awareness, Trust, Levels of Automation, Adjustable autonomy, Human-Machine Teaming.

1 Introduction

Technology has always been tied to human activity. The need for this relationship is self-evident because technology is not developed for its own sake, but rather for human ends in settings that inevitably include humans as beneficiaries and facilitators. In short, we might say that “no technology is an island” [1]. For this reason, providing support for *interdependence* within mixed groups of people and technological components is not a mere nice-to-have but rather an essential requirement of smooth-running work and play involving any technology. Despite this, consideration for interdependence is often deferred until it is too late in the design process, if not ignored altogether.

One of the reasons people fail to take the importance of interdependence into account has to do with the inaccurate or imprecise terminology we frequently use to describe technology. For example, the term “autonomy” (as used today and for the foreseeable future) is at best a “wishful mnemonic” [2] and at worst a costly (and potentially deadly) misconception [3]. It drives the single-minded pursuit of the illusory always-just-out-of-reach chimera of flawless performance in every situation needed for the goal of a perfectly independent system, rather than the more practical but less glamorous goal of creating machines that can work interdependently with and for people [1].

Misunderstanding human-machine interdependence inevitably leads to misguided arguments, mischaracterized results, misinformed claims, and misinterpreted conclusions.

Several research efforts have made strides in re-framing the problem in a way that highlights the needs and contributions of people whenever technology is applied. User-centered design (e.g., [4]) and human-centered approaches (e.g., [5]) have led the way in this respect. Unfortunately, relatively little has changed in the decades since these approaches were first defined. Despite the many improvements that efforts to apply these approaches have provided, the evolution of new technology remains largely technology-driven. Why is this so?

In our experience, one of the biggest reasons is that the process of accounting for the needs and contributions of people remains more of an art than a science. Acknowledging the essential role that people play in successful technology deployment is only the first step. Going further, many well-reasoned approaches fail because they fail to bridge the gulf that extends an elegant theory into the realm of implementation. AI and system engineers cannot live by abstract design concepts alone, but require guidance and tools that allow them to map general principles to specific situations [6]. Engineers, who may not be experts on human-machine teaming, will probably get very limited mileage out of high-level concepts like the infuriatingly vague requirement of keeping humans “in-the-loop.” The gap between developers writing code and theoreticians providing guidance divorces research from real-world technology design and implementation.

Another significant challenge is the breadth and complexity of domains to be mastered to ensure effective human-machine teaming. The topic spans a wide range of disciplines, each with their own theories, principles, and terminology. Consensus is difficult because researchers with different backgrounds tend to talk crossways to each other. Without a framework that can span multiple domains and connect a range of concepts in workable synergy, it is difficult to translate advances from one area to another.

2 Interdependence as an Integrative Framework for Teamwork

In line with what is suggested in the title of the workshop where the ideas in the present chapter were first presented (*Toward the Science of Interdependence for Autonomous Human-Machine Teams*), we believe that the key to understanding *human-machine teamwork* is understanding *interdependence* [7]. Though teamwork comes in many shapes and sizes, we have learned from long experience that recognizing, supporting and managing interdependence is the common ingredient that transforms capable individual performers into great teammates, no matter what the task or situation.

The same message is echoed across a range of research domains throughout the research literature. Organizational theorists Malone and Crowston specifically define

effective coordination in terms of “managing dependencies between activities” [8]. Human teamwork specialists have likewise identified appropriate interdependence among team members as a defining feature of a well-oiled team [9]. Human-machine researchers such as Paul Feltovich have identified interdependence as the very essence of joint activity [10].

We all know this intuitively. But how can we apply our intuitions in a practical manner?

2.1 The Challenge



Fig. 1. Some of the many concepts that make up the conceptual space of human-machine teamwork and technology in general.

As we have argued above, a scientific approach to human-machine teamwork requires a framework for understanding and associating the concepts that make up this complex multi-disciplinary domain. Fig. 1 shows some of the more common concepts used when discussing technology, particularly in the domain of human-machine teaming. The goal of our research is to advance the conversation on how such concepts might fit together, enabling research results from each of these fields to contribute their unique perspectives comfortably and compatibly. We propose that “interdependence” is not only the key principle that unlocks an understanding of teamwork overall but also the means of framing these seemingly disparate research perspectives as parts of a common endeavor. We will present our argument for this claim by unfolding a series of concept maps describing the framework, followed by a discussion of how this framework relates to the terminology and perspectives of colleagues working on similar problems.

2.2 What criteria define joint activity that is teamwork?




Fig. 2. There are two criteria defining joint activity as teamwork: The work must be interdependent, and the participants must intend to work together.

The first concept we will introduce is that of “joint activity,” inspired by the work of the distinguished linguist Herbert Clark [11]. To better intuit the concept of joint activity, consider an example of playing the same sheet of music as a solo versus a duet. Although the sheet of music used by the performer is identical in both situations, the processes involved in performance are different due to the interdependence relationship that governs the “joint activity” of the two artists [12].

We define “*joint-work activity*” as a special kind of “joint activity.” As depicted in the concept map shown in Fig. 2, there are two basic criteria for the kind of “joint-work activity” that we are concerned with in this chapter, a kind of joint-work activity we might rightfully refer to as “teamwork”:

1. The participants intend to work jointly toward their common goal;
2. The work itself is to be carried out in a manner that puts the participants in a state of interdependence.

In the example above, the soloist deliberately intended to work alone and the work itself was not carried out in an interdependent fashion. On the other hand, to play a duet, the two players must consciously work together in an interdependent and mutually supportive manner, thus, the duet is joint-work activity.

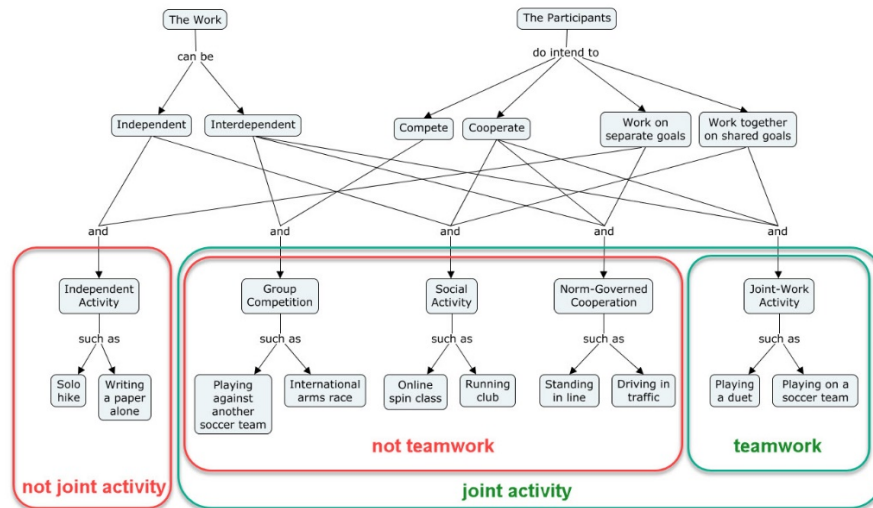


Fig. 3. Some examples of activity that are not teamwork and some that are.

Often it helps to consider what things are *not* to better understand what they *are*. Fig. 3 extends our concept map with examples that draw finer distinctions between teamwork from independent activity or other kinds of joint activity:

- *Independent activity.* Clearly, some work is independent. Those who engage in independent work have no intention to work 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 to advance a win-win situation that will advance the prospects of both teams but rather are consciously trying to work *against* each other to produce a win-lose outcome selfishly favoring only their own team. While teamwork is crucial *within* each team—and while minimal cooperation in keeping the rules and showing common courtesy is needed in order to play a *fair* and enjoyably *sociable* game—the nature of the joint activity itself mitigates against *cross-team* teamwork of the sort that would help the other team score points for themselves.
- *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. Social activity of this sort is not teamwork in the sense we have defined it because the state of interdependence between club members is

primarily social and only *indirectly* impacts the performance of the work itself.

- *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 situation, they rely for the most part on traffic laws and norms (e.g., driving courteously) to manage their interdependence with other drivers. We exclude cooperation of this type—the “do no harm” type of cooperation that is merely intended to provide a level playing field whereby all drivers can get to their own separate destinations—from our definition of teamwork.
- *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.

In summary, our two criteria of interdependent and intentionally shared work exclude looser forms of joint activity, such as social activity and norm-governed cooperation, that do not resemble teamwork of the sort we are discussing. While social activity, norm-governed cooperation, and competition could be seen as examples of *joint activity*, they do not rise to the level of what most people would call teamwork. That said, it is not uncommon for looser forms of joint activity to develop features that border on or may even cross over into joint-work activity. For example, two drivers may use eye contact and hand signals to help each other avoid collisions while changing lanes in traffic. Similarly, although physical workouts are, strictly speaking, an individual matter, you might ask someone to spot you in a gym and running club members might offer you encouragement, which can indeed help your performance. The fluid nature of joint activity sometimes makes it challenging to understand teamwork.

Fig. 3 does not exhaust the possibility for additional nuanced forms of joint and independent activity. There are no doubt interesting examples we have not covered.¹ However, 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.

¹ Examples might include asynchronous vs. synchronous collaborative work (e.g., software development teamwork), independent work on shared goals (e.g., independent charitable organizations trying to eradicate poverty, reduce death and illness, or reduce hunger), and cooperation established through the (non-intentional) establishment of norms (e.g., cow paths).

2.3 What makes teamwork a special kind of joint-work activity?

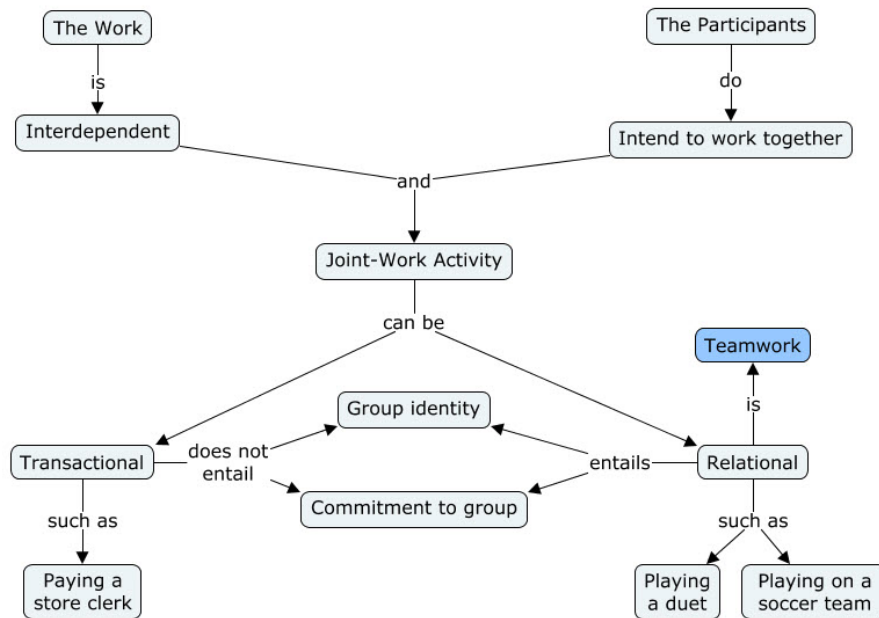


Fig. 4. There are two additional factors that help differentiate teamwork as a special kind of joint activity: group identity and commitment to the group.

Even within joint-work activity, there are additional factors that go further in distinguishing teamwork from more ordinary kinds of joint-work activity, as depicted in Fig. 4:

1. Team members tend to *share a group identity*. In sports, this identity may be made salient by things like uniforms, team names, slogans, and mascots. But these kinds of things in and of themselves do not create group identity, they simply reinforce the subjective perceptions of group identity held by team members.
2. Team members *share a commitment to their group*; a commitment that favors their group over others when decisions must be made. Sometimes that commitment leads to the sacrifice of individual goals and preferences when they conflict with team objectives. We might call that trait “team loyalty.”

These factors are not found in all types of joint-work activity. For example, it would be odd to call temporary, transactional forms of joint activity such as paying a store clerk, teamwork. In transactional forms of joint activity there is usually no persistent shared identity and generally no mutual commitment beyond what the immediate situation demands. (Of course, in everyday life, some people go out of their way to be

helpful and friendly with people they may never meet again, just because they know that valuing others in this way makes life more pleasant for everyone.)

In summary, teamwork is a special type of joint activity. It requires parties that intend to work together in an interdependent manner. The members of the team tend to share a form of group identity, unique to the team, and a commitment of loyalty to the team that influences their decision-making. While this background understanding is helpful in understanding what we mean—and what we *don't* mean—when we speak of teamwork, we need to dig deeper in order to apply the concept effectively to the design and use of technology.

2.4 Where does interdependence come from?

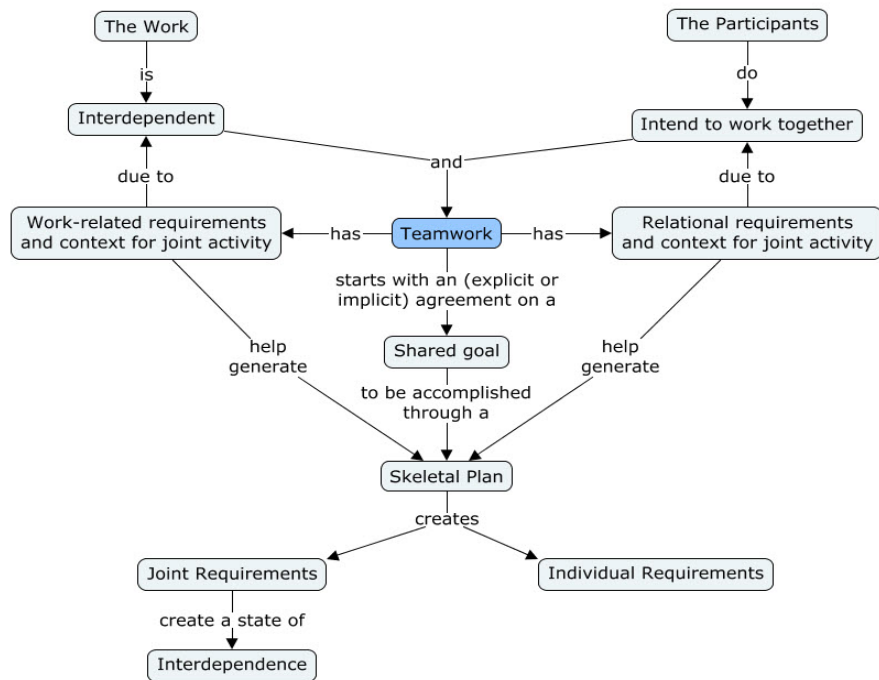


Fig. 5. A state of interdependence arises when the joint requirements of the skeletal plan are being executed. Joint requirements stem both from *work-related requirements* due to interdependence inherent in the nature of the work itself and *relational requirements* due to interdependence created by the intention of the parties to work together.

Now that we have laid out the basic features that distinguish teamwork as a type of joint activity, we can start discussing interdependence. Fig. 5 extends our concept map for the discussion. Teamwork starts with a shared goal, either implicit or explicit. This

goal is what you intend to work together for, which leads to what we might call a “skeletal plan” [13]. In most of what we do in life, we start with a general, skeletal plan of what we need to do and then flesh out the plan or change it incrementally as we go along. There is no teamwork without both a shared goal and a shared plan, be they ever so simple or implicit. Sharing a general goal with someone else’s — like world peace — is not enough. It is not teamwork until we begin doing something about it together through the execution of a skeletal plan.

Skeletal plans to support teamwork contain two kinds of requirements: individual requirements and joint requirements. *Individual* requirements are what one team member must do *independently*. For example, to get some piece of work done you might need to be able carry things or to fly an airplane on your own. The *joint* requirements within a plan create a state of *interdependence* between some set of participants with respect to some part of the work to be performed. For example, the plan may require you to lift your side of a large table that is being moved, or to continually monitor your teammates so you can stay aware of any problems that crop up and provide help as needed.

2.5 Where does the skeletal plan come from?

We have just argued that the state of interdependence that exists between teammates is generated during execution of the joint requirements of a shared skeletal plan. But where does the skeletal plan come from? In brief, there are three necessary elements (see Fig. 5):

1. *Shared goals*. One or more shared goals give direction and purpose to the plan.
2. *Work-related requirements and context*. It is well and good to define the goals of the work, but to be able to plan successfully, we must know something about the requirements and context of the work itself. With respect to work *requirements*, we need to know, for example, whether tasks need to be done in a particular order, in a particular way, and whether they can or must be performed by more than one person. With respect to work *context*, we must know, for example, if the task is performed differently in daytime than it is in nighttime, or under different weather conditions, or when available team members have different capabilities.
3. *Relational requirements and context*. As an example of a relational *requirement*, even though the work itself may not strictly require more than one person to perform it, the plan may take into account that although assigning two people to the task may seem a short-term loss of efficiency, it may strengthen the skills or increase the comradery between them in the long run. More formal relational requirements are organizational structures such as the chain of command. Examples of relational *context* can be simple things like non-availability of a teammate or more nuanced contexts such

as knowing that two of the people on the team are not getting along well today.

When these three elements are combined to form a skeletal plan, joint requirements generally result. As the work progresses, changes to any aspects of the goals, requirements, and context are likely to affect the interdependence among the teammates, which can redefine what counts as good teamwork. This coupling of goals, requirements and context is another aspect of teamwork that makes it complex.

2.6 How does teamwork relate to taskwork?

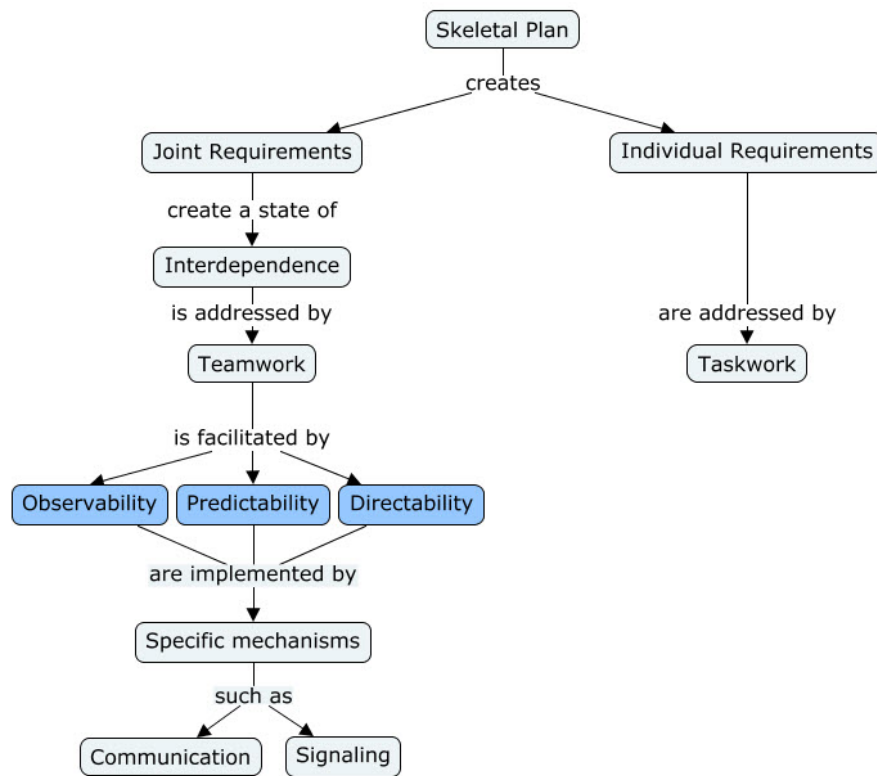


Fig. 6. Teamwork is facilitated by observability, predictability and directability.

Fig. 6 extends our concept map to differentiate between teamwork and taskwork. Specific actions and supportive behaviors that facilitate teamwork address the state of interdependence between participants in a joint activity that has been created as joint requirements are in the process of being satisfied. An example of teamwork is when two individuals carry a large table together. On the other hand, taskwork addresses requirements that can be handled by a single individual — for example, hammering a nail. Of course, these are not rigid categories: in theory an individual whose arms are

long and strong enough might be able to move a large table alone. Likewise, someone hammering a nail might sometimes need someone to hold the nail. Thus, we might say, more generally, that every item of taskwork has the potential to develop into a form of teamwork and vice versa.

As another example of the fluidity of the distinction between taskwork and teamwork, note that the same action sometimes can be performed in support of either taskwork or teamwork. For example, opening a door for oneself is simple taskwork, but opening a door to help a friend is a form of teamwork. The fact that what was originally envisioned as individual taskwork may suddenly become a joint-work activity is yet another example of what makes teamwork difficult to characterize.

Another thing that makes characterizing teamwork difficult is the fact that team members sometimes invest in joint efforts that require more labor but seem to have no immediate benefit when compared with the option of doing the work alone. Participants might justify their “wasteful” behavior by pointing to the possibility of greater rewards or reduced risk mitigation in the long term. To understand this situation, consider the short-term irrationality of buying insurance. Because people pay for insurance up front and the chances of an immediate disaster are small, it is likely that in the short run you will end up spending more in your premiums than you get back in your claims. However, when trouble strikes, insurance benefits provide financial relief, potentially forestalling fiscal ruin. In the same way, although teamwork may sometimes be less efficient than solo performance, it helps to ensure resilience when challenges arise. For example, a good copilot will monitor the pilot, verify whether the flight trajectory is appropriate, and be ready to act if anything goes wrong. However, if everything goes smoothly, there is no measurable benefit to the outcome (except the confidence provided by having a reliable backup). Similarly, even when an inattentive or incapable co-pilot is on duty, there may still be no measurable effect on the outcome when the pilot’s performance is flawless. Sometimes teamwork benefits and failures are only exposed when things do not go as planned. For this reason, the value of teamwork should be assessed not only with respect to those instances where it facilitates *actual* failure recovery, but also in its anticipatory, protective function in mitigating the *potential* risks of failure or allowing *potential* exploitation of unforeseen opportunities in the long-run. Determining the sweet spot in the cost-benefit tradeoffs that govern opportunity exploitation and risk mitigation is yet another challenge for understanding and evaluating teamwork.

By way of additional clarification, Fig. 5 and Fig. 6 capture the fact that the word “teamwork” is used in different ways. For example, sometimes it refers generally to the type of activity in which participants work together interdependently on shared goals, as shown in Fig. 5. This sense of “teamwork” *defines* the requirements. At other times, it refers to the specific actions and supportive behaviors that facilitate the type of activity, as displayed in Fig. 6. This sense of “teamwork” defines how the requirements are *met*. Below we explore this second sense, a sense that is at the heart of the teamwork design challenge.

2.7 What kinds of support are needed to facilitate teamwork?

We have already seen that teamwork generates joint requirements. Now we seek to understand what kinds of support are needed to facilitate the satisfaction of those joint requirements. In brief, providing effective support for team members as they fulfill joint requirements is the *sine qua non* of designing and developing successful human-machine team performance.

In our view, non-trivial teamwork is facilitated to the degree that joint activity is supported by some combination of human effort and technology helps in three things: observability, predictability, and directability (see figure 6).

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.

These three interdependence relationships are consistent with long standing principles in human-centered design [14]. The importance of these three interdependence relationships can be seen throughout the automation literature with many references to observability (often referred to as transparency) e.g., [15]–[18], predictability e.g., [16], [18]–[20], and directability e.g., [16], [21].

As an example of how observability, predictability, and directability facilitate teamwork, let's imagine a group working together to search a building. Team members support observability when they share their current location with others. They support predictability when they inform the team that they are heading to the second floor. They are supporting directability when they are able to ask someone to check the stairwell and also when they allow themselves to be directed by someone else when appropriate.

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 of the three kinds of support provides a unique form of facilitation that can become almost essential in a given setting.

For example, consider a situation where you have agreed to meet a new friend for coffee, without specifying a time. If they were simply observable, through a phone app that allows you to track their location, you could just monitor their position in real time and wait to see if they eventually arrive. Though simply *observing* might eventually work if you were not in any kind of a hurry and knew that if you waited long enough they would eventually arrive, your job would be much easier if instead you had agreed beforehand to meet at a specific time. Agreeing to a specific time allows you to *predict* their arrival time. That way you don't need to start observing until a few minutes prior to your appointment.

If you began the habit of meeting every week at the same time, your experience would allow you to improve your predictions. For instance, you might notice your friend consistently arrives fifteen minutes late. Your past *observations* can thus be used to adjust your future *predictions*.

Finally, consider the additional usefulness of *directability*. Your friend calls you on the phone in advance and asks you to meet him at a different coffee shop this time, thus influencing your future actions. This example shows how past and current *observations*, shared intentions (enabling *prediction*), and *directability* work in tandem to facilitate coordination within teams. Each of these three forms of support has its own cost and its own value in a given situation. Generally speaking, using all three forms together is more effective and reliable than relying on only one.

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. In order to achieve robust and resilient teamwork, as much support for observability, predictability and directability should be provided as can be leveraged effectively in the performance of relevant joint requirements

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. Herbert Clark's joint action ladder [12, p. 147] is a reminder that an interaction requires complementary capabilities by both participants in the interaction. In other words, a signal cannot be successfully "observed" unless the receivers are able to attend, perceive and interpret them.

2.8 How does teamwork continually adjust over time?

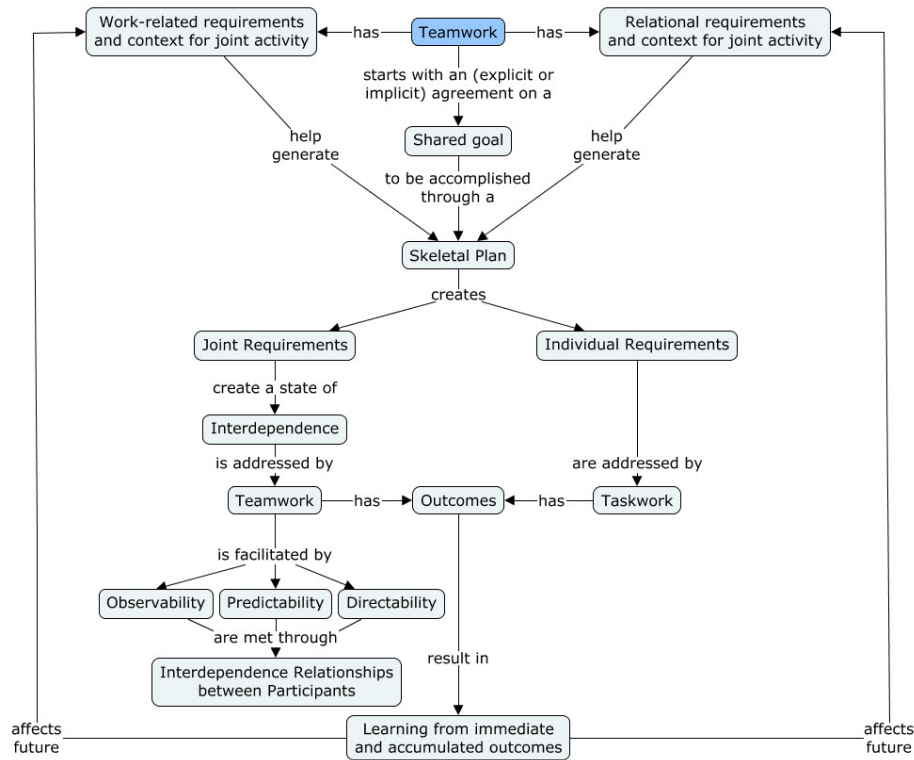


Fig. 7. All plans are tentative, thus activity, including teamwork, needs to adjust over time. The ability to adjust fluently is one of the most important characteristics of effective teams.

As the joint activity progresses, the teamwork and the taskwork will result in *both interim and final* outcomes. Fig. 7 extends our concept map to include outcomes, feedback, and learning. Effective teams will learn from their experiences. Learning from both immediate outcomes and the results of accumulated outcomes over time can provide feedback that will affect both the work and relational requirements and context. Future decision priorities and the skeletal plan itself can be adjusted in light of these results.

Fig. 7 provides a basic summary of our proposed conceptual framework. It emphasizes the critical role that interdependence plays in joint activity generally, while emphasizing why it is even more crucial in teamwork. Interdependence relationships among team members stem from their participation in fulfilling joint requirements in the skeletal plan. The skeletal plan is formed and adjusted by taking the goals, the nature of the work, and the participants' relationships into account. The teamwork necessitated

by work to fulfill joint requirements is facilitated by specific mechanisms that support observability, predictability and directability. To be successful, these mechanisms should be compatible with the complementary abilities of team members to both communicate and understand the signals used. Learning from immediate and accumulated outcomes allows requirements, contexts, goals, and plans to be refined and improved as the joint activity progresses. Another complexity in discussing teamwork is that the term *teamwork* is used both to describe a certain type of joint activity and also to describe the processes used to facilitate successful accomplishment of such activity.

3 How does the framework help us to understand the broader world of human-machine teamwork?

We now return the conceptual soup of different approaches to human-machine teamwork shown in Fig. 1. We will apply our proposed framework to show selected examples of how these seemingly disparate research perspectives can be seen as parts of a common endeavor.

3.1 How does interdependence relate to situation awareness?

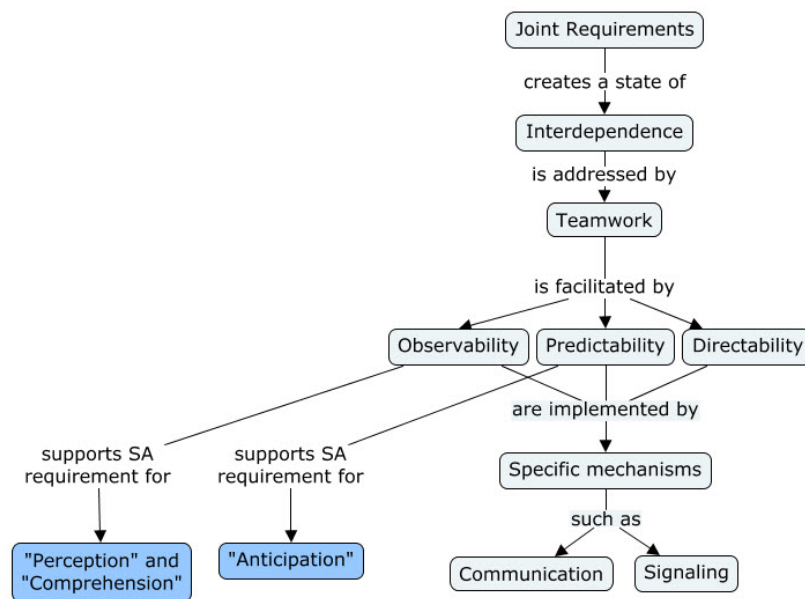


Fig. 8. Situation awareness (SA) is supported through observability and predictability requirements. Failure to meet those requirements will lead to poor situation awareness.

The theory of situation awareness, as outlined by Endsley and Kiris [22], describes the role of perception and projection in teamwork, a good match for the respective concepts of observability and predictability in our framework. The joint and individual requirements determine what information is relevant. Support for situation awareness must be configured in a way that takes both observability and predictability into account. Providing more support than is needed for a given task will be unnecessary, potentially annoying, distracting or overwhelming. Providing less than is needed will be potentially detrimental. Situation awareness, even when properly matched to the interdependence needs of joint requirements, can be hampered if the specific mechanisms used to convey it are poorly implemented for the situation at hand.

The challenge for a designer is knowing what comprises the “situation.” It is equally important for each teammate to understand the “situation” of other teammates when teamwork is actually underway. Teammates need to understand what other team members are aware of or need to be made aware of so they can communicate effectively. One proposed approach to designing for situation awareness is to employ goal-directed *task* analysis [23]. This is a logical starting point, since it is the joint activity that is the basis for the team’s formation. It is also true that shared goals play an important role in determining situation awareness needs. However, we argue it is not the tasks generally but rather the interdependence relationships engendered in the performance of joint requirements that determine the situation awareness needs of each teammate. These interdependence relationships derive from a combination of the work requirements, the participants relationships, the shared goals, and the skeletal plan.

3.2 How does interdependence relate to the “levels of automation” and “adjustable autonomy” approaches?

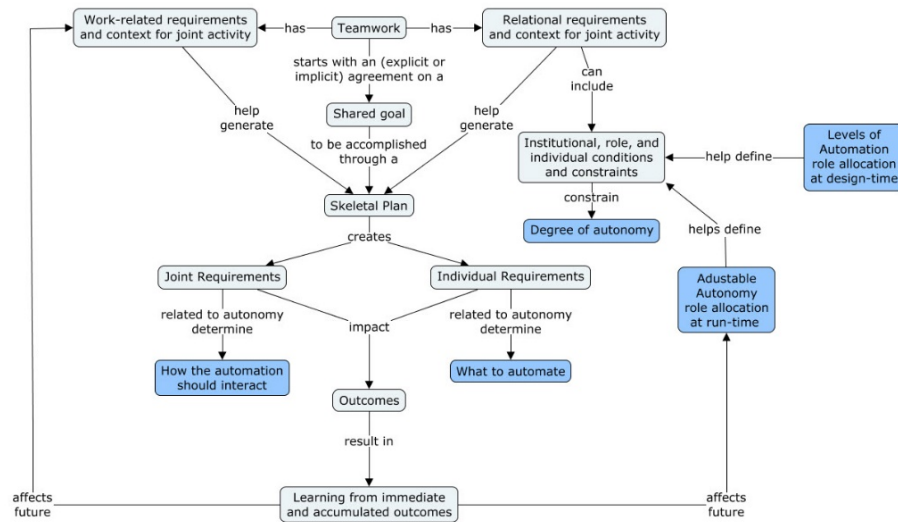


Fig. 9. Levels of automation help define some of the relational requirements, which, in turn, constrain the degree of autonomy for various elements in a given system. Joint and individual requirements related to autonomy determine what to automate and how the automation should interact.

The idea of *levels of automation* (LOA) is one of the most pervasive concepts in the domain of human-machine teaming. It has been viewed as an approach to designing systems [24] or a framework to help designers make design choices [25], [26]. LOA is often used to help designers make appropriate role allocation decisions at design-time that will use combinations of human and machine capabilities to best advantage. In turn, the degree of autonomy exhibited by a machine, described simply, is a manifestation of the opportunities and constraints for action delegated to it in its particular role, as seen in Fig. 9. In essence, a role is a label for a set of tasks that define what an actor is responsible for—and, usually implicitly, what responsibilities should be left to others. In LOA, task assignments may thus be seen a simple mechanism for defining relational boundaries between different team members. These decisions help answer the question of “what to automate” [25].

Adjustable Autonomy (AA) is somewhat similar philosophically to LOA, except that it goes beyond *design-time analysis* to allow for dynamically adjusting allocation strategies *at run-time*. LOA and AA are similar in that they both constrain the degree of autonomy by defining roles that divvy up individual responsibilities.

However, while LOA and AA help in the design of *taskwork* to satisfy *individual* requirements, it says less about the problem of how to design *teamwork* so that the interdependence created by *joint* requirements can be addressed. It is important to realize that all “intermediate levels” of automation are joint-work activity, not cleanly separable functions to be allocated in isolation [1]. For this reason, it is important that designers to build algorithms not just to do individual work, but to support joint-work by supporting interdependence relationships through specific teamwork mechanisms. Increased effectiveness in human-machine systems hinges not merely on trying to make machines more “independent” through increasing levels of automation but also on striving to make them better team players with humans and other machines with which they interact [16].

LOA and AA decisions constrain teaming options and shape the potential interactions between people and technology. This can impose (often unanticipated) joint requirements which can have negative side effects and impede effective teaming. The human factors community has a long history of documenting this result e.g., [27], [28]. Given the critical nature of effective teaming, particularly in the envisioned sophisticated domains, we argue effective teaming should take a prime position in the design of automation. As a critical factor in the design of technology [7], the interdependence of participants in joint activity should be shaping decisions about autonomy, not the other way around.

3.3 How Are Trust Decisions Made?

Trust is an important aspect of teamwork and it is an increasingly important part of technology. As technology moves out of factories and laboratories and into the world, the importance of trust will continue to grow. To understand trust, it is important to understand that trust is relational [29]. In order to establish, develop and maintain appropriate trust, technology needs to be endowed with appropriate support for teamwork mechanisms that support interdependence relationships, such as observability, predictability and directability [30].

Trust, whether between people or between people and machines, is always exploratory and context-dependent. As Hoffman states:

Active exploration of trusting-relying relationships cannot and should not be aimed at achieving single stable states or maintaining some decontextualized metrical value, but must be aimed at maintaining an appropriate and context-dependent expectation. [31, p. 157]

As technology increases in sophistication and complexity, it will become more and more challenging for people to establish trust. People cannot simply observe a complex system in a single circumstance and be confident it can handle all situations. When the goal is a joint solution, with people and technology combining to produce something greater than either alone, it becomes even more challenging. For example, doctors will

not blindly accept AI medical decisions without understanding something about the process behind the decision. Thus, today's sophisticated technology will need support mechanisms that allow them to engage in interdependence relationships. Observability, predictability and directability allow the trustor to develop appropriately calibrated trust relationships with the trusted party.

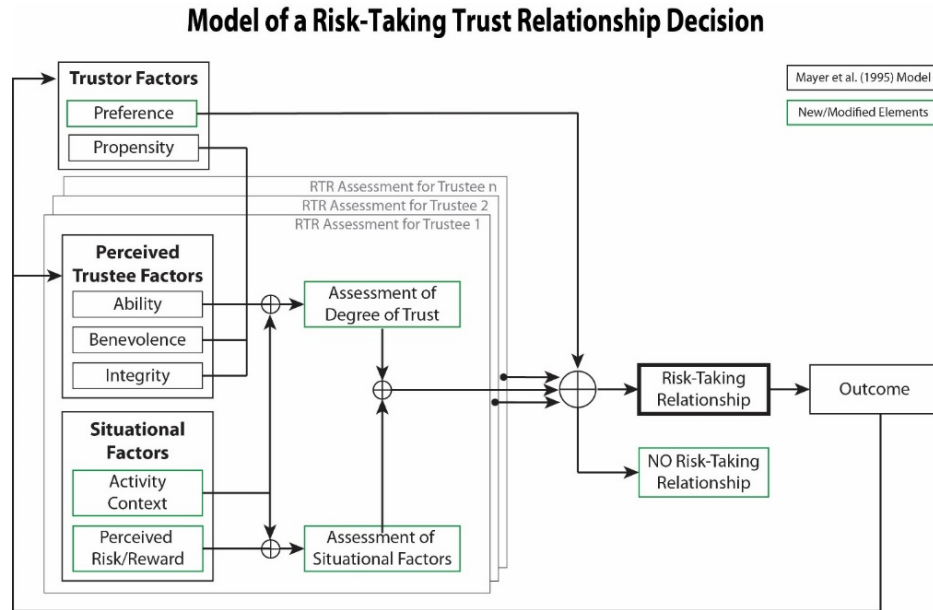


Fig. 10. Modified model of risk-taking relationship (from [30]).

Before associating trust with our framework, we need to review the concept of trust, using a modified version of the popular Mayer trust model [29], shown in Fig. 10. The Mayer model distinguished the trustor factors and trustee factors that influence trust. Most critically, Mayer emphasizes that trust does not exist until the trustor engages in what Mayer calls a *risk-taking relationship*. These relationships are influenced by feedback from outcomes to calibrate trust over time. A more detailed description of our extended model of trust has been published elsewhere [30].

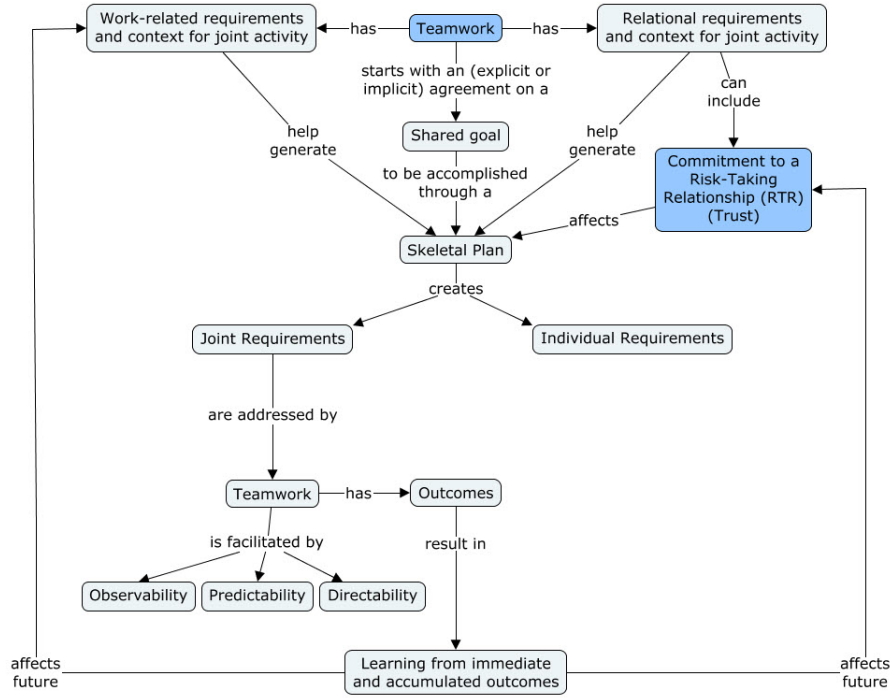


Fig. 11. Trust within the interdependence framework.

Since trust is about a risk-taking relationship, it is therefore just another aspect of relational context that affects the skeletal plan. Moreover, trust in our modified Mayer model is calibrated over time through feedback from outcomes, just as in our framework as shown in Fig. 11. Trust feedback is largely related to the teamwork outcomes facilitated by observability, predictability and directability. Further discussion on trust can be found in our previous works [30], [32].

4 Discussion

The purpose of the proposed framework is to help ground the conversation about human-machine teaming in easy-to-understand concepts that address each of the key research issues. It advances the science of interdependence by clarifying:

- what we mean by interdependence,
- where it comes from,
- how to identify it in a specific application,
- how to facilitate the management of it, and
- how, in working together, each of these factors impact human-machine team outcomes in particular ways.

Significantly, and in line with the objective we outlined at the beginning of the chapter, we have shown examples of how “interdependence,” as we think of it, is not only the key principle that unlocks an understanding of teamwork overall but also the means of framing seemingly disparate research perspectives as parts of a common endeavor.

In instances where we may have mischaracterized or oversimplified the mapping of other research perspectives to our framework, we would welcome feedback from the research community, so our misconceptions can be corrected. With respect to other important research perspectives that we have not directly discussed in this chapter, we hope to be able to work cooperatively with the research community to see whether these ideas can be further extended to encompass the “world of teamwork” so that our somewhat fragmented field can be further unified.

Before concluding, we would like to address two questions that introduce additional benefits of an interdependence-based approach like the one we have described above:

- How does an interdependence-centric framework help researchers generalize results?
- How does an interdependence-centric framework help to explain experimental results?

4.1 How Does an Interdependence-Centric Framework Help Researchers Generalize Results?

A significant feature of the scientific method is that rather than merely confining its interest to explaining specific instances, it seeks to understand laws and principles that may govern a variety of situations. In brief, the ability to generalize and extend the results of one experiment to a broad class of problems is at the heart of science.

In our own research, we have been impressed by repeated findings by our own team and by others that interdependence analysis (IA) is a highly efficacious approach that can be applied to a wide range of socio-technical systems involving people and machines. The Coactive Design method [7], [33] introduced interdependence analysis and it has been extended in more recent work [34]:

- IA has been applied to ground robotics [35], aerial robotics [36], humanoid robotics [37], and software agents [38].
- IA has proven effective in several domains including disaster response [39], [40], military applications [41], [42], space applications [43], network security [44] system design [45] and the design of large scale unmanned aerial vehicle operation centers.
- IA can be applied in the formative design stage [46], throughout an iterative design and re-design process [37], [40], and as an analysis tool for existing systems or proposed conceptual designs [42]. It helps both the

development of proper behaviors [37], [47] and development of appropriate interfaces [44], [48].

Besides the examples of generalization mentioned above, the practice of IA has helped researchers identify common patterns of failure and their solutions. For example, one common failure pattern identified in different projects is what we call *breaking what is working*. This phenomenon occurs when the addition of new technology disrupts existing information flows and hinders performance. The use of IA has made it clear that this problem is often due to impediments to observability that emerge when automation “hides” aspects of the work being performed that need to be observable to team members. This problem cannot be identified simply by looking at the taskwork that the newly added technology performs, but rather must be understood through an examination of interdependence relationships that are created through joint requirements such as, in this case, observability. Evidence for these claims was found by Johnson et al. [33] in the disaster response domain, in which automation of a grasping task inhibited the human’s ability to interpret the situation and recover from small deviations. Beierl and Tschirley [42] found similar issues when they used IA to review proposed solutions for integrating robots into Marine Fire Teams. The proposed introduction of a machine into the Fire Team inhibited some of the natural communication channels used between team members.

A second example of such a pattern is what we call *black box help*. This is similar to *breaking what is working*, in that lack of observability is generally the issue. However, instead of interrupting an *existing* workflow by preventing teammates from viewing the actions of something that was previously observable, it arises when *new* technology and workflows are added that hinder observability of the new technology’s actions by others. The problems generated by the opaque nature of black boxes have spawned programs to try to address the issue, such as Defense Advanced Research Projects Agency Explainable Artificial Intelligence program (DARPA XAI), which notes that “the effectiveness of these systems is limited by the machine’s current inability to explain their decisions and actions to human users.”² While traditional human factors analysis has a long history of identifying this as a problem *retrospectively*, IA provides an *upfront* analysis solution where the problem and potential options for solving it can be identified early in the design phase [49].

A third example of how common patterns and solutions can be brought to light through IA is what we call the *incomplete solution*. This is when technology is injected to assist people who are overloaded. However, when the new technology cannot address all aspects of the work, often including implicit requirements that were not identified during the design process, it becomes more of a burden than a help. The IA conducted by Zach [41] analyzed the effort to leverage unmanned systems to provide intelligence, surveillance, and reconnaissance to Marine Corps tactical units. It identified five missing feedback loops which demanded more from the marines due to the

² <https://www.darpa.mil/program/explainable-artificial-intelligence>

lack of capabilities on the unmanned system. The IA conducted by Beierl & Tschirley [42] identify the unmanned system’s lack of ability to distinguish “the enemy” as a key inhibitor in its effectiveness. The IA conducted by Johnson et al. [34] showed how current automated traffic avoidance systems lack the ability to address uncooperating traffic. In each of these cases, the problem was not found by looking at what that automated system was doing, but by looking at how that work was interdependent with what the people interacting with the system, its “teammates,” would be doing.

In all the ways we have just mentioned, IA represents progress in improving the degree to which the scientific method can be applied to human-machine teamwork: identifying generalizable laws, principles, patterns, and solutions that facilitate meaningful progress in our common research field.

4.2 How Does an Interdependence-Centric Framework Help Explain Experimental Results?

Though space prevents an exhaustive survey in the present chapter, we will provide a few examples of how an interdependence-based framework helps explain many common results in human-machine teaming studies.

As a first example, we note that while the promise of technology has nearly always been that higher levels of automation (LOA) will make our lives easier, reduce workload, and improve performance, experimental evidence often suggests otherwise. Endsley’s [26] computer-based dynamic control task experiment showed that LOA involving joint human-computer selection had no significant impact on performance, as compared to purely human selection. She also found that operator ability to recover from automation failures was substantially improved with lower LOA, and operators were actually hindered when assistance was provided with higher level cognitive functions. Similarly, Calhoun and Draper [50] found it took significantly longer to complete the re-routing task in their multi-UAV testbed with the high LOA. Li and Wickens’ [51] experiment with remote operation of a robot echoed this pattern of high LOA performing worse. Many different reasons have been given for such results, but we are convinced that they can frequently be explained through the lens of interdependence. For example, we have previously demonstrated that increasing the LOA without adequately addressing interdependence relationships can disrupt natural coordination (i.e., *breaking what was working*) and result in a negative inflection in performance as automation increases at the expense of observability and predictability[52]. Endsley [26] speculates this is the case suggesting the results are indicative of a lower level of direct feedback experienced when not actually performing a task implementation. Our analysis explains some otherwise surprising LOA results that have reported by others [53], [54].

As a second example, one of the most common experimental observations is that “humans tend to be less aware of changes in environmental or system states when those changes are under the control of another agent (whether that agent is automation or another human) than when they make the changes themselves” [25]. From our

perspective, this result can be explained straightforwardly from the fact that many systems are not observable or predictable, making the system opaque to the user (i.e., *black box help*), inhibiting situation awareness and maintenance of common ground [52]. One of the most common results found in level of automation studies is that situation awareness significantly decreases with “higher” levels of automation [24]. This effect can be mitigated by properly designing support for interdependence (i.e., observability, predictability, directability).

In this regard, it should be no surprise that a breakdown in lower-level team functions, like observability and predictability, has a negative impact on team performance. It is particularly apparent when teams face off-nominal situations. Although higher levels of automation can beneficially reduce workload during normal operating conditions, it can also impede situation awareness (i.e., *breaking what was working*), making it difficult for the human-machine team to adapt to unexpected events. In a remote robot arm experiment, Kaber et al. [24] did in fact show that increasing LOA reduced time-to-complete task for normal operations, but the opposite was true for time-to-recover during failures. Li and Wickens [51] also showed that failure detection is poorer when technology is under the control of automation, which is noted as a consistent pattern in Onnasch et al.’s survey which concluded that “automation helps when all goes well, but leaving the user out of the loop can be problematic because it leads to considerable performance impairment if the automation suddenly fails” [55]. We prefer a more specific description than the general statement that the human is “out of the loop” to explain these findings. More specifically, we posit that lack of support for interdependent work (i.e., observability, predictability, directability) in the automated solution diminishes the ability of the person to recover from failure. Conducting a proper IA can go further and identify the specific joint requirements that were not adequately supported by facilitating mechanisms.

Parasuraman et al. [25] rightfully acknowledge the challenges of function allocation when they state, “The performance of most tasks involves interdependent stages that overlap temporally in their processing operations.” Their four-stage model included sensing, perception, decision-making and response selection. It is clear that what you perceive depends on what you sense, what you decide depends on what you perceive and what you select depends on what you decide. What is less obvious is that what you need to decide on may shape your perception or shift your attention. This means failure anywhere can have ripple effects on system outcomes, highlighting the critical role interdependence plays in understanding experimental results in human-machine teaming. In order to properly interpret experimental outcomes, a thorough interdependence analysis is essential. This interdependence analysis has an added benefit of potentially providing new ways to instrument and measure teamwork.

5 Conclusion

We have proposed and illustrated that a correct understanding of *interdependence* is key to characterizing human-machine teamwork in an understandable, actionable, and

generalizable manner. Our framework proposes three principal forms of support that can be used to facilitate effective, interdependent teamwork: observability, predictability, and directability. Though researchers have sometimes used different words to describe these three forms of teamwork support, results from our own studies and those of others has provided convincing evidence of their primacy, usefulness, and generalizability. We have argued that an interdependence-centric framework enables researchers to predict and explain experimental results in terms of interlocking relationships among well-defined operational principles. We are hopeful that continued discussions of this framework will provide a sound and practical foundation for developing a deeper science of interdependence.

Acknowledgements

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR001120C0037. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Defense Advanced Research Projects Agency (DARPA). Thanks for Kathleen, Robert, Thomas, and Samuel Bradshaw for their helpful comments and suggestions.

References

- [1] M. Johnson and A. H. Vera, “No AI is an Island: The Case for Teaming Intelligence,” *AI Mag.*, vol. Spring, pp. 16–28, 2019.
- [2] D. McDermott, “Artificial intelligence meets natural stupidity,” *ACM SIGART Bull.*, 1976.
- [3] J. M. Bradshaw, R. R. Hoffman, M. Johnson, and D. D. Woods, “The Seven Deadly Myths of ‘Autonomous Systems,’” *IEEE Intell. Syst.*, vol. 28, no. 3, pp. 54–61, May 2013.
- [4] D. A. Norman and S. W. Draper, Eds., *User Centered System Design*. Taylor & Francis, 1986.
- [5] C. Billings, “Human-centered aircraft automation: A concept and guidelines,” 1991.
- [6] R. R. Hoffman and S. V. Deal, “Influencing versus Informing Design, Part 1: A Gap Analysis,” *IEEE Intell. Syst.*, vol. 23, no. 5, pp. 78–81, 2008.
- [7] M. Johnson, J. Bradshaw, P. Feltovich, C. Jonker, B. van Riemsdijk, and M. Sierhuis, “The Fundamental Principle of Coactive Design: Interdependence Must Shape Autonomy,” in *Coordination, Organizations, Institutions, and Norms in Agent Systems VI*, vol. 6541, M. De Vos, N. Fornara, J. Pitt, and G. Vouros, Eds. Springer Berlin / Heidelberg, 2011, pp. 172–191.
- [8] T. W. Malone and K. Crowston, “The interdisciplinary study of coordination,” *ACM Comput. Surv.*, vol. 26, no. 1, pp. 87–119, 1994.
- [9] E. Salas, M. A. Rosen, C. S. Burke, and G. F. Goodwin, “The wisdom of collectives in organizations: An update of the teamwork competencies,” in

Team effectiveness in complex organizations: Cross-disciplinary perspectives and approaches, New York: Routledge/Taylor & Francis Group, 2009, pp. 39–79.

- [10] P. J. Feltovich, J. M. Bradshaw, W. J. Clancey, and M. Johnson, "Toward an Ontology of Regulation: Socially-Based Support for Coordination in Human and Machine Joint Activity," in *Engineering Societies in the Agents World VII*, vol. Lecture No. G. O'Hare, M. O'Grady, A. Ricci, and O. Dikenelli, Eds. Heidelberg, Germany: Springer, 2007, pp. 175–192.
- [11] H. H. Clark and S. Brennan, "Grounding in communication," *Perspect. Soc. Shar. Cogn.*, pp. 127–149, 1991.
- [12] H. H. Clark, *Using language*. Cambridge [England] ; New York: Cambridge University Press, 1996.
- [13] S. W. Tu, M. G. Kahn, M. A. Musen, L. M. Fagan, and J. C. Ferguson, "Episodic Skeletal-Plan Refinement Based on Temporal Data," *Commun. ACM*, vol. 32, no. 12, pp. 1439–1455, Jan. 1989.
- [14] C. Billings, *Aviation automation: The search for a human-centered approach*. 1997.
- [15] J. G. J. Gao and J. Lee, "Extending the decision field theory to model operators' reliance on automation in supervisory control situations," *IEEE Trans. Syst. Man, Cybern. - Part A Syst. Humans*, vol. 36, no. 5, pp. 943–959, 2006.
- [16] G. Klein, D. D. Woods, J. M. Bradshaw, R. R. Hoffman, and P. J. Feltovich, "Ten Challenges for Making Automation a 'Team Player' in Joint Human-Agent Activity," *IEEE Intell. Syst.*, vol. 19, no. 6, pp. 91–95, 2004.
- [17] N. B. Sarter and D. D. Woods, "How in the world did we ever get into that mode? Mode error and awareness in supervisory control," *Hum. Factors J. Hum. Factors Ergon. Soc.*, vol. 37, no. 1, pp. 5–19, 1995.
- [18] E. L. Wiener, "Human factors of advanced technology (glass cockpit) transport aircraft," (*Nasa-Cr-177528*), no. June, p. 222, 1989.
- [19] A. Kirlik, R. a Miller, and R. J. Jagacinski, "Supervisory control in a dynamic and uncertain environment II: A process model of skilled human environment interaction," *IEEE Trans. Syst. Man. Cybern.*, vol. 23, pp. 929–952, 1993.
- [20] E. Rovira, K. McGarry, and R. Parasuraman, "Effects of imperfect automation on decision making in a simulated command and control task," *Hum. Factors*, vol. 49, no. 1, pp. 76–87, 2007.
- [21] K. L. Myers and D. N. Morley, "Directing Agent Communities: An Initial Framework," in *Proceedings of the IJCAI Workshop on Autonomy, Delegation, and Control: Interacting with Autonomous Agents*, 2001.
- [22] M. R. Endsley and E. O. Kiris, "The Out-of-the-Loop Performance Problem and Level of Control in Automation," *Hum. Factors Ergon. Soc.*, vol. 37, no. 2, pp. 381–394, 1995.
- [23] M. R. Endsley, B. Bolté, and D. G. Jones, *Designing for Situation Awareness: An approach to user-centered design*. Taylor & Francis, 2003.
- [24] D. B. Kaber, E. Onal, and M. R. Endsley, "Design of automation for telerobots and the effect on performance, operator situation awareness, and subjective workload," *Hum. Factors Ergon. Manuf.*, vol. 10, no. 4, pp. 409–430, 2000.

- [25] R. Parasuraman, T. Sheridan, and C. Wickens, "A model for types and levels of human interaction with automation," *Syst. Man Cybern. Part A, IEEE Trans.*, vol. 30, no. 3, pp. 286–297, 2000.
- [26] M. R. Endsley, "Level of automation effects on performance, situation awareness and workload in a dynamic control task," *Ergonomics*, vol. 42, no. 3, pp. 462–492, 1999.
- [27] N. B. Sarter, D. D. Woods, and C. E. Billings, "Automation Surprises," in *Handbook of human factors and ergonomics*, Second., G. Salvendy, Ed. Wiley, 1997, pp. 1926–1943.
- [28] E. L. Wiener and R. E. Curry, "Flight-deck automation: promises and problems," *Ergonomics*, vol. 23, no. 10, pp. 995–1011, 1980.
- [29] R. C. Mayer, J. H. Davis, and F. D. Schoorman, "An Integrative Model Of Organizational Trust," *Acad. Manag. Rev.*, vol. 20, no. 3, pp. 709–734, 1995.
- [30] M. Johnson and J. M. Bradshaw, "The role of interdependence in trust," in *Trust in Human-Robot Interaction*, C. S. Nam and J. B. Lyons, Eds. London, UK: Academic Press, 2020, pp. 379–403.
- [31] R. R. Hoffman, "A taxonomy of emergent trusting in the human-machine relationship," in *Cognitive Systems Engineering: The Future for a Changing World*, 2017, pp. 137–164.
- [32] R. R. Hoffman, M. Johnson, J. M. Bradshaw, and A. Underbrink, "Trust in automation," *IEEE Intell. Syst.*, vol. 28, no. 1, 2013.
- [33] M. Johnson, J. M. Bradshaw, P. J. Feltovich, C. M. Jonker, B. M. van Riemsdijk, and M. Sierhuis, "Coactive Design: Designing Support for Interdependence in Joint Activity," *J. Human-Robot Interact.*, vol. 3, no. 1, pp. 43–69, 2014.
- [34] M. Johnson, M. Vignati, and D. Duran, "Understanding Human-Autonomy Teaming through Interdependence Analysis," in *Symposium on Human Autonomy Teaming*, 2018.
- [35] J. Carff, M. Johnson, E. M. El-Sheikh, and J. E. Pratt, "Human-robot team navigation in visually complex environments," in *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2009*, 2009.
- [36] M. Johnson, J. Carff, and J. E. Pratt, "Coactive Design For Human-MAV Team Navigation," in *International Micro Air Vehicle conference and competition*, 2012.
- [37] M. Johnson *et al.*, "Team IHMC's Lessons Learned from the DARPA Robotics Challenge Trials," *J. F. Robot.*, vol. 32, no. 2, pp. 192–208, 2015.
- [38] J. M. Bradshaw, P. J. Feltovich, and M. Johnson, *Human-agent interaction*. 2011.
- [39] M. Johnson *et al.*, "Team IHMC's Lessons Learned from the DARPA Robotics Challenge Trials," *J. F. Robot.*, vol. 32, no. 2, pp. 192–208, Mar. 2015.
- [40] M. Johnson *et al.*, "Team IHMC's Lessons Learned from the DARPA Robotics Challenge: Finding Data in the Rubble," *J. F. Robot.*, vol. 34, no. 2, pp. 241–261, Mar. 2017.
- [41] M. S. Zach, "Unmanned tactical autonomous control and collaboration coactive design," Calhoun, the NPS Institutional Archive, 2016.

- [42] C. Beierl and D. Tschirley, "Unmanned tactical autonomous control and collaboration situation awareness," 2017.
- [43] W. Zachary, M. Johnson, R. Hoffman, T. Thomas, A. Rosoff, and T. Santarelli, "A Context-based Approach to Robot-human Interaction," *Procedia Manuf.*, vol. 3, 2015.
- [44] J. M. Bradshaw *et al.*, "Sol: An Agent-Based Framework for Cyber Situation Awareness," *KI - Künstliche Intelligenz*, vol. 26, no. 2, pp. 127–140, 2012.
- [45] K. H. Lee, K. W. L. Chua, D. S. M. Koh, and A. L. S. Tan, "Team Cognitive Walkthrough: Fusing Creativity and Effectiveness for a Novel Operation," 2019, pp. 117–126.
- [46] K. H. Lee, K. W. L. Chua, D. S. M. Koh, and A. L. S. Tan, "Team Cognitive Walkthrough: Fusing Creativity and Effectiveness for a Novel Operation," in *Advances in Intelligent Systems and Computing*, 2019, vol. 824, pp. 117–126.
- [47] T. Koolen *et al.*, "Summary of Team IHMC's virtual robotics challenge entry," in *IEEE-RAS International Conference on Humanoid Robots*, 2015, vol. 2015-Febru, no. February, pp. 307–314.
- [48] K. Chua Wei Liang, M. Johnson, T. Eskridge, and B. Keller, "AOA: Ambient Obstacle Avoidance Interface," in *The 23rd IEEE International Symposium on Robot and human Interactive Communication*, 2014, pp. 18–23.
- [49] M. Johnson, J. M. Bradshaw, and P. J. Feltovich, "Tomorrow's Human–Machine Design Tools: From Levels of Automation to Interdependencies," *J. Cogn. Eng. Decis. Mak.*, p. 155534341773646, Oct. 2017.
- [50] G. Calhoun and M. Draper, "Effect of level of automation on unmanned aerial vehicle routing task," *Proc. Hum.*, 2009.
- [51] H. Li and C. Wickens, "Stages and levels of automation in support of space teleoperations," *Hum. Factors ...*, 2014.
- [52] M. Johnson, J. M. Bradshaw, P. J. Feltovich, C. Jonker, B. van Riemsdijk, and M. Sierhuis, "Autonomy and interdependence in human-agent-robot teams," *Intell. Syst. IEEE*, vol. 27, no. 2, pp. 43–51, 2012.
- [53] D. B. Kaber and M. R. Endsley, "The effects of level of automation and adaptive automation on human performance, situation awareness and workload in a dynamic control task," *Theor. Issues Ergon. Sci.*, vol. 5, no. 2, pp. 113–153, 2004.
- [54] M. S. Whalley *et al.*, "Flight test results for a new mission-adaptive autonomy system on the rascal juh-60a black hawk," in *American Helicopter Society 72th Annual Forum*, 2016.
- [55] L. Onnasch, C. D. Wickens, H. Li, and D. Manzey, "Human Performance Consequences of Stages and Levels of Automation," *Hum. Factors*, vol. 56, no. 3, pp. 476–488, May 2014.