Uncertainty Scope and Sphere of Influence

Matthew Johnson[1], Jeffrey M. Bradshaw[1], Paul J. Feltovich[1], Catholijn Jonker[2], Birna van Riemsdijk[2] and Maarten Sierhuis[3]

[1]Florida Institute for Human and Machine Cognition (IHMC), Pensacola, FL, USA

[2]EEMCS, Delft University of Technology, Delft, The Netherlands

[3]PARC, Palo Alto, California, USA

{mjohnson, jbradshaw, pfeltovich}@ihmc.us,

{c.m.jonker, m.b.vanriemsdijk}@tudelft.nl, maarten.sierhuis@parc.com

ABSTRACT
In this paper, we present two novel concepts for understanding interdependence in a human-robot system; *uncertainty scope* and *sphere of influence*. Uncertainty scope is about the magnitude and range of the uncertainty introduced by a change in the human-machine system. Sphere of influence provides a more detailed understanding of who is interdependent and how they are interdependent. We use Interdependence Analysis Tables which are unique in the way they model the interdependence among participants in the joint activity. We have demonstrated how these new concepts can be applied using our experimental domain and demonstrated how they can lead to greater insight in understanding the experimental results.

Categories and Subject Descriptors
H.1.2 [**Models and Principles**]: User/Machine Systems – *human factors.*

General Terms
Design, Human Factors.

Keywords
Interdependence, Collaboration, Human-Robot Interaction, Joint Activity, Coactive Design, BW4T

1. INTRODUCTION
In contrast to the nearly exclusive emphasis of the past on military and industrial applications, today's goal for robotics also envisions collaborative or cooperative systems that enhance human experience and productivity in everyday life. This goal captures the expectations of future robotic systems in virtually every kind of application, and the more sophisticated roles they are expected to play [2].

Paradoxically, however, although people are increasingly looking for "collaborative" or "cooperative" robots, these terms are still difficult to define—and even more challenging to map to engineering guidelines. So, we come to the question: exactly what makes a collaborative or cooperative robot? The Coactive Design approach [8] suggests that support for interdependence is both the distinguishing feature of collaborative robots and the central organizing principle informing interaction design.

In order to study interdependence in the context of a human-robot system, we developed a simple joint activity testbed in which to analyze joint activity [9]. We have begun a series of experiments using this testbed with the goals of developing a richer understanding of interdependence, developing analysis tools and methodologies, and eventually producing principles and guidelines to assist system developers as they make design choices in the envisioned sophisticated human-robot systems of the future. The hypothesis for our last experiment [8] was that as autonomy increased, system performance and user preference would increase—up to a point. We predicted that at some point, the benefits from increasing autonomy would be outweighed by its disadvantages. In this case, the major disadvantage was the reduced user visibility ("opacity") of what the system was doing as autonomy increases. Our expectations were borne out, as shown in Figure 1, demonstrating both the anticipated increase in opacity and inflection point (increase followed by a decrease) in user preference. This inflection in performance is due to the competing factors of reduction in workload and increase in opacity as autonomy is increased. While the predictions held, they cried out for explanation of what specifically caused the reversal and why it occurred between treatments 3 and 4.

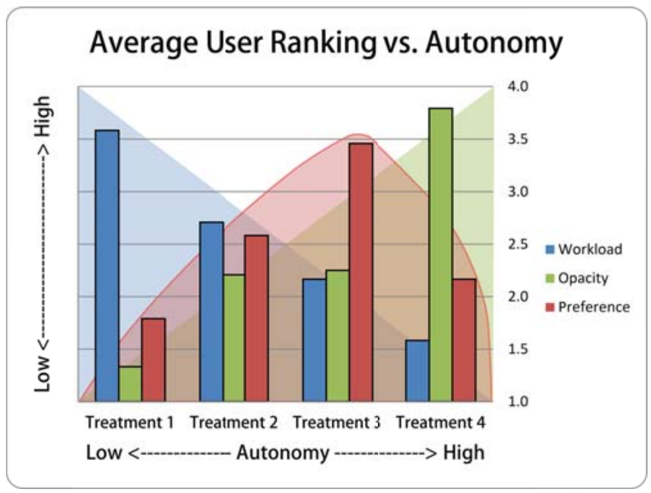

Figure 1 **Exit survey results of the participant's rankings for the autonomy treatments. The results show the expected decrease in workload expected from increased autonomy. The results show reduced awareness (increased opacity) with increased autonomy as we predicted.**

In this paper, we present two novel concepts for understanding interdependence in a human-robot system; *uncertainty scope* and *sphere of influence*. Along with these concepts we present novel analysis techniques for examining each. We apply these concepts to help understand the results in Figure **1**, both demonstrating the technique and providing a deeper analysis of the results. Before we introduce the concepts, we will discuss some background and related work as well as provide a brief description of the BW4T domain.

2. BACKGROUND AND RELATED WORK

Traditional approaches to human-robot systems design usually employ task decomposition and allocation. The most well-known example of this approach is supervisory control [15], in which people allocate tasks to one or more machines and then monitor their performance. To address requirements for variable task allocation in different situations, there has been interest dynamic and adaptive function allocation—also known in slightly different forms as adjustable autonomy, dynamic task allocation, sliding autonomy, flexible autonomy, or adaptive automation.[1] In each case, the system must decide at runtime which functions to automate and to what level of autonomy [12]. The challenge with these approaches is identifying what to adjust and when to adjust it. This challenge is further complicated because it is often difficult to predict the impact a change may have on the system as a whole in a given context. March and Simon point out that "one peculiar characteristic of the assignment problem is that, if taken literally, problems of coordination are eliminated" [10]. This is because approaches based on allocation unrealistically tend to ignore what March and Simon describe as "the contingent character of activities" [10]. Any significant form of collaboration cannot be fully addressed through mere task decomposition and allocation. Such an approach succeeds only when all subtasks can be addressed *independently* of one another. However, it is the joint nature of key tasks that defines the heart of collaborative activity—and it is the effective management of *interdependence* that makes such work possible. Therefore, effective management of systems with autonomy requires an understanding of the impact a change may have on the interdependence in the human-machine system.

Past research on interdependence in the social sciences includes the work of Thompson [16]. The three types of interdependence he identified (pooled, sequential and reciprocal) have some relevance for human-machine design, but are insufficient to cover the nuances of collaboration. Other research communities have developed Hierarchical Task Analysis (HTA) [1] as a method for identifying and decomposing complex tasks. Cognitive Task Analysis (CTA) [4][13] has extended this methodology to include a representation of the knowledge and reasoning required to perform tasks. Goal-Directed Task Analysis (GDTA) [5] is a type of CTA which includes situation awareness requirements. These approaches, while useful, typically focus only on tasks. Interdependence in a team can be due to more than just the task.

3. THE DOMAIN

In order to follow the analysis, it is necessary to provide an understanding of the BW4T domain. BW4T is a multi-player simulation environment for joint activity (Figure 2). The task environment is composed of nine rooms containing a random assortment of blocks and a drop off area for the goal. Each player controls one simulated robot. This robot can be moved between rooms to pick up and drop off blocks. In one recent experiment, teams were composed of two players: one human and the other a software agent. The two players worked toward a shared team goal, namely to pick up and deliver colored blocks to a drop zone in a specified order. Neither player can see what the other is doing. Players can only see objects within their current room, and only one player is allowed inside a given room at a time. Human players control their own robot to move around and deliver blocks. They also control their agent partner through menu selection of commands, allowing the agent teammate to be directed in order to provide assistance toward achieving the team goal.

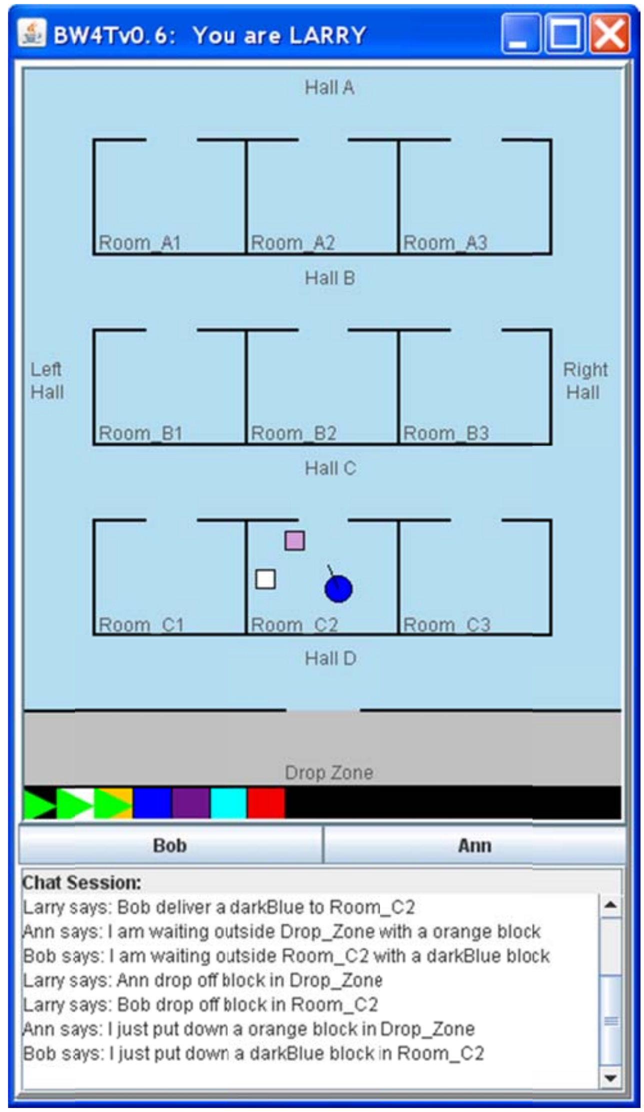

Figure 2 BW4T Interface for a single human player with one agent partner. Rooms contain blocks which must be delivered according to the goal sequence specified below the map.

For this experiment, we designed our agent player to work "perfectly," meaning that they never failed to perform the task assigned by the human teammate. The algorithm chosen as the basis for the agent behavior is based on the most common approach observed in human players. The algorithmic solution is shown on the left side of Figure 3. The main goal (a color sequence) is composed of several subgoals (individual colors). To achieve any given subgoal, you simply find the block of the appropriate color and deliver it. The overall task (yellow box in Figure 3) can be thought of as being composed of several "find" tasks (red box in Figure 3) and several "deliver" tasks (blue box in Figure 3). These tasks are in turn composed of decision and action primitives. The action primitives include going to a room,

[1] For more detailed critiques of these approaches, see [3] and [7].

entering the room, going to a block, picking up a block, and putting down a block. The two main decisions are: 1) whether to look for a block or to deliver a block, and 2) which room to go to in order to look for a block.

In order to compare the effects of increasing and decreasing the autonomy of the agent players, we needed to define different autonomy treatments. Additionally, we needed some way to provide ordinal relationships between the treatments so we could define which treatment had "more" autonomy. The vertical black lines or bands in Figure 3 are used to indicate the portion of the algorithm that activity covers "autonomously." Within the black band, the agent can be considered at Sheridan's [19] highest level, as the agent will perform everything necessary to complete the task specified by the band. Outside the band the agent is at the lowest level and is completely reliant on the human for all decisions and actions. The neglect tolerance [15], meaning how long the agent can perform its duties effectively without human intervention, is indicated by the length of the band. Note that the length of the band covers a portion of the algorithm and does not directly correlate to length of time since some tasks take longer than others. However, longer bands cover more sections of the algorithm, thus, in general, they entail more autonomy.

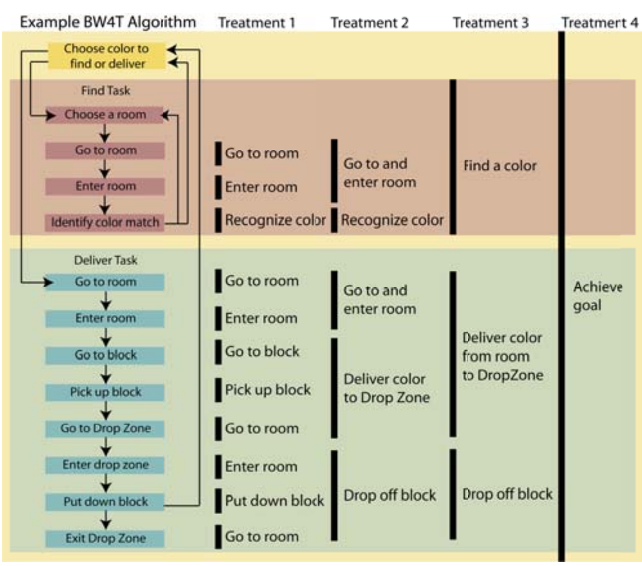

Figure 3 Defining Autonomy Treatments for BW4T. The overall task (yellow box) can be thought of as being composed of several "find" tasks (red box) and several "deliver" tasks (blue box). These tasks are in turn composed of decision and action primitives.

Treatment 1 requires the human player to direct the agent player using only the action primitives. This is very much like teleoperation. Our remaining three treatments were carefully designed to isolate different aspects. Treatment 2 simply combines action primitives, but does not include any decisions. Treatment 3 mainly adds a single decision (which room to search). Treatment 4 adds a second decision (whether to find or deliver) which makes the task "fully autonomous". No coordination or collaboration support is provided other than task assignment and reporting of the completion status after the task is completed.

4. Interdependence Analysis Techniques

We now present two novel concepts for understanding interdependence in a human-robot system; uncertainty scope and sphere of influence. Along with these concepts we present analysis techniques for each. Application of these techniques is explain using the BW4T domain described in Section 3.

4.1 Uncertainty Scope

Uncertainty scope is about the magnitude and range of the uncertainty introduced by a change in the human-machine system. This uncertainty can come from a variety of sources including decisions, iterative tasks, dynamic environments, adversarial environments, and agent competency. The scope is about the potential for entropy or unpredictability in the system. The magnitude of the scope is related to the number of options and their probabilities [18]. The range of the scope is about the span of activity over which that uncertainty is distributed. Uncertainty that affects only isolated aspects of a task has a smaller uncertainty scope than uncertainty which affects many aspects of a task. Both magnitude and range play a role in overall scope. In general, activities with broader uncertainty scope are more likely to have unexpected negative side effects.

The uncertainty's impact is realized in different ways. Often it results in suboptimal choices due to the lack of information. It can also lead to redundant behavior and inefficiency, resource conflict, blocking behavior, delays due to hesitation and reduced confidence or trust in the system. All of these have the potential to have a negative impact on system performance and user satisfaction.

It is important to note that the uncertainty scope does not necessarily correspond to what people sometimes call the "level of autonomy", though it often may. This is because the amount of activity does not always directly relate to uncertainty. For example, a lot of predictable activity can have less uncertainty than a single unpredictable act.

Uncertainty scope can be useful during the design stages to compare different design choices. It can also be used to analyze existing systems to better understand the sources of uncertainty in a system and the potential impact of that uncertainty. This could be used to help decide what to change or improve in the system. It can also highlight the aspects of the system where more advanced collaboration techniques might be necessary to compensate for the uncertainty. Uncertainty scope could also be valuable for real-time adjustable autonomy strategies that have to make decisions about what to adjust and when to adjust it.

4.1.1 Analysis Technique

In order to analyze a system's uncertainty, we suggest a rough estimate using something similar to Shannon's entropy [18] from information theory. Shannon's uncertainty measures the average ambiguity of the received signal. We are interested in the average ambiguity in an autonomous block of activity. Another way to view it is that we are interested in determining the average missing information content for a particular autonomy configuration. The magnitude of the uncertainty is related to the number of options and their probabilities. For example, a predictable action adds no uncertainty. Similarly, a sequence of predictable actions adds no uncertainty. However, a decision does add some uncertainty based on the probability distribution. In general, the more options available and the more equally distributed the probabilities, the greater the magnitude of uncertainty. Iterative cycles can also add some uncertainty if the number of cycles is a variable dependent on circumstance. This technique can typically provide a rough estimate of the magnitude of uncertainty. More rigorous numerical methods may produce improved quantitative metrics, but it is unclear whether the higher cost of such techniques would provide results of greater value.

We must also consider the range of affect the uncertainty will have on the system. This includes considering the number of activities affected as well as the number of team members affected. This will be easier explained with the following example.

4.1.2 *Application Example*

In our experimental domain, agent competence is not a factor and only decisions and iterative cycles play a role in uncertainty. Note that in real world applications, this would not be the case due to potentially frail autonomy and dynamic or adversarial environments.

Consider treatment 2 in Figure 3, which basically combines some action primitives into a single sequential task (blue grouping in Figure 4). The scope is limited to a single sequential process without any possible variation and thus no uncertainty.

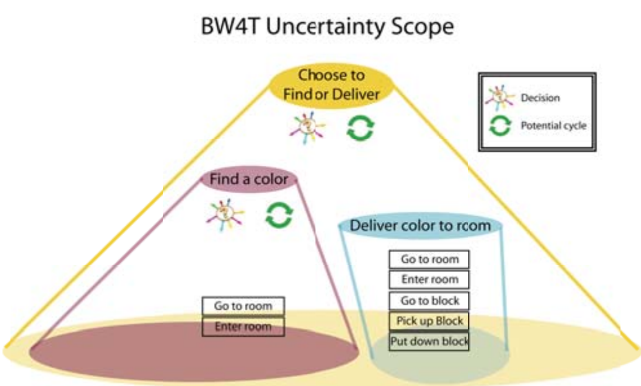

Figure 4 Relative sizes of Uncertainty Scope for different aspects of BW4T algorithm. The delivery task has more activity, but less uncertainty than the find task, while the choice to find or deliver is broader in scope than both.

Now consider treatment 3 in Figure 3, in which the ability to find a color is automated. Though this makes use of fewer action primitives than the delivery task (red grouping in Figure 4), it includes a decision about which room to search which adds potential variability. Even though the algorithm is static and easily understood, there is still variability in context because the human will be unaware of the agent's location and its current knowledge about block locations, both of which effect the decision (i.e., contingent characteristics of activity [13]). There are nine possibilities for which room to search. The task to find a color also includes a cyclical process to allow searching repeatedly until the color is found. This also adds uncertainty, because the human will be unaware if the agent will complete the task in one cycle or nine (number of rooms in BW4T). While automating the find task increases the magnitude of the uncertainty, the range is still fairly small (the individual search task). The only impact on other's activities is the potential for room occupancy conflicts, and the probability of that is fairly low (1:9).

Treatment 4 is much broader in scope (yellow in Figure 4), though it merely includes one additional decision. The scope of this decision ranges across the entire task domain and encompasses the decision in treatment 3, as shown in Figure 4. The range includes goal selection which is inherently coupled to the other team members.

One of the results from our experiment [8] was that treatment 4 had three times the number of occurrences of players blocking each other from entering rooms, as shown in Figure 5. Uncertainty scope can help explain why there were more resource conflicts. Both treatments introduce autonomous decisions that could lead to resource synchronization conflicts. However, treatment 4 has a much broader uncertainty scope. This increase in uncertainty leads to less predictability which we attribute as the cause for the increased conflicts.

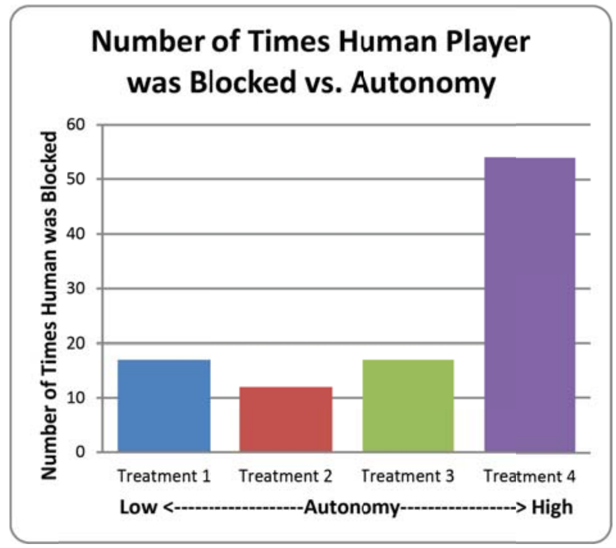

Figure 5 Number of times a human player is blocked by their agent teammate during BW4T experiment

4.2 SPHERE OF INFLUENCE

The scope of the uncertainty provides some intuition about the magnitude and range of uncertainty a change in autonomy might have, but it does not provide specific details about how each team member's activities are specifically affected by that uncertainty.

Sphere of influence is about identifying the specific interdependencies among participants in a joint activity. The concept is being introduced as a way to capture how interconnected team members are, as well as how interconnected any aspect of the task is to the team members. Besides identifying how parties are connected, it also provides some idea of how much they are connected, though it is not a quantitative measure. Sphere of influence is more than just a single party's dependence. It is about the reciprocal nature of the relationship. It also represents more than strict dependence because it also considers soft interdependencies [10].

In some sense, sphere of influence helps define what constitutes common ground in joint activity. It maps out the knowledge and understanding that the participants must develop and hold in common. It can also help define the coordination behaviors and needs of a human-robot system as autonomy is being developed. Rather than focusing solely on supplementing limitations, sphere of influence also helps identify the potential opportunities in a system through inclusion of soft interdependencies. This can potentially provide guidance to designers trying to develop novel supportive behaviors in their advanced collaborative systems.

Much like uncertainty scope, sphere of influence is potential useful both at design time and at run-time. While is provides less information about the magnitude of uncertainty, it provides much more specific detail about how the human-machine system is interdependent.

4.2.1 Interdependence Analysis Table

To assess and visualize the sphere of influence we developed an analysis methodology that we call the *Interdependence Analysis* (IA) Table. It is similar in spirit to Goal Directed Task Analysis [5] in that it extends task analysis with relevant situation awareness information upon which the task is dependent. We extend these types of analysis tools by:

- allowing for more types of interdependence than just task dependency
- allowing for competency assessment
- allowing for soft constraints
- allowing for potentially concurrent activity
- representing other participants in the activity, either individually, as categories or even by roles
- representing team interdependencies including interdependence across goals or based on organizational structure

The IA process begins like traditional HTA, decomposing the task to an appropriate level of granularity. The tasks column of Figure 6 shows a hypothetical example plan for using a particular tool to analyze a rock sample on location. We then add a task dependencies column to capture requirements in a manner similar to CTA or GDTA. However, we do not limit this to informational needs and include such things as sensing needs, perception needs, decision needs and action needs. This enables consideration for supplementing team members in any of these areas. Our Figure 6 example only lists a single task dependency per task item, but there can be multiple dependencies. Next we add the team interdependencies.

The team interdependencies column is composed of several columns. The first column represents the perspective of the individual we are focusing on, which we call the focal entity (FE). The remaining columns represent the other participants in the joint activity. The columns can be specific individuals, categories, or even roles. While specific individuals may be useful on small-scale teams, we expect roles to be particularly valuable for larger scale teams.

The FE column in Figure 6 is focusing on a robot that is charged with analyzing a particular rock sample. In this column we use colors to represent the self-sufficiency (i.e., capability) and self-directedness (i.e., authority) though one could potentially use a sub column for each. Green indicates that robot is able to meet the requirements of the associated task dependency. For example, the FE robot is able to navigate to the tool room without any assistance. Yellow indicates the robot is able to meet the requirement but it is not 100% reliable. For example, in order to get to the sample the robot may need to traverse a complex outdoor environment so it's reliability at meeting this requirement might be less than 100%. Orange indicates the robot can contribute to the task but cannot do it on its own. In our example, the robot may need to load a heavy tool that it is not capable of lifting on its own, but can lift with assistance. Finally, red indicates that the robot cannot meet the requirement. While the FE robot can recognize the tool with some reliability less than 100%, it has no capability of recognizing the rock sample on its own.

The other columns under team interdependence address the teammates. In this example, there are three; a robot caddie (which we will simply refer to as caddie to distinguish it from the FE robot) to assist with lifting, a human operator capable of driving the robot remotely, and a geology specialist. The color scheme used is different than that of the first column. For the teammates, green indicates that the FE does not need anything from this teammate. Yellow represents soft interdependence, meaning the focal robot does not require help as a necessity but could *benefit* in some way from this team member. For example, since the robot has less than perfect object recognition, any of the human team members could provide assistance. Soft dependencies are not limited to supplementing poor performance. They can also indicate useful support for perfectly capable individuals. For example, the robot can navigate to the tool room on its own, but the operator may be able to provide useful information about a passageway being blocked. Orange is a stricter constraint in that it requires a contribution to the effort. In our example, this is shown as the caddie being required to contribute to lifting the heavy tool. Red indicates complete dependence, such that the teammate must meet the requirement completely. In our example, the geologist is required to recognize the rock sample of interest.

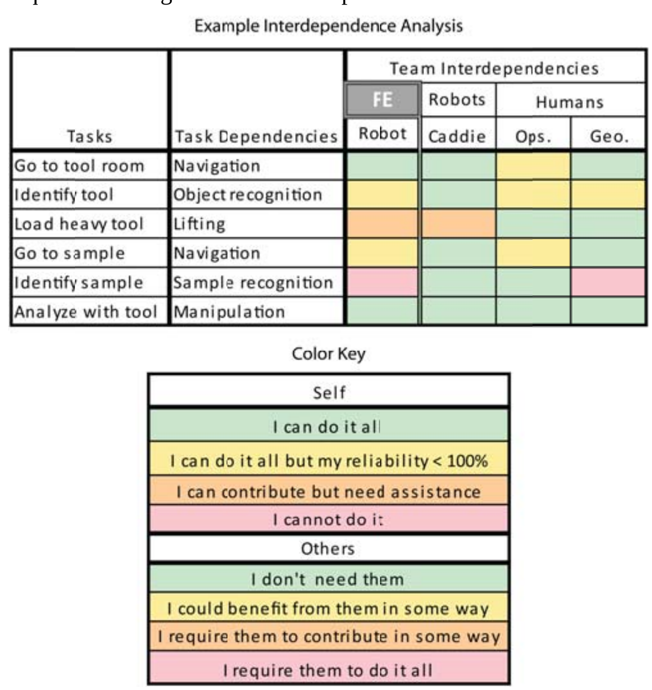

Figure 6 Example Interdependence Analysis Table for a hypothetical rock sample analysis task

The IA Table provides a way to map and visualize the sphere of influence. It shows who is interdependent and how they are interdependent. It is useful for comparing changes in autonomy, which we will show in the next section. Although we focus on an individual, we will still consider the team as a whole. We can also focus on different individuals by creating additional, thus providing different perspectives. IA extends traditional analysis tools by including the aspects of behavior that distinguish it as collaborative or cooperative.

4.2.2 *Application Example*

Referencing Figure 3 as explanation of the treatments, we will now apply these concepts to analyze the results from the experiment [8]. Specifically, we are looking to understand why the opacity change in treatment 4 was significantly different then treatment 3 and why it had such a negative impact on the system, as demonstrated by the subjective assessment (Figure 1) and quantitative error measures such as inefficiency as shown in (Figure 7).

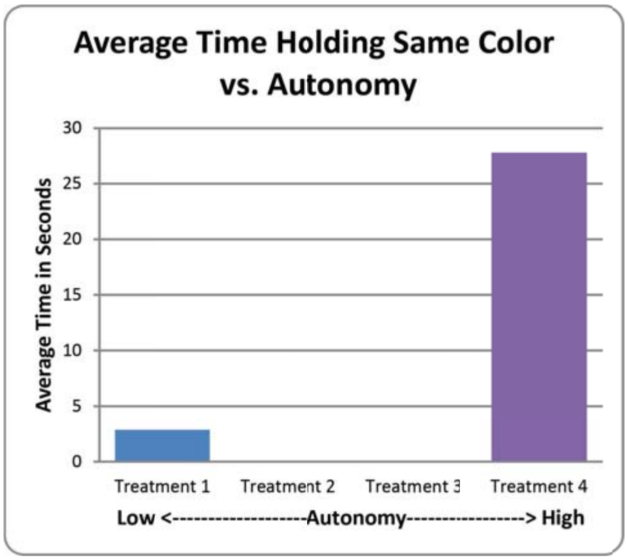

Figure 7 Average time holding the same color (inefficiency)

Space precludes a full presentation of the full IA Table, so we will select two specific pieces that highlight the differences. Consider the Deliver-color-from-room-to-room task in treatment 3. We represent this in Table 1. The FE column indicates that this task and all of its nuances can be done completely autonomously (i.e. all the actions and decisions are green). The only interdependencies are resource synchronization to coordinate entering and leaving rooms, indicated by yellow shading. This is also capable of being done independently since the player can observe the status of the door themselves. However, performance can be improved if teammates coordinate the use of rooms, which are restricted to one player at a time.

Now consider the Choose-to-find-or-deliver task in treatment 4. Table 2 shows that the task itself has fewer sub-tasks, but that the task has more interdependencies. The FE column indicates that this task can be done autonomously (i.e. the decision can be made), but that not all of the information is attainable by the agent. Since teammates are not directly observable in this domain, the status and intention is not available once automated. No coordination mechanisms were included in this example, so sharing world state is also inhibited. Interestingly, both of these are available in treatments with less autonomy because the human is making the decision for the agent players. As we automate these decisions we relinquish control to some extent and can potentially decouple the natural mechanisms that allow us to deal with these interdependencies.

Table 1 Interdependence Analysis Table for the automation of the delivery-color-from-room-to-room task in treatment 3 of the BW4T domain algorithm.

Tasks	Task Dependence	Team Interdependence	
		FE (Agent)	Human
Go the room	Knowing where room is located		
	Knowing how to get there		
	Going to a room		
Enter the room	Synchronization door access		
	Enter a room		
Go the block	Identifying colored blocks		
	Going to block		
Pick up block	Synchronization door access		
	Enter a room		
Go the drop zone	Knowing where room is located		
	Knowing how to get there		
	Going to a room		

Table 2 Interdependence Analysis Table for the automation of the choose-to-find-or-deliver task in treatment 4 of the BW4T domain algorithm

Tasks	Task Dependence	Team Interdependence	
		FE (Agent)	Human
Choose to find vs. deliver a colored block	Knowing the current goal status		
	Knowing the location of the color block		
	Knowing what teammates are currently finding		
	Knowing what teammates are currently delivering		
	Knowing which subgoal is most appropriate		
	Decide to find or deliver		

Comparing Table 1 and Table 2 it should be evident why treatment 4 has a higher average time holding the same color (i.e. redundant activity) shown in Figure 7. This is a direct result of poor goal coordination due to the reduced awareness of the teammate's status as a result of the automation of the Choose-to-find-or-deliver task.

The Coactive design approach highlights the reciprocal and mutually constraining nature of joint activity, so it is often useful to look at a task from each player's perspective. Table 3 compares treatment 3 to treatment 4 for both the agent and human perspectives. In treatment 3 the robot is dependent on the human for quite a bit, but it is irrelevant because the human is responsible for the decision. Treatment 4 allows the robot to make the decision, but in doing so inhibits situation awareness in ***both*** players by distributing it between the two players without appropriate coordination mechanisms. Not only is the agent less capable of making the proper decision due to the inability to obtain appropriate situation awareness information, but the human suffers in the same way. Had we provided the human with a global viewpoint, enabling the missing situation awareness information, this would have been an asymmetric relationship.

Table 3 Interdependence Analysis of the human and agent player and comparison between autonomy treatments in the BW4T domain

Tasks	Task Dependence	Team Int.		Team Int.	
		Agt.	Hum.	Hum.	Agt.
Treatment 3 Choose to find vs. deliver a colored block	Knowing the current goal status				
	Knowing the location of the color block				
	Knowing what teammates are currently finding				
	Knowing what teammates are currently delivering				
	Knowing which subgoal is most appropriate				
	Decide to find or deliver				
Tasks	Task Dependence	Team Int.		Team Int.	
		Agt.	Hum.	Hum.	Agt.
Treatment 4 Choose to find vs. deliver a colored block	Knowing the current goal status				
	Knowing the location of the color block				
	Knowing what teammates are currently finding				
	Knowing what teammates are currently delivering				
	Knowing which subgoal is most appropriate				
	Decide to find or deliver				

Although not a factor for this experiment, real world systems will also need to address frail autonomy or limitations in competence typical of real world robotic systems. Table 4 shows an example of a robot with imperfect perceptual abilities (indicated by yellow). This is another aspect not covered by traditional analysis tools. The yellow in the human column indicates the potential for the other human to assist in this task. The table provides insight into the challenges in implementing such supporting behavior. In this case, the human can match colors, but to provide this assistance they would need to know which color is desired by the individual and be able to receive the precepts about the colors in the room, which is not currently available to them, as indicated by the red in Table 4.

Table 4 Interdependence Analysis to identify the sphere of influence for the supporting behavior of identifying colors in the BW4T domain

Tasks	Task Dependence	Team Interdependence	
		FE (Agent)	Human
Finding color in a room	Sensing color blocks in the room		
	Knowing which color is desired		
	Correlating sensing with desired color		

5. CONCLUSION

As part of the Coactive Design Analysis approach, we have introduced the concepts of *uncertainty scope* and *sphere of influence*. Uncertainty scope is about the magnitude and range of the uncertainty introduced by a change in the human-machine system. Sphere of influence provides a more detailed understanding of who is interdependent and how they are interdependent. The IA Table is unique in the way it models the interdependence among participants in the joint activity. It also allows for reciprocal perspective comparisons.

We have demonstrated how these new concepts can be applied using our experimental domain and demonstrated how they can lead to greater insight in understanding the experimental results. These concepts and techniques can be used to guide a designer toward what needs to be coordinated and why. In some sense, it helps define what constitutes common ground in joint activity, the knowledge and understanding that the participants must develop and hold in common. They also help define the coordination behaviors and needs of a human-robot system as autonomy is being developed. The techniques can also be used to better understand the impacts of runtime adjustments to autonomy in other approaches such as adjustable autonomy and dynamic functional allocation. Lastly, rather than focusing solely on supplementing limitations, sphere of influence also helps identify the potential opportunities in a system through inclusion of soft interdependencies. If we truly desire to develop more collaborative human-robot interaction, we must develop a richer understanding of interdependence and these concepts are a step toward this goal.

6. REFERENCES

[1] Annett, J. (2003) Hierarchical Task Analysis. *Handbook of Cognitive Task Design.* Lawrence Erlbaum Associates: London, pp. 17 – 35.

[2] Bradshaw, J. M., Paul Feltovich, and Matthew Johnson. Human-Agent Interaction. In Handbook of Human-Machine Interaction, edited by Guy Boy. Ashgate, 2011, pp. 283–302.

[3] Bradshaw, Jeffrey M., Paul Feltovich, Hyuckchul Jung, Shri Kulkarni, William Taysom, and Andrzej Uszok. "Dimensions of adjustable autonomy and mixed-initiative interaction." In Agents and Computational Autonomy: Potential, Risks, and Solutions. Lecture Notes in Computer Science, Vol. 2969, edited by Matthias Nickles, Michael Rovatsos and Gerhard Weiss, 17-39. Berlin, Germany: Springer-Verlag, 2004.

[4] Crandall, B., G. Klein, and R. R. Hoffman, Working Minds: A Practitioner's Guide to Cognitive Task Analysis Bradford Books, 2006.

[5] Endsley, M. R., B. Bolté, and D. G. Jones, Designing for Situation Awareness: An approach to user-centered design: Taylor & Francis, 2003.

[6] Johnson, M., J. M. Bradshaw, P. J. Feltovich, C. M. Jonker, B. van Riemsdijk, and M. Sierhuis, "Analyzing Autonomy and its relation to Interdependence in Human-Machine Systems," *IEEE Intelligent Systems* vol. Mar/Apr, 2012.(in press.)

[7] Johnson, Matthew, Jeffrey M. Bradshaw, Paul J. Feltovich, Robert R. Hoffman, Catholijn Jonker, Birna van Riemsdijk, and Maarten Sierhuis. Beyond Cooperative Robotics: The Central Role of Interdependence in Coactive Design. IEEE Intelligent Systems, May/June 2011 (vol. 26 iss. 3), pp. 81-88.

[8] Johnson, M., J. M. Bradshaw, P. J. Feltovich, C. M. Jonker, B. van Riemsdijk, and M. Sierhuis, "The Fundamental Principle of Coactive Design: Interdependence Must Shape Autonomy.," in *COIN@AAMAS&MALLOW'10*, 2010, pp. 172-191.

[9] Johnson, M., C. Jonker, B. v. Riemsdijk, P. J. Feltovich, and J. M. Bradshaw, "Joint Activity Testbed: Blocks World for Teams (BW4T) " in *Engineering Societies in the Agents World X*, 2009.

[10] March, J.G., H. A. Simon, and H. S. Guetzkow, *Organizations*, 2nd ed. Cambridge, Mass., USA: Blackwell, 1993.

[11] Olsen, D. R. and M. Goodrich, "“Metrics for Evaluating Human-Robot Interaction," in *Proceedings of PERMIS*, 2003.

[12] Parasuraman, P., T. Sheridan, and C. Wickens, "A model for types and levels of human interaction with automation," *Systems, Man and Cybernetics, Part A, IEEE Transactions on*, vol. 30, pp. 286-297, 2000.

[13] Schraagen, J. M., S. F. Chipman, and V. L. Shalin, "Cognitive Task Analysis," Lawrence Erlbaum Associates, Inc., 2009.

[14] Shannon, C. E. "A Mathematical Theory of Communication." vol. July and Oct: The Bell System Technical J. 27, 1948, pp. 379-423 and 623-656.

[15] Sheridan, T.B., *Telerobotics, automation, and human supervisory control*. Cambridge, Mass.: MIT Press, 1992.

[16] Thompson, J.D., Organizations in action; social science bases of administrative theory. New York,: McGraw-Hill, 1967.