

Analyzing Social Situations for Human-Robot Interaction

Alan R. Wagner

*Georgia Institute of Technology
College of Computing, Atlanta GA*

Technology Square Research Building, Room S27
85 Fifth St. NW,
Atlanta, GA. 30308

alan.wagner@cc.gatech.edu

Ronald C. Arkin

*Georgia Institute of Technology
GVU Center, College of Computing, Atlanta GA
and LAAS/CNRS Toulouse, France*

Technology Square Research Building, Room
S27
85 Fifth St. NW,
Atlanta, GA. 30308

arkin@cc.gatech.edu, rarkin@laas.fr

Abstract – This paper presents an algorithm for analyzing social situations within a robot. We contribute a method that allows the robot to use information about the situation to select interactive behaviors. This work is based on interdependence theory, a social psychological theory of interaction and interpersonal situation analysis. Experiments demonstrate the utility of the information provided by the situation analysis algorithm and of the value of this method for guiding robot interaction. We conclude that the situation analysis algorithm offers a viable, principled, and general approach to explore interactive robotics problems.

Keywords – *Human-robot interaction, social situation, interdependence, situation analysis.*

Biographical notes:

Alan R. Wagner received his M.S. degree in computer science from Boston University and his B.A. in psychology from Northwestern University. He is currently pursuing his doctoral degree under the supervision of Dr. Ronald Arkin at the Georgia Institute of Technology. His ongoing thesis research explores the development of a general framework for characterizing the trustworthiness of a robot's social relationships and the social situations it encounters.

Ronald C. Arkin holds the rank of Regents' Professor and is the Director of the Mobile Robot Laboratory in the School of Interactive Computing at the Georgia Institute of Technology. Dr. Arkin's research interests include behavior-based reactive control and action-oriented perception for mobile robots and unmanned aerial vehicles, hybrid deliberative/reactive software architectures, robot survivability, multiagent robotic systems, biorobotics, human-robot interaction, robot ethics, and learning in autonomous systems. He has over 130 technical publications in these areas. He is the Series Editor for the MIT Press book series Intelligent Robotics and Autonomous Agents. He was elected a Fellow of the IEEE in 2003.

1 Introduction

Sociologists and social psychologists have long recognized the importance of the situation as a determining factor of interpersonal interaction (Kelley et al., 2003; Kelley & Thibaut, 1978; Rusbult & Lange, 2003). Solomon Asch, a renowned psychologist, stated that, “most social acts have to be understood in their setting and lose meaning if isolated.” (as cited in Kelley & Thibaut, 1978 pg. 4). If a goal of artificial intelligence is to understand, imitate, and interact with humans then researchers must develop theoretical frameworks that will allow an artificial system to, (1) understand the situation-specific reasons for a human’s social behavior, and (2) consider the situation’s influence on the robot’s social behavior. Understanding human interactive behavior is critical as it implies that the robot will then be capable of predicting and planning for future interactions and their consequences. Recognition of the situational impacts on a robot’s own interactive behavior is similarly necessary if robots will be expected to operate in the presence of humans in social settings such as the home or the workplace.

This paper contributes an algorithm for extracting situation-specific information and uses this information to guide interactive behavior. For our purposes, a social situation describes the environmental factors, outside of the individuals themselves, which influence interactive behavior. The objectives of this paper are to 1) introduce the human-robot interaction community to the ideas of interdependence theory; 2) present a novel algorithm for situation analysis developed by the authors from interdependence theory that provides a robot with information about its social environment; and 3) demonstrate that the algorithm

provides information that can be profitably used to guide a robot's interactive behavior. Simulation experiments accomplish the final objective. These simulations first demonstrate that the algorithm is applicable to robotics problems involving collaborations among humans and robots and then examine the algorithm's effectiveness across a wide expanse of social situations.

Consider, as a running example, an industrial accident involving a toxic spill and injured victims. A teleoperated robot is assigned to rescue victims and an autonomous robot operates simultaneously to cleanup the spill. During the cleanup, both the human and the robot will select behaviors directed towards the effort. Perhaps due to the properties of the spilled material, the victims need to be cleaned before being rescued. In this case, the success of the cleanup depends entirely on both robots working together. Alternative chemical spills will allow the robot and the human to operate in an independent manner, with victims being rescued separately from the cleanup. In either case, the situation should influence the autonomous robot's decision to coordinate its cleanup behavior with the human or to operate independently. Moreover, the effectiveness of the cleanup will depend on the robot's ability to characterize the situation and to use this characterization to select the appropriate behaviors.

The remainder of this paper begins by first summarizing related research. Next, our algorithm is described, followed by a set of experiments used to examine the algorithm. This article concludes with a discussion of these results and directions for future research.

2 Related work

Many researchers have explored human-robot interaction within a single social situation. Breazeal examines situations involving emotive dialogue between a human and a robot (Breazeal, 2002). Pineau et al. explore an assistive situation concerning elderly residents of a retirement home and a robot (Pineau, Montemerlo, Pollack, Roy, & Thrun, 2003). Several researchers have explored interactive situations involving museum tour guides (see Fong, Nourbakhsh, & Dautenhahn, 2003 for a review). We, however, currently know of no direct consideration of the theoretical aspects of social situations as applied to interactive robots.

Social psychologists, on the other hand, have long considered the situation-specific aspects of interpersonal interaction (Kelley & Thibaut, 1978). The use of social situations for examining social interaction is widespread within both neuroscience (Sanfey, Rilling, Aronson, Nystrom, & Cohen, 2003) and experimental economics (Berg, Dickhaut, & McCabe, 1995). Interdependence theory is a social psychological theory developed by Kelley and Thibaut as a means for understanding and analyzing interpersonal situations and interaction (Kelley & Thibaut, 1978). Interdependence theory began as a method for investigating group interaction processes and evolved over the authors' lifetimes into a taxonomy of social situations categorizing interpersonal interactions (Kelley et al., 2003; Kelley & Thibaut, 1978). It is often described as one of the most influential theories for exploring interpersonal relationships and has been characterized by some as a type of social exchange theory (Sears, Peplau, & Taylor, 1991). The term interdependence specifies the extent to which one individual of a dyad influences the other. Using interdependence

theory as a basis of analysis, psychologists have recently developed an atlas of interpersonal situations that maps social situations to a multi-dimensional interdependence space (Kelley et al., 2003). The social situations that occupy this space are not ad hoc constructions. Rather, they represent real situations experienced by real people in the world (Kelley, 1979). Some situations, such as the prisoner's dilemma, have been the focus of intense research involving human subjects spanning decades (Axelrod, 1984). Thus, if we expect robots to interact with untrained people in real world environments, it is important for robots master these situations.

3 Situation-based human-robot social interaction

Interdependence theory underlies our framework for situation-based human-robot interaction. The following section briefly summarizes the aspects of interdependence theory that are used in this work. Next, an algorithm, which uses aspects of interdependence theory to produce information about social situations is detailed. Afterwards, we develop a complete computational process by which a robot can use perceptual information to guide interactive behavior.

3.1 Interdependence theory

Interdependence theory is based on the claim that people adjust their interactive behavior in response to their perception of a social situation's pattern of rewards and costs. Thus, each choice of interactive behavior by an individual offers the possibility of specific rewards and costs—also known as outcomes—after the interaction. Interdependence theory represents

social situations computationally as an outcome matrix (figure 1). Outcome matrices are the social psychological equivalent to game theory's normal form game. An outcome matrix represents a social situation by expressing the outcomes afforded to each interacting individual with respect to the pairs of behavior choices selected by the dyad. Figure 1 shows the outcome matrix for our toxic spill cleanup example. In the dependent situation in figure 1, the robot receives outcome equal to the number of hazards if both the robot and the human cooperate and choose to cleanup hazards, but receive a zero outcome if they do not cooperate. Critics of interdependence theory often state that (1) it ignores the non-economic aspects of interpersonal interaction such as altruism and (2) that it assumes people are rational, outcome maximizers. Kelley responds to these criticisms directly, stating that the non-economic aspects of interaction can also be included in a description of a person's outcomes and that the theory does not presume either rationality or outcome maximization (Kelley, 1979). Rather, as will be explained shortly, individuals often transform social situations to include the irrational aspects of socialization such as emotion or social bias.

[Figure 1 about here]

[Figure 2 about here]

Kelley and Thibaut conducted a vast analysis of both theoretical and experimental social situations and were able to generate a space that mapped particular social situations to the dimensional characteristics of the situation (Kelley & Thibaut, 1978). This interdependence space (figure 2 depicts three of the four dimensions) is a four dimensional

space consisting of: (1) an interdependence dimension, (2) a correspondence dimension, (3) a control dimension, and (4) a symmetry dimension. The interdependence dimension measures the extent to which each individual's outcomes are influenced by the other individual's actions in a situation. In a low interdependence situation, for example, each individual's outcomes are relatively independent of the other individual's choice of interactive behavior. A high interdependence situation, on the other hand, is a situation in which each individual's outcomes largely depend on the action of the other individual. Correspondence describes the extent to which the outcomes of one individual in a situation are consistent with the outcomes of the other individual. If outcomes correspond then individuals tend to select interactive behaviors resulting in mutually rewarding outcomes, such as teammates in a game. If outcomes conflict then individuals tend to select interactive behaviors resulting in mutually costly outcomes, such as opponents in a game. Control describes the way in which each individual affects the other's outcomes in a situation. In some situations individuals must exchange action for reaction, such as situations involving buying and selling. Alternatively, some situations demand that individuals coordinate their actions to produce a result, as in the rescue of a victim that is too heavy to be saved by one individual alone. Symmetry describes the balance of a situation's outcomes in favor of one individual over another. In a symmetric situation, both individuals have equal influence over their partner's outcomes. Asymmetric situations, on the other hand, place more influence over the situation's outcomes in one individual than in the other.

A matrix's location in interdependence space provides important information relating to the situation. For example, in a situation of low interdependence the robot should generally select the behavior that maximizes its own outcome, because its choice of action will not have a large impact on the outcome of its partner. We term the process of deconstructing a matrix into its interdependence space dimensions *situation analysis*. As will be demonstrated, the information provided by situation analysis can be used to profitably guide interactive behavior selection by a robot.

3.2 The situation analysis algorithm

Situation analysis is a general technique we developed from interdependence theory to provide a robot with information about its social situation. As an algorithm, it can be used in an on-line or an off-line manner to provide information about any social situation represented by an outcome matrix. Thus, in theory, a robot could use situation analysis as a tool to investigate potential social situations it might encounter or situations that have occurred in the past among others. The input to the algorithm is an outcome matrix representing the social situation. The algorithm outputs a tuple, $\langle \alpha, \beta, \gamma, \delta \rangle$, indicating the situation's location in the four dimensional interdependence space. Situation analysis involves 1) deconstructing the outcome matrix into values representing the variances in outcome and 2) the generation of the dimensional values for the interdependence space. Box 1 describes situation analysis algorithmically.

[Box 1 about here]

[Figure 3 about here]

The first step is matrix deconstruction. This procedure iteratively separates the values in the input or raw outcome matrix into three separate matrices (figure 3) (Kelley & Thibaut, 1978). The Bilateral Actor Control (BAC) matrix represents the variance in outcome resulting from the robot's own interactive decisions. This matrix thus quantifies the robot's control over its own outcomes. The Mutual Partner Control (MPC) matrix, on the other hand, represents the variance in outcome resulting from a partner's interactive decisions and thus quantifies a partner's control over the robot's outcomes. Finally, the Mutual Joint Control (MJC) matrix represents the variance in outcome resulting from both the robot's and its partner's joint interactive decisions. In other words, the MJC matrix describes how each individual is affected by his, her, or its joint actions. As depicted in figure 3, all outcome variance occurs in the BAC matrix when deconstructing an independent situation. This procedure results in values for variables BC , PC , JC individually representing the variance of both the robot's and the human's outcomes in the situation. The subscripts in this figure denote the variance of the outcome for the robot (R) and the human (H) respectively.

[Table 1 about here]

Once the variances for the situation have been computed these values can be used to calculate the situation's location in interdependence space. This is accomplished using equations (1-4) from table 1. Equations (1) and (2) are from (Kelley & Thibaut, 1978). Equations (3) and (4) are contributions of this work. Equation (3) subtracts the outcome

resulting from joint action by the individual's from the outcome resulting from partner and individual control. This value is then normalized. Equation (4) subtracts one individual's control over their own outcomes from the other individual's control. This value is normalized with respect to both individual's outcomes. These values constitute the tuple $\langle \alpha, \beta, \gamma, \delta \rangle$, the situation's location in interdependence space.

3.3 Using situation analysis to select interactive behaviors

The situation analysis algorithm presented above begs several questions. Notably, 1) how are the outcome matrices created? 2) How is the location in interdependence space used to control a robot's behavior? 3) Does knowing a situation's location in interdependence afford valuable information for determining which behavior to select? This section addresses each of these questions in turn.

The creation of outcome matrices that accurately reflect a robot or agent's social environment is a current topic of investigation for several research groups. Vorobeychik, Wellman, and Singh, for instance, have explored the use of machine learning techniques to determine the outcome values in an outcome matrix (Vorobeychik, Wellman, & Singh, 2005). Nevertheless, the absence of a general approach for creating outcome matrices has not impeded their use in numerous fields. Neuroscience researchers, for example, use the value of money to directly populate the outcome matrix (see Sanfey, Rilling, Aronson, Nystrom, & Cohen, 2003 as an example). Moreover, a great deal of work has considered the challenge of representing uncertainty within an outcome matrix (Osborne & Rubinstein,

1994). Results range from probabilistic expectations over the utility values to cumulative probability functions that model normative human responses capturing several types of psychological phenomena (Kahneman & Tversky, 1992). Moreover, often the actual values within the cells of a matrix are less important than the relation of one cell to another cell. For example, it is typically more valuable to know which action in an outcome matrix provides maximal reward than it is to know the actual value of the reward provided. We therefore assume that a method for creating the outcome matrix from a social situation exists and that the outcome matrix created accurately reflects the social situation including its uncertainty. For the experiments conducted as part of this research, the number of hazards and victims perceived is used to construct the outcome matrix (figure 1). These matrices expand upon the human-robot cleanup situation described previously. In these examples, both the human and the robot select either an action to rescue a victim or to cleanup a hazard. The outcome for each pair of selected actions, in this case, is a function of the number of victims and hazards in the environment. The functions in figure 1 were selected to give the autonomous robot a preference for cleanups and the teleoperated robot a preference for victims. Preferences such as these might result from the configuration of each robot. In the independent situation, for example, if the robot chooses to cleanup a hazard and the human chooses to rescue a victim, then the human obtains an outcome equal to the number of victims and the robot obtains an outcome equal to the number of hazards. In the dependent condition, on the other hand, positive outcome is only obtained if both the robot and the human select the same action. A situation such as this could occur if victims must be cleaned prior to being rescued.

[Table 2 about here]

Before discussing how this information is used to control a robot's behavior, we consider strategies by which the outcome matrix can be directly used to select actions. The most obvious method for selecting an action from an outcome matrix is to simply choose the action that maximises the robot's outcome. We term this strategy *max_own*. Alternatively, the outcome matrix can be transformed to create a new, different matrix that the robot uses to select a behavior. Table 2 lists several different methods for transforming an outcome matrix. In the case of *max_other* the partner's outcome values are swapped with the robot's outcome values. The *max_joint* transformation, on the other hand, replaces the robot's outcomes with the sum of the robot and its partner's outcome. Once an outcome matrix has been transformed, the *max_own* strategy is used to select an action. This simple technique of transforming the outcome matrix and then using the *max_own* strategy to select a behavior serves as a control strategy and has the benefit of changing the character of the robot's response without consideration of the actual actions involved.

Because the situation analysis algorithm simply provides information, this information could theoretically be used in many different ways to aid action selection. For instance, rules could directly map a situation's location to a particular **action**. Alternatively, the information could be used to select transformations (table 2). One advantage of the latter method is that it does not require knowledge of the actions available to the robot. Rather, the situation's interdependence space location is used to alter the *character* of the robot's response *independent of interactive actions available*. Another advantage of this approach

is that, one can test a specified set of transformations at a given location to determine which transformation is best at that location. In this manner, a mapping of interdependence space location to transformation can be developed which is independent of the individuals interacting and the actions available. As will be discussed in the next section, our initial step for this research was creating this mapping of situation location to transformation.

Finally, does knowing a situation's location in interdependence space afford valuable information? We approached this question empirically by performing two experiments in simulation. The first experiment investigates the value of this information in a practical scenario. The second experiment considers the value of knowing the situation's location over the entire interdependence space.

3.4 Mapping a situation's location to a transformation

A mapping from a situation's location to a transformation can be described formally as the function $f : L \rightarrow T$ where L is the interdependence space location and T is the space of possible transformations. We subdivide the interdependence space into three areas of interest to robotics researchers, namely high interdependence ($\alpha_R \geq 0.75$) and low correspondence ($\beta \leq 0$), high interdependence ($\alpha_R \geq 0.75$) and high correspondence ($\beta > 0$) and low interdependence if $\alpha_R < 0.75$. These areas are abbreviated as l_{hl}, l_{hh}, l_l respectively. The area l_{hl} represents situations in which the robot's outcomes greatly depend on its partner but the robot and the human do not select actions towards the same goal, potentially resulting in poor outcomes for the robot. The area l_{hh} , on the other hand,

describes situations in which the robot's outcomes also greatly depend on its partner and both the robot and the human select actions towards the same goal. Finally, the area l_l represents the location of situations in which the robot's outcomes do not greatly depend on its partner. Thus $L = \{l_{hl}, l_{hh}, l_l\}$ describes the domain of f . The codomain of f is $T = \{ \text{max_own}, \text{min_own}, \text{max_other}, \text{min_other}, \text{max_joint}, \text{min_joint}, \text{max_diff}, \text{min_diff}, \text{min_risk} \}$ (see table 2 for descriptions), the set of transformations considered as part of this work.

Given the preceding description, the challenge then is to determine for each location in L which transformation from T results in the greatest overall net outcome. To do this we created a random matrix and then used the situation analysis algorithm to determine the matrix's location in interdependence space until we had 1000 matrices in each area l_{hl}, l_{hh}, l_l . Random matrices consisted of an empty matrix populated with random numbers between 0 and 24. Next, for every matrix in each area l_{hl}, l_{hh}, l_l , we iterated through the set T altering the matrix according to the transformation's specification (table 2). Afterward, a simulated robot selects the action from the transformed matrix that maximizes its outcome. The robot's simulated partner also selects an action from the original matrix that maximizes its outcome. Finally, the robot's outcome resulting from the action pair (as dictated by the original matrix) is recorded. Figure 7 in section 4.1 graphically depicts this procedure and the other experimental procedures used.

[Table 3 about here]

Table 3 presents the mean outcome resulting from each transformation at each location. The transformation that results in the greatest mean outcome for each location is shaded. Note that the difference in mean outcome for several of the transformations is not great. This lack of difference reflects the similarity of the transform in the particular area of interdependence space. More importantly, it foreshadows the need of a robot to interact with its partner in a variety of situations located at different positions in interdependence space in order to determine the partner's transformation preference or type. The table indicates that *max_own*, *max_joint*, and *min_risk* are the best transformations of the group of possible transformations in low interdependence, high interdependence/high correspondence, and high interdependence/low correspondence situations respectively. From this data the function f mapping interdependence space location to transformation

takes the following form, $f(l^*) = \begin{cases} \textit{max_own} & l^* = l_l \\ \textit{max_joint} & \text{for } l^* = l_{hh} \\ \textit{min_risk} & l^* = l_{hl} \end{cases}$ where l^* is the interdependence

space location generated by the situation analysis algorithm. This function can also be visualized as the decision tree in figure 4.

[Figure 4 about here]

3.5 A computational process for situation analysis

Assuming that outcome matrices can be generated and given the mapping from interdependence location to transformation developed in the preceding section, a computational process can be developed that selects a robot's behavior from its perception

of the situation. This computational process is depicted in figure 5. The right side of this figure depicts a stepwise procedure for generating interactive action from perception. The first step is the creation of an outcome matrix. In our experiments, these either were derived perceptually using matrices in figure 1, or generated by populating an empty matrix with random values. The next two steps consist of the situation analysis algorithm described in section 3.2, which results in an interdependence space tuple. This tuple is then mapped to a transformation using the function f (also depicted in figure 4). The transformation is used to transform the original matrix in the next step. The transformation process results in the construction of an outcome matrix on which the robot can act—the effective situation (Kelley & Thibaut, 1978). In the final step, the robot selects the action in the effective situation that maximizes its own outcome. The left side of figure 5 depicts an example run through the procedure. The next section discusses our empirical examination of this process.

[Figure 5 about here]

4 Experiments and Results

The preceding discussion has described *how* an outcome matrix can be mapped to a location in interdependence space and *how* information about the matrix’s location can be used to select a robot’s interactive action. We have not yet shown, however, that the information afforded by the situation analysis algorithm results in better interactive behavior on the part of the robot. The experiments presented in this section, therefore, examine the value of the information generated by the situation analysis algorithm. Value

here is operationalized as increase in net outcome. Both experiments test the hypothesis that the use of the situation analysis algorithm will result in an increase in net outcome when compared to alternative control strategies. The first experiment uses the computational process from figure 5 to guide a simulated robot's action selection in the cleanup and rescue example described in section 1. The second experiment generalizes the results from the first experiment to the entire interdependence space and compares the algorithm to a larger number of control strategies.

4.1 Situation analysis in practice

To revisit the scenario described in the first section, a teleoperated robot attempts to rescue victims of an industrial accident while an autonomous robot works to cleanup a spill. We considered two scenarios in simulation: one involving greater dependence (high interdependence condition) and another involving little dependence (low interdependence condition). Notionally, because of the properties of the chemical the high interdependence condition requires that the victims be cleaned before being rescued. Thus, in this condition, the robots must both cooperate in order to complete the rescue task successfully. In the low interdependence condition, both robots can operate independently of one another. This scenario is based on the well-studied foraging problem in robotics (Arkin, 1998). Figure 6 depicts the layout. Potential victims and hazards for cleanup are located within a disaster area. A disposal area for hazardous items is located towards the bottom and a triage area for victims is located to the right.

[Figure 6 about here]

This experiment compares the net outcome obtained by both robots as well as the number of victims rescued and hazards cleaned in four separate conditions. In the experimental conditions, the autonomous robot used the computational process depicted in figure 5 to select its action. In the control conditions, the autonomous robot consistently selected the behavior that maximized its own outcome without consideration of its partner (*max_own*). The experimental and control condition were explored in both high interdependence situations and low interdependence situations. A high interdependence situation was created by populating the dependent outcome matrix from figure 1. Similarly, a low interdependence situation was created by populating the independent outcome matrix from the figure 1. Thus, the experiment consisted of the following four conditions: high interdependence-situation analysis, high interdependence-control strategy, low interdependence-situation analysis, low interdependence-control strategy. In all conditions, the teleoperated robot selected the behavior that maximized its own outcome without consideration of its partner (*max_own*). The primary author controlled the teleoperated robot. Because the teleoperated robot employs a static strategy, experimenter bias is eliminated.

Figure 7 describes the experimental procedure used (middle procedure). First, a random number of victims and hazards were generated. Next, the victims and hazards were randomly placed in the environment. In the low interdependence condition, the autonomous robot perceives the number of victims and hazards and uses the independent matrix from figure 1 to create its outcome matrix. In the high interdependence condition, the autonomous robot uses the dependent matrix to create its outcome matrix. The outcome

matrix is then tested using the situation analysis algorithm and the control strategy. The behaviors that the robot selects are actually collections of actions that direct the robot to locate the closest attractor, pickup the attractor, transport the attractor to a disposal area where it is dropped off and finally return to a staging area. The *MissionLab* mission specification system was used. *MissionLab* is a graphical software toolset that allows users to generate mobile robot behavior, test behaviors in simulation, and execute collections of behaviors on real, embodied robots (MacKenzie, Arkin, & Cameroon, 1997).

[Figure 7 about here]

We conducted thirty trials in each of the four conditions. In these experiments, interaction occurs when both individuals (autonomous robot and teleoperated robot, or both simulated robots) are presented with an outcome matrix and simultaneously select actions from the matrix receiving the outcome that results from the action pair. We recorded the number of victims rescued and the hazards collected after each trial. We predicted that the situation analysis algorithm would outperform the control strategy in the dependent condition but not in the independent condition. Independent situations, by definition, demand little consideration of the partner's actions. Thus, in these situations, the autonomous robot's performance is not affected by the actions of the partner. Dependent situations, on the other hand, demand consideration of the partner, and we believed that our algorithm would aid performance in these conditions.

Figure 8 illustrates the results from the cleanup and rescue experiment. The left two bars portray the results for the independent situation. In these conditions, the autonomous robot forages for hazards to cleanup and the human-operated robot forages for victims.

Thus, in all of the 30 trials each robot retrieves either a victim or a hazard. As predicted, both robots faired equally well in this condition.

[Figure 8 about here]

In the dependent condition, the best possible score was thirty. The autonomous robot's use of situation information results in ten additional victims being rescued. Thus, as predicted, in the dependent condition the autonomous robot's use of situation information affords better performance than the robot that does not consider the situation. In this case, the information provided by our algorithm indicates to the autonomous robot that its outcomes for this situation rely on collaboration with its human-operated partner. The control strategy, on the other hand, fails to consider the partner's role even though the situation demands collaboration, hence resulting in poorer performance.

Overall, this experiment demonstrates that the information resulting from an analysis of the social situation can improve a robot's ability to perform interactive tasks similar to collaborative foraging. The algorithm we have proposed successfully uses perceptual stimuli in the environment to produce information about the social situation. Minimally, we have shown the feasibility of our approach and the potential importance of situational considerations in human-robot interaction, ideas which have not been investigated as a part of HRI in the past. Nevertheless, the results of this experiment are limited in several ways. First, the situations encountered as part of the experiment are derived from a limited portion of the interdependence space. Second, only a single control strategy was considered. The next experiment generalizes these results to the entire interdependence space and considers additional controls.

4.2 Situation analysis over the entire interdependence space

Whereas the previous experiment only explored high interdependence or low interdependence outcome matrices, this experiment considers outcome matrices from every corner of the interdependence space. We examine the algorithm's performance over thousands of different matrices representing a broad spectrum of the interdependence space. Because of time-constraints, it was not possible to test each of these matrices using interaction between a human and a robot. Rather, the human was replaced with an agent that selected the behavior that maximized its own outcome without consideration of its partner (*max_own*). The strategy employed by the human in the first experiment and the agent in this experiment were identical.

For this experiment, we also compare the algorithm's performance to four different control strategies. For the first control strategy, the autonomous robot consistently selected the behavior that maximized its own outcome without consideration of its partner (*max_own*). For the second control strategy, the autonomous robot consistently selected the behavior that minimized the difference of its and its partner's outcome (*min_diff*). For the third control strategy, the autonomous robot consistently selected the behavior that maximizes the sum of its and its partner's outcome (*max_joint*). For the final control strategy, the autonomous robot consistently selected the behavior that resulted in the greatest guaranteed outcome (*min_risk*).

Figure 7 describes the experimental procedure used (right procedure). First, a random matrix is created from an empty matrix populated with random numbers between 0 and 24.

The random matrix in this case does not have actions assigned. Hence, these matrices are abstract in the sense that the rewards and costs are associated with selecting one of two non-specified actions. Once a matrix is created, it is presented to both the simulated robot and the agent. Both simultaneously select actions from the matrix receiving the outcome that results from the action pair. The simulated robot uses either situation analysis or a control strategy to determine which action to select from the matrix. This experiment was conducted as a numerical simulation and hence did not occur in a robot simulation environment. In other words, the simulated robot in this case was an agent that selects an action in accordance with the strategy dictated by the experimental condition, but did not actually have to perform the action in an environment. Consequentially, this experiment did not require perceptual generation of the outcome matrix and the actions selected by the agents did not affect the environment.

In order to ensure coverage over the entire space, we examined one hundred trials each consisting of 1000 randomly generated outcome matrices. We recorded the outcome obtained by each individual for the pair of actions selected. We predicted that the net outcome received by the simulated autonomous robot would be significantly greater when the robot used the computational process from figure 5 when compared to the controls. We reasoned that, on average, the information provided by situation analysis would be valuable to the robot for its selection of its behavior. We thus hypothesized that the use of this information would result in a greater net outcome than the control strategies.

Figure 8 presents results for this experiment. The second bar from the left depicts the net outcome using our algorithm. The next four bars to the right indicate the net outcome

for the control conditions. Our algorithm significantly outperforms the controls in all four conditions ($p < 0.01$ two-tailed, for all). The maximum possible outcome for a robot with complete a priori knowledge of all of its partner's actions is also depicted to the left for reference.

[Figure 9 about here]

The results confirm our prediction that use of the situation analysis algorithm results in greater net outcome than does the use of the control strategies. The graph also indicates that our procedure outperforms several different control strategies. Furthermore, the results show that our procedure is beneficial on average to an agent or robot that will face many different social situations from unique locations in the interdependence space. Still, the algorithm performs far below the maximum possible. Better performance could likely be achieved by increasing the size of the domain and codomain of f , the mapping from interdependence space location to transform (from section 3.4). In this work, we subdivided the interdependence space into three areas, denoted l_{hl}, l_{hh}, l_l . Greater subdivision of the space would make better use of the information provided by the situation analysis algorithm. We also limited the number of transformations considered to nine. Additional transformations would increase the algorithm's performance if a novel transformation outperformed all other transformations at some location in the space.

The value of the situation analysis algorithm, as presented in this paper, stems from the very fact that it knows nothing of its interactive partner. The computational process does not assume anything about the partner. Rather it operates only on the information available within the outcome matrix. This is in contrast to game theory, which operates on the

presumption of the partner's rationality (Osborne & Rubinstein, 1994). We expect that the performance of this approach would increase drastically as additional, partner specific, information is provided.

5 Conclusions

This paper has introduced a method for capturing information about social situations and for using this information to guide a simulated robot's interactive behavior. We have presented an algorithm for situation analysis and a computational process for using the algorithm. Our approach is derived from the social psychological theory of interdependence and has close ties to the psychology of human-human interaction (Kelley & Thibaut, 1978). The value of knowing a situation's location in interdependence space has been highlighted with experiments indicating that, on average, this information can aid in selecting interactive actions and that in some situations this information is critical for successful interaction and task performance.

One limitation of our approach is that it requires that the robot's and its partner's utilities as well as the actions available to both individuals be represented in an outcome matrix. Nevertheless, researchers are developing methods to create these outcome matrices automatically (Vorobeychik, Wellman, & Singh, 2005). We also do not address the challenge of managing uncertainty in this article. Much work has already addressed this topic with respect to the outcome matrix (Osborne & Rubinstein, 1994). The uncertainty present in the outcome matrix will result in similar uncertainty in the situation's location in interdependence space.

We have presented one method for using information about a situation's location to guide behavior selection. Our method relates the matrix's location to a transformation of the matrix. For the most part, we have not used all of the information available. We did not, for example, explore the effect of a situation's symmetry on the behavior of the robot. Symmetry describes the balance of control that the robot or its partner has over the other. The value of this dimension could play an important role in determining behavior. We intend to explore this possibility as part of future work. Moreover, we have assumed throughout that the partner consistently selects the *max_own* transformation. The exploration of different partner types will also be the fruits of future work. Additional avenues for future work will also focus on extending these results to real robots. We believe that the embodiment afforded by a real robot will present both new challenges and new opportunities.

In summary, it is our contention that this approach offers a general, principled means for both analyzing and reasoning about the social situations faced by a robot. The development of theoretical frameworks that include situation-specific information is an important area of study if robots are expected to move out of the laboratory and into one's home. Moreover, because this work is based on research which has already been validated for interpersonal interaction, we believe that it may eventually allow an artificial system to reason about the situation-specific sources of a human's social behavior.

Acknowledgments

We would like to thank Zsolt Kira, Yoichiro Endo, and Patrick Ulam for their many helpful comments.

References

- Arkin, R. C. (1998). *Behavior-based robotics*. Cambridge, MA: The MIT Press.
- Axelrod, R. (1984). *The evolution of cooperation*. New York: Basic Books.
- Berg, J., Dickhaut, J., & McCabe, K. (1995). Trust, reciprocity, and social history. *Games and Economic Behavior*, 10, 122-142.
- Breazeal, C. L. (2002). *Designing sociable robots*. Cambridge, MA: The MIT Press.
- Fong, T., Nourbakhsh, I., & Dautenhahn, K. (2003). A survey of socially interactive robots. *Robotics and Autonomous Systems*, 42, 143-166.
- Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk, *Econometrica*, 47, 263-291.
- Kelley, H. H. (1979). *Personal relationships: Their structures and processes*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Kelley, H. H., Holmes, J. G., Kerr, N. L., Reis, H. T., Rusbult, C. E., & Lange, P. A. M. V. (2003). *An atlas of interpersonal situations*. New York, NY: Cambridge University Press.
- Kelley, H. H., & Thibaut, J. W. (1978). *Interpersonal relations: A theory of interdependence*. New York, NY: John Wiley & Sons.
- MacKenzie, D., Arkin, R., & Cameroon, J. (1997). Multiagent mission specification and execution. *Autonomous Robotics*, 4(1), 29-52.

- Osborne, M. J., & Rubinstein, A. (1994). *A course in game theory*. Cambridge, MA: The MIT Press.
- Pineau, J., Montemerlo, M., Pollack, M., Roy, N., & Thrun, S. (2003). Towards robotic assistants in nursing homes: Challenges and results. *Robotics and Autonomous Systems*, *42*, 271-281.
- Rusbult, C. E., & Van Lange, P. A. M. (2003). Interdependence, interaction and relationship. *Annual Review of Psychology*, *54*, 351-375.
- Sanfey, A. G., Rilling, J. K., Aronson, J. A., Nystrom, L. E., & Cohen, J. D. (2003). The neural basis of economic decision-making in the ultimatum game. *Science*, *300*, 1755-1758.
- Sears, D. O., Peplau, L. A., & Taylor, S. E. (1991). *Social psychology*. Englewood Cliffs, New Jersey: Prentice Hall.
- Vorobeychik, Y., Wellman, M. P., & Singh S. (2005). Learning payoff functions in infinite games. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence, Edinburgh, Scotland*, 977-982.

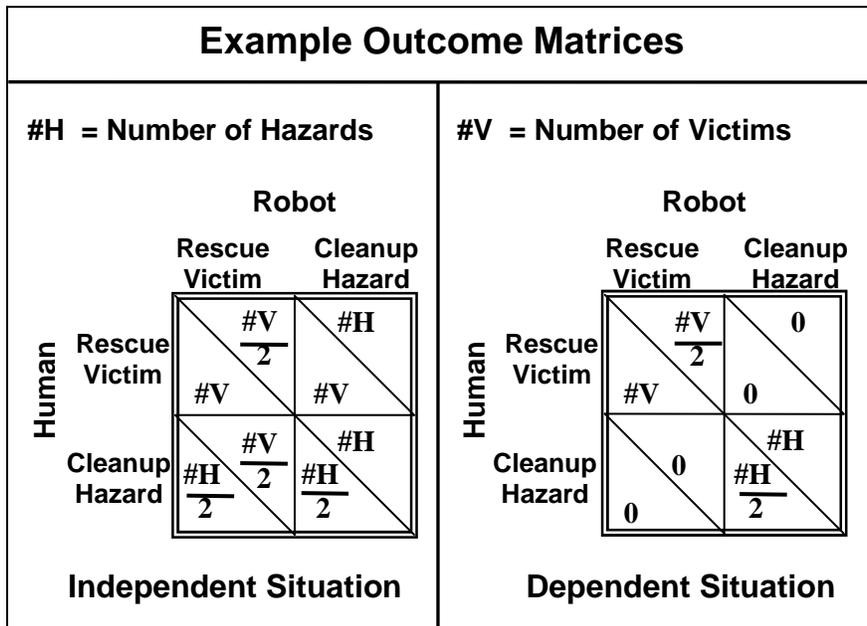


Figure 1. This figure depicts two example outcome matrices for the cleanup of a toxic spill and the rescue of victims by a human and a robot. During any one interaction, both individuals choose to either rescue a victim or clean up a hazard. The outcomes resulting from each pair of choices are depicted in the cells of the matrix. The human's outcomes are listed below the robot's outcomes. In the leftmost matrix, the outcomes for the human and the robot are independent of the other's action selection. In the rightmost matrix, the outcomes of the human and the robot largely depend on the other's action selection.

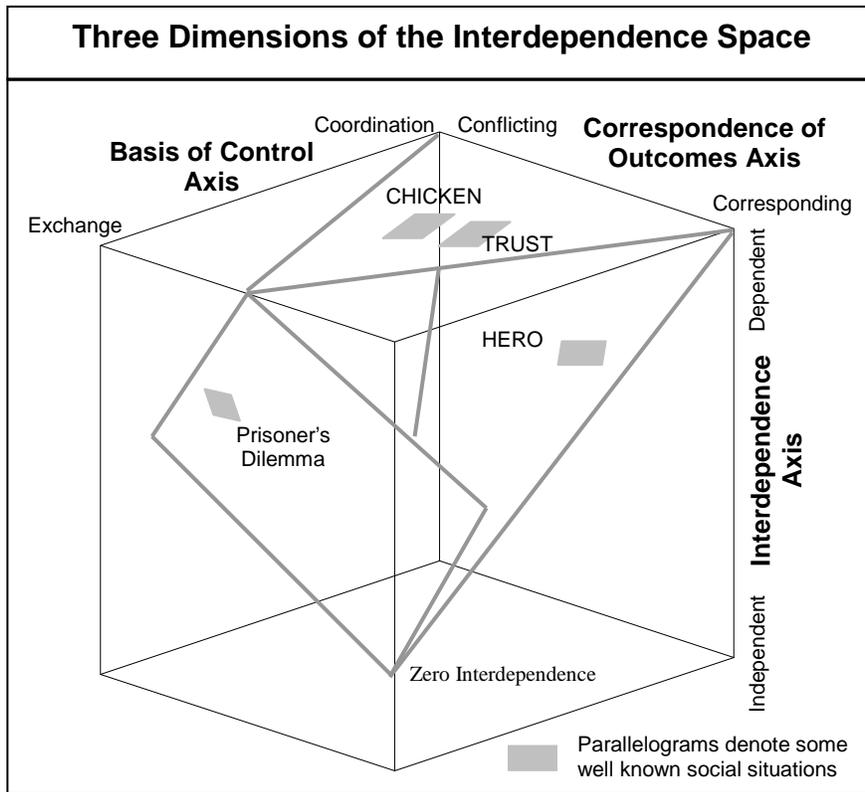


Figure 2. Three dimensions of interdependence space are depicted above (Kelley et al., 2003). Interdependence theory represents social situations computationally as an outcome matrix within this interdependence space. The dimensions depicted above are interdependence, correspondence, and basis of control. Planes within this space denote the location of some well-known social situations, including the prisoner's dilemma game, the trust game, and the hero game. A matrix's location allows one to predict possible results of interaction within the situation.

The Situation Analysis Algorithm

Input: Outcome matrix O .

Output: Interdependence space tuple $\langle \alpha, \beta, \chi, \delta \rangle$.

1. Use procedure from figure 3 to deconstruct the outcome matrix.
2. Use the equations in table 1 to calculate the dimensional values for the interdependence space tuple.
3. **Return** the tuple

Box 1. An algorithm for the analysis of a social situation.

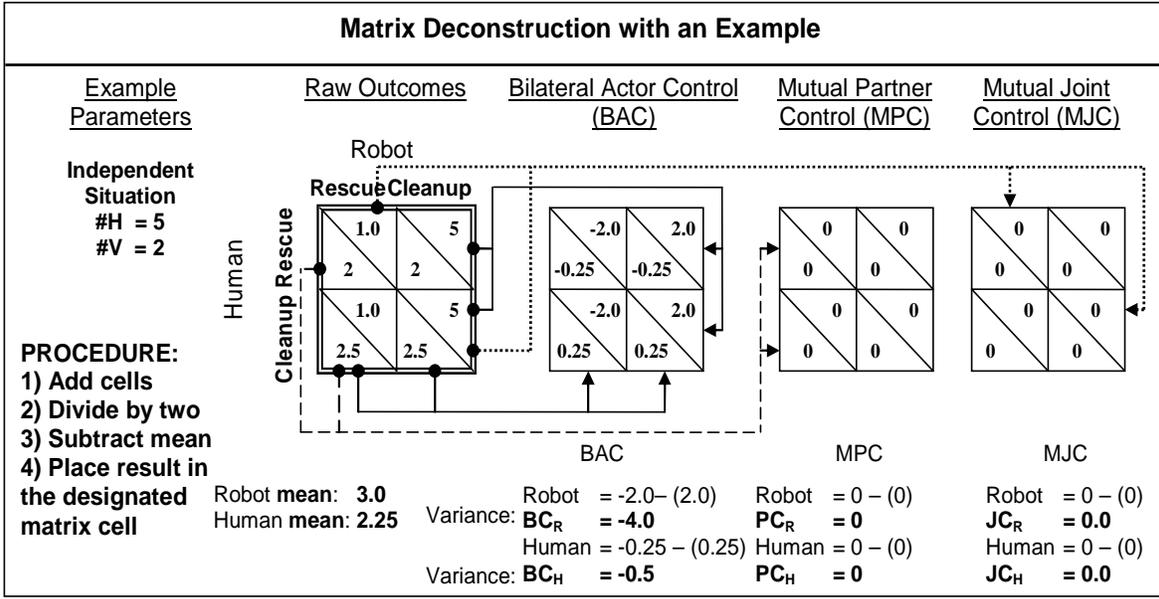


Figure 3. The procedure (from Kelley & Thibaut, 1978) for deconstructing a social situation is presented above. This procedure is an analysis of variance of the outcome matrix that deconstructs the raw outcome matrix into three new matrices (the BAC, MPC, and MJC) representing different forms of control over the situation's outcomes. The outcome values for each of these three matrices are produced from the raw outcome matrix by iteratively 1) adding the noted cells, 2) dividing by the number of actions, and 3) subtracting the individual's mean outcome value. The variances of each matrix type are generated by calculating the outcome range for each choice of behavior and each individual. Because this example is of an independent situation, the MPC and MJC matrices do not vary.

Table 1. Calculation of the interdependence space dimensions given the variances from figure 3.

Dimension	<p style="text-align: center;">Computation ((1) and (2) are from (Kelley & Thibaut, 1978), (3) and (4) were developed by the authors)</p>
Interdependence (α_R, α_H)	$\alpha_R = \frac{(PC_R^2 + JC_R^2)}{(BC_R^2 + PC_R^2 + JC_R^2)} \quad (1)$ <p>Calculate separately for each individual. Range is from 0 for independent situations to +1 for dependent situations.</p>
Correspondence (β)	$\beta = \frac{2(BC_R PC_H + BC_H PC_R + JC_R JC_H)}{(BC_R^2 + PC_R^2 + JC_R^2 + BC_H^2 + PC_H^2 + JC_H^2)} \quad (2)$ <p>Calculate once for both individuals. Range is from -1 for a situation in which the dyad's outcomes conflict to +1 for a situation in which the dyad's outcomes correspond.</p>
Basis of Control (γ)	$\gamma = \frac{4(\sigma - \nu)}{(Sum(sit)^2)} \text{ where} \quad (3)$ $\sigma = (JC_R + JC_H)^2 + (JC_R - JC_H)^2$ $\nu = (BC_R + PC_H)^2 + (BC_H + PC_R)^2 + (BC_R - PC_H)^2 + (BC_H - PC_R)^2$

	Calculate once for both individuals. Range is from -1 for a situation controlled by exchange and to +1 for a situation controlled by coordination. $Sum(sit)$ is a cell by cell sum of the matrix.
Symmetry (δ)	$\delta = \frac{(BC_R^2 + PC_H^2 + JC_R^2) - (BC_H^2 + PC_R^2 + JC_H^2)}{(BC_R^2 + PC_R^2 + JC_R^2 + BC_H^2 + PC_H^2 + JC_H^2)} \quad (4)$ <p>Calculate once for both individuals. Range is from -1 for an asymmetric situation in which individual R depends on H to +1 for an asymmetric situation in which individual H depends on R. The value of 0 denotes a symmetric situation (i.e. mutual dependence).</p>

Table 2. A list of several simple matrix transformations. The list is not exhaustive.

Transformation name	Transformation mechanism	Social character
<i>max_own</i>	No change	Egoism —the individual selects the action that most favors their own outcomes
<i>max_other</i>	Swap partner's outcomes with one's own	Altruism —the individual selects the action that most favors their partner
<i>max_joint</i>	Replace outcomes with the sum of the individual and the partner's outcome	Cooperation —the individual selects the action that most favors both their own and their partner's outcome
<i>max_diff</i>	Replace outcomes with the difference of the individual's outcome to that of the partner	Competition —the individual selects the action that results in the most relative gain to that of its partner
<i>min_diff</i>	Maximize the value of the action that has the minimal difference to that of the partner.	Fairness —the individual selects the action that results in the least disparity
<i>min_risk</i>	Maximize the value of the action that has the greatest minimal outcome	Risk-aversion —the individual selects actions that result in the maximal guaranteed outcomes

Table 3. The cells denote the mean outcome obtained by the transformation at each location. The shaded cells indicate the mean of the best transformation. The confidence interval is included for all values.

Low interdependence		High interdependence/high correspondence		High interdependence/low correspondence	
Transformation	Mean outcome	Transformation	Mean outcome	Transformation	Mean outcome
<i>max_own</i>	13.47 ± 0.46	<i>max_own</i>	15.01 ± 0.39	<i>max_own</i>	14.27 ± 0.41
<i>min_own</i>	10.36 ± 0.46	<i>min_own</i>	8.75 ± 0.40	<i>min_own</i>	7.712 ± 0.38
<i>max_other</i>	11.67 ± 0.43	<i>max_other</i>	15.10 ± 0.36	<i>max_other</i>	7.80 ± 0.37
<i>min_other</i>	11.86 ± 0.43	<i>min_other</i>	10.52 ± 0.42	<i>min_other</i>	12.94 ± 0.42
<i>max_joint</i>	12.90 ± 0.43	<i>max_joint</i>	16.03 ± 0.34	<i>max_joint</i>	13.40 ± 0.42
<i>min_joint</i>	11.16 ± 0.44	<i>min_joint</i>	9.55 ± 0.41	<i>min_joint</i>	10.52 ± 0.43
<i>max_diff</i>	11.41 ± 0.46	<i>max_diff</i>	10.41 ± 0.43	<i>max_diff</i>	9.93 ± 0.47
<i>min_diff</i>	12.08 ± 0.42	<i>min_diff</i>	12.48 ± 0.43	<i>min_diff</i>	12.10 ± 0.41
<i>min_risk</i>	13.08 ± 0.41	<i>min_risk</i>	14.82 ± 0.38	<i>min_risk</i>	14.79 ± 0.37

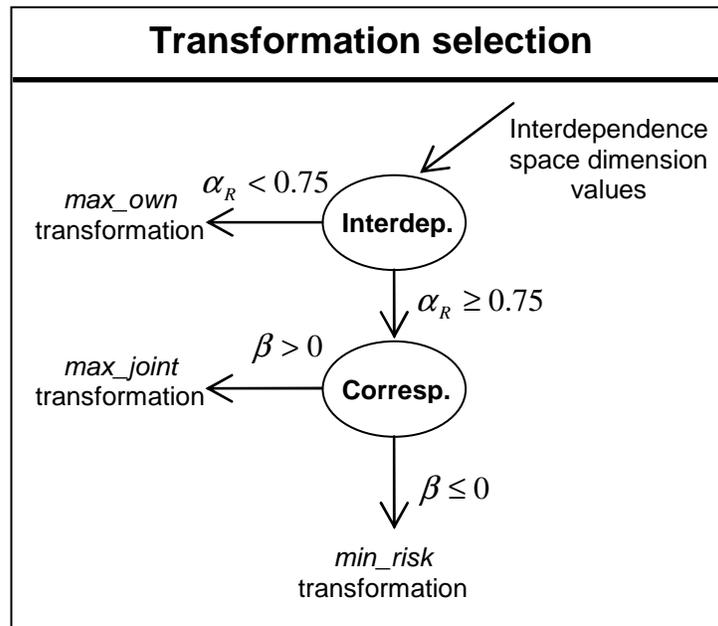


Figure 4. A mapping of interdependence space location to outcome matrix transformation.

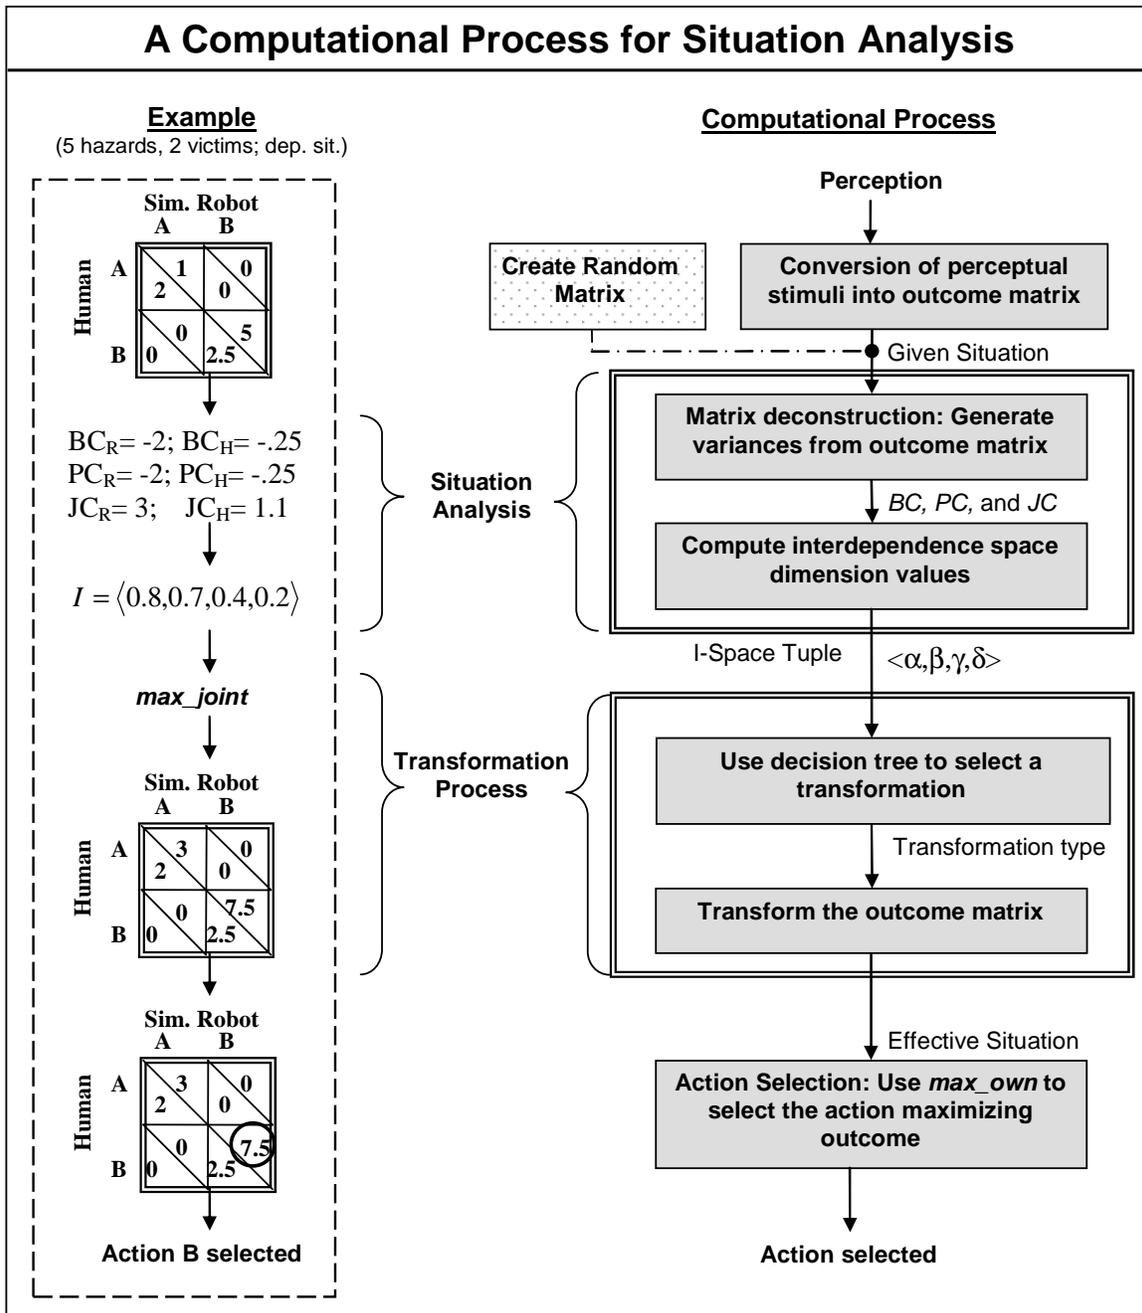


Figure 5. This figure depicts the algorithmic process contributed by this work. The process consists of six steps. The first step generates an outcome matrix. The second step analyzes the matrix's variances. The third step computes the situation's interdependence space dimensions. These two steps constitute the process of situation analysis. The fourth step selects a transformation and in the fifth step, the transformation is applied to the outcome matrix resulting in the effective situation. Steps 4 and 5 constitute the transformation process. Finally, an action is selected.

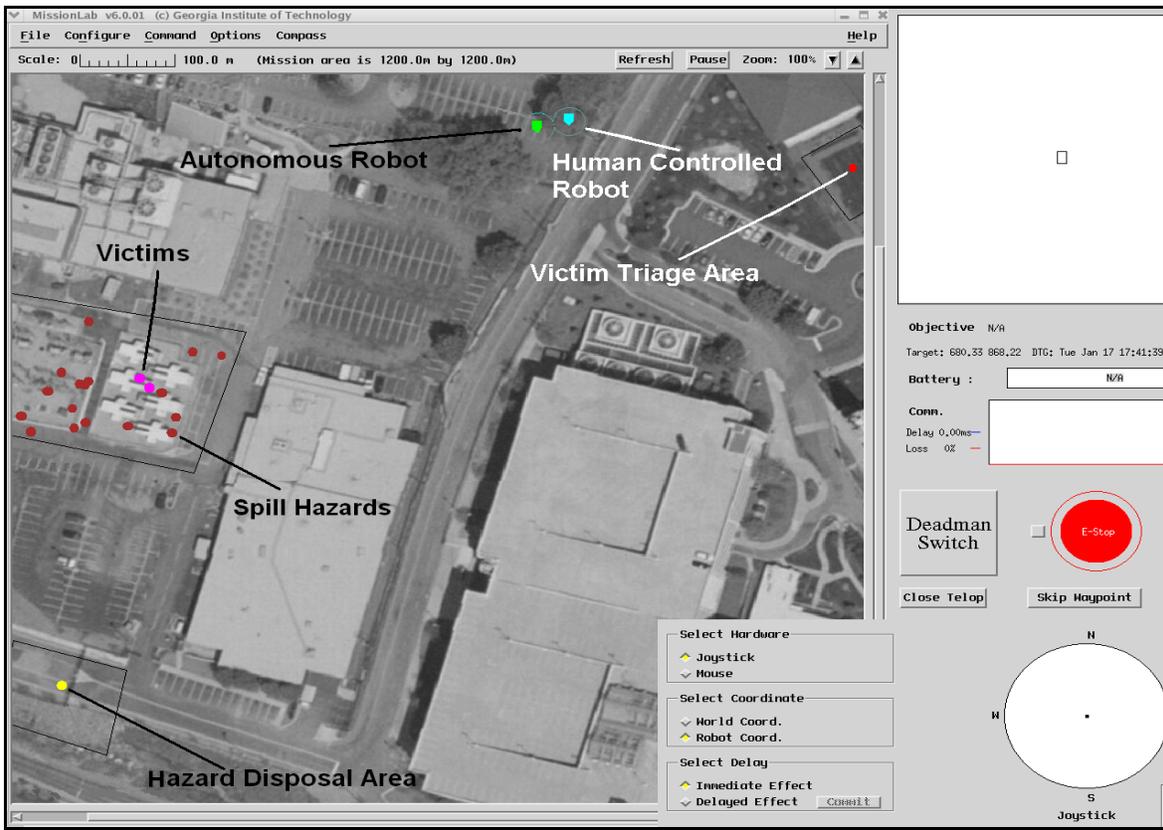


Figure 6. The simulation environment used for the cleanup and rescue experiment is depicted above. The experiment required that a teleoperated robot rescue victims while an autonomous robot performs a cleanup. Experimental conditions included independent versus dependent situations and the use of our situation analysis algorithm versus a control strategy. The teleoperation interface used by the human is depicted the right.

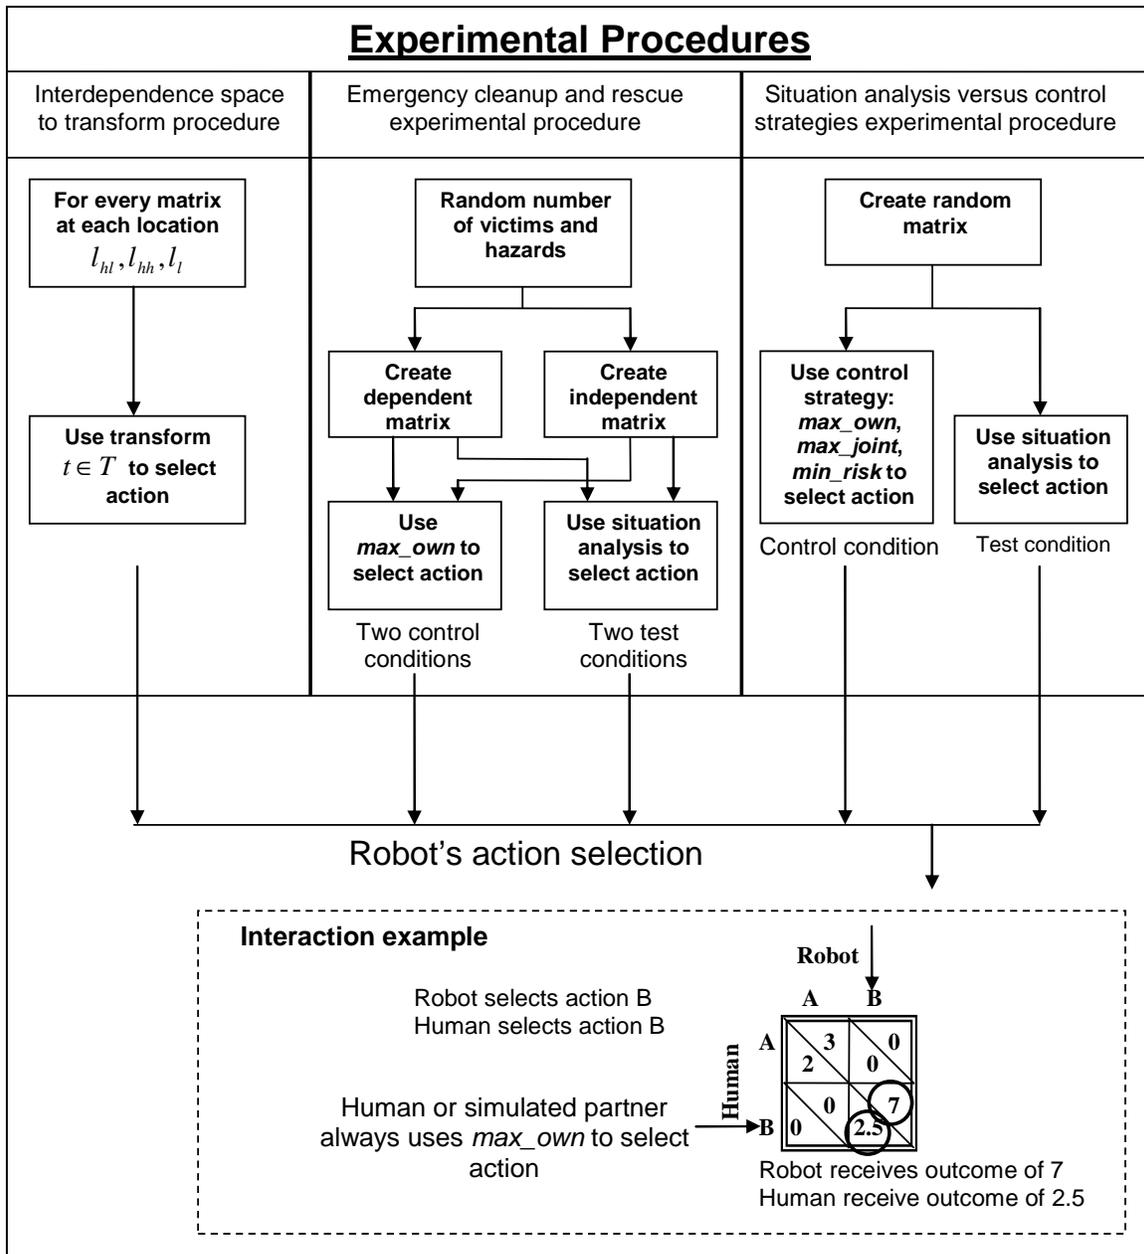


Figure 7. The procedures used to create and use outcome matrices are depicted above. The left side details the procedure used to generate table 3. This procedure first iterates through all matrices in each areas l_{hl}, l_{hh}, l_l and then iterates through the set of transformations to produce the matrix the robot will use to select actions. The middle procedure first creates a random number of victims and hazards. Next, an independent and dependent matrix is created from the number of victims and hazards. Finally, in the control conditions, *max_own* is used to select an action. In the test procedure, situation analysis is used to select an action. The right most procedure, first generates a random matrix and then transforms the matrix with respect to a control matrix or uses situation analysis. The robot selects an action from the transformed matrix. The interaction example at the bottom denotes the method used to determine how much outcome each individual receives from the presentation of an outcome matrix.

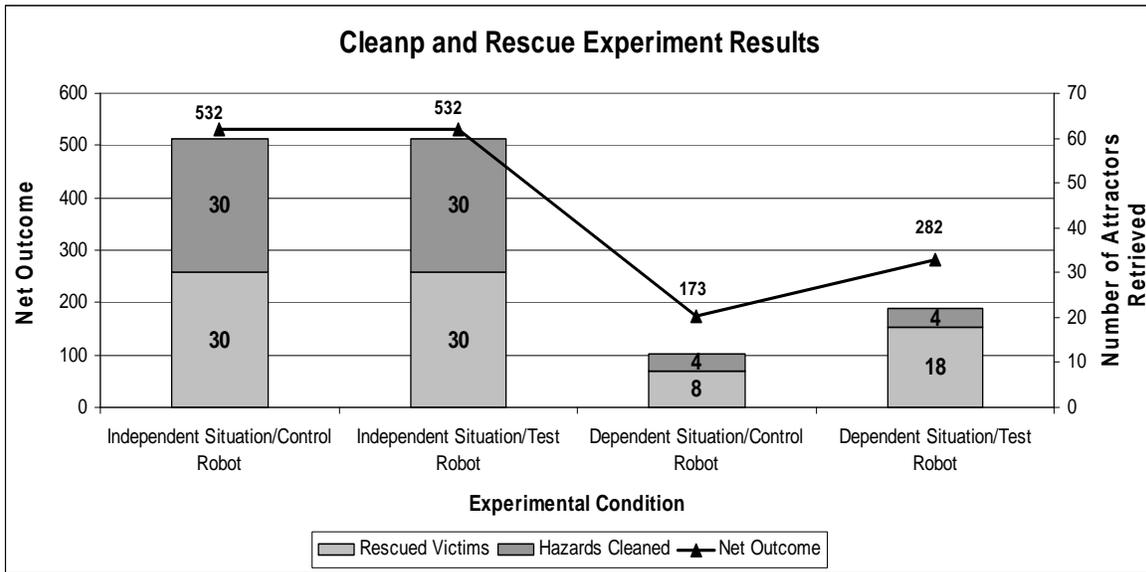


Figure 8. Results for the cleanup and rescue experiment are presented above. The line graph portrays the net outcome for each condition. The bars depict the number of hazards and victims retrieved. Hazards cleaned are shown above the number of victims rescued. The left two bars and line points depict the independent conditions for both the test and the control robot. In these conditions both the control and test robot perform equally well. The right two bars and line points examine the dependent situation. In this situation the test robot outperforms the control robot.

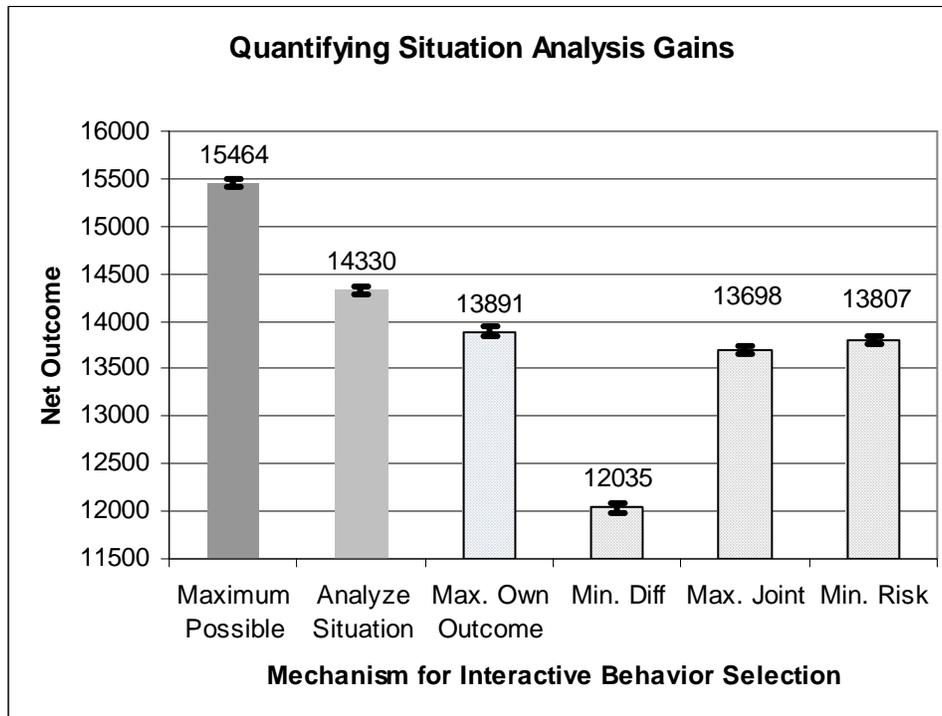


Figure 9. Results of this second experiment are presented above. The second bar from the left indicates the net outcome when our procedure is used. The next four bars are the controls for the experiment. Error bars

indicate 95% confidence interval. Analyzing the situation resulted in the greatest net outcome of when compared to the control strategies. The leftmost bar portrays the maximum possible net outcome.