

A Framework for Situation-based Social Interaction

Alan R. Wagner

Georgia Institute of Technology

GVU Center, College of Computing, Atlanta, GA

alan.wagner@cc.gatech.edu

Ronald C. Arkin

Georgia Institute of Technology

GVU Center, College of Computing, Atlanta GA and

LAAS/CNRS Toulouse, France

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

Abstract – This paper presents a theoretical framework for computationally representing social situations in a robot. This work is based on interdependence theory, a social psychological theory of interaction and social situation analysis. We use interdependence theory to garner information about the social situations involving a human and a robot. We also quantify the gain in outcome resulting from situation analysis. Experiments demonstrate the utility of social situation information and of our situation-based framework as a method for guiding robot interaction. We conclude that this framework offers a principled, general approach for studying interactive robotics problems.

Index Terms – *Human-Robot Interaction, social situation, interdependence, social development.*

I. INTRODUCTION

Many scientists have recently come to recognize the social aspects of intelligence [1]. In contrast to purely cognitive intelligence, which is most often described by problem solving ability and/or declarative knowledge acquisition and usage, social intellect revolves around an individual's ability to effectively understand and respond in social situations [2]. Compelling neuroscientific and anthropological evidence is beginning to emerge supporting theories of social intelligence [3, 4]. From a roboticist's perspective, it then becomes natural to ask how this form of intelligence could play a role in the development of an artificially intelligent robot. As an initial step, one must first consider which concepts are most important to social intelligence.

Social interaction is one fundamental concept [5]. Psychologists define *social interaction* as influence—verbal, physical, or emotional—by one person on another [6]. Furthermore, researchers describe the results of social interaction as a function not only of the individuals involved but also of the social situation [7]. For our purposes, a social situation describes the environmental factors, outside of the individuals themselves, which influences interactive behavior.

Sociologists and social psychologists have long recognized the importance of the situation as a determining factor of interpersonal interaction [7-9]. 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.

This paper contributes a theoretical framework that allows a robot to manage both of these challenges. A general,

established, computational representation for interactive situations that is not tied to specific social environments or paradigms is presented [9]. Moreover, we contribute an algorithm for extracting situation-specific information from this representation and for using this information to guide interactive behavior. Preliminary simulation results examining the framework's effectiveness across a wide expanse of social situations are provided. Finally, we demonstrate that this situation-based framework is applicable to robotics problems involving collaborations among humans and robots.

Consider, as a running example, a human and a robot attempting to cleanup a toxic waste spill—a task of great significance for modern day robots. 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 or of the cleanup environment itself, the robot and the human may need to coordinate their behavior in order to successfully accomplish the cleanup as a team. Alternative scenarios will allow the robot and the human to collaborate in an independent manner. In either case, the situation should influence the robot's decision to coordinate its cleanup behavior with the human or to operate independently. Moreover, the effectiveness of the cleanup will be mitigated by the robot's ability to characterize the situation and to use this characterization to select the appropriate social behaviors.

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

II. RELATED WORK

Human-robot interaction (HRI) is a subfield of AI that combines aspects of robotics, human factors, human computer interaction, and cognitive science [10]. The details of how and why humans and robots interact are focal research areas within HRI [11]. Typically, HRI research explores mechanisms for interaction, such as gaze following, smooth pursuit, face detection, and affect characterization [12, 13].

Many researchers have explored human-robot interaction within a single social situation. Breazeal examines situations involving emotive dialogue between a human and a robot [12]. Pineau et al. explore an assistive situation concerning elderly residents of a retirement home and a robot [14]. Several researchers have explored interactive situations involving museum tour guides [11]. Others have considered interactive situations necessary for search-and-rescue robots [11]. Multi-

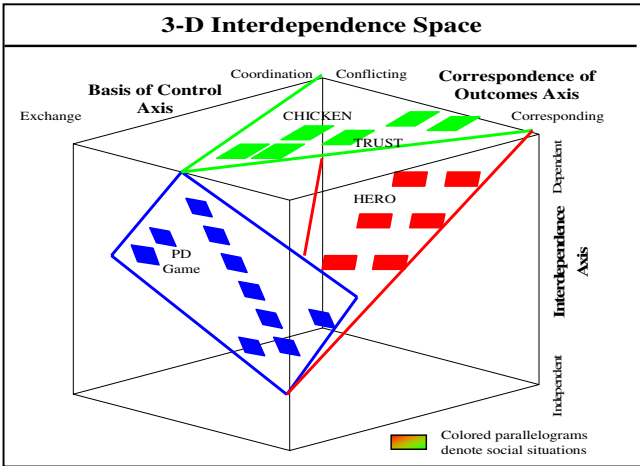


Fig. 1 Three dimensions of interdependence space are depicted above. Interdependence theory represents social situations computationally as an outcome matrix. Planes within this space denote the location of some well-known social situations, including the prisoner’s dilemma game, the trust game, and the hero game. A matrix’s location allows one to predict possible results of interaction within the situation

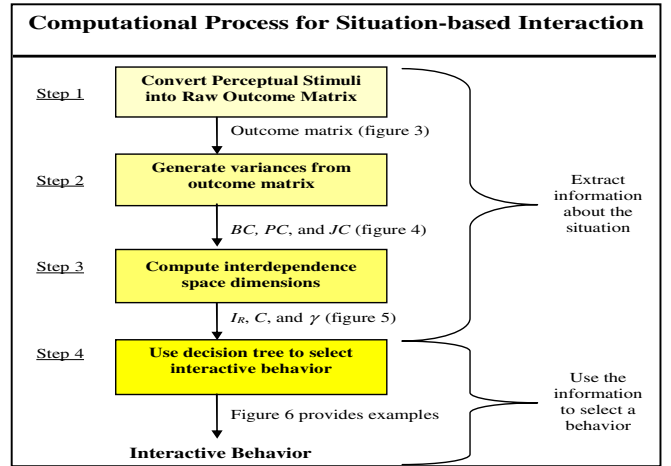


Fig. 2 The algorithmic process contributed by this work is depicted above. The process consists of four steps. The first step generates an outcome matrix. The second step analyzes the matrix’s variances. The third step computes the situation’s interdependence space dimensions. The final step selects an action based on the situation’s position in interdependence space.

agent researchers and sociologists have also explored several different simulated social situations [15]. 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 [9]. The use of social situations for examining social interaction is widespread within both neuroscience [16] and experimental economics [17]. Interdependence theory is a social psychological theory developed as a means for understanding and analyzing interpersonal situations and interaction [9]. The term interdependence describes the extent to which one individual of a dyad influences the other. Expanding upon this theory, psychologists have recently developed an atlas of interpersonal situations that maps social situations to a multi-dimensional interdependence space (fig. 1) [8]. Moreover, these social situations are not ad hoc constructions. Rather, they represent real situations experienced by real people in the world. It is thus important that robots master them. Interdependence theory underlies our framework for situation-based social interaction.

II. A FRAMEWORK FOR SITUATION-BASED SOCIAL INTERACTION

Our situation-based framework translates a robot’s perception of a social situation into action in four steps: (1) matrix construction, (2) situation analysis, (3) interdependence space mapping, and (4) action selection. The first three steps generate information about the situation. The final step uses this information. Figure 2 depicts the computation process. The following section briefly summarizes the aspects of interdependence theory that were used for this work, before describing the algorithm for situation-based social interaction in detail.

A. Interdependence Theory

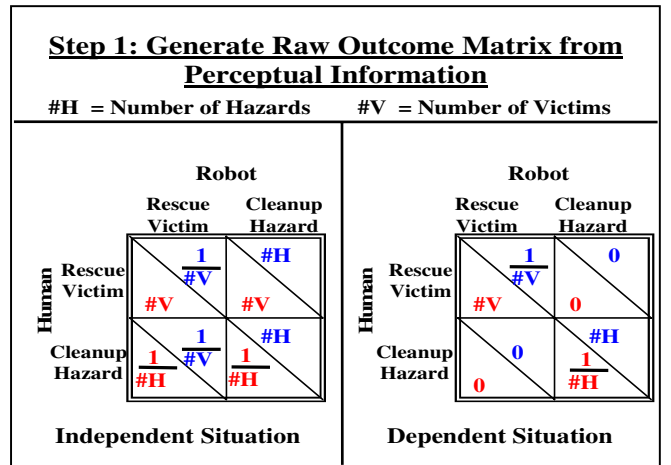


Fig. 3 This figure depicts example outcome matrices for the cleanup of a toxic spill and the rescue of victims by a human and a robot. During any one interaction both individuals choose to either rescue a victim or cleanup a hazard. The outcomes resulting from each pair of choices is depicted in the matrix. The human’s outcomes are listed in red 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 depend entirely on the other’s action selection.

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 after the interaction. Outcome is a term used by psychologists to describe value of the rewards minus the costs. It is therefore important for robots interacting with humans to consider the outcomes their choice of interactive behavior will produce for the human.

Interdependence theory represents social situations computationally as an outcome matrix (fig. 3). Outcome matrices are the social psychological equivalent to the normal form game within game theory. An outcome matrix represents

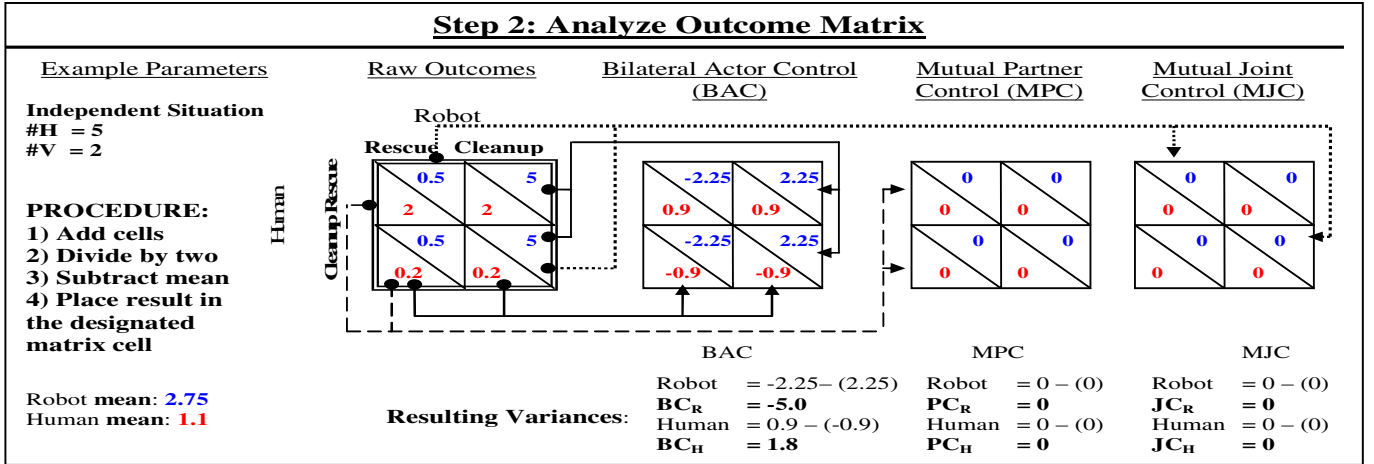


Fig. 4 The procedure (from [9]) for analysing a social situation is presented above. This procedure is an analysis of variance of the outcome matrix that deconstructs the raw outcome matrix into three new matrices (the BAC, MPC, and MJC) representing different forms of control over the situation's outcomes. 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 two, and 3) subtracting the individual's mean. The variances are calculated 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.

a social situation in terms of the outcomes afforded to each interacting individual with respect to pairs of behavior choices selected by the dyad.

The interdependence space (figure 1 depicts the three dimensions used in this study) 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 the each individual's outcomes depend on the action of the other individual. Correspondence describes the extent to which the outcomes of the individuals in a situation are consistent with the outcomes of other individuals. 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 individual's 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.

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 select the behavior that maximizes its own outcome, because its choice of action will not have a large impact on the outcome of its

Step 3: Compute the interdependence space dimensions	
Interdependence (I_R, I_H)	$I_R = \frac{(PC_R^2 + JC_R^2)}{(BC_R^2 + PC_R^2 + JC_R^2)} \quad (1)$ Calculated separately for each individual. The range is from 0.00 for completely independent to +1.00 for completely dependent.
Correspondence of Outcomes (C)	$C = \frac{2(BC_R PC_H + BC_H PC_R + JC_H JC_R)}{(BC_R^2 + BC_H^2 + PC_R^2 + PC_H^2 + JC_R^2 + JC_H^2)} \quad (2)$ Calculated once for both individuals. Range is from -1.00 for conflicting interaction to +1.00 for corresponding interaction.
Basis of Control (γ)	$\gamma = \frac{4(\sigma - \nu)}{(\text{Sum}(\text{sit}))^2} \quad (3)$ $\sigma = (JC_R + JC_H)^2 + (JC_R - JC_H)^2$ $\nu = (BC_R + PC_H)^2 + (BC_H + PC_R)^2 + (BC_R - PC_H)^2 + (BC_H - PC_R)^2$ Calculated once for both individuals. Range is from -1.00 for exchange interaction to +1.00 for coordination interaction. Sum(sit) is a cell by cell summation of the matrix.

Fig. 5 The portrayed equations compute the situation's dimensions in interdependence space. The inputs to these equations are the variances resulting from step two. The equation for calculating interdependence I and correspondence C are from [9]. The authors developed the equation for basis of control.

partner. The process of deconstructing a matrix into its interdependence space dimensions provides additional information about the social situation that can then be used to guide interactive behavior selection by a robot. Interdependence theory has developed the computational mechanisms for handling steps (2) and (3) from figure 2. One contribution of this paper is to complete the algorithmic process by developing the mechanisms for steps (1) and (4) and to apply this framework to a representative robotics problem.

B. A Computational Process for Situation-based Interaction

The four steps employed to translate the perception of a social situation into action have already been mentioned. The first step requires the construction of an outcome matrix representing the situation. Matrix construction describes the process by which the robot transforms perceptual and/or internal state information such as motivations and

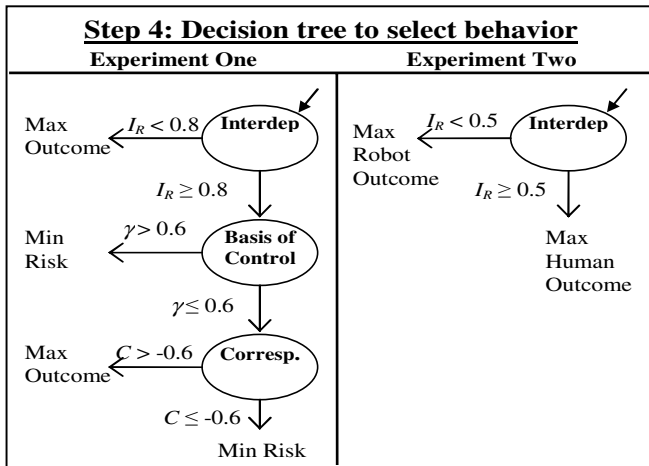


Fig. 6 Simple decision trees select either the outcome matrix action that maximizes the robot's own outcome or minimized risk. The values in the tree for experiment one were optimized for outcome generation. Max outcome selects the action with the greatest potential outcome for oneself. Min risk selects the action the greatest minimum outcome. If both actions have the same minimum outcome then the one with the greatest average is selected.

predilections into an outcome matrix. Interdependence theory terms this process the transformation process [7-9]. The transformation process results in the construction of an outcome matrix on which the robot can act. In some cases, raw perceptual stimuli can be directly used to construct the outcome matrix. As depicted in figure 3, the number of hazards and victims perceived is used to construct the outcome matrix for this work. These matrices expand upon the human-robot cleanup situation described previously. In these examples, both the human and the robot selected 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 3 were selected to give the autonomous robot a preference for cleanups and the human 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. Alternatively, the values within an outcome matrix can be generated from actual data (as they have been in psychological studies involving humans) or can be theoretically derived. 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.

The next step in the process depicted in figure 2 is situation analysis. Situation analysis involves the deconstruction of the raw outcome matrix into values representing the variances in outcome. Situation analysis is accomplished by using the procedure in figure 4. This procedure iteratively separates the values in the input or raw

outcome matrix into three separate matrices. 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 4, all outcome variance occurs in the BAC matrix when deconstructing an independent situation. The procedure for computing the variance of the social situation results in values for variables BC , PC , JC representing the variance of both the robot's and the human's outcomes in the situation. The subscripts denote the variance of the outcome for the robot and the human respectively.

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 (1-3) from figure 5. Equations (1) and (2) are from reference [9]. Equation (3) is a contribution of this work.

Finally, the interdependence space dimension values are used to guide action selection (fig. 6). As mentioned above, social situations are represented as an outcome matrix with respect to a pair of potential actions available to each individual. In the running example, the robot would choose either to rescue a victim or to cleanup a hazard. As depicted in figure 6, the decision trees use the interdependence space values to determine how the robot selects an action. The specifics of these decision trees will be discussed in the following section. These trees represent simple heuristics for selecting actions based on a situation's location in interdependence space. Decision trees were used because of their ease of implementation. Alternatives such as neural networks are also possible.

III. EXPERIMENTS

We conducted simulation experiments to test the proposed framework. These experiments focus on the possibility of capturing and using information about social situations to select behaviors. The experiments explore the framework in two ways. The first experiment examines the generality of a situation-based approach by testing the system's response to thousands of randomly generated abstract outcome matrices representing a broad spectrum of social situations. Because of time-constraints, it was not possible to test each of these randomly generated matrices using interaction between a human and a robot. The second experiment, therefore, focuses on a limited number of different social situations, but examines the response of the framework to these situations in detail using human-robot interaction and grounds the experiment in a typical robotics problem.

A. Experiment One

As mentioned above, the purpose of this experiment is to quantify the net outcome gains resulting from information generated about the situation over large portions of the interdependence space and to compare the results of our algorithm to other potential methods of interactive behavior selection. Because of time constraints, this experiment did not involve a human or a perceptually generated outcome matrix. Rather, in this case two simulated robots are directly presented a randomly generated outcome matrix within the interdependence space. Thus, in order to explore the entire space, this experiment bypasses the first step in the algorithm from figure 2, the conversion of perceptual information into a raw outcome matrix. Each randomly generated outcome matrix represents an abstract situation in the sense that the rewards and costs are associated with selecting one of two non-specified actions. Thus, these abstract situations are entirely general, but could easily be concretized by relating each matrix to specific actions and the outcomes to perceptual stimuli.

The experimental procedure first required the creation of a random outcome matrix. This matrix is then presented to the two simulated robots. One robot, the test robot, employed either our algorithm for situation-based interaction or a control strategy to select one of the abstract actions. The second robot consistently selected the action that maximizes its own outcome without consideration of its partner. The decision tree for experiment one is depicted in figure 6. Action selection by both robots occurs simultaneously. The outcome received by the test robot during each interaction is added to a running sum for the entire trial. A single trial consisted of 1000 randomly generated outcome matrices. One hundred trials were conducted for this experiment.

We explored three different control strategies for selecting actions as part of this experiment: (1) always select the action that maximized one’s own outcome without consideration of the partner, (2) select the action that maximizes both one’s own outcome and the outcome of the partner, and (3) select the action that minimizes the risk of losing outcome. The independent variable in this experiment is the type of strategy used by the test robot for selecting actions. The dependent variable then is the resulting net outcome for the test robot. We hypothesized that a situation-based approach would result in the greatest net outcome.

B. Experiment Two

In the second experiment, we used the *MissionLab* behavior specification system to explore the computational process described in figure 2. *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 [18].

In this experiment, an autonomous robot and a human teleoperated robot are attempting to cleanup toxic spill hazards and rescue injured victims. In order to complete this task both robots must forage for attractors such as victims or hazards to cleanup. Foraging, a well-studied problem in robotics, requires that a robot search for and retrieve attractor objects [19].

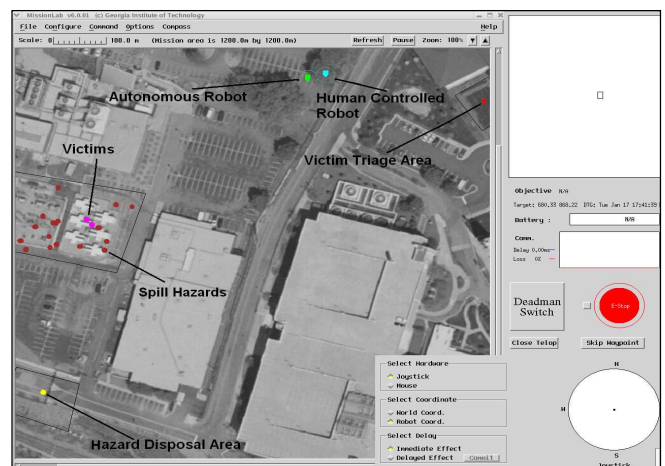


Fig. 7 This figure depicts the *MissionLab* toolset with the simulation environment used for experiment two. The experiment explored a foraging task in an urban environment. The teleoperation interface used by the human is depicted the right.

Figure 7 depicts the layout for this experiment. Potential victims and hazards for cleanup are located within a disaster area. A disposal area for hazard items is located towards the bottom and a triage area for victims is located to the right. For this experiment, outcome matrices are constructed from perceptual information about the situation, namely the number of victims and hazards perceived (fig. 3). 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.

This experiment compared our algorithm for capturing and using information about social situations (fig. 2) to a strategy that does not use situation information and focuses solely on maximizing the robot’s immediate outcome. Independent and dependent situations were investigated. Independent situations allow the human teleoperated robot and the autonomous robot to forage independently (fig. 3 left). Dependent situations (fig. 3 right), on the other hand, require that both the autonomous robot and the human collaborate to rescue a victim or cleanup a hazard. In all conditions, the teleoperated robot selected the interactive behavior that maximizes its own outcome without consideration of its partner. In the experimental condition, our algorithm is used in conjunction with the decision tree for experiment two depicted in figure 6. This decision tree represents a simple heuristic that guides the autonomous robot to select the interactive behavior that maximizes its own outcome if the situation is determined to be more independent than dependent. If, on the other hand, the situation is determined to be more dependent than independent, the autonomous robot selects the behavior that maximizes its partner’s outcome. Alternatively, in a control condition the autonomous robot always selects the behavior that maximizes its own outcome without consideration of the partner. Thus, in this control condition both robots employ a selfish outcome maximizing strategy.




Fig. 8 Results for experiment one are presented. The leftmost bar (grey) depicts the maximum possible net outcome. The second bar from the left (blue) indicates the net outcome when our procedure is used. The next three bars (red) depict the controls for the experiment. The rightmost bar (green) depicts the partner's outcome for the experimental condition. Error bars indicate Standard Error of the Mean.

We conducted thirty trials in each of the four conditions for this experiment: (1) independent situation and control robot, (2) independent situation and experimental robot, (3) dependent situation and control robot, and (4) dependent situation and experimental robot. A random number of victims and hazards were created for each trial and were randomly placed within the disaster area.

IV. RESULTS

Figure 8 presents results for the first experiment. The second bar from the left (blue) depicts the net outcome using our algorithm. The next three bars to the right (red) indicate the net outcome when for the control conditions. Our algorithm significantly outperforms the controls in all three conditions ($p < 0.04$ for all). The maximum possible outcome is also depicted for reference (gray).

Although significant, the difference in net outcome between our situation-based algorithm and the outcome maximizing control strategy is not large. This difference reflects the simplicity of the decision trees used to select actions. The decision tree uses the information about the situation provided by the previous steps of our algorithm to select the best action. Additional effort could have been devoted to the construction of a better method for selecting interactive behaviors based on information about the situation, including mechanisms for modelling the partner and predicting its action selections. However, the purpose of this experiment was to examine the effect of having information about the situation over a spectrum of social situations. Thus, any statistically significant difference suffices and work devoted to optimizing action selection based on information about the situation is left for future work.

The results from this experiment demonstrate that our situation-based approach to interaction improves performance (as measured by outcome) over a broad expanse of outcome




Fig. 9 Results for experiment two are presented. The left two bars depict the independent conditions for both the test and the control robot. In the independent situation, the autonomous robot forages for hazards and the teleoperated robot forages for victims. The right two bars portray the dependent situation. Hazards cleaned (green) is depicted above of victims rescued.

matrices. Our procedure of computationally analyzing and mapping the outcome matrix is thus not limited to certain specific social situations and these results serve as evidence of the potential broad applicability of this work. Moreover, this experiment indicates that the information provided by an analysis and mapping of the outcome matrix is useful, on average, for a robot selecting interactive actions.

Figure 9 depicts the results from the foraging experiment. The left two bars portray the results for the situation in which the autonomous robot and the teleoperated robot were independent. In this case, the autonomous robot forages for hazards to cleanup and the human-operated robot forages for victims. Thus, all of the 30 trials each robot retrieves either a victim or a hazard. Because each trial consisted of a random number of victims and hazards, some trials resulted in zero victims or hazards. In this case, both robots foraged for the same object. Both robots faired equally well in the independent situation. Thus in a situation in which each individual's performance is independent, neither needs to consider the effect of their action on the other while completing the task.

In the dependent condition, however, the autonomous robot's use of situation information affords better performance than the robot that does not consider the situation. In the test condition, 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 decision tree therefore directs the autonomous robot to perform the action that will maximize its partner's outcome, i.e. forage for victims, if any exist. The control strategy simply maximizes the autonomous robot's own outcome without consideration of the partner even though the situation demands collaboration, hence resulting in poor performance.

Overall, the second experiment demonstrates that the information resulting from an analysis of the social situation

can improve a robot's ability to perform interactive tasks such as collaborative foraging. Moreover, the algorithm we have proposed can successfully use perceptual stimuli in the environment to produce information about the social situation. These results were demonstrated on a problem traditionally of interest to robotics researchers [19]. Still, these experiments merely skim the surface of potential avenues for research and future work will be devoted to testing this framework on real robots in real social situations.

V. CONCLUSIONS

This paper has introduced mechanisms for capturing information about social situations and for using this information to guide interactive behavior. We have presented a situation-based framework for human-robot interaction derived from the social psychological theory of interdependence. This approach includes representations for social situations as outcome matrices and provides the computational tools for mapping situations to their location in an interdependence space. 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.

For an autonomous robot interacting in a complex dynamic environment, such as search and rescue, it is important to provide the robot with methods for understanding and extracting information from its social environment. These methods are in turn expected to result in better performance on a variety of social tasks, such as collaboration. Moreover, the experiments presented in this paper do not use all of the information provided by our procedure. Information pertaining to correspondence, basis of control, and symmetry could clearly provide a great deal of additional knowledge, perhaps leading to better selection of interactive behaviors by a robot, and will be one direction for future work.

The work presented in this paper represents an ongoing research effort. Future work will 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. We also intend to explore the process for transforming a perceived social situation in greater detail.

It is our contention that this framework 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 will be expected to move out of the laboratory and into one's home. Moreover, because this framework 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.

ACKNOWLEDGMENT

The authors would like to thank Zsolt Kira, Yoichiro Endo, and Patrick Ulam for their many helpful comments. This paper is dedicated to the memory of Mark D. Grimaldi whose interaction always resulted in laughter.

REFERENCES

- [1] R. W. Byrne and A. Whiten, "Machiavellian intelligence," in *Machiavellian Intelligence II: Extensions and Evaluations*, A. Whiten and R. W. Byrne, eds., Cambridge University Press, Cambridge, 1997, pp. 1-23.
- [2] N. K. Humphrey, "The social function of intellect," in *Growing Points in Ethology*, P. P. G. Bateson and R. A. Hinde, eds., 1976, pp. 303-317.
- [3] R. Bar-On, D. Tranel, N. L. Denburg, and A. Bechara, "Exploring the neurological substrate of emotional and social intelligence," *Brain*, vol. 126, 2003, pp. 1790-1800.
- [4] T. J. Bergman, J. C. Beehner, D. L. Cheney, and R. M. Seyfarth, "Hierarchical Classification by Rank and Kinship in Baboons," *Science*, vol. 302, 2003, pp. 1234-1236.
- [5] R. A. Hinde, "Can Nonhuman Primates Help Us Understand Human Behavior?" in *Primate Societies*, B. B. Smuts, D. L. Cheney, R. M. Seyfarth, R. W. Wrangham, and T. T. Struhsaker, eds., University of Chicago Press, Chicago, 1987, pp. 413-420.
- [6] D. O. Sears, L. A. Peplau, and S. E. Taylor, *Social Psychology*. Prentice Hall, Englewood Cliffs, New Jersey, 1991.
- [7] C. E. Rusbult and P. A. M. V. Lange, "Interdependence, Interaction and Relationship," *Annual Review of Psychology*, vol. 54, 2003, pp. 351-375.
- [8] H. H. Kelley, J. G. Holmes, N. L. Kerr, H. T. Reis, C. E. Rusbult, and P. A. M. V. Lange, *An Atlas of Interpersonal Situations*. Cambridge University Press, New York, NY, 2003.
- [9] H. H. Kelley and J. W. Thibaut, *Interpersonal Relations: A Theory of Interdependence*. John Wiley & Sons, New York, NY, 1978.
- [10] E. Rogers and R. Murphy, "Human-Robot Interaction Final Report for DARPA/NSF Study on Human-Robot Interaction," 2001, pp.
- [11] T. Fong, I. Nourbakhsh, and K. Dautenhahn, "A survey of socially interactive robots," *Robotics and Autonomous Systems*, vol. 42, 2003, pp. 143-166.
- [12] C. L. Breazeal, *Designing Sociable Robots*. The MIT Press, Cambridge, MA, 2002.
- [13] B. Scassellati, "How Developmental Psychology and Robotics Complement Each Other," in *NSF/DARPA Workshop on Development and Learning*. Michigan State University, Lansing, MI, 2000., 2000.
- [14] J. Pineau, M. Montemerlo, M. Pollack, N. Roy, and S. Thurn, "Towards robotic assistants in nursing homes: Challenges and results," *Robotics and Autonomous Systems*, vol. 42, 2003, pp. 271-281.
- [15] W. S. Bainbridge, E. E. Brent, K. M. Carley, D. R. Heise, M. W. Macy, B. Markovskiy, and J. Skvoretz, "Artificial Social Intelligence," *Annual Review of Sociology*, vol. 20, 1994, pp. 407-436.
- [16] A. G. Sanfey, J. K. Rilling, J. A. Aronson, L. E. Nystrom, and J. D. Cohen, "The Neural Basis of Economic Decision-Making in the Ultimatum Game," *Science*, vol. 300, 2003, pp. 1755-1758.
- [17] J. Berg, J. Dickhaut, and K. McCabe, "Trust, Reciprocity, and Social History," *Games and Economic Behavior*, vol. 10, 1995, pp. 122-142.
- [18] R. C. Arkin and T. Balch, "AuRA: Principles and Practice in Review," *Journal of Experimental and Theoretical Artificial Intelligence*, vol. 9, 1997, pp. 175-189.
- [19] R. C. Arkin, *Behavior-Based Robotics*, 2 ed. The MIT Press, Cambridge, MA, 1999.