



Quantifying Human Decision-Making: Implications for Bidirectional Communication in Human-Robot Teams

Kristin E. Schaefer¹(✉), Brandon S. Perelman¹, Ralph W. Brewer¹,
Julia L. Wright¹, Nicholas Roy², and Derya Aksaray³

¹ US Army Research Laboratory, Aberdeen Proving Ground,
Adelphi, MD 21005, USA

kristin.e.schaefer-lay.civ@mail.mil

² MIT, Cambridge, MA 02139, USA

³ University of Minnesota, Minneapolis, MN 55455, USA

Abstract. A goal for future robotic technologies is to advance autonomy capabilities for independent and collaborative decision-making with human team members during complex operations. However, if human behavior does not match the robots' models or expectations, there can be a degradation in trust that can impede team performance and may only be mitigated through explicit communication. Therefore, the effectiveness of the team is contingent on the accuracy of the models of human behavior that can be informed by transparent bidirectional communication which are needed to develop common ground and a shared understanding. For this work, we are specifically characterizing human decision-making, especially in terms of the variability of decision-making, with the eventual goal of incorporating this model within a bidirectional communication system. Thirty participants completed an online game where they controlled a human avatar through a 14×14 grid room in order to move boxes to their target locations. Each level of the game increased in environmental complexity through the number of boxes. Two trials were completed to compare path planning for the condition of known versus unknown information. Path analysis techniques were used to quantify human decision-making as well as provide implications for bidirectional communication.

Keywords: Human-robot teaming · Bidirectional communication
Decision-making

1 Introduction

Human-agent teaming is a critical area of research because technological advancements are reaching the point where machines are able to make both independent and interdependent decisions. Due to these advancements, human team member roles are transitioning to more communication-based interactions supporting larger goals and intentions rather than direct control or teleoperation of the system [1]. One key limitation in the development of effective teaming has been the process of building a shared understanding of the mission space, whereby robotic team members can quickly and

accurately understand the human's intent and behaviors [2]. This is important because successful collaborative partnerships require communication, cooperation, and coordination between the acting members as they work towards a common goal [3], and distinct from prior work in supervisory control, which has often focused on ensuring the human has an accurate model of the robot's behavior. In a truly collaborative context, it is crucial that the robot also have situation awareness of the human partner.

1.1 Bidirectional Communication

Bidirectional communication is an area of current research that aims to improve development of common ground and shared understanding. It is especially critical for the transformation of a robot from tool to team member [4] because it allows for joint decision-making and development of shared mental model representations [5], and knowledge transfer through communication supports shared situation awareness [6, 7]. Increased autonomy, especially the capability for independent and interdependent decision-making in complex environments, supports this need for the development of bidirectional communication. In order to effectively communicate, it is important for both human and robotic team members to understand the decisions and decision-making processes within their team. The interpretation of an interaction or actions of a robotic team member can be directly influenced by the person's expectations for the interaction. Similarly, if the human teammate's actions and behaviors deviate from the robot's expectations, there will be a degradation in trust. Thus, if a robot team member can interpret and correctly predict the actions and behaviors of the human, then the robot can react accordingly. Bidirectional communication can also help both the team members understand when a decision may intuitively counter their own ideas or models by providing reasoning information that defines the appropriateness of its decisions, thus updating the team member's mental model and expectations for the task. In addition, the mode of communication and feedback capabilities have an effect on trust development in human-agent teams [8].

1.2 Human Decision-Making

While there are many types of decisions, this paper focuses on spatial decision-making. Human spatial decision-making is characterized by the ability to rapidly produce robust solutions to complex problems. For example, the Traveling Salesman Problem (TSP) requires participants to connect nodes, representing cities, to create the shortest tour among the nodes. While its instructions are simple, the TSP is NP-hard, and brute force solutions require calculating $(n - 1)!/2$ tours where n is the number of nodes. Despite the computational complexity, humans produce near-optimal solutions to this problem in linear time [9–12] using a combination of global and local spatial heuristics [13–15]. Due to the quality of the solutions and the speed at which they are produced, the decision-making mechanisms humans use to solve these problems are an area of study for the AI community, but the underlying mechanisms still remain unknown.

Naturalistic spatial decision-making tasks allow these mechanisms to be supported by guidance from top-down cognitive processing systems [16–19]. Humans are capable of adapting paths easily to mission requirements during naturalistic, real-world tasks,

such as mission planning for unmanned aerial vehicles [20]. Yet, recent research demonstrated that aggregate human solutions tend to converge on not just one but several solution groups, each characterized by a distinct spatial mental model [21]. The adaptive nature of these top-down processes permits mission-dependent flexibility in the spatial decision-making process, and this characteristic has implications for bidirectional communication in human-agent teams.

Shared Mental Models. Understanding decision-making helps to classify the mental model for a task, which then guides expectations for interaction. Spatial mental models are mental representations of the environment [22–24], and weightings of the importance of features in those representations relative to goals [22]. Spatial mental models directly impact a solution to a given spatial decision-making problem, as well as their evaluations of solutions generated by other humans and algorithms. This has direct implications for all manner of human-agent teaming problems. For example, in collaborative spatial decision-making, an algorithm may propose a route to a human who can either accept or choose to replan it. This replanning or retasking degrades performance and situation awareness, and increases workload [25]. In addition, divergence between the human team members’ spatial mental models and the actions taken by an intelligent agent can reduce predictability and degrade trust [26–28]. Conversely, spatial mental models that are similar to an agent’s suggestion can improve agent trust, and increase the rate of acceptance for that solution [21]. This is especially true for cases where agents are unable to articulate their reasoning for producing solutions that may contradict human teammates’ spatial mental models. Therefore, this area represents a potential target for future research in bidirectional communication for the purpose of achieving consensus between human spatial mental models and intelligent agent problem solving mechanisms.

Implications for Agent Development and Teaming. Knowing how humans make decisions could help a robot to derive a model of the team member’s planning model, which allows the robot to infer future human behavior, and provides the needed context to communicate its state and goals in the same representation as the human. Such an extrapolation could greatly reduce the need for explicit communication (e.g., it might suffice for a robot to observe nodding or a hand cue to infer how human will act next). Moreover, matching both representation and goals could result in an efficient search over the action space (e.g., a robot could narrow its search space based on expected human behavior). For intractable problems, knowing the optimal solution may not be possible. In that case, knowing how humans solve a problem could be a benchmark when developing robot algorithms. Further, understanding the limitations of a human can help teaming in such a way that the robot can take the initiative of being the main actor (e.g., computing a plan) in cases where the human is limited.

1.3 Current Work

The main objective of this research is to understand similarities and differences among human spatial decision-making processes as they apply to future human-agent teaming. When developing new spatial planning algorithms for robotic systems that will be collaborating with people to complete a task (e.g., moving objects around a room), it is

important to characterize and compare the behavior of each of the agents under different conditions. Further, since each agent applies its own spatial mental model or algorithm to solve a given problem, in order to achieve robust collaboration and teamwork it is critical to recognize how the decision-making processes of each agent will handle increasing environmental complexity and uncertainty. Where disparities exist between the resultant robot and human behaviors, bidirectional communication can be used as a means to achieve an optimal solution collaboratively. A first step toward achieving this goal is to characterize human spatial decision-making behaviors in the proposed tasks. For this study, an online game was developed to assess human spatial decision-making processes involved with controlling a virtual avatar through a virtual room with the purpose of pushing virtual boxes from a set of start locations to a set of end locations. The design of the study was such that each level represented an increase in environmental complexity, and the two conditions represented an increase in task difficulty based on the availability of planning information.

2 Methodology

2.1 Participants

Thirty participants between the ages of 18 and 60 were recruited. This age restriction was selected to reduce variance in participants' spatial abilities. In prior studies involving tasks requiring spatial working memory, age-related cognitive decline reduces navigation speed [29] and overall task performance [30–32].

2.2 Game Development and Task

A Java Applet was built around similar game dynamics as the puzzle game Sokoban [33]. Sokoban, developed by Thinking Rabbit game studio in 1982, is a logic puzzle designed for the user to push objects (stones, boxes, etc.) around a playing field to a goal area in the fewest moves possible. For our study, the main game space for all levels was a 14×14 square grid surrounded by a brick wall on all sides. The grid space was developed to match the laboratory facilities at MIT to allow for future comparison of human and robot decision-making. The difference in design between this application and the original Sokoban game was that typically levels only had minimal number of solutions, while the open area of this playing field made for exponentially more paths to reach a solution.

There were nine levels (Level 0 through Level 8) that increased in the number of boxes from two to 10 boxes (Fig. 1). The number and location of the boxes and target locations were devised in such a way to represent increasing environmental complexity. In order to investigate the variability in human decision-making, different patterns, clusters, and spatially distant blocks were used while choosing the initial and target locations of the blocks. The game levels were designed in such a way that the optimal (or very close to optimal) solution was not obvious. This helps to identify variability in human decision-making behaviors.

The start location of the avatar (represented by a person) was always in the same start corner. The avatar could only move up, down, left, or right (no diagonals) and could only push (not pull) the boxes. Therefore, the initial placement of the boxes were located a certain distance from the boundaries to ensure the existence of a solution. In order to avoid infeasibility (e.g., deadlock), the exterior 2 cells were intentionally left




Fig. 1. Each game level (a)–(h) represents increasing environmental complexity. For Condition 2 (Unknown Planning Information), the entire game board was transposed over the y-axis making each condition directly comparable but unique.

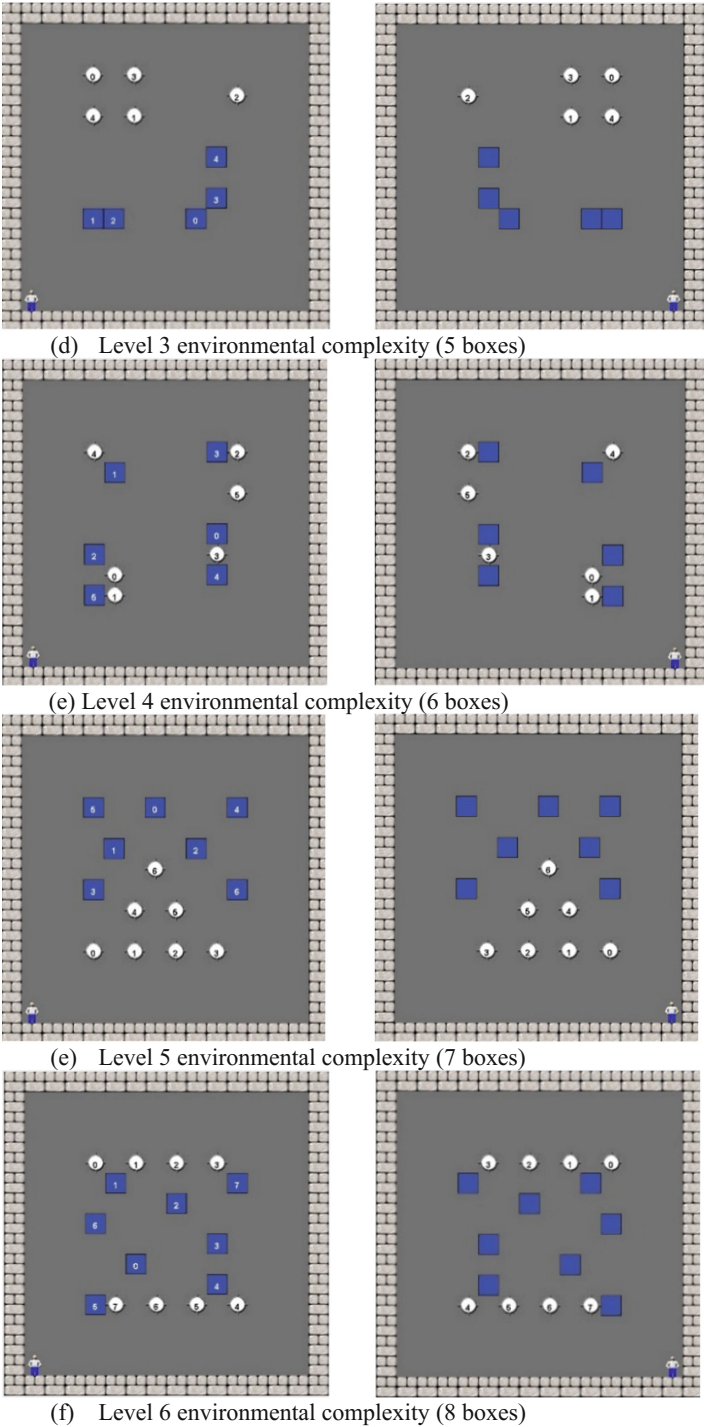


Fig. 1. (continued)

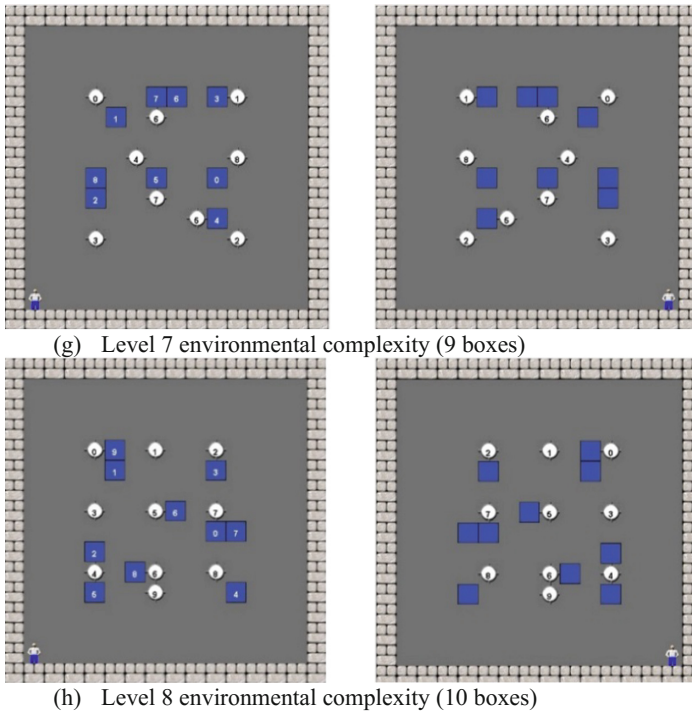


Fig. 1. (continued)

blank and boxes were not placed in corridor like shapes. An undo option to backup through previous moves, as well as a reset level option were available so that it was always possible to reach a solution.

The overall goal was to move all boxes from their initial locations to their target locations. To this end, there were two main criteria to determine the overall trajectory, the sequence the boxes should be moved and calculating the shortest path from a box's initial location to its target location. Participants completed two conditions representing increasing task difficulty. For Condition 1 (Known Planning Information), all boxes and target locations were known, such that participants had to control the human avatar through the virtual environment and push Box 1 to Target Location 1. Participants were instructed that the numbering on the boxes and target locations were only there to inform which box was connected with which target location. They could complete the task in any order. For Condition 2 (Unknown Planning Information), the boxes were unlabeled however all the target locations were labeled. The box numbers only became visible once the human avatar pushed the box to a new grid square location.

2.3 Design

The experimental design was a 2 condition (known versus unknown planning information) \times 9 environmental complexity (game levels ranging from 2–10 boxes)

within-subjects design. Participants completed Condition 1, Levels 0 through 8, followed by Condition 2, Levels 0 through 8. There were two main hypotheses.

Hypothesis 1: Overall path length, number of movements, and time to completion will be longer when the environment is more complex (i.e., more boxes) and when the task is more difficult (i.e., changes in the amount of previously known information about the task).

Hypothesis 2: Based on our prior research, there will be more than one “human” way of making decisions to plan a path and move the boxes. The Algorithm for finding the Least Cost Areal Mapping between Paths (ALCAMP) [34] will be used to quantify the divergence among participants’ solutions in Condition 1 and Condition 2 for each level. Higher levels of environmental complexity associated with the levels will produce greater divergence (i.e., variability) in solutions, and we expect that the availability of planning information (manipulated in each Condition) will also impact the divergence among solutions in each level.

2.4 Analysis

Performance. Specific decision-making time and movements were recorded for all levels across both trials. Decision-making times included planning time (i.e., time to first movement), total completion time (i.e., time till last box was placed on the correct target), and action time (i.e., total completion time minus planning time). The actual decisions were analyzed by looking at the total number of moves to complete each level, and nearest-neighbor analysis. The nearest-neighbor analysis calculated the number of participants who first moved to the closest box to the start location compared to another box. This analysis provides insight into whether or not they used a local versus global strategy.

Variability. A novel approach was used to characterize the variability among participants’ solutions for each of the spatial problems in each level. The variability among solutions is important to understand differences between individuals’ spatial solutions and to determine predictability in decision-making behaviors. Thus, increasing solution variability corresponds to decreasing predictability for both human- and algorithm-produced solutions. To measure the variability, participants’ solutions were pooled within each level representing environmental complexity and condition of task difficulty. ALCAMP [34] was used to compare all solutions within each pool in a pairwise-exhaustive fashion. The resultant values of this analysis reflect the divergence between the pair of paths as measured by Euclidian divergence among the grid squares, such that large values indicate that the two solutions are different and small values indicate that they are similar. Finally, these values were used to populate a symmetric dissimilarity matrix through a distance matrix. The mean of the upper or lower triangle indicates the average dissimilarity among all of the solutions for a given level and condition – the higher the value, the greater the variability among participants’ solutions to that problem. This technique has been used in previous research to infer consensus in spatial decision-making processes [11, 21].

3 Results

3.1 Performance Analysis

Total Completion Time. A 2 condition (known vs unknown boxes) \times 9 levels of environmental complexity (levels 0–8) repeated measures within-subjects ANOVA was conducted to assess total completion time. There was a main effect of condition, $F(1,24) = 15.98, p = .001, d = 1.03$, where Condition 1 (known planning information) was longer, $M = 49.894$ s, $SE = 2.553$, than Condition 2 (unknown planning information), $M = 44.049$ s, $SE = 1.898$. There was a main effect of level, $F(8, 17) = 78.89, p < .001, d = 1.35$, whereby increased environmental complexity led to increased completion time, and an interaction, $F(8, 17) = 8.821, p < .001, d = 0.41$, see Fig. 2. These results show that increasing the amount of information available to an agent can increase processing times despite providing important cues to objects in the environment.

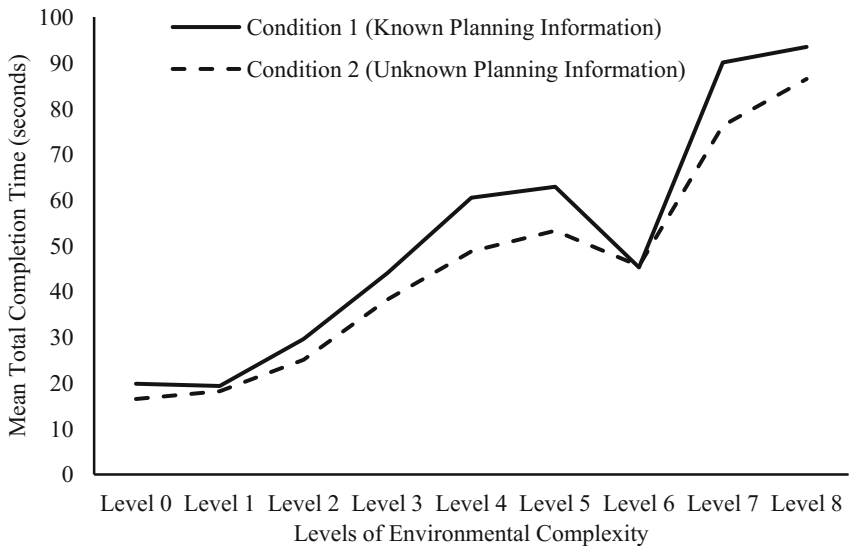


Fig. 2. Total completion time (seconds) on each game level representing environmental complexity (game level) for the known and unknown planning information conditions

Paired samples t -tests were conducted for each level of environmental complexity to compare total completion time in the known and unknown planning information conditions. There was a significant difference in scores for Levels 0 ($p = .020$), 2 ($p = .008$), 3 ($p = .009$), 4 ($p < .001$), 5 ($p = .004$), and 7 ($p < .001$), whereby completion time was significantly longer for Condition 1 (known planning information) than Condition 2 (unknown planning information). Results are reported in Table 1. These results show that the effect between the two conditions collapsed completely on Level 6, and partially on Level 8, likely owing to characteristics of those specific environments.

Table 1. Paired samples *t*-tests for total completion time

Game level	Mean time 1	SD 1	Mean time 2	SD 2	<i>t</i>	<i>p</i>	<i>d</i>
0	19.851	6.827	16.513	5.294	<i>t</i> (26) = 2.479	.020	0.48
1	19.342	6.401	18.211	4.015	<i>t</i> (27) = 1.272	.214	0.24
2	29.614	10.303	25.077	5.351	<i>t</i> (27) = 2.851	.008	0.54
3	44.052	11.258	38.273	7.160	<i>t</i> (27) = 2.830	.009	0.53
4	60.510	16.714	48.775	1.4152	<i>t</i> (27) = 4.373	<.001	0.83
5	62.897	22.463	53.272	13.360	<i>t</i> (26) = 3.117	.004	0.60
6	45.275	11.520	45.611	12.612	<i>t</i> (26) = -.139	.891	-0.03
7	90.077	26.434	76.256	20.313	<i>t</i> (27) = 5.855	<.001	1.11
8	93.486	29.564	86.447	27.756	<i>t</i> (28) = 1.893	.069	0.35

Note. Time is in seconds

Planning and Action Time. In order to determine whether the source of the total completion time effects described above were due to differences in planning or action, we split the total completion time into a planning phase (duration between trial presentation and the first move) and action time (the remainder of the total completion time – planning time). A 2 condition × 9 levels of environmental complexity repeated measures within-subjects ANOVA was conducted to assess planning time. There was a main effect of condition, $F(1,15) = 24.33, p < .001, d = 2.006$, where Condition 1 (known planning information) was longer, $M = 3.123$ s, $SE = 0.501$, than Condition 2 (unknown planning information), $M = 1.375$ s, $SE = 0.129$. There was a marginal main effect of level, $F(8, 120) = 1.967, p = .056, d = 0.028$, whereby increased environmental complexity led to increased planning time, and an interaction, $F(8, 120) = 2.776, p = .007, d = .038$. Paired samples *t*-tests were conducted for each level of environmental complexity to compare planning time in the known and unknown planning information conditions. There was a significant difference in scores for Levels 1 ($p = .005$), and Levels 2 through 8 ($p < .001$), whereby planning time was significantly longer for Condition 1 (known planning information) than Condition 2 (unknown planning information). Results are reported in Table 2.

Table 2. Paired samples *t*-tests for planning time

Game level	Mean time 1	SD 1	Mean time 2	SD 2	<i>t</i>	<i>p</i>	<i>d</i>
0	3.230	1.281	2.765	1.341	<i>t</i> (24) = 1.591	.125	0.32
1	3.608	3.213	1.891	0.948	<i>t</i> (25) = 3.077	.005	0.60
2	2.686	1.326	1.369	0.732	<i>t</i> (24) = 5.548	<.001	1.11
3	4.326	3.542	1.268	0.574	<i>t</i> (25) = 4.627	<.001	0.91
4	6.629	6.468	1.436	0.694	<i>t</i> (28) = 4.452	<.001	0.83
5	5.099	4.601	1.496	0.694	<i>t</i> (27) = 4.405	<.001	0.83
6	3.155	2.036	1.453	0.790	<i>t</i> (25) = 4.563	<.001	0.89
7	5.124	3.815	1.211	0.509	<i>t</i> (24) = 5.227	<.001	1.05
8	6.134	5.246	1.713	0.907	<i>t</i> (26) = 4.542	<.001	0.87

Note. Time is in seconds




Fig. 3. Mean task times across levels of environmental complexity for planning time (a) and action time (b) where the solid black line represents Condition 1 (Known Planning Information) and the dashed line represents Condition 2 (Unknown Planning Information)

These results show that, while planning time generally increased with environmental complexity when boxes were known (though the layout of the environment clearly played a role as well, as shown by the dip in planning time for Level 6), planning time essentially dropped to floor when boxes were unknown. One interpretation of this result is that participants adopted a very simple local decision-making

strategy when the box numbers were unknown, as opposed to the global search performed when all planning information was presented. In addition, the paired samples t -tests showed that the effect of uncertainty on planning requires a minimum amount of environmental complexity (i.e., number of boxes) to manifest.

To assess action time, a 2 condition \times 9 levels of environmental complexity repeated measures within-subjects ANOVA was conducted to assess action time. There was a main effect of condition, $F(1,29) = 5.683$, $p = .024$, $d = .359$, where Condition 1 (known planning information) was longer, $M = 51.040$ s, $SE = 4.202$, than Condition 2 (unknown planning information), $M = 48.464$ s, $SE = 4.172$. There was a main effect of level, $F(8, 232) = 99.828$, $p < .001$, $d = .365$, whereby increased environmental complexity led to increased completion time. There was not a significant interaction, $p = .088$. Figure 3 depicts the mean planning times and mean action times for both conditions across all levels of environmental complexity.

Paired samples t -tests were conducted for each level of environmental complexity to compare action time in the known and unknown planning information conditions. There was a significant difference in scores for Level 4, $t(29) = 2.217$, $p = .035$, $d = 0.40$; Level 5, $t(29) = 2.045$, $p = .050$, $d = 0.37$; Level 6, $t(29) = -2.116$, $p = .043$, $d = -0.39$, and Level 7, $t(29) = 2.291$, $p = .029$, $d = 0.04$. These results show that the interaction effect between information availability and environmental complexity nearly disappears when removing the variance attributable to the planning phase of problem solving. Thus, these results taken together show that the differences reflect differences in planning for spatial problem solving rather than the action of actually moving the avatar to solve the problem.

Number of Moves. A 2 condition \times 9 levels of environmental complexity repeated measures within subjects ANOVA was conducted to assess number of moves needed to complete the task. There was a marginal main effect of condition, $F(1,21) = 3.85$, $p = .063$, $d = .337$, where Condition 2 (unknown planning information) required more moves, $M = 117.64$, $SE = 1.595$, than Condition 1 (known planning information), $M = 115$, $SE = 1.590$. There was a main effect of level, $F(1, 21) = 2707.23$, $p < .001$, $d = 2.48$, whereby increased environmental complexity led to increased completion time. There was significant interaction, $F(1, 21) = 4.33$, $p = .050$, $d = .042$. Paired samples t -tests were conducted for each level of environmental complexity to compare total number of moves in the known and unknown planning information conditions. There was a significant difference in scores for Level 1, $t(26) = -3.389$, $p = .002$, $d = -.65$; and a marginal significant difference in scores for Level 8, $t(28) = -2.028$, $p = .052$, $d = -.38$. The result of this analysis indicate that the conditions did not have a meaningful impact on the number of moves required to complete each problem. Rather, the number of moves was mostly impacted by the environmental complexity.

Nearest Neighbor Analysis. In order to test the hypothesis that participants relied more heavily on local decision-making heuristics when boxes were unknown, we calculated the number of participants who first moved the box closest to the starting position (*nearest neighbor preference*), at each level of environmental complexity for each trial condition. A high percentage of participants exhibiting nearest neighbor preference would indicate a very simple local decision-making strategy. A lower percentage would indicate that participants employed a more global strategy. Figure 4

shows the percentage of participants who interacted with the closest box first for each level, in both experimental condition.




Fig. 4. Nearest neighbor analysis indicates the percentage of trials on which participants visited the box closes to the starting location first.

Participants' nearest neighbor preference was near-ceiling when the boxes were unknown. When the boxes were known in advance, the percentage of participants showing a nearest-neighbor preference decreased with increasing environmental complexity. These results, taken together, substantiate the hypothesis that participants typically employed global decision-making strategies when the box identities were known, but resorted to local decision-making strategies in the absence of that information.

3.2 Variability Analysis

Using the aforementioned procedure for calculating the average divergence in each condition and level, we see that variability in participants' solutions increases roughly linearly with increasing environmental complexity, once complexity reaches a certain threshold (in this case, it appears to be 5 boxes). A 2 condition (Known vs. Unknown planning information) \times 9 environmental complexity (Levels 0–8) ANOVA revealed significant main effects of both Conditions, $F(1, 7812) = 227.33, p < .001$, and Level, $F(8, 7812) = 2305.92, p < .001$, as well as an interaction effect between these two variables, $F(8, 7812) = 28.27, p < .001$ (see Fig. 5). Post hoc analysis using Tukey's HSD tests showed significant effects between Conditions for levels 0 ($p = .022$), 2 ($p < .001$), 4 ($p < .001$), 5 ($p < .001$), 7 ($p < .001$) and 8 ($p < .001$). In all of these




Fig. 5. The divergence value indicates the number of grid points by which solutions differed or variance in solutions. The average divergence among participants' solutions show very little differences in solutions for lower levels of environmental complexity (Levels 0–2) but a large increase in the number of possible solutions starting at Level 3.

cases, participants' solutions exhibited greater average divergence when the boxes were known (i.e., Condition 1) versus unknown (Condition 2). Note that the dip in divergence in both conditions on Level 5 (7 boxes) likely reflects characteristics of that particular environment.

4 Discussion

Communication is described by the research community as a reciprocal process where teammates send and receive information that form and reform the team's attitudes, behaviors, and cognition [35] whereby a shared body of knowledge can be used to develop shared expectations, allowing for improved team performance without explicit coordination [36, 37]. In the past, a variety of methods to facilitate bidirectional communication have been explored. For example, transparent user displays can convey agent intent [2], as well as its goals, reasoning, and projected outcomes [38–40]. While a multimodal approach to communication can reduce workload and degraded situation awareness [41] by using both implicit (nonverbal – behaviors, actions) and explicit (voice, natural language or auditory) communication modalities [42]. Considerable efforts have gone into determining the type, amount, modality, and rate by which information should be communicated between team members (e.g., Situation Awareness agent-based Transparency Model [40]). But perhaps key to the development of effective bidirectional communication within human-agent teams is the need for team members to be aware of goals, reasoning, actions, and projected outcomes of their teammates [43–46]. Therefore, a major step to developing appropriate bidirectional

communication is being able to quantify human behavior across tasks. If human behavior does not match the robots' models or expectations, there can be a degradation in trust that can impede team performance and may only be mitigated through explicit communication.

4.1 General Discussion

This was the first study in a set of studies using this paradigm. It was designed to advance the technical capabilities of a robot to more accurately perceive and interpret human team member behavior, and to develop appropriate bidirectional communication required for future collaborative tasking. By first looking at quantifying human behavior, we can provide a foundation for understanding how human expectations for planning and spatial task solutions are formed. This is essential for future teaming because when human expectations do not match robot behaviors then degradations in trust can occur. Therefore, quantifying the decision space can provide insights into identifying when and how bidirectional communication could mitigate divergences in human and robot team behaviors.

Human Performance. The results of the present study showed that completion times generally increased with increasing environmental complexity. Furthermore, participants generally took longer to complete the levels when the box contents were known. Separating participants' solution times into planning times and action times showed that the majority of this discrepancy between the two information availability conditions was due to differences in time spent in planning rather than action. When the box numbers were visible to participants, participants took longer to begin moving than when the box numbers were not known, and we believe this time was spent analyzing the environment and planning their moves. Curiously, this increase in planning time did not translate to increased efficiency, as participants' solutions did not vary between the two information availability conditions in terms of the number of moves. Generally, this result indicates that perfect world knowledge did not improve performance, and actually reduced the speed with which participants completed each level. Understanding variance in planning and completion times can provide insights into situations that may require more explicit communication between team members to clarify the underlying reasoning process for the decision being made, as well as help to determine timing associated with providing feedback to a team member.

Global versus Local Decision-Making and Implications for Bidirectionality. In order to quantify decision-making behaviors, as well as further investigate the planning time difference described above, we performed a simple analysis to determine whether participants were using a local decision-making heuristic - nearest neighbor. The nearest neighbor analysis showed that nearly all participants employed a nearest neighbor heuristic when the box numbers were unknown, visiting the closest box first. When box numbers were known, participants appeared to increasingly leverage global decision-making strategies. This result, taken in the context of the performance results, shows that perfect world knowledge, which facilitates global decision-making processes, does not produce any marked advantage in efficiency or speed over simple local decision-making heuristics for these problems. This is important for bidirectional

communication, as it shows that more complex decision-making algorithms may produce only marginal performance gains over simpler algorithms, at a cost of being far more difficult to explain to human teammates and the computational complexity of the algorithm itself.

Predictability of Decisions. An important part of human-agent teams is the extent to which agents can predict one another's actions. This can be viewed as a function of the number of different solutions that a group of agents will produce, or that a stochastic algorithm will produce on successive runs. Greater differences among solutions indicates that those solutions will be harder for teammates to predict, whereas if all teammates' solutions converge to only a few possibilities it will be easier to predict their actions. In order to examine the predictability of human solutions, we calculated the mean pairwise divergence among all solutions to each of the problems, in each condition. The main effect of environmental complexity was characterized by a general increase in divergence with increasing environmental complexity, once the complexity of the environment increased beyond a threshold; in the present study, the threshold was five boxes. This means that most humans will make similar solutions when environmental complexity is low suggesting that additional explicit communication may not be needed since the likelihood that expectations will match behaviors is high. However, when the environmental complexity reaches a set level, the number of possible solutions and variance between those solutions greatly increases leading to more unpredictable human behavior. These results also showed that with the exception of Level 6, after this threshold participants solutions diverged more when the box numbers were not shown. One interpretation of this finding, in light of the previous results, is that participants' reliance on local decision-making heuristics when the box numbers were shown reduced the variance across their solutions.

Summary Discussion. The results described above, taken together, show that local decision-making heuristics are sufficient for this task, as global processing takes longer, does not improve performance or efficiency, and increases the divergence among participants' solutions. These solutions would thus be harder for teammates to predict. Beyond a certain level of environmental complexity, bidirectional communication becomes increasingly important because the range of possible solutions to a given problem increases substantially. In these cases, bidirectional communication will be necessary to promote shared situation awareness and trust, and to facilitate fluid, flexible interaction between humans and non-human intelligent agents.

4.2 Implications on Algorithm Development

Bidirectional communication has various impacts on the development of algorithms for the robot decision-making in collaborative missions. Specifically, shared understanding of the mission goals and mental models could minimize uncertainty in team decision making, and result in more predictable consequences. From the algorithmic perspective, predictable results after taking actions would greatly reduce online computations such as replanning because less deviations from the original plan would be observed. Also, understanding human decision-making and limitations could help developing robots that can autonomously decide when and how to help humans. For example,

robot might explicitly offer help when human spends a lot of time for planning the move, or robot could infer a specific part of the problem (e.g., furthest area from the human) and start working on that to shrink the decision problem of human. Overall, both implicit (e.g., posture, gesture) or explicit (e.g., natural language, feedback through displays) communication play an important role when developing decision-making strategies for robots that are expected to operate with humans in complex missions.

Acknowledgment. Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-10-2-0016. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

References

1. Scholtz, J.: Theory and evaluation of human robot interactions. In: 36th Annual Proceedings of the International Conference on System Sciences. IEEE, Hawaii (2003)
2. Schaefer, K.E., Straub, E.R., Chen, J.Y.C., Putney, J., Evans, A.W.: Communicating intent to develop shared situation awareness and engender trust in human-agent teams. *Cogn. Syst. Res.* **46**, 26–39 (2017)
3. Groom, V., Nass, C.: Can robots be teammates? Benchmarks in human-robot teams. *Interact. Stud.* **8**(3), 483–500 (2007)
4. Oh, J., et al.: Integrated intelligence for human-robot teams. In: Kulić, D., Nakamura, Y., Khatib, O., Venture, G. (eds.) ISER 2016. SPAR, vol. 1, pp. 309–322. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-50115-4_28
5. Kaupp, T., Makarenko, A., Durrant-Whyte, H.: Human-robot communication for collaborative decision making – A probabilistic approach. *Robot. Auton. Syst.* **58**(5), 444–456 (2010)
6. Lyons, J.B.: Being transparent about transparency: a model for human-robot interaction. In: Proceedings of the AAAI Spring Symposium, Palo Alto, CA. AAAI (2013)
7. Sycara, K., Sukthankar, G.: Literature Review of Teamwork Models (Report No. CMU-RI-TR-06-50). Carnegie Mellon University, Pittsburgh (2006)
8. Schaefer, K.E., Chen, J.Y.C., Szalma, J.L., Hancock, P.A.: A meta-analysis of factors influencing the development of trust in automation: Implications for understanding autonomy in future systems. *Hum. Factors* **58**(3), 377–400 (2016)
9. Dry, M., Lee, M.D., Vickers, D., Hughes, P.: Human performance on visually presented traveling salesperson problems with varying numbers of nodes. *J. Probl. Solving* **1**(1), 20–32 (2006)
10. MacGregor, J.N., Ormerod, T.: Human performance on the traveling salesman problem. *Percept. Psychophys.* **58**(4), 527–539 (1996)
11. Perelman, B.S., Mueller, S.T.: Identifying mental models of search in a simulated flight task using a pathmapping approach. In: Proceedings of the 18th International Symposium on Aviation Psychology (2015)
12. Pizlo, Z., Stefanov, E., Saalweachter, J., et al.: Traveling salesman problem: a foveating pyramid model. *J. Probl. Solving* **1**(8), 83–101 (2006)

13. Best, B.J., Simon, H.A.: Simulating human performance on the traveling salesman problem. In: Proceedings of the 3rd International Conference on Cognitive Modeling, Groningen, Netherlands, pp. 42–49. ICCM (2000)
14. Perelman, B.S., Mueller, S.T.: Considerations influencing human TSP solutions and modeling implications. In: Reitter, D., Ritter, F.E. (eds.) Proceedings of the 14th International Conference on Cognitive Modeling. University Park, PA (2016)
15. Kong, X., Schunn, C.D.: Information seeking in complex problem solving. In: Proceedings of the 8th International Conference on Cognitive Modeling, Oxford, UK, pp. 261–266. ICCM (2007)
16. Blum, A., Chawla, S., Karger, D.R., et al.: Approximation algorithms for orienteering and discounted-reward TSP. *SIAM J. Comput.* **37**, 653–670 (2007)
17. Ragni, M., Wiener, J. M.: Constraints, inferences and the shortest path: Which paths do we prefer? In: Proceedings of the Annual Meeting of the Cognitive Science Society, pp. 2216–2221. *CogSci*, Sapporo (2012)
18. Tenbrink, T., Seifert, I.: Conceptual layers and strategies in our tour planning. *Cogn. Process.* **12**, 109–125 (2011)
19. Tenbrink, T., Wiener, J.: The verbalization of multiple strategies in a variant of the traveling salesman problem. *Cogn. Process.* **10**, 143–161 (2009)
20. Perelman, B.S., Mueller, S.T.: Examining memory for search using a simulated aerial search and rescue task. In: Proceedings of the 17th International Symposium on Aviation Psychology, Dayton, OH (2013)
21. Perelman, B.S., Evans III, A.W., Schaefer, K.E.: Mental model consensus and shifts during navigation system-assisted route planning. In: Proceedings of the Human Factors and Ergonomics Society, 61(1), 1183–1187. HFES, Austin (2017)
22. Brunyé, T.T., Taylor, H.A.: Extended experience benefits spatial mental model development with route but not survey descriptions. *Acta Psychol.* **127**, 340–354 (2008)
23. Jahn, G., Johnson-Laird, P.N., Knauff, M.: Reasoning about consistency with spatial mental models: hidden and obvious indeterminacy in spatial descriptions. In: Freksa, C., Knauff, M., Krieg-Brückner, B., Nebel, B., Barkowsky, T. (eds.) *Spatial Cognition 2004*. LNCS (LNAI), vol. 3343, pp. 165–180. Springer, Heidelberg (2005). https://doi.org/10.1007/978-3-540-32255-9_10
24. Tversky, B.: Spatial mental models. *Psychol. Learn. Motiv.* **27**, 109–145 (1991)
25. Clare, A.S., Maere, P.C.P., Cummings, M.L.: Assessing operator strategies for real-time replanning of multiple unmanned vehicles. *Intell. Decis. Technol.* **6**, 221–231 (2012)
26. Chen, J.Y.C., Barnes, M.J.: Human-agent teaming for multirobot control: a review of human factors issues. *IEEE Trans. Hum.-Mach. Syst.* **44**(1), 13–29 (2014)
27. Hancock, P.A., Billings, D.R., Schaefer, K.E.: Can you trust your robot? *Ergon. Des.* **19**(3), 24–29 (2011)
28. Osofsky, S., Schuster, D., Phillips, E., Jentsch, F.G.: Building appropriate trust in human-robot teams. In: Proceedings of the AAAI Spring Symposium. AAAI, Stanford (2013)
29. Zakzanis, K.K., Quintin, G., Graham, S.J., Mraz, R.: Age and dementia related differences in spatial navigation within an immersive virtual environment. *Med. Sci. Monit.* **15**(4), 140–150 (2009)
30. Moffat, S.D., Elkins, W., Resnick, S.M.: Age differences in the neural systems supporting human allocentric spatial navigation. *Neurobiol. Aging* **27**(7), 965–972 (2006)
31. Yuan, P., Daugherty, A.M., Raz, N.: Turning bias in virtual spatial navigation: age-related differences and neuroanatomical correlates. *Biol. Psychol.* **96**, 8–19 (2014)

32. Zahr, N.M., Rohlfing, T., Pfefferbaum, A., Sullivan, E.V.: Problem solving, working memory, and motor correlates of association and commissural fiber bundles in normal aging: a quantitative fiber tracking study. *NeuroImage* **44**(3), 1050–1062 (2009)
33. Jansson, E.S.V.: TSBK07 Computer Graphics Project Report for NQ Sokoban (2016). <https://caffeineviking.net/share/papers/rnqsok.pdf>
34. Mueller, S.T., Perelman, B.S., Veinott, E.S.: An optimization approach for measuring the divergence and correspondence between paths. *Behav. Res. Methods* **48**(1), 53–71 (2016)
35. Salas, E., Shuffler, M.L., Thayer, A.L., Bewell, W.L., Lazzara, E.H.: Understanding and improving teamwork in organizations: a scientifically based practical guide. *Hum. Resour. Manag.* **54**(4), 599–622 (2015)
36. Cooke, N.J., Gorman, J.C., Myers, C.W., Duran, J.L.: Interactive team cognition. *Cogn. Sci.* **37**(2), 255–285 (2013)
37. MacMillan, J., Entin, E.E., Serfaty, D.: Communication overhead: the hidden cost of team cognition. In: Salas, E., Fiore, S.M. (eds.) *Team Cognition: Understanding the Factors that Drive Process and Performance*, pp. 61–82. American Psychological Association, Washington, DC (2004)
38. Mercado, J.E., Rupp, M.A., Chen, J.Y.C., et al.: Intelligent agent transparency in human-agent teaming for Multi-UxV management. *Hum. Factors* **58**(3), 401–415 (2016)
39. Selkowitz, A.R., Larios, C.A., Lakhmani, S.G., Chen, J.Y.C.: Displaying information to support transparency for autonomous platforms. In: Savage-Knepshield, P., Chen, J.Y.C. (eds.) *Advances in Human Factors in Robots and Unmanned Systems*, pp. 161–173. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-41959-6_14
40. Chen, J.Y.C., Procci, K., Boyce, M., et al.: Situation awareness-based agent transparency (ARL-TR-6905). US Army Research Laboratory, Aberdeen Proving Ground, MD (2014)
41. Barber, D., Lackey, S., Reinerman-Jones, L., Hudwon, I.: Visual and tactile interfaces for bi-directional human robot communication. In: *Proceedings for the International Society for Optics and Photonics, SPIE Defence, Security, and Sensing conference*. SPIE, Baltimore (2013)
42. Lackey, S., Barber, D., Reinerman, L., Badler, N.I., Hudson, I.: Defining next-generation multi-modal communication in human robot interaction. In: *Proceedings of the Human Factors and Ergonomics Society*, vol. 55, no. 1, pp. 461–464 (2011)
43. Rau, P.P., Li, Y., Li, D.: Effects of communication style and culture on ability to accept recommendations from robots. *Comput. Hum. Behav.* **25**(2), 587–595 (2009)
44. Selkowitz, A.R., Lakhmani, S.G., Larios, C.N., Chen, J.Y.C.: Agent transparency and the autonomous squad member. In: *Proceedings of the Human Factors and Ergonomics Society*, vol. 60, no. 1, pp. 1319–1323. Sage, Los Angeles (2016)
45. Stowers, K., Kasdaglis, N., Rupp, M., Chen, J., Barber, D., Barnes, M.: Insights into human-agent teaming: intelligent agent transparency and uncertainty. In: Savage-Knepshield, P., Chen, J.Y.C. (eds.) *Advances in Human Factors in Robots and Unmanned Systems*, vol. 499, pp. 149–160. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-41959-6_13
46. Wright, J.L., Chen, J.Y.C., Barnes, M.J., Hancock, P.A.: Agent reasoning transparency: the influence of information level on automation-induced complacency (ARL-TR-8044). US Army Research Laboratory, Aberdeen Proving Ground, MD (2017)