Toward Initiative Decision-Making for Distributed Human-Robot Teams

Shu D. Jiang
School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA, USA

Jonathan Odom
Georgia Tech Research Institute
Georgia Institute of Technology
Atlanta, GA, USA

ABSTRACT

The brief history of human-robot teams can be traced through the changing perspective of a robot's role within the team, which has evolved from being treated as a tool to a recent shift toward the desire to have the robot act as an equal partner. While researchers have made tremendous strides in recent years, "making robots into team players" [1] that can work with humans as peers still presents a multitude of challenges. One key characteristic of a synergistic team is the ability to intervene or backup each other as necessary (e.g., when the other is underperforming). Hence, in this article, we formulate distributed human-robot teamwork in the framework of mixed-initiative interaction, which is an interaction strategy that lets the best-suited member of the team to perform the work by allowing team members to interleave their contributions to the overall team performance through opportunistic seizure and relinquishment of task initiatives. Specifically, this paper aims to address the issue of initiative decision-making – that is, *when* should a robot take over (relinquish) control from (to) a human teammate.

CCS CONCEPTS

• Computing methodologies → Intelligent agents • Computing methodologies → Multi-agent planning

KEYWORDS

Mixed-initiative interaction, human-robot team, search and rescue

ACM Reference format:

Shu Jiang and Jonathan Odom. 2018. Toward Initiative-Decision Making for Distributed Human-Robot Teams. In *Proceedings of the 6th Annual International Conference on Human-Agent Interaction (HAI'18), December 15-18, 2018,* Southampton, United Kingdom. ACM, 7 pages. https://doi.org/10.1145/as-assigned-during-rightsreview

ACM ISBN 978-1-4503-5953-5/18/12...$15.00
https://doi.org/10.1145/3284432.3284467

1 INTRODUCTION

Recent years have seen an increase in the use of robots in hazardous emergency response situations (i.e., potentially harmful to first responders) that range from natural disasters (e.g., Fukushima nuclear plant meltdown) to terrorist attacks (e.g., the World Trade Center (WTC) disaster). In these situations, human-robot teams are employed in a manner where the robots are controlled by one or more operators at a remote location, away from the danger zone. However, this distance creates a disconnect between the human and robot that presents some unique challenges for effective collaboration within the human-robot team (e.g., situational awareness, time delay). Moreover, once a disaster occurs, the human condition and physical site exacerbate the issues of fostering effective human-robot teamwork. Typical disaster sites caused by earthquakes are permeated with rubble piles, confined spaces, and unstable structures can greatly impair a robot's mobility and perceptual capabilities [2]. In addition, critical emergency situations also put first responders under constant stressors (e.g., time pressure and high-stake risks) that can cause fatigue and make them error-prone while operating and/or supervising robots [3].

Despite both the human and the robot having their own respective limitations when operating under the extreme conditions of disaster response missions, they each also have a set of complementary skills that if interleaved properly, can enable the human and robot to collaborate as an effective team. Several models have been proposed for the design of human-robot systems. The earliest and most primitive form of human-robot teamwork is teleoperation, where the robot complements the human's physical limitations and extends her physical presence. Supervisory control is then proposed to alleviate the significant workload of teleoperation by having the human act as a supervisor, who monitors the robot and intervenes whenever necessary. However, humans are found to be negligent as a supervisor [4]. In recent years, new models of human-robot systems have emerged that have both the human and the robot playing active roles in the overall performance of the team [5-13]. These advancements can be contributed to the adoption of theories and models of teamwork from the areas of human-automation interaction, human-computer interaction, psychology, cognitive science, and human teamwork [14-20].

However, many challenges still remain for "making robots into team players" [1]. One important aspect of distributed human-robot teamwork is where the robot can freely intervene and take over control of the task from the human teammate when it is deemed necessary to do so. This aspect of teamwork is especially important during critical disaster response missions where the human operator is prone to errors [3]. Furthermore, the behavior of backing up fellow team members has been identified as one of the "Big Fives" core components of effective teamwork [21]. Hence in this article, we formulate human-robot teamwork in the framework of mixed-initiative interaction, which is an interaction strategy that allows the human and robot to work together to achieve a common goal in a way that exploits their complementary capabilities through efficient interleaving of contributions [22-24]. The basic idea is to interleave team members' contributions to the overall team performance through opportunistic seizure and relinquishment of task initiatives.

While there are many issues when designing effective mixed-initiative human-robot teamwork, this paper focuses on the issue of initiative decision-making, or what is the computational mechanism for the robot to determine *when* it is appropriate to take control or to defer to its human teammate. The goal is toward building mixed-initiative human-robot teamwork, where the robot can freely intervene and take over control of the task from (or relinquish control to) the human teammate when it is deemed necessary to do so. In the next section, we briefly survey the state of practice of human-robot teamwork. In section 3, we present the interactive partially observable Markov decision process (I-POMDP) as the computational framework for initiative decision-making. Section 4 presents a victim search task as an application of initiative decision-making along with the simulation results of the task to illustrate the effectiveness of initiative decision-making mechanism. The last section concludes the paper and proposes the necessary future work.

2 RELATED WORK

The brief history of human-robot teams can be traced from the changing perspective of a robot's role within the team, which has evolved from being treated as a tool to a recent shift toward the desire to have the robot act as an equal partner/peer. Teleoperation is an early example of the teaming perspective where the robots are treated as tools. On the other extreme of the spectrum is supervisory control, where the robot is fully autonomous and able to carry out the designated task without human intervention [28]. In this paradigm, the robot is seen as a subordinate and the human as the supervisor, who can intervene to assist the robot whenever it is needed. However, humans are found to be easily bored and negligent as a supervisor [4].

While teleoperation and supervisory control are still active areas of research, the perspective of human-robot teamwork has shifted toward the middle between teleoperation and supervisory control, where the control of the task is shared between the human and robot. A notable shared control strategy is the autonomy-centered approach [25, 26], where the basic tenet is to let the autonomy of the robot change as the situation evolves. Before a system can adjust its autonomy based on situational demands, there first needs to be well-defined levels of autonomy. However, there are no widely accepted levels of autonomy for human-robot teamwork. Furthermore, Baker and Yanco [27] noted that human operators rarely change autonomy modes even when it would improve performance. Another critical issue of autonomy-centered approach is the issue of control decision authority [28] – that is, who decides when the control function must be shifted (e.g., from human to robot), a function that is distinct from the "direct controlling" function [29]. Principles of human-centered automation requires the human operator to be the final authority and that only she can decide when and how automation is changed [30]. However, Moray et al. [31] argued that final authority for decisions and action must be allocated to automation in time-critical situations.

Recent research efforts in human-robot teaming have also started to take advantage of the insights (i.e., theories and models) drawn from research in cognitive science, psychology, and human teamwork for building effective human-robot teams where the robot is envisioned as a partner. Hoffman and Breazeal [6] presented a human-robot collaboration architecture based on joint intention theory [15]. This theory predicts that effective teamwork requires team members to maintain a set of shared beliefs, demonstrate joint intention toward a shared goal and to provide mutual support [15]. The resulting collaborative system is able to dynamically assign sub-tasks between team members while taking into consideration the collaborator's abilities and the current task state [6]. Several works have drawn on recent findings in cognitive science on joint action. Sebanz et al. [16] identified several cognitive mechanisms that are involved in joint action: 1) joint attention, 2) action observation, 3) task-sharing, 4) action coordination, and 5) agency in joint action. Mutlu et al. explored how these mechanisms might be used to design collaborative robot behaviors through gaze cues to improve joint attention and action observation to monitor task breakdowns [7].

Perspective taking has been shown to occur in various collaborative situations, which is the ability of people to take one another's perspective, and may be used to predict what other people will do [8]. Trafton et al. [8] presented a computational cognitive model of perspective taking for human-robot teamwork. Ros et al. [10] used perspective taking for ambiguity resolution for ambiguous descriptions generated by the human partner. Research on mirror neurons [17] has inspired models of human-robot collaboration that involved a robot imitating or simulating the behaviors of its partner in order to make inferences about her actions [7]. Gray et al. proposed a model of teamwork that enabled a robot to observe the actions of its partner by using self as the simulator and make inferences about the partner's beliefs states to anticipate and offer help to the human partner as needed [11].

Recent work has also started to take advantage of the findings from research into human team interaction. Shah et al. presented a robot plan execution system that uses insights from human-human teaming such as the use of explicit and implicit coordination behaviors; the executive system is able to choose and schedule the robot's action, adapt to the human operator, and minimize the human's idle time [32]. Nikolaidis and Shah applied a shared mental model, a key teamwork process of human teams [19], for effective human-robot teaming [12]. Nikolaidis and Shah also presented a computational formulation of the robot's inter-role knowledge based on cross-training methodology, a widely used human team training strategy [20], for human-robot teamwork [13].

While significant advances have been made in human-robot teamwork, multitudes of challenges still remained. One such challenge is the backup behavior within the team, which is defined as *"the provision of task-efforts to another when there is recognition that there is a problem in the team"* [21]. The backup behavior would allow the robot to take over control of the task from the human teammate when it is necessary to do so, which is especially important during critical disaster response missions where the human operator is prone to errors [3]. Thus, our research aims to explore the use of backup behavior within a distributed human-robot team. The hypothesis is that *backup* behavior would afford a human-robot team fewer errors and greater adaptability. This work formulates the *backup* behavior in the framework of mixed-initiative human-robot interaction [33], which we examine in more detail in the following section.

3 COMPUTATIONAL MODEL OF INITIATIVE DECISION-MAKING

This section presents a computational model of initiative decision-making as the first step toward building effective mixed-initiative teams. The goal of mixed-initiative teaming strategy is to interleave contributions of team members in an effective manner to achieve a common goal. Interleaving of contributions reflects the basic idea of mixed-initiative interaction by letting the best operator perform the task. While effective interleaving of contributions is a desired manifestation of the mixed-initiative system, opportunistic intervention is a requirement to achieve such a desired effect. Hence, initiative decision-making is concerned with the question of *when* should a robot seize (relinquish) initiative from (to) a human teammate. Or more specifically, what is the appropriate reasoning mechanism the robot should employ to evaluate whether or not it should take control (relinquish) control from (to) the human operator? In the context of distributed human-robot teams, initiative decision-making can simply be viewed as an arbiter that switches between actions of human and robot over the course of a task, Figure 1.

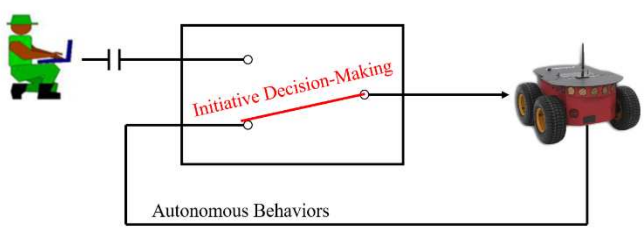

Figure 1: Initiative Decision-Making for Distributed Human-Robot Teams

However, "inescapable uncertainties" about human intention and the benefit of intervention at different times is an inherent issue for real world operations [34]. Furthermore, inappropriate timing of intervention could have severe consequences beyond mission failure, depending on the criticality of the mission. That means a successful mixed-initiative human-robot team would require an inherent capability of the mixed-initiative system for recognizing the opportunity (i.e., when) to assist the human operator (or ask for help) during a collaborative task in a timely manner. Consequently, effective teamwork necessitates each team member to factor fellow teammates' decision-making into consideration when deliberating among action choices. Thus, in the context of a human-robot team, the robot needs to take the human partner's actions into consideration when deciding whether to seize initiative from the human teammate. Hence, the issue of initiative decision-making is essentially a partially observable stochastic game.

There are two general perspectives governing approaches employed to the problem of decision-making in stochastic games [35]: objective and subjective. Objective approaches try to find plans for all agents centrally, which are then distributed and executed by the agents independently. Decentralized POMDP (Dec-POMDP) is a common objective approach to planning for a team of cooperative agents. Dec-POMDP is a generalization of the single-agent POMDP to multi-agent systems by considering joint actions and observations [35]. The goal is to find a joint policy that maximizes the expected cumulative reward under the assumption that the agents are fully cooperative. While Dec-POMDP is a natural framework for multi-robot teamwork, we argue that it may not be an appropriate for human-robot teamwork since the human might not execute a policy faithfully even if one can be assigned to them.

On the other hand, subjective approaches for multi-agent decision-making reason from the perspective of one particular agent. The subjective approaches consider each agent independently and have each agent maintain an explicit models of the other agents [35]. In essence, the goal of each agent is to take the best actions by predicting the actions of other agents in the environment. Interactive POMDP (I-POMDP) is a subjective approach that extends POMDP to multi-agent settings by Gmystrasiewicz and Doshi [36] to include the concept of agent models into the state space, which they termed the *interactive* state space. As a result, in addition to the beliefs about the physical environment, the interactive state space can also include the preferences, capabilities, and beliefs of other agents. Hence, I-POMDP allows planning from the robot's perspective while taking into account the beliefs and actions of the human teammate.

Formally, an I-POMDP for agent i is defined as a tuple

$$IPOMDP_i = < IS_i, A, T_i, \Omega_i, O_i, R_i >$$

where [36]:

- IS_i is a set of interactive states defined as $IS_i = S \times M_j$, where S is the set of states of the physical environment, and M_j is the set of possible models of agent j
- $A = A_i \times A_j$ is the set of joint moves of all agents

- T_i is the transition model
- Z_i is observation space
- O_i is an observation function
- R_i is defined as $R_i: IS_i \times A \rightarrow \mathcal{R}$

Solving I-POMDP is computationally very expensive. In addition to the complexity issues that it inherits from POMDP, I-POMDP also suffers from the curse of nested beliefs, where the agent's belief include the beliefs of other agents and their beliefs about the agent's beliefs, and so on [37]. Fortunately, for our application of initiative decision-making in distributed human-robot teams, we can take advantage of the partial observability of human actions to simplify the problem. In distributed human-robot teams, the human operator's actions are directly (albeit partially) observable to the robot since their commands are sent to the robot to be executed, Figure 1. The physical environment is then changed by the robot, which either carries out the human's desired action, modifies the action, or takes over as necessary. Hence, modeling the beliefs of the human teammate can be replaced with the directly observed actions of the human for initiative decision-making.

With the simplified I-POMDP as our computational framework, the algorithm for initiative decision-making is shown in Figure 2. The algorithm takes as input the model of the decision problem and outputs the initiative decision – that is, to seize, relinquish, or follow initiative. The algorithm first generates an initiative policy (using a POMDP solver [38]) in the form of a policy graph, where each node consists of an action and edges represent transitions based on observations. At every time step, the task state and human action are estimated, which form the observation state that is used to query the policy for the initiative decision. The initiative decision is then executed by the robot through either carrying out the human command or acting on its own initiative.

Algorithm 1 Initiative Decision-Making

Input:	IS, A, O, T, Z, R
Output:	a_R *// initiative decision*

// get policy in the form of a policy graph
$\pi = generatePolicy(IS, A, O, T, Z, R)$
while !endOfTask **do**
    $a_H \leftarrow estimateHumanAction()$
    $s \leftarrow estimateTaskState()$
    // observation = {task state} x {human action}
    $z \leftarrow \{s, a_H\}$
    // query the policy for initiative decision
    $a_R \leftarrow \pi(z)$ *// e.g., follow or seize initiative*
    $execute(a_R)$
end while

Figure 2: Algorithm for Initiative Decision-Making

The algorithm above presents the high-level steps of initiative decision-making, where the detailed implementation of each step depends on the specific task of concern and the type of initiative that the robot is allowed to take. We previously defined initiative as "*an element of a task that can range from low-level motion control of the robot to high-level specification of task goals*" [39]. Consequently, the output of *estimateHumanAction()* would depend on the initiative type, which can be high-level verbal commands from human or low-level joystick inputs. One example of low-level type of initiative is the safety-initiative, where the robot is allowed to protect itself by engaging the obstacle avoidance behavior. For this safety-initiative, the function would estimate whether the human is driving forward, backward, turning left or right from the joystick inputs. On the other hand, for higher level initiative such as a search task, the *estimateHumanAction()* function would estimate high level behaviors such as whether the human is searching a local area or moving to new areas. Furthermore, different type of initiatives can be combined to form a more comprehensive backup behaviors that involves different type of initiatives. For instance, safety initiative and search initiative can be combined to result in robot that has the initiative to protect itself and the initiative to take over search task from the human operator when necessary to do so.

Similarly, while the initiative decision space consists of a small set of actions, e.g., {seize, follow, relinquish}, the result of $execute(a_R)$ can be very different depends on the type of initiative. For instance, for the safety initiative, seize means the robot would engage the avoid obstacle behaviors; and for search initiative, seize when human is spending too much time searching a local area means that the robot would move to a new area. This, in effect, separates initiative decision-making from the actual low-level implementation of behavioral actions (e.g., obstacle avoidance), which simplifies problem formulation and render it more generalizable to different kinds of task and initiative types.

3 EXPERIMENT AND RESULTS

3.1 Initiative Decision-Making for Victim Search

The Kobe earthquake and Oklahoma City bombing motivated the development of robots for humanitarian efforts in search and rescue of trapped victims and propelled the emergence of urban search and rescue (USAR) as an important area of research for robotics [3]. These efforts led to the first use of robots for search and rescue at the World Trade Center disaster in 2001 [3]. While the use of robots at disaster sites has found tremendous success, significant challenges still remain. For instance, researchers have found that 50% of the terminal failures in disaster robots are caused by human error [40]. We posit that with the presented mixed-initiative framework, these terminal failures can be drastically reduced and mitigated by allowing the robot to take initiative when human error occurs during disaster rescue operations.

Hence, as an illustrative example of initiative decision-making, we examine a disaster response scenario where the distributed human-robot team is tasked to search within an unknown environment for disaster victims without a-priori information about their locations. In order to backup the human teammate effectively, the robot needs to decide at every moment in time whether to follow human commands or to take control to act autonomously. To formulate this initiative decision problem in the framework presented in the previous section, we start with the interactive state, which we simplified in this work to consists of the physical state of the search task, S, and human actions, A_H:

$$IS = \{S\} \times \{A_H\}$$

Based on behavioral ecology literature of human foraging behaviors [41], human search behaviors can broadly be categorized into two high level behaviors: exploitation and exploration. Exploitation is the behavior of thoroughly searching a local area, which is limited by the robot's visual sensory range for victim detection and identification. The goal of exploitation is to thoroughly search a local area to ensure completeness before leaving to search a distant area. However, the law of diminishing returns states that the probability of finding a victim in one area decreases as the area is exhaustively exploited. Hence, the exploration behavior is required to move the robot to a new area when it is no longer productive to stay in one area. The goal of exploration is to cover as much area as fast as possible by continually moving to the next best view location, or a location that provides the most information gain. While exploration ensures efficient area coverage in a given time interval, it can overlook some areas that might contain victims. An efficient search behavior is then simply the efficient tradeoff between exploiting one locale and exploring others. Consequently, we defined the human action space for the search task as

$$A_H = \{exploit, explore\}$$

Furthermore, the state of the search task is defined in terms of exploitation state, S_e, and victim status, S_v:

$$S = \{S_e\} \times \{S_v\}$$

The exploitation state S_e tries to capture whether the robot is spending sufficient time searching a local area before moving on to a new area. In the search task, S_e is defined as the percentage of the robot's surrounding area that has been searched (i.e., looked at using its onboard camera) for potential victims. The objective of search is to find a victim, hence, the state of whether a victim is found, S_v, is part of the state space. For this example, exploitation state is discretized into 10 levels, while victim-found status is a Boolean variable indicating whether a victim has been found. As a result, the interactive state becomes:

$$IS = \{A_H\} \times \{S_e\} \times \{S_v\}$$

At every moment in time, the robot needs to decide whether to follow human initiative (i.e., follow human command) or to seize initiative from a human (i.e., take control from a human). For instance, if the human is spending too much time exploiting a local area, the robot might seize the initiative to explore new areas. Hence, the initiative decision space represents the action space of the robot:

$$A_R = \{follow, seize\}$$

The observation space is the same as the interactive state space, which includes the state of exploitation, victim status, and human action. The transition model describes how the interactive state changes after a specific action in executed; for this example, we assume human action has a tendency to stay the same from one time step to the next and the level of exploitation state can only change one level at a time. Furthermore, we assume the exploitation behavior has a higher probability of finding a victim than the exploration behavior. Lastly, the goal of search is to find a victim, hence a positive reward is given when a victim is found, i.e., when $S_v = 1$:

$$R(S_v = 1) = 10$$

3.2 Experimental Setup

The operating hypothesis of this work is that the performance of a distributed human-robot team would be improved when the robot is allowed to back up, or take over control opportunistically from the human teammate. To validate the initiative decision-making for the victim search task, a simulation experiment was conducted to compare the search performances of different search behaviors. The experimental search environment is shown in Figure 3 in the Gazebo simulation framework. The area of the search environment is approximately 60m by 60m. Ten simulated victims are distributed across the environment along with blue markers for ease of victim detection (since victim detection is not the primary concern of this research). The robot used for the experiment is a simulated Pioneer 3-AT, a four-wheel drive mobile platform, equipped with a camera with a field of view of 60 degrees and a laser range-finder. The robot has neither the prior knowledge of the environment nor the number and locations of the simulated disaster victims.

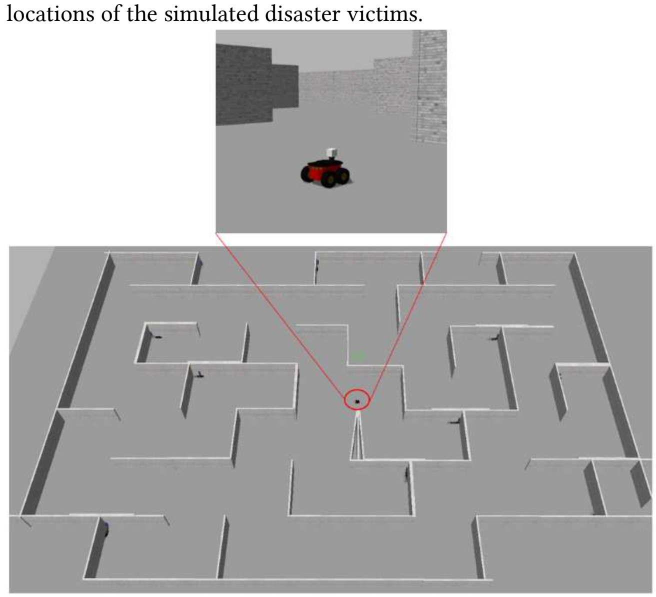

Figure 3: Environment for Victim Search in Gazebo Simulation

The experiment consists of two basic steps. First, we simulated three human search models, with different level of expertise at victim search, that generates joystick commands to control the robot. The three simulated human search performance profiles are:

1. *Novice* – this profile simulates a human operator who searches an environment by moving the robot in random directions; this constitutes a poor human search performance, which would present opportunities for the robot to seize initiative.
2. *Expert* – this profile simulates a human operator who knows the victim locations, hence able to search the environment optimally; ideally, the robot would not need to intervene in this case.
3. *Intermediate* – this profile simulates an average human operator, who acts as a *Novice* operator half of the time and an *Expert* operator the other half of the time

Second, we compare the condition of ***mixed-initiative,*** where the robot is allowed to take initiative to intervene and take over from the (simulated) human, to the condition of *teleoperation*, where the robot does not take initiative and simply follows the (simulated) human's joystick commands. When taking initiative, without prior knowledge of the environment and victim locations, the robot utilizes a frontier-based exploration strategy that seeks the most gain in area coverage of the unknown environment when deciding where to move next [39]. In effect, this search behavior would always move the robot toward the largest frontier in current area covered, where a frontier is defined as the boundary between explored and unexplored areas.

Our hypothesis is that by allowing the robot to seize initiative opportunistically to backup the human operator, the performance of the human-robot system would be better than the case where the robot does not take such initiatives. Furthermore, we also hypothesize that the number of interventions would decrease as the competency of the human operator is increased. In our case, that is, the robot would intervene most often with the *Novice* human operator and least frequently with the *Expert* operator.

3.3 Results

Using thirty Monte-Carlo runs, the results of the experiment are summarized in Figures 4 and 5. Figure 4 shows the performances of the search task of the human-robot system with the three different types of human operators, for both the *mixed-initiative* and *teleoperation* conditions respectively, in terms of the number of victims found and the time at which each victim was found. First, as expected, regardless of the teaming strategy, the human-robot system with the *Novice* human operator has the worst performance while the system with the *Expert* human has the best performance. Second, the performances of the human-robot systems are significantly improved for both the *Novice* ($p = 0.0025$) and *Intermediate* ($p = 0.027$) human operators. Lastly, there is no significant difference ($p = 0.565$) in performances between the human-robot systems with and without robot initiative for the *Expert* human operator. This makes sense since the robot would not need to intervene the *Expert* human operator who is able to perform the search task optimally. This is also evident in Figure 5, where the result shows that the robot rarely intervened when the *Expert* operator was in control of the robot. On the other hand, we found that the robot intervened most frequently when the *Novice* human was operating the robot. The results illustrated that the robot was able to seized initiative appropriately to improve the performance of the human-robot team.

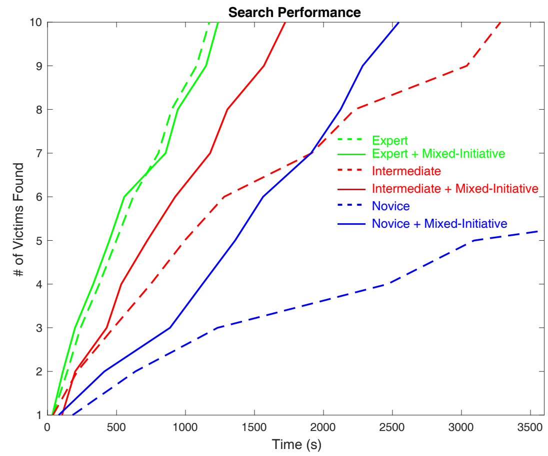

Figure 4: Search Performances of the Human-Robot Team with Different Levels of Search Expertise and Initiative Conditions

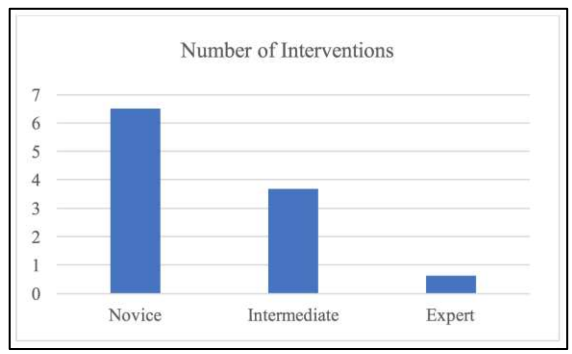

Figure 5: The Average Number of Times that the Robot Seized Initiative for Operators of Different Search Expertise

4 DISSCUSSIONS AND FUTURE WORK

Motivated by the importance of backup behaviors in human teams, this work addressed the question of when such behaviors should be engaged by a robot while working with a remote human teammate. We formulated this problem of initiative decision-making in the framework of I-POMDP, which offers a subjective approach in solving stochastic games. A simulation experiment of a search task was conducted to illustrate the effectiveness of the initiative decision-making mechanism, which allows the robot to seize initiative from a simulated human teammate. The results of the experiment validated our hypothesis that when the robot is allowed to opportunistically intervene and take over from a poorly-performing teammate, the overall performance of team can be improved.

While this experiment is limiting due to the nature that simulated human search behaviors were used, it nonetheless demonstrated the effectiveness of the mixed-initiative strategy in improving the team performance when the robot is able to backup a poorly-performing teammate. Our next step is to further validate our hypotheses of backup behavior and the model of initiative decision-making with human subject studies. Furthermore, this work did not address the issue of trust when the robot is allowed to seize initiative from a human teammate. Trust is an important issue of human-robot teamwork, and robot taking control from human could exacerbate the issue. Hence, another future work is to study *how* to seize initiative after the decision to intervene is made such that trust can be maintained.

REFERENCES

[1] G. Klein, D. D. Woods, J. M. Bradshaw, R. R. Hoffman, and P. J. Feltovich, "Ten challenges for making automation a" team player" in joint human-agent activity," *IEEE Intelligent Systems,* vol. 19, no. 6, pp. 91-95, 2004.

[2] J. L. Burke, R. R. Murphy, M. D. Coovert, and D. L. Riddle, "Moonlight in Miami: Field study of human-robot interaction in the context of an urban search and rescue disaster response training exercise," *Human–Computer Interaction,* vol. 19, no. 1-2, pp. 85-116, 2004.

[3] R. R. Murphy, "Human-robot interaction in rescue robotics," *IEEE Transactions on Systems, Man, and Cybernetics,* 2004.

[4] P. J. Hinds, T. L. Roberts, and H. Jones, "Whose job is it anyway? A study of human-robot interaction in a collaborative task," *Human-Computer Interaction,* vol. 19, no. 1, pp. 151-181, 2004.

[5] T. Fong, C. Thorpe, and C. Baur, "Collaboration, dialogue, human-robot interaction," in *Robotics Research*: Springer, 2003, pp. 255-266.

[6] G. Hoffman and C. Breazeal, "Collaboration in human-robot teams," in *Proc. of the AIAA 1st Intelligent Systems Technical Conference, Chicago, IL, USA*, 2004.

[7] B. Mutlu, A. Terrell, and C.-M. Huang, "Coordination mechanisms in human-robot collaboration," in *Proceedings of the Workshop on Collaborative Manipulation, 8th ACM/IEEE International Conference on Human-Robot Interaction*, 2013.

[8] J. G. Trafton, N. L. Cassimatis, M. D. Bugajska, D. P. Brock, F. E. Mintz, and A. C. Schultz, "Enabling effective human-robot interaction using perspective-taking in robots," *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on,* vol. 35, no. 4, pp. 460-470, 2005.

[9] M. Berlin, J. Gray, A. L. Thomaz, and C. Breazeal, "Perspective taking: An organizing principle for learning in human-robot interaction," in *Proceedings of the National Conference on Artificial Intelligence*, 2006, vol. 21, no. 2, p. 1444: Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999.

[10] R. Ros *et al.*, "Which one? grounding the referent based on efficient human-robot interaction," in *RO-MAN, 2010 IEEE*, 2010, pp. 570-575: IEEE.

[11] J. Gray, C. Breazeal, M. Berlin, A. Brooks, and J. Lieberman, "Action parsing and goal inference using self as simulator," in *Robot and Human Interactive Communication, 2005. ROMAN 2005. IEEE International Workshop on*, 2005, pp. 202-209: IEEE.

[12] S. Nikolaidis and J. Shah, "Human-robot teaming using shared mental models," *ACM/IEEE HRI,* 2012.

[13] S. Nikolaidis and J. Shah, "Human-robot cross-training: Computational formulation, modeling and evaluation of a human team training strategy," in *Proceedings of the 8th ACM/IEEE international conference on Human-robot interaction*, 2013, pp. 33-40: IEEE Press.

[14] T. B. Sheridan, "Function allocation: algorithm, alchemy or apostasy?," *International Journal of Human-Computer Studies,* vol. 52, no. 2, pp. 203-216, 2000.

[15] P. R. Cohen and H. J. Levesque, "Teamwork," *Nous,* pp. 487-512, 1991.

[16] N. Sebanz, H. Bekkering, and G. Knoblich, "Joint action: bodies and minds moving together," *Trends in cognitive sciences,* vol. 10, no. 2, pp. 70-76, 2006.

[17] V. Gallese and A. Goldman, "Mirror neurons and the simulation theory of mind-reading," *Trends in cognitive sciences,* vol. 2, no. 12, pp. 493-501, 1998.

[18] E. Salas, C. S. Burke, and J. A. Cannon‐Bowers, "Teamwork: emerging principles," *International Journal of Management Reviews,* vol. 2, no. 4, pp. 339-356, 2000.

[19] J. E. Mathieu, T. S. Heffner, G. F. Goodwin, E. Salas, and J. A. Cannon-Bowers, "The influence of shared mental models on team process and performance," *Journal of applied psychology,* vol. 85, no. 2, p. 273, 2000.

[20] C. E. Volpe, J. A. Cannon-Bowers, E. Salas, and P. E. Spector, "The impact of cross-training on team functioning: An empirical investigation," *Human Factors: The Journal of the Human Factors and Ergonomics Society,* vol. 38, no. 1, pp. 87-100, 1996.

[21] E. Salas, D. E. Sims, and C. S. Burke, "Is there a “Big Five” in teamwork?," *Small group research,* vol. 36, no. 5, pp. 555-599, 2005.

[22] J. Allen, C. Guinn, and E. Horvtz, "Mixed-initiative interaction," *IEEE Intelligent Systems and their Applications,* vol. 14, no. 5, 1999.

[23] G. Tecuci, M. Boicu, and M. T. Cox, "Seven aspects of mixed-initiative reasoning: An introduction to this special issue on mixed-initiative assistants," *AI Magazine,* vol. 28, no. 2, p. 11, 2007.

[24] E. Horvitz, "Principles of mixed-initiative user interfaces," *SIGCHI conference on Human factors in computing systems,* 1999.

[25] G. Dorais, R. P. Bonasso, D. Kortenkamp, B. Pell, and D. Schreckenghost, "Adjustable autonomy for human-centered autonomous systems," in *Working notes of the Sixteenth International Joint Conference on Artificial Intelligence Workshop on Adjustable Autonomy Systems*, 1999, pp. 16-35.

[26] M. A. Goodrich, D. R. Olsen, J. W. Crandall, and T. J. Palmer, "Experiments in adjustable autonomy," in *Proceedings of IJCAI Workshop on Autonomy, Delegation and Control: Interacting with Intelligent Agents*, 2001, pp. 1624-1629.

[27] M. Baker and H. A. Yanco, "Autonomy mode suggestions for improving human-robot interaction," in *SMC (3)*, 2004, pp. 2948-2953.

[28] T. Inagaki, "Adaptive automation: Sharing and trading of control," *Handbook of cognitive task design,* vol. 8, pp. 147-169, 2003.

[29] T. B. Sheridan, "Adaptive automation, level of automation, allocation authority, supervisory control, and adaptive control: Distinctions and modes of adaptation," *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on,* vol. 41, no. 4, pp. 662-667, 2011.

[30] C. E. Billings, *Aviation automation: The search for a human-centered approach*. 1997.

[31] N. Moray, T. Inagaki, and M. Itoh, "Adaptive automation, trust, and self-confidence in fault management of time-critical tasks," *Journal of Experimental Psychology: Applied,* vol. 6, no. 1, p. 44, 2000.

[32] J. Shah, J. Wiken, B. Williams, and C. Breazeal, "Improved human-robot team performance using chaski, a human-inspired plan execution system," in *Proceedings of the 6th international conference on Human-robot interaction*, 2011, pp. 29-36: ACM.

[33] S. Jiang and R. C. Arkin, "Mixed-initiative human-robot interaction: definition, taxonomy, and survey," in *Systems, Man, and Cybernetics (SMC), 2015 IEEE International Conference on*, 2015, pp. 954-961: IEEE.

[34] E. J. Horvitz, "Reflections on challenges and promises of mixed-initiative interaction," *AI Magazine,* vol. 28, no. 2, p. 3, 2007.

[35] F. A. Oliehoek and A. Visser, "A decision-theoretic approach to collaboration: Principal description methods and efficient heuristic approximations," in *Interactive collaborative information systems*: Springer, 2010, pp. 87-124.

[36] P. J. Gmytrasiewicz and P. Doshi, "A framework for sequential planning in multi-agent settings," *J. Artif. Intell. Res.(JAIR),* vol. 24, pp. 49-79, 2005.

[37] T. N. Hoang and K. H. Low, "Interactive POMDP Lite: Towards Practical Planning to Predict and Exploit Intentions for Interacting with Self-Interested Agents," in *IJCAI*, 2013, pp. 2298-2305.

[38] A. R. Cassandra, *Exact and approximate algorithms for partially observable Markov decision processes*. Brown University, 1998.

[39] S. Jiang and R. C. Arkin, "SLAM-based spatial memory for behavior-based robots," *IFAC-PapersOnLine,* vol. 48, no. 19, pp. 195-202, 2015.

[40] R. R. Murphy, *Disaster robotics*. MIT press, 2014.

[41] R. L. Goldstone and B. C. Ashpole, "Human foraging behavior in a virtual environment," *Psychonomic bulletin & review,* vol. 11, no. 3, pp. 508-514, 2004.