

# Human prior knowledge estimation from movement cues for information-based control of mobile robots during search

RAFAL KRZYSIAK, University of California, Merced, USA and Northern Illinois University, USA

SACHIT BUTAIL, Northern Illinois University, USA

Robotic search often involves teleoperating vehicles into unknown environments. In such scenarios, prior knowledge of target location or environmental map may be a viable resource to tap into and control other autonomous robots in the vicinity towards an improved search performance. In this paper, we test the hypothesis that despite having the same skill, prior knowledge of target or environment affects teleoperator actions, and such knowledge can therefore be inferred through robot movement. To investigate whether prior knowledge can improve human-robot team performance, we next evaluate an adaptive mutual-information blending strategy that admits a time-dependent weighting for steering autonomous robots. Human-subject experiments show that several features including distance travelled by the teleoperated robot, time spent staying still, speed, and turn rate, all depend on the level of prior knowledge and that absence of prior knowledge increased workload. Building on these results, we identified distance travelled and time spent staying still as movement cues that can be used to robustly infer prior knowledge. Simulations where an autonomous robot accompanied a human teleoperated robot revealed that whereas time to find the target was similar across all information-based search strategies, adaptive strategies that acted on movement cues found the target sooner more often than a single human teleoperator compared to non-adaptive strategies. This gain is diluted with number of robots, likely due to the limited size of the search environment. Results from this work set the stage for developing knowledge-aware control algorithms for autonomous robots in collaborative human-robot teams.

CCS Concepts: • Human-centered computing → HCI theory, concepts and models; Empirical studies in HCI; • Computer systems organization → Robotics.

Additional Key Words and Phrases: Human-robot interaction, Information based control, Teleoperation, Prior knowledge, Search and Rescue

## ACM Reference Format:

Rafal Krzysiak and Sachit Butail. 2024. Human prior knowledge estimation from movement cues for information-based control of mobile robots during search. 1, 1 (December 2024), 29 pages. <https://doi.org/10.1145/nnnnnnnn>.

## 1 INTRODUCTION

Knowledge about a situation plays an important role in how we respond and learn from it [30], [59]. Indeed, when making decisions under time constraint, humans are expected to possess and utilize prior knowledge towards an optimal solution [56]. Examples of tasks that involve critical decision making in the real world include search and rescue, where prior knowledge about missing person behavior or terrain [61] may lead to quicker times to find [2], medical surgery where knowledge

---

Authors' addresses: Rafal Krzysiak, University of California, Merced, 5200 North Lake Road, Merced, California, USA, 95343 and Northern Illinois University, 590 Garden Road, DeKalb, Illinois, USA, 60115, [rkrzysiak@ucmerced.edu](mailto:rkrzysiak@ucmerced.edu); Sachit Butail, [sbutail@niu.edu](mailto:sbutail@niu.edu), Northern Illinois University, 590 Garden Road, DeKalb, Illinois, USA, 60115.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2024 Association for Computing Machinery.

XXXX-XXXX/2024/12-ART \$15.00

<https://doi.org/10.1145/nnnnnnnnnnnnnn>

about system capabilities can aid surgical tasks that require precision [58, 67], and satellite control where knowledge about contact dynamics between two satellites may help in maintaining pose stability during servicing and rendezvous [14, 45].

Determining how that prior knowledge affects human actions can play an integral role in building efficient human-robot interaction (HRI) systems, where robots must respond intelligently to human behavior. Inferring human prior knowledge indirectly through actions can advance control blending strategies in HRI systems [10, 22] and enable new control sharing architectures for large collaborative systems [4, 16]. For example, in the case of multi robot systems where a human may operate one of many robots [38, 42], the ability to detect, infer, and act on human prior knowledge provides a meaningful way to make other autonomous robots respond to human actions directly in the field [35, 66]. However, communicating prior knowledge directly to robots (for example, by responding to survey questions [7]) poses additional burden on the human and makes an unreasonable assumption that prior knowledge remains constant over the course of the mission. In this context, inferring prior knowledge directly from the motion of the teleoperated robot can enable new human-aware swarm control strategies in the field [32, 40].

Prior knowledge may not necessarily translate to skill [3, 47]. For example, a skillful drone operator may still methodically explore a new environment looking for a missing person. Accordingly, their search strategy would be different compared to someone who knows the region or high-probability target locations. Indeed, it is possible that existence of prior knowledge could lead to differences in intent and therefore actions during a search operation. The underlying hypothesis of this paper is that different search strategies translate to different actions and therefore movement of the teleoperated robot on the field. Furthermore, we expect that such differences in movement can provide the means to inform the actions of other autonomous robots towards a more effective human-robot team.

Inferring skill from actions is not new. Studies in driver modeling, aircraft piloting, and telerobotics have shown that skill and intent can be inferred from observing the trajectories of the vehicle or robotic manipulator [11, 54, 64, 65]. Compared to using physiological data [33], estimating human knowledge from robot movement has the potential for tighter integration of autonomous inputs with human commands in the field, and fewer concerns for privacy. With respect to teleoperation for example, it has been shown that teleoperators perform differently depending on skill in operating the robot. Experiments within a variety of environments and setups [23, 27, 34, 55, 69] have shown that tasks such as obstacle avoidance, navigation, and control depend on the experience with the system at hand, camera placement, and presence of haptic feedback. To the best of our knowledge, whether similar differences in teleoperation exist with respect to prior knowledge of the mission have not been explored.

With respect to search by teleoperating a mobile robot, prior knowledge may be related to environmental map and target location. This is similar to search scenarios where searchers are either aware of the target's (missing person or animal) behavior, and therefore likely regions to look for, or the local terrain [13, 41], or both. In addition to mission performance, we expect that prior knowledge can also affect teleoperation actions and in turn robot movement on the field. If inferred reliably, the estimated prior knowledge can then be used to drive control actions of other autonomous robots in the field.

Accordingly, in this paper, we address the following questions: (i) does prior knowledge affect teleoperated robot movement during a search scenario? (ii) what robot movement cues can be extracted from limited observation to infer prior knowledge reliably? and (iii) whether such prior knowledge inference can be used to weight information-based objectives in shared control of autonomous robots assisting in search within a human-robot team?

We address the first question by performing laboratory experiments where participants are asked to search for a missing target as information about target location and environment is varied. In addition to tracking robot movement we survey participant workload and their reactions to possible inaccuracies in prior knowledge. Prior knowledge is inferred by analyzing features of robot movement that are found to be statistically different between conditions. In particular, we optimize for trajectory length that provides the maximum difference, in terms of Wasserstein distance, between probability distributions of select features conditioned on prior knowledge. These features are then used in a simulation study of human-robot teams where the autonomous robots observe and adaptively assist the teleoperated robot moving per experimental trajectories in searching for the missing target. In particular, we expand an information-based control strategy that blends two mutual information objectives for determining the control actions of autonomous robots. This involves an adaptive strategy that performs real-time weighting of objectives for autonomous robots competing between assisting the teleoperator or searching independently. Performance of this approach is measured in terms of the average time to find the target, fraction of trials in which the human-robot team is able to find the target sooner than the teleoperated robot only, and the average time gained in those trials where the team found the target earlier than a single teleoperated robot. We simulate the adaptive search strategy informed by three different sets of movement cues with two- and three-member human-robot teams. Their performance is compared with three non-adaptive strategies where autonomous robots: (i) move randomly in the environment, (ii) search independently and (iii) only assist the human teleoperated robot.

Our experimental results show that prior knowledge affects search behavior and robot movement on the ground. In particular, robot distance moved, time spent staying still (also called freezing time), speed, and turn rate, all differ significantly based on prior knowledge. Out of these, distance moved and freezing time calculated for 15- and 30-second segments show largest difference between probability distributions arising from different types of prior knowledge. When these values are used in real time to infer prior knowledge and update the weighting of objectives, they register the most improvement with a one-human one-autonomous robot team in terms of early finding of the target when the human teleoperator had no knowledge of target location. When the number of autonomous robots are increased to two, there is no significant gain in performance over independent search, likely due to additional collision avoidance strategies taking over the navigation in the relatively small and cluttered search environment.

This paper is organized as follows: Section 2 highlights related work on the role of human skill in search and rescue, and inferring human intent and strategies for human-robot teaming. In Section 3 we provide the preliminary material on mutual information based control and information blending strategy for robotic assistance during search. The experiments are presented next in Section 4 with details on the experimental setup and results comparing average values of features related to robot movement cues as tracked across different conditions. Section 5 describes the prior knowledge classification based on movement cues and sets up the adaptive search strategy utilizing such cues. This is followed by a study built on experimental data that simulates human-robot teams comprising autonomous and teleoperated robots. We conclude in Section 6 with a discussion of the main results.

## 2 RELATED WORK

The proposed work of inferring prior knowledge and acting on it belongs to the general body of literature focusing on the role of humans in search and rescue missions, and intent estimation for human-robot teaming to find missing persons.

## 2.1 Role of human skill in search and rescue

In a missing person search, humans are expected to assimilate and utilize complex spatial information about maps and missing person behavior in limited time [53]. Within the mission itself, humans are expected to reliably identify victims [6], and navigate successfully in low-light environments with stairs. Humans are also expected to be situationally aware of the mission at hand, ask questions to reduce ambiguity [8], and quickly adapt to changing requirements [51]. Extracting and quantifying these aspects of human involvement for shared autonomy approaches where robots and humans work together is an open area of research [9, 39].

## 2.2 Inferring human intent and strategies towards human-robot teaming in search

Human intent inference plays an integral role in human-robot interaction [25, 43], enabling robot interaction strategies that can respond to human actions naturally in real time. Intent inference has been possible by modeling human dynamics as Markov processes [29], Gaussian processes [65], and neural networks [63]. Many of the underlying models rely on movement or behavioral cues as observations to predict human actions [62]. Within the search and rescue applications, human intent or actions are generally measured in terms of where they may search next. For example, Ognibene et al. [46] has proposed an active intention recognition paradigm that enables robots to perceive responders' intentions, thereby enhancing joint exploration strategies. They use Monte-Carlo Tree Search algorithms in partially observable environments to predict the movements of human responders and identify critical areas for exploration. In terms of predicting rescuer's actions and goals, Guo et al. use transfer learning and attention-based Long-short-term memory (LSTM) networks in urban search and rescue missions to predict navigation and triage strategies [19]. Heintzman et al. integrate predictive models of lost person behavior, anticipated human searcher trajectories, and UAV sensor data using a Gaussian process model to optimize search paths that significantly lower target location uncertainty compared to lawn mower sweeps [24]. The mutual information based strategy for independent search used in this work is similar to the approach in [24] in the sense that maximizing mutual information amounts to minimizing uncertainty. A key distinction however is the use of particle filters which relaxes the Gaussian assumption on target location making it possible to search within obstacle-ridden environments that can give rise to multi-modal distributions. Recent modeling efforts also focus on modeling and inferring human trust in a mobile robot team using dynamic Bayesian networks as they perform search [15] or area coverage [68] tasks. In [68], measured trust is used to either alter the number of interventions or increase the coverage area assigned to individual robot.

## 3 PRELIMINARIES

### 3.1 Mutual information based control

Information theory provides a meaningful objective for robot search strategies in the form of mutual information [26, 28]. For example, maximizing information gain, as opposed to reaching a goal location, provides a task-invariant objective for missions that can differ greatly in terms of environments and target locations. To understand mutual information, we briefly introduce information entropy, defined as the amount of uncertainty within a random variable  $X \in \mathcal{X}$  and is calculated as  $H(X) = -\sum_x p(X = x) \log p(X = x)$ . Correspondingly, joint entropy  $H(X, Y) = -\sum_{x,y} p(x, y) \log p(x, y)$  and conditional entropy is  $H(X|Y) = -\sum_{x,y} p(x, y) \log p(x|y)$ . Mutual information quantifies the information shared between two random variables and is defined as  $I(X; Y) = H(X) + H(Y) - H(X, Y) = H(X) - H(X|Y)$ .

In our context, the random variable  $\theta_k \in \Theta \subset \mathbb{R}^2$  represents the unknown target location at time step  $k$ , and  $Z_k \in \mathcal{Z}$  represents measurements taken by the robot such as range and bearing.

Denoting predicted location of the target before the next measurement is taken as  $\theta_k^-$  and possible target measurement as  $Z_k^-$ , the mutual information  $I(\theta_k^-; Z_k^-)$  represents the information gained about the target by taking that measurement [32]. Because the robot has different options in terms of where to take the next measurement at any time step, maximizing this mutual information serves as a meaningful, trajectory independent, way to achieve the objective of finding the target. Specifically, the control input at each time step is determined as [26]

$$\mathbf{u}_k = \max_{\mathbf{u} \in U} I(\theta_k^-; Z_k^-), \quad (1)$$

where  $U$  is the range of control inputs to choose from. In our case, for example,  $U$  represents the combination of speed and turn rate for moving a differentially driven ground robot.

Implementing a mutual information based control in an obstacle rich environment involves calculating non-Gaussian joint probability distributions such as  $p(\theta_k^-, Z_k^-)$  at every time step. A particle filter, which is a sequential Monte Carlo technique, where the probability density functions are represented as discrete distributions of instances of state called particles and associated weights, is best suited for handling nonlinear, non-Gaussian estimation problems [1]. Specifically, denoting the state  $\mathbf{X}_k \in \mathbb{R}^3$ , of a robot moving in two dimensions in terms of position  $(x_k, y_k)$  and orientation  $\psi_k$ , the probability of  $\tilde{\mathbf{X}}_k$  in a particle filter can be written as  $p(\tilde{\mathbf{X}}_k) = \sum_{q=1}^{N_p} w_q \delta(\mathbf{X}_{k,q} - \tilde{\mathbf{X}}_k)$ , where  $N_p$  is the number of particles,  $w_q$  is the weight of the  $q$ -th particle and  $\delta$  denotes the Dirac delta function. The weights are updated at each step using a likelihood function, which relates measurements to state as  $w_{k,q} = p(Z_k | \mathbf{X}_{k,q})$ , where  $\mathbf{X}_{k,q}$  denotes the  $q$ -th particle for the estimate at time-step  $k$ .

### 3.2 Weighted strategy for robot assistance during search

In a search scenario where both human teleoperated and autonomous robots are involved, human prior knowledges serves as a resource that can be utilized to improve performance. This is explored in [32] where a weighted strategy is used to control autonomous robots who search along with a reference robot modeled to represent human search. The reference robot uses (1) to search a target with some prior knowledge about its location. Simulations are performed as the prior knowledge accuracy is varied by changing the center of a Gaussian probability distribution representing the believed target location across multiple scenarios. This work proposes a information blending strategy where the normalized mutual information with respect to the predicted locations of the teleoperated robot and the target are blended in varying proportions to see its effect on search performance. In doing so, this approach assumes that a) the human teleroperator searches in regions proximal to it and b) searching close to the teleoperated robot is equivalent to maximizing relevant information content that may be implicitly available to the teleoperator. Furthermore, maximizing mutual information with respect to teleoperated robot location serves as a proxy to take a measurement near it.

The information blending approach suggests picking a control input for autonomous robots that maximizes a weighted normalized mutual information based objective [32]

$$\mathbf{u}_k = \max_{\mathbf{u} \in U} \left[ \alpha \widehat{I}(\theta_k^-; Z_k^{\theta,-}) + (1 - \alpha) \widehat{I}(\mathbf{X}_k^{\hbar,-}; Z_k^{\hbar,-}) \right], \quad (2)$$

where  $Z_k^{\theta,-}$  denotes the predicted target measurement,  $\mathbf{X}_k^{\hbar,-}$  denotes teleoperated robot's predicted location, and  $Z_k^{\hbar,-}$  denotes the predicted teleoperated robot measurement obtained upon applying control input  $\mathbf{u}$ ; the quantity  $\widehat{I}(\theta_k^-; Z_k^{i,\theta,-}) = I(\theta_k^-; Z_k^{i,\theta,-}) / H(Z_k^{i,\theta,-})$ , for example, represents the mutual information between the predicted target measurement and predicted target location and measurement normalized with respect to the entropy of the target measurement. This normalization

ensures that different entropy values of target and teleoperated robot measurements, due to possibly different dynamics, do not bias the mutual information comparisons;  $\alpha \in [0, 1]$  is the weighting parameter, where  $\alpha = 1$  implies that the autonomous robots will search for the target independent of the teleoperated robot, and  $\alpha = 0$  implies that autonomous robots will search near the teleoperated robot. This is because with bearing measurements, minimizing the uncertainty with respect to a position amounts to coming close to it [26]. Note that measurements are obtained only if the target is within the sensor range and teleoperated robot is within the communication range. Furthermore, an information blending approach seeks to maximize information through two complementary sources rather than weigh between two opposing strategies. In other words, one may search independently near the teleoperated robot.

Because it is derived directly from mutual information, such a strategy provides a natural abstraction of human robot interaction in multi-robot settings. Realistic simulations in two different scenarios show that the search performance depends on the weighting parameter  $\alpha$ , and the accuracy of prior knowledge of the reference robot. Furthermore, an information-blending strategy performs better than mixed initiative strategy, where robots switch roles, and similar to linear blending of control inputs in scenarios where prior knowledge is inaccurate. Unlike the latter, however, where different control inputs with dissimilar effects are combined, information blending attempts to combine mutual information gain from different strategies thereby making it more amenable for linear blending. Building on these results, a natural next question is what would  $\alpha$  represent in an adaptive strategy? In this work, we argue that  $\alpha$  can be a measure of prior knowledge related to the mission. In particular, with  $\alpha$  as a measure of prior knowledge, the autonomous robot will be able to effectively weigh between exploration, or searching the environment independently and exploitation, which here implies contributing to the perceptual range of the teleoperated robot.

#### 4 EXPERIMENTS

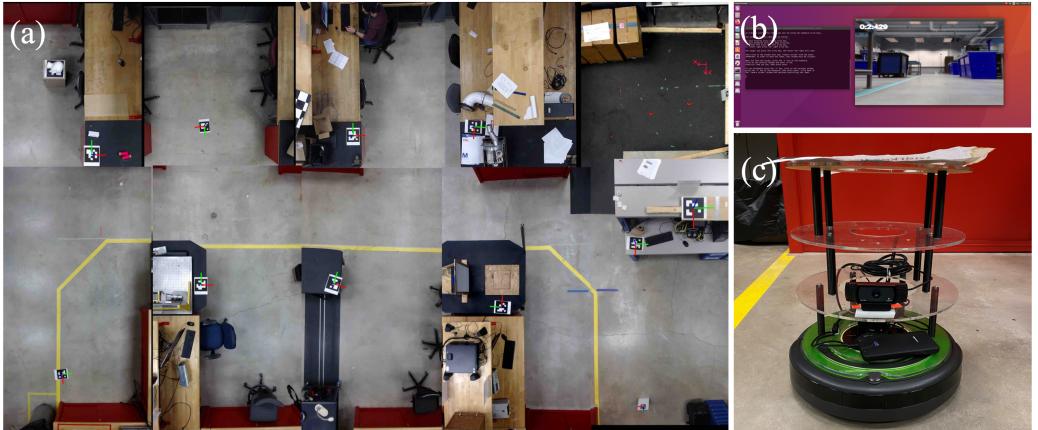


Fig. 1. Mosaic image (a) created using all ten overhead camera images of the search environment. Screen capture (b) of first person view as streamed from the camera on the robot to a participant in the control room, and the teleoperated robot (c) with an onboard camera. A fiducial marker placed on top of the robot allowed it to be tracked at all times.

In this section, we describe the experimental study performed to artificially create prior knowledge and study its effect on the movement of the teleoperated robot movement and physical and

mental workload of the teleoperator. By altering one of the experimental conditions we additionally investigate the effect of inaccuracy in prior knowledge on robot movement. The human-subject experiments focus on teleoperating a robot to find a missing target in a  $9 \times 18$  m room. The experiments are motivated by the following hypotheses: H1: The time to find a missing target depends on the knowledge of environment and target location; H2: The movement of the teleoperated robot depends on the participants' prior knowledge about environment and target location.

#### 4.1 Setup

The experimental setup consisted of a differentially driven ground robot that was teleoperated in an indoor search environment and tracked using overhead cameras. The ground robot (iRobot Create 2) had a webcam (Logitech, C920, 78 degrees horizontal field of view) attached using a serial connection to a microcomputer (Raspberry Pi 4) shown in Figure 1c. The microcomputer was powered with a 10000 mAh portable power bank. The indoor search environment was a 9 meter wide and 18 meter long robotics laboratory room situated within the Engineering academic building. The room consists of six large laboratory workbenches, several chairs, and tables that serve as obstacles (Figure 1). Ten webcams (Logitech C920) mounted on the ceiling 4 meters high were used to track the position and orientation of the robot, in real time, with a custom tracking system (programmed in Python with OpenCV), implemented on a dedicated tracking computer (Ubuntu Linux 18.04 operating system, 16 GB Memory, 3.4 GHz processor).

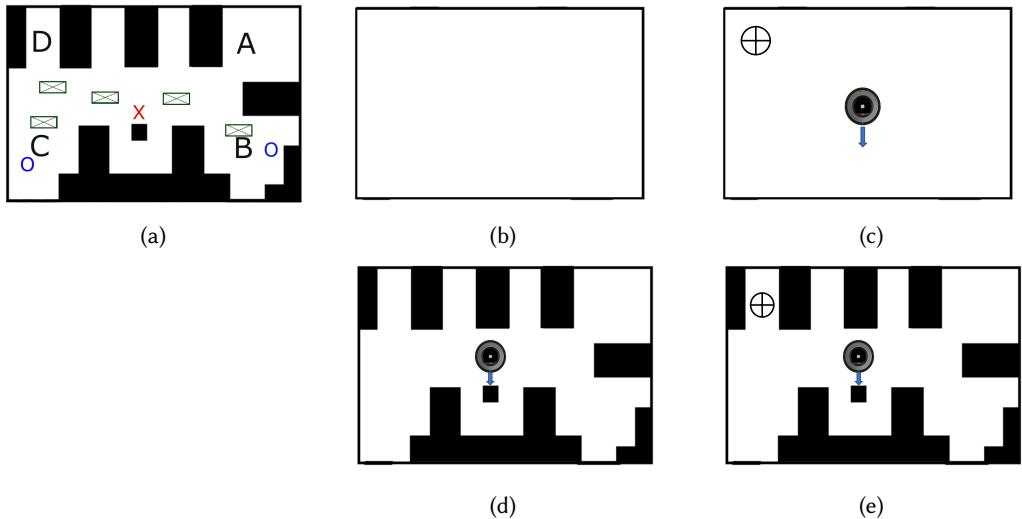


Fig. 2. Schematic of the search environment (a). The red 'X' is the robot location for conditions  $xMyT$ ,  $yMxT$ , and  $yMyT$ . For condition  $xMxT$  only, the blue 'O' marks one of two robot locations and the green rectangles denote the boxes placed within the lab. Sample maps provided to the participant during each condition (b-e), where (b) is for (No Map, No Target), (c) is for (No Map, Yes Target), (d) is for (Yes Map, No Target), and (e) is for (Yes Map, Yes Target). Whenever applicable, a high-probability region for target location was shown as a circle with cross-hair.

The robot was teleoperated from a desktop computer (Ubuntu Linux 16.04 operating system, 16 GB Memory, 3.4 GHz processor) that had a wide display monitor ( $2560 \times 1080$  pixels resolution) and was located in a smaller ( $4 \text{ m} \times 6 \text{ m}$ ) control room adjacent to, but isolated from, the search

environment. The display provided a first person view from the camera onboard the robot, along with a timer placed at the top left corner (Fig. 1b). A WiFi router (ASUS, Nighthawk) placed in the control room and a receiver (TP-Link, Archer T3U Plus) plugged into the robot was used to establish a strong wireless connection between the desktop computer and the robot. Further details on teleoperation and overhead tracking are provided in the supplementary material.

#### 4.2 Design

The experimental conditions were designed to study the effects of prior knowledge of target location and environmental map on search behavior. In a repeated measures design, we let each participant find a missing target as their prior knowledge of map and environment was varied. To further investigate if inaccuracy in prior knowledge of target location played a role in search behavior, one of the conditions was altered to have the target located in a place different from that indicated on the map provided prior to the trial. The following experimental conditions were considered (Fig. 2):

xMxT or No Map, No Target, where the participant was provided neither a map of the environment, nor a target location. The robot was initially placed at one of the corners facing the wall in the room (B or C in Fig. 2a), with the target placed on one of the corners on the far side of the environment. To further ensure that the participant did not associate any environmental features with the layout in subsequent conditions, several  $0.3 \times 0.75 \times 1$  m sized boxes were placed at random locations within the environment.

xMyT or No Map, Yes Target, where the participant was given a map with only the outer boundaries of the environment, initial robot position and orientation, and target location with a “high probability”. The target was always placed in the high probability location.

yMxT or Yes Map, No Target, where the participant was given a fully detailed map of the environment with the location and orientation of the robot, but the target location was not provided.

yMyT or Yes Map, Yes Target, where the participant was given a fully detailed map of the environment with the location and orientation of the robot, and a “high probability” target location. The target location could be accurate ( $yMyT_A$ ) or inaccurate ( $yMyT_I$ ) in which case the target was placed in a location different from the high-probability region.  $yMyT_I$  was tested for about half of the participants.

Participants for this study ( $N = 30$ , 5 Females,  $22 \pm 3.3$  years old), were recruited through flyers posted throughout the campus and through email announcements. Data from one participant had to be excluded owing to the robot battery getting discharged in the middle of a trial, leaving a total of  $N=29$  participants. The experiment was approved by the Institutional Review Board at Northern Illinois University under protocol # HS21-0372. Exclusion criteria included that participants be above 18 years of age, had not visited the robotics engineering laboratory in the past year, and had not been part of any training in that room before that for more than six hours. The last two exclusion criteria ensured that participants were not familiar with the laboratory layout and could therefore not possess prior knowledge of the environment.

#### 4.3 Procedure

Prior to arrival, each participant was assigned a random identifier to record all information anonymously. Upon arrival, participants were asked to go through the consent form and sign if they wish to proceed. They were then shown the robot they will be teleoperating, how it will be communicating with the computer through a microcomputer, and that they will see a first person view from a webcam attached on the robot. They were also shown the fiducial marker on top of the robot which will assist in tracking their movements through the environment.

Next, participants were trained to operate the robot while in the control room. A training time of 2-5 minutes was allocated to each participant during which they were encouraged to maneuver the robot from within the control room steering it past objects and also leave the room into the hallway if they desired. A participant was free to end the training at any time after two minutes when they felt comfortable. The goal of the training phase was to ensure that any changes in robot navigation observed during the experiment were only because of knowledge of the unknown environment or target location, and not from operating skill.

Before starting the experiment the participant was shown the target (a stuffed toy) and was informed that they will have to find that target located in a different room as soon as possible. A stopwatch was displayed on the left corner of the first person view (Fig. 1b) to keep participants motivated to find the target as soon as possible. The target was placed in a random location (A,B,C or D in Fig. 2) for every trial except the first, when the target was only placed on B or C. Once found, participants were instructed to enter a randomly generated 4-digit numeric code, different for each trial, on the target to successfully terminate the search. Each participant was asked to find the target four times, once for each condition, in the order xMxT, xMyT, yMxT, yMyT. We selected this order, as opposed to counterbalancing conditions, to avoid participants having knowledge of the map before they were supposed to, such as for example if yMyT were to occur before xMyT. In each condition, once the target was found, a NASA TLX questionnaire [5] was prompted on the screen posing questions related to workload related to the task.

After all conditions within the experiment were completed, participants were requested to fill a post-experiment questionnaire regarding the levels of trust they placed in the experimenter when told about the high probability target location, prior teleoperating and video gaming experience, and degree of delay they experienced in operating the robot. The first two questions in the post-experiment questionnaire were asked to determine the role that trust in experimenter played in regulating target prior knowledge. In other words, a high trust in the experimenter would ensure that the participants believed that the target was located where the experimenter told them, and any change in behavior during the yMyT<sub>I</sub> condition would be due to ultimately realizing that the location pointed out on the map was inaccurate. Finally, participants were compensated \$10 for their time and requested not to share the knowledge of the environment with others.

#### 4.4 Teleoperation viability

Table 1. Comparison of post-experiment questionnaire responses. Question responses were scaled between 1-7, where for questions 1 and 2, 1 represented low, and 7 represent high; for questions 3 and 4, 1 represented rare, and 7 represented frequent; and for question 5, 1 represented high latency and 7 represented low latency. For reference, the second trial was No Map, Yes Target (xMyT) and the fourth trial was Yes Map, Yes Target (yMyT).

Question	Value
How would you describe your trust in the experimenter the first time (second trial) ...	6.3103 ± 1.3655
How would you describe your trust in the experimenter the second time (fourth trial) ...	6.2759 ± 1.4367
How would you describe your past experience with teleoperation	2.7931 ± 1.897
How would you describe your past experience with video games	5.3103 ± 1.9839
How would you rate the latency of the video stream ...	4.2414 ± 1.3537

Table I shows average responses to post experimental survey asking participants to rate their trust in the experimenter as they indicated the “high probability” target locations, prior experience with teleoperation and video games, and latency of the video stream. Participants indicated high

trust levels in the accuracy of the target locations indicated, had little experience with teleoperation, albeit much more experience with video games, and considered the video stream to be low latency.

#### 4.5 Data Analysis

Because condition yMyT had half of the trials altered to put the target at a different location than that indicated on the map, data from the corresponding condition, yMyT<sub>I</sub>, was temporally split and included with yMyT<sub>A</sub> towards increased sample size. Specifically, the trajectory in yMyT<sub>I</sub> was split temporally into yMyT<sub>I<sub>1</sub></sub> and yMyT<sub>I<sub>2</sub></sub> to mark the time when the participant first reached the target location indicated on the map (see supplementary material). The first part, yMyT<sub>I<sub>1</sub></sub>, of the trajectories was then combined with yMyT<sub>A</sub> for increased sample size. We computed the following measures based on trajectory data to analyze the differences between conditions:

- (1) *Performance* was calculated in terms of time to find the target and total distance travelled by the robot. To compensate for the effect of obstacles in xMxT, we teleoperated the robot five times across the environment from different locations with and without obstacles. This gave us a factor of 1.135 that denoted the additional time spent navigating the environment due to obstacles. Distance travelled and time to find the target in xMxT were divided by this scaling factor to ensure a meaningful comparison with the remaining conditions (see Supplementary material).
- (2) *Robot speed and turn rate* were calculated to measure teleoperation efficiency [11] and motivate the selection of movement cues. These were calculated from filtered trajectory data  $\hat{x}_k$ ,  $\hat{y}_k$ , and  $\hat{\psi}_k$  as

$$\begin{aligned}\hat{v}_k &= \frac{1}{\Delta t} \sqrt{(\hat{x}_k - \hat{x}_{k-1})^2 + (\hat{y}_k - \hat{y}_{k-1})^2}, \\ \hat{\omega}_k &= \frac{\Delta \hat{\psi}_k}{\Delta t}\end{aligned}\quad (3)$$

where  $\hat{v}_k$ ,  $\hat{\omega}_k$  denote the robot linear and angular speeds calculated as the difference between estimated position and orientation on successive time steps with a  $\Delta t=0.5s$ .

- (3) *Situational awareness* seeking behaviors were quantified in terms of tendency to stay in one place, turn in place [57], stay still, and frequency of stops. Specifically, we calculated the number of time steps spent staying in place when the estimated robot speed  $\hat{v}_k \leq 0.1$  m/s. This was further divided into time spent turning in place calculated as the number of time steps when the robot estimated speed  $\hat{v}_k \leq 0.1$  m/s and robot estimated turn rate  $\hat{\omega}_k \geq 0.1$  rad/s; time spent staying still (freezing) was calculated as  $\hat{v}_k \leq 0.1$  m/s and robot estimated turn rate  $\hat{\omega}_k \leq 0.1$ . We expected participants without any prior knowledge would stay in place, turn, or simply stay still more than when they knew the environment layout or target location.

Statistical analyses were performed using Friedman non-parametric repeated measures tests after verifying non-normality of data. In each test, prior knowledge about map or target location was the independent variable and the corresponding measure (e.g. time to find, speed, etc.) as the dependent variable. Significance was noted for  $p$  value less than 0.05. Posthoc pair-wise comparisons were performed with Bonferroni correction. All analyses were performed in MATLAB.

#### 4.6 Experimental results

*Existence of prior knowledge affected mental and physical workload and effort.* Figure 3 shows results of non-parametric repeated measures Friedman tests with Bonferroni correction for questions on the NASA-TLX survey taken after every trial. We find that the existence of prior knowledge

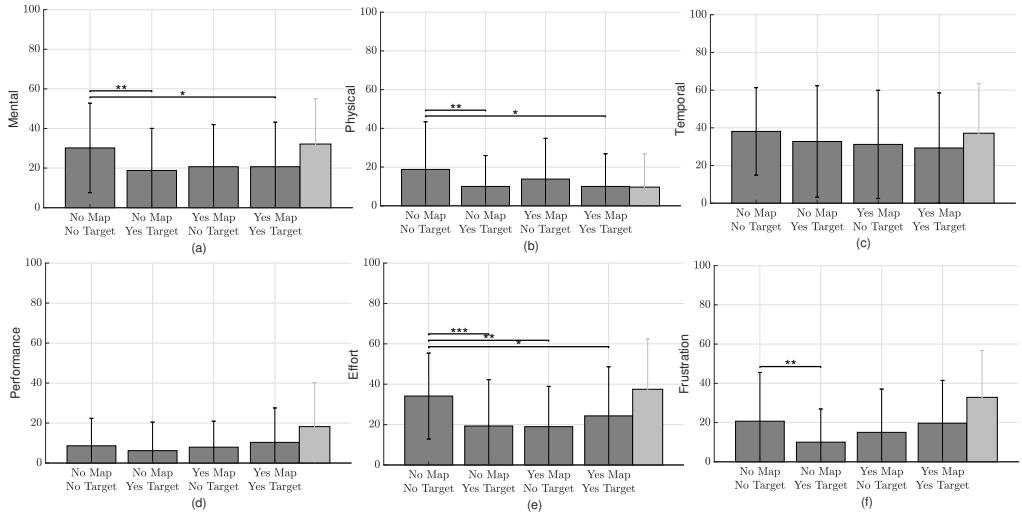


Fig. 3. Responses to questions in the NASA-TLX survey for all participants. Question responses were scaled between 0-100, where 0 represents Low and 100 represents high. For question 4 (Performance), 0 represents perfect and 100 represents failure. For (Yes Map, Yes Target) darker bars denote responses for all participants (yMyT) and lighter bars denote the responses for those participants only whose prior knowledge of target location was inaccurate (yMyT<sub>I</sub>). Statistical comparisons are only shown for darker bars. Overhead bars indicate pairwise difference  $p < 0.05^*$ ,  $p < 0.01^{**}$ ,  $p < 0.01^{***}$  if significant difference was found in Friedman non-parametric repeated measures test.

had a significant effect on responses to five out of six questions within the survey, that is across all questions except performance. For all participants, post-hoc comparisons for mental demand revealed that all participants found (No Map, No Target) more mentally demanding than (No Map, Yes Target) and (Yes Map, Yes Target). The same was observed in the case of physical demand. In the case of effort, post-hoc comparisons revealed that participants who did not have any knowledge (No Map, No Target) felt that they put more effort in the search of the missing target in the environment than all other conditions (No Map, Yes Target; Yes Map, No Target; Yes Map, Yes Target). Finally, participants felt more frustrated when they did not have any knowledge than when they had knowledge of target location only (No Map, Yes Target). Wilcoxon rank sum test between participants who had different types of target knowledge (15 accurate samples versus 14 inaccurate samples with an expected rank  $W = 225$  if values are not significantly different) in the yMyT condition revealed that participants who had inaccurate knowledge of the target (yMyT<sub>I</sub>, light colored bars in Figure 3) felt that the task was more mentally demanding ( $W = 170, p = 0.014$ ), required more effort ( $W = 163, p = 0.007$ ) and caused more frustration ( $W = 155, p = 0.002$ ) than when they had accurate knowledge (yMyT<sub>A</sub>, not shown explicitly). When only considering the subset of participants who only had accurate knowledge in (Yes Map, Yes Target), that is yMyT<sub>A</sub>, we still found a significant effect on five out of six questions, albeit these comprised a different set (not shown here).

*Qualitative analysis of trajectory data.* Figure 4 shows trajectory data, and robot speeds and turn rates from each condition of a randomly selected participant. A visual inspection reveals that the participant spent much of the time exploring the environment and looking for the target in (No Map, No Target). However, looking at (No Map, Yes Target) (second row from top), when the

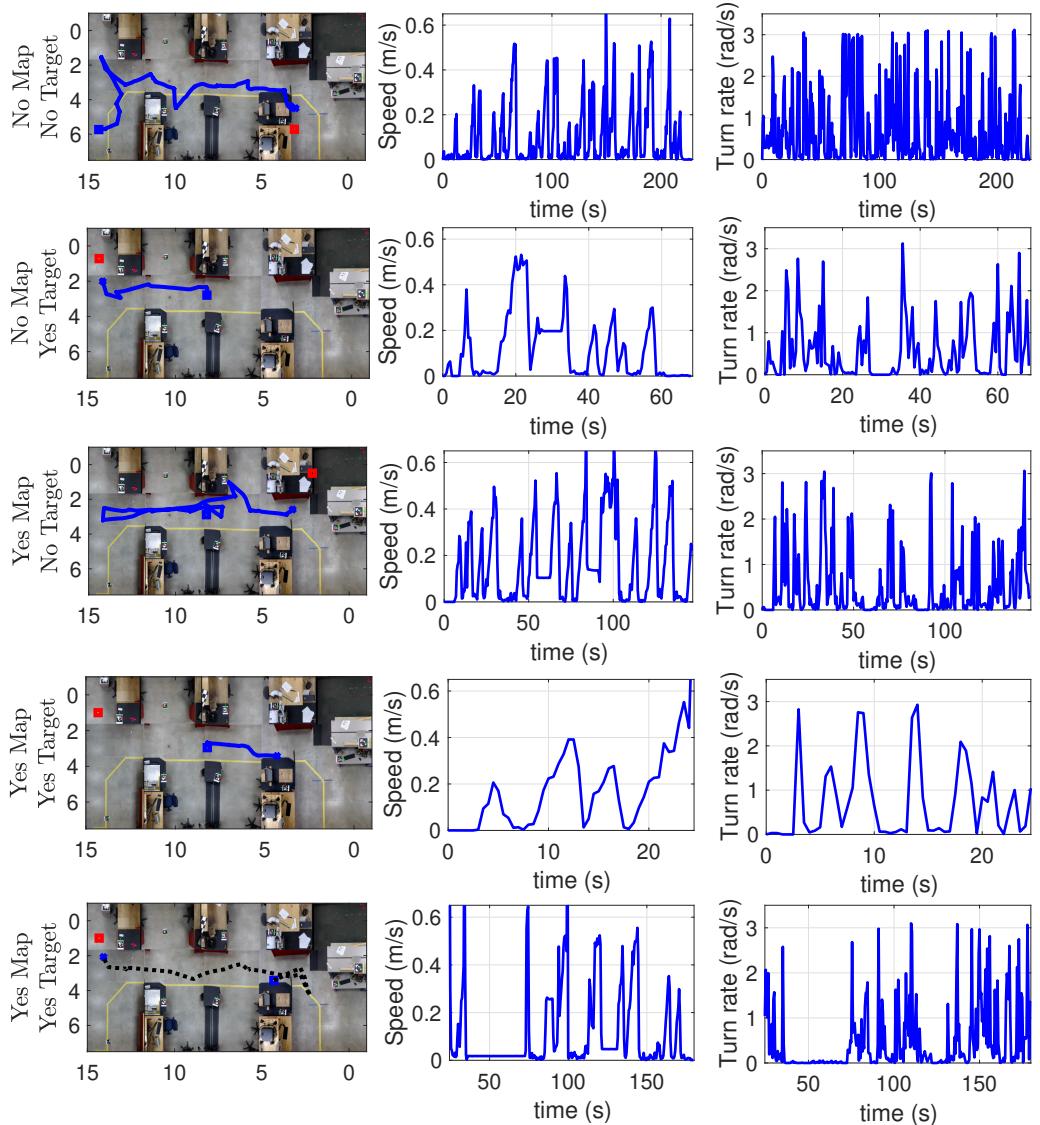


Fig. 4. Robot position, speed, and turn rate for a single participant who was informed of a target location different from actual one in condition 4. The first column shows the robot trajectory within the search environment and columns 2 and 3 show robot speed and turn rate over time. Rows 1-3 denote conditions xMxT, xMyT, and yMxT, and rows 4 and 5 denote trajectory data before  $yMyT_{I_1}$  and after  $yMyT_{I_2}$  the split for corresponding to the moment when the participant realizes that the target location knowledge was inaccurate.

participant knew target location only, the trajectory is directed towards the true target location accompanied by fewer turns. These observations are corroborated when comparing the trajectories of all participants in each condition (Fig. 5). Specifically, participants had fewer turns when they

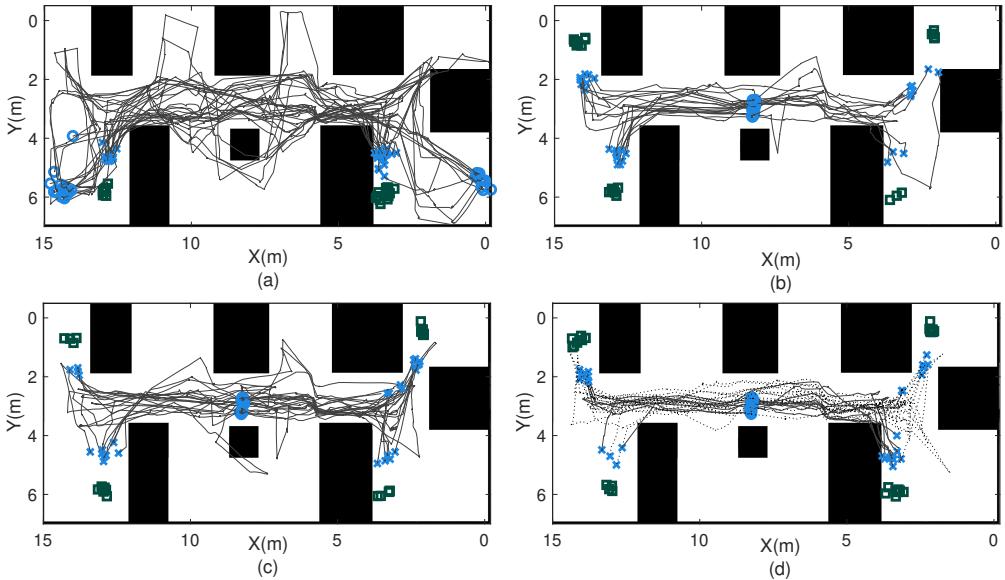


Fig. 5. Robot trajectories for all participants in conditions 1-4 with (a) for (No Map, No Target), (b) for (No Map, Yes Target), (c) for (Yes Map, No Target), and (d) for (Yes Map, Yes Target). In (d), dotted trajectories denote when participants did not have an accurate knowledge ( $y_{MyT_I}$ ) of the target location. Target locations are shown as green squares, and trajectory start and end locations are shown as blue circles and blue crosses.

knew the target location (including an inaccurate location) than when they only knew the map. (See supplementary material for sample overhead videos for each condition.) Robot speed for when target location was known (Fig. 4, rows 2 and 4) exhibits frequent delays in movement with longer delays in (No Map, Yes Target) compared to (Yes Map, Yes Target) prior to realizing that the target was not located where it was indicated. Furthermore, the number of instances where the robot start-and-stopped appear to be more in xMxT compared to xMyT, and in xMxT compared to yMyT; in each case the participants did not have prior knowledge about the target location compared to the next condition. These visual analyses set the stage for statistical comparisons for testing our hypotheses next.

*Search performance depends on prior knowledge.* Expectedly, the level of prior knowledge had significant effect on time to find the target ( $\chi^2(3, 29) = 63.08, p < 0.001$ , Fig. 6a). Post-hoc pair-wise comparisons revealed that the participants without prior knowledge of map and target (No Map, No Target) took longer time (scaled) to find the target than when they knew either the location of the target, or the map, or both (No Map, Yes Target; Yes Map, No Target; Yes Map, Yes Target). Furthermore, knowing target location further reduced the time to find if map was known (Yes Map, No Target, versus Yes Map, Yes Target). The distance the robot traveled across all conditions also revealed a significant effect ( $\chi^2(3, 29) = 48.60, p < 0.001$ , Fig. 6b). Post-hoc pairwise comparisons revealed that the participants who did not have any knowledge (No Map, No Target) traveled more distance than when they knew the target location or map or both (No Map, Yes Target; Yes Map, No Target; Yes Map, Yes Target). We also see that knowing target location and map (Yes Map, Yes Target) led to less distance travelled than when only map was known (Yes Map, No Target).

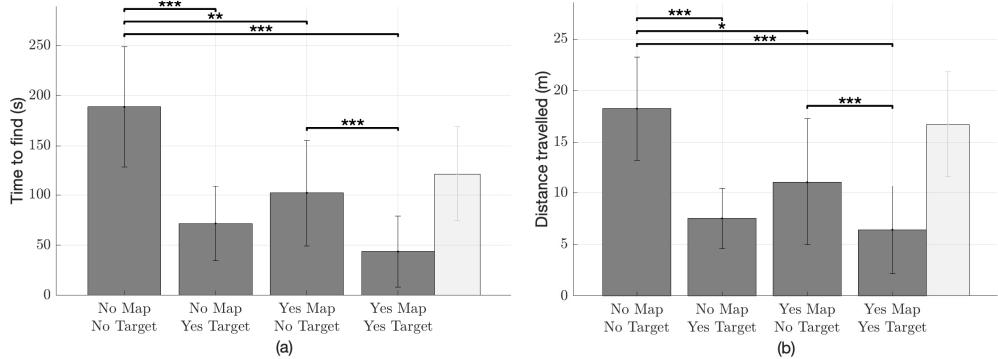


Fig. 6. Data showing time to find the target in seconds (a) and total distance travelled by the teleoperated robot in meters (b). Time to find and distance travelled for xMxT (No Map No Target) are scaled by a factor of 1.135 to account for the presence of obstacles in the environment. Data is shown for the four different conditions corresponding to the four different types of prior knowledge a participant was given prior to the experiment. For (Yes Map, Yes Target) darker bars denote data from yMyT<sub>A</sub> and yMyT<sub>I<sub>1</sub></sub>, and lighter bars denote the data from yMyT<sub>I<sub>2</sub></sub>. Statistical comparisons are only shown for darker bars. Overhead bars indicate pairwise difference  $p < 0.05^*$ ,  $p < 0.01^{**}$ ,  $p < 0.001^{***}$  if significant difference was found in Friedman non-parametric repeated measures test.

*Search behavior depends on prior knowledge.* With respect to fraction of time staying in place we found significant effect of prior knowledge ( $\chi^2(3, 29) = 10.77, p = 0.013$ ) with post-hoc comparisons revealing that the robot stayed more in place when no prior knowledge was present (No Map, No Target) than when both target location and map was known (Yes Map, Yes Target). Looking at different components of the time spent in place, Figures 7a and b compare the fraction of time spent turning in place, and fraction of the time spent staying still (freezing). Whereas the fraction of time spent turning in place produced no significant effect ( $\chi^2(3, 29) = 2.669, p = 0.445$ , Fig. 7a), freezing revealed a significant effect ( $\chi^2(3, 29) = 18.60, p < 0.001$ , Fig. 7b). Post-hoc pairwise comparisons revealed participants spent more time freezing when map knowledge was absent (No Map, No Target; No Map, Yes Target) than when both target and map knowledge was present (Yes Map, Yes Target). Frequency of stops made during a trial were not found to significantly depend on the prior knowledge ( $\chi^2(3, 29) = 6.558, p = 0.087$ ).

Teleoperated robot speed significantly depended on prior knowledge ( $\chi^2(3, 29) = 22.77, p < 0.001$ , Fig. 7c). Post-hoc pair-wise comparisons revealed that when both target location and map was known (Yes Map, Yes Target), the robots were faster than when only map was known (Yes Map, No Target) or when neither map nor target location was known (No Map, No Target). Having only target location knowledge (No Map, Yes Target) resulted in a higher speed than when neither map nor target knowledge was known (No Map, No Target). Robot turn rate also registered a significant effect ( $\chi^2(3, 29) = 10.613, p = 0.014$ , Fig. 7d). Post-hoc pair-wise comparisons revealed that when both target location and map was known (Yes Map, Yes Target), teleoperated robots had a higher turn rate than when only target location was known (No Map, Yes Target).

*Inaccuracy in prior knowledge affected some search behavior.* The data to investigate the effect of inaccuracy of prior knowledge were drawn from yMxT, yMyT<sub>A</sub> combined with yMyT<sub>I<sub>1</sub></sub>, and yMyT<sub>I<sub>2</sub></sub> to create three independent factors: No target, Accurate target, and inaccurate target. For comparison, note that the lighter bars in Figure 7 denote yMyT<sub>I<sub>2</sub></sub>. Friedman repeated measures

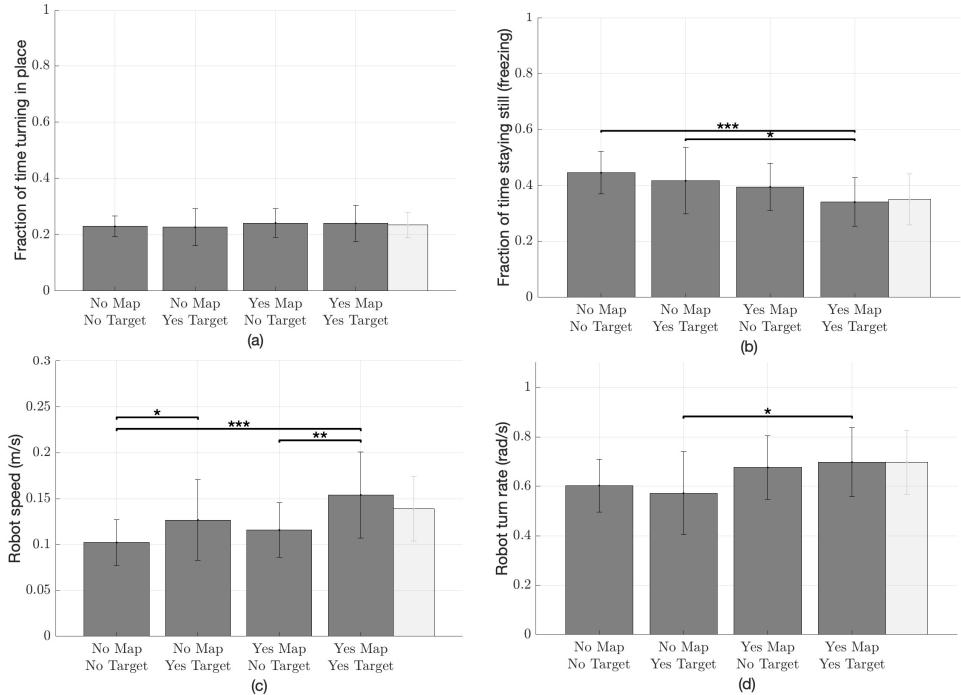


Fig. 7. Data showing fraction of time the teleoperated robot spent turning in place (a), and fraction of time the robot spent staying still (b), robot speed (c) and absolute robot turn rate (d). Overhead bars indicate pairwise difference. Data is shown for the four different conditions corresponding to the four different types of prior knowledge a participant was given prior to the experiment. For (Yes Map, Yes Target) darker bars denote data from yMyT<sub>A</sub> and yMyT<sub>I<sub>1</sub></sub>, and lighter bars denote the data from yMyT<sub>I<sub>2</sub></sub>. Statistical comparisons are only shown for darker bars. Overhead bars indicate pairwise difference  $p < 0.05^*$ ,  $p < 0.01^{**}$ ,  $p < 0.01^{***}$  if significant difference was found in Friedman non-parametric repeated measures test.

with accuracy of prior knowledge of target location as the independent factor revealed significant effect of prior knowledge accuracy on robot speed ( $\chi^2(2, 14) = 9.571, p = 0.008$ ) but not on turn rate ( $\chi^2(2, 14) = 1.00, p = 0.606$ ). Post-hoc comparisons put robot speed with accurate knowledge to be higher than when target location was not known (Accurate target versus No target). Although staying in place was found to be dependent on the accuracy of prior knowledge ( $\chi^2(2, 14) = 7.00, p = 0.030$ ), neither fraction of time turning in place ( $\chi^2(2, 14) = 3.00, p = 0.223$ ), nor fraction time staying still ( $\chi^2(2, 14) = 1.85, p = 0.395$ ) was found to significantly depend on the accuracy of prior knowledge about target location.

The experimental results can be summarized as follows:

- Participants experienced more mental and physical workload when no prior knowledge was available compared to when at least target knowledge was available. Participants also felt that they exerted significantly more effort when no prior knowledge was available than when any knowledge (Map or Target) was provided.
- Performance in terms of time to find the target was significantly affected by prior knowledge (H1) with map or target knowledge significantly reducing the time to find compared to when no such knowledge was available.

- Search behavior was also significantly affected by prior knowledge (H2) with participants spending significantly less time staying in a place with no movement when they knew both the target location and the map than when they didn't know the map.
- The robot was teleoperated at a high speed and turn rate when prior knowledge was available (H2). Speed was significantly lower when target location was unknown and turn rate was significantly lower when map was unknown.
- While target location played a role in the speed of the teleoperated robot, it was not affected when the participants realized that the target location may be inaccurate.

## 5 PRIOR KNOWLEDGE BASED CONTROL OF AUTONOMOUS ROBOTS

In this section we describe the prior knowledge inference from movement cues and perform simulations to evaluate control strategies that utilize such inference. First, building on the experimental results, we identify movement cues that can be used to classify prior knowledge about target location and environment. Next, we present a Bayesian inference strategy to directly calculate the weighting parameter within the mutual information objective function that adaptively blends two complementary search strategies for autonomous robots who are searching for the same target as the teleoperated robot. We finally use this framework to test hypothesis H3: a human-robot team, where autonomous robots apply an adaptive strategy that involves inferring and acting on prior knowledge performs better than other strategies where such inference is not used.

### 5.1 Inferring prior knowledge from movement cues

The dependence of various movement cues on prior knowledge motivates the possibility of directly inferring such knowledge from teleoperated robot motion. To make such inference practical for use in a control strategy, however, such inference should be possible from a smaller observation window. Therefore, we sought to quantify the difference between distributions of a movement cue (feature) conditioned on the type of prior knowledge. Specifically, denoting a feature observed for past  $\tau$  time steps by  $f_\tau$ , we used experimental data to calculate  $p(f_\tau|\mathcal{K})$  where  $\mathcal{K}$  denotes the prior knowledge state. In our case, prior knowledge is represented by a random variable that can take values in the sample space  $\{xMxT, xMyT, yMxT, yMyT\}$ . We limited the observation window  $\tau$  to range between 5 and 30 seconds to allow a usable strategy to develop within reasonable time as the teleoperated robot is observed. When calculating these distributions, data from the last 5 seconds of the trial was ignored to avoid considering behaviors after the target was likely spotted by the participant. Motivated by experimental results, where we see dependence of distance, speed, turn rate, and freezing time on prior knowledge, probability distributions were built using features such as average speed, distance travelled, and freezing, calculated on  $\tau$ -second sections within a moving window ending at the current time step.

To achieve a robust classification of prior knowledge, only those features were selected for inference that had maximum difference in probability distributions conditioned on the two extreme scenarios  $xMxT$ , and  $yMyT$ , in terms of the Wasserstein distance. Although we calculate probability of all four levels of prior knowledge, we selected the two extreme scenarios to maximize the distinguishability between having no prior knowledge and having all available prior knowledge. The Wasserstein distance, also known as the earth-mover's distance, is a distance metric between probability distributions that calculates the cost of turning one distribution, viewed as a pile of sand, into another [48]. Figure 8 shows that we attain the maximum Wasserstein distance between  $p(f_\tau|\mathcal{K} = xMxT)$  and  $p(f_\tau|\mathcal{K} = yMyT)$  for freezing and distance covered, with an observation time of 30 and 15 seconds respectively. We therefore select these two features as movement cues to infer prior knowledge. To investigate if combining features provided better performance, we additionally

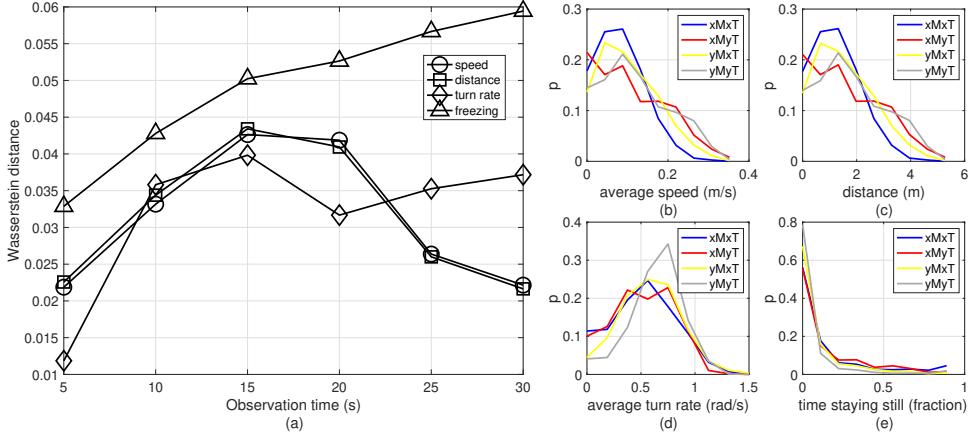


Fig. 8. Wasserstein distance between conditional probability distributions (a) of having complete or no prior knowledge for four different movement features as a function of observation time. Sample distributions for all four types of prior knowledge considered in this study for 15-second observation time (b)–(e).

select a combined two-dimensional feature set of freezing and distance covered with an observation time of 15 seconds.

A natural choice to set the weighting parameter would be the probability  $p(\mathcal{K} = \text{xMxT}|f_\tau)$  of not having prior knowledge of target location or environment conditioned on observations of the feature  $f_\tau$ . This would ensure that the parameter (a) stays between 0 and 1 permitting a linear blending and (b) directly relates to the need for additional assistance in the form of increased perceptual range from the autonomous robot. In particular, if  $p(\mathcal{K} = \text{xMxT}|f_\tau)$  is low, implying that the teleoperator has good prior knowledge of target location or environmental map or both, it may be efficient to assist the teleoperated robot (low  $\alpha_k$ ) by staying close by and increase its field of view; conversely, a high value of  $p(\mathcal{K} = \text{xMxT}|f_\tau)$  implies that the human robot may be searching inefficiently and therefore an effective strategy would be to search independently (high  $\alpha_k$ ). Given the observation of a particular feature, this probability can be calculated using Bayes' rule as

$$p(\text{xMxT}|f_\tau) = \frac{p(f_\tau|\text{xMxT})p(\text{xMxT})}{\sum_i p(f_\tau|\mathcal{K}_i)p(\mathcal{K}_i)}, \quad (4)$$

where  $p(\text{xMxT}|f_\tau)$  is abbreviated from  $p(\mathcal{K} = \text{xMxT}|f_\tau)$ , and  $p(\mathcal{K}_i), i = 1, \dots, 4$  denotes the probability of prior knowledge being in one of the four possible states so that for example  $p(\mathcal{K}_1) = p(\mathcal{K} = \text{xMxT})$ .

For an adaptive control strategy, the parameter can be time-varying as  $\alpha_k$ , and can be set to  $p(\mathcal{K} = \text{xMxT}_k|f_\tau^k)$  with  $f_\tau^k$  denoting the collection of feature measurements all the way up to  $k$ ; the value of  $\alpha_k$  is then recursively updated as

$$\alpha_k = \frac{p(f_{\tau,k}|\text{xMxT}_k)\alpha_{k-1}}{\sum_i p(f_{\tau,k}|\mathcal{K}_{i,k})p(\mathcal{K}_{i,k-1})}, \quad (5)$$

where the normalization factor is calculated by estimating and adding the probability for each of the four states at each time step. The value  $p(\mathcal{K}_{i,k-1}), i = 1, \dots, 4$  are all set to  $1/4$  for  $k = 1, \dots, \tau$ , until enough observations are available to make an inference.

## 5.2 Controlling autonomous robots in response to teleoperated robot movement

Compared to a control strategy that presupposes prior knowledge, we now propose a mutual information based objective function that responds to varying prior knowledge over the course of the experiment. In particular  $\alpha_k$ , same for all autonomous robots, is computed at every time step and used to determine the control input that optimizes the blended mutual information objective

$$\mathbf{u}_k = \max_{\mathbf{u} \in U} \left[ \alpha_k \widehat{I}(\boldsymbol{\theta}_k^-; \mathbf{Z}_k^{\theta, -}) + (1 - \alpha_k) \widehat{I}(\mathbf{X}_k^{\hbar, -}; \mathbf{Z}_k^{\hbar, -}) \right], \quad (6)$$

where again, as before, the normalized mutual information is computed based on predicted measurements obtained by applying control input  $\mathbf{u}$ . To further assess the effect of number of robots, we conduct the simulation study with one and two autonomous robots.

## 5.3 Simulation setup

The simulation setup consisted of human-robot teams where one robot (teleoperated robot) followed a predetermined trajectory from the actual experiments as the autonomous robots adapt their search behavior according to the control law (6). The autonomous robots themselves were simulated to match the ground robot platform used in the experiments with a 360-degree field of view. Specifically, the robots were assigned differential drive dynamics, a bearing-only sensor with a limited range of 1.5 m (such as calibrated omnidirectional camera), and a collision handing algorithm that turned the robots away from the point of collision.

Specifically, the position,  $(x_k, y_k)$ , and orientation,  $\psi_k$ , of an autonomous robot  $i$  was updated at every time step  $k$  as

$$\begin{aligned} x_{k+1}^i &= x_k^i + v_{k+1}^i \cos \psi_k^i \Delta t \\ y_{k+1}^i &= y_k^i + v_{k+1}^i \sin \psi_k^i \Delta t \\ \psi_{k+1}^i &= \psi_k^i + \Omega_{k+1}^i \Delta t \end{aligned} \quad (7)$$

where  $\Omega_{k+1}^i \in \Psi$  denotes the turn rate and  $v_{k+1}^i \in V$  denotes the speed, with  $\Psi = [-0.25, 0.25]$  rad/s and  $V = [0, 0.833]$  m/s denoting the range of possible turn rates and speeds. At each step, solving (6) gave the value of  $\Omega_{k+1}^i$  and  $v_{k+1}^i$ . The length of the simulation time step  $\Delta t = 0.5$  second was set to match the observation sampling rate from the experiments.

Collision handling was performed so that upon a collision at an angle  $\gamma_k$  with respect to heading  $\psi_k$ , the robots changed their instantaneous turn rate  $\Omega_k$  to  $-1.5\gamma_k$ , thus turning them opposite to the angle at which a collision is detected.

Autonomous robots took target (if it was visible) and teleoperated robot location measurements using a bearing-only sensor with measurement noise set to a zero mean Gaussian random variable with 0.1 rad standard deviation. We note that, even though the teleoperated robot location was available directly from the overhead tracker (or by communication), it was converted into a bearing measurement to assign an information value to a measurement taken at the teleoperated robot location.

All autonomous robots ran a particle filter similar to the one in [32] with 1200 particles representing a combined state consisting of self pose  $\mathbf{X}_k^i = [x_k^i, y_k^i, \psi_k^i]$ , two dimensional target location  $\boldsymbol{\theta}_k$  and two dimensional teleoperated robot location  $\mathbf{X}_k^{\hbar}$ . State estimates were updated using likelihood functions for target  $p(Z_k^{\theta}|\mathbf{X}_k)$ , and teleoperated robot  $p(Z_k^{\hbar}|\mathbf{X}_k)$  measurement both of which were modeled as Gaussian density functions when a measurement was obtained (note the change in font of  $Z_k^{\hbar}, Z_k^{\theta}$  from bold to normal indicating that they are scalar for our setup). The bearing-only sensor obtained a measurement only if the target was within 1.5 meters visible range selected to be slightly more than the distance of 1.33 m at which the 4-digit numeric code on the target can be

identified by human participants. When a measurement was not obtained, the likelihood function for target measurement excluded all particles that lied within its sensor range thus pushing target estimates further into the unexplored region.

Autonomous robots shared their estimates of target and human teleoperated robot location through a combined likelihood function. Specifically, when there are  $n_r > 1$  autonomous robots, each autonomous robot  $i$  updates its weights as

$$p(Z^{i,\theta}, Z^{i,\hbar} | \mathbf{X}_k^i) = \prod_{j=1}^{n_r} \left[ p(Z_k^{j,\theta} | \mathbf{X}_k^j) \cdot p(Z_k^{j,\hbar} | \mathbf{X}_k^j) \right], \quad (8)$$

indicating that all robots are able to communicate with each other. This is a viable setup in an indoor environment even in the absence of an overhead tracker where robots are able to perform localization and mapping and communicate wirelessly through radio. In a larger GPS-denied environment, the number of robots that share measurements may be limited based on wireless communication range.

To evaluate the adaptive control strategy (6) we adopted a cross-validation approach where  $\alpha_k$  is computed from probability distributions generated from data from 90% of the participants and then used to control the autonomous robots in the remaining 10%. Human-robot teams with two and three robots (with one and two autonomous robots respectively) were simulated to search the same environment for the target.

Based on our analysis, we evaluated the robot assistance control strategy based on prior knowledge built on a) distance travelled with observation time  $\tau = 15$  seconds, b) freezing with observation time  $\tau = 30$  seconds, and c) distance travelled and freezing with observation time  $\tau = 15$  seconds. The target was considered found by an autonomous robot if it came within its visual sensor range. Because we were comparing similar setups, search time in condition xMxT was no longer divided by a scaling factor. Furthermore, to ensure that autonomous robots were not favored in terms of finding the target earlier than the teleoperated robot when the corresponding participant whose trajectory was being followed were simply entering the 4-digit numeric code, the simulation time (also the search time for the teleoperated robot) for comparison was reduced by five seconds. The performance of a control strategy was measured in terms of the average time to find the target, the fraction of simulations where the human-robot team found the target earlier than a single human, and the average time gained over those trials. The adaptive search strategies based on different movement cues were compared with three other strategies where the autonomous robots (i) search randomly, (ii) search independently ( $\alpha = 1$ ) and (iii) search always assisting the teleoperated robot ( $\alpha = 0$ ).

With respect to hypothesis H3, we expect that when the teleoperated robot is moving slowly (less distance travelled over the observation window of 15 seconds), then it would imply a high probability of not knowing where the target or map is so that  $\alpha_k = p(xMxT_k | f_\tau^k)$  close to 1; the accompanying autonomous robots will search for the missing target independently and sometimes further away from the teleoperated robot. Conversely, we expect that a fast moving teleoperated robot would imply a low  $\alpha_k$  close to 0 with the autonomous robots searching in a region close by. This should result in higher exploration when the teleoperated robot is slow and higher exploitation when the teleoperated robot is moving fast.

## 5.4 Simulation results

Figures 9 and 10 show sample simulation trajectories of one and two autonomous robots along with  $\alpha_k$  values inferred on the basis of distance travelled for each condition. In most instances,  $\alpha_k$  is dynamic suggesting that movement behaviors evolve over time for the same levels of prior

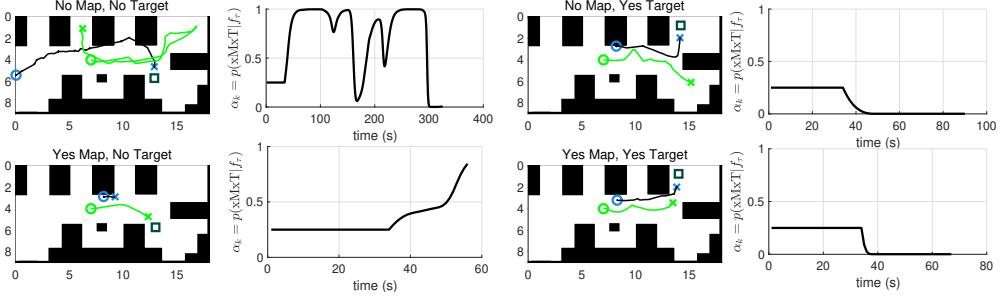


Fig. 9. Sample trials showing the teleoperated (black) and one autonomous robots (green) trajectories. Start and end positions are shown as circle and cross. Target location is shown as a square. The instantaneous  $\alpha_k$  values computed based on distance travelled by the teleoperated robot in the past 15 seconds are shown next to each trajectory plot.

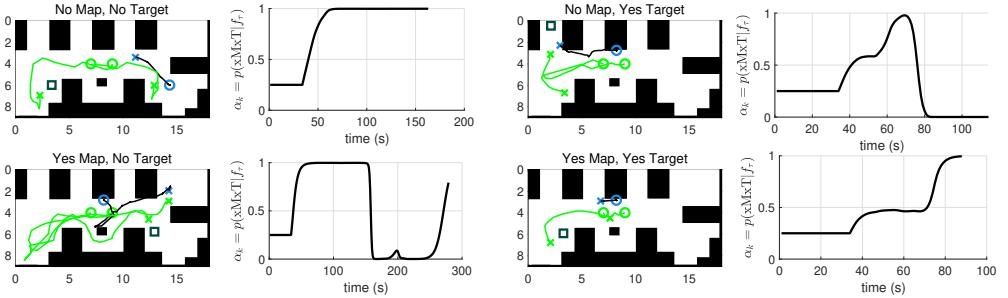


Fig. 10. Sample trials showing the teleoperated (black) and two autonomous robots (green) trajectories. Start and end positions are shown as circle and cross. Target location is shown as a square. The instantaneous  $\alpha_k$  values computed based on distance travelled by the teleoperated robot in the past 15 seconds are shown next to each trajectory plot.

knowledge during a search mission. In terms of prior knowledge inference, we note that  $\alpha_k$  is closer to 1 for a majority of the time for No Map, No Target condition, which is expected because the teleoperator has no knowledge of either. This in turn leads to the autonomous robots generally searching for the target away from the teleoperated robot. When two autonomous robots are searching along with the human teleoperator in this sample trial, it results in an early finding of the target, likely because of the high amount of exploration. At the same time, because  $\alpha_k$  is inferred based on movement data, it may be close to 1 even if the target knowledge is known (as in Yes Map, Yes Target in Fig. 10), when the teleoperated robot is moving slowly. We also note that when  $\alpha_k$  is closer to 0, the robots tend to stay close to the human but not necessarily searching in the same corner (one autonomous robot, Fig. 9, No Map, Yes Target) which exemplifies the capability of an information-based search which seeks to lower the overall target uncertainty, as opposed to a distance based search, where an autonomous robot would simply stay close.

Figures 11 and 12 summarize the performance of the two- and three-member human-robot team for a set of hundred experiments each consisting of four conditions, tested on randomly selected three participants whose data was not used to build the pdfs  $p(f_t|\mathcal{K})$ . This resulted in a total of 1200 simulations per team (300 per condition). We measure search performance in terms of the average time to find the target (Fig. 11a and Fig. 12ba), fraction of trials where the human-robot

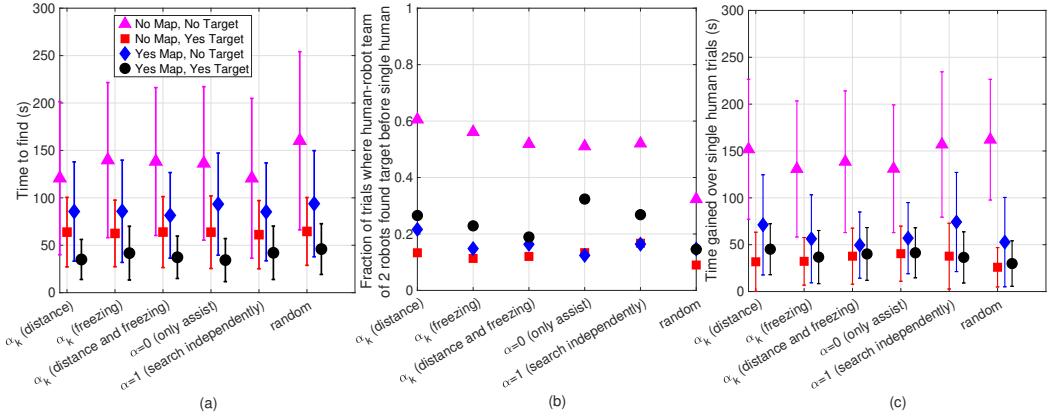


Fig. 11. Average  $\pm$  standard deviation of time to find the target (a) over all trials by the human robot team of 2 robots (one teleoperated, one autonomous); fraction of trials (b) where the human robot team found the target before the single human teleoperated robot and (c) the time saved by the team over the single robot. Adaptive search strategies are denoted by  $\alpha_k$ , and search strategies that involve mutual information based control with a fixed objective (assist or search independently) are denoted by the value assigned to  $\alpha$ .

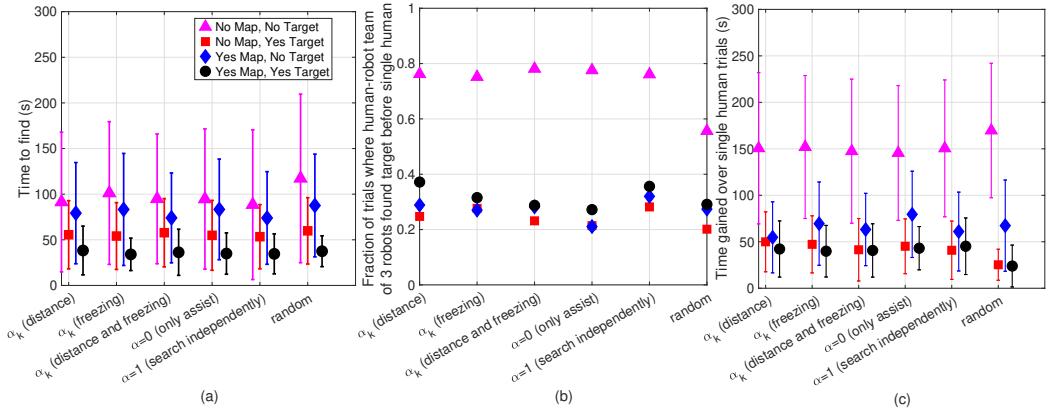


Fig. 12. Average  $\pm$  standard deviation of time to find the target (a) over all trials by the human robot team of 3 robots (one teleoperated, two autonomous); fraction of trials (b) where the human robot team found the target before the single human teleoperated robot and (c) the time saved by the team over the single robot. Adaptive search strategies are denoted by  $\alpha_k$ , and search strategies that involve mutual information based control with a fixed objective (assist or search independently) are denoted by the value assigned to  $\alpha$ .

team found the target before single human did (Fig. 11b and Fig. 12b), and time gained over those trials (Fig. 11c and Fig. 12c).

For a two-robot team (one human teleoperated and one autonomous, (Fig. 11)), we find that the time to find the target depends on the prior knowledge associated with the teleoperated robot. If prior knowledge was absent, all strategies perform similar except the random search which took a longer time to find. When evaluating search performance in terms of how the team does in comparison to a single robot (the value of assistance), we see that search performance of

autonomous robots depends on the type of prior knowledge associated with the teleoperated robot, with the largest gain registered when there was no knowledge of the target location or environment (xMxT=triangles). In terms of fraction of trials where the human-robot team found the target prior to a single human, the team where autonomous robots operated on prior knowledge based on distance outperformed all other strategies when target location was not known. Specifically, when there was no knowledge of target location (xMxT=triangles, yMxT=diamonds), the team where autonomous robot acted on prior knowledge inferred from distance had the best performance followed by the team that searched on the basis of freezing time; robot that searched randomly had the lowest performance for all conditions except when only the map was known; in that particular case, the team where the robot fully assisted the teleoperated robot performed worse. When both target location and map was known, the autonomous robot that always assisted had the best performance followed by the team where the robot searched independently or one that acted on prior knowledge based on distance.

In terms of time gained, all strategies had similar gain in search time with most time gained when target and environment knowledge was absent. Specifically, searching with an autonomous robot gave more than 100 seconds lead time when no prior knowledge was available and up to 45 seconds lead time when all prior knowledge was available. For the latter case, the highest time gained at 45.2 seconds was when the robot searched based on prior knowledge inferred from distance and the lowest time gained was when the robot searched randomly at 29.8 seconds.

For a three-robot team (one teleoperated and two autonomous), we expectedly see lower times to find the target compared to two-robot team. When comparing with a single human trials where the participant had no prior knowledge, all mutual information based control strategies perform at the same level (Fig. 12b and c), with adaptive search performing the best when informed by a combined features of distance and freezing, only slightly better than full assist and independent search. When all the prior knowledge was available, the adaptive search strategy based on distance performed the best and slightly better than independent search.

In terms of time gained among trials where the autonomous robots found the target earlier than the single human, all strategies gained approximately 150 seconds when there was no prior knowledge on target location and environment. When target location was known, the random search had the least gain in time compared to other mutual information based search strategies.

## 6 DISCUSSION

Search is a time-intensive task where members of search team may possess varying degrees of prior knowledge about the environment and the missing target. When searching for a missing target using teleoperated robots, autonomous robots can intelligently contribute to the search if they are able to infer human prior knowledge from the robot movement in the field.

Whereas there are several works that have utilized intent [25, 43] or trust [15, 68] inference to improve the performance of a human-robot team, inference of prior knowledge has not been explored. When such inference is available they are integrated into the robot control strategy through mixed initiative approaches [21], increased interventions [68], or risk minimization [24]. In this context, this work presents two contributions: (i) we conduct an experimental study that investigates the role of prior knowledge on search performance and behavior as a ground robot is teleoperated through a obstacle ridden environment; the results of this study allows us to isolate key movement cues that can be used to infer prior knowledge, and (ii) we test, validate, and compare an adaptive search strategy that blends mutual information combining independent search and human assistance.

Our experimental results indicate that prior knowledge of target location and environment can influence teleoperated robot movement in the field. Prior knowledge also affected mental and

physical demands and the perceived effort and frustration during a search task. Moving further, integrating this prior knowledge into an adaptive information blending strategy for an autonomous robot accompanying a teleoperated robot showed improvement in search efficiency when target location was unknown in terms of being able to find the target earlier than a human when compared to blind assistance, independent search or random search strategies. Increasing the number of autonomous robots in the same environment appears to dilute this gain in performance by adaptive strategies over blind assist and independent search.

### 6.1 Effect of Prior knowledge

*Performance.* Having prior knowledge of target location helped improve search performance in terms of time to find the target as well as the distance travelled. As expected, knowing target location beforehand produced lower times than when no such knowledge was available. When target location was known, the least time to find and distance travelled was attained when a map of the environment was available suggesting an efficient navigation around the environment. This result is consistent with the lost-person literature where it is generally agreed that searchers who have knowledge about the missing person behavior can influence the mission performance [49]. When target location was known, knowledge of the map reduced search time. At the same time, when target location was not known, knowledge of the environment did not significantly reduce the time to find. It is possible that for such a difference to manifest in an exploration-heavy mission, the environment should be even more complex. Our search environment had many obstacles, and we expected to see a significant effect of environmental knowledge on performance. However, based on the straight trajectories in subsequent trials, it appears that participants soon realized that they could explore certain areas between the benches and corners without actually entering them.

*Behavior.* Per our expectations we found that participants spent less time staying in place when they had prior knowledge about target location and the environment, than when they knew neither. Looking closely at conditions where the only difference was knowledge of target location, it is unlikely that knowing where to search reduced these behaviors. This was somewhat unexpected since part of being situationally aware is to know where to go and the lower search times from (No Map, No Target) condition to (No Map, Yes Target) suggested that participants may have stayed still much less. A likely explanation is that even though participants stayed still as often when target location was known they were much faster in between such times, more confident of their path. This may be seen through the differences in the speed probability distributions in Figure 8 for these two conditions. Specifically, we note that (No Map, Yes Target) has a heavier tail towards high speeds and lower probability of smaller speeds compared to (No Map, No Target). Participants were found freezing (staying still) more often when they did not know the map. It is also unlikely that they spent this time staying still reading the map given to them in the control room because it was very sparse during the first two tasks. On the other hand, it is likely that freezing was more indicative of an effort to gain situational awareness in our experiments possibly trying to remember their position with respect to the environment or locations they may have already visited.

Teleoperated robot speed increased with prior knowledge of target location. Differently, however, robot turn rate did not reduce with target knowledge. It is possible that as participants sped up with increased knowledge sudden stops and starts registered as high turn rates in the tracked data. We note that this is less likely due to tracking errors, which were small (see Supplementary material), and more likely a result of operating a robot with inertial mass. While some of these differences in robot movements could be reduced by designing better control gains, we note that in a real-world setting, this is an entirely expected phenomenon [60]. Situations where robot movement diverges significantly from user input may affect human perception of the robot even if they continue to

teleoperate it to search [52]. Knowledge of the environment resulted in higher turn rates when searching for a known target location. Again, this is possibly due to starts and stops at high speeds registering as sudden turns in the robot movement.

*Accuracy of target prior knowledge.* With respect to accuracy of target location, variations in speed as the trial progressed suggest that participants moved with frequent stop-and-go motions accompanied by high turn rate immediately after they realized that their prior knowledge about target knowledge was inaccurate. That these changes are most likely due to inaccuracy in prior knowledge, and not due to mistrust in the experimenter, is supported by participant response to the post-experiment survey questions where they indicated by their similar responses to the two questions related to trust in experimenter when they indicated the high-probability region. We also see that the trajectories (Figure 5) for such trials have more instances of moving in corners close to the possible target location. It is likely that this was due to repeated observations made by the participants who were trying to confirm that they properly checked the region. A similar behavior has been reported in reading studies when individuals are faced with information that is inconsistent with their beliefs [50]. We didn't find any difference in robot speed likely because (a) the participant knew the map well enough to not pause for situational awareness and (b) the participant could still easily guess the target location in the relatively small search environment by process of elimination.

## 6.2 Differences in perceived workload

Participants perceived a difference in all but one components of the NASA-TLX workload assessment. In particular, participants perceived differences in mental demand, physical demand, temporal demand (how rushed they felt), effort (how hard did they work to accomplish the level of performance), and frustration as a function of prior knowledge.

With respect to mental demand, participants felt that searching for a missing target with no prior knowledge in an unknown environment was more mentally demanding than when they knew where target was located. In addition to searching without any prior knowledge, participants in this task also had to navigate throughout the entire search environment with more obstacles. A significant drop in mental demand as participants proceeded through the conditions also suggests that familiarity with the environment played an important role. At the same time, because such a drop was not witnessed when participants had no target but good environmental knowledge suggests that familiarity with the environment was not the only factor at play. Finally, any possible frustration participants may have felt due to not finding the target in the expected high-probability location in (Yes Map, Yes Target) did not contribute to high overall mental demand.

These trends were somewhat mirrored in physical demand and effort with the additional difference that participants felt that they worked less hard in searching for a missing target in a known environment compared to in an unknown environment. As with mental demand, participants who found their prior knowledge to be inaccurate contributed to a sudden rise in these factors. It is possible that the additional time spent searching for the target amidst obstacles in the first condition created an impression of high physical demand and the need for extra effort, even though the experimental setup remained the same throughout conditions. Frustration among participants was only reduced when they had accurate prior knowledge of the target location. The reason we do not see a further lowering of frustration when participants had both target and environment knowledge is likely due to placement of target in a different location than where they expected.

### 6.3 Search performance with adaptive strategies

Two-member human robot teams comprising one teleoperated who did not know apriori where the target was located and one autonomous robot that adaptively blended its mutual information based objective function performed better than teams where the autonomous robot searched randomly or did not adapt their movement to prior knowledge. The largest improvements were found to be for when the teleoperator had no knowledge of target location.

When target location was known, the full-assist strategy worked best as it is the most likely strategy to effectively increase the field of view once the teleoperated robot reaches close to the target. The gain in performance by adaptive strategies over non-adaptive ones, especially independent search was diluted when the number of robots increased. This is likely because of the size and structure of the environment which made it relatively efficient for three robots searching independently to find the target. In particular, the relatively small environment coupled with visible sensing range of the robots made it easier for independently searching robots to quickly resolve the target location estimate to one of the four corners in the environment, where the target was always placed.

We also find that the full assist strategy no longer proved to have the best gain in performance when everything was known. This is likely because with more robots collision avoidance behaviors take precedence over search strategies. Among the three adaptive strategies, the one where weighting was inferred based on distance did generally better than the other two. This is possibly due to freezing time calculated over a larger time window of 30 seconds making it less responsive than a distance based inference which was calculated over a window of 15 seconds; a combined feature of distance and freezing would be suboptimal compared to distance or freezing at their respectively different observation periods.

### 6.4 Limitations

Limitations of this study include a relatively small environment which took on an average about two minutes to search by the participants; a larger search time could have revealed the behaviors more distinctly. Specifically, in a larger building with multiple rooms or multiple floors, prior knowledge about the map could have resulted in quicker navigation to different areas revealing differences in turn rate in addition to speed. While the inference of prior knowledge from movement data presents an opportunity for application in larger and more complex settings, using the same features across different environments may entail similarity in terms of obstacle density, size and type of robots.

A second limitation is the possibility that despite the presence of randomly placed obstacles, some participants may have understood the environment better than the others during the xMxT condition. Such differences in prior knowledge can be obtained in future by posing survey questions prior to every trial within an experimental session.

Another limitation is that the search environment was adequately lighted which may not realistically represent a real-world search. A low light environment could also have increased search time and added false positives, making the search process more complicated, and possibly revealing additional differences in search strategies. The generalizability of these results to broader search scenarios remains to be tested. In this context, virtual reality or augmented reality setups [17, 34, 44] may prove useful in creating a large variety of situations and robots including ground, walking, underwater, and aerial robots.

A fourth limitation of this study is the qualitative assessment of prior knowledge through movement cues. This is most evident in the constantly changing values of  $\alpha_k$  as inferred during a trial. Without a deeper assessment and quantification of how knowledge states evolve during a search mission it becomes difficult to say if the dynamic  $\alpha_k$  are a true reflection of prior knowledge

state of a teleoperator or a result of how the distributions of movement cues such as distance and freezing time overlap for all knowledge states.

A fifth limitation of this study is related to how we measure situational awareness and trust in the robot. While general measures of situational awareness that can provide continuous measurements involve eye tracking or physiological metrics [12], we favored a surrogate measure with the goal of directly inferring prior knowledge state from robot movement cues. Regarding trust in the robot, the only question in the post-experiment survey that measured if participants rated the performance of the teleoperation well was the question on latency. A broader set of questions adapted from validated metrics and measures [18, 20, 36] could have helped identify behavioral correlates of trust in teleoperation.

## 6.5 Conclusion and future directions

Results from this study show the potential of inferring meaningful information about human prior knowledge from robot movement during search with a teleoperated robot. Despite the same skill, teleoperated robot movement and operator workload was found to differ as a function of prior knowledge. Furthermore, limited time observation of select movement features were found to aid in autonomous robot assistance during a mutual-information based search. The response of human operators to the control strategies as proposed here remains to be evaluated. In particular, future work can compare an information-blending strategy with blind followership, independent search, and mixed initiative strategies in terms of human situational awareness, workload and trust.

The results of this study also set the stage for a data-driven dynamical model of human search that incorporates realistic representations of prior knowledge. A reliable dynamical model of human search can be used to test complex hypotheses in human-robot teams [37]. Direct inference of human prior knowledge from robot movement has applications in human swarm robotics [31] where autonomous robots could adopt leader-follower strategies based on field observations of teleoperated robot movement [32, 37].

## ACKNOWLEDGMENTS

This research was partially supported by National Science Foundation under grant # IIS-2033918, Illinois Space Grant Consortium under Subaward NNX15AI05H, and a Great Journeys Assistantship to Rafal Krzysiak. The authors gratefully acknowledge Zachary Taylor for developing the marker based tracking system, and Di'Quan Ishmon and Arunim Bhattacharya for help in performing experiments towards error analysis of the tracking system and calibrating the robot movement.

## REFERENCES

- [1] M. Sanjeev Arulampalam, Simon Maskell, Neil Gordon, and Tim Clapp. 2002. A tutorial on particle filters for online nonlinear/nongaussian bayesian tracking. *IEEE Transactions on Signal Processing* 50, 2 (2002), 723–737. <https://doi.org/10.1109/78.978374>
- [2] Frédéric Bourgault, Tomonari Furukawa, and Hugh F. Durrant-Whyte. 2006. Optimal Search for a Lost Target in a Bayesian World. In *Field and Service Robotics: Recent Advances in Research and Applications*, Shin'ichi Yuta, Hajima Asama, Erwin Prassler, Takashi Tsubouchi, and Sebastian Thrun (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 209–222. [https://doi.org/10.1007/978-3-540-34828-7\\_11](https://doi.org/10.1007/978-3-540-34828-7_11)
- [3] Saniye Tugba Bulu and Susan Pedersen. 2012. Supporting problem-solving performance in a hypermedia learning environment: The role of students' prior knowledge and metacognitive skills. *Computers in Human Behavior* 28, 4 (2012), 1162–1169. <https://doi.org/10.1016/j.chb.2012.01.026>
- [4] Jennifer L. Burke, Robin R. Murphy, Michael D. Coover, and Dawn L. Riddle. 2004. Moonlight in Miami: A field study of human-robot interaction in the context of an urban search and rescue disaster response training exercise. *Human-Computer Interaction* 19, 1-2 (2004), 85–116. <https://doi.org/10.1080/07370024.2004.9667341>
- [5] Alex Cao, Keshav K Chintamani, Abhilash K Pandya, and R Darin Ellis. 2009. NASA TLX: Software for assessing subjective mental workload. *Behavior Research Methods* 41, 1 (2009), 113–117. <https://doi.org/10.3758/BRM.41.1.113>

[6] Jennifer L Casper and Robin R Murphy. 2002. Workflow study on human-robot interaction in USAR. In *Proceedings of the IEEE International Conference on Robotics and Automation*, Vol. 2. IEEE, 1997–2003. <https://doi.org/10.1109/ROBOT.2002.1014834>

[7] Yuhang Che, Heather Culbertson, Chih-Wei Tang, Sudipto Aich, and Allison M Okamura. 2018. Facilitating human–mobile robot communication via haptic feedback and gesture teleoperation. *ACM Transactions on Human-Robot Interaction (THRI)* 7, 3 (2018), 1–23. <https://doi.org/10.1145/3243503>

[8] Hugh Cottam, Nigel Shadbolt, John Kingston, Howard Beck, and Austin Tate. 1995. Knowledge level planning in the search and rescue domain. In *Research and Development in Expert Systems XII*, Vol. 9. SGES Publications, 309–326.

[9] Jeffrey Delmerico, Stefano Mintchev, Alessandro Giusti, Boris Gromov, Kamilo Melo, Tomislav Horvat, Cesar Cadena, Marco Hutter, Auke Ijspeert, Dario Floreano, et al. 2019. The current state and future outlook of rescue robotics. *Journal of Field Robotics* 36, 7 (2019), 1171–1191. <https://doi.org/10.1002/rob.21887>

[10] Anca Dragan and Siddhartha Srinivasa. 2014. Integrating human observer inferences into robot motion planning. *Autonomous Robots* 37, 4 (2014), 351–368. <https://doi.org/10.1007/s10514-014-9408-x>

[11] Nima Enayati, Giancarlo Ferrigno, and Elena De Momi. 2018. Skill-based human–robot cooperation in tele-operated path tracking. *Autonomous Robots* 42, 5 (2018), 997–1009. <https://doi.org/10.1007/s10514-017-9675-4>

[12] Mica R Endsley. 2021. A Systematic Review and Meta-Analysis of Direct Objective Measures of Situation Awareness: A Comparison of SAGAT and SPAM. *Human Factors* 63, 1 (2021), 124–150. <https://doi.org/10.1177/0018720819875376>

[13] Don Ferguson. 2008. GIS for wilderness search and rescue. In *ESRI Federal User Conference*, Vol. 2012. 10.

[14] Angel Flores-Abad, Ou Ma, Khanh Pham, and Steve Ulrich. 2014. A review of space robotics technologies for on-orbit servicing. *Progress in Aerospace Sciences* 68 (2014), 1–26. <https://doi.org/10.1016/j.paerosci.2014.03.002>

[15] Maziar Fooladi Mahani, Longsheng Jiang, and Yue Wang. 2021. A bayesian trust inference model for human–multi-robot teams. *International Journal of Social Robotics* 13, 8 (2021), 1951–1965. <https://doi.org/10.1007/s12369-020-00705-1>

[16] Dieter Fox, Jonathan Ko, Kurt Konolige, Benson Limketkai, Dirk Schulz, and Benjamin Stewart. 2006. Distributed multirobot exploration and mapping. *Proc. IEEE* 94, 7 (2006), 1325–1339. <https://doi.org/10.1109/JPROC.2006.876927>

[17] Juan C García, Bruno Patrão, Luís Almeida, Javier Pérez, Paulo Menezes, Jorge Dias, and Pedro J Sanz. 2015. A natural interface for remote operation of underwater robots. *IEEE Computer Graphics and Applications* 37, 1 (2015), 34–43. <https://doi.org/10.1109/MCG.2015.118>

[18] P Goillau, C Kelly, M Boardman, and Emmanuelle Jeannot. 2003. *Guidelines for trust in future atm systems-measures*. Technical Report 1. European Organization for Safety of Air Navigation.

[19] Yue Guo, Rohit Jena, Dana Hughes, Michael Lewis, and Katia Sycara. 2021. Transfer learning for human navigation and triage strategies prediction in a simulated urban search and rescue task. In *IEEE International Conference on Robot & Human Interactive Communication (RO-MAN)*. IEEE, 784–791. <https://doi.org/10.1109/RO-MAN50785.2021.9515526>

[20] Yaohui Guo and X Jessie Yang. 2021. Modeling and predicting trust dynamics in human–robot teaming: A Bayesian inference approach. *International Journal of Social Robotics* 13, 8 (2021), 1899–1909. <https://doi.org/10.1007/s12369-020-00703-3>

[21] Benjamin Hardin and Michael A Goodrich. 2009. On using mixed-initiative control: A perspective for managing large-scale robotic teams. In *Proceedings of the ACM/IEEE International Conference on Human Robot Interaction*. 165–172. <https://doi.org/10.1145/1514095.1514126>

[22] Kelsey P Hawkins, Shray Bansal, Nam N Vo, and Aaron F Bobick. 2014. Anticipating human actions for collaboration in the presence of task and sensor uncertainty. In *IEEE International Conference on Robotics and Automation*. IEEE, 2215–2222. <https://doi.org/10.1109/ICRA.2014.6907165>

[23] Hooman Hedayati, Michael Walker, and Daniel Szafir. 2018. Improving collocated robot teleoperation with augmented reality. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*. 78–86. <https://doi.org/10.1145/3171221.3171251>

[24] Larkin Heintzman, Amanda Hashimoto, Nicole Abaid, and Ryan K Williams. 2021. Anticipatory planning and dynamic lost person models for human–robot search and rescue. In *IEEE International Conference on Robotics and Automation*. IEEE, 8252–8258. <https://doi.org/10.1109/ICRA48506.2021.9562070>

[25] Guy Hoffman, Tapomayukh Bhattacharjee, and Stefanos Nikolaidis. 2023. Inferring Human Intent and Predicting Human Action in Human–Robot Collaboration. *Annual Review of Control, Robotics, and Autonomous Systems* 7 (2023), 73–95. <https://doi.org/10.1146/annurev-control-071223-105834>

[26] Gabriel M. Hoffmann and Claire J. Tomlin. 2010. Mobile sensor network control using mutual information methods and particle filters. *IEEE Trans. Automat. Control* 55, 1 (2010), 32–47. <https://doi.org/10.1109/TAC.2009.2034206>

[27] Farrokh Janabi-Sharifi and Iraj Hassanzadeh. 2010. Experimental analysis of mobile-robot teleoperation via shared impedance control. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 41, 2 (2010), 591–606. <https://doi.org/10.1109/TSMCB.2010.2073702>

[28] Brian J. Julian, Sertac Karaman, and Daniela Rus. 2014. On mutual information-based control of range sensing robots for mapping applications. *International Journal of Robotics Research* 33, 10 (2014), 1375–1392. <https://doi.org/10.1177/0278364014531872>

0278364914526288

[29] Richard Kelley, Alireza Tavakkoli, Christopher King, Monica Nicolescu, Mircea Nicolescu, and George Bebis. 2008. Understanding human intentions via hidden markov models in autonomous mobile robots. In *Proceedings of the ACM/IEEE International Conference on Human Robot Interaction*. 367–374. <https://doi.org/10.1145/1349822.1349870>

[30] ShinWoo Kim and Bob Rehder. 2011. How prior knowledge affects selective attention during category learning: An eyetracking study. *Memory & cognition* 39, 4 (2011), 649–665. <https://doi.org/10.3758/s13421-010-0050-3>

[31] Andreas Kolling, Phillip Walker, Nilanjan Chakraborty, Katia Sycara, and Michael Lewis. 2016. Human Interaction with Robot Swarms: A Survey. *IEEE Transactions on Human-Machine Systems* 46, 1 (2016), 9–26. <https://doi.org/10.1109/THMS.2015.2480801>

[32] Rafal Krzysiak and Sachit Butail. 2021. Information-Based Control of Robots in Search-and-Rescue Missions With Human Prior Knowledge. *IEEE Transactions on Human-Machine Systems* 51, 1 (2021), 52–63. <https://doi.org/10.1109/THMS.2021.3113642>

[33] Dana Kulic and Elizabeth A Croft. 2007. Affective state estimation for human–robot interaction. *IEEE Transactions on Robotics* 23, 5 (2007), 991–1000. <https://doi.org/10.1109/TRO.2007.904899>

[34] Daniel Labonte, Patrick Boissy, and François Michaud. 2010. Comparative analysis of 3-D robot teleoperation interfaces with novice users. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 40, 5 (2010), 1331–1342. <https://doi.org/10.1109/TSMCB.2009.2038357>

[35] Gregory Lemasurier, Gal Bejerano, Victoria Albanese, Jenna Parrillo, Holly A Yanco, Nicholas Amerson, Rebecca Hetrick, and Elizabeth Phillips. 2021. Methods for Expressing Robot Intent for Human–Robot Collaboration in Shared Workspaces. *ACM Transactions on Human–Robot Interaction (THRI)* 10, 4 (2021), 1–27. <https://doi.org/10.1145/3472223>

[36] Michael Lewis, Katia Sycara, and Phillip Walker. 2018. The role of trust in human–robot interaction. In *Foundations of Trusted Autonomy*, Reid D. Abbass H., Scholz J. (Ed.). Springer International Publishing, 135–159. [https://doi.org/10.1007/978-3-319-64816-3\\_8](https://doi.org/10.1007/978-3-319-64816-3_8)

[37] Mengxi Li, Minae Kwon, and Dorsa Sadigh. 2021. Influencing leading and following in human–robot teams. *Autonomous Robots* 45, 7 (2021), 959–978. <https://doi.org/10.1007/s10514-021-10016-7>

[38] Yugang Liu and Goldie Nejat. 2013. Robotic urban search and rescue: A survey from the control perspective. *Journal of Intelligent & Robotic Systems* 72, 2 (2013), 147–165. <https://doi.org/10.1007/s10846-013-9822-x>

[39] Yugang Liu and Goldie Nejat. 2013. Robotic urban search and rescue: A survey from the control perspective. *Journal of Intelligent & Robotic Systems* 72 (2013), 147–165. <https://doi.org/10.1007/s10846-013-9822-x>

[40] Matteo Luperto, Michele Antonazzi, Francesco Amigoni, and N Alberto Borghese. 2020. Robot exploration of indoor environments using incomplete and inaccurate prior knowledge. *Robotics and Autonomous Systems* 133 (2020), 103622. <https://doi.org/10.1016/j.robot.2020.103622>

[41] Ashish Macwan, Goldie Nejat, and Beno Benhabib. 2011. Target-motion prediction for robotic search and rescue in wilderness environments. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 41, 5 (2011), 1287–1298. <https://doi.org/10.1109/TSMCB.2011.2132716>

[42] Keiji Nagatani, Seiga Kiribayashi, Yoshito Okada, Kazuki Otake, Kazuya Yoshida, Satoshi Tadokoro, Takeshi Nishimura, Tomoaki Yoshida, Eiji Koyanagi, Mineo Fukushima, et al. 2013. Emergency response to the nuclear accident at the Fukushima Daiichi Nuclear Power Plants using mobile rescue robots. *Journal of Field Robotics* 30, 1 (2013), 44–63. <https://doi.org/10.1002/rob.21439>

[43] Manisha Natarajan, Esmaeil Seraj, Batuhan Altundas, Rohan Paleja, Sean Ye, Letian Chen, Reed Jensen, Kimberlee Chestnut Chang, and Matthew Gombolay. 2023. Human–robot teaming: grand challenges. *Current Robotics Reports* 4, 3 (2023), 81–100. <https://doi.org/10.1007/s43154-023-00103-1>

[44] Curtis W. Nielsen, Michael A. Goodrich, and Robert W. Ricks. 2007. Ecological interfaces for improving mobile robot teleoperation. *IEEE Transactions on Robotics* 23, 5 (2007), 927–941. <https://doi.org/10.1109/TRO.2007.907479>

[45] Mitsuhige Oda, Kouichi Kibe, and Fumio Yamagata. 1996. ETS-VII, space robot in-orbit experiment satellite. In *Proceedings of IEEE International Conference on Robotics and Automation*. IEEE, 739–744. <https://doi.org/10.1109/ROBOT.1996.503862>

[46] Dimitri Ognibene, Lorenzo Mirante, and Letizia Marchegiani. 2019. Proactive intention recognition for joint human–robot search and rescue missions through monte-carlo planning in pomdp environments. In *International Conference on Social Robotics, ICSR*. Springer, 332–343. [https://doi.org/10.1007/978-3-030-35888-4\\_31](https://doi.org/10.1007/978-3-030-35888-4_31)

[47] Yasuhiro Ozuru, Kyle Dempsey, and Danielle S McNamara. 2009. Prior knowledge, reading skill, and text cohesion in the comprehension of science texts. *Learning and Instruction* 19, 3 (2009), 228–242. <https://doi.org/10.1016/j.learninstruc.2008.04.003>

[48] Victor M Panaretos and Yoav Zemel. 2019. Statistical aspects of Wasserstein distances. *Annual Review of Statistics and its Application* 6 (2019), 405–431. <https://doi.org/10.1146/annurev-statistics-030718-104938>

[49] Ken Phillips, Maura J Longden, Bil Vandergraff, William R Smith, David C Weber, Scott E McIntosh, and Albert R Wheeler III. 2014. Wilderness search strategy and tactics. *Wilderness & Environmental Medicine* 25, 2 (2014), 166–176.

<https://doi.org/10.1016/j.wem.2014.02.006>

[50] David N Rapp. 2008. How do readers handle incorrect information during reading? *Memory & Cognition* 36, 3 (2008), 688–701. <https://doi.org/10.3758/MC.36.3.688>

[51] Jennifer M Riley and Mica R Endsley. 2004. The hunt for situation awareness: Human-robot interaction in search and rescue. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 48. SAGE Publications Sage CA: Los Angeles, CA, 693–697. <https://doi.org/10.1177/15419312040480038>

[52] Maha Salem, Gabriella Lakatos, Farshid Amirabdollahian, and Kerstin Dautenhahn. 2015. Would you trust a (faulty) robot? Effects of error, task type and personality on human-robot cooperation and trust. In *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 1–8. <https://doi.org/10.1145/2696454.2696497>

[53] Yangming Shi, John Kang, Pengxiang Xia, Oshin Tyagi, Ranjana K Mehta, and Jing Du. 2021. Spatial knowledge and firefighters' wayfinding performance: A virtual reality search and rescue experiment. *Safety Science* 139 (2021), 105231. <https://doi.org/10.1016/j.ssci.2021.105231>

[54] V Smrz, J Boril, I Vudarcik, and M Bauer. 2019. Utilization of Recorded Flight Simulator Data to Evaluate Piloting Accuracy and Quality. In *2019 New Trends in Aviation Development (NTAD)*. IEEE, 164–169. <https://doi.org/10.1109/NTAD.2019.8875570>

[55] Hyoung Il Son, Antonio Franchi, Lewis L Chuang, Junsuk Kim, Heinrich H Bulthoff, and Paolo Robuffo Giordano. 2013. Human-centered design and evaluation of haptic cueing for teleoperation of multiple mobile robots. *IEEE Transactions on Cybernetics* 43, 2 (2013), 597–609. <https://doi.org/10.1109/TCYB.2012.2212884>

[56] Mark Steyvers, Michael D Lee, and Eric-Jan Wagenmakers. 2009. A Bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology* 53, 3 (2009), 168–179. <https://doi.org/10.1016/j.jmp.2008.11.002>

[57] Richard T Stone, Michael Dorneich, Stephen Gilbert, and Elease McLaurin. 2013. Human Differences in Navigational Approaches during Tele-Robotic Search. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 57. SAGE Publications Sage CA: Los Angeles, CA, 625–629. Issue 1. <https://doi.org/10.1177/1541931213571134>

[58] Russell H Taylor, Arianna Menciassi, Gabor Fichtinger, Paolo Fiorini, and Paolo Dario. 2016. Medical robotics and computer-integrated surgery. In *Springer Handbook of Robotics*. Springer, 1657–1684. [https://doi.org/10.1007/978-3-319-32552-1\\_63](https://doi.org/10.1007/978-3-319-32552-1_63)

[59] Shirlene Wade and Celeste Kidd. 2019. The role of prior knowledge and curiosity in learning. *Psychonomic Bulletin & Review* 26, 4 (2019), 1377–1387. <https://doi.org/10.3758/s13423-019-01598-6>

[60] Toshihiro Wakita, Koji Ozawa, Chiyomi Miyajima, and Kazuya Takeda. 2005. Parametric versus non-parametric models of driving behavior signals for driver identification. In *International Conference on Audio-and Video-Based Biometric Person Authentication*. Springer, 739–747. [https://doi.org/10.1007/11527923\\_77](https://doi.org/10.1007/11527923_77)

[61] David Waller, Jack M Loomis, Reginald G Golledge, and Andrew C Beall. 2000. Place learning in humans: The role of distance and direction information. *Spatial Cognition and Computation* 2, 4 (2000), 333–354. <https://doi.org/10.1023/A:1015514424931>

[62] Sarah E Walsh and Karen M Feigh. 2022. Understanding human decision processes: Inferring decision strategies from behavioral data. *Journal of Cognitive Engineering and Decision Making* 16, 4 (2022), 301–325. <https://doi.org/10.1177/15553434221122899>

[63] Weitian Wang, Rui Li, Yi Chen, Yi Sun, and Yunyi Jia. 2021. Predicting human intentions in human–robot hand-over tasks through multimodal learning. *IEEE Transactions on Automation Science and Engineering* 19, 3 (2021), 2339–2353. <https://doi.org/10.1109/TASE.2021.3074873>

[64] Wenshuo Wang, Junqiang Xi, and Huiyan Chen. 2014. Modeling and recognizing driver behavior based on driving data: A survey. *Mathematical Problems in Engineering* 2014 (2014), 245641. <https://doi.org/10.1155/2014/245641>

[65] Zhikun Wang, Katharina Mülling, Marc Peter Deisenroth, Heni Ben Amor, David Vogt, Bernhard Schölkopf, and Jan Peters. 2013. Probabilistic movement modeling for intention inference in human–robot interaction. *The International Journal of Robotics Research* 32, 7 (2013), 841–858. <https://doi.org/10.1177/0278364913478447>

[66] Ryan Wegner and John Anderson. 2006. Agent-based support for balancing teleoperation and autonomy in urban search and rescue. *International Journal of Robotics and Automation* 21, 2 (2006), 120–128. <https://doi.org/10.2316/Journal.206.2006.2.206-2796>

[67] Adam Williams, Bijo Sebastian, and Pinhas Ben-Tzvi. 2019. Review and analysis of search, extraction, evacuation, and medical field treatment robots. *Journal of Intelligent & Robotic Systems* 96, 3 (2019), 401–418. <https://doi.org/10.1007/s10846-019-00991-6>

[68] R Oliver Zanone and Javad Mohammadpour Velni. 2023. Trust-based Performance Optimization for Human-Swarm Collaboration. *IFAC-PapersOnLine* 56, 3 (2023), 571–576. <https://doi.org/10.1016/j.ifacol.2023.12.085>

[69] Mojtaba Zarei, Navid Kashi, Ahmad Kalhor, and Mehdi Tale Masouleh. 2020. Experimental study on shared-control of a mobile robot via a haptic device with an optimal velocity obstacle based receding horizon control approach. *Journal of Intelligent & Robotic Systems* 97, 2 (2020), 357–372. <https://doi.org/10.1007/s10846-019-01023-z>