Spatial Tasking in Human-Robot Collaborative Exploration

Yves Georgy Daoud

CMU-RI-TR-22-70

December 6, 2022

The Robotics Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA

Thesis Committee:
Wennie Tabib, Co-Chair
Nathan Michael, Co-Chair
Henny Admoni
Tejus Gupta

Submitted in partial fulfillment of the requirements for the degree of Master of Science in Robotics.

Abstract

This work develops a methodology for collaborative human-robot exploration that leverages implicit coordination. Most autonomous single- and multi-robot exploration systems require a remote operator to provide explicit guidance to the robot team. Few works consider how to integrate the human partner alongside robots to provide guidance in the field. A challenge for collaborative human-robot exploration is the efficient communication of goals from the human to the robot. For applications like search and rescue, which are time-sensitive and high-stress endeavors, robot assistants must provide value for human partners to motivate deployment in the field.

In this thesis, we develop a methodology that implicitly communicates a region of interest from a helmet-mounted depth camera on the human's head to the robot and an information gain-based exploration objective that biases motion planning within the viewpoint provided by the human. We also study the human perception of robot efficiency in a search task to better understand how to design and develop robots that explore alongside human partners. The result is an aerial system that safely accesses regions of interest that may not be immediately viewable or reachable by the human. The approach is evaluated in simulation and with hardware experiments in a motion capture arena. Our findings from the subject study suggest that users' trust in the robot is highly dependent on mission success. Additionally, participants perceived that they had more control over the robot's actions when the mission was successful, and the robot as being more in control when the mission failed. The results highlight the importance of designing robust robots with transparent behaviors for successful human-robot collaboration.

Acknowledgments

I would first like to thank my co-advisors Wennie Tabib and Nathan Michael. Thank you both for providing me with the tools, knowledge, and guidance necessary to conduct my research. It has been a pleasure being advised by you. I appreciate every time you reoriented me to the bigger picture of what we are trying to solve whenever I faced obstacles in my research. Thank you for believing in me and taking me under your wing as a master's student at RISLab. I would also like to thank Henny Admoni for helping me see the problem from a human-robot interaction perspective, and taking the time to go over the analyses with me.

To every graduate student that has mentored me along this journey, including Kshitij Goel, Tejus Gupta, and Alex Spitzer, thank you. Kshitij for constantly providing critique and suggestions on the approach, theory, and experiments for the SSRR work, Tejus for being on my committee and providing high-level feedback, and Alex for onboarding me in the lab.

I am grateful for every person and friend I have met along the way at RI and CMU. You have each indirectly contributed to my well-being and motivated me in hard times to keep going. Thank you for every coffee chat, pool match, walk outside, and new cuisine that you have introduced me to. Thank you to every staff member behind the curtains who has made this journey possible.

Finally, I thank my parents for believing in me and supporting me from overseas. This thesis is dedicated to my mom and dad, for being resilient in life and teaching me the everyday values that I adhere to. Thank you.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Recent years have seen a rise in the use of robots for Search and Rescue (SAR) missions [15]. Within the context of robotics deployment in *a priori* unknown environments, the spectrum of autonomy has varied from teleoperated semi-autonomous robots [17], to fully autonomous deep learning-based solutions [34]. Such unmapped environments present a variety of challenges for human-driven SAR. Underground environments, such as mines and caves, contain small passageways (see fig. 1.1a) that are difficult for humans to traverse. Tight passageways pose a risk to humans, because they may become trapped. For example, in 2009 a man perished in Utah's Nutty Putty Cave after becoming trapped upside-down in a narrow passage. Caves also have technically difficult sections, such as the one shown in fig. 1.1b, which requires the use of both hands to operate special equipment and limits the operator's degrees of freedom. Given these challenges, robot-assisted SAR that: (1) allows hands-free task allocation, (2) supports guided exploration, (3) alleviates the human's cognitive load, and (4) requires minimal supervision is crucial to enable efficient human-robot collaboration.

State-of-the-art exploration methodologies leverage the human as an *operator* outside of the exploration environment instead of directly engaging them side-by-side with robots [37, 39, 43]. Modeling the human as a *collaborator* instead of an *operator* in a shared workspace for exploration enables more efficient distributed exploration and useful emergent robot behaviors [19]. In this work, a collaborative human-robot exploration system is developed to explore 3D unstructured environments (fig. 1.2a) by communicating the field of view (FoV) of the human to the robot (fig. 1.2b) and

(a) Image reproduced from Tabib et al. [40].

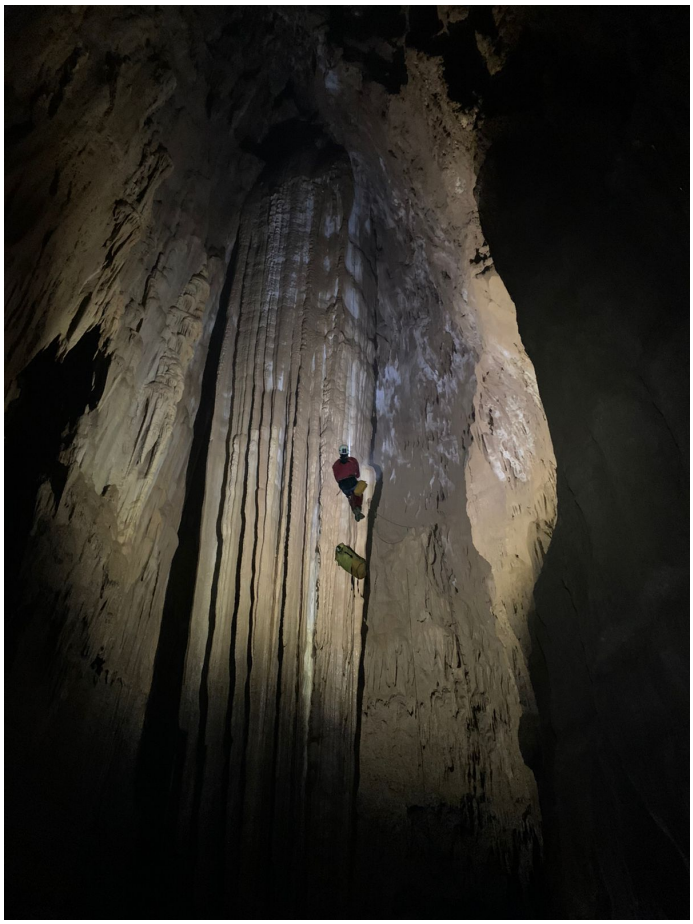

(b) Image courtesy of R. Arja.

Figure 1.1: Challenges found in cave environments. A caver can be seen crawling and mapping a tight passageway in fig. 1.1a, while fig. 1.1b shows a caver rappelling down a rope.

having the robot use this region of interest (ROI) to bias motion plans to acquire views of areas occluded to the human (fig. 1.2c).

Explicit robot tasking is impractical [14] during time-sensitive human-robot collaborative exploration (e.g. cave search and rescue) if humans must reduce their operational tempo [21, 25, 47], so implicit communication of spatial goals is imperative. State-of-the-art exploration objectives reduce environmental uncertainty without providing the flexibility to prioritize ROIs. To address these gaps in the state of the art, this work presents a collaborative human-robot exploration system that: (1) leverages implicit communication to spatially task an aerial system to regions occluded to the human, and (2) develops an information-gain based objective function inspired by the active object reconstruction literature [12] to bias motion planning within the ROI specified by the human. The approach is evaluated with real-time simulations and real-world hardware experiments in a motion capture arena.

In addition to mitigating the environmental challenges of SAR missions through efficient spatial tasking of the robot, we must study the human-robot interaction under the time-sensitive and high-stress conditions that these environments present [7, 31]. In particular, Casper and Murphy [7] find that the physical (e.g. lack of sleep) and psychological (e.g. stress and fatigue) difficulties, especially during the initial critical hours of the search task, take a toll on first responders. In light of these stressors, trust is an important factor to study, since a high amount of cognitive load results in a degradation of trust perception and task performance [9]. Failure to trust the robot partner might affect the human's reliance on the system [29], resulting in a failure to complete the SAR task. The idea of trust has been studied heavily in the context of HRI [38, 44]. This work is unique from the trust literature in three ways. First, few studies have considered the effect of human trust in an SAR context. Second, the relation between task efficiency and perceived trust in SAR is yet to be understood. Finally, to the best of our knowledge, the dynamics between the human and the robot within a shared autonomy scenario and its effect on trust has not been explored. To that extent, this paper studies the human's perception of a robotic partner in a simulated search scenario and the effect of task efficiency on perceived trust.

(a)

(b)

(c)

Figure 1.2: (a) A human-robot team is tasked with exploring a cave. (b) The human implicitly conveys a region of interest to the robot by means of transmission of their current viewpoint. (c) The robot plans a path to areas of the environment that are occluded to the human.

Chapter 2

Background and Related Work

Human-robot collaboration for search and exploration is a fairly new area combining different aspects of field robotics and HRI, mainly lying at the intersection of two areas: human-robot collaborative exploration and the relationship between human factors and robotic performance. In this section, we review and contrast related works with the method detailed in this paper.

With regards to collaborative human-robot exploration, few prior works study the use of implicit coordination. Govindarajan et al. [19] achieve coordination through a distributed strategy that assigns robots to homotopy classes that are complementary to the ones being traversed by the human. It is assumed that a blueprint of the environment is available to identify homotopy classes before operation. In contrast, the proposed approach does not assume prior information on the environment layout. A motion primitive-based planner is leveraged to maximize information gain, which takes the human's view into account, and drives the robot to explore regions occluded to the human; therefore, prior environment knowledge is not required. Within the context of multitasking, implicit communication has been used to augment human situational awareness via a robotic system. Bentz et al. [3] leverage head tracking while a human performs an arbitrary number of complex tasks and fit the data to a visual interest function. An aerial robot uses the visual interest function to provide camera views that augment the human's situational awareness. This methodology does not directly translate to the exploration context because the visual interest function, which effectively rates the utility of a viewpoint, is difficult to specify

before or during exploration. Instead, the proposed approach uses the notion of potential information gain over a discrete set of candidate viewpoints to drive the robot towards the ROI. Reardon et al. [36] leverage augmented reality to share information between a robot and human cooperatively exploring in the field. The goal is to influence the behavior of the human teammate in the human-robot cooperative exploration task by sharing information about the robot's current plan, the task state, and communicating future actions. In contrast, the proposed approach develops a methodology to influence the robot's behavior depending on actions taken by the human. This implicit coordination is desirable in applications like search and rescue where the robot is expected to adapt to the human's operational tempo. Caltieri and Amigoni [6] propose a semi-supervised autonomy stack for human-robot interaction during SAR operations. Simulation experiments with recruited participants show that high-level commands, which are defined as direction and area of interest, reduce the human's workload and increase both the area covered and number of victims found as compared to a direct teleoperation method. However, the study in [6] treats the human as an operator outside the field. In contrast, our work enables the human to be involved in the mission next to the robots, while also providing high-level guidance. Using a different modality for issuing commands, Gromov et al. [20] develop a method that utilizes gestures to direct the vehicle on where to go. However, in SAR, the human may require the use of their hands in order to negotiate the challenging terrain, such as the one seen in fig. 1.1b. In contrast, we provide the human with hands-free task allocation using head orientation. In addition, their approach is a direct and instantaneous low-level control of the robot as opposed to the collaborative exploration system we develop in this work. Aggravi et al. [1] develop an approach for visiting a list of known targets in the context of reconnaissance by a multi-agent team consisting of one human and at least one robot. To maintain connectivity during the patrol, the user is equipped with haptic feedback to notify them when they deviate from the team formation. Consequently, this approach places the burden of maintaining connectivity on the user, slowing them down and increasing their cognitive load. In time-critical tasks like SAR, the robot must keep up with the human's cadence or it will not provide value to the mission. To address this drawback, our approach allows for robot operation with minimal supervision from their human partner.

In order to enable human-robot collaborative exploration, many motion planning objectives have been proposed for active mapping. For example, frontier-based objectives utilize the distance to the boundary between unknown and known space to drive the robot's exploration [48]. For multi-agent operation, prioritization between frontiers is utilized to assign agents towards complementary regions of the environment [5]. However, deployments of this idea and its variants have been limited to multiple robots [13, 37], with the human largely supplying spatial goals explicitly when desired [33]. Further, information-theoretic objectives utilize the expected change in the entropy of the map due to candidate sensor measurements to drive viewpoint selection in both 2D [4, 24] and 3D environments [8]. While real-world operation has been shown using single [41] and multiple aerial robots [16], these techniques have not been leveraged in collaborative human-robot exploration. A key capability missing in these objectives is the ability to prioritize spatial ROIs. Towards imposing such spatial constraints, volumetric next-best-view (NBV) selection methods have been proposed for the active object reconstruction problem where the viewpoint is generated to focus on high-fidelity reconstruction of a single object [28, 45]. Delmerico et al. [12] propose several variants of information gain objectives that are either counting-based [45], probabilistic [28], or a combination. However, a method to apply these objectives in the collaborative human-robot exploration system is lacking in the literature. To this end, this work proposes and evaluates an Occlusion-Aware Volumetric Information (OAVI) objective that extends the work of Delmerico et al. [12] to the collaborative human-robot exploration problem. We further contrast it to ROI-constrained Cauchy-Schwarz Quadratic Mutual Information (ROI-CSQMI), an extension of [8] developed in this work, which applies the human's FoV as a spatial constraint.

Regarding the study of robot performance in accordance with human factors, Gombolay et al. [18] evaluate the effect of task scheduling with regards to user preferences in the context of human-robot teaming. Their experiments show that human subjects prefer working with a robot which accounts for their personal choices. However, this preference does not hold when the team's efficiency is at stake, resulting in the participants preferring a robot that helps them complete the task faster. The relationship between trust and robot performance in a shared control framework is yet to be answered. Thus, in this work, we investigate how task efficiency affects the

human's perception of trust in a robot that follows their guidance. Baraglia et al. [2] compare different strategies for robot-assisted task completion in an attempt to answer when and how should a robot take the initiative to assist a human. Their findings suggest that users prefer the robot to help when requested rather than react when it perceives that the human is in need of help. This is in line with our methodology; the robot operates as an autonomous agent navigating the environment, then proceeds to bias its exploration policy towards the human's POV when requested. Moreover, our work looks into the implications of failure on trust within a human-robot shared control scheme. Lyons et al. [30] explore the effect of unexpected robot behavior on trust, and ways to repair that trust using explanations from the robot. The study simulates a SAR scenario where a human operator provides a goal for a fully autonomous robot to search for victims in. Participants are first shown a video of a robot following its expected path commanded by the human operator. The second video depicts the robot deviating from its original path and going into a separate area, therefore violating the operator's instructions. Trust and responsibility attribution were measured using a 7-point Likert scale questionnaire both before and after the robot's unexpected behavior. The study found that the trust score decreased when the robot swerved from its expected plan due to its unpredictability. In addition, it was found that participants believed that the robot should be held accountable for the unexpected behavior rather than the human. We expect to see this decrease in trust as well in our study when the robot fails to complete the mission. While the work in [30] treats the robot as fully autonomous, our study aims to explore the repercussions of task efficiency on trust in the case where there is shared control between the human and the robot. Instead of treating the user as an observer, we place control in the user's hands by enabling them to guide the robot's actions, which results in shared blame attribution.

Chapter 3

Collaborative Human-Robot Exploration via Implicit Coordination

This chapter details the collaborative human-robot exploration method. The human and robot incrementally build a shared map of the environment using range measurements, while the robot uses the occupancy, ROI, and distance information within the shared map for motion planning. We first describe the shared map representation.

3.1 Shared Map Representation

The shared map is modeled as a global 3D occupancy grid (OG) map, $\mathbf{m} = \{m_1, ..., m_{|\mathbf{m}|}\}$. Each cell m_i contains a tuple of three scalar features: (1) the probability of occupancy (o_i), (2) a boolean indicating if the cell is in the ROI (b_i), and (3) the distance of the cell from the closest obstacle (d_i). Each cell m_i is initially presumed unknown ($o_i = 0.5$), considered outside the ROI ($b_i = 0$), and assumed to be at an infinite distance from the closest obstacle ($d_i \to \infty$). The range measurements at time t are denoted by $\mathbf{z}_t^h$ for the human and $\mathbf{z}_t^r$ for the robot. It is assumed that the global position and orientation of these sensors are perfectly known.

For the cells within the FoVs of $\mathbf{z}_t^h$ and $\mathbf{z}_t^r$, the probability of occupancy o_i is updated using the standard log-odds update [42]. However, the ROI values b_i are

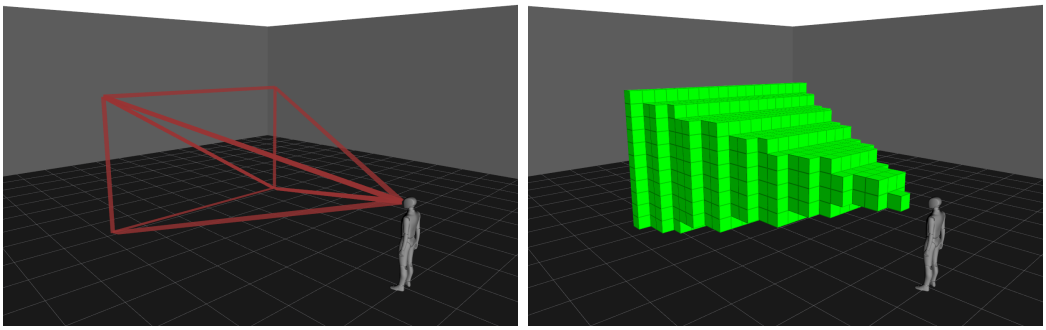

(a) Human's FoV (b) Voxels within the FoV

Figure 3.1: (a) The human's field of view (FoV) is shown in red and used to determine which (b) cells in the global occupancy map are within the ROI (shown in green).

only set to 1 within the FoV of $\mathbf{z}_t^h$. The FoV is mathematically modeled using the fusion of two triangles in 2D and two tetrahedrons in 3D built from the sensor's intrinsic matrix (see fig. 3.1a). A subset of cells corresponding to inliers of the FoV are shown in fig. 3.1b. This subset is extracted via inlier queries with respect to the tetrahedrons on the centers of all cells in the shared map. The distance values d_i are updated for the cells raycasted by both $\mathbf{z}_t^h$ and $\mathbf{z}_t^r$ with the Euclidean distance from the nearest occupied cell. We utilize the approximation by Delmerico et al. [12] that extends the rays behind a hit cell, m_{hit}, and populates the distance value at the current time, d_i^t, for the remaining raycasted cells:

$$d_i^t = \begin{cases} \|f(m_{\text{hit}}) - f(m_i)\|_2, & \text{if } \|f(m_{\text{hit}}) - f(m_i)\|_2 < d_i^{t-1} \\ d_i^{t-1}, & \text{otherwise} \end{cases} \tag{3.1}$$

where d_i^{t-1} is the previously stored distance in cell m_i and $f : \mathbb{Z}_+ \to \mathbb{R}^n$ is a function that converts a cell index to the cell position in the world frame. $n = 2$ or $n = 3$ depending on the dimensionality of the map representation.

After obtaining the first observation from the human and updating the shared map, the robot iteratively performs a two-step process: updating the map and selecting the next best action. The rate of this process is specified by the user prior to operation. The space of candidate actions used for action selection is generated using a library of forward-arc motion primitives for a depth camera as presented in [41]. The best

primitive is chosen by maximizing the information gain over this discrete action space, which is computed at the end viewpoint of each motion primitive. We contribute one information gain objective function (OAVI) and contrast to a baseline information-theoretic objective function (CSQMI) as well as an extension of it (ROI-CSQMI) in this work.

3.2 ROI-constrained CSQMI (ROI-CSQMI)

The ROI-CSQMI information gain objective function modifies the CSQMI objective function proposed by Charrow et al. [8] by imposing a spatial constraint corresponding to the FoV of the range measurement from the human $\mathbf{z}_t^h$. The original formulation proceeds as follows; first, for a candidate viewpoint, a beam-based measurement model is used to determine which cells in the current map will be observed via raycasting. Second, the CSQMI objective is computed treating the raycasted cells independently of each other. Lastly, the CSQMI contributions from all cells are added to determine the total utility of the viewpoint. Our contribution lies in modifying the second step using the ROI information stored in the shared map representation (section 3.1).

Let the set of cells raycasted by the candidate measurement $\hat{\mathbf{z}}_t^r$ be denoted by $\hat{\mathbf{m}}$. This set is a subset of the current map, $\hat{\mathbf{m}} \subseteq \mathbf{m}$. To impose the spatial constraint corresponding to the FoV of $\mathbf{z}_t^h$, we leverage the information in $\mathbf{m}$ to obtain a new set of raycasted cells within the ROI denoted as $\hat{\mathbf{c}}$:

$$\hat{\mathbf{c}} = \{m_i \in \hat{\mathbf{m}} \,|\, b_i = 1, i \in \{1, \ldots, |\hat{\mathbf{m}}|\}\}. \tag{3.2}$$

The ROI constraint is imposed at the raycasting step by finding the subset of cells that lie within the human's ROI. The CSQMI objective for the candidate measurement $\hat{\mathbf{z}}_t^r$, $I_{\mathrm{CS}}[\hat{\mathbf{m}}; \hat{\mathbf{z}}_t^r]$, is computed using $\hat{\mathbf{m}}$ via the equations derived in [8]. Note that when $\hat{\mathbf{c}} = \hat{\mathbf{m}}$, the ROI-CSQMI objective function is equivalent to the CSQMI objective function.

For a 2D environment with an obstacle (fig. 3.2a), the map $\mathbf{m}$ is shown in fig. 3.2b after adding one range measurement $\mathbf{z}^h$ from the human who is located at the bottom center of the environment. The grey, black, and white cell colors denote unknown, occupied, and free state, respectively. Figure 3.2c shows the heatmap for

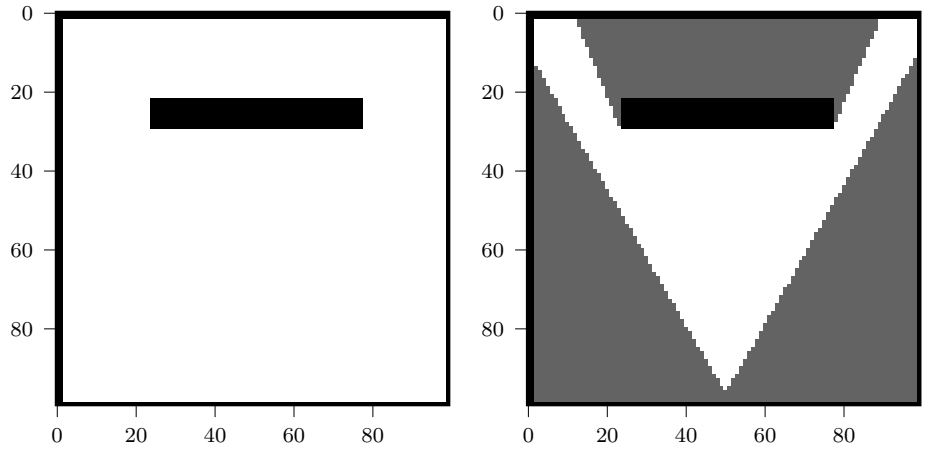

(a) Environment (b) Map after Human Observation

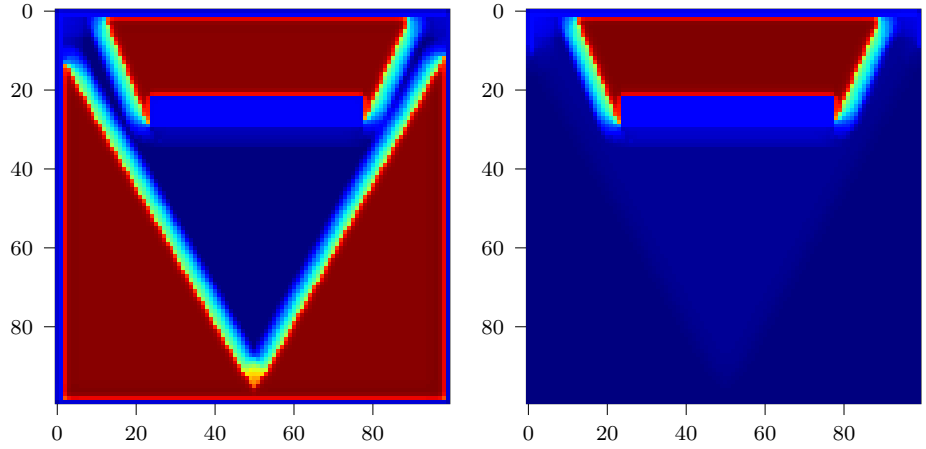

(c) CSQMI (d) ROI-CSQMI

Figure 3.2: Comparison of the information gain objectives using a 2D numerical example. For the environment in (a) and human at (50, 0), the map updated after one $\mathbf{z}^h$ (section 3.1) is shown in (b). The CSQMI objective from [8] shown in (c) does not account for $\mathbf{z}^h$, while the ROI-CSQMI objective in (d) places higher weights in the occluded region.

$I_{\text{CS}}[\hat{\mathbf{m}}; \hat{\mathbf{z}}_t^r]$, which treats all unknown space equally. Figure 3.2d shows the heatmap for $I_{\text{CS}}[\hat{\mathbf{c}}; \hat{\mathbf{z}}_t^r]$, which prioritizes regions within the FoV of $\mathbf{z}^h$. The spatially constrained ROI-CSQMI objective enables implicit coordination between the human and robot during exploration by placing higher weight on views that intersect the ROI. However, there are two drawbacks to this modification: (1) the objective weighs all cells within the occluded region equally, as opposed to the regions closer to the obstacle within the human's FoV, and (2) once the robot enters the ROI during exploration, it is unlikely that it will exit it. The OAVI objective, presented next, alleviates these drawbacks.

3.3 Occlusion-Aware Volumetric Information (OAVI)

The proposed information-gain objective function, OAVI, is inspired by [12] and modifies the uncertainty-aware, I_{UA}, the ROI, I_{ROI}, and the proximity-aware, I_{PA} metrics.

The uncertainty-aware metric I_{UA} measures the uncertainty of the cell and accounts for potential occlusions:

$$I_{\text{UA}}(m_i) = H(m_i) P_{\text{V}}(m_i). \tag{3.3}$$

$H(m_i)$ is Shannon's entropy [10] of cell m_i, and $P_{\text{V}}(m_i)$ is the likelihood that the cell is visible from the current sensor pose. The result is shown in fig. 3.3a for the 2D map in fig. 3.2b and illustrates high weights in the unknown space.

The ROI metric, I_{ROI}, biases the objective values towards the ROI. We employ the information stored in the shared map (section 3.1) to mark the contribution of cells in the ROI towards I_{ROI} as 1. For the other regions of the map, the contribution is set to a user-specified value, $\alpha_{\text{ROI}} < 1$:

$$I_{\text{ROI}}(m_i) = \begin{cases} 1, & \text{if } b_i = 1 \\ \alpha_{\text{ROI}}, & \text{otherwise} \end{cases} \tag{3.4}$$

Intuitively, α_{ROI} controls the weight given to the unknown regions of the environment

outside the ROI. A non-zero α_{ROI} enables the robot to explore unknown regions after prioritizing occluded regions within the ROI. This metric is shown in fig. 3.3b with $\alpha_{\text{ROI}} = 0.15$ for the 2D map in fig. 3.2b.

The same modification is applied for the proximity-aware metric, I_{PA}, which utilizes the distance values, d_i, in the cells:

$$I_{\text{PA}}(m_i) = \begin{cases} d_{\max} - d_i, & \text{if } o_i = 0.5 \text{ and } d_i \leq d_{\max} \\ \alpha_{\text{PA}}, & \text{otherwise} \end{cases} \tag{3.5}$$

where d_{max} is the max sensor range, d_i is the distance from cell m_i to the closest raytraced occupied cell, and $\alpha_{\text{PA}} \in [0, 1)$ is a tunable parameter. This modification produces a gradient that places higher weights on cells close to an observed surface (e.g. see fig. 3.3c where $\alpha_{\text{PA}} = 0.10$).

The final information gain I_{OAVI} is defined as the cumulative product of each metric over the raycasted cells $\hat{\mathbf{m}}$ corresponding to the robot measurement at the viewpoint $\hat{\mathbf{z}}_t^r$:

$$\begin{aligned} I_{\text{OAVI}}[\hat{\mathbf{m}}; \hat{\mathbf{z}}_t^r] &= \sum_{i \in [1, |\hat{\mathbf{m}}|]} I_{\text{OAVI}}(\hat{m}_i) \\ &= \sum_{i \in [1, |\hat{\mathbf{m}}|]} I_{\text{UA}}(\hat{m}_i) I_{\text{ROI}}(\hat{m}_i) I_{\text{PA}}(\hat{m}_i). \end{aligned} \tag{3.6}$$

Figure 3.3d illustrates the heatmap corresponding to $I_{\text{OAVI}}[\hat{\mathbf{m}}; \hat{\mathbf{z}}_t^r]$. Note the gradient behind the obstacle in OAVI, which has the effect of weighting the viewpoints that observe occluded regions more heavily, and contrast this with the uniform weighting of ROI-CSQMI in fig. 3.2d.

3.4 Experimental Design and Results

The approach is evaluated in simulation and with real-world hardware experiments. The experiment begins when the human transmits the pose of their helmet-mounted range sensor with the corresponding pointcloud to the robot partner. Only one instance of these pose and pointcloud pairs is transmitted for both simulation and hardware experiments. The proposed methodology may allow for multiple pose and

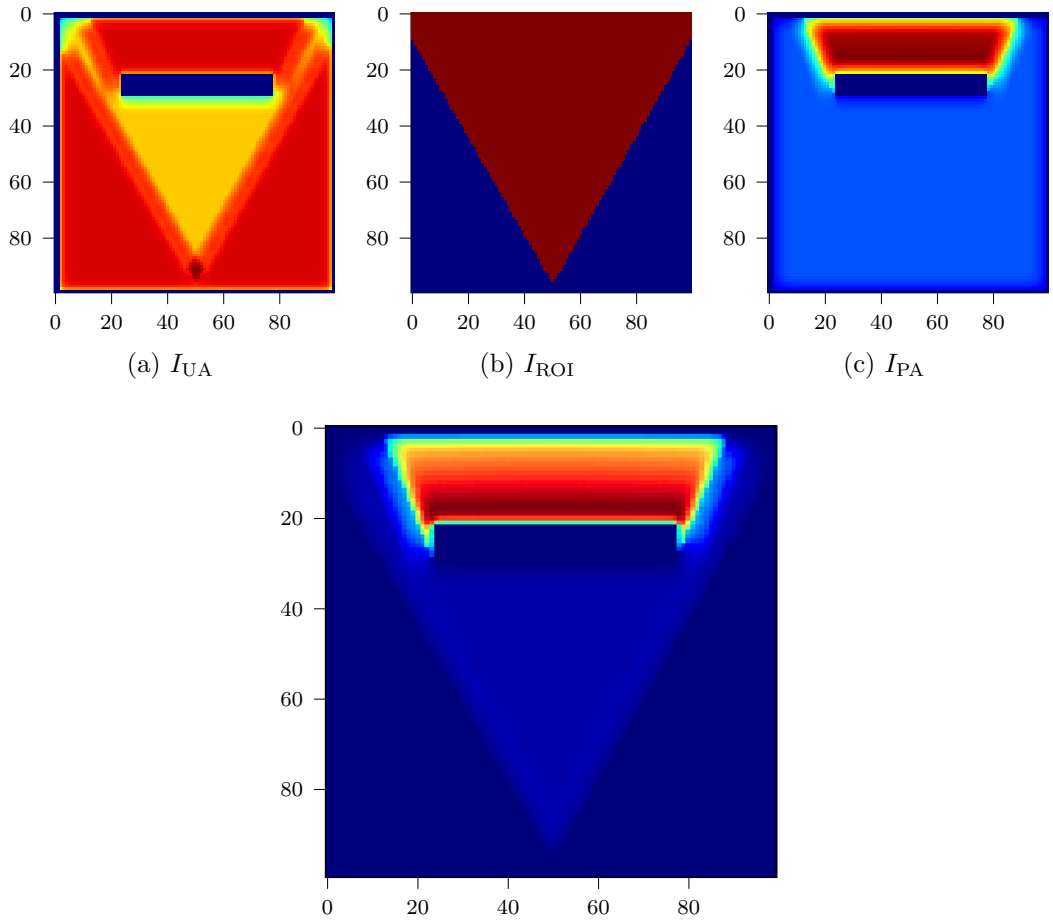

(a) I_{UA} (b) I_{ROI} (c) I_{PA}

(d) OAVI

Figure 3.3: Heatmaps for the OAVI objective and its constituent terms (section 3.3) over the 2D map shown in fig. 3.2b. Compared to ROI-CSQMI in fig. 3.2d, the OAVI objective function in fig. 3.3d exhibits a gradient biasing the exploration to focus on the occluded region closer to the human's FoV first.

pointcloud pairs, but this is left as future work. When the robot receives data from the human, exploration begins.

The OAVI approach is compared against the ROI-CSQMI and CSQMI approaches. Two quantitative and one qualitative measures are used to evaluate performance. The two quantitative evaluations measure the entropy of the map and ROI over time. The qualitative evaluation plots the evolution of the robot's trajectory over time.

Parameter	**Simulation**	**Hardware**
Robot sensor range	5 m	2 m
Robot sensor downsampling	2×	2×
Human sensor range	10 m	6 m
Human sensor downsampling	4×	4×
Human FoV percentage	40%	30%
Mapping frequency	10 Hz	10 Hz
Voxel resolution	0.3 m	0.2 m
Grid bounding box	30 × 30× 10 m	4 × 5× 2 m
Planning frequency	1 Hz	1 Hz
Number of motion primitives	21	15
Max. forward velocity	$0.75\,\mathrm{m\,s^{-1}}$	$0.40\,\mathrm{m\,s^{-1}}$
Max. yaw rate	$0.25\,\mathrm{rad\,s^{-1}}$	$0.25\,\mathrm{rad\,s^{-1}}$
OAVI α_{ROI}	0.10	0.10
OAVI α_{PA}	0.15	0.15

Table 3.1: List of parameter values used in the simulation and hardware experiments.

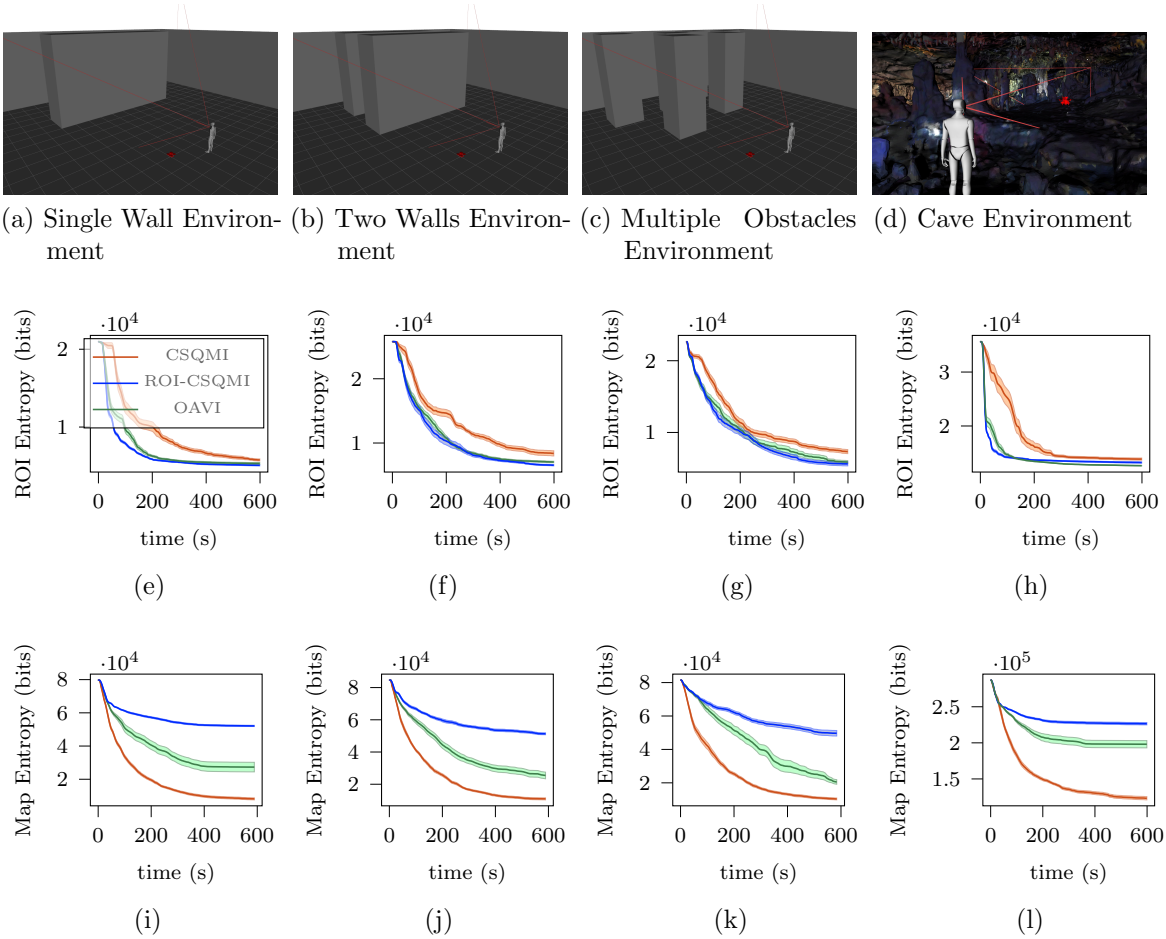

Figure 3.4: (a)–(d) simulation environments, (e)–(h) ROI entropy plotted as a function of time and (i)–(l) map entropy plotted as a function of time for the CSQMI, ROI-CSQMI, and OAVI exploration variants. 30 trials are run for each exploration variant and simulation environment. Note that ROI-CSQMI and OAVI explore the human's FoV 3× faster than CSQMI while CSQMI reduces the total map uncertainty faster. OAVI reduces the map uncertainty 56% more than ROI-CSQMI.

3.4.1 Simulation Experiments

Simulations in four environments (fig. 3.4) are conducted to evaluate the approaches developed in this work against the baseline approach. The simulation environments consist of a single wall, two walls, multiple obstacles, and the cave environment from fig. 1.2 with the same human-robot placement. In each environment the goal is for the robot to obtain views in regions occluded to the human. These environments are selected to highlight the merits and drawbacks of the information gain objectives.

In each environment, the human faces the area of interest. The human's FoV, which is the FoV of a simulated depth camera on the human's head, is shown as red lines in fig. 3.1a. The robot is placed at a randomly selected location within a $4 \times 4\,\text{m}$ box around the human's starting position. After the human transmits their pointcloud observation and pose to the robot, the robot updates its onboard map according to section 3.3. Each exploration variant is run for 30 trials per environment for a total of 360 trials over all environments and variants. Exploration is terminated after 10 min resulting in a total of 60 h of simulations.

The entropy of the ROI is plotted over time for each environment in figs. 3.4e to 3.4h and the entropy of the map (including the ROI) is plotted over time in figs. 3.4i to 3.4l. In analyzing the performance of figs. 3.4e to 3.4h one can see that OAVI and ROI-CSQMI decrease the uncertainty of the ROI approximately $3\times$ faster than CSQMI. ROI-CSQMI slightly outperforms OAVI because the mutual information of a view entirely outside of the ROI is zero, which means that the robot will not select actions outside the ROI. In contrast, OAVI tends to drive the robot outside the ROI after sufficient views of the ROI have been acquired. When analyzing the map entropy over time in figs. 3.4i to 3.4l, one can see that the final map entropy of the OAVI approach at 600 s is on average 56% lower than the ROI-CSQMI approach across environments. The baseline CSQMI approach outperforms the other approaches because it selects views that maximize the mutual information between the map and sensor without consideration for the ROI.

From these results, we arrive at the following conclusions: first, the baseline CSQMI approach is not well suited for collaborative human-robot exploration because it does not bias the exploration towards the ROI; second, the ROI-CSQMI approach is ideal for a leader-follower exploration strategy because it selects motion plans

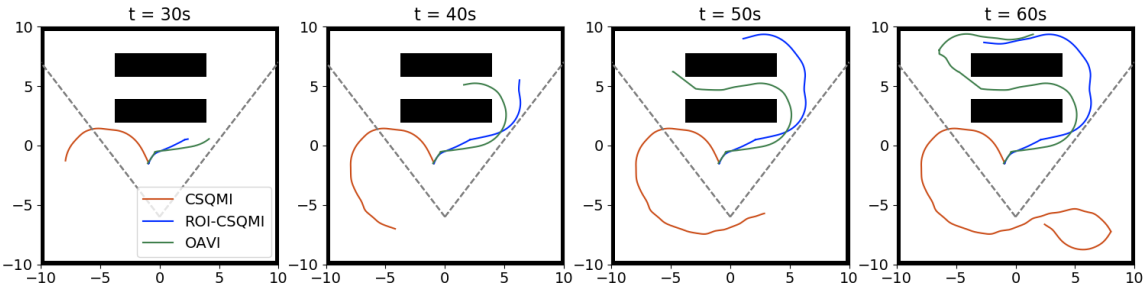

Figure 3.5: Top-down snapshots of the trajectory taken by the robot for the three approaches in the two walls environment with the human's FoV drawn in gray dashed lines. CSQMI proceeds to explore the unknown regions outside of the human's FoV, while the ROI-constrained CSQMI and OAVI prioritize the ROI first. As opposed to ROI-CSQMI, the gradient in the OAVI approach (see fig. 3.3d) pushes the robot to explore the occluded region closest to the human first.

that are restricted within the ROI; and third, the OAVI approach is ideal for a collaborative framework where the robot biases views within the ROI first and then selects observations outside the ROI.

Figure 3.6: (Left) Aerial robot and (Right) helmet for the human partner used in the hardware experiments.

Figure 3.5 plots the top-down views of the trajectories taken by the robot for the three approaches in the two walls environment (shown in fig. 3.4b). The evolution of the trajectory for $t = \{30, ..., 60\}$s in fig. 3.5 demonstrates that OAVI first explores the occluded region closest to the human observer and then proceeds to the second

obstacle after the robot has updated its distance field. In comparison, ROI-CSQMI does not incorporate a measure of the distance to obstacles so it selects actions that maximize the mutual information between the potential observation and ROI, while CSQMI explores areas outside of the human's ROI for the first 60 s because it does not have a notion of the human's ROI.

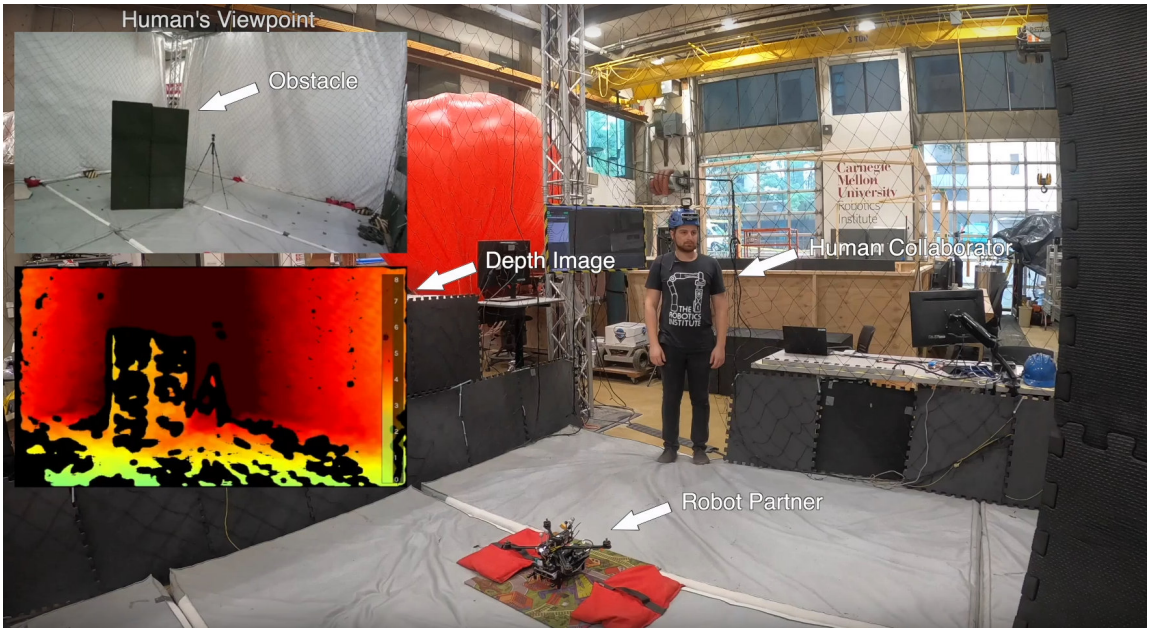

Figure 3.7: A human-robot team explores an environment inside a motion capture arena, with an obstacle in front the human requiring the robot to provide complementary views.

3.4.2 Hardware Experiments

Experiments are run inside a motion capture arena to validate the proposed approach against the baselines in the real-world. The human is equipped with a helmet-mounted Intel RealSense D455 depth camera (see fig. 3.6). The robot, of dimensions 0.6m x 0.3m x 0.3m and a mass of 2.5kg, is equipped with a downward-facing mvBlueFox-MLC200wC color global shutter camera and Sunex DSL219D-650-F2.0 lens to estimate state. A forward-facing Intel Realsense D455 is used to estimate depth. Cree Xlamp XM-L2 High Power LEDs (Cool White 6500K) are mounted to the front and underneath the system to provide illumination for color camera images.

The mvBlueFox operates at 60Hz and the Realsense D455 operates at 15Hz. The laser power of the IR projector on the D455 is increased from the default 150mW to 300mW to increase the depth accuracy in darkness. The vehicle is equipped with two computers that communicate via Ethernet link: a Gigabyte Brix 8550U and an NVidia TX2 with a J120A Auvidea carrier board. The TX2 runs the mvBlueFox driver, communication to and from the Pixracer, state estimation, and control. The Gigabyte Brix runs mapping, collision avoidance, motion planning, and the Intel Realsense D455 driver. The drone frame is a Lumenier QAV-RXL 10" FPV quadcopter frame. The ESC is an Aikon AK32 55A 3-6S BLHeli32 4-in-1 ESC and the motors are T-Motor F100 Cinematic KV1100 with 9" Advanced Precision Composites propellers. Teleoperation commands are sent via an RC Transmitter to a TBS Crossfire Micro RX V2 receiver.

The set of parameters used in the hardware experiments (see fig. 3.7) are listed in table 3.1. Each approach is run once starting from the same initial robot pose and a fixed helmet orientation for a total time of 2 min.

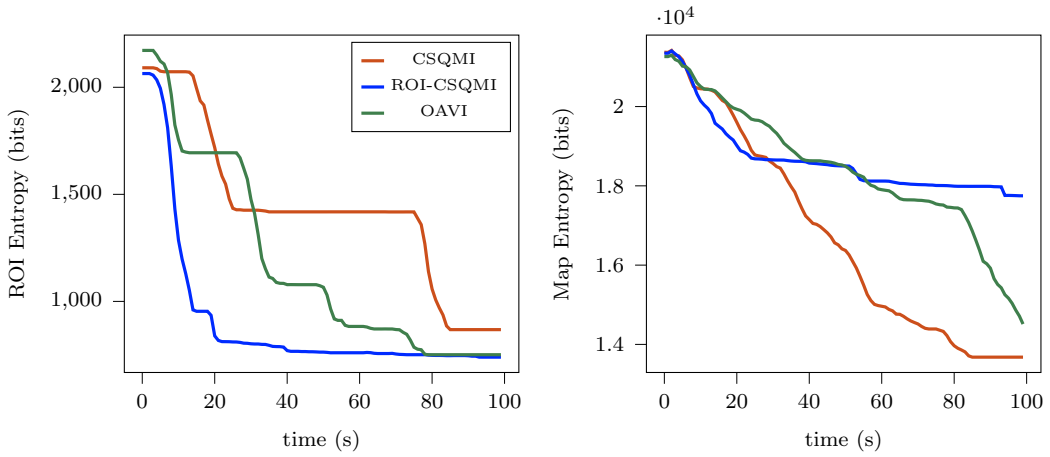

(a) ROI Entropy

(b) Map Entropy

Figure 3.8: ROI and map entropy as a function of time for the three approaches. The baseline CSQMI approach minimizes the total map entropy, while its extension ROI-CSQMI prioritizes the ROI. OAVI successfully reduces the uncertainty in the ROI first, followed by an exploratory behavior. A video of the experimental setup and the three exploration approaches can be found at https://youtu.be/7jgkBpVFIoE.

The ROI and map entropy are plotted over time in fig. 3.8. The baseline CSQMI

approach reduces the total uncertainty in the environment fastest fig. 3.8b. ROI-CSQMI explores the ROI twice as fast as OAVI, and 4× as fast as CSQMI (see fig. 3.8a) but does not select actions outside of the ROI once it reaches the ROI. This behavior yields the blue plateau in the map at $t = 50\,\text{s}$. In contrast, OAVI explores the rest of the unknown environment as shown in the final map (fig. 3.9), reducing the map entropy by 75% more than ROI-CSQMI.

Approach	**Planning Time**
CSQMI	0.027 ± 0.02s
ROI-CSQMI	0.022 ± 0.02s
OAVI	0.036 ± 0.03s

Table 3.2: Planning times onboard the robot's computer during hardware experiments show the computational-efficiency of the proposed approach.

To demonstrate computational-efficiency of the proposed approach, we record the planning times taken by the action generation, scoring, and best primitive selection modules onboard the robot's computer. The results in table 3.2 show close planning times between the three approaches, allowing our planner to run at up 15 Hz.

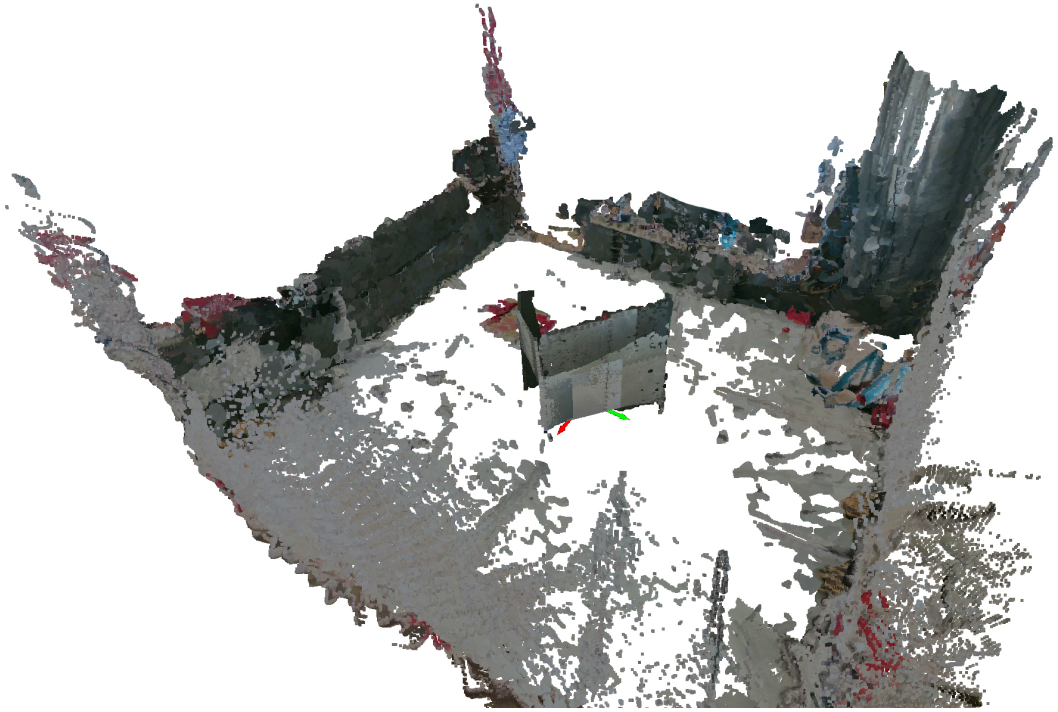

Figure 3.9: Reconstructed point cloud map of the $50\,\mathrm{m}^3$ environment from the OAVI hardware trial.

Chapter 4

Human Preferences for Robot Autonomy in a Search Task

This chapter explores the human perception of robot efficiency in a search task to better understand how to design and develop robots that explore alongside human partners. We detail the user study where participants were tasked with guiding a robot to find a hidden object in a simulated environment.

4.1 Methodology

This section details the design of the user study conducted to evaluate the effect of task efficiency on human trust in a collaborative human-robot search task. We first describe the experimental setup design.

4.1.1 Experimental Setup

Each participant is tasked with finding an object inside a simulation environment as soon as possible. The object is hidden behind one of three obstacles presented to the user, to which they do not have access, as seen in fig. 4.1. This is motivated by a scenario in which a first responder (e.g. firefighter) is unable to access a location due to some hazard (e.g. room on fire). Assisting each user is an aerial robot partner. The user is able to guide the robot to explore a particular area of interest

by employing our previously developed Occlusion-Aware Volumetric Information (OAVI) exploration policy [11]. Participants have control over a first-person simulated human on a computer where they are able to rotate and change their point of view (POV). By looking at one of the three obstacles and pressing a button, participants are able to guide the robot's exploration into their POV visualized in fig. 4.2. The simulation is developed in C++ and Python within the Robotic Operating System (ROS) framework and RViz.

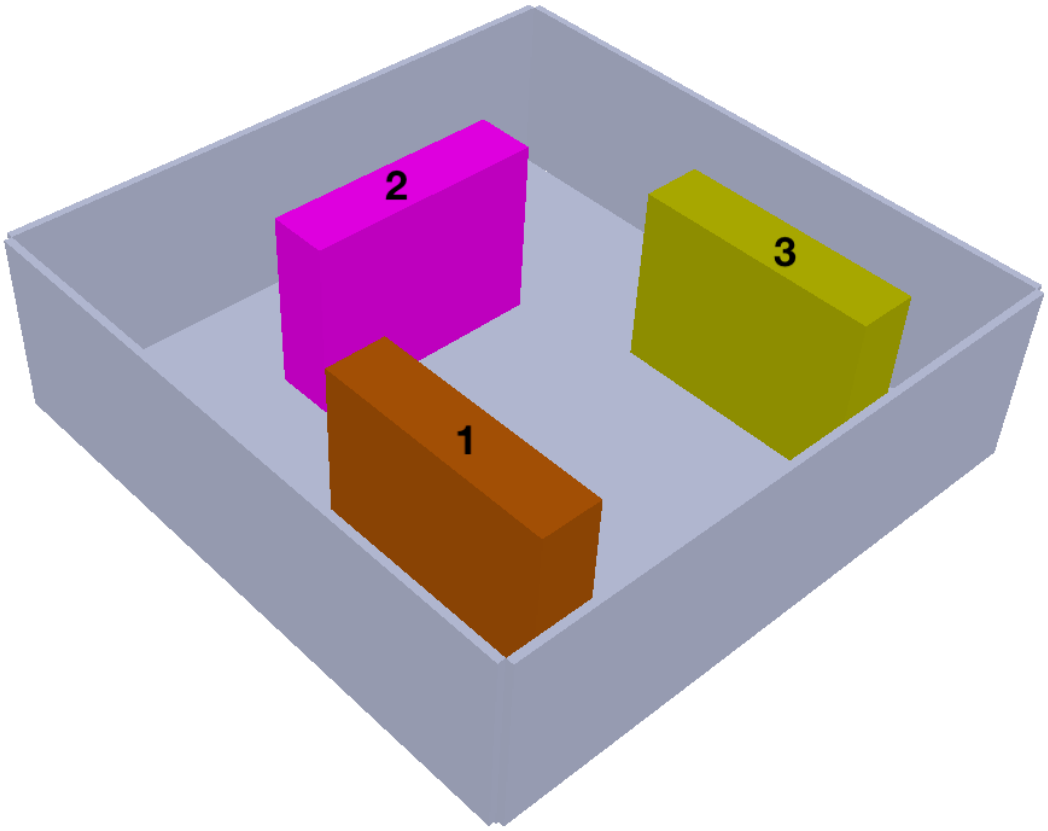

Figure 4.1: Simulation environment with the three obstacles behind which the object will be randomly placed. Users do not have access to the occluded regions beyond the obstacles.

The study is a within-subjects experiment where each participant conducts two experiments. In the first experiment, the object is placed behind one of the obstacles

and the robot is able to find the object within the time limit of 3 minutes. As soon as the robot finds the object, the mission time is displayed on the screen and the trial is ended. During the second experiment, the object is removed from the environment without telling the participant. This results in the robot never finding the object within the time limit and the task failing. The order of the experiments is shuffled between participants such that the expectation bias is minimized.

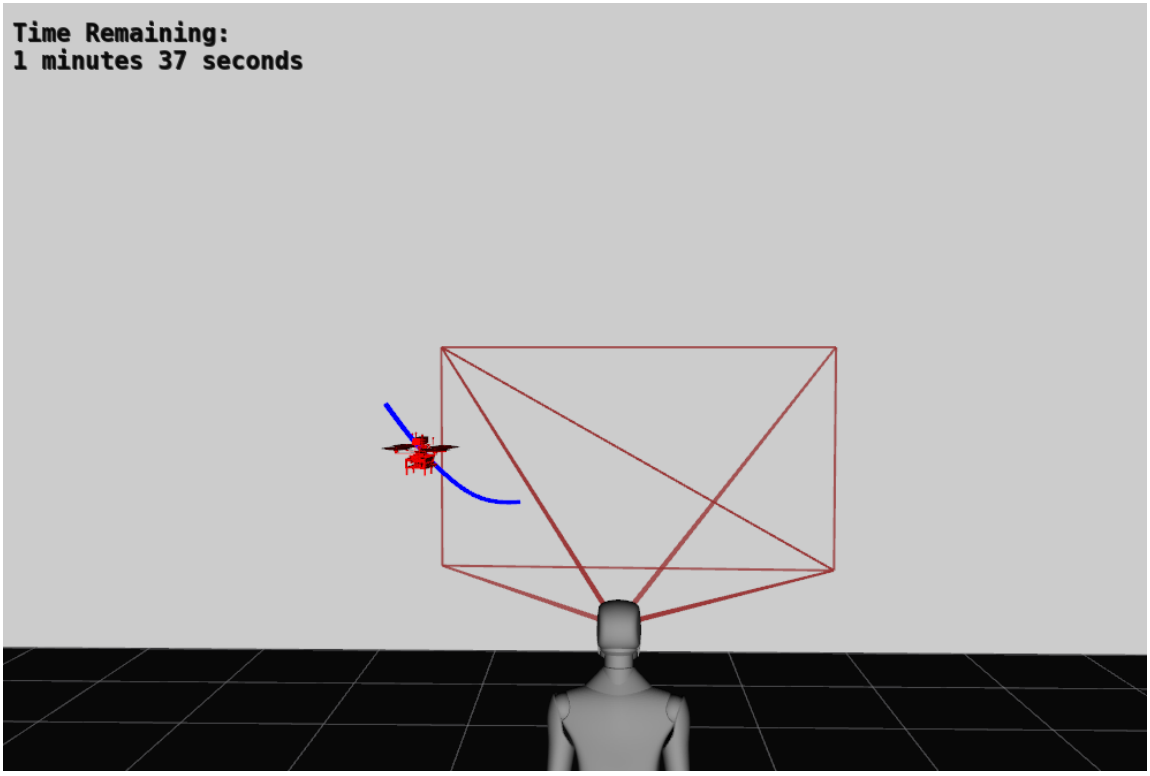

Figure 4.2: User's perspective modeled by two tetrahedrons built from the intrinsic matrix of a pinpoint camera model. The aerial robot can be seen exploring within the POV, while the mission timer is displayed on the top-left of the GUI.

4.1.2 Measures

In order to measure trust as a function of task efficiency participants are presented with a set of subjective questions. After each experiment, the user will be asked to rate the statements in table 4.1 on a 5-point Likert scale, with 1 signifying a low score and 5 signifying the highest score a user could give. The choice of questions is

motivated by studies on perceived behavioral control over another agent [46], measures of perceived trust and acceptance in social robots [22], and literature on trust in automation adapted to fit the human-robot collaboration context [23, 29].

	Question	**Min. score = 1**	**Max. score = 5**
Q1	How satisfied were you with the robot's performance?	Not satisfied	Very satisfied
Q2	How much do you trust the robot?	I don't trust the robot	I trust the robot
Q3	Were you frustrated with the robot's behavior?	Very frustrated	Not frustrated at all
Q4	Would you use the robot again to assist you in the search?	Never	Absolutely
Q5	Who do you think had more control you or the robot?	The robot had more control	I had more control
Q6	Did the robot's actions make sense to you?	They didn't make any sense	The actions made complete sense

Table 4.1: Post-experiment questionnaire that participants were asked to fill, with the 5-point Likert scale extremities labels.

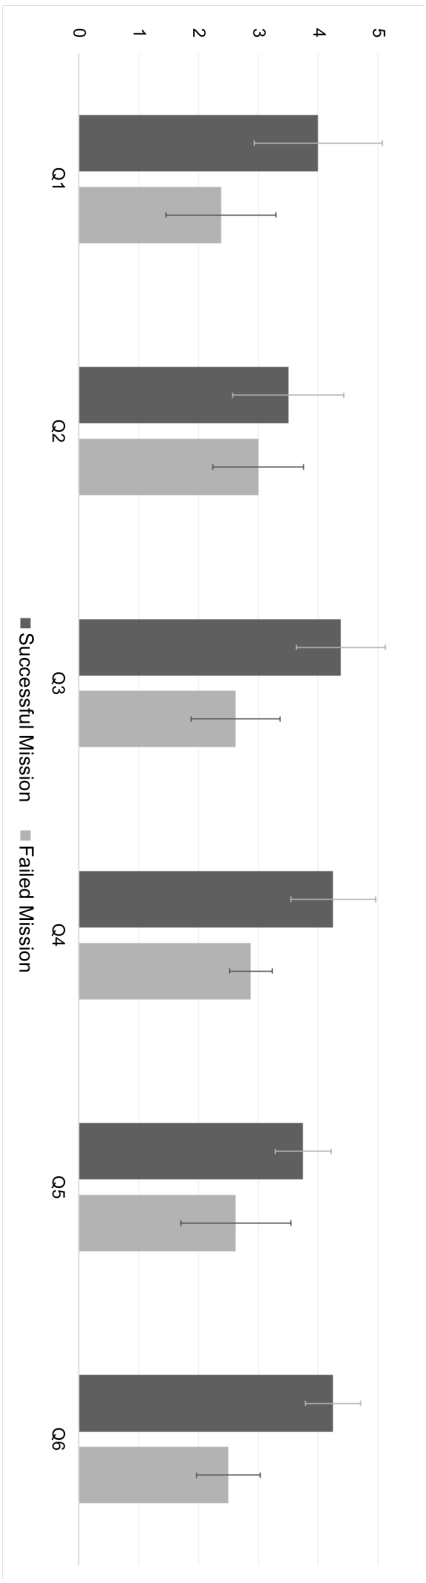

Figure 4.3: Results of the questionnaire for the experiment where the robot was successful in finding the object and the experiment where the mission failed. The data is shown for each individual question with the standard deviation visualized as a capped two-sided line.

4.1.3 Hypotheses

- **H1** Trust in the robot that did not find the object will be lower than that of the robot that successfully completed the mission.
- **H2** Participants perceived that they had more control over the robot when the mission was successful, compared to a belief that the robot was more in control when the mission failed.
- **H3** Trust is higher when the robot's actions made sense to the participants, while it is lower when the robot acted contrary to their intuition.

4.2 Results

This section presents the participants' demographics, results, and analysis for the data collected during the user study.

4.2.1 Participants

The study was conducted with 8 participants, mostly graduate students, from the Robotics Institute at Carnegie Mellon University. Half of the participants self-identified as male and the other half self-identified as female. Participants were first briefed about the study and signed an informed consent form as part of the university's Institutional Review Board (IRB) protocol and regulations. The participants were then given the chance to get accustomed to the simulation environment and control of the robot. The results of the post-study questionnaires are presented in fig. 4.3 for the two experiments where in the successful case the robot found the object, and in the failure case the robot was unable to find to object within the time limit.

4.2.2 Trust Measure

In order to obtain less noisy results, we average the first four questions embodying trust into a single trust measure. The resulting bar plot is shown in fig. 4.4. Cronbach's alpha is used to test for internal consistency of the trust questionnaire resulting in $\alpha = 0.8049$, indicating good internal consistency and the reliability of the questionnaire

averaging as a measure of trust.

Figure 4.4: Trust measure used for this study as the average of the results from questions 1 through 5 in table 4.1.

4.2.3 Analysis of Hypothesis H1: Trust and Task Efficiency

To investigate the relationship between the dependent variable (DV) of trust and the independent variable (IV) of task efficiency, we ran an independent samples t-Test on the trust measure from fig. 4.4. Consistent with our hypothesis **H1**, participants rated their trust in the robot higher when the robot succeeded at the mission ($M = 4.03$, $SD = 0.59$) compared to when the robot failed to complete the mission ($M = 2.72$, $SD = 0.43$), $t(7) = 5.69$, and $p = 0.0007$. The results are intuitive; when the mission fails, participants are dissatisfied with the robot's performance (Q1), report high frustration with its behavior (Q3), and lower intent to use the robot again for assistance in a search task (Q4) according to results reported in fig. 4.3. This is in line with the observation from several other studies, including the work by Khavas

et al. [26], which states that there exists a strong correlation between the performance of a robotic agent and the level of trust in human-robot teammates.

4.2.4 Analysis of Hypothesis H2: Control and Task Efficiency

To examine the interrelation between the perception of control DV and task efficiency IV, we ran an independent samples t-test on the data from the fifth question representing the user's perception of the level of autonomy in fig. 4.3. Participants perceived that they had more control over the robot's planner when the mission was successful ($M = 3.75$, $SD = 0.46$), and believed that the robot was more in control in the case where the team did not find the object resulting in mission failure ($M = 2.63$, $SD = 0.92$), demonstrating statistical significance ($t(7) = 2.55$, $p = 0.038$) in support of **H2**. The two experiments were conducted with a robot employing the same exploration policy and consequently producing the same behavior. Similar lines of research such as the work by Kim and Hinds [27] have found that people will attribute more blame to an autonomous robot and less to themselves and their co-workers. In accordance with what our findings suggest as well, they remark that users do not give credit to the robot when a mission was successful.

4.2.5 Analysis of Hypothesis H3: Trust and Transparency

To test **H3**, a Pearson's correlation coefficient was computed to assess the linear relationship between the two DVs transparency (Q6 in table 4.1) and trust measure from fig. 4.4. There was a positive correlation between the two variables ($r(15) = 0.77$, $p = 0.0004$) as seen in fig. 4.5, validating that intuitive robot behavior is akin to higher perceived trust from their human partner. While the robot's exploration policy, and accordingly its actions, did not change between the two trials, participants trusted the robot whose actions were regarded as rational more than the robot whose actions seemed confusing to them. Other studies on trust and transparency have reached similar conclusions, including the work by Nesset et al. [32] which finds that high transparency in a robot system can help users calibrate their trust levels towards it. Particularly in the field of aerial robotics within HRI, Okamura and Yamada [35]

emphasize the importance of maintaining system transparency for proper user trust in the autonomous drone. This evidence supports the hypothesis **H3**.

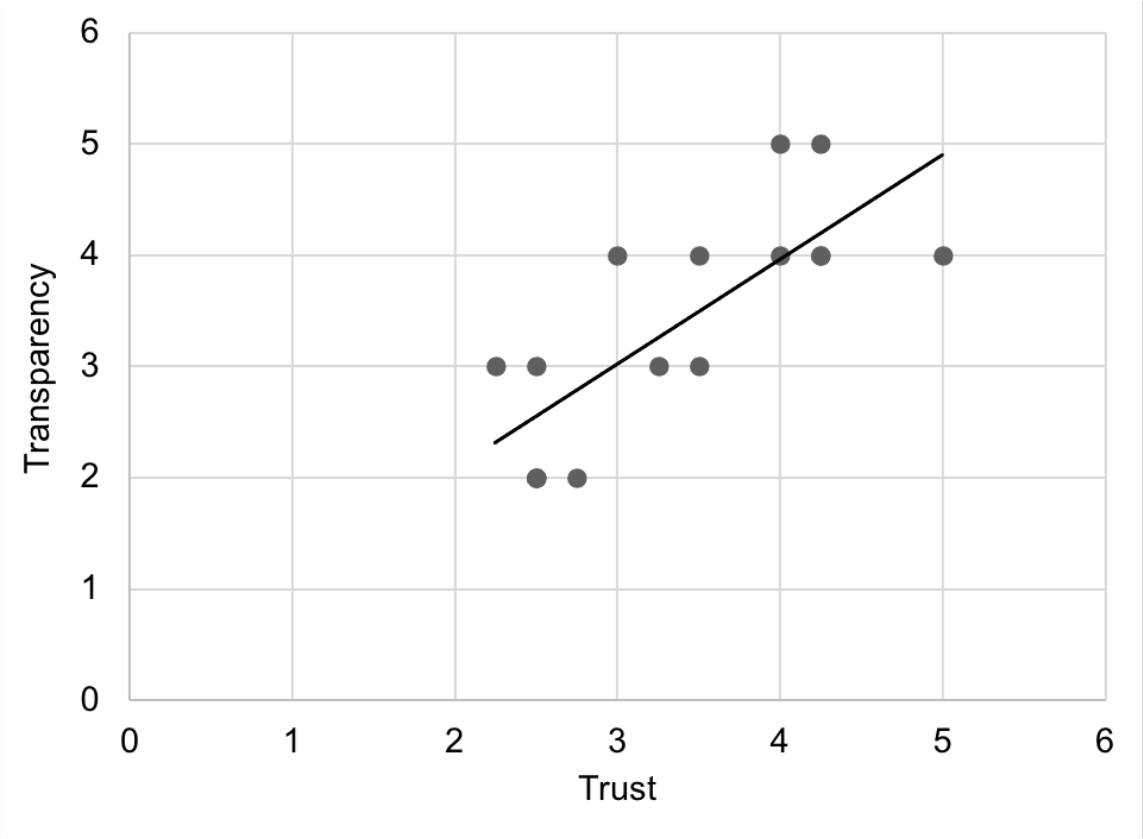

Figure 4.5: Scatter plot of the collected data from the trust measure in fig. 4.4 against the last survey question measuring transparency, with a line fit showcasing the positive correlation between the two variables.

Chapter 5

Conclusion

This work presented a methodology for collaborative human-robot exploration with implicit coordination. The approach developed in this work, OAVI, is an information-gain objective function inspired by active reconstruction techniques. The proposed approach was compared against an information-theoretic exploration baseline, CSQMI, and an extension to this baseline, ROI-CSQMI, which applies a spatial constraint to bias actions within the human's FoV. Comparing these approaches in simulation and hardware yields the following conclusions: (1) the baseline CSQMI approach is not well-suited to the collaboration paradigm detailed in this paper because it has no notion of the human's ROI and cannot bias motion plans to reduce uncertainty towards the human's FoV; (2) the ROI-CSQMI approach is ideal for a leader-follower exploration strategy because it selects motion plans that are restricted within the ROI; and (3) the OAVI approach is ideal for the collaborative human-robot exploration paradigm outlined in this paper because it causes the robot to select views within the ROI first and then explore outside the ROI when a sufficient number of views within the ROI have been collected.

Additionally, we investigated the relationship between trust and task efficiency in a collaborative human-robot search task. Given the complexity of SAR missions, understanding the interaction between a human and a robot teammate is crucial for a successful endeavor. To that extent, we conducted a user study in simulation where participants are tasked with finding an object of interest within the environment with the assistance of an aerial robot. The study was structured so that the participants

succeeded once and failed the other time at detecting the hidden object within the time limit. Participants completed a questionnaire at the end of the experiment. The data collected from the questionnaires were used to quantify how trust, transparency, perception of control, and task efficiency are interrelated. The data collected during the user study revealed three key findings for human-robot collaboration with shared control. First, trust was found to be higher for a robot that successfully accomplished the mission, as compared with lower trust in a robot that failed in finding the object. Therefore, task efficiency is an important factor to optimize for in a human-robot team. Second, participants felt more in control of the robot's actions when the team succeeded at finding the object and perceived the robot as having more control when the team failed at the mission. These findings highlight the relationship between blame attribution and perception of control; a human will perceive the robot as having control and therefore place blame on the robot when a mission fails, even when the robot's performance doesn't change. Third, users trusted a robot whose actions made sense to them more than a robot that they did not understand. Thus, robots' behaviors should be designed to be intuitive and transparent for acceptance and trust in a human-robot collaborative system.

In future work, we aim to deploy the exploration system for longer durations and with a moving human collaborator in outdoor, field environments. Further, we will relax assumptions on perfect knowledge of human-robot poses and their relative transforms in the world frame. One of the limitations of the user study is the participants were all from the Robotics Institute, which means they all have prior exposure to robots. In addition, trust was only measured post-experiments after the mission time ran out. In the future, tracking trust over time would provide a better understanding of the human's perception of the robot in relation to task efficiency and system transparency.

Bibliography

[1] Marco Aggravi, Giuseppe Sirignano, Paolo Robuffo Giordano, and Claudio Pacchierotti. Decentralized control of a heterogeneous human-robot team for exploration and patrolling. *IEEE Transactions on Automation Science and Engineering*, pages 1–17, 2021. doi: 10.1109/TASE.2021.3106386. 2

[2] Jimmy Baraglia, Maya Cakmak, Yukie Nagai, Rajesh Rao, and Minoru Asada. Initiative in robot assistance during collaborative task execution. In *2016 11th ACM/IEEE international conference on human-robot interaction (HRI)*, pages 67–74. IEEE, 2016. 2

[3] William Bentz, Sahib Dhanjal, and Dimitra Panagou. Unsupervised learning of assistive camera views by an aerial co-robot in augmented reality multitasking environments. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3003–3009. IEEE, 2019. 2

[4] Frederic Bourgault, Alexei A Makarenko, Stefan B Williams, Ben Grocholsky, and Hugh F Durrant-Whyte. Information based adaptive robotic exploration. In *2002 IEEE/RSJ IROS*, volume 1, pages 540–545. IEEE, 2002. 2

[5] Wolfram Burgard, Mark Moors, Cyrill Stachniss, and Frank E Schneider. Coordinated multi-robot exploration. *IEEE Transactions on robotics*, 21(3):376–386, 2005. 2

[6] Alain Caltieri and Francesco Amigoni. High-level commands in human-robot interaction for search and rescue. In *Robot Soccer World Cup*, pages 480–491. Springer, 2013. 2

[7] Jennifer Casper and Robin R. Murphy. Human-robot interactions during the robot-assisted urban search and rescue response at the world trade center. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 33(3): 367–385, 2003. 1

[8] Benjamin Charrow, Sikang Liu, Vijay Kumar, and Nathan Michael. Information-theoretic mapping using cauchy-schwarz quadratic mutual information. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4791–4798. IEEE, 2015. (document), 2, 3.2, 3.2, 3.2

[9] Fang Chen, Jianlong Zhou, Yang Wang, Kun Yu, Syed Z Arshad, Ahmad Khawaji, and Dan Conway. Trust and cognitive load. In *Robust Multimodal Cognitive Load Measurement*, pages 195–214. Springer, 2016. 1

[10] Thomas Cover. *Elements of information theory*. John Wiley & Sons, 1999. 3.3

[11] Yves Georgy Daoud, Kshitij Goel, Nathan Michael, and Wennie Tabib. Collaborative human-robot exploration via implicit coordination. *arXiv preprint arXiv:2209.09294*, 2022. 4.1.1

[12] Jeffrey Delmerico, Stefan Isler, Reza Sabzevari, and Davide Scaramuzza. A comparison of volumetric information gain metrics for active 3d object reconstruction. *Autonomous Robots*, 42(2):197–208, 2018. 1, 2, 3.1, 3.3

[13] Mihir Dharmadhikari, Huan Nguyen, Frank Mascarich, Nikhil Khedekar, and Kostas Alexis. Autonomous cave exploration using aerial robots. In *2021 International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 942–949. IEEE, 2021. 2

[14] José Enrique Domínguez-Vidal, Iván J Torres-Rodríguez, Anaís Garrell, and Alberto Sanfeliu. User-friendly smartphone interface to share knowledge in human-robot collaborative search tasks. In *2021 30th IEEE International Conference on Robot & Human Interactive Communication (RO-MAN)*, pages 913–918. IEEE, 2021. 1

[15] Daniel S Drew. Multi-agent systems for search and rescue applications. *Current Robotics Reports*, 2(2):189–200, 2021. 1

[16] Kshitij Goel, Wennie Tabib, and Nathan Michael. Rapid and high-fidelity subsurface exploration with multiple aerial robots. In Bruno Siciliano, Cecilia Laschi, and Oussama Khatib, editors, *Experimental Robotics*, pages 436–448, Cham, 2021. Springer International Publishing. ISBN 978-3-030-71151-1. 2

[17] Kshitij Goel, Yves Georgy Daoud, Nathan Michael, and Wennie Tabib. Hierarchical collision avoidance for adaptive-speed multirotor teleoperation. *arXiv preprint arXiv:2209.08413*, 2022. 1

[18] Matthew Craig Gombolay, Cindy Huang, and Julie Shah. Coordination of human-robot teaming with human task preferences. In *2015 AAAI Fall Symposium Series*, 2015. 2

[19] Vijay Govindarajan, Subhrajit Bhattacharya, and Vijay Kumar. Human-robot collaborative topological exploration for search and rescue applications. In *DARS*, pages 17–32. Springer, 2016. 1, 2

[20] Boris Gromov, Luca M Gambardella, and Alessandro Giusti. Robot identification and localization with pointing gestures. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3921–3928. IEEE,

2018. 2

[21] Morris Gu, Elizabeth Croft, and Akansel Cosgun. Ar point&click: An interface for setting robot navigation goals. *preprint arXiv:2203.15219*, 2022. 1

[22] Marcel Heerink, Ben Kröse, Vanessa Evers, and Bob Wielinga. Assessing acceptance of assistive social agent technology by older adults: the almere model. *International journal of social robotics*, 2(4):361–375, 2010. 4.1.2

[23] Jiun-Yin Jian, Ann M Bisantz, and Colin G Drury. Foundations for an empirically determined scale of trust in automated systems. *International journal of cognitive ergonomics*, 4(1):53–71, 2000. 4.1.2

[24] Brian J Julian, Sertac Karaman, and Daniela Rus. On mutual information-based control of range sensing robots for mapping applications. *IJRR*, 33(10): 1375–1392, 2014. 2

[25] Shanker Keshavdas and Geert-Jan M Kruijff. Functional mapping for human—robot collaborative exploration. *International Journal of Computers and Applications*, 35(3):125–135, 2013. 1

[26] Zahra Rezaei Khavas, S Reza Ahmadzadeh, and Paul Robinette. Modeling trust in human-robot interaction: A survey. In *International Conference on Social Robotics*, pages 529–541. Springer, 2020. 4.2.3

[27] Taemie Kim and Pamela Hinds. Who should i blame? effects of autonomy and transparency on attributions in human-robot interaction. In *ROMAN 2006-The 15th IEEE international symposium on robot and human interactive communication*, pages 80–85. IEEE, 2006. 4.2.4

[28] Simon Kriegel, Christian Rink, Tim Bodenmüller, and Michael Suppa. Efficient next-best-scan planning for autonomous 3d surface reconstruction of unknown objects. *Journal of Real-Time Image Processing*, 10(4):611–631, 2015. 2

[29] Michael Lewis, Katia Sycara, and Phillip Walker. The role of trust in human-robot interaction. In *Foundations of trusted autonomy*, pages 135–159. Springer, Cham, 2018. 1, 4.1.2

[30] Joseph B Lyons, Izz aldin Hamdan, and Thy Q Vo. Explanations and trust: What happens to trust when a robot partner does something unexpected? *Computers in Human Behavior*, 138:107473, 2023. 2

[31] Marta Makara-Studzińska, Krystyna Golonka, and Bernadetta Izydorczyk. Self-efficacy as a moderator between stress and professional burnout in firefighters. *International journal of environmental research and public health*, 16(2):183, 2019. 1

[32] Birthe Nesset, David A Robb, José Lopes, and Helen Hastie. Transparency in hri: Trust and decision making in the face of robot errors. In *Companion of the*

2021 ACM/IEEE International Conference on Human-Robot Interaction, pages 313–317, 2021. 4.2.5

[33] Yashodhan Nevatia, Todor Stoyanov, Ravi Rathnam, Max Pfingsthorn, Stefan Markov, Rares Ambrus, and Andreas Birk. Augmented autonomy: Improving human-robot team performance in urban search and rescue. In *2008 IEEE/RSJ IROS*, pages 2103–2108. IEEE, 2008. 2

[34] Farzad Niroui, Kaicheng Zhang, Zendai Kashino, and Goldie Nejat. Deep reinforcement learning robot for search and rescue applications: Exploration in unknown cluttered environments. *IEEE Robotics and Automation Letters*, 4(2): 610–617, 2019. 1

[35] Kazuo Okamura and Seiji Yamada. Calibrating trust in human-drone cooperative navigation. In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 1274–1279. IEEE, 2020. 4.2.5

[36] Christopher Reardon, Kevin Lee, John G Rogers, and Jonathan Fink. Communicating via augmented reality for human-robot teaming in field environments. In *2019 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, pages 94–101. IEEE, 2019. 2

[37] S. Scherer, V. Agrawal, G. Best, C. Cao, K. Cujic, and R. Darnley et al. Resilient and modular subterranean exploration with a team of roving and flying robots. *Field Robotics*, 2:678–734, 2022. doi: https://doi.org/10.55417/fr.2022023. 1, 2

[38] Mareike Schüle, Johannes Maria Kraus, Franziska Babel, and Nadine Reißner. Patients' trust in hospital transport robots: Evaluation of the role of user dispositions, anxiety, and robot characteristics. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 246–255. IEEE, 2022. 1

[39] Wennie Tabib, Kshitij Goel, John Yao, Mosam Dabhi, Curtis Boirum, and Nathan Michael. Real-time information-theoretic exploration with gaussian mixture model maps. In *Robotics: Science and Systems*, 2019. 1

[40] Wennie Tabib, Kshitij Goel, John Yao, Curtis Boirum, and Nathan Michael. Autonomous cave surveying with an aerial robot. *IEEE Transactions on Robotics*, 2021. 1.1a

[41] Wennie Tabib, Kshitij Goel, John Yao, Curtis Boirum, and Nathan Michael. Autonomous cave surveying with an aerial robot. *IEEE Transactions on Robotics*, pages 1–17, 2021. doi: 10.1109/TRO.2021.3104459. 2, 3.1

[42] Sebastian Thrun. Probabilistic robotics. *Communications of the ACM*, 45(3): 52–57, 2002. 3.1

[43] M. Tranzatto, F. Mascarich, L. Bernreiter, and C. Godinho et al. Cerberus: Autonomous legged and aerial robotic exploration in the tunnel and urban

circuits of the darpa subterranean challenge. *Field Robotics*, 2:274–324, 2022. 1

[44] Diede PM Van der Hoorn, Anouk Neerincx, and Maartje MA de Graaf. " i think you are doing a bad job!" the effect of blame attribution by a robot in human-robot collaboration. In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, pages 140–148, 2021. 1

[45] J Irving Vasquez-Gomez, L Enrique Sucar, Rafael Murrieta-Cid, and Efrain Lopez-Damian. Volumetric next-best-view planning for 3d object reconstruction with positioning error. *International Journal of Advanced Robotic Systems*, 11 (10):159, 2014. 2

[46] Viswanath Venkatesh. Determinants of perceived ease of use: Integrating control, intrinsic motivation, and emotion into the technology acceptance model. *Information systems research*, 11(4):342–365, 2000. 4.1.2

[47] Nils Wilde, Alexandru Blidaru, Stephen L Smith, and Dana Kulić. Improving user specifications for robot behavior through active preference learning: Framework and evaluation. *The International Journal of Robotics Research*, 39(6):651–667, 2020. 1

[48] Brian Yamauchi. Frontier-based exploration using multiple robots. In *Proceedings of the second international conference on Autonomous agents*, pages 47–53, 1998. 2