

Human-Centered Autonomy for UAS Target Search

Hunter M. Ray^{*†}, Zakariya Laouar^{*}, Zachary Sunberg and Nisar Ahmed

Abstract—Current methods of deploying robots that operate in dynamic, uncertain environments, such as Uncrewed Aerial Systems in search & rescue missions, require nearly continuous human supervision for vehicle guidance and operation. These methods do not consider high-level mission context resulting in cumbersome manual operation or inefficient exhaustive search patterns. We present a human-centered autonomous framework that infers geospatial mission context through dynamic feature sets, which then guides a probabilistic target search planner. Operators provide a set of diverse inputs, including priority definition, spatial semantic information about ad-hoc geographical areas, and reference waypoints, which are probabilistically fused with geographical database information and condensed into a geospatial distribution representing an operator’s preferences over an area. An online, POMDP-based planner, optimized for target searching, is augmented with this reward map to generate an operator-constrained policy. Our results, simulated based on input from five professional rescuers, display effective task mental model alignment, 18% more victim finds, and 15 times more efficient guidance plans than current operational methods.

I. INTRODUCTION

Uncrewed Aerial Systems (UAS) are revolutionizing public safety as they provide new perspectives and rapid access to first responders in a variety of tasks including search & rescue (SAR), firefighting, water rescue, and law enforcement [1]. Operation of UAS often requires up to three team members to respectively pilot the vehicle, act as a technical specialist or observer, and coordinate platform integration across a mission. UAS autonomy can enable single-pilot operations with greater operator mobility and situational awareness by substantially reducing high task loads present in these settings. However, the unpredictable nature of public safety incidents puts challenging requirements on an autonomous system’s adaptability to changing circumstances.

Autonomous systems must seamlessly fit into teams to be effective, taking unquestioning direction from an operator and rapidly processing novel information, such as new search requirements, to dynamically (re)program their goals. For a UAS involved in SAR, this direction involves the operator defining their preferences about where and how to search specific areas based on their experience and knowledge, which characterizes a portion of their mental model. However, the unstructured nature of these preferences must be condensed into an accessible form for computationally constrained planning and informed execution.

Authors are with the Ann And H.J. Smead Aerospace Engineering Sciences Department, 429 UCB, University of Colorado Boulder, Boulder CO 80309, USA. Correspondence: hunter.ray-1@colorado.edu.

^{*} These authors contributed equally to this work

[†]Rescuer, UAS Supervisor, Boulder Emergency Squad, 3532 Diagonal Highway Boulder, CO 80301.

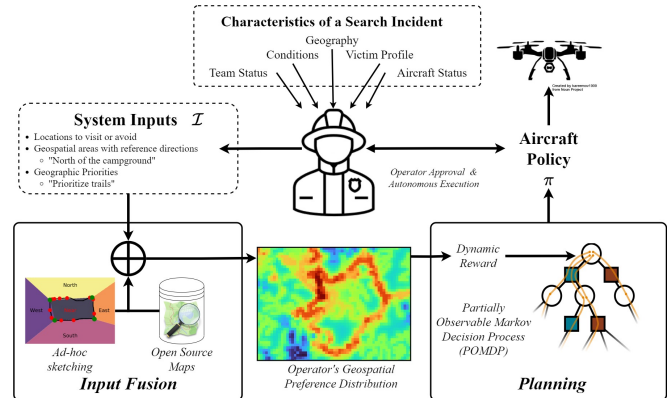


Fig. 1: Search & rescue incidents require operators to fuse multiple sources of information to direct an aircraft. Our human centered architecture fuses a variety of unstructured operator inputs to inform an optimal planning process that generates a policy for autonomous execution.

Building upon our experience in rescue operations, we develop a human-centered autonomous framework, shown in Figure 1, which enables dynamic planning for UAS SAR missions. Current methods of automated flight use an operator’s inputs, \mathcal{I} , as a deterministic policy, π , to follow a set of waypoints or execute an exhaustive search over a polygonal area [1]. However, these \mathcal{I} also contain valuable context, which can infer higher level mission goals, and thereby define extensive autonomous behavior. With this new information, the autonomous agent must effectively balance the operator’s direction while fulfilling the primary task of searching for the victim.

We build upon current modes of input through an augmented set of \mathcal{I} - priority definition, spatial semantic observations over ad-hoc geographical areas, and example waypoints - which informs an operator’s geospatial preferences over a mission area. These inputs are processed using a probabilistic model and fused via Bayesian inference to estimate a geospatial distribution reflecting operator preferences. Our new, model-driven, approach is necessary as no two incidents are exactly the same and attempting to systematically learn overall preferences based on prior events is challenging due to multiple communication modalities, key nuances across incidents, and limited data availability. Instead, we dynamically fuse and then embed operator preferences into a reward model for a Partially Observable Markov Decision Process (POMDP) to generate an online policy.

In summary, our key technical contributions include: (1) a new interaction paradigm that enables operators to provide diverse modes of input, which inform a nuanced and complex set of preferences on autonomous behavior; (2) a probabilistic model that rapidly aligns task mental models for geospatial feature prioritization between autonomy and

expert using sparse training data and operator-augmented feature vectors; (3) a general POMDP formulation for collaborative target search using adaptive reward; and (4) proof-of-concept user study with five rescuers that leverages a realistic search scenario and export informed truth model to evaluate our approach in simulations against an operational baseline.

II. PROBLEM STATEMENT AND BACKGROUND

Our work is motivated with a hypothetical incident, derived from real-life scenarios, without loss of generality. In this scenario, a middle-aged man, a frequent fly-fisherman, was reported missing by his wife after he failed to return from a daytime trip to the Walker Ranch Open Space in Boulder County, CO. The victim’s vehicle was found at the southern trailhead. The first arriving unit includes a UAS team, which is tasked with performing a preliminary search of the area.

A. Problem Statement

A search operation takes place over a larger state space, $S \subseteq \mathbb{R}^f$, representing f mission influencing factors, including static (geography) or dynamic (team locations) variables. The operator acts as a means of fusing the contextual and environmental aspects of an evolving operation. They provide a set of inputs, \mathcal{I} , that correspond to their preferences over locations to search. \mathcal{I} must be interpreted along with aircraft performance constraints, such as limited battery life, to generate a guidance policy, π , that effectively searches an area. The aircraft can perform its internal navigation and control to follow π while looking for infrared hotspots or visual clues as to a static victim’s location. This detection can rely on computer vision or the operator to supervise video feeds, however this is outside this paper’s scope.

B. Related Work

Given the information in the motivating mission, an officer directing operations gains awareness of the victim’s physical condition, potential equipment, last-known-point, mindset, and behavior. These factors of their evolving mental model are then grounded in the mission area, which informs the resulting search methods and tasking [2]. If the UAS management system can minimize the need for input and monitoring, the operator can move across the environment and maintain situational awareness, resulting in a faster rescue time. However, this requires interpreting and then incorporating the mental model into the overall autonomy architecture.

While robots interpret mental models in multiple ways [3] and teams that use them improve their performance [4], they often rely on static, pre-defined architectures [5],[6]. Autonomous behavior has also been effectively guided by operators using a variety of methods, including implementing a POMDP to couple the operator’s inputs with the search task planning [7], leveraging set plays [8], or using Bayesian priors on geological knowledge to define operator preferences [9]. While these methods effectively leverage the operator’s knowledge to guide behavior, they require limited, specialized, or mission specific inputs. We aim to decouple

the fusion of the operator’s inputs with the vehicle planning and execution to enable a more flexible, ‘plug-and-play’, approach that accommodates diverse underlying autonomous architectures, such as interchangeable input types, fusion methods, or planners.

An operator’s preferences have also been learned using Inverse Reinforcement Learning (IRL). However, whereas IRL relies on observed expected behaviour, such as near optimal reference trajectories, to directly infer reward functions [10], we infer higher-level preference distributions by fusing multi-modal inputs and readily available geographic database information. This distribution is factored into a reward function, complementing a planner’s baseline performance with end user expertise, addressing the challenges found with rewards designed by engineers [11],[12].

Leveraging autonomy to aid rescuers in efficiently searching an area provides obvious benefits. Recent work in automated UAS flight planning [13] provides effective search patterns (e.g. lawnmower) often used by U.S. maritime SAR personnel [14]. While this can be improved with more explicit modeling of target location uncertainty [15], obtaining an accurate and informed prior is challenging. The fusion of our inputs into a preference distribution could be applied as a target location estimate, however, we aim to preserve mission level uncertainty of target location while accounting for dynamic operator preferences.

POMDPs are effective at modeling these competing objectives (finding the victim and satisfying their operator) in their ability to reason over aleatoric and epistemic uncertainty [16]. POMDPs have already been applied to SAR, but these explicitly model domain information such as cell signals and crowd reports [17]. We don’t assume this information is available and instead use a POMDP to model expert-provided domain knowledge as dynamic reward, allowing observations to inform and shape the target belief over time. This allows us to flexibly balance the operator’s preferences against the agent’s primary goal of finding the victim.

Methods for searching cluttered environments, such as kitchens, have applied POMDP planners. These complex spaces have been navigated by condensing them into multiple resolutions [18], incorporating natural language from the user [19], or sequentially decluttering the space [20]. While these studies focus on close quarters, we investigate planning over multiple square kilometers and flexibly integrating expressive human input from semantic language and sketches.

Compared to previous work, we rely on the operator to interpret the vast set of unstructured data involved in a SAR mission, which is conveyed through intuitive methods and fused with geographical data. During execution, our approach allows for *human-on-the-loop* collaboration, where the human can provide inputs at any time during the search, but the system can still function competently without input. This paradigm enables the autonomous system to effectively balance the overarching goal of finding the victim while being informed and guided by dynamic operator preferences.

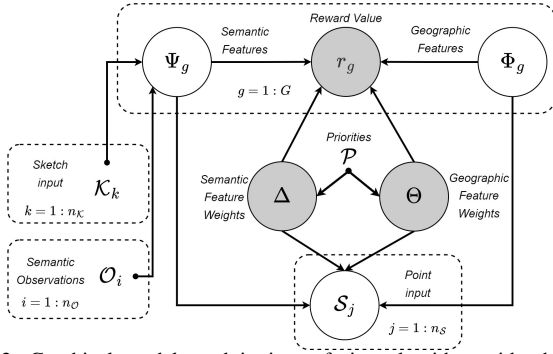


Fig. 2: Graphical model used in input fusion algorithm with observed variables (white) and unobserved variables (grey).

III. METHODOLOGY

A two-step approach is implemented towards interpreting user inputs and generating an aircraft policy, as shown in Figure 1. An operator's inputs are captured and fused via a probabilistic graphical model to define a geospatial preference distribution. This distribution augments a POMDP reward function, generating a dynamic policy to inform aircraft guidance.

A. Input Fusion

We define a grid $\mathcal{G} \subset S$, which discretizes the operational region within $S \subseteq \mathbb{R}^2$ into an $n \times m$ grid, \mathcal{G} , with areas $g \in \mathcal{G}$ of uniform resolution. Each g contains a set of defining features compactly represented as vectors. This includes a static set of geographic features, $\Phi \in \mathbb{R}^{n_\Phi}$, which describe how g is related to the built and natural environment including trails or buildings. Each g is augmented with a set of user-defined features, $\Psi \in \mathbb{R}^{n_\Psi}$, which describe g 's spatial context within a particular mission, such as its proximity to a staging area or location within a specific map sector.

We assume that the operator's preferences can be modeled as a set of visitation preferences over Φ and Ψ , which are represented as an unnormalized set of associated feature weights for geographic, $\theta \in \mathbb{R}^{n_\theta}$, and semantic, $\Delta \in \mathbb{R}^{n_\Delta}$, components. For a given g with operator reward, r , the associated set of features, Φ_g and Ψ_g , are modeled to be linearly related to their respective weights, θ and Δ ,

$$r_g = \theta^T \Phi_g + \Delta^T \Psi_g \quad (1)$$

This linear relation enables the different features to be appropriately accounted for while being flexible to changing n_Ψ . To account for uncertainty from imperfect mapping, limited or mistaken operator input and other unknown error sources, we seek to infer a probability distribution over r_g , $p(r_g|\mathcal{I})$. This inference is accomplished through the graphical model shown in Figure 2, which the operator engages with through a set of three possible inputs.

We define the set of \mathcal{I} to include $n_\mathcal{P}$ priorities \mathcal{P} , $n_\mathcal{O}$ semantic geospatial observations \mathcal{O} , and $n_\mathcal{S}$ waypoints, \mathcal{S} . All three types of \mathcal{I} can be provided at the start and while it is amenable to modification during the mission this is not explicitly shown in our experiments. Each \mathcal{P} acts as a positive, equally weighted prior on the respective geographic or semantic feature's weight. An input of \mathcal{O} is built with a combination of a 2D convex sketch, \mathcal{K} , and a geospatial

semantic reference label, where each additional \mathcal{O} increases the dimensions of Δ and Ψ . Finally, the set of \mathcal{S} includes locations for the aircraft to visit ($S = 1$) and those that should be avoided ($S = 0$).

Within the context of our reference mission, these inputs may take the following form: $\mathcal{P} = [\text{Trails, Section A}]$, $\mathcal{O} = [\text{"Search north of the Bridge", "Search inside section A"}]$, $\mathcal{K} = [\text{Bridge, Section A}]$, $\mathcal{S}_{\text{visit}, (S=1)} = [\text{Set of waypoints along a river}]$, $\mathcal{S}_{\text{avoid}, (S=0)} = [\text{Set of waypoints near structures}]$.

While the set of \mathcal{P} and \mathcal{O} initialize the estimates and add features to the environment, the observable set of \mathcal{S} help tie the known feature components Φ and Ψ to their respective unknown weightings θ and Δ . We assume that each g must reach a specific threshold of an operator's optimal positive or negative preference for it to be provided as a reference. We therefore model $p(S_g|\theta, \Delta, \Phi_g, \Psi_g)$ as a logistic function in equations 2 and 3 with r_g defined as in equation 1,

$$p(S_g = 1|\theta, \Delta, \Phi_g, \Psi_g) = \frac{\exp(r_g)}{1 + \exp(r_g)}, \quad (2)$$

$$p(S_g = 0|\theta, \Delta, \Phi_g, \Psi_g) = \frac{1}{1 + \exp(r_g)}. \quad (3)$$

To define Φ , geographic features can be extracted using publicly available government and open source datasets from across the world. A geographical information system software, such as ArcGIS, can integrate a subset of information for a specific area, such as a county or an agency's jurisdiction. For this work, we consolidate information on roads, trails, structures, stream lines, water bodies, and tree canopy. We aim to create a set of Φ_g that accurately maps an operator's perspective of their respective value, which includes each g 's proximity to a relevant feature. Therefore for each g with a given resolution, we determine the distance, d , to the closest respective feature, f , informing an adjacency metric, calculated as $\Phi_{f,g} = \exp(\frac{-d}{\text{resolution}})$. The geoprocessing can be performed offline and saved into a resolution specific database.

To define Ψ , prior work [21]–[24] is leveraged that probabilistically relates a geospatial semantic label to a given sketch, \mathcal{K}_k . To reduce computational loads, we sequentially reduce the number of vertices to five points that maximize the overall area to form the resulting convex polytope. The available semantic labels includes a comprehensive set of canonical bearing labels {"N", "NE", "E", "SE", "S", "SW", "W", "NW"} and discrete ranges {"Inside", "Near", "Outside"}. Given a particular grid area g , with a center location $x_g \in \mathbb{R}^2$ we can approximate the likelihood of the semantic label using the softmax function, where each class contains a set of parameters $w \in \mathbb{R}^2$ and $b \in \mathbb{R}^1$, which constrain their boundaries along the sketch border [23]. Once defined, a Monte Carlo approximation is used to correlate the softmax classes with specific semantic labels [24], resulting in $p(\text{label}|\text{class})$. We therefore define each Φ_g as $p(\text{label}|g)$, the probability of the given grid point being represented by a certain label,

$$\Psi_{i,g} = p(\text{label}|g) = \sum_{\text{class}} p(\text{label}|\text{class})p(\text{class}|g) \quad (4)$$

Having defined the specific components of the graphical model, inference must be performed over $p(\theta, \Delta | \Phi, \Psi, \mathcal{P}, \mathcal{S})$ to find the resulting reward. The weights, θ and Δ , are each represented as multivariate Gaussians with respective means μ_θ, μ_Δ and covariances $\Sigma_\theta, \Sigma_\Delta$, allowing a quick approximation of the distribution through the Laplace approximation [25]. We take the resulting expected value of θ and Δ to estimate the posterior reward:

$$\hat{r}_g = \mu_\theta^T \Phi_g + \mu_\Delta^T \Psi_g \quad (5)$$

Additional derivations and details on model inference are discussed in the extended version of our paper [26].

B. Adaptive-Reward Target Search POMDP

The objective of the planning module is to search for a static target with an unknown location while accounting for operator preferences. We do this by formulating and then solving a general target search POMDP that leverages the operator’s geospatial preference distribution to find the target within the previously defined $n \times m$ grid, \mathcal{G} . We introduce our POMDP formulation below:

A *Partially Observable Markov Decision Process* (POMDP) is a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{O}, \mathcal{T}, \mathcal{R}, \mathcal{Z}, \gamma)$, where: \mathcal{S}, \mathcal{A} , and \mathcal{O} are finite sets of states, actions and observations, respectively, $\mathcal{T} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is the transition probability function, $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the immediate reward function, $\mathcal{Z} : \mathcal{S} \times \mathcal{A} \times \mathcal{O} \rightarrow [0, 1]$ is the probabilistic observation function, $\gamma \in [0, 1]$ is the discount factor [16].

State space \mathcal{S} : Four components: 2D robot position, 2D target position, and battery life of the search agent. Lastly, the state is augmented with a $|n \times m|$ bitvector to encode whether a robot has visited each grid cell.

Observations \mathcal{O} : The target location uncertainty is represented as a one-hot vector encoding of target occupancy in the \mathcal{M} cells adjacent to the true target cell. Additionally, the agent may receive a *target-not-found* observation represented as a vector of zeros, leading to a cardinality of $|\mathcal{M} + 1|$.

The observation function, \mathcal{Z} , models the agent’s ability to accurately observe the target, informed by factors such as sensor capabilities or terrain complexity. This nuance can be expressed by choosing a Manhattan distance, D_{obs} , to express how close the agent must be to receive a positive target observation. The agent receives a noisy observation of the target at each timestep. If within D_{obs} , the agent can receive one of the following three observations given the true location of the target. The agent can receive a true positive observation with probability Z_{true} . The agent can also receive a *proximal observation* which returns a target observation near the true target location. Each cell defined as proximal to the true target cell will receive this observation with probability Z_{prox} . Lastly, the agent can receive a false negative observation with probability Z_{neg} as shown in Figure 3. Otherwise, the agent receives the *target-not-found* observation with 100% probability.

In the example shown in Figure 3, the agent receives a positive observation only if adjacent to the true target. Thus, \mathcal{O} is presented with $\mathcal{M} = 5$ and \mathcal{Z} being distributed over the true target cell and the 2 adjacent cells.

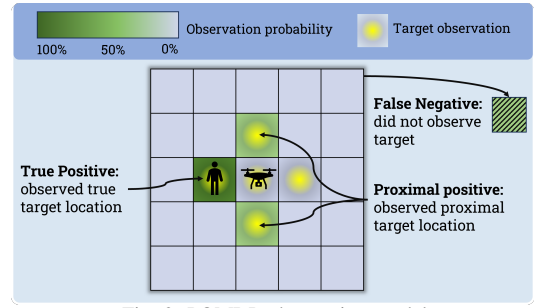


Fig. 3: POMDP observation model.

Transitions \mathcal{T} : The agent can deterministically move in four cardinal directions: *up*, *down*, *left*, *right* as well as *stay* in the grid cell. As the target is assumed to be static, their position remains the same. The state’s battery component decreases a fixed B_{cost} amount each timestep. When the agent transitions from state s to s' , the occupancy history is updated, ensuring a record of g that the agent has visited.

Rewards $\mathcal{R}(s, a, s')$ = $R_{time} + R_{target} + R_{op}$

- R_{time} : Time penalty for each step taken.
- R_{target} : Reward collected for finding the target, i.e. if the agent and target positions in s' are equivalent.
- R_{op} : Each grid cell g has a corresponding reward \hat{r}_g obtained from the operator’s geospatial preference distribution. To incorporate operator preferences, R_{op} will only be collected if the robot location in s' has yet to be visited: $R_{op} = \hat{r}_g \times \mathbb{1}_{visited(s')}$.

To incentivize the agent to travel to the regions that have a concentrated belief over the target location, the reward for finding the target needs to be higher than the highest value in the operator’s geospatial preference distribution. If this distribution provides higher reward regions, the agent will more likely seek out those regions, ignoring valuable information about the target’s potential location.

Termination: A *terminal* or *absorbing* state is a state from which no further transitions to other states are possible. Conditions for termination are: (i) the current robot position equals the target position, and (ii) the current battery equals the battery required to return to the initial state, with the initial battery life B_{max} tuned to represent operational constraints. Modeling the battery life as a terminal state ensures that the aircraft can return to base before the battery depletes and, in doing so, incentivizes exploration closer to base.

C. Planning

Our POMDP formulation results in a very large state space ($|\mathcal{S}| = B_{max} \times (n \times m)^2 \times 2^{(n \times m)}$) to plan over. As SAR requires fast and efficient planning, offline POMDP solvers are too slow for practical application. While amenable to any online Monte Carlo based solver, leverage POMCP [27] and capture the belief over target location using a particle filter to overcome scalability challenges associated with planning in high dimensional spaces.

Rollout Policy: A simple and effective rollout policy is a practical method for evaluating nodes of an online POMDP solver. Our rollout policy relies on an MDP abstraction of the

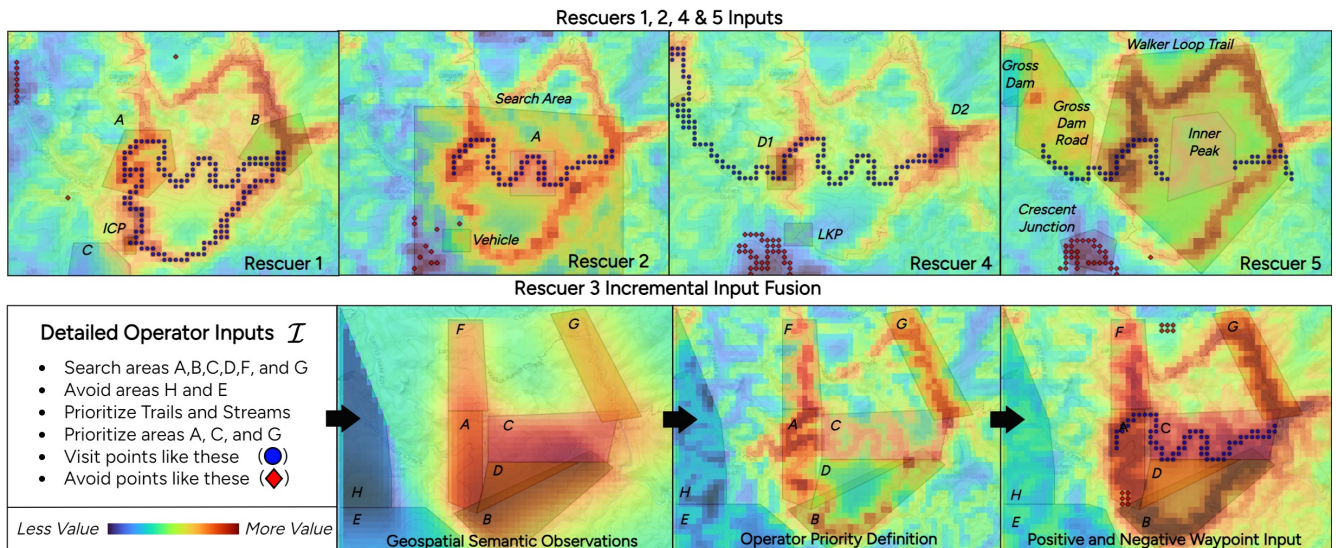


Fig. 4: Inputs from five rescuers overlay the resulting geospatial preference distribution. The lower row shows the progressive addition of rescuer 3’s inputs and resulting concentration of reward around trails and streams within prioritized areas. Note that in execution, inputs are fused simultaneously.

larger POMDP by removing observations and only considering the agent’s position, battery, and a randomly sampled target position. We solve the abstraction using discrete value iteration upon problem initialization. During the execution of the POMDP, this policy is evaluated for candidate methods of exploring the state space and, once completed, collects the full reward including the operator’s preference distribution. While not accounting for agent’s current belief or operator preferences, this approach retains acceptable grid exploration with low computational cost.

IV. USER EVALUATION RESULTS AND DISCUSSION

In this section, our approach is validated against an operational baseline in a realistic simulation environment. Results show that without loss of generality with respect to operator input, our approach can reliably interpret and plan over an operator’s inputs to more efficiently find a missing person than current operational standards for autonomous flight.

To that end, inputs from five first responders were collected from the Boulder Emergency Squad (BES), a volunteer technical rescue team in Boulder, CO that has been using UAS for over eight years in SAR, firefighting, and law enforcement operations [1]. Subjects were ages 24-61 and were either UAS pilots or had supplementary accreditation in search management, reflecting subject matter expertise. Each rescuer provided a set of inputs, \mathcal{I} , based on the Lost Fisherman scenario discussed previously. Their inputs informed randomized simulations with multiple vehicle launch locations and target positions, whereby we compare our system’s performance against an operational baseline.

A. Data Collection

To collect subject data, the researcher first walked subjects through a testing scenario detailing the inputs they could provide. Subjects were guided towards appropriate inputs when necessary, such as using the priority definition when they wanted certain features to be focused on. To replicate

an intuitive interaction method uninhibited by software limitations, inputs were collected on a paper map, which was then translated into the reward model and inference routine post-hoc. Figure 4 shows, for illustration, the progressive fusion of rescuer 3’s inputs and the final fusion results for rescuers 1, 2, 4 and 5. All rescuers provided unique inputs, although often promoted similar features, such as trail intersections and primary streams. In the lower part of Figure 4, rescuer 3’s inputs build the reward map with their preferences concentrating in the prioritized areas and where the trail travels close to streams.

B. Input Fusion Validation

Before evaluating the overall framework, accurate fusion of an operator’s inputs must be validated by aligning the algorithm’s estimate with their true geographic and feature preferences. This is accomplished with an error metric that compares the uncertainty-weighted reward estimate of 21 distributed locations against an operator’s quartile ranking. In addition, the algorithm’s positive and negative feature prioritization is computed to understand the underlying motivation for alignment. This evaluation adapts a search engine ranking metric known as normalized Discounted Cumulative Reward [28], which normalizes the estimated ranking order by the subject-provided ranking to provide an optimal value of 1. This metric is repetitively evaluated with a set of weights sampled from $p(\theta, \Delta | \Phi, \Psi, \mathcal{P}, \mathcal{S})$. Additional details on these metrics is included in the extended version of our paper [26].

Input validation results are shown in Table I, where each metric is compared to a random baseline, shown in parentheses, which is used as no comparable baseline exists. Results across all subjects show very low error values. Positive feature rankings mostly show good alignment, though alignment rank suffered if limited inputs were provided or subjects provided feature weightings that were close together, such as for Rescuer 4. Negative feature rankings proved to be more effective, thereby highlighting that the system can accurately discriminate features that should be avoided.

TABLE I: Alignment results (with random baseline)

Rescuer	Error (Optimal = 0)	Pos. Feature Rank (Optimal = 1)	Neg. Feature Rank (Optimal = 1)
1	0.015 (2.95)	0.93 (0.63)	1.21 (3.66)
2	0.015 (2.19)	0.97 (0.89)	1.07 (1.21)
3	0.013 (2.41)	0.95 (0.72)	1.28 (2.17)
4	0.025 (2.41)	0.84 (0.55)	1.85 (2.02)
5	0.009 (2.06)	0.94 (0.76)	1.47 (1.81)

C. Search and Rescue Simulation


Based on the value maps generated from the rescuers’ inputs, the planner is simulated and its performance is compared to operational methods. The environment consists of a 44×59 grid, with each grid cell spanning 100×100 meters. Simulations are initialized across four expert-informed start locations and sample target locations from a truth model, shown in our extended version [26], which was informed by BES Chief Andy Amalfitano, an expert in SAR.

Baseline Implementation: The baseline search algorithm receives the operator’s sketch set \mathcal{K} , positive ("Inside" and "Near") observations, \mathcal{O} , and waypoints \mathcal{S}_{visit} . This approach reflects currently available autonomous modes used in operations [1]. The search agent starts at one of four locations and follows a shortest path algorithm to traverse all waypoints. Once complete, the agent travels to each sketched polygon and executes a lawnmower search pattern of the space. If the target is still not found, the agent redirects to the exhaustive search stage and conducts a lawnmower search over the entire environment.

Our Approach: The POMDP described in Sec. III-B is solved with the following details. We define a strict target observation condition $D_{obs} = 1$, defining positive observations of the target only when within 1 grid cell. If within D_{obs} , the UAS receives a true positive observation with $Z_{true} = 80\%$ and a proximal positive observation with $Z_{prox} = 10\%$ for the two adjacent cells in the direction of the target. We set our reward model’s running cost $R_{time} = -1$ and the reward for finding the target $R_{target} = 1000$, the \hat{r}_g leveraged from the geospatial preference distribution. Finally, we set the UAS initial battery $B_{max} = 1000$.

The resulting UAS behavior is represented in Figure 5. Our approach efficiently explores high value areas, autonomously searching in and outside of the defined sectors and waypoints, eventually finding the target 2.6 times faster. Our approach is comprehensively evaluated against the baseline using two metrics. The first, *Localization Ratio* defines the ratio of simulation runs that find the target over runs that do not find the target within the given battery life. The second metric, *Reward/Timestep* expresses the discounted reward accumulated per timestep. If the agent fails to find the target under the battery constraints, the simulation is terminated. For each of the four start locations, 100 simulations are performed with the overall reported mean and standard error of the mean reported in Table II.

Our method results in statistically significant ($p < 0.05$) improvement from the baseline approach for all rescuers’ input and metrics. Evaluating the localization ratio, our method finds the target 18% more than the baseline within



(a) Our approach planning over the underlying geospatial preference distribution. (b) Baseline approach planning over raw operator inputs.

Fig. 5: Example trajectories of our planner and the baseline simulated using rescuer 1 input data from Figure 4. Our approach finds the target in 329 timesteps, baseline finds the target in 858 timesteps. Trajectory opacity decreases with search time.

the available battery life. A two-tailed Binomial test ($N = 400$) is performed to find statistical significance. More importantly, the system collects 15.4 times more reward per timestep, reflecting efficient path planning that follows the operator’s geospatial context drawn from a limited set of inputs. Statistical significance for this metric is evaluated using a two sample Z-test ($N = 400$).

TABLE II: Simulation Results - 100 runs per start position

Rescuer	Localization Ratio		Reward/Timestep	
	Ours	Base	Ours	Base
1	75.3%	54.8%	4.56 ± 1.21	0.490 ± 0.043
2	54.5%	23.3%	2.05 ± 0.60	0.195 ± 0.027
3	63.0%	46.3%	4.72 ± 0.97	0.173 ± 0.023
4	44.3%	25.5%	4.70 ± 1.36	0.101 ± 0.015
5	51.0%	45.8%	1.85 ± 0.72	0.201 ± 0.027
Average	57.6%	39.1%	3.58 ± 0.67	0.232 ± 0.067

V. CONCLUSION

We introduced a human-centered autonomy framework tailored to UAS deployed in dynamic and uncertain environments, particularly in SAR missions. Our approach hinges on inferring geospatial context and operator preferences using minimal operator inputs to guide a probabilistic target search policy. An operator defines priorities, spatial semantic observations over ad-hoc geographical areas, and waypoints as inputs, which provide valuable context of the operator’s mental model. The system infers a reward function from these inputs and plans online using this reward signal and its own belief about target location. Effective alignment of inferred and true operator preferences is proved through a spatial error metric and feature ranking. Simulated Monte Carlo trials revealed that our inference and planning pipeline significantly outperforms an operational baseline in both target localization and reward accumulation.

While our application focuses on SAR, the ability to estimate an operators geospatial preferences for guided planning shows promise in enabling expert-directed autonomy in any data-driven environment. However, our method is limited to areas with mapped geography. Future work will seek ways to address this limitation, such as incorporating online aerial perception techniques, and deploying fielded hardware experiments with vehicle and interface integration.

ACKNOWLEDGEMENTS

We thank the members of the Boulder Emergency Squad for their thoughtful contributions and input.

REFERENCES

- [1] H. Ray, R. Singer, and N. Ahmed, "A Review of the Operational Use of UAS in Public Safety Emergency Incidents," in *Proc. of the International Conference on Unmanned Aircraft Systems (ICUAS)*, Dubrovnik, Croatia: IEEE, 2022, pp. 922–931.
- [2] W. B. Rouse and N. M. Morris, "On looking into the black box: Prospects and limits in the search for mental models," *Psychological Bulletin*, vol. 100, pp. 349–363, Nov. 1986.
- [3] A. Tabrez, M. B. Luebbbers, and B. Hayes, "A Survey of Mental Modeling Techniques in HumanRobot Teaming," *Current Robotics Reports*, vol. 1, pp. 259–267, Dec. 2020.
- [4] F. Gervits, D. Thurston, R. Thielstrom, T. Fong, Q. Pham, and M. Scheutz, "Toward Genuine Robot Teammates: Improving Human-Robot Team Performance Using Robot Shared Mental Models," in *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems*, Auckland, New Zealand: IFAAMAS, May 2020, pp. 429–437.
- [5] M. Scheutz, S. A. DeLoach, and J. A. Adams, "A Framework for Developing and Using Shared Mental Models in Human-Agent Teams," *Journal of Cognitive Engineering and Decision Making*, vol. 11, no. 3, pp. 203–224, Sep. 2017.
- [6] S. V. Albrecht and P. Stone, "Autonomous Agents Modelling Other Agents: A Comprehensive Survey and Open Problems," *Artificial Intelligence*, vol. 258, pp. 66–95, May 2018.
- [7] S. Jamieson, J. P. How, and Y. Girdhar, "Active Reward Learning for Co-Robotic Vision Based Exploration in Bandwidth Limited Environments," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, May 2020, pp. 1806–1812.
- [8] C. A. Miller, "Delegation and intent expression for human-automation interaction," in *Proc. of the International Conference on Human-Computer Interaction in Aerospace*, ser. HCI-Aero '14, New York, NY: Association for Computing Machinery, 2014.
- [9] A. Arora, P. M. Furlong, R. Fitch, S. Sukkarieh, and T. Fong, "Multi-modal active perception for information gathering in science missions," *Autonomous Robots*, vol. 43, pp. 1827–1853, Oct. 2019.
- [10] S. Arora and P. Doshi, "A survey of inverse reinforcement learning: Challenges, methods and progress," *Artificial Intelligence*, vol. 297, p. 103 500, Aug. 2021.
- [11] D. Hadfield-Menell, S. Milli, P. Abbeel, S. J. Russell, and A. Dragan, "Inverse Reward Design," in *Advances in Neural Information Processing Systems*, vol. 30, Curran Associates, Inc., 2017.
- [12] S. Booth, W. B. Knox, J. Shah, S. Niekum, P. Stone, and A. Allievi, "The Perils of Trial-and-Error Reward Design: Misdesign through Overfitting and Invalid Task Specifications," in *Proc. of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, Washington, D.C, Feb. 2023.
- [13] D. Kingston, S. Rasmussen, and L. Humphrey, "Automated UAV tasks for search and surveillance," in *2016 IEEE Conference on Control Applications (CCA)*, IEEE, 2016, pp. 1–8.
- [14] C. Thomas, "US coast guard addendum to the US NSS to the IAMSAR," *United States Coast Guard, Tech. Rep. COMDTINST MI6130*, vol. 2, 2013.
- [15] F. Bourgault, T. Furukawa, and H. F. Durrant-Whyte, "Optimal search for a lost target in a bayesian world," *Field and Service Robotics: Recent Advances in Research and Applications*, pp. 209–222, 2006.
- [16] M. J. Kochenderfer, T. A. Wheeler, and K. H. Wray, *Algorithms for decision making*. MIT press, 2022.
- [17] C. A. Baker, S. Ramchurn, W. L. Teacy, and N. R. Jennings, "Planning search and rescue missions for UAV teams," in *Proceedings of the Twenty-second European Conference on Artificial Intelligence*, 2016, pp. 1777–1778.
- [18] K. Zheng, Y. Sung, G. Konidaris, and S. Tellex, "Multi-resolution POMDP planning for multi-object search in 3D," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2021, pp. 2022–2029.
- [19] A. Wandzel, Y. Oh, M. Fishman, N. Kumar, L. L. Wong, and S. Tellex, "Multi-object search using object-oriented POMDPs," in *2019 International Conference on Robotics and Automation (ICRA)*, IEEE, 2019, pp. 7194–7200.
- [20] Y. Xiao, S. Katt, A. ten Pas, S. Chen, and C. Amato, "Online planning for target object search in clutter under partial observability," in *2019 International Conference on Robotics and Automation (ICRA)*, IEEE, 2019, pp. 8241–8247.
- [21] N. R. Ahmed, E. M. Sample, and M. Campbell, "Bayesian Multicategorical Soft Data Fusion for HumanRobot Collaboration," *IEEE Transactions on Robotics*, vol. 29, pp. 189–206, Feb. 2013.
- [22] L. Burks, H. M. Ray, J. McGinley, S. Vunnam, and N. Ahmed, "HARPS: An Online POMDP Framework for Human-Assisted Robotic Planning and Sensing," *IEEE Transactions on Robotics*, 2022, In Press.
- [23] N. Sweet and N. Ahmed, "Structured synthesis and compression of semantic human sensor models for Bayesian estimation," in *2016 American Control Conference (ACC)*, Boston, MA: IEEE, Jul. 2016, pp. 5479–5485.
- [24] C. L. Burks, "Active Collaborative Planning and Sensing in Human-Robot Teams," PhD thesis, University of Colorado Boulder, Boulder, CO, May 2020.
- [25] C. M. Bishop, *Pattern Recognition and Machine Learning*. Cambridge U.K.: Springer, 2006.
- [26] H. M. Ray, Z. Laouar, Z. Sunberg, and N. Ahmed, *Human-centered autonomy for UAS target search*, 2023. arXiv: 2309.06395 [cs.LG].

- [27] D. Silver and J. Veness, "Monte-carlo planning in large POMDPs," *Advances in neural information processing systems*, vol. 23, 2010.
- [28] K. Järvelin and J. Kekäläinen, "Cumulated gain-based evaluation of IR techniques," *ACM Transactions on Information Systems*, vol. 20, pp. 422–446, Oct. 2002.