

Collaborative Multi-Robot Multi-Human Teams in Search and Rescue

Ryan K. Williams*

Virginia Tech[†]
rywilli1@vt.edu

James McClure

Virginia Tech
mcclurej@vt.edu

Larkin Heintzman

Virginia Tech
hlarkin3@vt.edu

Tianzi Wang

Virginia Tech
tianzi@vt.edu

Nicole Abaid

Virginia Tech
nabaid@vt.edu

Nathan Lau

Virginia Tech
nkclau@vt.edu

Amanda Hashimoto

Virginia Tech
ahashimo@vt.edu

Chinmaya Patnayak

Virginia Tech
cpatnayak@vt.edu

Akshay Kumar

Virginia Tech
akshaykumark@vt.edu

ABSTRACT

Robots such as unmanned aerial vehicles (UAVs) deployed for search and rescue (SAR) can explore areas where human searchers cannot easily go and gather information on scales that can transform SAR strategy. Multi-UAV teams therefore have the potential to transform SAR by augmenting the capabilities of human teams and providing information that would otherwise be inaccessible. Our research aims to develop new theory and technologies for field deploying autonomous UAVs and managing multi-UAV teams working in concert with multi-human teams for SAR. Specifically, in this paper we summarize our work in progress towards these goals, including: (1) a multi-UAV search path planner that adapts to human behavior; (2) an in-field distributed computing prototype that supports multi-UAV computation and communication; (3) behavioral modeling that yields spatially localized predictions of lost person location; and (4) an interface between human searchers and UAVs that facilitates human-UAV interaction over a wide range of autonomy.

Keywords

Search & Rescue, Autonomy, Lost-Person Modeling, GIS, Visualization.

INTRODUCTION

According to the National Crime Information Center's statistics, nearly 650,000 missing people were recorded in the United States in 2016 (*2016 NCIC Missing Person and Unidentified Person Statistics 2016*), and approximately 100,000 of these cases resulted in searches in urban or wilderness settings. For wilderness searches, the International Search and Rescue Incident Database (based on approximately 50,000 records) reports an overall mortality rate of

*corresponding author

[†]<http://caslab.ece.vt.edu>. This work was supported by the National Science Foundation through grant # CNS-1830414.

9%, with 5% never located by searches and an additional 24% found injured. Search efforts usually involve multiple agencies and may exceed 100 personnel working over days. In this paper, we describe our vision and current work to redefine lost person search-and-rescue (SAR) by enabling teams of human searchers and unmanned aerial robots to collaborate towards improving search outcomes and reducing human effort.

In typical searches, a human assigns tasks to searchers based on a computed probability of area (POA), the probability that a lost person resides in a given area. As time elapses, this area grows geometrically, favoring for quick resolution. Although trained searchers are the gold standard for tracking lost people in the wilderness, SAR would greatly benefit from teams of *autonomous* unmanned aerial vehicles (UAVs) able to augment human abilities. Beyond their ability to search inaccessible areas, multi-UAV teams in SAR offer an exciting testbed to study human-robot interactions (HRI). Since lost person behavior is understood with respect to both geophysical and transient landscape features, mathematical models that include features measurable by UAVs are expected to enrich a search coordinator's ability to calculate POA with certainty. However, how information would flow between UAVs and humans is a very complex problem: would a UAV team act like a group of remotely-controlled searchers or would it exploit its observations and known models to search autonomously? What is the balance of collaboration and autonomy in multi-robot multi-human (MRMH) search teams?

To ground the discussion, consider a team of aerial robotic systems that are deployed into a search environment, along with a team of human searchers who vary in capabilities across a range of SAR tasks (e.g., mobility, visual identification, etc.). Assume that based on knowledge of the lost person(s) and expert input from search leaders, an estimated probability map describing the existence of lost person(s) in the environment is initially available as input to our proposed solution. Finally, given the large quantity of data produced by real-time search, assume that the human-robot search activities are supported by three tiers of computation: (1) embedded processing per-robot; (2) edge computation local to the search environment and available to all robots; and (3) off-site cloud computing. In this context, *the overarching problem posed by MRMH-SAR is to determine, with minimal deployment cost, maximal search effectiveness, and optimal computational efficiency, a time-varying policy that deploys human-robot teams to measure an environment in search of lost persons*. This problem is uniquely challenging for a robot team as it requires a careful *balance* of coordination over multiple, potentially vast, spatiotemporal scales, efficient decision-making that accounts for robot, human, and sensor heterogeneity, minimally invasive yet effective human-robot collaboration, and novel remote sensing and estimation techniques.

Robotic Teams With Humans in the Loop As a relatively new research area, results in robotic planning with humans in the loop are sparse compared to more general planning methods. Recent results include synthesis methods for human-in-the-loop control protocols (Feng et al. 2016), improved human-robot team performance with human-inspired plan execution (Shah et al. 2011), team organizations for improved performance in human-robot teams (Lewis et al. 2010; Kruijff et al. 2014; Nourbakhsh et al. 2005), analyses of human vs. algorithmic planning in SAR (Chien et al. 2010), and anticipatory planning (Koppula et al. 2016) (see taxonomy (Yanco and Drury 2004) and survey (Goodrich and Schultz 2007) for more details). The related question of how one person can control a team of robots simultaneously is also an open and active area of research. For ground and air vehicles, research laboratories are devoting tremendous effort to identifying salient cues that control collective behavior in robots. However, extant research in human-robotic interaction typically focuses on very specific control functions rather than an overarching problem in an application domain. Specifically, research has not looked into human-robot interaction from the perspective of hierarchical levels of control, i.e., the relationships between task generation (what), task assignment (who), and path planning (how).

Autonomous UAVs in SAR While UAVs are currently used heavily for surveillance in military and civilian police operations, they are most often operated using remote controls and are seldom applied for acquiring human data in peaceful scenarios. At present, the use of UAVs in SAR are: (1) mapping; (2) victim search; (3) target observation; (4) delivery; and (5) communication relay (De Cubber 2017), with recent examples of aerial victim identification (Vempati et al. 2015), and important applications of manually operated UAVs in real-world searches (NBC 2018). Of particular note, a recent commercial application (*DroneSAR* 2020) offers path planning algorithms for UAV control tailored to SAR, but these paths are generated offline using static maps and cannot be updated in real-time. With that in mind, an international group of UAV owners and pilots have assembled to assist local law enforcement with SAR (*DroneSAR* 2020). While this group constitutes a powerful human resource, enabling *autonomous* UAVs in SAR would potentially improve searches as UAVs increase access to search data.




Figure 1. Human-robot collaboration in SAR.

OUR VISION

Our vision for SAR is a MRMH control system supported by a distributed computing infrastructure (Figure 1) that will be driven by innovation in three intersecting areas:

- **Multi-Robot Planning and Control.** A multi-UAV control system capable of augmenting SAR teams, with theoretical progress in: (i) risk-aware multi-robot planning with agent-based behavioral modeling for human-in-the-loop (HITL) control; (ii) long-term and online task assignment for search tasking, on demand human control, and distributed computing; and (iii) connectivity planning for flexible human-robot collaboration.
- **Human-Robot Interaction.** An interactive visualization system enabling human teams to co-generate, co-assign, and co-perform critical tasks with the multi-UAV control system, with theoretical progress in: (i) interactive visualization for autonomous systems operating under highly variable environmental/resource constraints; (ii) supporting human interactions with robotic exploration and exploitation; and (iii) design of reinforcement learning systems for promoting autonomous capabilities through human use.
- **Distributed Computation.** A distributed computational infrastructure for supporting autonomous decision-making and control, with theoretical progress in: (i) strategic utilization of embedded computing devices subject to power and network constraints; (ii) scalability of distributed decision-making to support collaborative autonomous control of UAVs; and (iii) opportunistic utilization of network and compute resources.

CURRENT RESEARCH

This paper presents our work in progress to achieve our above vision for SAR. Specifically, we describe our latest research findings in developing autonomous control of UAVs, edge computing in wilderness, lost-person behavioural modeling, and a web-based user interface for SAR professionals.

Autonomous path planning

We present a framework to plan a set of paths for a team of UAVs to autonomously gather information about the environment and search for a lost person (Cangan 2019). To optimize measurements the UAVs take into account a lost person predictive model, topography in the search environment, and predicted human searcher trajectories. The lost person is modeled by taking into account prior beliefs as well as movement patterns informed from previous SAR missions, and is discussed in a later section. The dataflow diagram for the proposed solution is shown on the left in Figure 2. By utilizing the advantages afforded by UAVs, we aim to increase the *efficiency* of SAR operations by minimizing the total time required while maximizing the chances of locating a lost person with available resources (Cooper et al. 2003).

As part of optimizing UAV search trajectories, we must predict the paths of human searchers operating on the ground. In a land SAR mission, a particular area of land will have human searchers assigned to search it, which are referred to as *sectors* (Koester 2008). We assume that each searcher has an entry and exit point for each sector, based on the overall sequence of sectors to be searched. The human searcher model we use here has two modes, a waypoint following mode and a gradient following mode. In the waypoint following mode, each searcher is represented as a self-propelled particle moving towards a predefined set of waypoints. Once a searcher is within a known radius of the current waypoint, the next waypoint becomes the next target. In our case, the waypoints




Figure 2. (Left) Data flow diagram for the proposed human-robot planning framework. (Right) The lost person probability heatmap, with the searcher and UAV paths overlaid in blue and red, respectively.

are arranged to generate search paths, sometimes referred to as *lawn mowing* paths. To better reflect reality, each searcher path is also influenced by the terrain gradient. For example, if the terrain in a sector becomes too steep to climb, the searcher is forced to navigate around the obstacle. In these cases, our human searcher model will switch to the gradient following mode of operation. However, there are cases where searchers may get stuck in steep trenches or canyons, and thus to combat this we included a tenacity parameter that gradually increases, allowing searchers to escape from these scenarios (i.e., local minima). Shown on the right in Figure 2 are the resulting anticipated searcher paths in blue, along with the lost person probabilistic heatmap. Notice that the searchers' lawnmower paths are being influenced by the underlying terrain.

Moving forward, we require some quantification of the risk inherent in a trajectory based on the lost person model. In this case, the risk can also be restated as our *belief* of the lost person's location with some uncertainty associated. We will omit a full derivation of the risk function for brevity, the details of which can be found in (Cangan 2019), and instead skip to the final objective function result. We parameterize the UAV paths with Bézier curves, which allows us to represent the risk cost as a function of the sparse set of parameters $\lambda \in \mathbb{R}^{p \times D}$, where $p \in \mathcal{N}$, which define trajectories given by $\theta_\lambda(t)$. Formally, our goal is to find a set of parameters, λ , to minimize the risk metric (risk cost), that is:

$$\lambda^* = \arg \min_{\lambda \in \Lambda} \mathcal{R} [\theta_\lambda(t)] \quad (1)$$

Where $\mathcal{R} [\theta_\lambda(t)]$ is the *risk* associated with a particular set of parameterized UAV trajectories, $[\theta_\lambda(t)]$. The resulting set of UAV trajectories would then be $\theta_{\lambda^*}(t)$. We consider several other constraints, such as smoothness, and total path length, however for brevity we omit the details, readers are referred to (Cangan 2019) for details.

Now to optimize UAV trajectories, given an initial sampling, utilizing a gradient descent approach operating on the objective function in (1). Specifically, we use an iterative approach, where in each iteration the gradient is computed about the current trajectory set, θ_λ , with respect to the current parameter set λ . The parameter set is then propagated at each iteration by following the direction of steepest descent as defined by the gradient of the objective functional. Since the objective functional is of large dimension and highly non-convex, we use an optimization method typically found in neural network-based machine learning applications, referred to as *Adam*. Adam features adaptive learning rate adjustment based on first and second moments of the gradient. The resulting UAV paths can be seen on the right in Figure 2 in red, the yellow circles indicate *no-fly* zones. Notice how the UAV paths tend to cover areas not already covered by human searchers, and high probability points according to the lost person probability heatmap.

For a quantitative evaluation of the proposed framework, we compare against three cases that relate to how SAR operations function: a) human searchers performing a search task without the assistance of UAVs, the most common circumstance currently (Van Tilburg 2017), b) human searchers performing a search with manual UAVs that follow the same path as the human searchers at a fixed height of 15m over the terrain, c) human searchers with fully autonomous UAVs that follow the shortest collision-free path, as computed our solution described above. See Table 1 for the results from these comparisons, where a lower risk cost is desirable, and the rows correspond to cases a), b), and c) respectively. From the results in Table 1, clearly the human searchers benefit, in terms of risk, from using even manually controlled UAVs because of the high altitude coverage advantage they offer. The proposed method, wherein UAVs plan paths to explicitly minimize risk, performs better than all others by complementing

Table 1. Quantitative comparison results.


Scenario	Risk Cost ($\times 10^{18}$)	% of Max	Planning Time
No UAVs	4.846	100	N/A
UAVs, RRT*	4.412	91.044	20.397s
UAVs, Manual	4.122	85.056	N/A
UAVs, Risk	3.389	69.937	1193.781

searchers' efforts and controlling altitude effectively, balancing UAV field-of-view and quality of measurement. The scenarios considered here are reasonably consistent with how searches are currently performed within the SAR community (Van Tilburg 2017), and the results in Table 1 indicate that using autonomous robots, here modeled as UAVs, can greatly improve the success of the search mission via minimizing risk.

Edge computing for autonomous systems

Path planning, cognition, data reduction and human interfacing define the primary computational needs for a multi-robot SAR mission. Approaches for path planning, as described in the earlier section, can be accelerated by leveraging highly parallel processing capabilities offered by the modern GPUs (Bialkowski et al. 2011). Advances in machine learning techniques and specialized hardware for accelerating use cases such as image-based scene understanding have made it possible to detect and localize cues on-the-fly. This work scopes the development of custom deep neural networks for isolation and detection of cues with an objective to aid the human rescuer's understanding of the environment with sparse features of interest. The specialized networks are trained on aerial inputs from cameras working in the visible and the infrared (thermal) spectrum. The deep neural networks are expected to eliminate the need for domain experts to manually define features critical for the success of such a mission. Human interfaces that allow the rescuer to sufficiently understand large amounts of sensor data and computational results and effectively provide their feedback to the system also forms a significant part of this research.

A homogeneous system of UAVs for such missions, presents some unique physical and computational challenges. A representative modern drone, housing a high resolution camera, offers 15 minutes of flight time on average. The flight time is further diminished by additional payloads including auxiliary cameras, sensors and powerful compute devices expected to satisfy computational needs as outlined above. The payloads not only add on weight to the UAVs, thereby expending more energy during hovering and flight, but also consume from the available on-board power. The conditions in which such missions are expected to be undertaken also introduces the problem of limited connectivity to cellular network and cloud services. This enforces locality of computation and storage for the mission onboard the robots in the system. Brief periods of flight time, limited on-board resources and payload capacities, and disrupted connectivity limits the effectiveness of a mission in a SAR setting in terms of real-time performance and quality of cognition.

**Figure 3. Proposed System specifications for 4 UAVs, the Backpack Computing Unit and the Handheld Tablet**

This work attempts to counter the aforementioned challenges by offloading computational elements away from the UAVs. A custom computing unit housed inside of a backpack to be carried by a human rescuer is core to this strategy, depicted in Fig. 3. In this system, the role of on-board (UAV) computers may be limited to aggregation, preprocessing and efficiently routing the sensor data to the backpack. Additional computational demands may be met by the backpack. This proposed wearable unit comprises of four NVIDIA Jetson Xavier AGX computing modules, networking hardware and power supply. The computing modules, with a combined processing power of 128 Tera AI Operations Per Second (TOPS), allows for fulfilment of the computational needs in real-time. The backpack also forms the networking backbone among the Jetson modules, the UAVs and the handheld device for the rescue personnel. Such provisions offered by the backpack allows for reduced task loads for the drone-mounted computer and subsequently, the on-board power footprint.

With the proposed distributed heterogeneous system, by offloading computational and storage units to a rescuer backpack, we estimate a 14% increase in flight times. The backpack compute cluster also offers 22 times the number of AI operations per second when compared to the on-board computing available on the drones in a homogeneous system. The choice of hardware allows for this computational gains at as much as 10 times lesser operations per watt consumed.

Lost-person behavioral modeling

We develop an agent-based model of lost person behavior based on pedestrian dynamics and the International Search & Rescue Incident Database (ISRID) to inform both human searchers and UAV teams. ISRID is a database comprising data from more than 50,000 lost person searches which is partitioned into more than 30 lost person types, for example, hikers, hunters, and people with dementia (Koester 2008). The discrete-time model defines a lost person or “agent” as a self-propelled particle moving on a landscape according to realizations of a random variable which defines its behavior. In particular, the behavior employed at every time step is an independent realization of a random variable whose states are six of the known lost person strategies defined by Robert Koester (Koester 2008). Evolving this model from the lost person’s last known point allows us to generate a family of trajectories that defines a spatial distribution where the lost person can be expected as a function of time. Furthermore, the model is tuned using statistics from ISRID, including distributions for mobility (i.e. the time a lost person is mobile) and horizontal distance from the last known point. Through systematically simulating a set of trajectories for all possible behavior random variables and then selecting an optimal behavior through fitting with ISRID data, we can compare the model predictions to datasets from real incidents for a wide range of lost person types. This approach will give an expectation for lost person position at an arbitrarily higher spatiotemporal resolution, which is critical for aerial vehicles to have temporally-resolved information about lost person dynamics for their control.

The agent is simulated on a gridded map moving at discrete time steps to any of the eight adjacent cells around its position. In the current version of the model, an agent samples from a distribution of six strategies: random walking (RW), route traveling (RT), direction traveling (DT), staying put (SP), view enhancing (VE), and backtracking (BT) (Koester 2008; Hashimoto and Abaid 2019). Each lost person type (LPT) has a probability mass function (PMF) that captures the probability of the agent to use a specific behavior at a certain time. The PMF is a 6D vector of percentages of each of the strategies which sums to one. For example, a LPT with a probability of (RW, RT, DT, SP, VE, BT) = ($\frac{1}{2}$, 0, $\frac{1}{6}$, $\frac{1}{3}$, 0, 0) has a 50% chance of random walking, a 17% chance of direction traveling, and a 33% chance of staying put at a given time step. Independent realizations of this distribution of behaviors are generated at each time step and the agent’s position is updated accordingly.

We simulate all possible LPTs for 100 replicates and three different initial positions on the map, which represent the point the lost person would have last been seen. The simulation time of each replicate depends on the mobility cumulative distribution function (CDF) for each type of lost person (e.g. child, hiker, or person with dementia) given in ISRID (Koester 2008).

To explore the validity of the model, we need a metric to compare the simulated data from the model to the statistics from ISRID. In his book, Koester has summarized the horizontal distances traveled by the lost person as a CDF for each LPT (Koester 2008). We can use this to compute a PMF and compare it to the PMF of horizontal distances from each initial position to the final position in our model using the Kullback-Leibler (KL) divergence as a metric. For two discrete distributions P and Q , the symmetric KL divergence is defined as:

$$D_{\text{SKL}}(P \parallel Q) = \sum_{x \in \mathcal{X}} P(x) \log \left(\frac{P(x)}{Q(x)} \right) + \sum_{x \in \mathcal{X}} Q(x) \log \left(\frac{Q(x)}{P(x)} \right) \quad (2)$$

where $x \in \mathcal{X}$ are realizations of the distributions (MacKay 2003). We can use this measure of differences to show which model LPT fits best with distributions from ISRID.

The lost person categories we have simulated include children from three age groups, hikers, hunters, people with dementia, snowboarders, alpine skiers, anglers, despondents, and workers. Here, we highlight the results for the category of people with dementia. For a given initial position (IC 1), the CDF of horizontal distance from the database (blue stars) is shown with the CDFs of each of the LPTs computed using 100 Monte Carlo replicates in the left panel of Figure 4. The ISRID CDF is resampled to match the size of the simulation data and both the ISRID and simulation CDFs are numerically differentiated to find PMFs. The KL divergences are then computed on these PMFs to compare the ISRID PMF to the simulated ones pairwise. The right panel of Figure 4 shows the PMF of the best fit behavior plotted against the PMF from ISRID. We can see that the two distributions follow each other fairly well for this initial condition. The behavioral distribution for the best fit lost person type is (RW, RT, DT, SP, VE, BT) = ($\frac{1}{6}$, 0, $\frac{5}{6}$, 0, 0, 0). This means that the model predicts that a person with dementia would use 17% random walking and 83% direction traveling. This finding is consistent with the empirical observation from (Koester 2008) which says that people with dementia primarily use direction traveling when lost.

This procedure can be extended to lost people in different categories to provide insight into the behavioral strategies used for a specific search. Also, in future work, we will apply a similar modeling framework to build a dynamic model of human searcher behavior based on individualized data acquired during real and simulated searches.




Figure 4. Example results for a lost person with dementia. (Left) CDFs for horizontal distance from initial point from ISRID (blue stars) and all the lost person types with the best fit highlighted as red circles; (Right) PMFs for horizontal distance from initial point from ISRID and calculated for the best fitting LPT.

SAR mission web application

We are developing a web-based application to integrate our research in autonomous path-planning, mobile computing, and behavioral modeling for SAR professionals to utilize the unique and novel capabilities intuitively for completing their missions. Presently, our web application integrated with an geographical information system (GIS; <https://www.arcgis.com>) can support mission command (aka "management") in generating search segments for tasking, estimating lost-person location (i.e., probability of area), and assigning search tasks to teams.

GIS-based search segment generation

Every mission involves assigning teams of professionals and volunteers to search the surrounding of Last Known Position (LKP), where the lost person was, or suspected to be, last seen. The mission command assumes this responsibility by dividing surrounding of the LKP, known as the search area, into "search segments" according to historical statistical guidelines (Koester 2008) with respect to the available information on the lost person. After dividing the search area, the mission command assigns each search team to a segment according to the POA, the likelihood that a lost-person resides in a given region. The current practice in dividing the search area into segments involves a SAR professional manually drawing the segments on a paper map or in a SAR software program (e.g., Caltopo: <https://caltopo.com/>) while following the aforementioned historical statistical guidelines with respect to the subject's type, LKP, Point Last Seen (PLS), lost duration, etc.

Our web application automates the segmentation of search areas based on the statistical guidelines to alleviate workload so that SAR professionals can spend more time on critical human-centric activities (e.g., gathering additional information from friends relatives of the lost-person). To automatically divide the search area into



Figure 5. The web interface includes two main panels: 1) a map panel providing functions of displaying and editing information about the searching area, segments, team assignment, and basic map manipulation. 2) a task management panel including functions of manipulation of lost person model display, task management, team assignment, and some assistant tools

segments, the web application employs a Voronoi partitioning algorithm (Lévy and Liu 2010) with three mainly key steps:

- **Step 1:** Computer number of segments. The human specifies or inputs start location, the search area (i.e., the *length* and *width* of the searching area), size of a search segment (typically set to $U_{area} = 40$ acres for a search team to cover within four hours). These inputs permit computation of the total number of the searching segments for the mission.
- **Step 2:** Generate the center point of each segment. Given the number of segments, an array of concentric circles, denoted by $C = \{c_1, c_2, c_3, \dots, c_n\}$, sharing the same center point as the searching area (Figure 5). We use R to denote the radius of the first circle (i.e., $r_{c1} = R$). Radius of each circle in the array is set to $r_{cn} = n \times R$. The largest circle is inscribed within the search area (i.e., $n \times R \leq \min(Length, Width)$). Hence, for a given searching area with known *length* and *width*, we can calculate the number n of circles. (2.2) Then, for each annulus between two circles, we evenly distribute an amount of K segment central points along a circle in the middle of the annulus (i.e., for the ring between the first and second circle ($r_{c1} = R, r_{c2} = 2R$), the middle of the annulus is a circle of radius $1.5 \times R$). For each annulus between the n circle and the $n + 1$ circle, we calculate the size of the annulus as $\pi \times [(n + 1) \times R]^2 - (n \times R)^2$, and set K as the amount of search segments of size U_{area} in the annulus, i.e., $K_n = \lfloor \frac{\pi \times [(n+1) \times R]^2 - (n \times R)^2}{U_{area}} \rfloor$. (2.3) Finally, we stretch the center points of the segments obtained in the inscribed circles of the search area to their coordinating points within the inscribed ellipse of the searching region.
- **Step 3:** Apply Voronoi partitioning to generate each segment polygon.

The user interface of the web application displays full search area in satellite image view that contains the Initial Planning Point (IPP) marked as a red triangle and the 25%, 50%, 75% POA rings marked in red circles (Figure 5). The white polygons denotes the search segments for assigning to the search teams. To accommodate unique characteristics of each search mission (e.g., volunteer skills, geography) that cannot be anticipated with the algorithm, the web application provides control inputs for the users to alter each search segment by dragging and dropping the vertices of polygon. The web application thus aims to alleviate mission command from the workload of dividing search area into segments while retaining the flexibility to customize segments as necessary.

Visualization of lost-person behaviors

The web application further supports mission command with visualizations of the lost-person behavioral model (as described in Section) to aid assignment of search segments to individual search teams, which include multiple searchers and, in the future, UAVs. As mention, the lost-person behavioral model estimates the POA for every

location inside the search area for a given time. The web application can currently visualize the POA values with single marker for the highest likelihood of the lost-person location (Figure 6), and series of markers for a trajectory of the lost person over a period time (Figure 7), or a heat map for a region (Figure 8).




Figure 6. The most probable location in a given lost duration is labeled by a filled circle with different size and color to denote probability.




Figure 7. Given a time interval, lost person location is estimated, sampled, and linked to a trajectory.

Each visualization can provide unique advantages in supporting mission command to prioritize the segments for search team assignment. A single marker of the highest likelihood of the lost person location might help focus the attention of the mission commander. The series of markers illustrating likely path being traversed by the lost person might help commander to assignment of search segment over time. Finally, heatmap presenting POA distribution over the entire search area help commander to understand the relative importance of the segments. These visualizations should enable the mission command to utilize the lost person behavior model intuitively for prioritizing search segments effectively, thereby improving the chances of finding the lost person faster.

Search task assignment

The web application also supports mission command in assigning search task for each search team. To make an assignment, the user can drag and drop each black search team icon on the top of the user interface onto a search segment. After finishing allocating every team to a segment, the user can automatically generate a task assignment form designed according to the one used by Commonwealth of Virginia Department of Emergency Management (Figure 9). The digital form for each search segment automatically contains available or previously inputted task information, including the coordinates of vertices for each segment, task type, IPP, base, etc. Further, each field




Figure 8. A heat map of the lost person model, where the color scale can be changed based on user preference. In a rainbow color scale, the red area has the most probability.

remains editable for users to add and modify information. The task assignment functionality should further alleviate task assignment workload for the mission command to spend time for other tasks.

Figure 9. The task assignment form.

Our current research activity is conducting an experiment recruiting SAR professionals and college students to assess how the different visualizations could potentially alter the mission command in prioritizing which segments to assign to the search teams. In addition to the evaluation of the lost-person model visualization, we are currently developing the software and user interface for processing and displaying the UAV data streams (e.g., GPS, waypoints, video, etc.). The web application server is also being tested and migrated as the backpack computing unit is being developed. Thus, the web application will soon enable SAR professionals to assign UAVs for searching segments and observe sensor data as well as visualizing lost-person behavior and generating task assignment for human teams.

CONCLUSION

In this paper we outlined preliminary results from work in progress towards collaborative multi-robot multi-human teams in SAR. In future work, our immediate goals include: (1) deriving new algorithms for human-robot search task generation that build trust between autonomous systems and humans; (2) developing methods for balancing computation between our computational backpack and UAVs; and (3) extensive field testing in mock searches organized by our project team and held in coordination with SAR professionals.

REFERENCES

- 2016 NCIC Missing Person and Unidentified Person Statistics (2016). URL: <https://www.fbi.gov/file-repository/2016-ncic-missing-person-and-unidentified-person-statistics.pdf/view>.
- Bialkowski, J., Karaman, S., and Frazzoli, E. (Sept. 2011). “Massively parallelizing the RRT and the RRT”. In: *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3513–3518.
- Cangan, B. G. (2019). “Risk-Aware Human-In-The-Loop Multi-Robot Path Planning for Lost Person Search and Rescue”. PhD thesis. Virginia Tech.
- Chien, S.-Y., Wang, H., and Lewis, M. (Sept. 2010). “Human vs. Algorithmic Path Planning for Search and Rescue by Robot Teams”. In: *Proc. Hum. Fact. Ergon. Soc. Annu. Meet.* 54.4, pp. 379–383.
- Cooper, D. C., Frost, J. R., and Robe, R. Q. (Dec. 2003). *Compatibility of Land SAR Procedures with Search Theory*. Tech. rep. Potomac Management Group, Inc.
- De Cubber, G. (2017). *Search and Rescue Robotics - From Theory to Practice*. InTech.
- DroneSAR (2020). URL: <http://dronesarpilot.com/>.
- Feng, L., Wiltsche, C., Humphrey, L., and Topcu, U. (Apr. 2016). “Synthesis of Human-in-the-Loop Control Protocols for Autonomous Systems”. In: *IEEE Trans. Autom. Sci. Eng.* 13.2, pp. 450–462.
- Goodrich, M. A. and Schultz, A. C. (Jan. 2007). “Human-robot Interaction: A Survey”. In: *Found. Trends Hum.-Comput. Interact.* 1.3, pp. 203–275.
- Hashimoto, A. and Abaid, N. (Oct. 2019). “An agent-based model of lost person dynamics for enabling wilderness search and rescue”. In: *ASME Dynamic Systems and Control Conference*.
- Koester, R. J. (2008). *Lost Person Behavior: A Search and Rescue Guide on Where to Look - for Land, Air and Water*. Charlottesville, VA: dbS Productions LLC.
- Koppula, H. S., Jain, A., and Saxena, A. (2016). “Anticipatory Planning for Human-Robot Teams”. en. In: *Experimental Robotics*. Springer Tracts in Advanced Robotics. Springer, Cham, pp. 453–470.
- Kruijff, G. J. M., Janicek, M., Keshavdas, S., Larochelle, B., Zender, H., N J J, Mioch, T., Neerincx, M. A., Diggelen, J. V., Colas, F., et al. (2014). “Experience in System Design for Human-Robot Teaming in Urban Search and Rescue”. In: *Field and Service Robotics*. Springer Tracts in Advanced Robotics. Springer, Berlin, Heidelberg, pp. 111–125.
- Lévy, B. and Liu, Y. (July 2010). “Lp Centroidal Voronoi Tessellation and Its Applications”. In: *ACM Trans. Graph.* 29.4.
- Lewis, M., Wang, H., Chien, S.-Y., Scerri, P., Velagapudi, P., Sycara, K., and Kane, B. (Oct. 2010). “Teams organization and performance in multi-human/multi-robot teams”. In: *2010 IEEE International Conference on Systems, Man and Cybernetics*. ieeexplore.ieee.org, pp. 1617–1623.
- MacKay, D. J. (2003). *Information theory, inference and learning algorithms*. Cambridge University Press.
- NBC (2018). *Virginia Police Use Drone to Find Missing Boy*. URL: http://www.nbcwashington.com/news/local/Virginia-Police-Use-Drone-to-Find-Missing-Boy_Washington-DC-446220423.html#ixzz4tPoUw0mi.
- Nourbakhsh, I. R., Sycara, K., Koes, M., Yong, M., Lewis, M., and Burion, S. (Jan. 2005). “Human-robot teaming for search and rescue”. In: *IEEE Pervasive Comput.* 4.1, pp. 72–79.
- Shah, J., Wiken, J., Williams, B., and Breazeal, C. (2011). “Improved Human-robot Team Performance Using Chaski, a Human-inspired Plan Execution System”. In: *Proceedings of the 6th International Conference on Human-robot Interaction*. HRI ’11. ACM, pp. 29–36.
- Van Tilburg, C. (2017). “First report of using portable unmanned aircraft systems (drones) for search and rescue”. In: *Wilderness & environmental medicine* 28.2, pp. 116–118.
- Vempati, A. S., Agamennoni, G., Stastny, T., and Siegwart, R. (Dec. 2015). “Victim Detection from a Fixed-Wing UAV: Experimental Results”. en. In: *Advances in Visual Computing*. Lecture Notes in Computer Science. Springer, Cham, pp. 432–443.
- Yanco, H. A. and Drury, J. (Oct. 2004). “Classifying human-robot interaction: an updated taxonomy”. In: *2004 IEEE International Conference on Systems, Man and Cybernetics (IEEE Cat. No.04CH37583)*. Vol. 3. ieeexplore.ieee.org, 2841–2846 vol.3.