Anticipatory Planning and Dynamic Lost Person Models for Human-Robot Search and Rescue

Larkin Heintzman, Amanda Hashimoto, Nicole Abaid, and Ryan K. Williams

***Abstract*—In this work, we consider the problem of planning paths for a team of autonomous unmanned aerial vehicles (UAVs) to assist search and rescue practitioners. To address the problem, we develop a fully integrated framework that includes information from all aspects of the search environment. We take into consideration lost person motion via a behavior-based predictive model, anticipated human searcher trajectories, as well as measurements from fixed field of view sensors on board UAVs. We use a metric of posterior risk as the optimization target as it is an indicator of improved situational awareness and the effectiveness of continuing search efforts. Monte Carlo simulations are presented to demonstrate the effectiveness of the proposed framework.**

I. Introduction

In 2018, over 650,000 lost persons were reported in the United States, and nearly 100,000 of these cases were answered with organized searches in either urban or wilderness environments by the search and rescue (SAR) community [1]. Search efforts usually involve a large team, in some cases over 100 people, working over several days to locate the lost person [1], [2]. There have been efforts to increase the effectiveness of SAR operations, taking into account advances in search theory [2] and informative path planning [3]. In addition, an unmanned aerial vehicle (UAV) recently was used in a SAR mission that allowed searchers to inspect an area that otherwise would have been dangerous and time consuming to traverse [4]. However, the UAV in [4] was a remotely controlled model, requiring a human pilot to provide continuous control and image processing. In this work, we propose a framework wherein autonomous UAVs, with knowledge of the search scenario, can be used to improve situational awareness and support of human SAR teams.

UAVs have been used extensively for aerial monitoring and object detection tasks in recent years, especially with the advent of lightweight imaging systems [4]. Authors of the recent work [5] consider a related problem, the goal being to detect and track multiple surface vehicles at sea with an autonomous UAV. The major differences compared to this work are that we focus on planning *prior* to detecting a target and effectively collaborating with a ground team. Another perspective is seen in [6] where authors consider the more general problem of modeling interactions between first responders and a team of autonomous UAVs, with a goal of UAV task assignment for various mission scenarios.

Larkin Heintzman and Ryan K. Williams are with the Department of Electrical and Computer Engineering, Amanda Hashimoto is with the Engineering Mechanics Program, and Nicole Abaid is with the Department of Mathematics, Virginia Polytechnic Institute and State University, Blacksburg, VA USA, E-mail: {*hlarkin3, ahashimo, nabaid, rywilli1*}*@vt.edu*.

This work was supported by the National Science Foundation under grant CNS-1830414.

Fig. 1. Data flowchart overview of the proposed algorithm.

While undoubtedly necessary, the work of [6] functions at the sector assignment level whereas we consider search paths and lost person modeling within a single sector. Another similar problem to ours is considered in [7], where autonomous agents attempt to locate a target and transmit video information back to a base station. However, the authors of [7] do not consider models of lost person behavior nor collaboration with human search teams.

A lost person (LP) is defined as a person who is unable to identify their present location with respect to known locations and has no effective method for reorientation [8]. This lack of ability to reorient drives people to use a variety of different recorded behaviors. In *Lost Person Behavior*, Robert Koester has defined lost persons' behavioral strategies seen in a collection of data from the more-than-50,000-incident International Search and Rescue Incident Database (ISRID) [9], [10]. These behavioral profiles, or lost person types (LPTs), will be used to inform a model of the LP moving in the wilderness. The book identifies more than 30 LPTs, based on activity or demographic information, and reports summarized statistics, including the horizontal distance found from the last known point and time an LP was mobile. In practice, searchers use these statistics to create probability maps, like the distance-ring model [11] an example of which is shown in Figure 3b. However, these tactics assume that by the time the search has started the subject has stopped moving [12]. Effectively, motion should be considered in order to create a probability map that evolves in time. Existing models of human behavior have been used to study pedestrian dynamics, including force-based models showing collision avoidance, vision-based guidance, or goal-oriented behaviors [13]–[15]. Fundamentally different from pedestrian dynamics, modeling the behavior of an LP depends heavily on the landscape and the person type, as evidenced by the statistics in ISRID. The Bayesian model of such behavior in [16] considers terrain, but not strategies or lost person types. Alanis et al. developed a mechanistic model for a hiker that takes into account terrain influence and behavior, but only simulated one LPT and

neglected characteristic LP behaviors, like route traveling on known paths [17]. We seek to use a different approach from the previous works, by proposing an improved version of the dynamic model of LP behavior from [18] which synthesizes information about the specific environment for a search, as well as characteristics of different LPTs drawn from a large database of search incidents.

In order to complement human searchers and improve situational awareness, we propose a framework to plan a set of paths for a team of UAVs to autonomously assist a human ground team and to reduce the risk inherent in SAR scenarios [19]. Using a probabilistic risk metric derived from a Gaussian process model (GPM) coupled with gradient path optimization over parameterized Bézier curves, UAVs take into account a predictive LP model, topography in the search environment, and anticipated human searcher trajectories (see Figure 1). By utilizing the advantages afforded by UAVs, we aim to increase the *efficiency* of SAR operations by maximizing situational awareness with available resources while minimizing the aerial time required [2]. Here we build upon our preliminary work [20] by incorporating a much more in-depth treatment of the terrain, LP model, and evaluation methods.

The remainder of this work is organized as follows, in Section II we cover the background machinery required, and formally state the problem considered in this work. In Section III we discuss the path planning and optimization algorithms. In Section IV, we discuss the Monte Carlo based LP motion model. In Section V, we present several simulations to compare and contrast the proposed solution with a variety of other configurations and options. In Section VI, we offer closing remarks and suggestions for future work.

II. Modeling

A. Searcher Trajectories

In a land SAR mission, a particular area of land will have human searchers assigned to it, referred to as *sectors* [9]. We assume that each searcher has an entry and exit point for each sector, based on the overall sequence of sectors to be searched. To anticipate ground team paths, we propose a human searcher model with two modes, a waypoint following mode and a gradient following mode. In waypoint mode, each searcher is represented by a particle moving through a predefined set of waypoints. Once a searcher reaches to within some tolerance of the current waypoint, the next waypoint becomes the target. In our case, the waypoints are arranged to generate *lawn mowing* search paths. To better reflect reality, each searcher is also influenced by the terrain gradient. If the terrain beneath a searcher becomes too steep to climb, the searcher switches to level-curve following to navigate around the obstacle. There are scenarios in which a searcher could become trapped in a local minima of the terrain, which we combat via a tenacity parameter that allows searchers to overcome steeper terrain over time. The tenacity approach helps model scenarios where terrain is too steep for human searchers to easily navigate, while also preventing searchers from getting stuck. Shown in Figure 5 are two examples of terrain/feature maps from real locations with anticipated searcher paths in grey being affected by the terrain gradient. The proposed searcher model provides a tractable anticipation of human paths as inputs for our planning framework (see Algorithm 4 in [19] for a complete statement of the human searcher model).

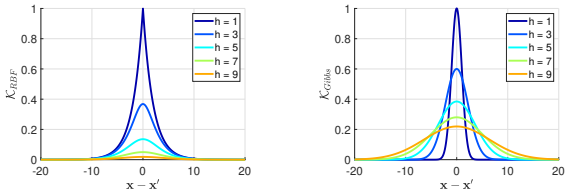

(a) (b)

Fig. 2. Showing in (a) a 2D measurement model that uses the RBF kernel, and in (b) a 2D measurement model that uses the Gibbs' kernel.

B. UAV Measurement Model

We use a GPM to combine the information from anticipated human searcher paths, UAV measurements, and the LP model (to be discussed in Section IV). Using a GPM allows us to model realistic data collection from UAVs via a *Gibbs'* kernel where the length scale is a function of altitude [21]–[23]. Here we model data collection from a downward facing sensor such as a camera, or any other fixed field of view (FOV) sensor that could be mounted on a UAV. With the proposed method, we can simulate data collection from any 3D point above the terrain. We use a Gibbs' kernel [24], as opposed to the commonly used radial basis function (RBF) kernel [25], because it allows for realistic modeling of *fixed* FOV sensors. The RBF kernel does not capture the change in FOV with changing altitude, thus the Gibbs' kernel is better suited here. The generic Gibbs' kernel is as follows:

$$\mathcal{K}_{\text{Gibbs}}(\mathbf{x},\mathbf{x}') = \prod_{d=1}^{D}\left(\frac{2l_d(\mathbf{x})l_d(\mathbf{x}')}{l_d(\mathbf{x})^2+l_d(\mathbf{x}')^2}\right)^{\frac{1}{2}}\cdot\exp\left(-\sum_{d=1}^{D}\frac{(x_d-x_d')^2}{l_d(\mathbf{x})^2+l_d(\mathbf{x}')^2}\right) \tag{1}$$

where D denotes the number of dimensions (3 in this case) with $\mathbf{x}=[x_1,x_2,x_3]^\top$, and $l_d(\mathbf{x}):\mathbb{R}^{3\times1}\to\mathbb{R}$ is an arbitrary function of $\mathbf{x}$ such that $l_d(\mathbf{x})>0$, $\forall x$. To adapt to our specific case of covariance calculation from a UAV, we take $l_1(\mathbf{x})=l_2(\mathbf{x})=f(x_3)$ and $l_3=\gamma$. Then the kernel is as follows:

$$\mathcal{K}_{\text{SAR}}(\mathbf{x},\mathbf{x}') = \frac{2f(x_3')f(x_3)}{f(x_3')^2+f(x_3)^2}\cdot\exp\left(-\frac{(x_1'-x_1)^2)}{f(x_3')^2+f(x_3)^2}-\frac{(x_2'-x_2)^2}{f(x_3')^2+f(x_3)^2}-\frac{(x_3'-x_3)^2}{2\gamma^2}\right) \tag{2}$$

where $\mathbf{x},\mathbf{x}'\in\mathbb{R}^{3\times1}$ are two arbitrary 3D points in the environment, γ is a tunable hyperparameter, and $f(\cdot)$ is the lengthscale function which in this case is the altitude above the terrain. Figure 2 has a visual comparison between the RBF and Gibbs' kernels for the 2D case at various altitudes; notice when the altitude decreases the width of measurement also decreases while the covariance increases, following intuition for a fixed FOV sensor.

C. Problem Statement

Given an environment of interest, a rectangular area defined by GPS referenced points, build a set of trajectories for a team of N aerial agents that effectively complement human searcher efforts and minimize risk. The risk metric in this case is the uncertainty in LP location, thus reducing risk increases environmental awareness and search effectiveness. The metric of risk includes information from all sources, the LP heatmap, anticipated human searcher paths, and UAV trajectories.

In more formal terms, generate trajectories, parameterized by $\lambda \in \mathbb{R}^{p\times 3}$ to generate $\theta_\lambda(t)$, for all UAVs that approximately minimizes the following:

$$\lambda^* = \underset{\lambda\in\Lambda}{\arg\min}\, \mathcal{R}\left[\theta_\lambda(t)\right] \tag{3}$$

where $\mathcal{R}$ is the metric of *risk* which quantifies our uncertainty in the location of the LP(s), the details of which are discussed in the next section. The result of solving this problem is a set of trajectories for all UAVs, $\theta_{\lambda^*}(t)$, that complement human searchers and improve situational awareness due to optimized UAV measurements.

III. Path Planning

A. Initial Robot Paths

Initial planning is necessary due to the potential of a UAV colliding with the terrain or obstacles. We compute the initial set from the sampling based approach of RRTs, in this case we use RRT* [19]. In our framework, we independently plan paths for all robots while avoiding inter-agent collisions, using the same sector entry and exit points as the human searchers, and quantify the information gathered from all paths using the GPM.

B. Risk Quantification

To optimize UAV paths, we must quantify the *risk* inherent in a given scenario and trajectory set. Risk indicates the uncertainty in LP location *given* the current system configuration and prior knowledge, making risk uniquely suited to replanning efforts as well. We use a GPM, and Gibbs' kernel, to calculate the mean and covariance of an occupancy map quantifying our *belief* in the LP location. Both humans and UAVs use the Gibbs' kernel to generate measurements, with a fixed altitude of 2m for the human searchers. The state, measurement, and inference points are in the form of:

$$\begin{aligned}\mathbf{x} &= \left[\mathbf{x}_{\text{MC}}^\top, \mathbf{x}_{\text{HS}}^\top, \mathbf{x}_{\text{RB}}^\top\right]^\top \\ \mathbf{y} &= \left[\mathbf{y}_{\text{MC}}^\top, \mathbf{y}_{\text{HS}}^\top, \mathbf{y}_{\text{RB}}^\top\right]^\top \\ \mathbf{x}^* &= \left[x_1^*, x_2^*, \ldots, x_n^*\right]^\top\end{aligned} \tag{4}$$

where $\mathbf{x}_{\text{MC}} \in \mathbb{R}^{n_{\text{MC}}\times 3}$ is the set of points from the LP heatmap above a tunable threshold $\tau \in [0,1]$, $\mathbf{x}_{\text{HS}} \in \mathbb{R}^{n_{\text{HS}}\times 3}$ is the set of human searcher paths, and $\mathbf{x}_{\text{RB}} \in \mathbb{R}^{n_{\text{RB}}\times 3}$ is the (initial) set of UAV trajectories. The term $\mathbf{x} \in \mathbb{R}^{l\times 3}$ is referred to as the state vector, only the UAV trajectory portion is optimized, and $l = n_{\text{MC}}+n_{\text{HS}}+n_{\text{RB}}$. As for the measurements, $\mathbf{y}_{\text{MC}} \in \mathbb{R}^{n_{\text{MC}}\times 1}$ is the set of log-odd probabilities for occupancy at each point in $\mathbf{x}_{\text{MC}}$, $\mathbf{y}_{\text{HS}} \in \mathbb{R}^{n_{\text{HS}}\times 1}$ is the set of log-odds for points in $\mathbf{x}_{\text{HS}}$, and similarly $\mathbf{y}_{\text{RB}} \in \mathbb{R}^{n_{\text{RB}}\times 1}$ is the set of log-odds for points in $\mathbf{x}_{\text{RB}}$. Lastly, $\mathbf{x}^* \in \mathbb{R}^{n\times 1}$ is the set of inference points, derived from discretizing the search sector based on resolution parameter δ_n, where the occupancy map is evaluated. The purpose of the occupancy map is to estimate the *posterior* probability $p(m|\mathbf{y},\mathbf{x})$ where m represents the map state for points in $\mathbf{x}^*$. We use the *log-odds* form of probabilities to simplify the coming calculations, defined for probability p as $\text{logit}(p) = \log(p/(1-p))$.

Using the terms in (4), we can calculate terms to create a metric of risk. The covariance of the state is calculated by $\mathbf{K}_{\mathbf{x}} = \mathcal{K}_{\text{SAR}}(\mathbf{x},\mathbf{x}) + \sigma^2\mathbf{I}$ where $\mathbf{I}$ is the square identity matrix of size l. Here $\sigma^2 \in \mathbb{R}^{l\times 1}$ is the measurement variance, another tunable parameter. Similarly, the covariance of the inference points is calculated by $\mathbf{K}_{\mathbf{x}^*} = \mathcal{K}_{\text{SAR}}(\mathbf{x}^*,\mathbf{x}^*)$, and cross-covariance between state and inference points $\mathbf{K}_{\mathbf{x}|\mathbf{x}^*} = \mathcal{K}_{\text{SAR}}(\mathbf{x},\mathbf{x}^*)$.

We can now calculate the mean and covariance of the occupancy map, in log-odds form, as $\mu_{\mathbf{x}^*} = \mathbf{K}_{\mathbf{x}|\mathbf{x}^*}^\top \mathbf{K}_{\mathbf{x}}^{-1}\mathbf{y}$ and $\mathbf{\Sigma}_{\mathbf{x}^*} = \mathbf{K}_{\mathbf{x}^*} - \mathbf{K}_{\mathbf{x}|\mathbf{x}^*}^\top \mathbf{K}_{\mathbf{x}}^{-1}\mathbf{K}_{\mathbf{x}|\mathbf{x}^*}$ respectively. Where $\mu_{\mathbf{x}^*}$ is the log-odds ratio of occupancy for each point in $\mathbf{x}^*$, this can be thought of as the log-odds result of combining all information sources. A value of $\mu_{x_i^*} = 0$ indicates a 0.5 probability that the point x_i^* is occupied by an LP. Similar to [26], we develop our risk metric as:

$$\mathcal{R} = \sum_{i=1}^{n} \frac{\Sigma_{\mathbf{x}_i^*}}{1+\mu_{\mathbf{x}_i^*}^2} \tag{5}$$

where $\Sigma_{\mathbf{x}_i^*}$ is the i^{th} diagonal element of the covariance matrix $\mathbf{\Sigma}_{\mathbf{x}^*}$.

C. Objective Function

As discussed, we have an initial set of trajectories that are optimized to reduce the *risk cost*. The goal of the optimization process is to generate UAVs trajectories that effectively complement the ground searchers, subject to constraints. The constraints are start and end points, smoothness, and time expenditure (path length). We parameterize UAV paths with Bézier curves, which enables evaluation of smoothness and path length [27] and provides a set of p parameters $\lambda \in \mathbb{R}^{p\times 3}$, which define trajectories $\theta_\lambda(t)$. Our goal is to find parameters, λ, to minimize the risk cost $\mathcal{R}\left[\theta_\lambda(t)\right]$. Incorporating constraints and writing as an objective functional:

$$\mathcal{F}\left[\theta_\lambda(t)\right] = \mathcal{R}\left[\theta_\lambda(t)\right] + \alpha_{\mathcal{L}}\mathcal{L}\left[\theta_\lambda(t)\right] + \alpha_{\mathcal{S}}\mathcal{S}\left[\theta_\lambda(t)\right] \tag{6}$$

where $\mathcal{L}\left[\theta_\lambda(t)\right]$ is the path length cost, and $\mathcal{S}\left[\theta_\lambda(t)\right]$ is the path smoothness cost, $\alpha_{\mathcal{L}}$ and $\alpha_{\mathcal{S}}$ are scaling terms for the path length and the path smoothness costs, respectively. The restated goal is to minimize the objective functional $\mathcal{F}\left[\theta_\lambda(t)\right]$ to maximize the efficiency of the planned UAV trajectories.

1) Trajectory Optimization:

$$\lambda_{i+1} = \lambda_i - \eta\nabla_\lambda\mathcal{F}\left[\theta_\lambda\right], \quad i = 1,2,\ldots \tag{7}$$

where η is a tunable descent rate parameter. Since the objective functional is of high dimension and non-convex, we use an optimization method typically found in neural network-based

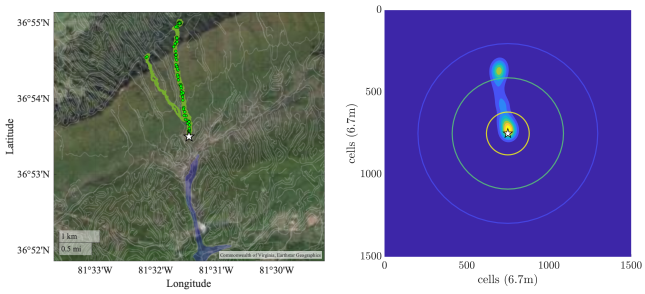

(a) Terrain Map (b) LP Heatmap and Ring Model

Fig. 3. Lost person model results. In (a), the linear features (grey lines), inaccessible areas (shaded blue), MC path trajectories (shaded green), final MC trajectory points (green circles), and IPP (white star) are overlaid on a satellite image of Hungry Mother State Park. In (b), the corresponding LP model heatmap of the final MC trajectory points and the circles representing the ring model for a hiker LP type are shown.

machine learning applications, referred to as *Adam* [28]. Even using Adam, the optimization process can be quite slow so, in the inference step, points from the grid are sorted based on the Morton Z-order space filling curve [29]. Using this technique allows us to approximate large blocks of the matrix as zeros as they correspond to spatially distant points.

IV. Lost Person Model

People lost in the wilderness use one or more typical behaviors to move around [8]. In this section, we present a model of how an LP uses these behaviors, which is an improved version of the model from [18]. The LP is modeled as a self-propelled agent moving in discrete time on a 2D grid. At every time step, the agent can move from its position to to any of its neighboring cells or the current cell. Here the map is informed not only by elevation, but also by USGS layers for roads, railroads, powerline easements, hiking trails, rivers, and lakes. These layers define the map's linear features as common routes often followed by an LP as a strategy for navigation. The boundaries of rivers and lakes are separated from the interior areas to simulate riverbank linear features and inaccessible areas. These added features provide a realistic simulation environment.

A. Behaviors

In the LP model, we have defined six possible strategies that are based on [9]: *Random Walking* (RW), *Route Traveling* (RT), *Direction Traveling* (DT), *Staying Put* (SP), *View Enhancing* (VE), and *Backtracking* (BT). Each LPT has a 6D vector PMF that captures the probability of the agent using each strategy over time. Independent realizations of this distribution are generated at each time step and the agent's position is updated based on the strategy. 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.

In a realization of the model, the initial and second position of the agent $x(1) \in \mathbb{N}^2$ and $x(2) \in \mathbb{N}^2$ are selected as two

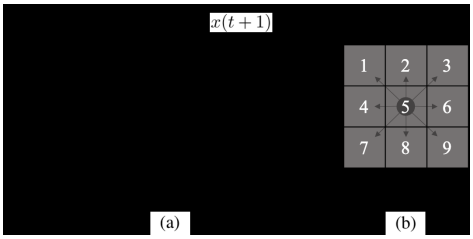

Fig. 4. The table in (a) lists the PMFs for each behavior strategy which are defined by the 9 ordered positions of the 3×3 grid, shown in (b), representing the agent's location in body coordinates relative to the entire map.

adjacent cells. At each time step t, the velocity is computed as $v(t) = x(t+1) - x(t)$, and we use a smoothing factor α to factor in the previous velocity to make the trajectories more realistic. The updated position at each t is computed as $x(t+1) = (2-\alpha)x(t) + (\alpha-1)x(t-1) + \alpha v(t)$. Placing the agent at the center of a 3×3 grid, in body coordinates, we generate the next step by selection of one of the six strategies. The PMFs for each behavior strategy, listed in Figure 4a, are defined by the order of the positions $(1, 2, 3, 4, 5, 6, 7, 8, 9)$ for the cells neighboring the agent, shown in Figure 4b. At each t, an independent realization of the PMF is generated, defining the strategy the LP will use to update its position $x(t+1)$ and velocity $v(t+1)$. The model is iterated for T time steps to generate the agent's entire trajectory.

B. Mobility

The length of each simulation, T, is dependent upon the LP mobility data from [9], given as CDFs of 25%, 50%, 75%, and 95% quartiles for each LPT's mobile time, with 50% as the median. We take the mobility CDF for each LPT and fit it with a Weibull distribution, defined as $\text{cdf} = 1 - e^{-(\frac{x}{a})^b}$ with a and b as fit parameters. Depending on the selection of a and b, the distribution can mimic shapes from exponential to Gaussian. Using MATLAB's curve fitting toolbox, we find the best fit parameters to define a Weibull PDF, that we then sample to find the simulation lengths for each LPT.

C. Map

All possible trajectories for a LP are simulated on a map generated from USGS geographic information system (GIS) data. The data provides map layers at a given resolution, roads, railroads, powerlines, trails, and water features. Along with structural linear features, additional paths are derived from elevation, such as drainages and mountain peaks, and are visible as critical points in the magnitude of the gradient of the elevation. After the gradient field's magnitude is computed and smoothed using the derivative of a Gaussian filter, linear features are found using the Canny Edge Detection method by locating edges in the local maxima and minima of the scalar field. By including these terrain features, the model from [18] becomes more realistic since landscape necessarily affects lost person behavior [8]. In Figure 3a, the combined map layers representing the linear features (grey lines) and inaccessible areas (shaded blue) are shown over a satellite image of the area.

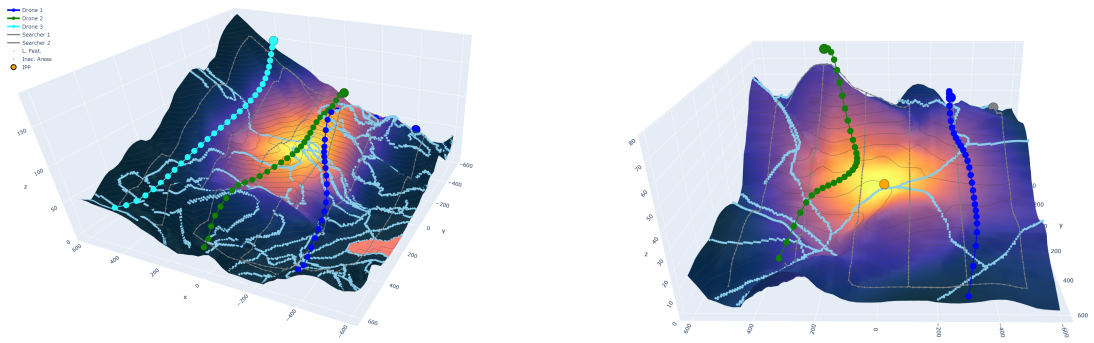

(a) H.M. State Park (b) Kentland Farms

Fig. 5. Showing in (a) a full pipeline test for the H.M. state park location with a 1.2km extent, and in (b) a full pipeline test for the Kentland Farms location with a 1.2km extent.

In an actual search, searchers need to know the initial location of the LP, whether it is the point last seen, given by an eyewitness, or the last known point where there is substantial evidence to place the subject. This point is called the initial planning point (IPP) and measures the progress of the search. In this model, the initial positions of the agent are assumed to be the IPPs.

D. Behavioral profiles

For this work, we simulate one type of LP, a hiker, based on the statistics from [9], [10]. A hiker was chosen due to the large amount of data for its type. In order to fit a hiker's behavioral profile, all permutations of the six behaviors are generated as LPT PMFs with each strategy weight as a multiple of $\frac{1}{6}$. We have a set of 462 LPTs with varying proportions of each behavior. Each LPT distribution is simulated for 100 iterations for each initial location. For each LPT and initial location, the total distance covered is determined, allowing us to compare the model performance when different behavioral profiles are used. To evaluate the model, we use the Kullback-Leibler (KL) divergence to compare the simulated data from the model to the statistics from ISRID. Using statistics on the distances from the IPP to each find location from [9], we compute a database PMF and compare it to the model PMF from simulation. For two discrete distributions P and Q, the symmetric KL divergence is defined as:

$$D_{\text{SKL}}(P \parallel Q) = \sum_{x \in \mathcal{X}} P(x) \log\left(\frac{P(x)}{Q(x)}\right) + \sum_{x \in \mathcal{X}} Q(x) \log\left(\frac{Q(x)}{P(x)}\right) \tag{8}$$

where $x \in \mathcal{X}$ are realizations of the distributions [30]. We can use this measure of differences to show which model LPT fits best with distributions from ISRID.

Using the KL divergence metric with the horizontal distance from the IPP, we can find the best behavioral profile. For a given initial position, the database CDF of horizontal distance is resampled to match the size of the simulation data and both the ISRID and simulation CDFs are numerically differentiated to find PMFs. Then we interpolate both the model and database distributions over their overlapping bins to ensure that neither distribution has any zeros, which may cause the KL divergence to go to infinity. The KL divergences are then computed on these PMFs to compare the ISRID PMF to the simulated ones, pairwise, allowing us to find the best fitting model LPT for a hiker type.

E. Generating the occupancy map

Once the behavioral profile is determined for the hiker, the model is initialized at an IPP and a probability field is generated for the lost person over the search area. We run 100 replicates of the Monte Carlo simulation and generate a probability distribution from the final locations of each simulation. This is used as the occupancy map for the path planning algorithm.

V. Simulations

A. Implementation

The components of our framework, LP model, human searcher model, and GPU gradient-based optimization were implemented in Python v3.6.8 and MATLAB 2020a. The code used to generate all results shown is available at our Gitlab[1], the ARCGIS layer information was generated with AGS Tools.[2]

B. Background

For the coming simulated experiments, the parameters used are shown in Table I. The terrain and features are generated using real data from ARCGIS services.[3] In Table I, η refers to the learning rate in (7), τ is the heatmap threshold, the FOVs are the detection radii for UAVs and searchers. We are comparing the LP model against the current SAR standard, the ring model, as shown in Figure 3b. The ring mobility distance values are listed in [9] for a typical lost hiker moving on hilly terrain. Later we will compare the proposed method against two other methods, basic lawn-mower sweeps (the *Sweep* method) and manually controlled drones (the *RC* method). In the sweep case, the UAV team traverses the environment in a set number of passes, with fixed height independent of the human searchers. The sweep case provides a comparison to a more brute force approach to searching. In the RC case, UAVs hover above searchers at a fixed height, due to the pilots' need to remain within line of sight.

[1]Gitlab link text: https://git.caslab.ece.vt.edu/hlarkin3/planning_llh_bgc
[2]AGS link text: https://git.caslab.ece.vt.edu/hlarkin3/ags_grabber
[3]The majority of layers were derived from https://hydro.nationalmap.gov

$x_{\min}, x_{\max}$:	$(-600, 600)$m	Sweep Alt. :	15m
$y_{\min}, y_{\max}$:	$(-600, 600)$m	Opt. Iter. :	100 epochs
η:	5	l_2:	5×10^{-3}
$\alpha_{\mathcal{C}}$:	1×10^{7}	$\alpha_{\mathcal{S}}$:	1×10^{3}
δ_n:	25	γ :	5
N (UAVs):	$\{1, 2, 3, 4, 5, 6\}$	S (Searchers):	$\{1, 2\}$
α :	0.55	T :	850
MC Reps :	100	LP type :	hiker
UAV FOV:	5m	S. FOV:	10m
Train Batch %:	0.90	Test Batch %:	0.10
Ring Mobi.:	$\{0.6, 1.8, 3.2\}$km	τ:	0.2
S. Sweep #:	3	UAV Sweep #:	4
H.M. Park Lat:	36.891640	H.M. Park Lon:	-81.524214
Kentland Lat:	37.197730	Kentland Lon:	-80.585233

TABLE I
SIMULATION PARAMETERS

We chose the Hungry Mother (H.M.) state park in Smyth County, VA as a test location because SAR practitioners frequently use the area to perform mock search training exercises. The park also presents a varied foothill environment with linear features and bodies of water. The second location is Kentland Farms, VA which was chosen as a flatter environment with fewer roads, for comparison. For the lost person model, we find a best fit behavioral profile for the simulated searches in each location. The profiles for H.M. and Kentland are $(\frac{1}{6}, \frac{1}{6}, 0, 0, \frac{1}{6}, \frac{1}{2})$ and $(0, 0, \frac{2}{3}, \frac{1}{6}, 0, \frac{1}{6})$, respectively.

C. Qualitative Results

Comparing the heatmaps generated via the LP model to the state-of-the-art ring model demonstrates the value of a dynamic LP model. In Figure 3b, the final locations are based on a specific behavior profile, and they are contained in a dramatically smaller area, as compared to the ring model. This improved specificity comes from the LP model taking into account the surrounding landscape and linear features, whereas the ring model is based solely on the distance from the IPP. The improved situational awareness afforded to the path planning algorithm from this model allows for more effective allocation of searchers.

Two example scenarios are shown in Figure 5. The terrain is displayed as a surface, upon which the LP heatmap is shown as the surface color. In addition to the heatmap, the terrain features are plotted on the surface as light blue lines and red polygons, indicating linear features and inaccessible areas respectively, and the IPPs are plotted as orange dots. In Figure 5, the UAV paths indicate effective planning in 3 dimensions as shown by the UAVs covering high probability points while controlling altitude. Note that the green UAV path in Figure 5b is drawn towards the high probability areas of the LP heatmap, whereas the blue UAV path is covering many less probable locations at a higher altitude. This clearly demonstrates that risk is being mitigated and the search task load is being distributed.

D. Quantitative Results

To evaluate the proposed method, we compare the risk metric and path length of all solutions. See Figure 6 for results from the risk metric in (6), where a lower value indicates lower cost. Each entry in the comparison plots was averaged over 10 iterations and the error bars indicate ± 1 standard deviation. The scale of plots in Figure 6 is due to weighting terms used in objective functional (6) (see Table I) which are necessary to match the scale of the risk metric itself.

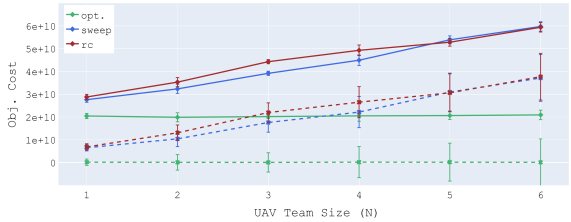

(a) Objective functional, H.M.

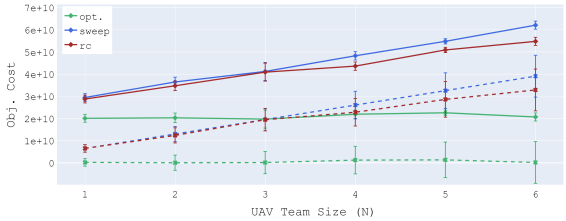

(b) Objective functional, Kentland

Fig. 6. Showing objective function costs under increasing UAV team sizes for the two locations. The solid lines indicate the overall cost value while the dashed lines show the length cost contribution.

By the comparisons shown, clearly the optimized case performs best for all team sizes examined. As expected the naïve methods, sweep and RC, perform significantly worse in overall cost, and note the major contribution of path length. In the optimized case, the overall cost remains fairly constant for both locations as each additional vehicle reduces the risk cost but also increases the path length cost. The specific performance of each case is highly location based, as shown in Figure 6b at $N = 6$ when the overall cost of the optimized case decreases due to the environment becoming saturated with aerial agents.

In the proposed method, the risk is maintained at a significantly lower level than all other cases whilst incurring a fraction of the path length penalties. This outcome is achieved by complementing searchers' efforts and controlling altitude effectively, balancing FOV and quality of measurement. The scenarios considered here are reasonably consistent with how searches are currently performed within the SAR community [4], and the results in Figures 6a and 6b indicate that using autonomous aerial agents, here modeled as UAVs, can greatly improve the effectiveness of search missions via risk minimization.

VI. CONCLUSIONS

In this work, we presented a framework to plan paths for autonomous UAVs to complement human searchers and improve situational awareness during a SAR scenario. Included in the framework are models of dynamic LP motion given environmental factors, anticipatory human searcher models, and expected measurements from an aerial fixed FOV sensor. We used a probabilistic risk metric to optimize UAV trajectories and assess performance in reference to typical SAR practices.

REFERENCES

[1] Federal Bureau of Investigation (FBI), "2018 NCIC Missing Person and Unidentified Person Statistics," National Crime Information Center, Tech. Rep., Mar 2018. [Online]. Available: https://www.fbi.gov/file-repository/2018-ncic-missing-person-and-unidentified-person-statistics.pdf/view

[2] D. C. Cooper, J. R. Frost, and R. Q. Robe, "Compatibility of Land SAR Procedures with Search Theory," Potomac Management Group, Inc., Tech. Rep., Dec 2003.

[3] A. Singh, A. Krause, and W. J. Kaiser, "Nonmyopic adaptive informative path planning for multiple robots," in *Twenty-First International Joint Conference on Artificial Intelligence*, 2009.

[4] C. Van Tilburg, "First report of using portable unmanned aircraft systems (drones) for search and rescue," *Wilderness & environmental medicine*, vol. 28, no. 2, pp. 116–118, 2017.

[5] F. S. Leira, H. H. Helgesen, T. A. Johansen, and T. I. Fossen, "Object detection, recognition, and tracking from UAVs using a thermal camera," *Journal of Field Robotics*, 2020.

[6] A. Agrawal, J.-P. Steghofer, and J. Cleland-Huang, "Model-Driven Requirements for Humans-on-the-Loop Multi-UAV Missions," *arXiv preprint arXiv:2009.10267*, 2020.

[7] J. Scherer, S. Yahyanejad, S. Hayat, E. Yanmaz, T. Andre, A. Khan, V. Vukadinovic, C. Bettstetter, H. Hellwagner, and B. Rinner, "An autonomous multi-UAV system for search and rescue," in *Proceedings of the First Workshop on Micro Aerial Vehicle Networks, Systems, and Applications for Civilian Use*, 2015, pp. 33–38.

[8] K. Hill, "The psychology of lost," *Lost person behavior*, p. 116, 1998.

[9] R. J. Koester, *Lost Person Behavior: A search and rescue guide on where to look - for land, air and water*. dbs Productions LLC, 2008.

[10] "International Search and Rescue Incident Database (ISRID)," http://www.dbs-sar.com/.

[11] R. Koester, *Lost person behavior: Instructor activity guide*. dbS Productions LLC, 2010.

[12] E. Sava, C. Twardy, R. Koester, and M. Sonwalkar, "Evaluating lost person behavior models," *Transactions in GIS*, vol. 20, no. 1, pp. 38–53, 2016.

[13] A. Seyfried, A. Schadschneider, U. Kemloh, and M. Chraibi, "Force-based models of pedestrian dynamics," *Networks and Heterogeneous Media*, vol. 6, no. 3, pp. 425–442, 2011.

[14] M. Moussaïd, D. Helbing, and G. Theraulaz, "How simple rules determine pedestrian behavior and crowd disasters," *Proceedings of the National Academy of Sciences*, vol. 108, no. 17, pp. 6884–6888, 2011.

[15] I. Karamouzas, B. Skinner, and S. J. Guy, "Universal power law governing pedestrian interactions," *Physical Review Letters*, vol. 113, no. 23, p. 238701, 2014.

[16] L. Lin and M. A. Goodrich, "A Bayesian approach to modeling lost person behaviors based on terrain features in wilderness search and rescue," *Computational and Mathematical Organization Theory*, vol. 16, no. 3, pp. 300–323, 2010.

[17] J. Alanis, M. Brown, J. Kitchens, J. Magana, C. Velastegui, M. Thakura, L. Arriola, B. Espinozaa, A. Murilloa, M. Rodriguez-Messana *et al.*, "Topography and behavior based movement modeling for missing hikers in land-wilderness settings," 2019.

[18] A. Hashimoto and N. Abaid, "An agent-based model of lost person dynamics for enabling wilderness search and rescue," ser. Dynamic Systems and Control Conference, vol. 2. American Society of Mechanical Engineers, 2019. [Online]. Available: https://doi.org/10.1115/DSCC2019-9222

[19] B. G. Cangan and R. Williams, "Risk-aware human-in-the-loop multi-robot path planning for lost person search and rescue," Ph.D. dissertation, Virginia Tech, 2019.

[20] B. G. Cangan, L. Heintzman, A. Hashimoto, N. Abaid, and R. K. Williams, "Anticipatory human-robot path planning for search and rescue," *arXiv e-prints*, pp. arXiv–2009, 2020.

[21] C. J. Paciorek and M. J. Schervish, "Nonstationary covariance functions for Gaussian process regression," in *Advances in neural information processing systems*, 2004, pp. 273–280.

[22] M. Heinonen, H. Mannerström, J. Rousu, S. Kaski, and H. Lähdesmäki, "Non-stationary Gaussian process regression with Hamiltonian Monte Carlo," in *Artificial Intelligence and Statistics*. PMLR, 2016, pp. 732–740.

[23] C. E. Rasmussen, "Gaussian processes in machine learning," in *Summer School on Machine Learning*. Springer, 2003, pp. 63–71.

[24] M. N. Gibbs, "Bayesian Gaussian processes for regression and classification," Ph.D. dissertation, Citeseer, 1998.

[25] D. Duvenaud, "Automatic model construction with Gaussian processes," Ph.D. dissertation, University of Cambridge, 2014.

[26] M. Popović, T. Vidal-Calleja, G. Hitz, I. Sa, R. Siegwart, and J. Nieto, "Multiresolution mapping and informative path planning for UAV-based terrain monitoring," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 1382–1388.

[27] S. Vincent and D. Forsey, "Fast and accurate parametric curve length computation," *Journal of graphics tools*, vol. 6, no. 4, pp. 29–39, 2001.

[28] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

[29] G. M. Morton, "A computer oriented geodetic data base and a new technique in file sequencing," 1966.

[30] D. J. MacKay, *Information theory, inference and learning algorithms*. Cambridge University Press, 2003.