

From Detection to Traversal: A Probabilistic Framework for UAS-Assisted Landmine Mapping and Circumvention

J. Josiah Steckenrider¹, Dongbin Kim², Pratheek Manjunath³,
 Jack Simpson³, Charles Rosbrook¹, and Micah Weiss³

Abstract—This research presents a robust probabilistic framework for minefield localization, mapping, and avoidance, addressing a technological gap in the field of aerial counterintelligence, while bypassing the well-established techniques of landmine detection. Our approach propagates the pose uncertainty matrix delivered by a drone’s flight controller’s Kalman filter to probabilistically estimate the location of detected mines. This probability map then seeds an artificial potential field path generator which creates a safe path for ground traversal by producing waypoints through the minefield. The system’s performance is evaluated in simulations and validated through flight trials, demonstrating its potential to improve the efficiency and safety of UAV-assisted minefield navigation and threat avoidance.

I. INTRODUCTION

A. Background and Motivation

Minefields pose a significant threat to both soldiers and civilians on and near modern battlefields. In recent years, as peer and near-peer adversaries have increased the use of landmines for both anti-vehicle and anti-personnel purposes, mapping and avoiding mines has never been more important. For example, it is estimated that approximately 30% of Ukraine was covered by minefields in early 2024, with landmine-caused civilian casualties numbering near 1,000 and rising [1]. Many of these devices are designed to be lethal against dismounted individuals, in violation of the 1997 Anti-Personnel Mine Ban Treaty [2]. Because manual minefield inspection and neutralization are inherently some of the most dangerous activities a soldier can undertake, the need for a robotic solution is both obvious and urgent. While the long-term solution to the landmine problem requires the methodical removal of each explosive device, a short-term necessity is the ability to map out and safely navigate through a minefield. Unmanned aerial vehicles (UAVs, or drones) are ideally suited for counterintelligence due to their ability to cover large areas from an aerial perspective.

B. Related Work

Robotic landmine detection is by no means a new area of research. For at least the past few decades, researchers and engineers have tasked robots and autonomous agents with

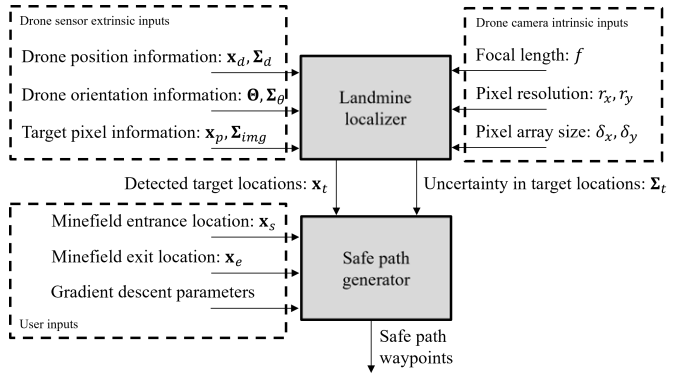


Fig. 1. Diagram of probabilistic aerial counterintelligence framework

identifying ground-based and subsurface explosives on active battlefields and formerly contested areas. The bulk of older work in this field utilized unmanned ground vehicles (UGVs) for mine detection, as they have historically been more capable of closely approaching and interrogating explosive devices [3]–[7]. In more recent years, the use of UAVs for counterintelligence efforts has increased, keeping pace with advancements in drone stability and controllability [8]–[14]. Many of these efforts focus on the low-level challenges involved in landmine detection. For example, some emphasize the design and construction of the robots [5], [6], [12], while others take a systems-level approach [10], [15], and still others focus on the sensing and detection task [8], [9], [13], [14]. Although some research addresses higher-level path planning and navigation objectives [4], [7], [11], [16]–[18], little is being done to develop autonomous end-to-end frameworks for aerial searching and mapping of a minefield to find safe routes through it.

The majority of existing UAV-based landmine detection techniques tend to overlook the mine localization problem—i.e., estimating the global position of a detected target given its position in the frame of a sensor onboard the robot. In fact, most of the aforementioned works neglect the issue altogether. A reasonable first-pass approximation is to assign the GPS coordinates of the searching drone to a mine when it is detected. However, this approach only works when the drone is at low altitudes directly above the mine. As the position and orientation of the drone change, this approach becomes insufficient. The most accurate estimation of target locations relies on precise estimation of the search agent’s position and orientation. Simultaneous localization

*This work was supported by US Army DEVCOM C5ISR

¹Department of Civil and Mechanical Engineering, United States Military Academy, West Point, NY, USA
 john.steckenrider@westpoint.edu

²College of Engineering, Technology, and Architecture, University of Hartford, West Hartford, CT, USA

³Department of Electrical Engineering and Computer Science, United States Military Academy, West Point, NY, USA

and mapping (SLAM) has been the standard approach for robot navigation and mapping of the surrounding environment, especially in GPS-denied indoor environments [19], [20]. Although it could provide a solution for minefield mapping, SLAM is computationally expensive and excessive for outdoor environments where GPS or visual/inertial navigation systems are available. Custom-built, application-specific frameworks can typically perform just as well or better, at a fraction of the computational cost.

In summary, while most related works focus on basic landmine detection and neglect higher autonomous functions, the research presented in this paper takes the opposite approach: we disregard the more primitive elements of mine detection that have been abundantly demonstrated over the past 20 years and instead focus on the novel contributions of localization, mapping, and path generation for countermine activity. There is a technological gap in a richly mathematical and theoretically robust minefield mapping framework, which we aim to fill with this research.

C. Original Contributions

There are two major original contributions of this work. The first contribution is specific and technical: we formulate and validate a probabilistic approach to estimating landmine locations from the air using multidimensional uncertainty propagation. The second contribution is more holistic: we present a mathematically cohesive method of generating a safe path through a minefield, where the only human input is the tuning of a few parameters. Under the proposed approach, a deployed drone with basic navigation and imaging sensors can autonomously operate an entire mission and deliver the waypoints of a safe path upon completion. Figure 1 illustrates the basics of the presented framework at a high level.

This paper does not address motion planning or obstacle avoidance for UAVs nor does it focus on landmine detection using electro-optical (EO) cameras. While sophisticated sensing techniques exist for landmine detection, our work abstracts this problem by using an EO camera and an orange traffic cone as a stand-in for a detected threat. This allows us to tackle the ensuing steps after a ground-based threat has been identified, rather than the detection process itself. The probabilistic path generation is to assist ground-based forces to maneuver around threats, and not for the UAS to follow.

This paper is organized as follows. Section II details the formulation of the probabilistic landmine localization problem and the natural implementation of the potential field algorithm to generate a safe path through a minefield. Evaluation of the system is conducted in simulation in section III, followed by an experimental validation in section IV. Finally, we conclude in section V with a discussion of the limitations of and future work related to this research.

II. LANDMINE LOCALIZATION AND PATH GENERATION

The classical approach to environment mapping utilizes sequential observations of a landmark or target as it is tracked through time and space, usually by an extended Kalman filter or similar estimator. While such methods are excellent for

converging certainty about the location of an environmental feature, they require accurate knowledge about not only the 6D pose of the agent but also its dynamics and control inputs. The proposed novel approach for landmine localization and mapping is formulated to eliminate the need for a temporally linked tracker/filter by aggregating independent observations made at each time step and weighting them by their uncertainties. Thus, the cornerstone of this technique is multidimensional uncertainty propagation. Rather than deliver a list of estimated coordinates of all unique detected landmines, this approach instead produces a probability density function (PDF) describing the probable locations of landmines. As section II-D will show, this lends naturally to gradient-based path generation.

A. Aerial 3D Geometry




Fig. 2. The three-dimensional geometry of aerial landmine localization.

For a static grounded target, we wish to estimate its global position $\mathbf{x}_t = [x_t \ y_t \ z_t]^\top$ relative to a Newtonian frame \mathcal{F} and the corresponding covariance matrix Σ_t which characterizes the uncertainty in that estimate. Let $\mathbf{x}_d = [x_d \ y_d \ z_d]^\top$ be the global 3D position of a drone observing the target (Fig. 2) and \mathbf{R} be the rotation matrix between the drone's non-Newtonian \mathcal{F}' frame and \mathcal{F} . The global location of the target can then be obtained by:

$$\mathbf{x}_t = \mathbf{x}_d + \mathbf{R}\mathbf{x}'_t, \quad (1)$$

where $\mathbf{x}'_t = [x'_t \ y'_t \ z'_t]^\top$ is the 3D position of the target in the \mathcal{F}' frame. The coordinates x'_t and y'_t can be obtained using the formula for 3D image projection with a pinhole camera model:

$$\begin{bmatrix} x'_t \\ y'_t \end{bmatrix} = \frac{z'_t}{f} \mathbf{x}_p, \quad (2)$$

where f is the focal length of the camera lens and $\mathbf{x}_p = [x_p \ y_p]^\top$ are the coordinates of the target in the image formed on the camera's image plane. These can be obtained by multiplying the target's pixel coordinates by the ratio of the usable physical size of the pixel array to the pixel resolution.

If the drone's camera parameters, a target's pixel coordinates, and the 6D position and orientation of the drone are all known, four unknown variables remain: x_t , y_t , z_t , and z'_t . Without further information, the above equations cannot be solved. To circumvent this problem, we enforce $z_t = 0$, a natural constraint consistent with a grounded

target.¹ Leveraging this assumption, z'_t can be found to be:

$$z'_t = -\frac{z_d}{\mathbf{R}_{3\mathbf{p}}}, \quad (3)$$

where \mathbf{R}_3 is the third row of the rotation matrix \mathbf{R} and $\mathbf{p} = \frac{1}{f}[x_p \ y_p \ f]^\top \equiv [p_1 \ p_2 \ p_3]^\top$. Following some substitution, it can be shown that

$$\mathbf{x}_t = \mathbf{x}_d - \frac{z_d}{\mathbf{R}_{3\mathbf{p}}}\mathbf{R}\mathbf{p}. \quad (4)$$

An additional assumption that should be noted here is that the origin of the drone's odometry frame is collocated with the focal point of the camera. While this assumption breaks down in practice (see section V), it simplifies subsequent analysis without loss of generality and with relatively minor consequences in application.

B. Uncertainty Propagation

Uncertainty in the target's position vector depends functionally on several variables, namely the contents of \mathbf{x}_d , \mathbf{R} , and \mathbf{p} and their covariances. The covariance matrix Σ_t describing the second moment of \mathbf{x}_t is approximated by:

$$\Sigma_t = \mathbf{J}_t \Sigma_{tot} \mathbf{J}_t^\top, \quad (5)$$

where

$$\begin{aligned} \mathbf{J}_t &= \begin{bmatrix} \frac{\partial \mathbf{x}_t}{\partial \mathbf{x}_d} & \frac{\partial \mathbf{x}_t}{\partial \mathbf{R}_1} & \frac{\partial \mathbf{x}_t}{\partial \mathbf{R}_2} & \frac{\partial \mathbf{x}_t}{\partial \mathbf{R}_3} & \frac{\partial \mathbf{x}_t}{\partial \mathbf{p}} \end{bmatrix} \\ &= \begin{bmatrix} 1 & 0 & \frac{-\mathbf{R}_1 \mathbf{p}}{\mathbf{R}_{3\mathbf{p}}} & \frac{-z_d \mathbf{p}^\top}{\mathbf{R}_{3\mathbf{p}}} & \mathbf{0} & \frac{z_d \mathbf{R}_1 \mathbf{p} \mathbf{p}^\top}{(\mathbf{R}_{3\mathbf{p}})^2} & \frac{z_d}{\mathbf{R}_{3\mathbf{p}}} (\frac{\mathbf{R}_1 \mathbf{p}}{\mathbf{R}_{3\mathbf{p}}} \mathbf{R}_3 - \mathbf{R}_1) \\ 0 & 1 & \frac{-\mathbf{R}_2 \mathbf{p}}{\mathbf{R}_{3\mathbf{p}}} & \mathbf{0} & \frac{-z_d \mathbf{p}^\top}{\mathbf{R}_{3\mathbf{p}}} & \frac{z_d \mathbf{R}_2 \mathbf{p} \mathbf{p}^\top}{(\mathbf{R}_{3\mathbf{p}})^2} & \frac{z_d}{\mathbf{R}_{3\mathbf{p}}} (\frac{\mathbf{R}_2 \mathbf{p}}{\mathbf{R}_{3\mathbf{p}}} \mathbf{R}_3 - \mathbf{R}_2) \\ 0 & 0 & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix} \end{aligned} \quad (6)$$

and

$$\Sigma_{tot} = \begin{bmatrix} \Sigma_d & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \Sigma_R & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \Sigma_p \end{bmatrix}. \quad (7)$$

In Eq. (7), Σ_d is the 3×3 covariance of the drone's position and Σ_R and Σ_p are defined as follows:

$$\Sigma_R = \mathbf{J}_R \Sigma_\Theta \mathbf{J}_R^\top \quad (8)$$

and

$$\Sigma_p = \mathbf{J}_p \Sigma_{img} \mathbf{J}_p^\top, \quad (9)$$

where Σ_Θ is the 3×3 covariance of the drone's Euler angle orientation and Σ_{img} is the 2×2 covariance of the target position on the image plane. The Jacobians \mathbf{J}_R and \mathbf{J}_p are:

$$\mathbf{J}_R = \begin{bmatrix} \frac{\partial R_{11}}{\partial \Theta} & \frac{\partial R_{12}}{\partial \Theta} & \frac{\partial R_{13}}{\partial \Theta} & \frac{\partial R_{21}}{\partial \Theta} & \frac{\partial R_{22}}{\partial \Theta} & \frac{\partial R_{23}}{\partial \Theta} & \frac{\partial R_{31}}{\partial \Theta} & \frac{\partial R_{32}}{\partial \Theta} & \frac{\partial R_{33}}{\partial \Theta} \end{bmatrix}^\top, \quad (10)$$

$$\mathbf{J}_p = \begin{bmatrix} \frac{\partial p_1}{\partial x_p} & \frac{\partial p_1}{\partial y_p} \\ \frac{\partial p_2}{\partial x_p} & \frac{\partial p_2}{\partial y_p} \\ \frac{\partial p_3}{\partial x_p} & \frac{\partial p_3}{\partial y_p} \end{bmatrix} = \begin{bmatrix} \frac{1}{f} & 0 \\ 0 & \frac{1}{f} \\ 0 & 0 \end{bmatrix}. \quad (11)$$

¹While this assumption strictly speaking limits the approach to planar terrains, mildly nonplanar ground surfaces can be handled so long as terrain variations are small in comparison with the sea-level altitude of the drone.

The vector $\Theta = [\theta \ \phi \ \psi]^\top$ contains the yaw, roll, and pitch of the drone, respectively.

C. Target Belief Distribution

It should be evident that the third row and column of Σ_t are zeros, consistent with the fact that $z_t = 0$ was an imposed constraint and therefore completely certain. As such, only the upper-left 2×2 submatrix is useful for characterizing the x - y uncertainty of the target's 2D location. By conducting a multidimensional uncertainty analysis in this way, each observation of a landmine in every frame contributes to an evolving PDF of landmine locations. This PDF is given by a multi-gaussian distribution defined as:

$$p_k(\mathbf{x}_t) = \frac{N_{1:k-1}}{N_{1:k}} p_{k-1}(\mathbf{x}_t) + \sum_{n=1}^{N_k} \frac{1}{N_{1:k}} \mathcal{N}(\mathbf{x}_t; \bar{\mathbf{x}}_{t_k}^{(n)}, \Sigma_{t_k}^{(n)}), \quad (12)$$

where $p_k(\mathbf{x}_t)$ is the PDF describing target belief at time step k and $\mathcal{N}(\mathbf{x}; \mu, \Sigma)$ denotes a normal distribution in \mathbf{x} with mean μ and covariance Σ . In this formulation, $N_{1:k-1}$ is the number of observations over all time steps from 1 to $k-1$, $N_{1:k}$ is the total number of observations through the current time step, $\bar{\mathbf{x}}_{t_k}^{(n)}$ is the n^{th} observed target location at time step k , and $\Sigma_{t_k}^{(n)}$ is the covariance of that observation as delivered by Eqs. (5-11).

As Fig. 3 shows, aggregating landmine observations into a multi-gaussian PDF, where each observation's covariance matrix is obtained by the uncertainty propagation described above, results in a probability distribution that naturally reflects the relative confidence of near and far observations. Per the figure, there is a radial structure to $p_{45}(\mathbf{x}_t)$ with higher certainty near the drone, demonstrating the consistency between intuition and mathematical formulation.

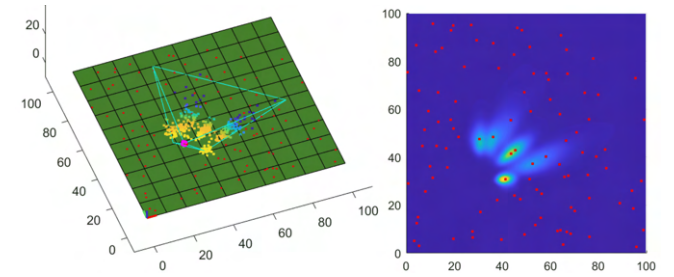


Fig. 3. Simulated example of $N_{1:k} = 360$ observations over $k = 45$ time steps. Left figure shows all observations scatter-plotted with each observation's color proportional to $\text{tr}(\Sigma_{t_k}^{(n)})$; right figure shows $p_{45}(\mathbf{x}_t)$ plotted as a heat map with true mine locations overlaid as red points.

D. Path Generation

The potential field algorithm is a technique for smoothly planning a path through obstacles which treats objects to avoid as repulsive field generators and destinations as attractive field generators [21]. Given some continuously differentiable surface \mathcal{S} which describes an artificial potential field as a function of an agent's position \mathbf{x} , the position at

step $k + 1$ can be determined from the position at k by gradient descent:

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \delta_k \nabla(\mathcal{S}(\mathbf{x}))|_{\mathbf{x}=\mathbf{x}_k}, \quad (13)$$

A key advantage of the proposed landmine mapping approach is that the multi-gaussian representation of belief lends to an efficient closed-form gradient-based path generation technique. Following a complete landmine mapping mission, the landmine location belief map constructed by observations over all time steps from all agents is a gaussian sum, consistent with Eq. (12), taking the following general form:

$$\mathcal{S}_m(\mathbf{x}) = \sum_{i=1}^N \frac{1}{N} \mathcal{N}(\mathbf{x}; \bar{\mathbf{x}}_i, \Sigma_i). \quad (14)$$

This mine surface $\mathcal{S}_m(\mathbf{x})$ can be thought of as an artificial gravitational potential field with “high” regions near areas of large target probability. In order to incentivize the gradient of this field toward the desired minefield exit location \mathbf{x}_e , a funnel surface $\mathcal{S}_f(\mathbf{x})$ with slope λ is formulated as follows:

$$\mathcal{S}_f(\mathbf{x}) = \lambda \|\mathbf{x} - \mathbf{x}_e\|. \quad (15)$$

Let $\mathcal{S}(\mathbf{x}) = \mathcal{S}_m(\mathbf{x}) + \mathcal{S}_f(\mathbf{x})$ as depicted in Fig. 4. This surface constitutes a combined artificial potential field which leads a “test charge” toward the exit while avoiding likely mine locations. The gradient of \mathcal{S} is given by:

$$\nabla(\mathcal{S}(\mathbf{x})) = \lambda \frac{\mathbf{x}}{\|\mathbf{x} - \mathbf{x}_e\|} - \sum_{i=1}^N \frac{1}{N} \mathcal{N}(\mathbf{x}; \bar{\mathbf{x}}_i, \Sigma_i) \Sigma_i^{-1} (\mathbf{x} - \bar{\mathbf{x}}_i). \quad (16)$$

Using the gradient descent algorithm of Eq. (13), a path can be generated to span from a starting location \mathbf{x}_s to the exit location \mathbf{x}_e , so long as there are no local minima in $\mathcal{S}(\mathbf{x})$. This caveat, however, is not guaranteed. To avoid “dead ends”, artificial landmine observations are placed at \mathbf{x}_k when $\|\mathbf{x}_k - \mathbf{x}_{k-1}\|$ falls below some threshold. Although this introduces a new parameter, the algorithm is relatively insensitive to reasonably small values (i.e. on the order of 1,000 times smaller than the search domain). Finally, as an additional constraint to ensure the safe path remains within the boundaries of the searched area, \mathbf{x}_k is clipped at the extremes of the search space on every iteration.

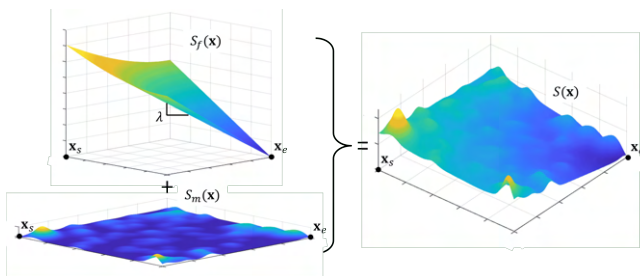


Fig. 4. Example of an artificial potential field generated by a simulated landmine search.

III. DEMONSTRATION OF CONCEPT: SIMULATION

A. Methodology

In order to test the proposed framework without the confounding factors of physical experimentation, a simulation environment was developed in MATLAB. Although the simulator does not yet incorporate flight physics, the effect on the search problem is negligible if desired flight speeds are low to moderate. Since the low-level problem of detecting landmines precedes the scope of this research, bright red squares are used as proxies for landmines (in practice, multi-spectral imaging is more commonplace for landmine detection). The simulator animates a bird’s-eye view of the search mission, showing the agent, its projected field of view (FOV) boundaries, true landmine locations, and the locations of transformed observations colored by their certainty level (see Fig. 3, left). Additionally, the searcher’s 2D FOV is converted to an image with the resolution of the simulated camera (Logitech C920 webcam) and given some gaussian blur of radius τ and noise to emulate a real imaging device. Some simple color-based thresholding is then implemented to binarize the image, from which binary object centroids are extracted. This flow is shown in Fig. 5.




Fig. 5. Color-based mine detection on simulated images.

To understand the effect of noise on the performance of the framework, 25 missions with uniformly randomly located mines were simulated over varying levels of uncertainty in Σ_{img} and Σ_{Θ} . These two covariance matrices are the most challenging to estimate in practice, as uncertainty in detected pixel coordinates and 3D drone angles is difficult to obtain experimentally. Thus, these covariances were modeled as identity matrices scaled by factors γ and α , respectively. The drone position covariance matrix Σ_d was given values typical of those delivered in the experiment by the flight controller’s Kalman filter (see section IV). For every simulated landmine observation and coordinate transformation, noise on the simulated drone’s 6D pose was drawn from a multivariate Gaussian distribution corresponding to these covariances. Each simulated mission was completed over a $W \times H$ rectangular search area by a parallel sweep (or “lawnmower”) pattern simply to ensure complete and uniform coverage of the minefield. Then, 200 entrance and exit points for the safe path were selected at edges of the minefield opposite each other at regular intervals 360° around the domain. For each safe path generated according to section II-D, the closest Euclidean distance d_{min} that the path came to any mine was recorded, in addition to the percent increase $\%L$ in total path length as compared to the minimum straight-line path from entrance to exit. These were then averaged over all 200 paths. The key parameters of these simulations are provided in Table I. The simulated camera had a focal length $f = 3.67$

mm, 640×480 resolution, diagonal FOV angle of 78° , and usable sensor size of 5.56 mm.

TABLE I
SIMULATION PARAMETERS

W	H	# of mines	Mine width	Altitude
50 m	50 m	100	32 cm	7 m
τ	Σ_{img}	Σ_Θ	λ	δ_k
2.5 pix	$\gamma\mathbf{I}$	$\alpha\mathbf{I}$	0.001	200

B. Results and Discussion

Figure 6 shows examples from three simulated missions of varying drone orientation uncertainties. The paths across the minefields are plotted over the PDFs of estimated mine locations, showing how the artificial potential field approach naturally avoids areas of high mine probability. It is clear from the figure that there are limits on path performance when the entropy (spread) of the distribution is too high or too low. In the former case, the gradient of $\mathcal{S}_m(\mathbf{x})$ is too shallow to reliably direct the path around mines. However, when the entropy is too low, the path is allowed to come much nearer to the mines before being directed away. These observations are further supported by Table II. It is important to note, however, that these trends are specific to the arbitrarily selected density of mines covering the minefield (4 mines per 100 square meters).




Fig. 6. Three example trials with increasing values of α . All 200 paths through the simulated minefields are shown overlaid on the mine heatmap $\mathcal{S}(\mathbf{x})$ with true mine locations designated by red points.

TABLE II
SIMULATION RESULTS (MAX VALUES OF RESPONSE VARIABLES IN BOLD)

		α					
		0.05	0.0158	0.005	0.00158	0.0005	
γ	0.01	d_{min} :	0.265	0.846	1.282	1.427	1.400
		$\%L$:	1.424	9.808	11.644	13.142	9.192
	0.0032	d_{min} :	0.242	0.673	1.260	1.255	1.182
		$\%L$:	1.738	11.955	12.245	10.319	8.805
	0.001	d_{min} :	0.321	0.674	1.259	1.092	0.945
		$\%L$:	1.972	12.210	14.666	14.480	10.334
	0.00032	d_{min} :	0.296	0.637	1.244	1.068	0.863
		$\%L$:	2.182	12.098	10.601	9.100	13.091
	0.0001	d_{min} :	0.285	0.784	1.242	0.976	0.764
		$\%L$:	1.867	8.593	9.708	10.853	8.349



Fig. 7. Flight test field: 16 orange traffic cones simulate landmines, with a drone at the bottom right

As the table shows, the average minimum distance to mines d_{min} and the percent path elongation $\%L$ depend more on α than on γ . This confirms that the proposed landmine estimation framework is more sensitive to error in drone orientation estimation than to pixel-wise error in image-based detection. Furthermore, the average minimum distance to mines is achieved at moderate levels of uncertainty in drone orientation, in keeping with the intuition provided by Fig. 6. There is somewhat of a trade-off between d_{min} and $\%L$: the further a path keeps away from mines, the longer it tends to be. However, since avoiding a mine is orders of magnitude more important than taking a shorter path, and considering the maximum percent elongation in the table is still less than 15%, not much is lost by designing for maximal d_{min} .

One final note is appropriate here. Although Σ_d , Σ_Θ , and Σ_{img} should be estimated experimentally to best characterize the sensors in use, there is a case to be made for tuning Σ_Θ and Σ_{img} heuristically. This is because $\Sigma_{img} = \gamma\mathbf{I}$ has relatively little effect on the shape of the safe path, and $\Sigma_\Theta = \alpha\mathbf{I}$ can be adjusted by α to obtain a surface $\mathcal{S}(\mathbf{x})$ optimally contoured for gradient-based path generation given a certain density of mines across the search domain. Tuning α will not significantly affect the *accuracy* of landmine estimation, but it can help achieve a mine PDF which is more conducive to route planning.

IV. EXPERIMENTAL VALIDATION AND VERIFICATION

A. Methodology

Several flight trials were conducted at the United States Military Academy at West Point to validate and verify the proposed framework using real data and devices (see Fig. 7). The drone was constructed with a Tarot 650 quadcopter frame, a Pixhawk 2.1 Blue Cube autopilot, a webcam, and an Odroid XU-4 onboard computer. The onboard computer ran ROS Noetic and OpenCV for drone flight path control and real-time color-based target detection, as described in Section III-A. The pixel coordinates of all detected objects in the drone's FOV were saved in JSON format. Also included were the drone's latitude and longitude (in meters Easting and Northing) and altitude, in addition to its Euler-angle orientation and the 6D pose covariance matrices. Data were saved and processed offline, where timestamps allowed for

interpolative synchronization.²

The framework was tested incrementally in experiments. First, a drone flew in a circle above a traffic cone with a constant yaw (i.e., always facing east) to isolate the effects of agent rotation from translation on landmine estimation. Then, the drone flew the same path with a constantly increasing yaw. Next, the drone flew a lawnmower pattern over a field containing 16 surrogate landmines with zero yaw, and finally, it flew a Lissajous search pattern where the heading of the drone faced in the direction of its trajectory (course over ground). Examples of these four trials are provided in Fig. 8. For these trials, $\alpha = 0.05$ and $\gamma = 0.0001$.

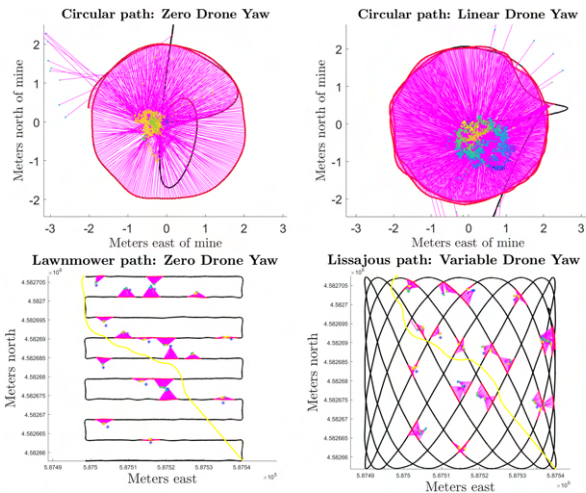


Fig. 8. Four sample trials. Black lines are drone paths, red dots are drone locations where targets were detected, multicolored dots are estimated true landmine locations (colored by certainty level), magenta lines connect drone locations where targets were detected to estimated true landmine locations, and blue asterisks are ground truth landmine locations.

B. Results and Discussion

The first two trials in Fig. 8 confirm the ability of the proposed formulations to both transform pixel coordinates of detected targets into the global frame and propagate the uncertainty in these observations consistent with their accuracy. This is demonstrated by the color of the transformed observations: those nearer to the true mine location are more yellow (certain) than those farther from it. Furthermore, there is a natural robustness to false positives in the proposed approach. Spurious detections get drowned out in $\mathcal{S}_m(\mathbf{x})$ both by the aggregation of many more true positives and by the higher certainty assigned to more accurate observations.

Figure 9 shows the mine PDFs generated by the data collected during the two full search missions. A few features are worth pointing out. First, there is a slight apparent northwest shift in the peaks of the “hot spots” of the PDF as compared to the true mine locations for the lawnmower pattern. This trend does not hold true for the Lissajous pattern. It is likely that this bias is an artifact of the small offset between the GPS and the camera focal point on the

²The separation of sensing from computation was purely logistical; data processing could take place in real time without significant changes.

drone, combined with the constant yaw of the agent in this experiment. Since the heading of the searcher did not change, the bias was consistent for all detections. However, since the angle of the drone was constantly changing for the Lissajous pattern, the bias varied.

Furthermore, note that the values of the peaks in the mine PDF delivered by the Lissajous pattern are somewhat inconsistent in comparison with the lawnmower pattern. This is because the density of the lawnmower path is more uniform across the search space, while the Lissajous curve sees increased coverage at the borders of the domain. Since the drone spent more time flying over the easternmost mine, more observations of it were acquired; this led to higher saturation of the PDF there. Practically, the effect of non-uniform coverage on the generated safe path is minimal.

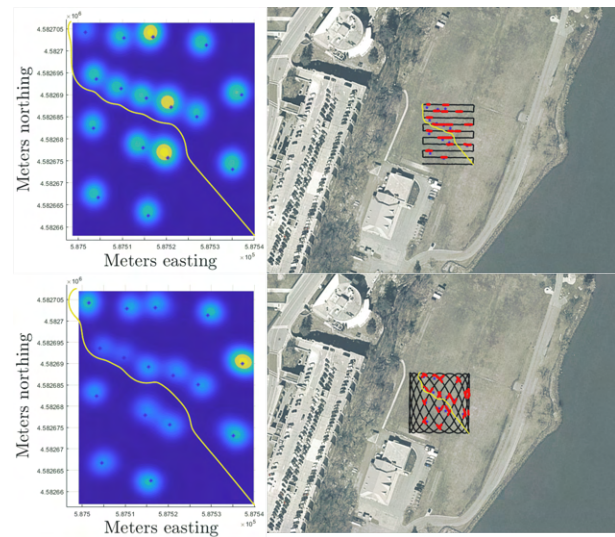


Fig. 9. Left: mine PDFs with safe ground traversal paths plotted in yellow and true mine locations denoted with blue asterisks. Right: satellite images of actual drone search paths (lawnmower and Lissajous) at the site of the experiments with red points indicating drone locations where mines were detected. The yellow lines show the safe traversal path on the ground, plotted from southeast to northwest corner.

V. CONCLUSIONS, LIMITATIONS, AND FUTURE WORK

This paper has presented a cohesive framework for autonomous aerial vehicles to search, map, and plot a safe path through a minefield with minimal need for user intervention. The validity of the approach has been confirmed through both simulation and experimentation, both of which support the effectiveness of the original formulations for probabilistically estimating landmines from the air. While the current phase of work processes data offline, the formulations are naturally hands-off and can readily be implemented in real time onboard unmanned aerial vehicles with minimal intervention.

There are a number of limitations to the framework that make excellent candidates for future research thrusts. First, the assumption of a coincident camera focal point and odometry frame on the autonomous agent can be addressed in future improvements to the formulation of the approach. Second, the equations presented in this paper were

derived from a first principles pinhole camera model. A more robust approach for future consideration would be to use an empirically derived intrinsic camera matrix obtained through calibration tests on real hardware. Third, the effects of unequally distributed mine observations as discussed in Section IV-B, could be mitigated by revising Eq. (14) to normalize for searcher coverage density. Last, although the framework is holistically designed to operate autonomously, the few parameters (mainly λ , α , and δ) that affect the quality of the safe path could be more carefully studied to develop automated tuning mechanisms.

REFERENCES

- [1] "Two years after the Russian invasion, land mines plague one-third of Ukraine." <https://www.nbcnews.com/investigations/two-years-russian-invasion-landmines-plague-one-third-ukraine-rcna138517>.
- [2] "Anti-Personnel Landmines Convention – UNODA." <https://disarmament.unoda.org/anti-personnel-landmines-convention/>.
- [3] H. Najjaran and A. A. Goldenberg, "Landmine detection using an autonomous terrain-scanning robot," *Industrial Robot*, vol. 32, no. 3, pp. 240–247, 2005.
- [4] K. L. Su, H. S. Su, S. W. Shiao, and J. H. Guo, "Motion planning for a landmine-detection robot," *Artificial Life and Robotics*, vol. 16, pp. 277–280, dec 2011.
- [5] M. A. Jaradat, M. N. BaniSalim, and F. H. Awad, "Autonomous navigation robot for landmine detection applications," *2012 8th International Symposium on Mechatronics and its Applications, ISMA 2012*, 2012.
- [6] V. Abilash and J. Paul Chandra Kumar, "Ardunio controlled landmine detection robot," *ICONSTEM 2017 - Proceedings: 3rd IEEE International Conference on Science Technology, Engineering and Management*, vol. 2018-January, pp. 1077–1082, jun 2017.
- [7] I. A. Hameed, "Motion planning for autonomous landmine detection and clearance robots," *2016 International Workshop on Recent Advances in Robotics and Sensor Technology for Humanitarian Demining and Counter-IEDs, RST 2016*, mar 2017.
- [8] J. Rodriguez, C. Castiblanco, I. Mondragon, and J. Colorado, "Low-cost quadrotor applied for visual detection of landmine-like objects," *2014 International Conference on Unmanned Aircraft Systems, ICUAS 2014 - Conference Proceedings*, pp. 83–88, 2014.
- [9] C. Castiblanco, J. Rodriguez, I. Mondragon, C. Parra, and J. Colorado, "Air Drones for Explosive Landmines Detection," *Advances in Intelligent Systems and Computing*, vol. 253, pp. 107–114, 2014.
- [10] Y. Ganesh, R. Raju, and R. Hegde, "Surveillance Drone for Landmine Detection," *Proceedings - 2015 21st Annual International Conference on Advanced Computing and Communications, ADCOM 2015*, pp. 33–38, aug 2016.
- [11] J. Colorado, C. Devia, M. Perez, I. Mondragon, D. Mendez, and C. Parra, "Low-altitude autonomous drone navigation for landmine detection purposes," *2017 International Conference on Unmanned Aircraft Systems, ICUAS 2017*, pp. 540–546, jul 2017.
- [12] R. K. Megalingam, D. Nagalla, P. R. Kiran, R. T. Geesala, and K. Nigam, "Swarm based autonomous landmine detecting robots," *Proceedings of the International Conference on Inventive Computing and Informatics, ICICI 2017*, pp. 608–612, may 2018.
- [13] L. S. Yoo, J. H. Lee, S. H. Ko, S. K. Jung, S. H. Lee, and Y. K. Lee, "A Drone Fitted with a Magnetometer Detects Landmines," *IEEE Geoscience and Remote Sensing Letters*, vol. 17, pp. 2035–2039, dec 2020.
- [14] Z. Kovács and I. Ember, "Landmine Detection with Drones," *Land Forces Academy Review*, vol. 27, no. 1, pp. 84–92, 2022.
- [15] P. Santana, J. Barata, H. Cruz, A. Mestre, J. Lisboa, and L. Flores, "A multi-robot system for landmine detection," *IEEE International Conference on Emerging Technologies and Factory Automation, ETFA*, vol. 1 2 VOLS, pp. 721–728, 2005.
- [16] C. Yan, X. Xiang, and C. Wang, "Towards real-time path planning through deep reinforcement learning for a uav in dynamic environments," *Journal of Intelligent & Robotic Systems*, vol. 98, pp. 297–309, 2020.
- [17] S. Aggarwal and N. Kumar, "Path planning techniques for unmanned aerial vehicles: A review, solutions, and challenges," *Computer Communications*, vol. 149, pp. 270–299, 2020.
- [18] A. C. Woods and H. M. La, "A novel potential field controller for use on aerial robots," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 49, no. 4, pp. 665–676, 2019.
- [19] K. Ebadi, L. Bernreiter, H. Biggie, G. Catt, Y. Chang, A. Chatterjee, C. E. Denniston, S.-P. Deschênes, K. Harlow, S. Khattak, L. Nogueira, M. Palieri, P. Petráček, M. Petřík, A. Reinke, V. Krátký, S. Zhao, A.-a. Agha-mohammadi, K. Alexis, C. Heckman, K. Khosoussi, N. Kottege, B. Morrell, M. Hutter, F. Pauling, F. Pomerleau, M. Saska, S. Scherer, R. Siegwart, J. L. Williams, and L. Carlone, "Present and future of slam in extreme environments: The darpa sub challenge," *IEEE Transactions on Robotics*, vol. 40, pp. 936–959, 2024.
- [20] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Transactions on Robotics*, vol. 32, no. 6, pp. 1309–1332, 2016.
- [21] M. Radmanesh, M. Kumar, P. H. Guentert, and M. Sarim, "Overview of path-planning and obstacle avoidance algorithms for uavs: A comparative study," *Unmanned Systems*, vol. 06, no. 02, pp. 95–118, 2018.