

# Resilient Terrain Navigation with a 5 DOF Metal Detector Drone

Patrick Pfreundschuh, Rik Bähnemann, Tim Kazik, Thomas Mantel, Roland Siegwart, and Olov Andersson

**Abstract**—Micro aerial vehicles (MAVs) hold the potential for performing autonomous and contactless land surveys for the detection of landmines and explosive remnants of war (ERW). Metal detectors are the standard detection tool but must be operated close to and parallel to the terrain. A successful combination of MAVs with metal detectors has not been presented yet, as it requires advanced flight capabilities. To this end, we present an autonomous system to survey challenging undulated terrain using a metal detector mounted on a 5 degrees of freedom (DOF) MAV. Based on an online estimate of the terrain, our receding-horizon planner efficiently covers the area, aligning the detector to the surface while considering the kinematic and visibility constraints of the platform. As the survey requires resilient and accurate localization in diverse terrain, we also propose a factor graph-based online fusion of GNSS, IMU, and LiDAR measurements. We validate the robustness of the solution to individual sensor degeneracy by flying under the canopy of trees and over featureless fields. A simulated ablation study shows that the proposed planner reduces coverage duration and improves trajectory smoothness. Real-world flight experiments showcase autonomous mapping of buried metallic objects in undulated and obstructed terrain.

## I. INTRODUCTION

Worldwide, landmines and ERW caused more than 7000 casualties in 2020. Member states of the Mine Ban Treaty reported that approximately 12 000 km<sup>2</sup> are contaminated, making vast areas uninhabitable [1]. Reducing suspected hazardous areas is challenging, as those areas are inaccessible. Additionally, detecting and mapping potential contamination requires expert training. As land clearance is expensive, it should be limited to minimal confirmed hazardous areas [2].

Recently, MAVs with various sensors have been investigated to accelerate land release. Drones can access any terrain without putting a person at risk and automatically collect geo-referenced data. In contrast to ground-based platforms, MAVs can inspect a surface contactless, relaxing traversability limitations [3] and eliminating the risk of accidental detonations. The essential sensor for humanitarian demining is the metal detector.

However, implementing a metal detector on a drone has failed so far because of the advanced flight capabilities necessary to perform practical surveys. Detection performance highly depends on the sensor head alignment with the soil surface [4]. Thus, the drone needs to maintain distance while aligning the sensor parallel with the surface. This requires a customized mechanical platform concept

This work was supported by Armasuisse W+T, Urs Endress Foundation's Project FindMine, a Wallenberg Foundation and WASP Postdoctoral Scholarship, and Swiss National Science Foundation's NCCR DfAB P3.

Authors are with Autonomous Systems Lab, ETH Zurich, e-mail: {patripfr, brik, tkazik, mantelt, rsiegwart, nandersson}@ethz.ch

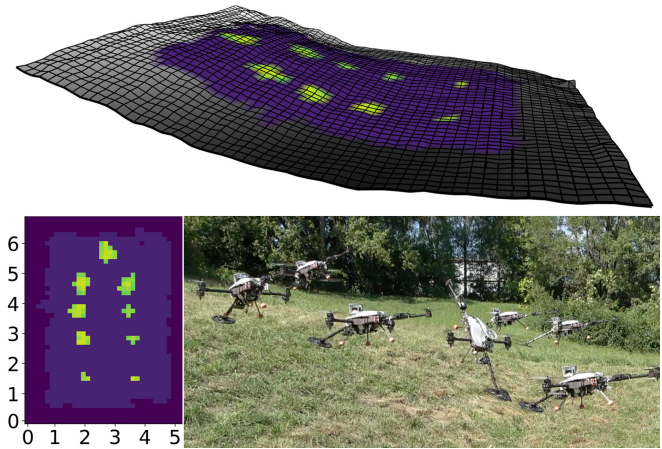


Fig. 1. *Top*: Metal detector signals visualized on a 2.5D terrain map. Colored cells were covered by the sensor. Brighter colors correspond to stronger signal responses from buried objects. *Bottom left*: Top-down 2D detection map. *Bottom right*: The system in action on undulated terrain.

and online perception to navigate resiliently in unknown, obstructed, undulated, outdoor terrain.

This work tackles this challenging surface tracking problem with a novel flying metal detector design. The system consists of a 5 DOF drone with a rigidly attached sensor. The drone can maneuver in all three translational directions and independently change its heading and pitch to align the sensor with the surface (Fig. 1). Simply aligning the sensor according to the surface normal is not ideal, as small changes in the normal can cause large motions of the platform as illustrated in Fig. 6. As this is inefficient, we propose a new local trajectory planner that trades off optimal alignment with minimal turning efforts and collision avoidance.

During the survey, the system maps the terrain with LiDAR. Precise surface tracking and complete coverage require accurate, smooth localization to keep the sensor close to the ground without touching it. The margins for inaccuracies are small, as they could result in a crash or invalid survey results, but ERW are often close to natural or human-made structures that cause GNSS outages or in feature-less areas which cause LiDAR inertial odometry (LIO) drift. We thus fuse these complementary modalities for resilient localization.

The main contributions of our work are:

- An autonomous system for a flying metal detector based on a 5 DOF MAV.
- A receding-horizon trajectory planner that optimizes terrain alignment in real-time while considering drone maneuverability and online map updates.
- Online factor graph GNSS, IMU, and LiDAR localization that is resilient to individual sensor degeneracy.

We show in simulation that our receding-horizon planning reduces coverage duration while maintaining sensor alignment. Real-world experiments show autonomous surveying of undulated and obstructed terrain and consistent localization even under the loss of GNSS or drift in LIO.

## II. RELATED WORK

### A. Flying Robotic ERW Detection

Several works towards ERW detection using aerial platforms exist. Some systems use ground penetrating radars [5]–[7] or cameras [8]. While these sensors can survey large areas from relatively far distances, standardization, and target detection are still being investigated. Yoon et al. [9] use a magnetometer but require special data processing to detect small objects. The most accepted sensor for deminers is the metal detector. Due to the limited sensing range of metal detectors, the sensor head must remain close and parallel to the surface for a strong sensor response [4]. Underactuated MAVs have coupled translational and rotational dynamics, which makes surface alignment challenging. The Mine Kafon Drone [10] proposed to solve this by mounting the metal detector on a gimbal but has not reported results yet. Our 5 DOF MAV with a rigidly mounted metal detector avoids the necessity of a gimbal as the vehicle pose directly controls the sensor orientation. Furthermore, our system is the first to perceive the environment online and avoid obstacles.

### B. Resilient Online State Estimation with GNSS and LIO

Both loosely- and tightly-coupled sensor fusion leverage the complementary nature of GNSS and LIO. The authors in [11] present a loosely-coupled approach that deploys a dual-factor graph to fuse the different modalities. In contrast, [12] uses an extended Kalman filter (EKF) to perform a tightly-coupled fusion of LiDAR, inertial, and GNSS measurements. LIO-SAM [13] uses a factor graph [14] to tightly fuse IMU and GNSS with scan-matching LIO. These approaches assume that the transformation between the local odometry and the GNSS frame is directly measured using a second GNSS receiver [11] or estimated in advance [12], [13] with an external pipeline that requires an initialization motion. Inspired by visual-inertial GNSS fusion [15]–[17], we formulate this transformation as an extrinsic online calibration. In contrast to the LIO-GNSS fusion frameworks above, we fuse the GNSS measurements into the odometry frame without requiring prior knowledge about the heading offset between the reference frames. Additionally, we define it as a transformation that varies over time to compensate for drift accumulating from LIO over periods of GNSS dropouts.

### C. Trajectory Planning on Surfaces

Finding the optimal flight trajectory for surface coverage in unknown 3D environments under visibility constraints is a challenging planning problem. Since the exact solution is computationally intractable, this work assumes a given coverage path for a region of interest. We then attempt to find a 5 DOF flight trajectory that accurately tracks the terrain surface to ensure good alignment for ERW detection, while,

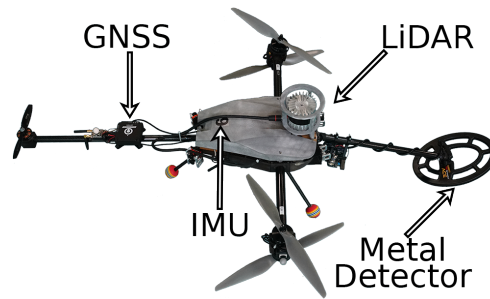


Fig. 2. The Voliro-T tricopter with a custom payload for autonomous navigation and metal detection.

in real-time, considering terrain visibility, speed, smoothness, and avoiding collisions. Earlier work addresses either planning on manifolds without considering coverage or coverage planning that either does not consider MAV kinematics or runs only offline. Manifold planning approaches include [18], which uses a model-predictive control approach, and [19], which proposes a reactive planner based on Riemannian motion policies. Visibility constraints are also included in [20] but do not consider coverage or run in real-time. On the other hand, [21] proposes a belief-space approach to optimize the coverage of free-flight MAV trajectories, but it is vulnerable to local minima and does not run onboard. Finally, [22] plans coverage patterns high above 3D terrain for remote sensing but does not consider MAV kinematics or appear suitable for real-time flight.

## III. PLATFORM SETUP

We built our system on top of a tricopter from Voliro AG. The platform, shown in Fig. 2, has fully actuated front propellers, allowing position control independent of the body pitch angle. The additional actuation also improves stable flight near surfaces where disturbances from propeller wash occur. We equipped the platform with an Ouster OS0-128 LiDAR with a 90° vertical field of view (FOV), allowing it to observe close terrain. The rear-facing FOV is partially obscured by the body. We also use an u-blox ZED-F9P multi-band RTK-GNSS receiver. A Bosch BMI085 IMU provides inertial measurements. Finally, a Garret ACE300i metal detector with 0.25 m major axis length is rigidly attached at the front. A custom microcontroller interface captures the digital output from the metal detector and sends it to the vehicle’s onboard computer. The software stack runs on a single-board PC with an AMD 4800U x86 CPU.

## IV. RESILIENT GNSS, IMU AND LiDAR FUSION

This section describes our GNSS and LIO fusion. Existing approaches [11]–[13] fuse the odometry into the GNSS frame, which requires prior knowledge of the yaw between the two reference frames. Our system directly operates in the LIO frame and estimates the alignment between the coordinate frames online. Thus our system is immediately deployable without prior initialization motion or sensing.

### A. Notation and Problem Definition

We define a fixed gravity-aligned odometry frame  $\mathbf{O}$  with origin at the initial pose of the IMU. The GNSS receiver

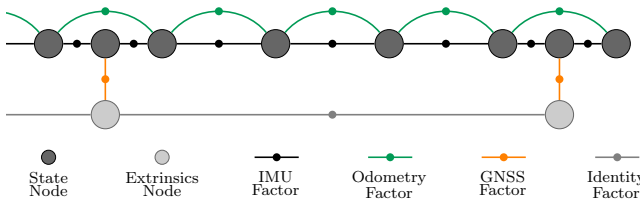


Fig. 3. The factor graph design. State nodes are connected through odometry and IMU factors, while extrinsic nodes are connected through identity factors. GNSS factors connect state and extrinsics nodes.

measures position with respect to an east-north-up gravity-aligned inertial frame  $\mathbf{I}$ , with origin at the position of the first measurement. The IMU frame serves as the robot body frame  $\mathbf{B}$ . The detector frame  $\mathbf{D}$  is placed in the center of the metal detector and aligned with the two axes of its elliptical coil. We define the robot's state as

$$\mathbf{x} = [\mathbf{R}_{OB}, \mathbf{O}\mathbf{p}_B, \mathbf{O}\mathbf{v}_B, \mathbf{B}\mathbf{b}^a, \mathbf{B}\mathbf{b}^g],$$

where  $\mathbf{R} \in SO(3)$  and  $\mathbf{p} \in \mathbb{R}^3$  represent the body orientation and position.  $\mathbf{v}_B \in \mathbb{R}^3$  is its linear velocity, while  $\mathbf{b}^a \in \mathbb{R}^3$  and  $\mathbf{b}^g \in \mathbb{R}^3$  indicate the accelerometer and gyro biases of the IMU. The transformation from the inertial to the odometry frame is  $\mathbf{T}_{OI} \in SE(3)$ . The relative position  ${}_{BP}\mathbf{p}$  between the GNSS receiver and the IMU is known.  $t^i, t^j, t^k$  are IMU, LiDAR, and GNSS measurements timestamps.

Given the set of LIO estimates, GNSS, and IMU measurements, the objective is to calculate the maximum a posteriori estimate for a set of states and extrinsics. We formulate this inference problem using the factor graph depicted in Fig. 3. We use a fixed window-length incremental smoothing and mapping (iSAM2) [23] solver to optimize the factor graph. In the following section, we present the individual factors.

### B. IMU Factor

Pre-integrated IMU factors summarize the IMU measurements between two nodes into a single factor [24]. We introduce nodes at GNSS and LiDAR timestamps, which arrive at a lower frequency. The IMU integration additionally serves as a high-frequency state estimate for the flight controller [25]. Initially, the vehicle is static to estimate an initial value for the IMU biases and gravity alignment.

### C. Odometry Factor

As a LIO pipeline, we use FAST-LIO2 [26], which operates on points directly. Raw points are favorable in environments with little geometric structure, as they can represent subtle features of the scene. FAST-LIO2 estimates a 6 DOF pose  $\mathbf{T}_{OB}^j$  at each  $t^j$ . We extract the relative transformation for the odometry factor as

$$\Delta\mathbf{T}^{j,j+1} = (\mathbf{T}_{OB}^j)^{-1}\mathbf{T}_{OB}^{j+1}.$$

We base the covariance estimation on the constraints of the registration problem, similar to [27], instead of assuming a fixed covariance as in [11], [13]. It computes as

$$\Sigma_{pos} = \mathbf{H}_{pos} = (\mathbf{J}_{pos}^T \mathbf{J}_{pos})^{-1}, \Sigma_{rot} = \mathbf{H}_{rot} = (\mathbf{J}_{rot}^T \mathbf{J}_{rot})^{-1},$$

where  $\mathbf{J}_{pos}$  and  $\mathbf{J}_{rot}$  correspond to the translational and rotational components of the measurement Jacobian in [26]. A tuning constant scales the covariance magnitude.

### D. Position Factor

The GNSS receiver measures the position with respect to the inertial frame  $\mathbf{I}$ . The measurement influences the odometry via the extrinsics  $\mathbf{T}_{OI}$ . The extrinsics' translation, roll, and pitch are directly observable, while yaw is only observable after the platform has moved sufficiently. We treat this transformation as a variable estimated in the factor graph. Position factors update the state and the extrinsics and use the covariance provided by the GNSS receiver. Their measurement model is

$$\mathbf{I}\mathbf{p} = (\mathbf{T}_{OI})^{-1} \cdot \mathbf{T}_{OB} \cdot \mathbf{B}\mathbf{p}.$$

We initialize the extrinsics with a high covariance on the yaw. Until the covariance of the yaw of the extrinsics is below a threshold, we artificially increase the GNSS measurement covariance. Hence, the GNSS measurements have negligible impact on the state during heading estimation. Once the covariance of the extrinsic yaw estimate decreases, the GNSS measurements increase their impact on the state.

## V. TRAJECTORY PLANNER FOR SURFACE TRACKING

The metal detector head has to be parallel to the ground surface to ensure reliable detection of buried objects. A fundamental part of the planning algorithm is an accurate local terrain representation. We use the 2.5D grid map framework [28] to fuse the LiDAR measurements into a map. The map also provides surface normals for alignment and a signed distance field (SDF) for collision avoidance. Second, we assume a 2D boustrophedon coverage path [29] given to serve as a global reference. The goal of the trajectory planner is to align the metal detector to the surface while safely tracking the reference.

Theoretically, our actuated 5 DOF drone can align with any surface by a combination of yaw  $\psi(t)$  and pitch  $\theta(t)$ . However, alternating surface normals result in many inefficient heading changes over uneven terrain, as seen in Fig. 6. Furthermore, uncontrolled yaw can cause the LiDAR to miss the terrain on the path ahead. Therefore, we need to plan a trajectory over position and orientation,  $(\mathbf{p}(t), \mathbf{R}(t))$ , that minimizes control effort (i.e. turning) to ensure smooth and efficient trajectories that satisfy ground alignment constraints, the kinematic limits for the MAV, and LiDAR visibility constraints to prevent flying into unobserved space.

To transform this complex optimization problem into a real-time solution, we first show how to simplify it for a 5 DOF MAV without obstacles before adding collision avoidance. Since the translational dynamics of the vectored thrust MAV are fast [30] compared to the maximum speed of 1 m/s allowed for reliable detection with a metal detector [31], we can assume that the coverage path can be tracked in position and just optimize rotations along the path.

At each discrete time step  $t^n$ , we solve a receding-horizon trajectory planning problem over the next  $h$  meters of the coverage path, which can be formalized as a trajectory

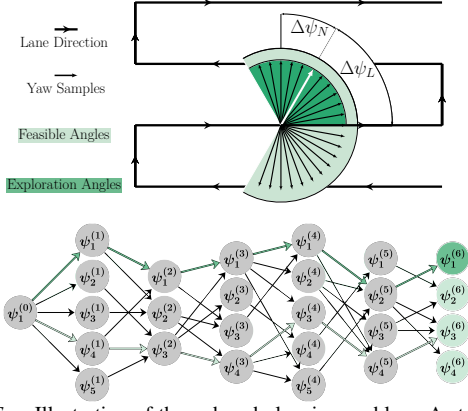


Fig. 4. *Top*: Illustration of the reduced planning problem. A state lattice is constructed from yaw samples along the path, which have to satisfy sensor visibility constraints (light green) of  $\pm 120^\circ$  around the direction of flight. Dark green samples are preferred as they ensure visibility of the next lane. *Bottom*: Illustration of the multitree. The shortest path is indicated in light green, but as solutions in the preferred sector exist, our algorithm selects the shortest path that ends in a preferred state, as shown in dark green.

optimization problem starting at  $t^n$ ,

$$\begin{aligned} \min_{\mathbf{R}(\cdot)} \quad & J = \int_{t^n}^{t^h} c(\dot{\mathbf{R}}(t)) dt & (1) \\ \text{subj. to} \quad & \dot{\mathbf{R}}(t) < \dot{\mathbf{R}}_{\max}, & \text{(kinematic limits)} \\ & \alpha(t) < \alpha_{\max}, & \text{(surface alignment)} \\ & |\Delta\psi_L(t)| < \Delta\psi_{L_{\max}}, & \text{(forward visibility)} \\ & \mathbf{R}(t^n) = \hat{\mathbf{R}}_{t^n}, & \text{(current orientation)} \\ & \mathbf{R}(t) \in SO(3), & \forall t \in [t^n, t^h], \end{aligned}$$

where  $t^h$  is the time at  $h$  meter into the path and  $\Delta\psi_L(t)$  is the angle between the planned heading and reference path direction as shown in Fig. 4. Finally,  $\alpha(t)$  is the ground alignment error defined as the shortest angle between the surface normal and the  $z$  axis of the detector frame shown in Fig. 5a.

Since the platform does not roll, the alignment error for a given yaw  $\psi$  is only a function of pitch  $\theta$ ,

$$\alpha(\theta) = \cos^{-1}(\mathbf{o}\mathbf{n} \cdot (\mathbf{R}_{OD}(\theta) \cdot \mathbf{D}\mathbf{z}_D)). \quad (2)$$

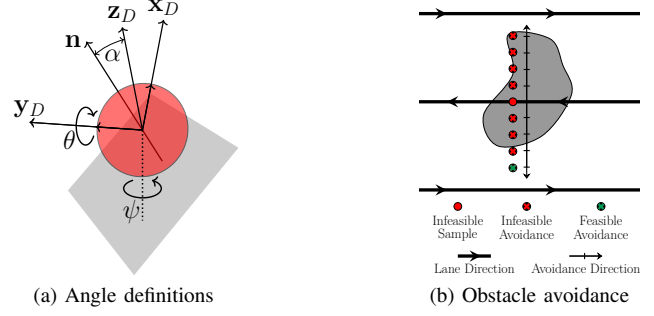
Following the proof in Appendix I, we can determine the optimal pitch  $\theta^*$  that minimizes  $\alpha$  for any candidate yaw,

$$\theta^*(\psi) = \tan^{-1} \left( \frac{n_x \cos(\psi) + n_y \sin(\psi)}{n_z} \right). \quad (3)$$

Fixing roll and minimizing the alignment error with an optimal pitch reduces the trajectory optimization problem in (1) to just solving for the optimal yaw trajectory  $\psi(\cdot)^*$  instead of the full rotation. Finally, aiming for a smooth trajectory in terms of yaw, we define the cost function to minimize yaw changes  $c(\dot{\mathbf{R}}(t)) = \dot{\psi}(t)$ .

#### A. Approximate Solution via Graph Search

At each discrete planner iteration  $t^n$ , we find an approximate solution to the planning problem in (1) by incrementally constructing a state lattice [32], and finding the minimum cost trajectory via graph search. The lattice is constructed



(a) Angle definitions  
(b) Obstacle avoidance

Fig. 5. *Left*: Illustration of the axes and angles to compute the optimal pitch angle. The grey square represents a surface cell. The red circle is the metal detector. *Right*: A grey obstacle invalidates the initial sample. We incrementally generate avoidance samples in the two perpendicular directions to the lane direction until a feasible trajectory is found.

by discretizing the 2D coverage path into reference positions. Given one discrete sample  $(x_{ref}^l, y_{ref}^l)$ , we calculate the 3D position using the current terrain map,

$$\mathbf{o}\mathbf{p}_D^l = [x_{ref}^l, y_{ref}^l, z_{ref}^l]^T + d \cdot \mathbf{o}\mathbf{n}. \quad (4)$$

Here  $d$  is the desired distance to the surface,  $\mathbf{o}\mathbf{n} \in \mathbb{R}^3$  is the surface normal, and  $z_{ref}$  is the surface elevation. We then sample yaw angles, as illustrated in Fig. 4, with a resolution of  $3^\circ$  to create a candidate tuple  $(\mathbf{o}\mathbf{p}_D^l, \psi_j^l)$ .

We attempt to connect each tuple to all feasible tuples in the previous step  $t^{l-1}$ , checking that the motion satisfies the constraints in (1), as well as checking collisions and traversability. The feasible yaw states and collision-free connections define a multitree  $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ , as shown in Fig. 4. Nodes  $\mathcal{V}_l$  at layer  $l$  represent the lattice states at the planning step  $t^l$ , and connect to  $\mathcal{V}_{l-1}$  through edges  $E(v) = \{e(v, v')\}_{v' \in N(v)}$ . The cost along each edge is the change in yaw  $c(v, v') = |\psi_{v'} - \psi_v|$ . The minimum cost path through the lattice is efficiently found via a breadth-first graph search. At each planning iteration, we only sample nodes for the new layer and reuse the previous tree.

#### B. Visibility Constraint

Our planner only allows paths through previously observed space to ensure a collision-free mission. To incentivize mapping of the area ahead, we modify the graph search to prioritize yaw candidates that enable the LiDAR to observe the *next* coverage path lane. To do so, we first solve the optimization problem (1) under the additional constraint

$$|\Delta\psi_N(t^h)| < e_{\max} \quad \text{(exploration constraint)}$$

where  $\psi_N$  is the direction that points towards the next lane, as seen in Fig. 4. This constraint is relaxed if no feasible solution with an exploration state was found, which can be efficiently implemented in the graph search approach by just checking the last layer in the search tree.

#### C. Collision Avoidance

Obstacles such as large rocks, trees, or poles may block parts of the coverage path. We identify traversable areas using a threshold on the slope of the terrain. Our algorithm pushes the path position out of the obstacle, as shown in



Fig. 6. Irregular, mostly flat simulated terrain. *Top*: The naive *Aligned-1* approach results in inefficient turning. *Bottom*: With the proposed planner, the system translates along the lane with only slight changes to its yaw.

Fig. 5b. For all infeasible reference positions  $\mathbf{p}^j$  in  $[t^n, t^h]$ , we iteratively sample avoidance positions orthogonal from the coverage path until the position is feasible. Solely shifting nodes keeps the graph structure intact. Using the graph search approximation above, we then plan a trajectory over this collision-free (sub-optimal) path.

## VI. EXPERIMENTAL RESULTS

In the following section, we first analyze the proposed planning approach quantitatively in a simulation environment. We then conduct real-world flight experiments to evaluate the sensor fusion. Finally, we validate the complete system by autonomously detecting metallic objects in the ground. We provide additional visualizations online<sup>1</sup>.

### A. Trajectory Planner for Surface Tracking

To evaluate our planner, we simulate the drone using a Gazebo model with the same dimensions, LiDAR sensor, and controller as the real platform. We deploy the system in a scaled-down model of a crater landscape [33] that exhibits different types of irregular terrain that we are interested in for ERW surveys. The  $25 \times 10\text{m}^2$  survey area starts in a primarily flat but irregular terrain (see Fig. 6), followed by a curved slope, and ends in an irregular, undulated area.

In addition to the proposed planner, we simulate three baseline approaches as an ablation study. The first baseline approach (*Aligned-1*) is a greedy surface tracker that only samples the two orientations that align perfectly with the surface at the next step, i.e., at  $h = 0.3\text{m}$ . The second baseline (*Aligned-6*) evaluates the optimal orientation for the next 6 steps. The same planning horizon as our proposed planner allows evaluating the performance gain from our planning strategy independent of its ability to look ahead. The third baseline (*Fixed Attitude*) is roughly equivalent to an underactuated quadcopter that ignores the slope and simply follows the terrain with a fixed height offset. We simulate this using the 5 DOF MAV but set  $\psi = \theta = 0$  and mount the metal detector in the middle of the platform. Note that this will result in improved results compared to a quadcopter, as the platform can translate without pitching. All approaches use ground truth odometry. We list the used parameters in Tab. I and present the results in Tab. II. To calculate alignment quality, we find the best-observed alignment  $\alpha_{\min}$  over all poses on the flight trajectory for each cell, and calculate the average over all cells, denoted as  $\overline{\alpha_{\min}}$ . Given the translational velocity constraint, the *Fixed*

<sup>1</sup><https://youtu.be/lv8ManXzV84>

TABLE I

| PARAMETERS USED FOR THE PRESENTED EXPERIMENTS |        |                         |             | Simulation      |                     | Real-World |                     |
|---|--------|-------------------------|-------------|-----------------|---------------------|------------|---------------------|
| Lane spacing                                  | 0.2 m  | $h$                     | 1.8 m       | $v_{\max}$      | 1 m/s               |            | 0.6 m/s             |
| Sample spacing                                | 0.3 m  | $\alpha_{\max}$         | $7.5^\circ$ | $\omega_{\max}$ | $60^\circ/\text{s}$ |            | $30^\circ/\text{s}$ |
| Grid cell size                                | 0.15 m | $\Delta\psi_{L_{\max}}$ | $120^\circ$ | $d$             | 0.15 m              |            | 0.18 m              |

TABLE II

| PLANNER PERFORMANCE IN SIMULATION. MEAN, STANDARD DEVIATION AND 95TH PERCENTILE ARE INDICATED BY $\bar{\cdot}$ , $\sigma$ AND $^{95\text{TH}}$ . |              |   |   |   |  |  |
|--|--------------|---|---|---|--|--|
| Method   | Duration [s] | $\overline{\Delta\psi_{t^n}} \pm \sigma$ [ $^\circ$ ] | $\overline{\Delta\theta_{t^n}} \pm \sigma$ [ $^\circ$ ] | $\overline{\alpha_{\min}}$ [ $^\circ$ ] | $\alpha_{\min}^{95\text{th}}$ [ $^\circ$ ] |  |
| Fixed Attitude   | 2132         | $0.0 \pm 0.0$   | $0.0 \pm 0.0$   | 11.38                                   | 24.29                                      |  |
| Aligned-1  | 2502         | $31.28 \pm 65.06$                                     | $2.70 \pm 2.55$   | 1.76                                    | 4.15                                       |  |
| Aligned-6  | 2469         | $28.07 \pm 58.15$                                     | $2.68 \pm 2.56$   | <b>1.74</b>                             | <b>4.14</b>                                |  |
| Proposed   | <b>2208</b>  | <b><math>4.17 \pm 17.17</math></b>                    | $2.94 \pm 3.64$   | 3.41                                    | 7.25                                       |  |

*Attitude* method is a lower limit on how fast the area can be covered. Our approach only takes 3.6% longer compared to this, while it covers the area in 10.6% less time than *Aligned-6* and has a much smoother trajectory in terms of yaw, as visible in Fig. 6. On average, the yaw change between subsequent target poses  $\Delta\psi_{t^n}$  is decreased by  $23.9^\circ$  compared to *Aligned-6*. As expected, our approach has a higher alignment deviation than the *Aligned* approaches. However, its 95th percentile stays below the intended limit of  $7.5^\circ$ , which is accurate enough to enable reliable inspection. It is worth noting that extending the planning horizon from *Aligned-1* to *Aligned-6* only marginally improves performance. Therefore, we conclude that the increased performance in our approach mostly stems from the proposed yaw sampling strategy. As the *Fixed Attitude* baseline does not track the orientation of the surface, the metal detector is not aligned in uneven areas, resulting in  $\alpha_{\min}^{95\text{th}} = 24.29^\circ$ , which is insufficient for reliable detection. Note that this result is a lower bound for a regular quadcopter, and would become even worse for more inclined areas. In contrast to our approach, all compared approaches traversed unobserved terrain multiple times.

### B. Resilient GNSS-IMU-LiDAR Localization

To evaluate the proposed sensor fusion approach's robustness to degeneracy, we deployed our system in an outdoor environment that contains both a GNSS-denied area and an area with severe LIO drift. LIO measurements arrive at 20 Hz. RTK-GNSS measurements arrive at 5 Hz. We use a window length of 5 s. A pilot steered the platform in

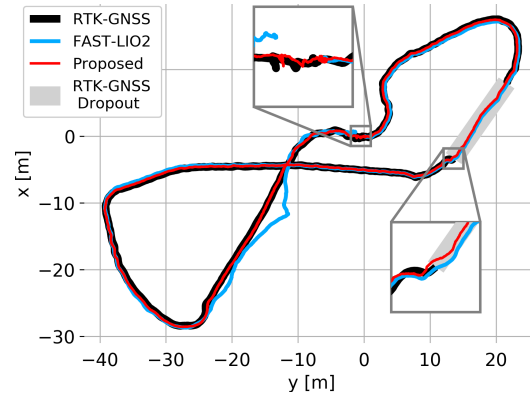


Fig. 7. Odometry plot. *Top detail*: The FAST-LIO2 trajectory is not closed due to accumulated drift. Start and end point of the proposed approach match. *Bottom detail*: Our approach estimates a consistent trajectory during RTK-GNSS loss and remains smooth upon recovery.

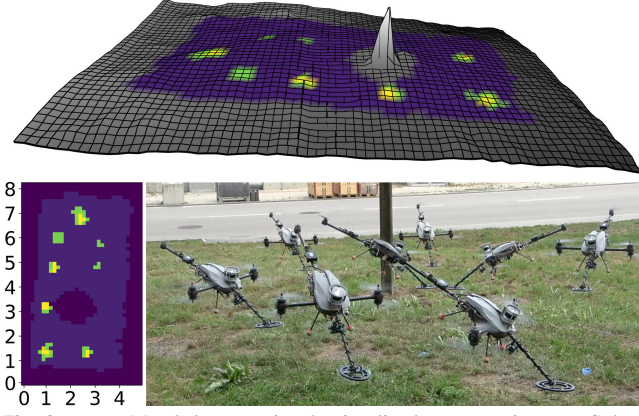


Fig. 8. *Top*: Metal detector signals visualized on a terrain map. Colored cells were covered by the sensor. *Bottom left*: 2D detection map. *Bottom right*: The system avoids the obstacle while maintaining detector alignment.

position control mode to trigger LIO and GNSS failure cases within one experiment. The controller uses the state estimation presented in Sec. IV. After takeoff, the extrinsics  $\mathbf{T}_{OI}$  converge after approximately 3.5 m. The drone flies through a dense tree canopy which causes RTK-GNSS loss, and later traverses flat grassland, which causes FAST-LIO2 to drift. The platform was started and landed in the same position with an accuracy of approximately 0.15 m.

As visualized in Fig. 7, the proposed approach retains a locally consistent state estimate while traversing through the RTK-GNSS-denied area that remains smooth upon signal recovery. FAST-LIO2 deviates 4.2 m from the RTK-GNSS trajectory where the point cloud registration is underconstrained. Since we calculate the LIO covariance based on the geometric structure, our approach implicitly handles the degeneracy and follows the RTK-GNSS trajectory. The distance between takeoff and landing position of our approach is 0.18 m. This is within our landing accuracy, while the final distance estimated by FAST-LIO2 is 1.52 m.

### C. Autonomous Search for Buried Metallic Objects

We evaluate the complete system in two real-world experiments. As metallic targets, we used gas cartridges with a diameter of 0.1 m. They were buried in the ground to be leveled with the surface to facilitate inspection of the results. To integrate the detector signal, we identify measured cells in the grid map by ray tracing from the sensor's bounding points and center to the map.

In the first experiment, we deploy the system in an undulated terrain covered by grass with a main slope of around  $20^\circ$  and steeper local inclinations along multiple axes. We placed 9 metallic targets. The platform covered an area of  $5.5 \times 4\text{m}^2$  in 317 s. As illustrated in Fig. 1, the planner generates trajectories that enable the sensor head to track the surface orientation. Additionally, despite the terrain steepness, the drone avoids collisions with any part of the body. We visualize the average sensor response per cell in Fig. 1. The 9 metallic objects are clearly visible. The lateral distance between the objects matches the actual distance of 1.65 m by visual inspection of the 2D detection map.

Due to the sensor's working principle, the patches with high-intensity are larger than the actual objects. One mea-

surement represents the whole area under the sensor head. Thus, even a tiny object can cause measurements of the size of the metal detector. The sensor used in this work does not expose raw signals and performs a calibration that the manufacturer does not specify. Thus, we observed that the sensitivity increased during the experiments. For future use, we propose to use a sensor that reports raw signals to avoid inconsistent sensitivity levels.

We performed a second experiment in an environment with a main slope of around  $8^\circ$ , which falls off slightly to both sides. A metal pole is present in the survey area. We placed 9 metallic targets. The platform covered an area of  $7 \times 3.5\text{m}^2$  in 373 s. The planner identifies the obstacle and maneuvers around it while maintaining the sensor aligned to the surface, as shown in Fig. 8. 8 metallic targets are visible in the generated map. The missed target was in an area close to the pole, which the flying detector avoided.

## VII. CONCLUSION

We presented a complete system for metal detection with a 5 DOF MAV capable of navigating in GNSS-denied and LIO degenerate environments while adjusting its coverage trajectory based on terrain, visibility, and obstacles. The platform's maneuverability shows superior surface tracking over an underactuated MAV. The proposed receding-horizon planner shows good terrain perception and reduces coverage time over compared approaches. Real-world experiments showcase that the system can autonomously survey challenging undulated terrain for ERW detection. Visualizing the metal detector response on the terrain map combined with the GNSS reference allows easy re-localization in hazardous areas for demining. In future work, the obstacle avoidance strategy, although sufficient in our experiments, could benefit from integration in a coverage re-planning layer. We also believe our resilient surface navigation holds promise for other surface tracking applications such as spraying, close-up imaging, or non-destructive testing.

## APPENDIX

We search  $\theta$  that minimizes the alignment error  $\alpha$  given  $\psi$ :  $\theta_{cand}^* = \arg \min_{\theta} \alpha(\theta) = \arg \max_{\theta} f(\theta)$  (5)

$$f(\theta) = \mathbf{o}\mathbf{n} \cdot (\mathbf{R}_{OD}(\theta, \psi) \cdot \mathbf{D}\mathbf{z}\mathbf{D}) \stackrel{\text{Roll}^{\pm 0}}{=} \sin(\theta) \underbrace{(n_x \cos(\psi) + n_y \sin(\psi))}_a + \cos(\theta)n_z \quad (6)$$

We find the critical point as:

$$f(\theta_c)' = \cos(\theta_c)a - \sin(\theta_c)n_z \stackrel{!}{=} 0 \\ \Leftrightarrow \theta_c = \tan^{-1}\left(\frac{a}{n_z}\right) = \tan^{-1}\left(\frac{n_x \cos(\psi) + n_y \sin(\psi)}{n_z}\right) \quad (7)$$

$$f(\theta)'' = -\sin(\theta)a - \cos(\theta)n_z \stackrel{!}{<} 0 \\ \begin{matrix} \cos(\theta) > 0 \\ \theta \in (-\pi, \pi) \\ \Leftrightarrow \end{matrix} -\tan(\theta)a < n_z^2 \quad (8)$$

Inserting (7) into (8) yields:

$$-\tan(\tan^{-1}\left(\frac{a}{n_z}\right))a = -a^2/n_z < n_z \stackrel{n_z > 0}{\Leftrightarrow} -a^2 < n_z^2 \quad (9)$$

Since  $f(\theta_c)' = 0$  and  $f(\theta_c)'' < 0$  we find  $\theta_{cand}^* = \theta_c$

## REFERENCES

- [1] International Campaign to Ban Landmines, “Landmine Monitor 2021,” 2021.
- [2] United Nations Mine Action Service (UNMAS), “IMAS 07.11 - Land Release,” 2019.
- [3] H. Balta, G. De Cubber, D. Doroftei, Y. Baudoin, and H. Sahli, “Terrain traversability analysis for off-road robots using time-of-flight 3d sensing,” in *7th IARP International Workshop on Robotics for Risky Environment-Extreme Robotics, Saint-Petersburg, Russia*, 2013.
- [4] European Committee for Standardization, “Humanitarian Mine Action - Test and evaluation - Metal Detectors,” 2003.
- [5] R. Bähmann, N. Lawrance, L. Streichenberg, J. J. Chung, M. Pantic, A. Grathwohl, C. Waldschmidt, and R. Siegwart, “Under the Sand: Navigation and Localization of a Micro Aerial Vehicle for Landmine Detection with Ground-Penetrating Synthetic Aperture Radar,” *Field Robotics*, vol. 2, no. 1, pp. 1028–1067, Mar. 2022.
- [6] Y. A. Lopez, M. Garcia-Fernandez, G. Alvarez-Narciandi, and F. L.-H. Andres, “Unmanned aerial vehicle-based ground-penetrating radar systems: A review,” *IEEE Geoscience and Remote Sensing Magazine*, 2022.
- [7] D. Sipos, P. Planinsic, and D. Gleich, “On drone ground penetrating radar for landmine detection,” in *2017 First International Conference on Landmine: Detection, Clearance and Legislations (LDCL)*, 2017, pp. 1–4.
- [8] J. Baur, G. Steinberg, A. Nikulin, K. Chiu, and T. S. de Smet, “Applying deep learning to automate uav-based detection of scatterable landmines,” *Remote Sensing*, vol. 12, no. 5, p. 859, 2020.
- [9] L.-S. Yoo, J.-H. Lee, Y.-K. Lee, S.-K. Jung, and Y. Choi, “Application of a Drone Magnetometer System to Military Mine Detection in the Demilitarized Zone,” *Sensors*, vol. 21, no. 9, p. 3175, Jan. 2021, number: 9 Publisher: Multidisciplinary Digital Publishing Institute.
- [10] M. Hassani. Mine Kafon. (2022, August 30).
- [11] J. Nubert, S. Khattak, and M. Hutter, “Graph-based Multi-sensor Fusion for Consistent Localization of Autonomous Construction Robots,” in *2022 International Conference on Robotics and Automation (ICRA)*, May 2022, pp. 10048–10054.
- [12] X. Li, H. Wang, S. Li, S. Feng, X. Wang, and J. Liao, “GIL: a tightly coupled GNSS PPP/INS/LiDAR method for precise vehicle navigation,” *Satellite Navigation*, vol. 2, no. 1, p. 26, Nov. 2021.
- [13] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and D. Rus, “Lio-sam: Tightly-coupled lidar inertial odometry via smoothing and mapping,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 5135–5142.
- [14] F. Dellaert and M. Kaess, 2017.
- [15] R. Mascaro, L. Teixeira, T. Hinzmann, R. Siegwart, and M. Chli, “GOMSF: Graph-Optimization Based Multi-Sensor Fusion for robust UAV Pose estimation,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, May 2018, pp. 1421–1428, iSSN: 2577-087X.
- [16] S. Boche, X. Zuo, S. Schaefer, and S. Leutenegger, “Visual-Inertial SLAM with Tightly-Coupled Dropout-Tolerant GPS Fusion,” Aug. 2022, arXiv:2208.00709 [cs].
- [17] W. Lee, K. Eckenhoff, P. Geneva, and G. Huang, “Intermittent GPS-aided VIO: Online Initialization and Calibration,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, May 2020, pp. 5724–5731, iSSN: 2577-087X.
- [18] G. Lu, W. Xu, and F. Zhang, “Model predictive control for trajectory tracking on differentiable manifolds,” *arXiv preprint arXiv:2106.15233*, 2021.
- [19] M. Pantic, L. Ott, C. Cadena, R. Siegwart, and J. Nieto, “Mesh manifold based riemannian motion planning for omnidirectional micro aerial vehicles,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4790–4797, 2021.
- [20] M. Watterson, S. Liu, K. Sun, T. Smith, and V. Kumar, “Trajectory optimization on manifolds with applications to quadrotor systems,” *The International Journal of Robotics Research*, vol. 39, no. 2-3, pp. 303–320, 2020.
- [21] B. Davis, I. Karamouzas, and S. J. Guy, “C-opt: Coverage-aware trajectory optimization under uncertainty,” *IEEE Robotics and Automation Letters*, vol. 1, no. 2, pp. 1020–1027, 2016.
- [22] Y. Choi, Y. Choi, S. Briceno, and D. N. Mavris, “Three-dimensional uas trajectory optimization for remote sensing in an irregular terrain environment,” in *2018 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2018, pp. 1101–1108.
- [23] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard, and F. Dellaert, “isam2: Incremental smoothing and mapping with fluid relinearization and incremental variable reordering,” in *2011 IEEE International Conference on Robotics and Automation*, 2011, pp. 3281–3288.
- [24] C. Forster, L. Carlone, F. Dellaert, and D. Scaramuzza, “On-manifold preintegration for real-time visual-inertial odometry,” *IEEE Transactions on Robotics*, vol. 33, no. 1, pp. 1–21, 2017.
- [25] V. Indelman, S. Williams, M. Kaess, and F. Dellaert, “Information fusion in navigation systems via factor graph based incremental smoothing,” *Robotics and Autonomous Systems*, vol. 61, no. 8, pp. 721–738, 2013.
- [26] W. Xu, Y. Cai, D. He, J. Lin, and F. Zhang, “Fast-lio2: Fast direct lidar-inertial odometry,” *IEEE Transactions on Robotics*, vol. 38, no. 4, pp. 2053–2073, 2022.
- [27] J. Zhang, M. Kaess, and S. Singh, “On degeneracy of optimization-based state estimation problems,” in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 809–816.
- [28] P. Fankhauser, M. Bloesch, and M. Hutter, “Probabilistic terrain mapping for mobile robots with uncertain localization,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 3, no. 4, pp. 3019–3026, 2018.
- [29] R. Bähmann, N. Lawrance, J. J. Chung, M. Pantic, R. Siegwart, and J. Nieto, “Revisiting boustrophedon coverage path planning as a generalized traveling salesman problem,” in *Field and Service Robotics: Results of the 12th International Conference*. Springer, 2021, pp. 277–290.
- [30] R. Watson, M. Kamel, D. Zhang, G. Dobie, C. MacLeod, S. G. Pierce, and J. Nieto, “Dry coupled ultrasonic non-destructive evaluation using an over-actuated unmanned aerial vehicle,” *IEEE Transactions on Automation Science and Engineering*, vol. 19, no. 4, pp. 2874–2889, 2021.
- [31] Garrett Electronics Inc. (2016) Ace 300i owner’s manual.
- [32] M. Pivtoraiko, R. A. Knepper, and A. Kelly, “Differentially constrained mobile robot motion planning in state lattices,” *Journal of Field Robotics*, vol. 26, no. 3, pp. 308–333, 2009.
- [33] R. Jaumann, D. A. Williams, D. L. Buczkowski, R. A. Yingst, F. Preusker, H. Hiesinger, N. Schmedemann, T. Kneissl, J. B. Vincent, D. T. Blewett, B. J. Buratti, U. Carsenty, B. W. Denevi, M. C. D. Sanctis, W. B. Garry, H. U. Keller, E. Kersten, K. Krohn, J.-Y. Li, S. Marchi, K. D. Matz, T. B. McCord, H. Y. McSween, S. C. Mest, D. W. Mittlefehldt, S. Mottola, A. Nathues, G. Neukum, D. P. O’Brien, C. M. Pieters, T. H. Prettyman, C. A. Raymond, T. Roatsch, C. T. Russell, P. Schenk, B. E. Schmidt, F. Scholten, K. Stephan, M. V. Sykes, P. Tricarico, R. Wagner, M. T. Zuber, and H. Sierks, “Vesta’s shape and morphology,” *Science*, vol. 336, no. 6082, pp. 687–690, 2012.