

MineInsight: A Multi-sensor Dataset for Humanitarian Demining Robotics in Off-Road Environments

Mario Malizia , Charles Hamesse , Ken Hasselmann ,
Geert De Cubber , Nikolaos Tsiogkas , Eric Demeester , Rob Haelterman 

Abstract—The use of robotics in humanitarian demining increasingly involves computer vision techniques to improve landmine detection capabilities. However, in the absence of diverse and realistic datasets, the reliable validation of algorithms remains a challenge for the research community. In this paper, we introduce MineInsight, a publicly available multi-sensor, multi-spectral dataset designed for off-road landmine detection. The dataset features 35 different targets (15 landmines and 20 commonly found objects) distributed along three distinct tracks, providing a diverse and realistic testing environment. MineInsight is, to the best of our knowledge, the first dataset to integrate dual-view sensor scans from both an Unmanned Ground Vehicle and its robotic arm, offering multiple viewpoints to mitigate occlusions and improve spatial awareness. It features two LiDARs, as well as images captured at diverse spectral ranges, including visible (RGB, monochrome), visible short-wave infrared (VIS-SWIR), and long-wave infrared (LWIR). Additionally, the dataset provides bounding boxes generated by an automated pipeline and refined with human supervision. We recorded approximately one hour of data in both daylight and nighttime conditions, resulting in around 38,000 RGB frames, 53,000 VIS-SWIR frames, and 108,000 LWIR frames. MineInsight serves as a benchmark for developing and evaluating landmine detection algorithms. Our dataset is available at <https://github.com/mariomlz99/MineInsight>.

Index Terms—Data sets for robotic vision, field robotics, computer vision for automation.

I. INTRODUCTION

LANDMINES remain one of the most dangerous legacies of conflicts, posing long-term threats to civilians and obstructing post-war recovery efforts. In many affected regions, the presence of antipersonnel (AP) and antitank (AT) landmines severely limits agricultural activities, infrastructure

development, and safe mobility. Traditional humanitarian demining methods, relying on human experience and metal detectors, are slow, hazardous, and costly. Consequently, robotic systems have been increasingly deployed [1] to accelerate landmine detection and improve operational safety. While many robotic systems rely solely on camera feedback for operators, researchers have increasingly explored diverse detection techniques [2]–[7] exploiting images captured across multiple spectral ranges. Existing landmine detection datasets (Table I) often suffer from significant limitations, such as a narrow range of explosive ordnance types, limited sensor diversity, or inadequate coverage of challenging environmental conditions. These constraints on existing datasets hinder the development and validation of novel detection algorithms. Moreover, the lack of a standardized storage format for data releases makes it difficult to ensure consistency and comparability across studies, highlighting the need for datasets that follow standardized practices established in off-road mobile robotics datasets [8]–[10]. Most datasets rely on one sensor, neglecting sensor fusion applications incorporating other spectral ranges such as visible short-wave infrared (VIS-SWIR) or long-wave infrared (LWIR). Additionally, these datasets often omit LiDAR data, which delivers essential depth information for navigating uneven terrain, detecting obstacles, and improving landmine localization. They do not adequately address occlusion, a major challenge when mines are covered with vegetation, debris, or soil [11]. To the best of our knowledge, the majority of existing datasets in the literature are recorded using Unmanned Aerial Vehicles (UAVs). However, in many scenarios, deploying Unmanned Ground Vehicles (UGVs) is more practical or necessary, highlighting a gap that we aim to fill. To address these gaps, we introduce **MineInsight**, a multi-sensor, multi-spectral dataset designed to advance robotic landmine detection. MineInsight integrates dual sensor perspectives, capturing both 2D and 3D information from a UGV and its robotic arm, enabling a wider spatial understanding of landmines in complex terrains. Our dataset features two LiDARs and a diverse set of multi-spectral sensors spanning monochrome, RGB, VIS-SWIR, and LWIR spectra.

The contributions of this paper are as follows:

- We introduce MineInsight, a dataset featuring a diverse array of inert AP and AT landmines, along with common items that might be found in natural environments, to better mimic realistic demining scenarios and challenge detection systems against false positives.
- Our dataset is, to the best of our knowledge, the first dataset that integrates dual sensor scans captured from

Manuscript received: March 20, 2025; Revised July 31, 2025; Accepted November 8, 2025.

This paper was recommended for publication by Editor Hyungpil Moon upon evaluation of the Associate Editor and Reviewers’ comments.

This work was supported by the Belgian Defense under Grant DAP 23/08. (Corresponding author: Mario Malizia) - Email: mario.malizia@mil.be.

M. Malizia is with the Department of Mechanics, Royal Military Academy, Av. de la Renaissance 30, 1000 Brussels, Belgium, and also with the ACRO Research Group, Department of Mechanical Engineering, Wetenschapspark 27, 3590 Diepenbeek, Belgium, and Flanders Make @ KU Leuven, B-3001 Heverlee, Belgium.

C. Hamesse and R. Haelterman are with the Department of Mathematics, Royal Military Academy, Av. de la Renaissance 30, 1000 Brussels, Belgium.

K. Hasselmann and G. De Cubber are with the Department of Mechanics, Royal Military Academy, Av. de la Renaissance 30, 1000 Brussels, Belgium.

N. Tsiogkas is with the ACRO Research Group, Department of Computer Science, KU Leuven, Wetenschapspark 27, 3590 Diepenbeek, Belgium, and also with Flanders Make @ KU Leuven, B-3001 Heverlee, Belgium.

E. Demeester is with the ACRO Research Group, Department of Mechanical Engineering, KU Leuven, Wetenschapspark 27, 3590 Diepenbeek, Belgium, and also with Flanders Make @ KU Leuven, B-3001 Heverlee, Belgium.

Digital Object Identifier (DOI): see top of this page.

TABLE I
COMPARISON OF DIFFERENT LANDMINE DETECTION DATASETS. MISSING ENTRIES INDICATE UNAVAILABLE DATA.
Note: "AVision" is the short form for AlliedVision.

Dataset	Total Landmines	Platform	Environment	Camera	LiDAR	Labels
De Smet et al. [12]	1	UAV	Grass, rubble	Parrot Sequoia Multispectral	-	-
Steinberg et al. [13]	1	UAV	Grass, rubble	Parrot Sequoia Multispectral	-	2D bounding boxes
Baur et al. [14]	1	UAV	Sand	Parrot Sequoia Multispectral	-	-
Leloglu et al. [15]	1	Fixed Tripod	Clay	ATOM 1024 IR, FLIR T 650 SC	-	-
Tamayo et al. [16]	1	UAV	Low vegetation, clay	Zenmuse XT IR	-	-
Vivoli et al. [17]	2	Handheld Support	Grass, gravel	iPhone 13	-	2D bounding boxes
MineInsight (Ours)	15	UGV	Low to high vegetation, leaves	AVision Alvium 1800 U-240C, AVision Alvium 1800 U-130, Sevensense Core Research, Teledyne FLIR Boson 640	Livox AVIA, Livox Mid-360	2D bounding boxes

both the UGV and its robotic arm. The robotic arm’s dynamic motion provides varying sensor viewpoints that effectively mitigate obstructions and enhance the detection of hidden landmines.

- We provide 2D bounding boxes refined with human supervision, serving as a benchmark for validating detection algorithms and supporting further research.

II. RELATED WORK

Landmine detection datasets, as illustrated in Table I, are relatively rare in the existing literature. This section offers an overview of the available ones, detailing their content and the conditions under which they were collected.

Binghamton University has contributed significantly to the field of landmine detection by releasing multiple datasets as part of the Scatterable Landmine Detection Project. These datasets include data from various UAV flights flown over different terrains. The project’s primary focus is the detection of PFM-1 landmines, together with their carrying case. We propose a categorization of the dataset releases into three groups (1–7, 8–9, and 10–24), facilitating a direct comparison between UAV platforms, sensors, and data collection methods across studies. The first group (Datasets 1-7), released by De Smet et al. [12], includes data collected using a DJI Phantom 4 Pro UAV and a DJI Matrice 600 Pro UAV, both equipped with a Parrot Sequoia Multispectral Camera. Flights were conducted over grass and rubble environments, resulting in seven high-resolution orthophotos. Steinberg et al. released the second group (Datasets 8-9) [13], which includes data collected from three flights over grass and four flights over rubble using a DJI Matrice 600 Pro UAV equipped with a Parrot Sequoia Multispectral Camera. Baur et al. later used the same data to explore the potential of deep learning-based landmine detection, with their experimental setup described in [3], where they address the challenges of detecting PFM-1 landmines in UAV-based imagery. Baur et al. [14] released the third group (Datasets 10-24), which comprises photogrammetric data collected using a DJI Phantom 4 UAV equipped with a Parrot Sequoia Multispectral Camera and processed with

Pix4DMapper. To introduce variability, the PFM-1 landmines were rotated, partially buried, or partially hidden in each orthophoto.

Kaya and Leloglu [18] conducted research on the detection of various landmine types, including M15, M16, M2, M48 metal, DM-11, and M14. These landmines, filled with wax or metal to simulate real thermal properties, were placed at varying depths (surface, semi-buried, and buried) to replicate real-world conditions. Experiments were conducted in two test areas with different soil types (clay and quasi-humid clay loam), cleared of vegetation to reduce thermal noise. As part of their study, the authors released a dataset [15] comprising 96 thermal images captured at 15-minute intervals over 24 hours. Data were collected with the ATOM 1024 infrared (IR) and FLIR T 650 SC cameras, which have spectral ranges of 8000–14000 nm and 7500–13000 nm, respectively. The cameras were positioned at fixed altitudes on a roof, pointing toward the test areas. However, the resulting dataset only includes DM-11 landmines.

Tamayo et al. [16] released a dataset for landmine detection containing thermographic and visible spectrum images of buried landmines at varying depths. The dataset includes 2,700 thermographic images, specifically designed to detect “leg-breaker” antipersonnel mines. These landmines were custom replicas, made from a PVC cylinder, a 5 ml syringe as a detonator, 20 g of nails, and 400 g of anthracite charcoal, chosen for its thermal properties similar to TNT. The experiments were conducted on a terrain with a low percentage of vegetation and clay soil. The mines were buried at various depths, and distributed across distinct zones to facilitate segmentation in image analysis. The dataset was collected using a Zenmuse XT IR camera with a spectral range of 7500–13000 nm, mounted on a DJI Matrice 100 UAV. Environmental parameters such as temperature, irradiance, wind speed, and relative humidity were recorded alongside the images to provide context for thermal variations in the dataset.

Vivoli et al. [17] released SULAND, a dataset for surface landmine detection. It consists of 33,541 RGB images captured using an iPhone 13 camera, with 10,169 images annotated

TABLE II
SPECIFICATIONS OF THE SENSORS USED FOR OUR DATA COLLECTION.

Notes: “AVision” is the short form for AlliedVision; “FoV” refers to the camera’s field of view; “Ch” denotes the number of channels.

Sensor	Type	Location	Resolution	FoV [°]		Ch	Spectral Range [nm]	Rate [Hz]
				H	V			
Livox Mid-360	LiDAR	Mobile Base	-	360.0	59.0	-	905*	10
Sevensense Core Research†	Monochrome Camera	Mobile Base	720 × 540	126.0	92.4	1	380 - 750	15
AVision Alvium 1800 U-240C	RGB Camera	Robotic Arm	1916 × 1216	103.6	76.7	3	300 - 1100	15
AVision Alvium 1800 U-130	VIS-SWIR Camera	Robotic Arm	1296 × 1032	22.7	18.2	1	400 - 1700	15
Livox AVIA	LiDAR	Robotic Arm	-	70.4	77.2	-	905*	10
Teledyne FLIR Boson 640	LWIR Camera	Robotic Arm	640 × 512	95.0	76.0	1	8000 - 13500	30

* The 905 nm value indicates the central wavelength of the LiDAR’s laser emitter, not a spectral range.

† For simplicity, this sensor is reported in one line, yet it includes five different cameras.

for object detection tasks. The sequences were collected under varying weather conditions, including cloudy, sunny, and shadowed settings, across grass and gravel environments. The targets selected for this study were the PFM-1 and PMA-2 landmines. Data were collected in diverse environmental settings, considering factors such as terrain type, weather conditions, and obstacles like bushes, rocks, and branches that could affect visibility.

With the increasing availability of services such as Roboflow [19], a wide variety of landmine detection datasets have become publicly accessible. These datasets include both real and simulated images captured using RGB or thermal cameras. However, they often lack essential details about the data collection process, such as camera specifications, environmental conditions, and experimental setup. This omission makes it challenging to replicate experiments or validate algorithms effectively.

Existing landmine detection datasets lack a standardized framework for data collection and release. We follow established practices from robotic datasets, adopting the Robot Operating System (ROS) 2 framework to maintain interoperability and reproducibility within the research community. No publicly available landmine detection dataset, to our knowledge, features a UGV-mounted robotic arm for data collection and features 3D sensors, such as LiDAR. Our dataset introduces these elements while also incorporating 15 landmine types, balanced between AP and AT. With the inclusion of items commonly found in vegetation and by capturing data in both daylight and nighttime conditions, we enable the analysis of the robustness and sensitivity of detection algorithms.

III. HARDWARE SETUP AND CALIBRATION

In this section, we describe the hardware configuration of our UGV, including the sensor suite installed on both the mobile base and the robotic arm. We then introduce the onboard computing architecture responsible for data processing and storage, outline the synchronization approach, and detail the calibration procedures for camera intrinsics and camera–LiDAR extrinsics.




Fig. 1. Illustration of the different sensors mounted on the UGV. Left - The UGV is ready for data collection. Top right - Sensors installed on the robotic arm. Bottom right: Sensors mounted directly on the mobile base.

A. Hardware setup

The UGV used for data collection, along with its sensor setup, is illustrated in Figure 1. We categorize the sensors based on their location into two groups: the mobile base sensor suite and the robotic arm sensor suite, described in Table II. In addition to the visible sensor suite, we mounted a Microstrain 3DM-GV7-AR IMU inside the UGV. The system relies on two onboard computers: a MYBOTSHOP Industrial PC, which we designate as our embedded computer, and an NVIDIA Jetson Orin (which we refer to as Orin). The embedded computer manages the control of the mobile base and robotic arm, as well as the processing of data streams from the base-mounted sensors. Meanwhile, the Orin is dedicated to data storage and processing for the sensors mounted on the robotic arm. We use a gigabit Ethernet switch to link the two computers and the robotic arm, and we ensure synchronization through hardware Precision Time Protocol (PTP). Both computers run Ubuntu Linux 22.04 LTS and ROS 2 Humble to ensure a common framework for data collection.



Fig. 2. Overview of all the objects included in the dataset. The objects on the left side of the black line are **landmines (L)**, while the objects on the right side of the black line are **common items (C)**. For what concerns the landmines (L) - **L1**: PFM-1 “Butterfly”, **L2**: PMN, **L3**: TC-3.6, **L4**: M35, **L5**: C-3 “Elsie”, **L6**: M6 (grey), **L7**: TMA-2, **L8**: MON-50, **L9**: MON-90, **L10**: M6 (blue), **L11**: TMM-1, **L12**: Type 72 (P), **L13**: TM-46, **L14**: PROM-1, **L15**: VS-50. For what concerns the common items (C) - **C1**: Soda metal can, **C2**: Disposable paper cup, **C3**: Sponge, **C4**: Plastic charger, **C5**: Metal pot, **C6**: Glass vinegar bottle, **C7**: Plastic water bottle (significantly crumpled), **C8**: Plastic shampoo bottle, **C9**: Plastic water bottle (partially crumpled), **C10**: Plastic water bottle (slightly crumpled), **C11**: Glass pepper dispenser, **C12**: Glass jar (grey cover), **C13**: Plastic chips bag (slightly crumpled), **C14**: Metal coke can, **C15**: Metal tuna can, **C16**: Glass beer bottle, **C17**: Glass jar (green cover), **C18**: Metal corn tin, **C19**: Plastic chips bag (significantly crumpled), **C20**: Plastic cup.

B. Synchronization

The Orin, the robotic arm, the two LiDARs, and the monochrome cameras are synchronized with the master clock of the embedded computer using hardware PTP. The RGB, VIS-SWIR, and LWIR cameras, all equipped with global shutters, are connected to the Orin via USB. Since the Jetson Orin is PTP hardware-synchronized with the embedded computer, the aforementioned USB-connected sensors remain software-synchronized relative to it.

C. Camera calibration

As listed in Table II, our setup comprises eight cameras, each requiring a distinct calibration approach based on its optical and operational properties. We calibrate the RGB and VIS-SWIR cameras using the Kalibr toolbox [20] [21]. For the LWIR camera, calibration is performed using the MRT Calibration Toolbox [22] and a custom-designed checkerboard pattern with metallic squares, inspired by the approach described in [23]. For the Sevensense Core Research module, we use the manufacturer’s parameters.

D. Camera and LiDAR calibration

Accurate extrinsic calibration between sensors ensures precise data fusion from multiple streams. We organize the calibration into two groups based on sensor placement. The first group includes the Livox Mid-360 LiDAR and each camera in the Sevensense Core Research module. The second group pairs the Livox AVIA LiDAR calibrated with the RGB, VIS-SWIR, or LWIR camera, respectively. We use the method described in [24] to perform targetless calibration, aligning each camera frame with the corresponding LiDAR scan.

IV. DATASET

In this section, we present the MineInsight dataset [25]. We describe the targets, detail the environmental conditions with varying vegetation and lighting, and outline the recorded data sequences, including the UGV and robotic arm configurations. Finally, we explain the reference data generation process, which combines an automated pipeline with human supervision to improve the quality of the annotations.

A. Targets

We placed a combination of AP and AT inert landmines and common items that might be found in natural environments across three tracks to construct a diverse dataset. The items were left overnight, allowing their temperatures to equilibrate with the ambient environment and better mimic realistic conditions. All the items included in the dataset are shown in Figure 2. We summarize the distribution of landmines and common items per track as follows:

Track 1 (25 targets): 15 landmines (11 AP, 4 AT) and 10 common items. **Track 2** (21 targets): 11 landmines (6 AP, 5 AT) and 10 common items. **Track 3** (21 targets): 11 AP landmines and 10 common items.

B. Environment

We selected three outdoor environments for data sampling, each consisting of an approximately 15 m long and 2.5 m wide track designated for inspection. The track boundaries were marked by ropes to indicate restricted access. The terrain was primarily covered with leaves, and the campaign took place under consistently cloudy conditions (no direct sunlight). We categorize the tracks as follows:

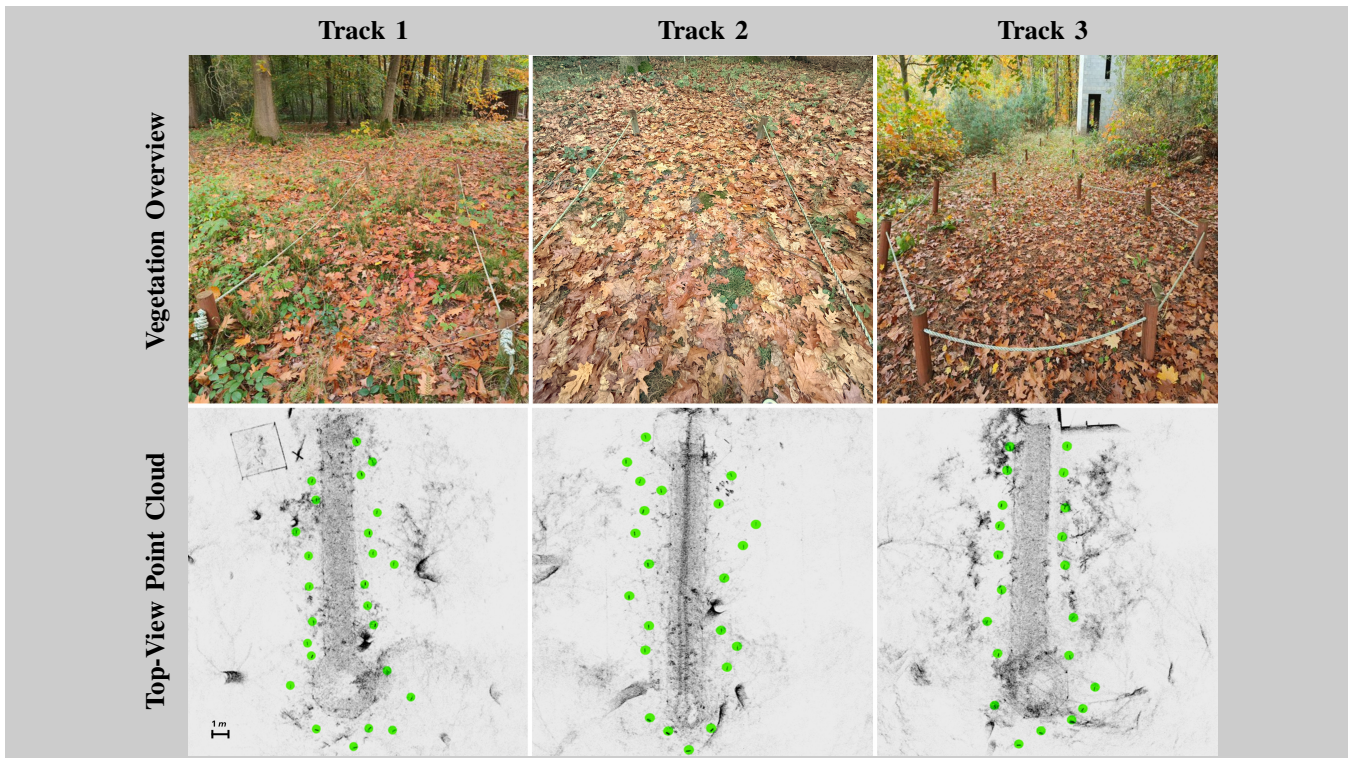


Fig. 3. Overview of the three tracks in the dataset. The first row displays an image of each track to show the condition of the vegetation. The second row presents a top-view point cloud representation, highlighting the locations of targets and illustrating their distribution across the tracks. Scale: 1 m (bottom left).

- **Track 1:** Forest floor with low to mid-high vegetation, mostly covered by freshly fallen deciduous leaves over moist clay-loam soil. Recorded between 13:00 and 14:15, with ambient air temperatures ranging from 12.8 °C to 13.6 °C. Conditions were overcast, giving diffuse light and a uniform thermal background;
- **Track 2:** Low to mid-high vegetation, with a more heterogeneous leaf litter than Track 1. The surface shows uneven layers of leaves with occasional gaps exposing the moist clay-loam soil beneath. Recorded between 15:15 and 16:05, with ambient air temperatures between 12.5 °C and 13.2 °C. Moisture variations created visible contrasts across modalities, with thermal differences between leaf-covered and exposed ground;
- **Track 3:** Denser and taller vegetation compared to the previous tracks, with soil partly visible under uneven leaf cover. Recorded between 17:30 and 18:00, under twilight conditions, with ambient air temperatures between 12.8 °C and 13.4 °C. Cooling soil and foliage created mixed thermal patterns. The RGB camera is unavailable for this track due to a hardware failure, although the twilight conditions would have limited its usefulness.

An overview of the vegetation conditions for each track, as well as the distribution of landmines and common items, is provided in Figure 3. In addition, we release a minute-by-minute climatology log covering both the target placement phase and the campaign itself. The log records atmospheric variables (temperature, humidity, wind, pressure, cloud coverage, precipitation) as well as soil temperatures at multiple

depths, providing the environmental context needed to interpret the thermal data.

C. Sequences

For each lane, we recorded three sequences: the reference sequence and two evaluation sequences. We use the reference sequence only for reference data generation, as detailed in section IV-D. The first and second evaluation sequences, referred to as Sequence 1 and Sequence 2, differ in the position of the robotic arm during data sampling. Figure 4 illustrates the arm’s position relative to the UGV, along with the rotation of wrist joint 1 and wrist joint 2 for each sequence. In **Sequence 1**, the UGV traverses the entire track back and forth while maintaining its position on the right side, at an average speed of 0.2 m/s. The robotic arm is fully extended vertically, with θ_1 tilted 45° downward. The sequence lasts about 3 to 4 minutes, depending on the track length. In **Sequence 2**, the UGV traverses the entire track back and forth while maintaining its position on the right side, at an average speed of 0.2 m/s. At approximately every meter, it stops to scan the environment on both sides using the robotic arm. Including these stops, the average speed for scanning the entire track is 0.054 m/s. The robotic arm remains fully extended in a vertical position, with both wrist joints 1 and 2 actuated. The end effector sweeps from left to right, with wrist joint 1 moving between $\theta_1 = 34^\circ$ and $\theta_1 = 70^\circ$, and wrist joint 2 moving between $\theta_2 = 55^\circ$ and $\theta_2 = 130^\circ$. The sequence lasts between 19 and 22 minutes, depending on the track. In total, we have recorded six sequences, providing a total




Fig. 4. *Left* - Arm stow and UGV motion: The manipulator is positioned relative to the UGV's body, with the vehicle's forward motion direction indicated by the red arrow. *Center* - End effector position in the first sequence: The arm is extended to show the initial placement and orientation of the end effector. *Right* - End effector position in the second sequence: The rotations of the arm's wrist joints are highlighted. The angles θ_1 and θ_2 correspond to the rotational movements of wrist joint 1 and wrist joint 2, respectively.



Fig. 5. *Left* - The track is ready to sample the reference sequence. *Right* - A MON-50 landmine and its corresponding AprilTag. The pose of each AprilTag, denoted as A_i , is evaluated, allowing determination of the position of the corresponding target T_i .

recording duration of approximately one hour. The reported frame counts correspond only to the sensors in the robotic arm sensor suite and are approximate: around 38,000 frames from the RGB camera, 53,000 frames from the VIS-SWIR camera, and 108,000 frames from the LWIR camera.

D. Reference data generation

As illustrated in Figure 6, targets can hardly be detected by the human eye in the images. Direct manual labeling is difficult, as many targets are barely visible in the raw imagery. Therefore, we implement an annotation pipeline that generates initial estimates using AprilTags [26], SLAM [27], ICP, and sensor kinematics, which are then adjusted by human intervention. We placed an AprilTag at each target location on every track and recorded a reference sequence. An example of a track prepared for the reference sequence, along with a sample target and its corresponding tag, is shown in Figure 5.

The main steps of the reference data generation process are the following:

- 1) **Target positioning in the reference sequence coordinate frame:** We detect AprilTags across all frames using the Sevensense Core Research module. With accurate intrinsic calibration of the camera (C_{SEV}), we determine each detected AprilTag's pose, including both position and orientation, relative to the camera frame. We denote this pose as ${}^{C_{SEV}}\mathbf{T}_{A_i} \in SE(3)$, where i indexes the detected AprilTag and $SE(3)$ represents the Special Euclidean group, which describes rigid body transformations in 3D space. Here, C_{SEV} refers to the Sevensense camera's coordinate frame.

Using the extrinsic calibration between the camera, the platform's LiDAR sensor (Livox Mid-360), and the associated IMU, we transform the AprilTag pose from the camera frame into the IMU coordinate frame as follows:

$${}^I\mathbf{T}_{T_i} = {}^I\mathbf{T}_L {}^L\mathbf{T}_{C_{SEV}} {}^{C_{SEV}}\mathbf{T}_{A_i} {}^{A_i}\mathbf{T}_{T_i} \quad (1)$$

where ${}^I\mathbf{T}_L$ and ${}^L\mathbf{T}_{C_{SEV}}$ denote the rigid body transformations from the LiDAR to the IMU and from the camera to the LiDAR, respectively. Additionally, ${}^{A_i}\mathbf{T}_{T_i}$ represents the transformation from the top left corner of the AprilTag to the actual target on the ground. The known offset along the stick plus a ground-plane intersection estimate is used to determine the actual target location from the AprilTag's pose.

We associate each image where a target is detected with the corresponding pose from the SLAM system, defined as the pose of the IMU in the world coordinate frame of the reference sequence, denoted ${}^{W_R}\mathbf{T}_I$. Consequently, we express the target poses in the reference sequence coordinate frame as:

$${}^{W_R}\mathbf{T}_{T_i} = {}^{W_R}\mathbf{T}_I {}^I\mathbf{T}_{T_i} \quad (2)$$

- 2) **Computing the transformation between the reference sequence coordinate frame and the evaluation sequence coordinate frame:**

We execute a point-to-plane ICP algorithm on the point clouds generated by running the SLAM system on both the reference sequence and the evaluation sequence. This process estimates the transformation from the reference sequence coordinate system (W_R), to the evaluation sequence coordinate frame W_E .

We denote this transformation as ${}^{W_E}\mathbf{T}_{W_R}$.

- 3) **Transforming the target locations into the evaluation sequence coordinate frame:**

After computing the transformation ${}^{W_E}\mathbf{T}_{W_R}$, we apply it to the target locations initially expressed in the reference coordinate frame W_R . This step allows us to represent the target positions within the evaluation sequence coordinate frame W_E .

The transformation of each target location is computed as:

$${}^{W_E}\mathbf{T}_{T_i} = {}^{W_E}\mathbf{T}_{W_R} {}^{W_R}\mathbf{T}_{T_i} \quad (3)$$

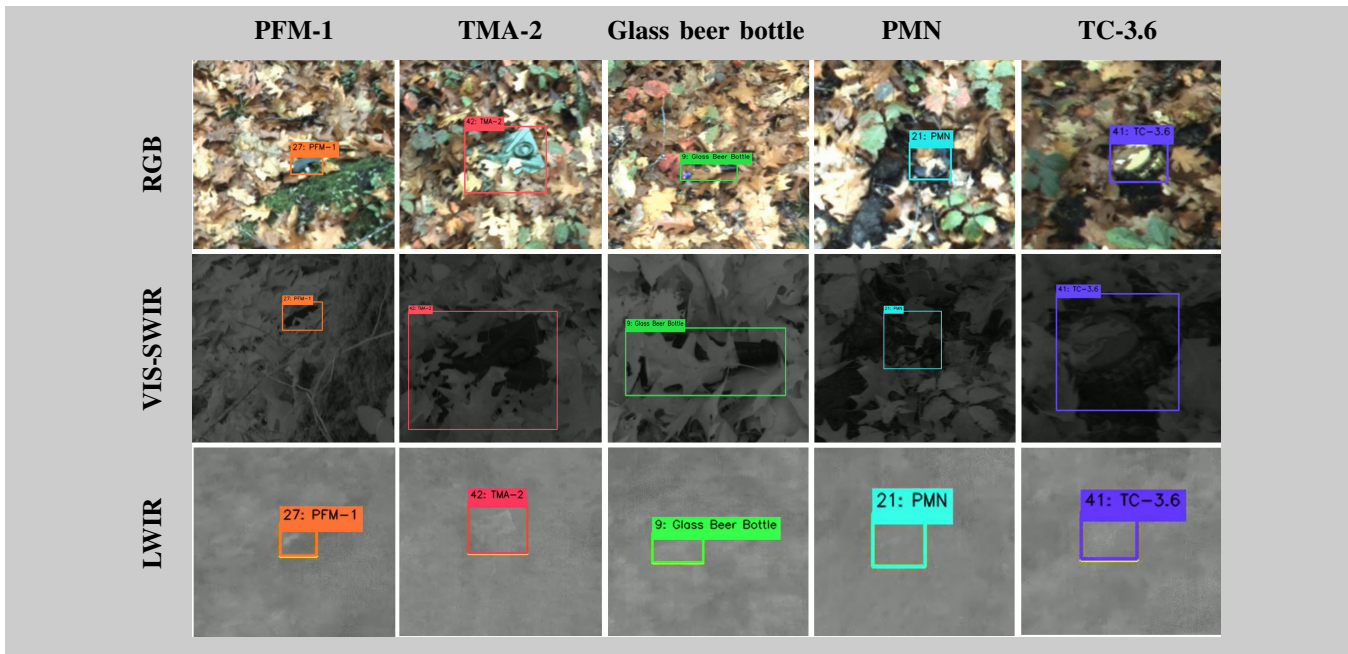


Fig. 6. Examples of five targets from the dataset, each captured by the robotic arm’s sensor suite in RGB, VIS–SWIR, and LWIR. Bounding boxes are obtained from the annotation pipeline and refined through human supervision to better match the visible extent of each object.

Here, ${}^{W_E}\mathbf{T}_{T_i}$ represents the pose of the i -th target in the evaluation sequence coordinate frame, while ${}^{W_R}\mathbf{T}_{T_i}$ denotes its corresponding pose in the reference frame.

4) Transforming into the UGV and camera coordinate frames:

To express the target poses in the UGV coordinate frame along its trajectory and in the camera coordinate frame, we apply the inverse transformations from Step 1, now within the evaluation sequence.

First, we transform the target pose from the evaluation sequence frame W_E to the IMU coordinate frame I :

$${}^I\mathbf{T}_{T_i} = ({}^{W_E}\mathbf{T}_I)^{-1} {}^{W_E}\mathbf{T}_{T_i} \quad (4)$$

In the case of the RGB camera, we reproject the target into the camera coordinate frame C_{RGB} as:

$${}^{C_{\text{RGB}}}\mathbf{T}_{T_i} = ({}^I\mathbf{T}_L {}^L\mathbf{T}_{C_{\text{RGB}}})^{-1} {}^I\mathbf{T}_{T_i} \quad (5)$$

where ${}^I\mathbf{T}_L$ represents the transformation from LiDAR to IMU, and ${}^L\mathbf{T}_{C_{\text{RGB}}}$ corresponds to the transformation from the RGB camera to the LiDAR frame. It incorporates the conversion from the RGB camera frame to the robot base and then to the Livox Mid-360 frame, including the necessary arm kinematics.

Finally, using the intrinsic and extrinsic calibration parameters of the respective camera, we project the transformed target locations onto the image plane to obtain the pixel coordinates of the center of each target, and apply a fixed-size bounding box according to the corresponding camera resolution.

5) Human-supervised adjustment of bounding boxes:

The automatically generated boxes serve as an initial guess. Annotators adjust them to better match the visible

extent of each target, mitigating errors propagated from AprilTag detection, pose estimation, or perspective distortions. RGB and VIS–SWIR annotations are manually revised. For LWIR, where signatures are often faint, adjustments rely on reprojections from the corrected RGB and VIS–SWIR annotations, with human input.

V. BASELINE EVALUATION

As a baseline, we trained YOLOv8 [28] on the SULAND dataset [17], which is the closest publicly available dataset to our scenario. Even with extensive augmentation strategies, the model did not transfer successfully to MineInsight and failed to produce reliable detections. Rather than serving as a performance benchmark, this outcome highlights the domain gap between existing datasets and MineInsight, driven by its clutter-rich environments with vegetation and distractor objects, as well as a broader variety of landmine types. This underlines MineInsight’s value as a benchmark for studying domain adaptation and multi-modal fusion in realistic demining scenarios.

VI. LIMITATIONS

MineInsight is, to the best of our knowledge, the first multi-modal dataset for landmine detection with RGB, VIS–SWIR, and LWIR sensing. Its current release, however, has some constraints. All data were collected during a single late-autumn campaign under overcast conditions. As a result, models trained on this dataset may face domain gaps when applied to different seasons, terrains, or weather conditions, such as sunny, snowy, arid, or grassy environments. The release also lacks RGB data for Track 3, though twilight conditions would in any case have limited its usefulness. While MineInsight surpasses existing public datasets in scale and

sensor diversity, the number of targets and total recording time remain modest relative to real-world scenarios. More broadly, the dataset highlights the domain-adaptation challenges faced when transferring models trained elsewhere to MineInsight. Finally, annotations are adjusted with human supervision; minor inaccuracies may remain, particularly for faint thermal signatures. These limitations motivate the future extensions outlined in the conclusion.

VII. CONCLUSION

We introduced MineInsight, a multi-sensor, multi-spectral dataset for robotic landmine detection in off-road environments. It combines dual viewpoints (UGV and robotic arm), multi-spectral imaging (monochrome, RGB, VIS-SWIR, LWIR), dual LiDARs, and synchronized data to address occlusions and sensor limitations. The dataset includes bounding boxes generated by an automated pipeline with a human-in-the-loop process to support the evaluation of detection algorithms, and it can also serve generic off-road robotics tasks such as navigation and mapping. In addition, the dataset includes minute-by-minute climatology measurements, which provide valuable context for exploiting LWIR data and studying thermal contrast under varying conditions. As noted in Section VI, the current release is limited to a single late-autumn campaign and may introduce domain gaps under other environmental conditions. Future work will expand MineInsight with recordings across varied terrains, climates, and seasons, additional sensors, new tracks, and further landmine types. We also aim to explore spectral fusion, landmine segmentation, and synthetic data for domain adaptation. Finally, the issue of RGB degradation at night, observed in Track 3, could be further studied with lighting systems to extend visibility.

ACKNOWLEDGMENT

The authors thank Alessandra Miuccio and Timothée Fréville for their support in the hardware and software design. They also thank Sanne Van Hees and Jorick Van Kwikenborne for their assistance in organizing the measurement campaign. Finally, they thank the Belgian Meteo Wing for providing the climatology study during the days of the test campaign.

REFERENCES

- [1] I. Nevliudov, D. Yanushkevych, and L. Ivanov, "Analysis of the state of creation of robotic complexes for humanitarian demining," *Technology audit and production reserves*, vol. 6, pp. 47–52, 2021.
- [2] M. O. Popov, S. A. Stankevich, S. P. Mosov, O. V. Titarenko, M. V. Topolnytskyi, and S. S. Dugin, "Landmine detection with UAV-based optical data fusion," in *IEEE International Conference on Smart Technologies (EUROCON)*, 2021, pp. 175–178.
- [3] 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: 859, 2020.
- [4] S. Colreavy-Donnelly, F. Caraffini, S. Kuhn, M. Gongora, J. Florez, and C. Parra, "Shallow buried improvised explosive device detection via convolutional neural networks," *Integrated Computer-Aided Engineering*, vol. 27, no. 4, pp. 403–416, 2020.
- [5] W. Lin and Q. Wang, "MineSeeker: A novel design for detecting explosive mines using low-resolution thermal imaging," in *IEEE International Conference on Imaging Systems and Techniques (IST)*, 2024, pp. 1–5.
- [6] M. Bajić and B. Potočník, "UAV thermal imaging for unexploded ordnance detection by using deep learning," *Remote Sensing*, vol. 15, no. 4: 967, 2023.
- [7] R. Mohamed Fikry and H. Kasban, "Deep neural networks for landmines images classification," in *International Conference on Advanced Intelligent Systems and Informatics (AISI)*, 2020, pp. 126–136.
- [8] M. Wigness, S. Eum, J. G. Rogers, D. Han, and H. Kwon, "A RUGD dataset for autonomous navigation and visual perception in unstructured outdoor environments," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 5000–5007.
- [9] P. Jiang, P. Osteen, M. Wigness, and S. Saripalli, "RELLIS-3D dataset: Data, benchmarks and analysis," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 1110–1116.
- [10] H. Teng, Y. Wang, X. Song, and K. Karydis, "Multimodal dataset for localization, mapping and crop monitoring in citrus tree farms," in *International Symposium on Visual Computing (ISVC)*, 2023, pp. 571–582.
- [11] J. Baur, K. Dewey, G. Steinberg, and F. O. Nitsche, "Modeling the effect of vegetation coverage on unmanned aerial vehicles-based object detection: A study in the minefield environment," *Remote Sensing*, vol. 16, no. 12: 2046, 2024.
- [12] T. de Smet, A. Nikulin, and J. Baur. (2020) Scatterable landmine detection project datasets 1-7. https://orb.binghamton.edu/geology_fac/4. Accessed: February 11, 2026.
- [13] G. Steinberg, J. Baur, A. Nikulin, and T. de Smet. (2020) Scatterable landmine detection project datasets 1-7. https://orb.binghamton.edu/geology_fac/11. Accessed: February 11, 2026.
- [14] J. Baur, T. De Smet, A. Nikulin, and G. Steinberg. (2020) Scatterable landmine detection project datasets 10-24. https://orb.binghamton.edu/geology_fac/15. Accessed: February 11, 2026.
- [15] U. M. Leloglu and S. Kaya, "Landmine thermal image series," *IEEE Dataport*, 2019.
- [16] A. Tenorio Tamayo, J. Forero Ramirez, B. Garcia, H. Loaiza-Correa, A. Restrepo-Girón, S. Nope Rodríguez, A. Barandica-López, and J. T. Buitrago, "Dataset of thermographic images for the detection of buried landmines," *Data in Brief*, vol. 49, p. 109443, 2023.
- [17] E. Vivoli, M. Bertini, and L. Capineri, "Deep learning-based real-time detection of surface landmines using optical imaging," *Remote Sensing*, vol. 16, no. 4: 677, 2024.
- [18] S. Kaya and U. M. Leloglu, "Buried and surface mine detection from thermal image time series," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 10, no. 10, pp. 4544–4552, 2017.
- [19] B. Dwyer, J. Nelson, and T. e. a. Hansen. (2024) RoboFlow (Version 1.0). Available: <https://roboflow.com/>. Accessed: February 11, 2026.
- [20] J. Rehder, J. Nikolic, T. Schneider, T. Hinzmann, and R. Siegwart, "Extending kalibr: Calibrating the extrinsics of multiple IMUs and of individual axes," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 4304–4311.
- [21] P. Furgale, J. Rehder, and R. Siegwart, "Unified temporal and spatial calibration for multi-sensor systems," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2013, pp. 1280–1286.
- [22] S. Schramm, J. Rangel, D. Aguirre Salazar, R. Schmoll, and A. Kroll, "Multispectral geometric calibration of cameras in visual and infrared spectral range," *IEEE Sensors*, vol. 21, no. 2, pp. 2159–2168, 2021.
- [23] R. Usamentiaga, D. Garcia, C. Ibarra-Castanedo, and X. Maldague, "Highly accurate geometric calibration for infrared cameras using inexpensive calibration targets," *Measurement*, vol. 112, pp. 105–116, 2017.
- [24] K. Koide, S. Oishi, M. Yokozuka, and A. Banno, "General, single-shot, target-less, and automatic LiDAR-camera extrinsic calibration toolbox," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 11 301–11 307.
- [25] M. Malizia, C. Hamesse, K. Hasselmann, G. De Cubber, N. Tsiogkas, E. Demeester, and R. Haelterman. (2025) MineInsight: A Multi-sensor Dataset for Humanitarian Demining Robotics in Off-Road Environments - GitHub Repository. Available: <https://github.com/marioml299/MineInsight>. Accessed: February 11, 2026.
- [26] E. Olson, "AprilTag: A robust and flexible visual fiducial system," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2011, pp. 3400–3407.
- [27] T.-M. Nguyen, D. Duberg, P. Jensfelt, S. Yuan, and L. Xie, "SLICT: Multi-input multi-scale surfel-based LiDAR-inertial continuous-time odometry and mapping," *IEEE Robotics and Automation Letters*, vol. 8, no. 4, pp. 2102–2109, 2023.
- [28] G. Jocher, A. Chaurasia, and J. Qiu, "YOLO by ultralytics," 2023, available: <https://github.com/ultralytics/ultralytics>.