

Reducing Object Detection Uncertainty from RGB and Thermal Data for UAV Outdoor Surveillance

Juan Sandino^{◇‡}, Peter A. Caccetta[‡], Conrad Sanderson^{†‡}, Frederic Maire[◇], Felipe Gonzalez[◇]

[◇] Queensland University of Technology, Australia

[†] Griffith University, Australia

[‡] Data61 / CSIRO, Australia

Abstract—Recent advances in Unmanned Aerial Vehicles (UAVs) have resulted in their quick adoption for wide a range of civilian applications, including precision agriculture, biosecurity, disaster monitoring and surveillance. UAVs offer low-cost platforms with flexible hardware configurations, as well as an increasing number of autonomous capabilities, including take-off, landing, object tracking and obstacle avoidance. However, little attention has been paid to how UAVs deal with object detection uncertainties caused by false readings from vision-based detectors, data noise, vibrations, and occlusion. In most situations, the relevance and understanding of these detections are delegated to human operators, as many UAVs have limited cognition power to interact autonomously with the environment. This paper presents a framework for autonomous navigation under uncertainty in outdoor scenarios for small UAVs using a probabilistic-based motion planner. The framework is evaluated with real flight tests using a sub 2 kg quadrotor UAV and illustrated in victim finding Search and Rescue (SAR) case study in a forest/bushland. The navigation problem is modelled using a Partially Observable Markov Decision Process (POMDP), and solved in real time onboard the small UAV using Augmented Belief Trees (ABT) and the TAPIR toolkit. Results from experiments using colour and thermal imagery show that the proposed motion planner provides accurate victim localisation coordinates, as the UAV has the flexibility to interact with the environment and obtain clearer visualisations of any potential victims compared to the baseline motion planner. Incorporating this system allows optimised UAV surveillance operations by diminishing false positive readings from vision-based object detectors.

I. INTRODUCTION

Recent advances in autonomous navigation of Unmanned Aerial Vehicles (UAVs)—also known as drones—have resulted in their gradual adoption in a set of civilian and time-critical applications such as surveillance, disaster monitoring, and Search and Rescue (SAR) [11], [20], [24], [25]. UAVs offer unique benefits such as compact sizes and low cost to scout outdoor and indoor environments, real-time telemetry and camera streaming to monitor challenging and otherwise inaccessible environments, extensive payload adaptability, and extensive possibilities to augment navigation capabilities through software [8], [16], [26], [31].

One critical challenge in deploying UAVs and robots in general into real-world and time-critical applications is the ever-presence of uncertainty. Factors that cause uncertainty are diverse, and they can be classified as external or internal. External factors come from sources beyond the scope of the UAV, such as poor weather and illumination conditions, strong

gusts, unknown situational-awareness of surveyed environments, and partial observability. Internal factors include sub-optimal camera calibration settings, low image resolution, noisy camera frames during streaming, or imperfect detection outputs from computer vision detectors. As shown in Fig. 1, uncertainty sources that are poorly managed can compromise the behaviour of UAVs and the flight mission itself [21]. Thus, it is essential to incorporate cognitive capabilities in UAVs to broaden their use in more real-world scenarios [12].

The elevated number of stranded people and human loss is a problem that is far from solved [1]. In Australia alone, an average of 38,000 people per year are reported missing and around 2% of them (or 720 persons) are never located [4]. In the event of an emergency—where time management plays a critical factor in the success of the rescue operation—the goal to identify and locate as many victims as quick as possible. Thus, UAV technology for autonomous navigation and victim detection in challenging environments could assist first-responders in locating as many victims as soon as possible.

Research works on applied decision-making theory in UAVs is extensive and indicates that using Partially Observable Markov Decision Processes (POMDPs) onboard UAVs can increase their cognitive capabilities for autonomous navigation



Fig. 1. Unmanned aerial vehicle (UAV) navigating in environments under uncertainty and partial observability. A small UAV with autonomous decision-making should be able to plan sequential sets of actions for optimal navigation trajectories, despite limitations from imperfect sensor data.

and object detection under uncertainty [5], [14], [27]. UAV frameworks for object detection and tracking have been tested in cluttered indoor environments and in the absence of Global Navigation Satellite System (GNSS) coverage [31], [32], [34], [35]. POMDPs have also been applied to solve multi-objective problems in UAVs, addressing tasks such as path planning, multiple object detection and tracking, and collision prevention [28], [29].

In time-critical applications such as SAR, real-time camera streaming is critical to comprehend the context of the environment [22]. However, drone pilots have a strong reliance on their communication systems to control most UAVs. If communication systems fail, the usability of the UAV could be seriously compromised [33]. Many approaches of POMDPs applied in UAVs for humanitarian relief operations have been tested in simulation [3], [36] and very few systems have been evaluated with real flight test using trivial targets [10].

Research efforts on onboard decision-making under object detection uncertainty from Convolutional Neural Network (CNN) models are scarce. Research conducted in [31], [32] described a framework and POMDP problem formulation for a SAR application in GNSS-denied environments with a sub 2 kg UAV. However, the framework was only tested in cluttered indoor scenarios.

This paper describes a modular UAV framework for autonomous onboard navigation in outdoor environments under uncertainty. The framework design aims to reduce levels of object detection uncertainty using a POMDP-based motion planner, which allows the UAV to interact with the environment to obtain better visual representations of detected objects. CNN-based computer vision inference and motion planning can be executed in resource-constrained hardware onboard small UAVs. The framework is tested with real flight tests with a simulated SAR mission, which consisted of finding an adult mannequin in an open area and close to a tree. Three flight modes are proposed to evaluate the feasibility of the framework for real-world SAR operations.

This paper extends the work in [30], [31], [32] with the following contributions: (1) an extension of their evaluated UAV framework—originally designed for navigation in GNSS-denied environments—for outdoor missions with GNSS signal coverage, and the design of a novel flight mode; (2) an additional validation of preliminary results of their proposed UAV framework with comprehensive real flight tests; and (3) a scalability approach of the framework by adapting a thermal camera and a custom object detector to locate victims using their heat signatures.

The rest of the paper is structured as follows. Section II details the UAV framework design for autonomous object detection in uncertain outdoor environments. Section III summarises the implemented probabilistic-based motion planner using a POMDP. The design of conducted experiments using real flight tests is presented in Section IV. Obtained results and discussion of performance indicators are provided in Section V. Conclusions and future avenues for research are discussed in Section VI.

II. FRAMEWORK DESIGN

The framework follows a modular system architecture for autonomous navigation onboard small UAVs as illustrated in Fig. 2. This design extends an existing UAV framework for autonomous navigation in cluttered environments under object detection uncertainty, tested in simulation and with real flight tests in a sub 2 kg quadcopter [32].

Fig. 2 illustrates the physical environment (or world) composed by the UAV frame and any attached payloads (i.e., RGB or thermal cameras), the victim and obstacles. Acquired camera frames represent the visual interface (also called observations) of the surveyed environment by the UAV. The UAV also contains the *autopilot*, which translates high-level action commands into low-level signals that control the UAV motors. The last hardware component of the UAV frame is a companion computer, which is allocated to execute software algorithms in dedicated modules for computer vision, mapping, and real-time path planning. Action commands from the planner are managed by the motion module, which interfaces with the flight controller of the autopilot.

The following subsections discuss each of the proposed framework components. The UAV framework used in this work is not limited to the hardware and software discussed below. Other UAV frame designs, payloads, autopilots, vision-based object detectors, planners, and software toolkits can also be implemented.

A. UAV Airframe and Payloads

The UAV airframe which offered the best combination between payload adaptability, size, and endurance for this research is a Holybro X500 quadrotor kit (Holybro, China). As shown in Fig. 3, key components utilised from the kit include a Pixhawk 4 autopilot, Pixhawk 4 GNSS receiver, 2216 KV880 brushless motors, 22.86 cm plastic propellers, and a 433 MHz Telemetry Radio. With dimensions of 41 cm \times 41 cm \times 30.0 cm, the UAV carries a four cell 5000 mAh LiPo battery, for an approximate flight autonomy of 12 min.

The companion computer is an Intel UP², chosen for its price tag, number of peripherals and CPU architecture. Key specifications include a 64-bit quad-core CPU at 1.1 GHz, 64 GB eMMC SSD, 8 GB RAM, four FL110 USB 3.0 connectors, two High-Speed UART controllers, and one mPCIe connector.

The proposed framework was tested using two Red, Green, Blue (RGB) cameras, namely an Arducam B019701 and a GoPro Hero 9. Thermal imagery is sourced from a FLIR Tau 2 connected to a ThermalCapture device for real-time frame streaming. The cameras, which can be interchangeably used in the proposed framework, are mounted onto an anti-vibration bracket, pointing to the ground and in parallel to Earth's nadir, as seen in Fig. 4. Core properties for the cameras can be found in Tab. IV in the Appendix.

B. Vision Module

This module consists of a deep learning object detector processing raw frames from the GoPro Hero 9 camera. Taking into account the performance limitations of running deep

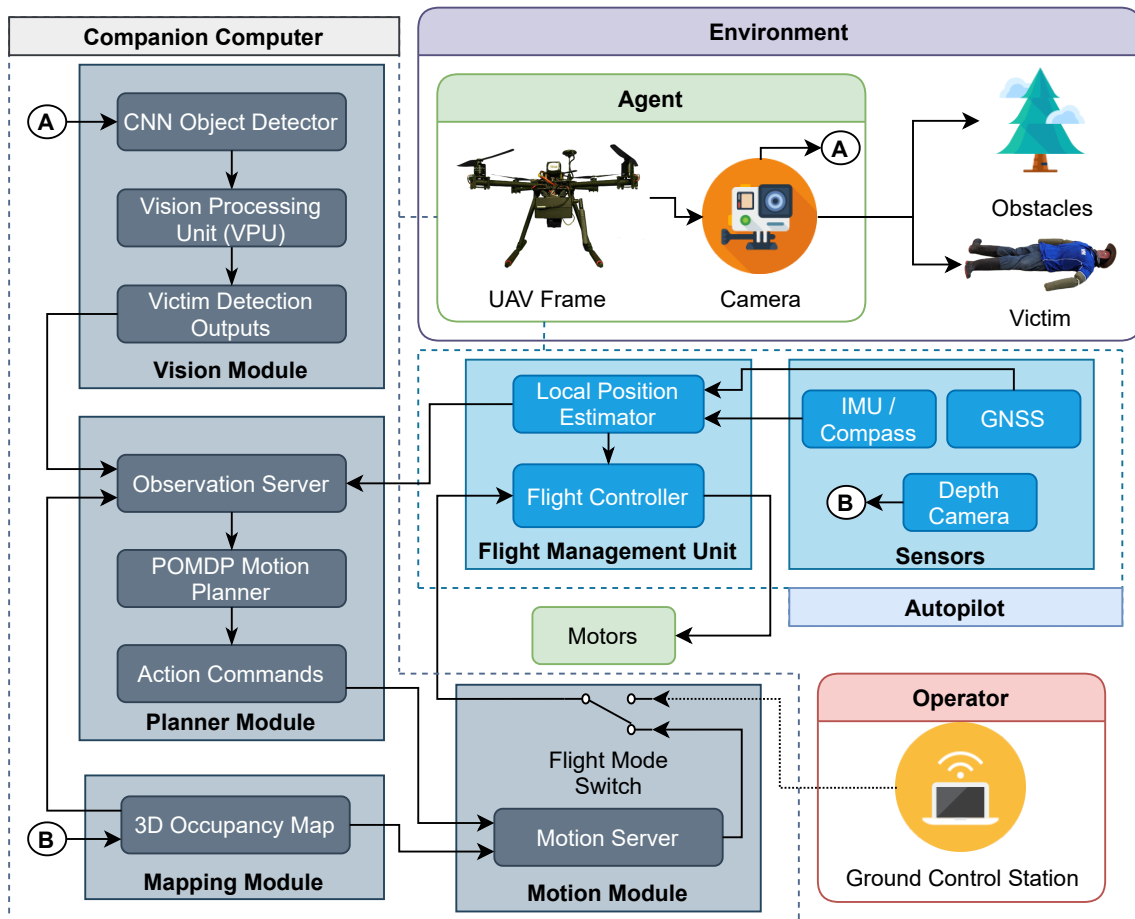


Fig. 2. Modular system architecture for autonomous navigation onboard UAVs in uncertain outdoor environments. The framework portrays the physical environment (or world) composed of the UAV frame, attached payloads, world obstacles, and the victim. A companion computer is attached to the UAV to execute software algorithms in dedicated modules for computer vision, mapping, real-time path planning, and a motion server that interfaces the companion computer with the UAV autopilot.

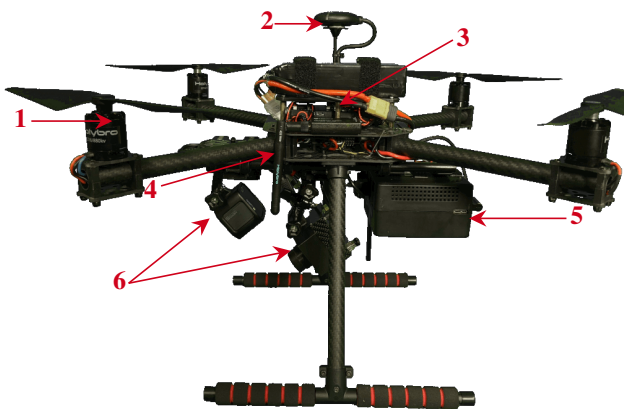


Fig. 3. Framework implementation in a sub 2 kg quadrotor UAV. Primary components include: (1) carbon fibre Holybro X500; (2) Pixhawk 4 GNSS receiver; (3) Pixhawk 4 autopilot; (4) 433 MHz telemetry radio; (5) Intel UP² companion computer; (6) payload.



Fig. 4. GoPro Hero 9 mounted onto an anti-vibration bracket, pointing to the ground and in parallel to Earth's nadir.

learning models in resource-constrained hardware, a Vision Processing Unit (VPU) is installed in the companion computer. Convolutional operations that normally run onboard a CPU or GPU are allocated to the VPU for inference of CNN models in resource-constrained hardware. In this implementation, the selected VPU is an Intel Movidius Myriad X, which is connected to the companion computer via the mPCIe slot. The detection module is programmed in Python and uses the OpenVINO library to optimise code instructions to load CNN models into the VPU.

The deep learning model architecture used to detect victims is an off-the-shelf Google MobileNet Single-Shot Detector (SSD) [7]. This model is deployed in Caffe [15] and tuned with pre-trained weights from the PASCAL VOC2012 dataset [9], scoring a mean average precision of 72.7%. The dataset covers up to 21 class objects (including persons). However, only positive detections for the class person are evaluated. Acquired camera frames are fitted into the input layer of the neural network (i.e., MobileNet SSD model) by downsizing the frames to 300×300 pixels.

C. Mapping Module

The Mapping module manages 3D occupancy maps, which are constituted by volumetric occupancy grids and displays the presence and localisation of objects in the surveyed environment. In this implementation, the 3D Occupancy Map are requested by the Motion and Planner modules to evaluate the presence of obstacles at selected position coordinates in the world coordinate frame. The maps are created through the use of the Octomap library [13].

D. Planner Module

The planner module computes the motion policy of the UAV and contains three primary components: (1) an observation server, which handles raw observations from the vision module (i.e., detected victims, confidence and victim coordinates), local position estimations of the UAV, and the state of the 3D occupancy map; (2) the POMDP motion planner, which calls the observation server every time the planner requires new observations; and (3) action commands computed by the motion policy of the planner and are read by the motion server. Complete details of the POMDP planner design can be found in Section III.

E. Communication Interface

The UAV framework runs on open-source software tools. The companion computer runs under Linux Ubuntu 18.04 LTS O.S. and the Robot Operating System (ROS) melodic. ROS is a middleware to communicate between the nodes of each module (following the architecture design from Fig. 2). The Pixhawk 4 autopilot is powered by PX4, which communicates with the companion computer through MAVROS, a ROS implementation of the MAVLink protocol, which is industry standard for UAV communication and control [18]. The POMDP solver implementation, which is described in Section III, also contains a ROS implementation to maximise the use of visualisation, telemetry and recording tools from ROS.

III. PLANNER DESIGN

This approach formulates the decision-making problem as a POMDP. The planner transmits UAV position commands to the motion planner derived from environment observations. The discussion presented in this section is adapted from [30] and only essential parts are shown in this paper for completeness.

With a taken action $a \in A$, the UAV receives an observation $o \in O$ encoded by the observation function $\mathcal{Z}(s', a, o) = \mathbb{P}(o | s', a)$. Every decision chain is then quantified with an estimated reward r , calculated using the reward function $R(a, s)$. A POMDP uses a probability distribution over the system states to model uncertainty of its observed states. This modelling is called the belief $b(H) = \mathbb{P}[s^1 | H], \dots, \mathbb{P}[s^n | H]$, where H is the history of actions, observations and rewards the UAV has accumulated until a time step t , or $H = a_0, o_1, r_1, \dots, a_{t-1}, o_t, r_t$.

The motion policy π of the UAV is represented by mapping belief states to actions $\pi : b \rightarrow A$. A POMDP is solved after finding the optimal policy π^* , calculated as follows:

$$\pi^* := \arg \max_{\pi} \left(\mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t R(S_t, \pi(b_t)) \right] \right) \quad (1)$$

where $\gamma \in [0, 1]$ is the discount factor and defines the relative importance of immediate rewards compared to future rewards. A given POMDP solver starts planning from an initial belief b_0 , which is usually generated using the initial conditions (and assumptions) of the flight mission.

A. Assumptions

In this implementation, the formulated problem for exploration and object detection (i.e., victims) using multi-rotor UAVs in outdoor environments assumes:

- An initial 3D occupancy map of the environment is pre-loaded to the planner before the UAV takes off.
- Observations come from processed camera frames (by the Vision module), the 3D occupancy map (by the Mapping Module), and the estimated local UAV position (by the Autopilot).
- Only a single, and static victim can be detected at the same time. If more victims appear on processed camera frames, the planner will only read data from the victim with the highest detection confidence values.
- The motion planner starts once the UAV reaches known position setpoint (i.e., at one of the corners of the surveyed area).
- The planner stops computing a motion policy once: (1) the UAV detects a victim whose detection confidence surpasses a set threshold; (2) the UAV explores the whole search area extent without finding any victims; or (3) the UAV exceeds the maximum flight time on air (because of low levels of battery power).

B. Actions

UAV actions are defined by seven position commands, namely *forward*, *backward*, *left*, *right*, *up*, *down*, and *hover*. UAV actions that are not included in the action space but are managed by the autopilot instead include *arm*, *disarm*, *take-off*, *return to launch and land*. Each action updates the position set point of the UAV in the world coordinate frame by calculating and applying a change of position δ .

TABLE I
APPLIED REWARD VALUES TO THE REWARD FUNCTION R , DEFINED IN
ALGORITHM 1.

Variable	Value	Description
r_{crash}	-50	Cost of UAV crash
r_{out}	-25	Cost of UAV breaching safety limits
r_{dte}	+25	Reward for detecting potential victim
r_{conf}	+50	Reward for confirmed victim detection
r_{action}	-2.5	Cost per action taken
r_{fov}	-5	Footprint overlapping cost

The magnitude for δ_x and δ_y depends on the estimated overlap value between camera frame observations:

$$\delta = l_{\text{FOV}}(1 - \lambda) \quad (2)$$

where l_{FOV} is the length of the projected camera's Field of View (FOV), and $\lambda \in [0, 1]$ is the desired overlap value.

C. States

A system state $s \in S$ is defined as:

$$s = (p_u, f_{\text{roi}}, f_{\text{dct}}, p_v, c_v) \quad (3)$$

where p_u is the position of the UAV in the world coordinate frame, f_{crash} is a flag raised when the UAV crashes with an obstacle, f_{roi} is a flag indicating whether the UAV is flying beyond the flying limits, f_{dct} is the flag raised if a potential victim is detected by the UAV. If $f_{\text{dct}} = \text{True}$, the position of the victim in the world coordinate frame is given in p_v , with detection confidence $c_v \in [0, 1]$. The system reaches a terminal state whenever $c_v \geq \zeta$, where ζ is the confidence threshold.

D. Transition Function

The motion dynamics of a multi-rotor UAV defines the transition from current to new states:

$$p_u(k+1) = p_u(k) + \Delta p_u(k) \quad (4)$$

where $p_u(k)$ is the position of the UAV at time step k , and $\Delta p_u(k)$ is the position change of the UAV between time steps. This formulation does not contain any actions for heading changes. However, Eqn. 4 can be expanded if required by adding the rotation matrix in multi-copters [6]. An illustration of a problem formulation including the rotation matrix can be found in [31].

E. Reward Function

The expected reward r after taking an action $a \in A$ from state $s \in S$ is calculated using the reward function $R(a, s)$ defined in Algorithm 1 and Table I. This function critically influences the UAV behaviour during flight missions, and its definition allows multi-objective task definition. A complete discussion on the design considerations of the reward function can be found in [30].

The order of the steps from Algorithm 1 classifies high-level tasks into two components. The first one is object detection and starts by evaluating any states which will negatively affect the integrity of the UAV, followed by states indicating positive

Algorithm 1 Reward function R for exploration and object detection in outdoor environments.

```

1:  $r \leftarrow 0$ 
2: if  $f_{\text{crash}}$  then
3:    $r \leftarrow r_{\text{crash}}$  ▷ UAV crashing cost
4: else if  $f_{\text{roi}}$  then
5:    $r \leftarrow r_{\text{out}}$  ▷ Beyond safety limits cost
6: else if  $f_{\text{dct}}$  then
7:    $r \leftarrow r_{\text{dte}}$  ▷ Detected object reward
8:    $r \leftarrow r + \left[ r_{\text{dte}} \cdot \left( 1 - \frac{z_u - z_{\text{min}}}{z_{\text{max}} - z_{\text{min}}} \right) \right]$  ▷ UAV altitude reward
9:   if  $c_v \geq \zeta$  and  $a = \text{Down}$  then
10:     $r \leftarrow r + r_{\text{conf}}$ 
11:   end if
12: else
13:    $r \leftarrow r_{\text{action}}$  ▷ Action cost
14:    $r \leftarrow r - \left[ r_{\text{dte}} \cdot \left( 1 - \frac{z_u - z_{\text{min}}}{z_{\text{max}} - z_{\text{min}}} \right) \right]$  ▷ UAV altitude cost
15:    $r \leftarrow r - \left[ r_{\text{dte}} \cdot \left( 1 - 0.5 \cdot 4^{d_v/d_w} \right) \right]$  ▷ Horizontal distance cost
16:    $r \leftarrow r + r_{\text{fov}} \cdot \varepsilon$  ▷ Footprint overlap cost
17: end if
18: return  $r$ 

```

victim detections. If the UAV detects a potential victim (Step 6), R calculates a linear function (Step 8) which returns increased reward values as the UAV gets closer to the minimum allowed altitude. A higher reward value is returned if a potential victim is confirmed (Step 9 and 10).

The second component of the algorithm addresses exploration. In case there are no detections, R applies a set of cost functions to encourage a greedy horizontal exploration of the environment. An exponential function in Step 14 calculates the Manhattan distance between the UAV and the victim d_v and the maximum exploration distance d_w which are defined as:

$$d_v = \sum_{i=1}^n |p_i - q_i|, \quad p_i=(x_u, y_u), \quad q_i=(x_v, y_v) \quad (5)$$

$$d_w = \sum_{i=1}^n |p_i - q_i|, \quad p_i=(x_{\text{max}}, y_{\text{max}}), \quad q_i=(x_{\text{min}}, y_{\text{min}}) \quad (6)$$

The overlap ε between the camera's current footprint and its correspondent location in the footprint map is defined as:

$$\varepsilon = \frac{\sum_{i=1}^n F_i(p_u)}{n}, \quad \varepsilon \in [0, 1] \quad (7)$$

where $F_i(p_u)$ are the pixel values of the projected FOV in the footprint map, and n is the total number of projected pixels in the footprint. A maximum overlap value of 1 indicates that such action will place the UAV to a fully previously explored area, and, as indicated in Step 16 of Algorithm 1, the whole penalty value r_{fov} will be added to the reward. A minimum value of 0 means that a given action will place the UAV in an unexplored area and no penalty will be added to the reward. Intermediate values of ε represent partial overlapping, adding a partial penalty value r_{fov} to the reward.

F. Observations

An observation $o \in O$ is defined as:

$$o = (o_{p_u}, o_{dtc}, o_{p_v}, o_{\zeta}, o_{obs}) \quad (8)$$

where o_{p_u} is the estimated position of the UAV by the autopilot; o_{dtc} is the flag triggered by potential victim detections received by the CNN model; o_{p_v} and o_{ζ} are the local position of the victim and the detection confidence respectively, both of them defined only if there are any positive detections; and o_{obs} is the flag triggered after processing the 3D occupancy map for any obstacles located in front of the UAV.

The detection confidence o_{ζ} measures the frequency of positive detections between the last two observation calls:

$$o_{\zeta} = \frac{\sum_{i=1}^n o_{dtc_i}}{n} \quad (9)$$

where n is the number of segmented frames between observation calls, and o_{dtc} is the flag indicating a positive detection per processed frame i .

G. Observation Model

This implementation uses Augmented Belief Trees (ABT) [19], an online POMDP solver that contains a model that generates T and Z using a modelled observation o given an action a and the next state s' . The variables contained in the generative model are the local position of the UAV s'_{p_u} , the local position of the victim s'_{p_v} and the detection confidence o_{ζ} .

Potential victim detections and their subsequent positioning estimations are conditioned by the camera pose at the UAV frame and its projected footprint of the environment. Specifically, if the 2D local position coordinates of the victim $s'_{p_v}(x, y)$ are within the projected footprint limits of the camera, the victim is assumed to be detected. This estimation is done by calculating the sum of angles between a 2D point (i.e., s'_{p_v}) and each pair of points that constitute the footprint boundaries (the footprint rectangular corners) [2]. The 2D projected footprint extent l of a vision-based sensor, illustrated in Fig. 5, can be calculated using:

$$l_{\text{top, bottom}} = s'_{p_u}(z) \cdot \tan\left(\alpha \pm \tan^{-1}\left(\frac{h}{2f}\right)\right) \quad (10)$$

$$l_{\text{left, right}} = s'_{p_u}(z) \cdot \tan\left(\beta \pm \tan^{-1}\left(\frac{w}{2f}\right)\right) \quad (11)$$

where s'_{p_u} is the UAV altitude, α and β are the camera's pointing angles from the vertical z and horizontal x axis of the World coordinate frame, w is the lens width, h is the lens height, and f is the focal length.

The footprint corners c from the camera's local coordinate frame I are translated to the world's coordinate frame W using:

$$\begin{bmatrix} c'(x) \\ c'(y) \end{bmatrix} = \begin{bmatrix} s'_{p_u}(x) \\ s'_{p_u}(y) \end{bmatrix} + \begin{bmatrix} \cos(\varphi_u) & -\sin(\varphi_u) \\ \sin(\varphi_u) & \cos(\varphi_u) \end{bmatrix} \begin{bmatrix} c(x) \\ c(y) \end{bmatrix} \quad (12)$$

where s'_{p_u} is the next UAV position state, and φ_u is the Euler yaw angle of the UAV. However, as no actions involve

adjusting the heading of the UAV mid-flight, and assuming yaw estimation errors are negligible, Eqn. 12 is simplified as:

$$\begin{bmatrix} c'(x) \\ c'(y) \end{bmatrix} = \begin{bmatrix} s'_{p_u}(x) + c(x) \\ s'_{p_u}(y) + c(y) \end{bmatrix} \quad (13)$$

The detection confidence o_{ζ} that comes as part of the output data from the CNN object detector is modelled using:

$$o_{\zeta} = \frac{(1 - \zeta_{\min})(d_{uv} - z_{\min} + \zeta_{\min})}{z_{\max} - z_{\min}} \quad (14)$$

where ζ_{\min} is the minimally accepted confidence threshold, z_{\max} and z_{\min} are the maximum and minimum UAV flying altitudes respectively, and d_{uv} is the Manhattan distance between the UAV and the victim.

IV. EXPERIMENTS

This research validated the proposed UAV framework with real flight tests on the sub 2 kg quadcopter shown in Section II-A. The tests were designed under a ground SAR application context, specifically, to locate a lost person last seen around a forest/bushland area. The subsections below present the location of conducted flights, environment setup, proposed flight modes for data collection and tuned hyperparameters of the online POMDP solver.

A. Location and Environment Setup

Flight tests were conducted at the Samford Ecological Research Facility (SERF), 148 Camp Mountain Road, Samford QLD 4520, Australia. As shown in Fig. 6, the 51 hectare property contains protected Dry Sclerophyll forest and grazing zones, where the latter ones were utilised to fly the UAV.

The delimited flying area covers a mostly flat grazing zone featuring buffel grassland, a five-metre tree, and a car purposely placed as an additional obstacle. Flight tests were conducted between 19 Jul 2021 and 8 Sep 2021, in a rich range of illumination and weather conditions. Weather conditions included tests under clear and partly cloudy skies, with calm

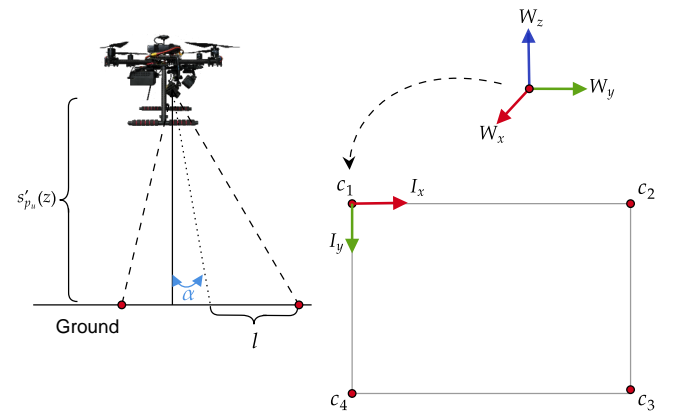


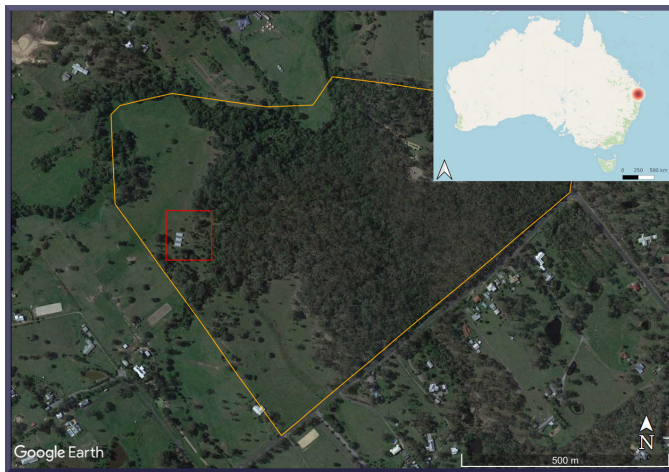
Fig. 5. Field of View (FOV) projection and footprint extent of a vision-based sensor. The camera setup on the UAV frame defines α as the pointing angle from the vertical (or pitch) and determines the coordinates of the footprint corners c .

and gusty winds from 6 km/h up to 24 km/h respectively. The range of recorded temperatures ranged from 14°C to 25°C.

This implementation employed a static adult mannequin posing as the victim to be found for safety reasons. The mannequin was placed at two predefined locations, as depicted in Fig. 7. The first location—referred from here as Location 1, or L1—is a trivial setup with the mannequin free of any nearby obstacles and entire visibility from downward-looking cameras. The second location—referred from here as Location 2, or L2—introduces a complex setup as the mannequin is placed nearby a tree which causes partial occlusion during the flight tests.

B. Flight Modes

The proposed UAV system is evaluated by collecting data using three flight modes: *mission*, *offboard*, and *hybrid*. The survey extent for these tests is delimited by a 6 m × 60 m



(a)



(b)

Fig. 6. Location of conducted flight tests at the Samford Ecological Research Facility (SERF), QLD, Australia. (a) SERF and surveyed area boundary extents (orange and red blobs respectively). (b) Aerial footage of surveyed area displaying buffel grassland and obstacles.

rectangular area, drafted in QGroundControl. Specific details of the survey pattern can be found in Fig. 15 and Tab. IV from the Appendix. A diagram illustrating the functionality of tested flight modes is shown in Fig. 8.

1) *Mission Mode*: When *mission* mode is activated, the UAV automatically follows a list of position and velocity waypoints which define the survey plan previously drafted in QGroundControl and uploaded to the autopilot before starting the flight operation. This flight mode is traditionally supported in many autopilots, and its out-of-the-box implementation serves as the planner baseline of this research. While mission



(a)



(b)

Fig. 7. Adult mannequin placed in the surveyed area as the victim to be found. (a) Trivial victim location (L1) with the mannequin fully exposed in the environment. (b) Complex victim location (L2) with the mannequin partly occluded by a five-metre tree.

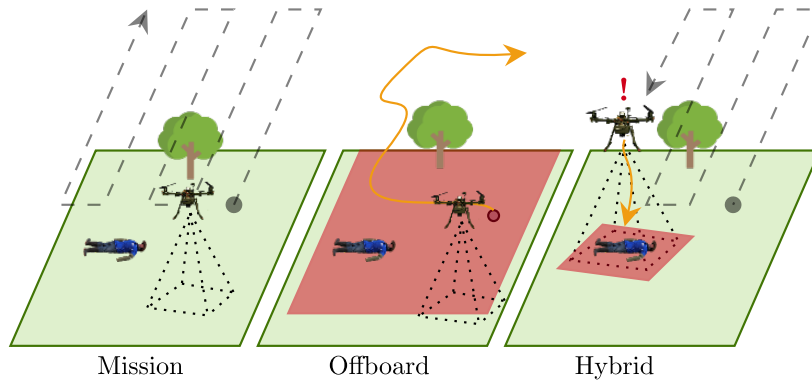


Fig. 8. Executed flight modes for exploration and object detection in outdoor environments. Mission mode is the baseline motion planner and lets the UAV survey the SAR emulated area by following a lawnmower pattern. Offboard mode runs the POMDP motion planner by populating an initial victim position belief across the entire flying area. Hybrid mode extends the functionality of mission mode by running the POMDP motion planner to inspect the area delimited by the camera's FOV.

mode is operated in the UAV, the object detector is running in parallel to record any positive detections while the UAV is navigating in the environment and completing the survey.

2) *Offboard mode*: *Offboard* mode offers autonomous navigation without a predefined survey plan of the environment. This flight mode internally executes the POMDP-based motion planner described in Section III by declaring as flight parameters the initial position waypoint where the UAV should begin the survey, and the global coordinates of the survey extents. The list of parameters can be found in Tab. V from the Appendix.

3) *Hybrid Mode*: This paper proposes the fusion of the provided capabilities between mission and offboard modes, in a flight mode denominated *hybrid*. The aim of this flight mode is to take advantage of the initial awareness and survey coverage coming from mission mode in outdoor environments with GNSS signal coverage, and the autonomous navigation capabilities of offboard mode. Instead of running the POMDP-based motion planner covering the entire extent of the surveyed area, in hybrid mode the survey extent is only limited by the extent of the camera's FOV. Once a first detection is received from the vision module, this flight mode triggers offboard mode, boots the motion planner and passes action commands to the autopilot until the POMDP solver reaches a terminal state (i.e. the UAV discards or confirms a victim). Afterwards, the UAV resumes its survey by triggering back mission mode. The process repeats itself with new detection outputs until the UAV completes the survey in mission mode.

C. POMDP solver

The navigation problem modelled as a POMDP is solved in real time through the use of the TAPIR toolkit [17]. TAPIR is coded using the C++ programming language and encapsulates the Augmented Belief Trees (ABT) solver [19] to calculate and update the motion policy online. ABT reduces computational demands by reusing past computed policies and updating the optimal policy if changes to the POMDP model are detected. Furthermore, formulated problems with ABT allow declaring continuous values for actions, states, and observations.

Once the motion server calls the motion planner after a first victim detection is received by the object detector, TAPIR is booted by calculating an offline policy for four seconds. Afterwards, the observation server retrieves an observation, updates the motion policy and takes the action that returns the highest expected reward. An idle period of 3.4 seconds is applied for the UAV to reach the desired position coordinate, and then, the process repeats itself by requesting a new observation from the observation server. The loop is broken once the detection confidence ζ exceeds a threshold. Specific parameters from the TAPIR toolkit and ABT solver are shown in Table V from the Appendix.

V. RESULTS AND DISCUSSION

The proposed UAV framework is evaluated through the performance indicators listed as follows: (1) Successful detections per flight mode; (2) Spatial distribution of recorded GNSS coordinates via heatmaps; (3) Elapsed time taken by the UAV to locate the victim per location; and (4) Scalability test using thermal imagery. Real flight demonstrations of the UAV framework can be found at https://youtu.be/U_9LbNXUwV0.

Accuracy metrics of victim detections were recorded using three variables: True Positives (TP), False Positives (FP), and False Negatives (FN). TP is defined here as the relative number of flight runs where the victim was successfully detected at the true location. FP is the relative number of flights which recorded victim locations in other areas than the true position of the victim. FN is the relative number of flights that did not detect the victim at their real location. In this context, a given flight test could report false positive detections and still detect the victim at the real location. A summary table of collected metrics is depicted in Tab. II.

Accuracy metrics of the proposed framework provided contrasting results at the tested victim locations. On flight tests with the victim placed in a trivial location (i.e., L1), the UAV achieved 100% of positive victim detections in mission mode, and 20% of those recorded GNSS coordinates of false victim locations. For tests in offboard and hybrid flight modes, the true positive rates decreased in comparison with mission

TABLE II

ACCURACY METRICS OF THE SYSTEM TO LOCATE A VICTIM AT TWO LOCATIONS (L1 AND L2). HERE, TP ARE TRUE POSITIVES, FP ARE FALSE POSITIVES, AND FN ARE FALSE NEGATIVES.

Flight Mode	Runs	TP (%)	FP (%)	FN (%)
Mission (L1)	5	100.0	20.0	0.0
Offboard (L1)	7	57.1	0.0	42.9
Hybrid (L1)	5	80.0	0.0	20.0
Mission (L2)	5	60.0	0.0	40.0
Offboard (L2)	5	80.0	0.0	20.0
Hybrid (L2)	7	71.4	28.6	28.6

mode. However, both setups achieved flight tests without any false positive readings. An illustration of the spatial distribution of recorded GNSS coordinates per flight mode for L1 is shown in Fig. 10.

Flight tests with the mannequin located in a complex location (i.e., L2) showed an overall improvement in TP rates for offboard and hybrid modes compared to the baseline planner (mission mode). For mission and offboard modes, there were no flights which reported FP detections even though the flight setup, payload, and object detector remained the same while testing the framework with the mannequin located in L1. Nevertheless, 28.6% of the flights in hybrid mode reported false positive victim locations. The rate of false positives for all flight tests were caused by limitations from the object detector by recording other objects as humans (i.e., from the car placed in the flight area). Conversely, most of the false negative records during the flight tests occurred from excessive vibration in the UAV frame caused by strong winds, as shown in Fig. 9.

A visual analysis of the spatial distribution of recorded GNSS coordinates during the flight tests is performed using heatmaps. The heatmaps, illustrated in Figs. 10 and 11, indicate a reduction in victim location uncertainty after operating the proposed UAV framework using offboard and hybrid modes.

The presented POMDP-based motion planner from [32]—scaled in this work for UAV navigation in outdoor environments—contributed on the reduction of object detection uncertainty in flight tests using offboard and hybrid modes. An example of how the UAV inspects an area to confirm whether a victim is truly located after receiving an initial detection is shown in Fig. 12. In the first time steps (Fig. 12a), a low confidence value of 19.51% is retrieved because of the few number of pixels representing the mannequin while surveying at 16 m, and partial occlusion from the nearby tree. After taking actions commands from the POMDP policy computed in the planner module, the UAV is positioned closer to the mannequin (i.e., 10 m) and with a better viewpoint of the scene, retrieving a confidence value of 90.0% (Fig. 12b). The traversed path by the UAV also suggests the capability of the UAV to adapt (or update) its motion policy while it interacts with the environment and receives new observations. Adjustments in the motion policy also occur from uncertainty sources, such as unexpected strong wind currents, oscillating GNSS signal errors, illumination changes, and false detections by the CNN model.

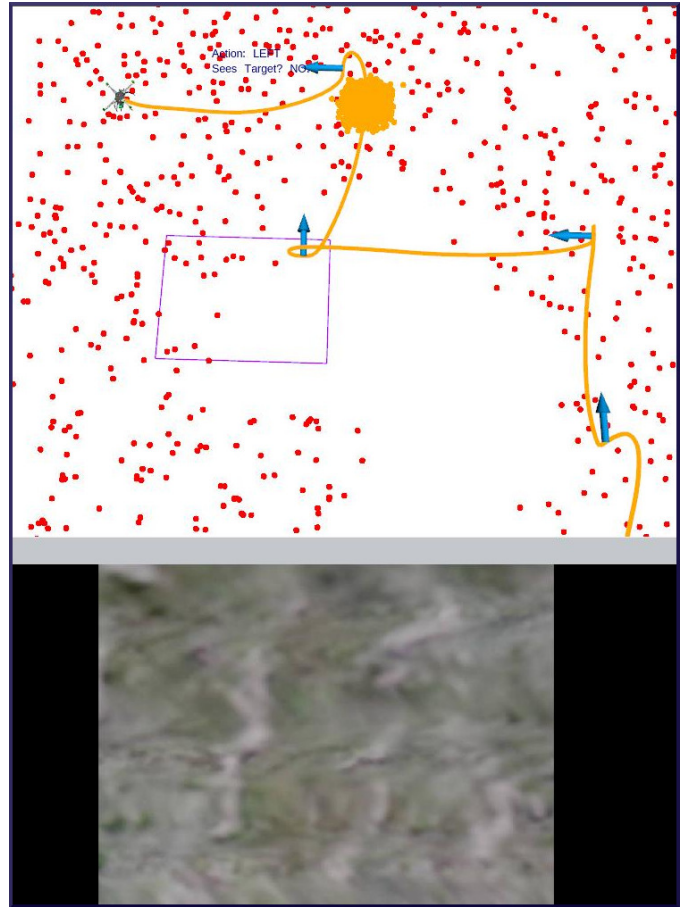


Fig. 9. Strong winds distorting RGB streaming in offboard mode. The top image illustrates the traversed path of the UAV and the probability distribution of the UAV and victim locations (orange and red points respectively). The bottom image shows the latest (distorted) streamed frame from the RGB camera.

This work also studied the speed of the proposed framework to find victims using mission, offboard, and hybrid flight modes. As presented in Tab. III, two primary insights are observed in elapsed times between victim locations and flight modes. The first observed trend defines higher values of standard deviation (especially in offboard mode) in flight tests with the victim at location L1. These values were caused by slight inconsistencies in the traversed path from the motion planner, as seen in Fig. 13. From the surveyed area, Location 1 is closer to the survey limits, whereas Location 2 is placed closed the upper centre. The motion planner lets the UAV move towards the centre of the surveyed area before exploring its corners. Despite test runs in offboard mode providing higher overall accuracy values than tests in hybrid mode, surveyed patterns could become suboptimal if a victim is located close to the boundaries of the delimited flying area.

The second observed trend is that hybrid mode recorded longer times for detecting and confirming the victim, regardless of its location in the surveyed area. This impact highly depends on the recall properties of the vision-based object detector. The number of false positive outputs while the UAV explores the



(a)



(b)

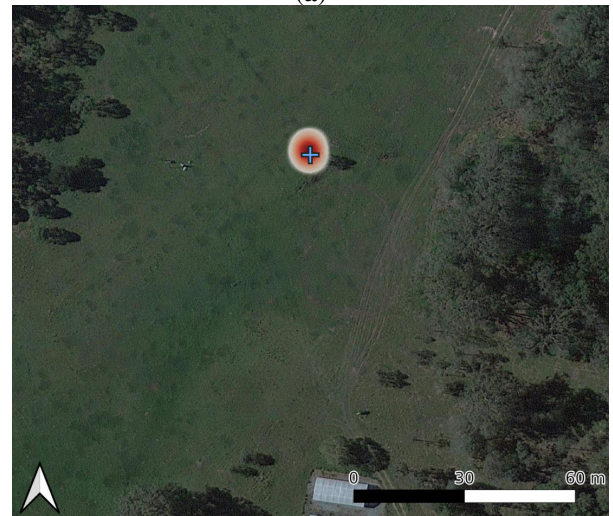


(c)

Fig. 10. Heatmaps of recorded GNSS coordinates in a trivial victim location (L1) using (a) mission, (b) offboard, and (c) hybrid flight modes.



(a)

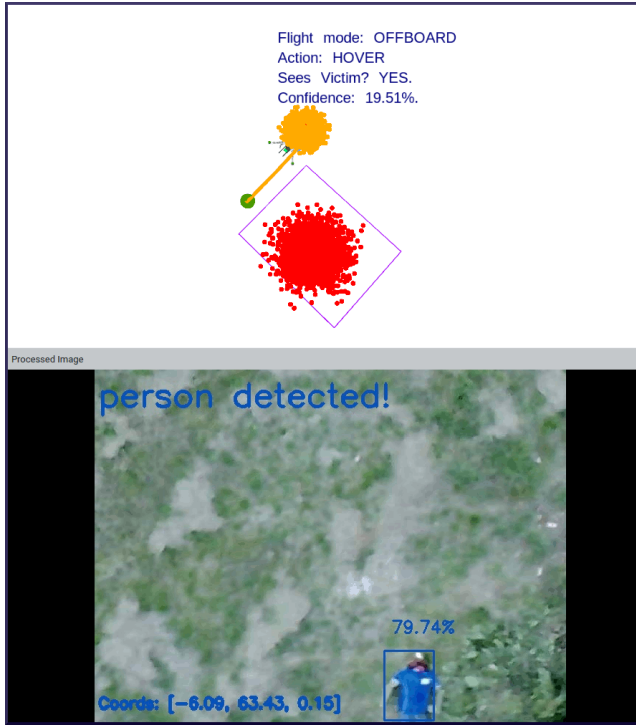


(b)

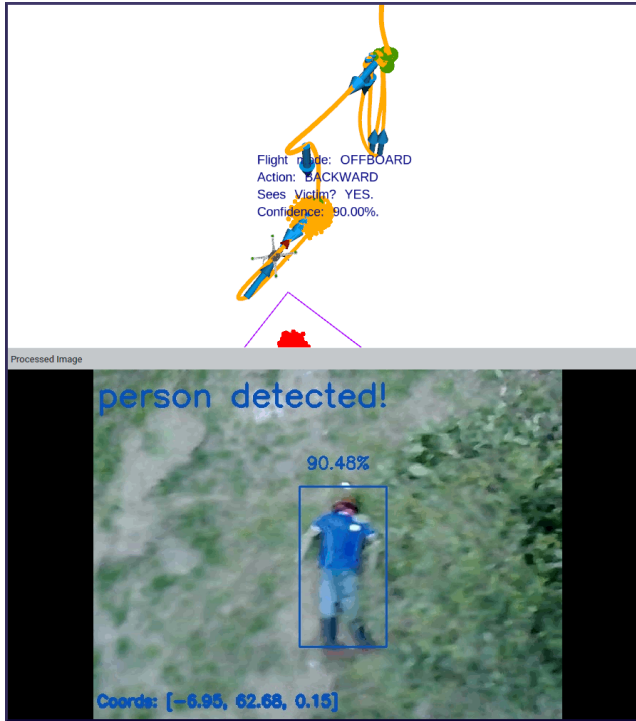


(c)

Fig. 11. Heatmaps of recorded GNSS coordinates in a complex victim location (L2) using (a) mission, (b) offboard, and (c) hybrid flight modes.



(a)



(b)

Fig. 12. Reduction of object detection uncertainty from RGB camera after executing the motion policy. (a) Initial detection of a potential victim with a confidence value of 19.51%. (b) Increased confidence value (90.0%) at subsequent detections after the UAV executes actions from the computed motion policy and gets closer to the victim.

TABLE III
ELAPSED TIME BY THE UAV TO LOCATE A VICTIM AT TWO LOCATIONS (L1 AND L2) PER FLIGHT MODE. HERE, SD STANDS FOR STANDARD DEVIATION AND SE STANDS FOR STANDARD ERROR.

Flight Mode	Mean (s)	SD (s)	SE (s)
Mission	169.67	—	—
Offboard (L1)	146.24	128.70	64.35
Hybrid (L1)	392.39	130.20	65.10
Offboard (L2)	148.30	25.58	12.79
Hybrid (L2)	289.26	76.84	34.36

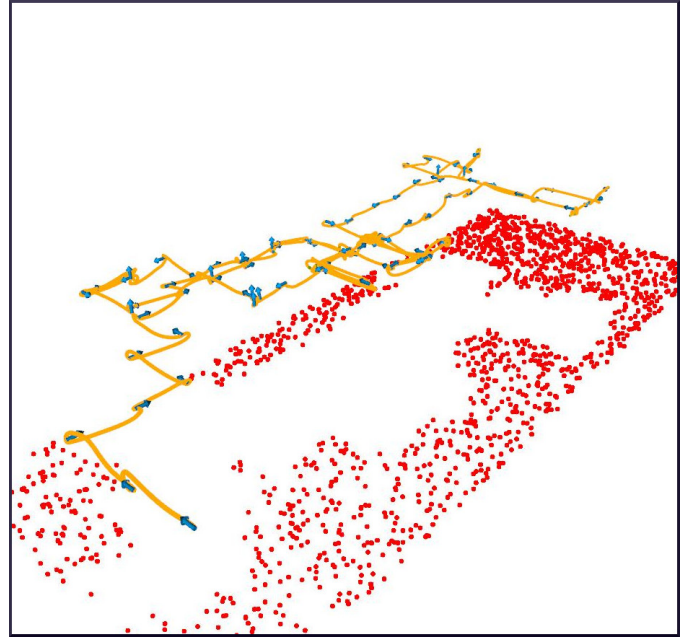


Fig. 13. Example traversed path in offboard mode while no victims are found. The UAV moves towards the centre of the surveyed area before exploring its corners.

environment defines the number of inspections, which will increase the flight time until the survey is complete. Using other object detectors tuned from airborne UAV datasets is expected to reduce the survey duration in hybrid mode, as these models should provide higher recall values than the MobileNet SSD detector implemented in this paper.

Limitations in the implemented object detector to output positive detections have conditioned the maximum altitude for the surveys, and consequently, impacted the overall speed of the system to complete the mission. Other restrictions are defined by the gimbal configuration, which scopes the footprint of the camera's FOV, and the image resolution of streamed frames. Therefore, reduced times to accomplish the victim finding mission can be accomplished by improvements in the object detector, higher image resolution and oblique camera angles that provide a wider viewpoint of the scene. Nonetheless, such improvements might involve an increment in computational power and further research should evaluate any impacts from better detectors and processing of high-resolution frames in embedded systems such as the UP².

The scalability test of the presented framework was achieved with preliminary flight tests using thermal imagery. The modular components replaced for this trial consisted of the UAV payload, the vision-based detector, and the flight plan generated by QGroundControl to match the desired overlap values previously tested with the RGB camera.

The object detector is a custom implementation of a TinyYoloV2 model architecture from Microsoft Azure Custom Vision services. A total of 5175 labelled images were used in the training process. The thermal dataset contains images of adults with a rich set of posing configurations: adults lying on the ground and waving their arms; adults standing up waving their arms; altitude and camera viewpoint variations from the UAV, ranging from 5 to 40 m, and gimbal angles of 45° and 0° from Nadir.

Thermal preview tests were conducted between 5:30 a.m. and 7:30 a.m. from 25 Aug 2021 to 23 Sep 2021, with ambient temperatures between 7°C and 14°C. An overview of the system navigating in hybrid mode to increase values of object detection confidence from streamed thermal frames can be found in Fig. 14. In this instance, the UAV starts inspecting the area covered by the FOV of the thermal camera (Fig. 14a). In subsequent time steps, the aircraft gets closer to have a better visual of the scene and confirm the presence of the victim by retrieving higher detection confidence values (Fig. 14b). A demonstrative preview video of flight operations with thermal imagery can be found at <https://youtu.be/yIPNBwNYtAo>.

The proposed UAV framework constitutes a novel approach to autonomous navigation for exploration and target finding under uncertainty. Real-world environments are full of uncertainties such as illumination conditions, strong wind currents, collision from static and dynamic obstacles, occlusion, and limitations in object detectors, which can negatively affect the success of the flight mission. While many approaches handle object detection uncertainty by fine-tuning CNN models with labelled airborne datasets [23], this research suggests augmenting the cognition power onboard UAVs from imperfect sensor and detection output observations. Moreover, augmenting autonomy capabilities in small UAVs might open more approaches to automated surveying in outdoor environments by using a swarm of UAVs which will not require permanent supervision by pilots to photo-interpret streamed camera frames to identify and locate potential victims.

Several benefits can be offered through the use of the framework in time-critical applications to SAR squads. A first assessment of the accessibility conditions of the surveyed environment, identification, localisation, quantification and conditions of victims and external hazards might be obtained rapidly thanks to real-time telemetry of processed camera frames. After completing a flight, the list of GNSS coordinates (depicted in Figs. 10 and 11) can be shared to SAR squads to coordinate better response strategies, as well as launching additional UAVs for critical deployment of medicines, food or water.

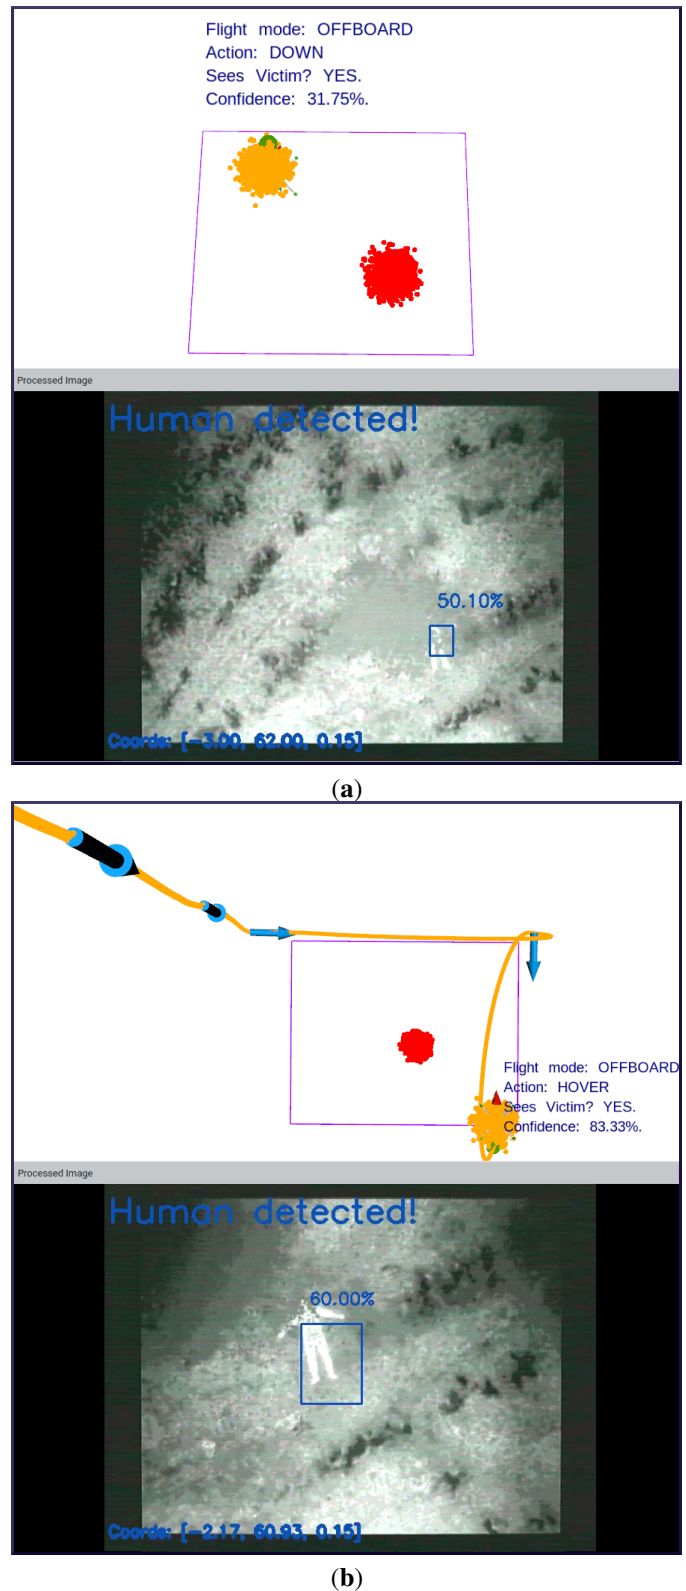


Fig. 14. Reduction of object detection uncertainty from thermal camera after executing the motion policy. (a) Initial detection of a potential victim with a confidence value of 31.75%. (b) Increased confidence value (83.33%) at subsequent detections after the UAV gets closer to the victim.

VI. CONCLUSIONS

This paper discussed a modular UAV framework for autonomous onboard navigation in outdoor environments under uncertainty. The system showed how levels of object detection uncertainty were substantially reduced by calculating a motion policy using an online POMDP solver and interacting with the environment to obtain better visual representations of potential detected targets. CNN-based computer vision inference and motion planning can be executed in resource-constrained hardware onboard small UAVs. The framework design was validated with real flight tests with a simulated SAR mission, which consisted in finding an adult mannequin in an open area and close to a tree. Collected performance indicators from three flight modes suggest that the system reduces levels of object detection uncertainty in outdoor environments whether information about the surveyed environment is available. This framework was also extended by adapting the payload with a thermal camera and converting a customised people detector from thermal imagery in the vision module.

The presented framework extends the contributions in [31], [30] by: (1) extending their tested UAV framework in GNSS-denied environments for outdoor missions with GNSS signal coverage with a novel flight mode (i.e., hybrid mode), (2) further validating preliminary results of the presented framework using real flight tests, and (3) demonstrating scalability opportunities of the modular framework design by adapting a thermal camera and custom object detector to locate victims using their heat signatures.

Future work should evaluate the performance of the UAV system using tailored object detectors from network architectures different than MobileNet SSD and high-resolution streamed frames. Further assessments with multiple victims and more complex environment configurations (e.g. slope terrain, type and number of obstacles) are encouraged. A comparison study of various online POMDP solvers, or other motion planners, could improve understanding the limits of using the ABT solver and POMDP solvers for motion planning under uncertainty.

APPENDIX

The survey design and flight plan parameters for conducted tests in *mission* and *hybrid* flight modes are shown in Fig. 15 and Tab. IV, respectively.

The set of hyper-parameters used in TAPIR and initial conditions to operate the UAV using offboard and hybrid modes are shown in Tab. V.

ACKNOWLEDGEMENTS

The authors wish to express their gratitude to Sharlene Lee-Jendili for her implementation of the thermal-based people detector used in this project. The authors acknowledge continued support from the Queensland University of Technology (QUT) through the Centre for Robotics. Special thanks to the Samford Ecological Research Facility (SERF) team (Marcus Yates and Lorrelle Allen) for their continuous assistance and equipment provided during the flight tests. The authors would also like to gratefully thank the QUT Research Engineering Facility (REF) team (Dr Dmitry Bratanov, Gavin Broadbent, Dean Gilligan) for their technical support that made possible conducting the flight tests.

TABLE IV
FLIGHT PLAN PARAMETERS BASED ON RGB CAMERA PROPERTIES.

Property	Value
UAV altitude	16 m
UAV velocity	2 m/s
Lens width	2.06 mm
Lens height	1.52 mm
Camera focal length	4.7 mm
Image resolution	640 by 480 px
Overlap	30%
Bottom right waypoint	-27.3892746°, 152.8730164°
Top right waypoint	-27.3887138°, 152.8730164°
Top left waypoint	-27.3887138°, 152.8727722°
Bottom left waypoint	-27.3892765°, 152.8727722°

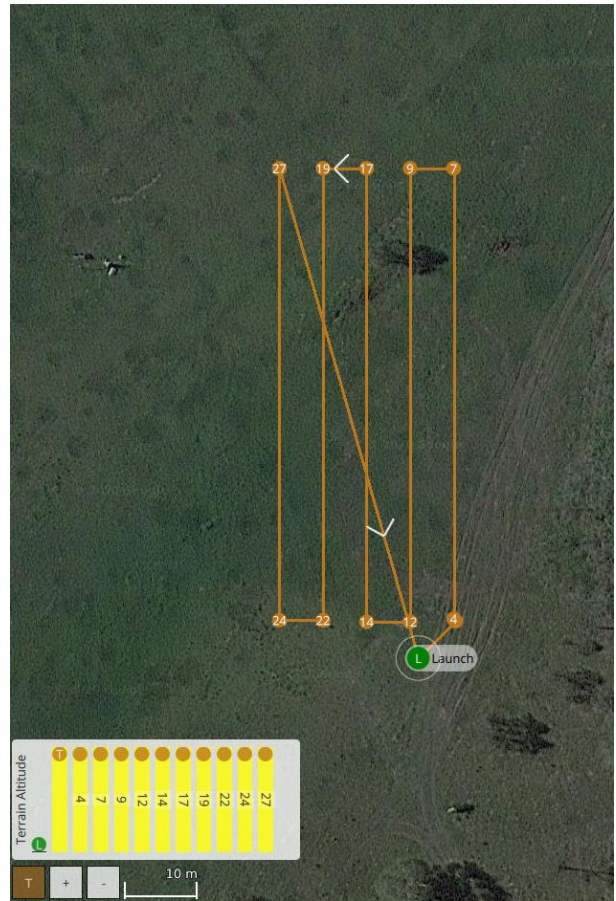


Fig. 15. Flight plan following a lawnmower pattern using the RGB camera properties from Tab. IV.

This research was funded by the Commonwealth Scientific and Industrial Research Organisation (CSIRO) through the CSIRO Data61 PhD and Top Up Scholarships (Agreement 50061686); the Australian Research Council (ARC) through the ARC Discovery Project 2018 “Navigating under the forest canopy and in the urban jungle” (grant number ARC DP180102250); and the Queensland University of Technology (QUT) through the Higher Degree Research (HDR) Tuition Fee Sponsorship. Special thanks to Hexagon through the Hexagon SmartNet RTK corrections service that enabled high accuracy surveying and positioning data using the EMLID Reach RTK receiver during the experimentation phase.

TABLE V
SET OF HYPER-PARAMETERS USED IN TAPIR AND INITIAL CONDITIONS TO OPERATE THE UAV IN OFFBOARD AND HYBRID FLIGHT MODES.

Variable	Description	Value
z_{\max}	Maximum UAV altitude	16 m
z_{\min}	Minimum UAV altitude	5.25 m
p_{u0}	Initial UAV position	(-27.3897972°, 152.8732300°, 20m)
φ_u	UAV Heading	0°
δ_z	UAV climb step	2 m
λ	Frame overlap	40%
α	Camera pitch angle	0°
ζ_{\min}	Minimum detection confidence	10%
ζ	Confidence threshold	85%
γ	Discount factor	0.95
Δt	Time step interval	4 s
t_{\max}	Maximum flying time	10 min

REFERENCES

- [1] Australian National Search and Rescue Council. Volume 2 - search and rescue operations. In *National Search and Rescue Manual*, chapter 3, pages 73–353. Canberra, Australia, 2020.
- [2] P. Bourke. Polygons and meshes. <http://paulbourke.net/geometry/polygonmesh>, 1997.
- [3] R. Z. B. Bravo, A. Leiras, and F. L. Cyrino Oliveira. The use of UAVs in humanitarian relief: An application of POMDP-based methodology for finding victims. *Production and Operations Management*, 28(2):421–440, 2019.
- [4] S. Bricknell. Missing persons: Who is at risk? Technical report, Australian Institute of Criminology, Canberra, Australia, 2017.
- [5] M. Chen, E. Frazzoli, D. Hsu, and W. S. Lee. POMDP-lite for robust robot planning under uncertainty. In *International Conference on Robotics and Automation*, pages 5427–5433, 2016.
- [6] A. Chovanová, T. Fico, L. Chovanec, and P. Hubinsk. Mathematical modelling and parameter identification of quadrotor (a survey). *Procedia Engineering*, 96:172–181, 2014.
- [7] Y. Chuanqi. Caffe implementation of Google MobileNet SSD detection network, with pretrained weights on voc0712 and map=0.727. <https://github.com/chuanqi305/MobileNet-SSD>, 2020. Accessed 2020-08-30.
- [8] M. Erdelj and E. Natalizio. UAV-assisted disaster management: Applications and open issues. In *International Conference on Computing, Networking and Communications*, pages 1–5, 2016.
- [9] M. Everingham, S. M. A. Eslami, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL visual object classes challenge: A retrospective. *International Journal of Computer Vision*, 111(1):98–136, 2015.
- [10] A. Gupta, D. Bessonov, and P. Li. A decision-theoretic approach to detection-based target search with a UAV. In *International Conference on Intelligent Robots and Systems*, pages 5304–5309, 2017.
- [11] L. Hanson and K. Namuduri. Real-world applications. In K. Namuduri, S. Chaumette, J. H. Kim, and J. P. G. Sterbenz, editors, *UAV Networks and Communications*, pages 194–213. Cambridge, UK, 2017.
- [12] M. Hassanalian, D. Rice, and A. Abdelkefi. Evolution of space drones for planetary exploration: A review. *Progress in Aerospace Sciences*, 97:61–105, 2018.
- [13] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard. Octomap: an efficient probabilistic 3D mapping framework based on octrees. *Autonomous Robots*, 34(3):189–206, 2013.
- [14] U. Ilhan, L. Gardashova, and K. Kilic. UAV using dec-POMDP model for increasing the level of security in the company. *Procedia Computer Science*, 102:458–464, 2016.
- [15] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *International Conference on Multimedia*, pages 675–678, 2014.
- [16] J. Jiménez López and M. Mulero-Pázmány. Drones for conservation in protected areas: Present and future. *Drones*, 3(1):10, 2019.
- [17] D. Klimentko, J. Song, and H. Kurniawati. TAPIR: a software toolkit for approximating and adapting POMDP solutions online. In *Australasian Conference on Robotics and Automation*, pages 1–9, 2014.
- [18] A. Koubaa, A. Allouch, M. Alajlan, Y. Javed, A. Belghith, and M. Khalgui. Micro air vehicle link (mavlink) in a nutshell: A survey. *IEEE Access*, 7:87658–87680, 2019.
- [19] H. Kurniawati and V. Yadav. An online POMDP solver for uncertainty planning in dynamic environment. In M. Inaba and P. Corke, editors, *Robotics Research. Springer Tracts in Advanced Robotics*, volume 114, pages 611–629. 2016.
- [20] S. Lee, D. Har, and D. Kum. Drone-assisted disaster management: Finding victims via infrared camera and lidar sensor fusion. In *3rd Asia-Pacific World Congress on Computer Science and Engineering*, pages 84–89, Nadi, Fiji, 2016.
- [21] S. Macdonald and A. Stevens. How to explore planets with drones. *Astronomy & Geophysics*, 59(3):3.18–3.22, 2018.
- [22] S. Mayer, L. Lischke, and P. W. Woźniak. Drones for search and rescue. In *1st International Workshop on Human-Drone Interaction*, pages 1–7, 2019.
- [23] P. Mittal, R. Singh, and A. Sharma. Deep learning-based object detection in low-altitude UAV datasets: A survey. *Image and Vision Computing*, 104:104046, 2020.
- [24] N. H. Motlagh, M. Bagaa, and T. Taleb. UAV-based iot platform: A crowd surveillance use case. *IEEE Communications Magazine*, 55(2):128–134, 2017.
- [25] G. Pajares. Overview and current status of remote sensing applications based on unmanned aerial vehicles (UAVs). *Photogrammetric Engineering & Remote Sensing*, 81(4):281–330, 2015.
- [26] M. G. Pensieri, M. Garau, and P. M. Barone. Drones as an integral part of remote sensing technologies to help missing people. *Drones*, 4(2):15, 2020.
- [27] C. Ponzonei Carvalho Chanel, A. Albore, J. T’Hooft, C. Lesire, and F. Teichteil-Königsbuch. Ample: an anytime planning and execution framework for dynamic and uncertain problems in robotics. *Autonomous Robots*, 43(1):37–62, 2019.
- [28] S. Ragi and E. K. P. Chong. UAV path planning in a dynamic environment via partially observable markov decision process. *IEEE Transactions on Aerospace and Electronic Systems*, 49(4):2397–2412, 2013.
- [29] S. Ragi and E. K. P. Chong. UAV guidance algorithms via partially observable markov decision processes. In K. Valavanis and G. Vachtsevanos, editors, *Handbook of Unmanned Aerial Vehicles*, chapter 73, pages 1775–1810. Dordrecht, 2015.
- [30] J. Sandino, F. Maire, P. Caccetta, C. Sanderson, and F. Gonzalez. Drone-based autonomous motion planning system for outdoor environments under object detection uncertainty. *Remote Sensing*, 13(21):4481, 2021.
- [31] J. Sandino, F. Vanegas, F. Gonzalez, and F. Maire. Autonomous UAV navigation for active perception of targets in uncertain and cluttered environments. In *IEEE Aerospace Conference*, pages 1–12, 2020.
- [32] J. Sandino, F. Vanegas, F. Maire, P. Caccetta, C. Sanderson, and F. Gonzalez. UAV framework for autonomous onboard navigation and people/object detection in cluttered indoor environments. *Remote Sensing*, 12(20):3386, 2020.
- [33] K. P. Valavanis and G. J. Vachtsevanos. Future of unmanned aviation. In K. P. Valavanis and G. J. Vachtsevanos, editors, *Handbook of Unmanned Aerial Vehicles*, chapter 126, pages 2993–3009. Dordrecht, 2015.
- [34] F. Vanegas, D. Campbell, M. Eich, and F. Gonzalez. UAV based target finding and tracking in GPS-denied and cluttered environments. In *International Conference on Intelligent Robots and Systems*, pages 2307–2313, 2016.
- [35] F. Vanegas and F. Gonzalez. Uncertainty based online planning for UAV target finding in cluttered and GPS-denied environments. In *IEEE Aerospace Conference*, pages 1–9, 2016.
- [36] S. Waharte and N. Trigoni. Supporting search and rescue operations with UAVs. In *International Conference on Emerging Security Technologies*, pages 142–147, 2010.