

# Receding Horizon Navigation and Target Tracking for Aerial Detection of Transient Radioactivity

Journal Title  
XX(X):1–15  
©The Author(s) 2016  
Reprints and permission:  
sagepub.co.uk/journalsPermissions.nav  
DOI: 10.1177/ToBeAssigned  
www.sagepub.com/

SAGE

Indrajeet Yadav<sup>1</sup>, Michael Sebok<sup>1</sup> and Herbert G Tanner<sup>1</sup>

## Abstract

The paper presents a receding horizon planning and control strategy for quadrotor-type micro aerial vehicle (MAV) to navigate reactively and intercept a moving target in a cluttered unknown and dynamic environment. Leveraging a lightweight short-range sensor that generates a point-cloud within a relatively narrow and short field of view (FOV), and an SSD-MobileNet based Deep neural network running on board the MAV, the proposed motion planning and control strategy produces safe and dynamically feasible MAV trajectories within the sensor FOV, which the vehicle uses to autonomously navigate, pursue, and intercept its moving target. This task is completed without reliance on a global planner or prior information about the environment or the moving target. The effectiveness of the reported planner is demonstrated numerically and experimentally in cluttered indoor and outdoor environments featuring maximum speeds of up to 4.5–5 m/s.

## Keywords

Receding Horizon Motion Planning, Target Tracking, Reactive Obstacle Avoidance, Aerial Radiation Detection

## 1 Introduction

### Overview

The work reported in this paper is motivated in part by application problems in the field of nuclear nonproliferation. In this context, there can be instances where one needs to quickly detect weak radiation sources that could be in transit through generally unknown, and sometimes GPS-denied environments. With MAVs being now a widely available and accessible technology, and with the possibility of mounting lightweight (albeit still of low efficiency) commercial off the shelf (COTS) radiation detectors on them, such a detection task is now possible using aerial means.

The approach to such problems advocated here utilizes an algorithmic pipeline that combines reactive receding horizon navigation with target tracking for MAVs, a visual-inertial (VI) state estimation algorithm, a SSD-MobileNetV2 based visual target identification and tracking, and a fast likelihood ratio-based binary decision making algorithm. Reactive planning is achieved through a new model predictive control (MPC)-type motion planner that fully incorporates the nonlinear vehicle dynamics of the MAV (cf. [Tanner and Piovesan \(2010\)](#)) and utilizes real-time sensor data in the form of a point-cloud generated by an onboard RGB-D camera, to select a (probabilistically optimal) safe path toward the target within the field of view. The reactive planner solves a multi-objective optimal control problem in real time, and thus balances detection accuracy and decision-making speed against platform flight agility limitations. The product of this optimization is a dynamically compatible minimum snap trajectory that fits along the kinematic reference path. Then a nonlinear geometric controller on the MAV tracks this trajectory in a receding horizon fashion.



**Figure 1.** Custom-built quadrotor capable of on-board planning, control, state estimation, ssd-mobilenet based target tracking motion planning and radiation measurement tracking the Jackal (ground-vehicle).

The planning and control architecture is implemented using low-cost COTS hardware, limited-range sensors, and computation devices with average capabilities. With these resources, the system has demonstrated speeds up to 4.5 m/s in cluttered indoor as well as outdoor environments. In the experimental results reported in this paper, the planner enables the MAV to intercept to static or dynamic ground targets. Such interception maneuvers have been shown to increase the likelihood of correct identification of the presence of a radioactive source on the target ([Sun and Tanner 2015](#)). In cases where the local information is insufficient to construct a feasible trajectory to the goal, the planner is designed to safely stop the vehicle. The features and capabilities of the reported reactive planning and control architecture is demonstrated in case studies of mobile

<sup>1</sup>University of Delaware, USA

### Corresponding author:

Indrajeet Yadav, University of Delaware, 126 Spencer Lab, Newark, DE-19716, USA.

Email: [indragt.yadav@udel.edu](mailto:indragt.yadav@udel.edu)

radiation detection, in which MAVs autonomously determine whether radiation sources are present on a ground vehicle moving in cluttered environments (Fig. 1).

## Background and Related Work

A quadrotor-type MAV is an inexpensive, lightweight, and agile sensor platform, suitable for many applications in areas including surveillance, aerial photography and mapping, precision agriculture, construction and defense. Although MAVs with various degrees of autonomy have been deployed in these application domains, the prevailing assumption is that the environment is at least partially known, so that a motion plan can be generated a priori and then used for steering the vehicle to a desired goal. To plan its motion in unknown environments, however, a robotic vehicle needs to construct and update a (local) environment map online; recent MAV literature addresses this problem using an onboard perception stack (Oleynikova et al. 2017; Hornung et al. 2013).

*Perception-Based Reactive Navigation* As long as vehicle dynamic constraints are respected, safe MAV navigation and aggressive maneuvering is possible by combining a reference trajectory generation process that splices together time polynomials between waypoint poses (keyframes) (Mellinger and Kumar 2011; Richter et al. 2016); a differential-geometric quadrotor point stabilization or tracking controller (Lee et al. 2010) then is utilized to transition between these keyframes. While platform-specific constraints on vehicle dynamics can be identified by experimental testing, ensuring safety during operation requires either complete prior knowledge of the environment, or some way of acquiring the missing information in real time through onboard sensing; both options involve several nontrivial and open research questions (Cadena et al. 2016).

Early work on online obstacle avoidance focused on building a new, or updating a prior, map of the environment. In this context, a continuous-time trajectory optimization using octomap (Hornung et al. 2013) is employed (Oleynikova et al. 2016), utilizing a local planner to re-generate a safe trajectory that assigns collision costs based on a (computationally expensive) Euclidean signed distance field (ESDF) map. In unknown environments, Gao and Shen (2016) construct a navigation approach utilizing online planning to produce a point-cloud map of the environment using a Velodyne 3D LiDAR, and use the map to find safe corridors through which their MAV plans its motion. To navigate to destinations outside the sensor range, a sequence of predefined waypoints is needed. Landry et al. (2016) assume knowledge of obstacle location and geometry, and locally decompose the available free space into convex polyhedra to generate safe paths for the vehicle. Variants of such approaches (Liu et al. 2017) consider the workspace represented as a 3D grid map with uniform voxels, which are used to create a convex safe flight corridor (SFC).

More recent reactive motion planning algorithms specifically for aerial vehicles include lightweight egospace-based algorithms extended to a quadrotor's configuration dynamics (Fragoso et al. 2018), or reactively sample safe trajectories in the field of view of the quadrotor, and decouple local obstacle avoidance from global guidance using a global

planner (Ryll et al. 2019). Sikang Liu et al. (2016) report on an impressive *receding horizon*-based approach to (local) planning that involves only limited onboard sensing, and which utilizes a local uniform resolution volumetric occupancy grid map and a cost map to find and navigate to safe frontier points (local goals that are closest to the global goal).

All aforementioned approaches either require some type of prior information and a global planner that generates a sequence of waypoints, or rely on (MAV payload-taxing) high-range sensors. They are particularly effective for navigating to static configurations; if, however, the MAV destination changes over time, or if the generation of a prior map is impossible, they cannot guarantee a priori the existence of a SFC, or ensure that intermediate waypoints will be within sensor range.

Yet even within a reactive navigation framework, some convergence guarantees can still be provided (Yadav and Tanner 2021a) based on appropriate extensions of the navigation function approach (Rimon and Koditschek 1992). Recent advances in such techniques have independently demonstrated success in online calculation of convergent vector fields using limited range onboard sensors (Vasilopoulos et al. 2020; Arslan and Koditschek 2019), where unknown obstacles are detected online using a deep neural network (DNN) trained on the predefined class of object geometries. Such approaches have been successfully tested on fully actuated or differential drive robots at relatively low speeds; their applicability to fast moving MAVs with under-actuated higher order dynamics remains to be demonstrated.

*Target Tracking* While the area of simultaneous localization, mapping, and moving object tracking (SLAMMOT) has been the focus of recent research efforts (Wang et al. 2007; Chojnacki and Indelman 2018), there are few implementations involving active visual-inertial target tracking approach (Chen et al. 2016) that have demonstrated impressive performance using a quadrotor. Conceptually closer to the one reported in this paper, is the aerial target tracking work of Thomas et al. (2017) who demonstrate tracking of a spherical rolling target. That approach employed a geometric technique similar to visual servoing for tracking the target combined with a receding horizon strategy that penalizes velocity and position errors. Alternatively, if the environment is known, a reactive motion planning scheme can utilize multi-objective optimization for obstacle avoidance and target tracking (Penin et al. 2018).

These approaches assume that the target is either known, detected via some type of a tag, or it can be localized using some form of visual servoing (which depends on the shape of the target). Alternative means of target detection include approaches based on convolution neural networks (CNNs), which have gained popularity in active quadrotor motion planning and target tracking with the advent of computationally efficient networks such as single shot detector (SSD) (Liu et al. 2016), you only look once (YOLO) (Redmon et al. 2016), and their variants, and seem particularly effective at high sensor speeds. For instance, DroNet architectures (Loquercio et al. 2018) are utilized in Drone Racing involving dynamically changing environments (Loquercio et al. 2020). In another related application of forest trail navigation, a multi-column DNN

is used to predict the turns in the trail and guide the quadrotor (Giusti et al. 2016), although that particular perception and control stack runs remotely on a laptop. Smolyanskiy et al. (2017) utilize a YOLO network but runs the planning and perception stack onboard the quadrotor on an NVIDIA TX1. A human gaze-driven quadrotor navigation utilizing a SSD network running on an NVIDIA TX2 features eye-tracking glasses along with a camera and an inertial measurement unit (IMU), which are combined to estimate the position of the human with respect to the quadrotor and steer the MAV (Yuan et al. 2019).

As far as the state of the art in experimental implementations is concerned, recent literature reports quadrotor navigation results at impressively high speeds (Liu et al. 2017; Mohta et al 2018; Ryll et al. 2019); yet most of the systems either involve a global planner (Ryll et al. 2019), or employ high-end and expensive sensors with extended range—e.g., Velodyne VLP-16 or Hokuyo UST-20LX LiDAR mounted on a gimbal to provide 270° FOV Liu et al. (2017); Mohta et al (2018), which can offer significantly more advanced perception capabilities compared to a 69.4° × 42.5° FOV sensor utilized in this work. In addition, the top speeds reported by Ryll et al. (2019) were achieved in open areas and with *average* flight speeds of 2.4 m/s. In the absence of a global planner, however, a planner with myopic vision cannot arbitrarily increase the speed in anticipation of an unseen obstacle. Moreover, the majority of experimental results reported relate to the case of converging to static final configurations.

Thus existing robotics literature on MAVs covers either safe and agile point-to-point navigation in unknown environments, or autonomously tracking moving target, but there is scarce, if any, work on the combined problem. In addition, the signal and data processing component, that is, what will be done with the measurement data once they have been collected, is typically an afterthought; not much effort is directed to designing the measurement process so that it serves the needs of decision-making based on these data.

## Paper Organization and Contributions

In the context of active sensing, this paper makes technical contributions to optimal sensor motion planning for active sensing. The innovations include techniques for adjustment of the optimization cost function, to manage balancing between competing mission objectives and fine-tune the aggressiveness of sensor maneuvers, and for parameterizing reference MAV trajectories in ways suitable for target tracking and receding horizon optimization. At a higher level, the approach of this paper is different in the sense that autonomy components of motion planning, safe navigation, target tracking and decision-making are integrated and co-designed to operate in concert with each other to yield significant performance improvements to fast aerial detection of weak radioactivity. This yields the first *real-time, sensor-based reactive navigation* and motion planning methodology applicable to problem instances involving either static or moving navigation goals. From an application standpoint, the planning and control architecture for aerial interception and target tracking is thus unique because:

1. it is applicable to cases with both static and moving goal configurations;
2. it generates trajectories in a new way by efficiently solving a multi-objective optimal control problem;
3. it is experimentally tested indoors and outdoors with MAV speeds of up to 4.5–5 m/s; and
4. it is demonstrated to be capable of solving a challenging radiation detection problem.

## 2 Overview of the Architecture

The block diagram of Fig. 2 shows the flow of information within the motion planning, state estimation, and platform control architecture. Starting with the RGB-D camera, a point-cloud is produced and then used to frame the obstacle-free portion of the workspace. On board state estimation is performed through a visual-inertial MSCKF navigation stack (Open-VINS (Geneva et al. 2020)). Given the free workspace and the state estimates, the receding horizon planner (i) generates a candidate minimum snap trajectory from the current position of the robot to each of the free points in the field of view based on Pontryagin’s minimum principle; (ii) selects the point in the FOV closest to the target (henceforth referred to as the *intermediate point*); (iii) assigns a cost to each trajectory as a weighted sum of the trajectory’s proximity to obstacles and the proximity of the trajectory’s end point to the intermediate point; and finally, (iv) selects the trajectory with the least such cost as the *reference trajectory* for the MAV to follow. The end point of this selected minimum-snap reference trajectory is referred to as the *local goal*. An initial segment of the reference trajectory that ends at the local goal is presented to a differential-geometric tracking controller, which initiates steering of the MAV to the designated local goal point within the FOV. In a typical receding horizon fashion, before the end of that initial segment is reached, the MAV uses an updated set of point-cloud data to generate a new reference trajectory, starting from a point on the current trajectory associated with the end of some designated control time horizon. Once in the neighborhood of this point, the MAV transitions smoothly to tracking of the new reference trajectory. This constitutes a *planning cycle* of duration  $\delta t$ . The planning, and concurrent trajectory tracking, cycle repeats until the final destination is reached—when the latter is static.

If the destination is a moving target, an SSD MobileNet-v2 based target tracker calculates the relative positions between the target and the MAV, which the planner utilizes to generate intercepting trajectories.

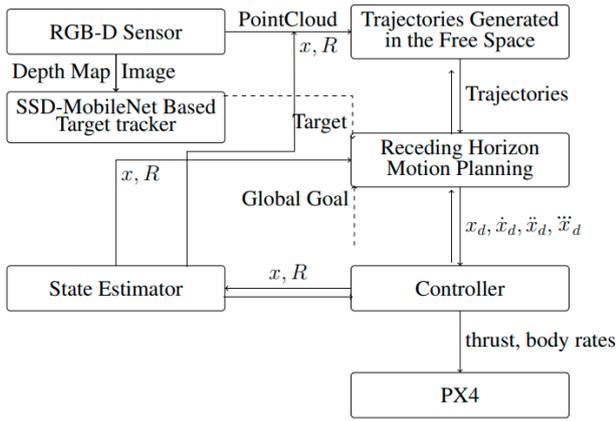
## 3 Radiation Detection Preliminaries

### Application Case Study Background

The robotics application that motivated this work was detecting\* weak mobile sources of radiation using aerial mobile sensors. One possibility for realizing an inexpensive

---

\*Note that the detection here refers to determining whether or not a *known* target carries a certain amount of radioactive material; it should not be confused with object detection in computer vision literature.



**Figure 2.** Block diagram of the motion planning and control architecture. The arrows indicate the direction of information flow. The dashed lines indicate that the input to the planner is either a user specified static goal or relative position of the target obtained from SSD-MobileNet.

fully autonomous aerial radiation detector is to utilize a quadrotor MAV, mount a COTS radiation counter, and allow it to maneuver itself in ways that allow it to determine as quickly as possible whether or not its given moving target is radioactive or not. While a network of static detectors can also be deployed to the detection, due to the strong (inverse square law) dependence of the process intensity on the distance between sensor and source (Nemzek et al. 2004), the static network will experience a dramatic decrease of its signal-to-noise-ratio (SNR) that will significantly reduce its detection accuracy and speed.

Detecting a mobile radioactive source based on a series of discrete events associated with the arrival of particles or gamma rays (indistinguishable from each other), can be formulated as an inhomogeneous likelihood ratio test (LRT) within the Neyman-Pearson framework (Pahlajani et al. 2014). In this framework, *analytical* (Chernoff) bounds on error probabilities can be derived in terms of relative sensor-source distance and used as proxies for the otherwise intractable true detection error probabilities (Pahlajani et al. 2014), to facilitate the design of optimal control laws that maximize detection performance in terms of speed and accuracy (Sun and Tanner 2015; Yadav and Tanner 2019).

### Neyman-Pearson based Radiation Detection

A sensor measurement event—for the case of a Gieger-Muller (GM) detector, for instance—is associated with the discrete random process where a gamma-ray (from the source or background) hits the sensitive area of the detector and records a “count.” This is mathematically modeled in a probabilistic setup (Pahlajani et al. 2014; Sun and Tanner 2015) as follows. Assume that there is a counting process  $N_t$  for  $t \in [0, T]$ , on a measurable space  $(\Omega, \mathcal{F})$ . In this context,  $N_t$  essentially represents the number of times gamma-rays have hit the radiation counter located at position  $x \in \mathbb{R}^3$ , up to (and including) time  $t \in [0, T]$ . This counting process is assumed to be a Poisson process (Snyder 1975). The sensor position  $x$  is known over the whole interval  $[0, T]$ .

The detector can make two competing hypotheses,  $H_0$  and  $H_1$ , each expressing an opinion as to whether its cumulative

count rate can be attributed to the presence of a radiation source of intensity  $a$  (in counts per unit time, i.e., counts per second (CPS)), located at a possibly time-varying position  $y \in \mathbb{R}^3$  which will be referred to as the *target*. Hypothesis  $H_0$  asserts that the target is benign, while hypothesis  $H_1$  claims that the target carries a source of intensity. The two hypotheses  $H_0$  and  $H_1$  correspond, respectively, to two distinct probability measures  $\mathbb{P}_0$  and  $\mathbb{P}_1$  on  $(\Omega, \mathcal{F})$ . With respect to measure  $\mathbb{P}_0$ , the process  $N_t$  is a Poisson process with intensity  $\beta(t)$ , i.e., the intensity of naturally occurring background radiation; with respect to  $\mathbb{P}_1$ , however, the same process is Poisson with intensity  $\beta(t) + \nu(t)$ , where  $\nu(t)$  represents the intensity of the source (whenever present) as perceived by the sensor at time  $t$ . The functions  $\beta(t)$  and  $\nu(t)$  defined on  $[0, T]$  are assumed to be bounded, continuous and strictly positive (Pahlajani et al. 2014).

Function  $\nu(t)$  should implicitly reflect the inverse square law dependence of the source intensity perceived by the sensor on the distance between the sensor and the source (Nemzek et al. 2004). If  $\chi$  denotes the sensor’s cross-section coefficient, one possibility for expressing the detector’s perceived intensity is in the form:

$$\nu(t) = \frac{\chi a}{2\chi + \|y(t) - x(t)\|^2} . \quad (1)$$

A test for choosing between  $H_0$  and  $H_1$  faces the risk of two types of errors. One of them is a *false alarm*, which occurs when the outcome of the test is in favor of  $H_1$  while  $H_0$  is instead the correct hypothesis; the other is a *missed detection* in which one sides with  $H_0$  when in fact  $H_1$  is true. In this setting, the optimal test for deciding between  $H_0$  and  $H_1$  is an LRT obtained as follows (Pahlajani et al. 2014). Let  $\tau_n$  for  $n \geq 1$  denote the  $n^{\text{th}}$  jump time of  $N_t$  (jumps occur when the sensor registers a count), and with the convention that  $\prod_{n=1}^0 (\cdot) = 1$ , let

$$L_T = \exp \left( - \int_0^T \nu(s) ds \right) \prod_{n=1}^{N_t} 1 + \frac{\nu(\tau_n)}{\beta(\tau_n)} . \quad (2)$$

be the *likelihood ratio*. Then for a specific fixed *threshold*  $\gamma > 0$ , the test

$$L_T \underset{H_0}{\overset{H_1}{\gtrless}} \gamma \quad (3)$$

is optimal in the (Neyman-Pearson) sense. This means that (3) minimizes the probability of a missed detection under a given upper bound constraint on the probability of false alarm. With  $\mu(t) \triangleq 1 + \frac{\nu(t)}{\beta(t)}$ , constants  $p \in (0, 1)$  and  $\eta \triangleq \log \gamma$ , and the quantity

$$\Lambda(p) \triangleq \int_0^T [\mu(s)^p - p\mu(s) + p - 1] \beta(s) ds , \quad (4)$$

one can now extract analytical expressions for Chernoff bounds on the probabilities of false alarm  $P_F$  and missed detection  $P_M$  (Pahlajani et al. 2014).

If an upper limit  $\alpha > 0$  is set on the bound on probability of false alarm, then there exists a unique solution  $p^* \in [0, 1]$  for which the tightest bound on the probability of missed detection can be obtained. The exponent in the bound on the probability of false alarm and missed detection, respectively,

is (Pahlajani et al. 2014)

$$\mathcal{E}_{\mathcal{F}} = \int_0^T [p^* \mu^{p^*} \log \mu - \mu^{p^*} + 1] \beta \, ds = -\log \alpha \quad (5)$$

$$\mathcal{E}_{\mathcal{M}} = \log \alpha + \Lambda'(p^*) \quad (6)$$

where the dependency of  $\mu$  and  $\beta$  on time is suppressed for clarity, and derivative  $\Lambda'(p) = \frac{\partial \Lambda}{\partial p}$  is expressed as

$$\Lambda'(p) = \int_0^T [\mu^p \log \mu - \mu + 1] \beta \, ds \quad (7)$$

For the optimal  $p^*$ , the  $\Lambda(p^*)$  and detection threshold  $\gamma$  are related in the form  $\gamma = \exp(\Lambda(p^*))$ .

### Problem Statement

Suppose now that the distance between target and sensor,  $\|y - x\|$ , is regulated by a control input  $u$  via dynamics of the form  $\dot{x} = f(x, u)$ ; then  $\nu$ , and consequently  $\mu$ , depend implicitly on  $u$  and an optimal control problem can be formulated as follows:

**Problem 1.** *Given current target and sensor locations  $y$  and  $x$ , respectively, find a control input  $u$  that optimizes  $\Lambda'(p^*)$  for a given upper limit  $\alpha$  on the bound on the probability of false alarm.*

Irrespective of whether  $\|y - x\|$  is deterministic or stochastic, it can be shown that the optimal sensor management strategy  $u$  for sensors is to close the distance between source and sensor as quickly as possible (Sun and Tanner 2015; Yadav and Tanner 2017).

## 4 Quadrotor Dynamics and Control

The MAV is modeled as a rigid body moving in SE(3). Let  $m$  and  $\mathbf{J}$  denote its mass and moment of inertia, respectively, and denote  $\mathbf{x} = (x, y, z)^\top$  and  $\mathbf{v} = \dot{\mathbf{x}}$  its 3D position and velocity in the inertial frame. Let  $\mathbf{R}$  be the rotation matrix from the body-fixed frame to the inertial frame, and  $\Omega$  be the MAV's angular velocity in the body-fixed frame. We denote  $\hat{\cdot}$  the skew symmetry operation, and write the gravitational acceleration vector and the unit vector along the inertial vertical as  $\mathbf{g} = (0, 0, g)^\top$ , and  $\mathbf{e}_3 = (0, 0, 1)^\top$ , respectively. The control inputs to the MAV are the (scalar) thrust  $f$  and the (vector) moment  $\mathbf{M}$ , both of which are considered bounded. The former is bounded in the interval  $(f_{\min}, f_{\max})$ ; the constraint on the latter, is assumed to translate to a bound on angular velocity  $\|\Omega\| \leq \Omega_{\max}$ . With these definitions and assumptions in place, the MAV dynamics is described as

$$\dot{\mathbf{x}} = \mathbf{v}, \quad m \dot{\mathbf{v}} = f \mathbf{R} \mathbf{e}_3 - m \mathbf{g} \quad (8a)$$

$$\dot{\mathbf{R}} = \mathbf{R} \hat{\Omega}, \quad \mathbf{J} \dot{\Omega} + \Omega \times \mathbf{J} \Omega = \mathbf{M} \quad (8b)$$

The control law for this system is designed based on a (differential) geometric method (Lee et al. 2010). To examine the particular control design employed here in more detail, consider a smooth position reference trajectory  $\mathbf{x}_d(t) \in \mathbb{R}^3$ . Based on this one can construct (Lee et al. 2010) a desired rotation matrix  $\mathbf{R}_d$  and angular velocity  $\Omega_d$  that are consistent with the reference trajectory  $\mathbf{x}_d(t)$ . The desired yaw angle  $\psi$  of the MAV is used to construct the desired MAV direction vector in the form  $\mathbf{b}_{1d} = (\cos \psi, \sin \psi, 0)^\top$ .

The tracking errors in position  $\mathbf{e}_x$ , velocity  $\mathbf{e}_v$ , orientation  $\mathbf{e}_R$ , and angular rate  $\mathbf{e}_\Omega$ , are expressed as

$$\mathbf{e}_x = \mathbf{x} - \mathbf{x}_d \quad \mathbf{e}_v = \dot{\mathbf{x}} - \dot{\mathbf{x}}_d \quad (9a)$$

$$\mathbf{e}_R = \frac{1}{2}(\mathbf{R}_d^\top \mathbf{R} - \mathbf{R}^\top \mathbf{R}_d) \quad \mathbf{e}_\Omega = \Omega - \mathbf{R}^\top \mathbf{R}_d \Omega_d \quad (9b)$$

Picking positive control gains  $k_x$ ,  $k_v$ ,  $k_R$  and  $k_\Omega$ , the control inputs can now be constructed as

$$f = -k_x \mathbf{e}_x - k_v \mathbf{e}_v + m \mathbf{g} + m \ddot{\mathbf{x}}_d, \quad (10a)$$

$$\mathbf{M} = k_R \mathbf{e}_R + k_\Omega \mathbf{e}_\Omega + \Omega \times \mathbf{J} \Omega, \quad (10b)$$

allowing one to achieve exponentially stable tracking behavior for initial attitude error less than  $\pi/2$  (cf. Lee et al. (2010)).

## 5 Reactive Receding Horizon Planning

### Representing the Free Space

Let  $\mathcal{V} \in \mathbb{R}^3$  denote the visible space within the FOV of the RGB-D sensor. This FOV area contained in  $\mathcal{V}$  is assumed to take the shape of a solid pyramid sector with its apex attached to the base frame of the sensor, with the depth direction of the sensor being aligned with the local  $x$  (heading) frame axis of the MAV. The base of the FOV pyramid is outlined by the maximum range  $R_{\max}$  of the sensor, while the side boundaries (polar and azimuth angles of the sector) are determined by the maximal viewing angles in the FOV of the sensor along the local  $y$  and  $z$  directions. Denote  $\phi_y$  and  $\phi_z$  those maximum viewing angles of the sensor at its apex along the local  $y$  and  $z$  directions, respectively. The motion planning process also takes as input a user-specified minimum range for perception,  $R_{\min}$ . Assuming now that the apex of the FOV pyramid is at local frame coordinates  $(0, 0, 0)^\top$ , any point within the FOV of the sensor can be expressed in the polar coordinates as  $(r, \theta, \varphi)^\top$  where  $R_{\min} \leq r \leq R_{\max}$ ,  $-\phi_y \leq \theta \leq \phi_y$  and  $-\phi_z \leq \varphi \leq \phi_z$ . By selecting a resolution  $\delta r$  on the range and  $\delta \theta$  on the viewing angles in both direction, the field of view of the sensor can be discretized and represented as an ensemble of points. Each point in this ensemble represents a potential desired next location for the MAV, should a feasible path to this location exist.

### Reference Trajectory Generation

The center of gravity (COG) of the MAV is designated as the initial point of any candidate reference trajectory. Using the (constant) transformation between COG and sensor frame (at the FOV apex), the coordinates of all the points and rays can be expressed relative to body-fixed COG frame of the vehicle.

Given the ensemble of points within the field of view, a minimum snap trajectory to each of these point is constructed. Note that the dynamics of a quadrotor MAV enjoys differential flatness properties (Mellinger and Kumar 2011), which ensure that all inputs and states can be written as functions of four (flat) outputs and their derivatives. The flat outputs are the Cartesian MAV position and its yaw angle, which can be brought together in a vector  $(\mathbf{x}, \psi)$ . Note that yaw  $\psi$  is decoupled from  $\mathbf{x}$  and can be steered independently.

The derivatives of  $\mathbf{x}$  that are of interest are velocity  $\mathbf{v} = \dot{\mathbf{x}}$ , acceleration  $\mathbf{a} = \ddot{\mathbf{x}}$ , and jerk  $\mathbf{j} = \dddot{\mathbf{x}}$ . In the flat

output space, the MAV position dynamics can therefore be assumed to match those of a quadruple integrator with state vector  $\mathbf{p}(t) = (\mathbf{x}^\top(t), \mathbf{v}^\top(t), \mathbf{a}^\top(t), \mathbf{j}^\top(t))^\top$ . The objective of the receding horizon planner now is to generate smooth trajectories  $\mathbf{p}_{[N]}$  and  $\mathbf{p}_{[N+1]}$ , defined over the course of two consecutive planning cycles each of duration  $\delta t$  and indexed  $N$  and  $N + 1$ ,

$$\begin{aligned} \mathbf{p}_{[N]}(t, t + \delta t) &= (\mathbf{x}_{d[N]}^\top, \mathbf{v}_{d[N]}^\top, \mathbf{a}_{d[N]}^\top, \mathbf{j}_{d[N]}^\top)^\top \\ \mathbf{p}_{[N+1]}(t + \delta t, t + 2\delta t) &= (\mathbf{x}_{d[N+1]}^\top, \mathbf{v}_{d[N+1]}^\top, \mathbf{a}_{d[N+1]}^\top, \mathbf{j}_{d[N+1]}^\top)^\top, \end{aligned}$$

which always remain within  $\mathcal{F}$  and satisfy the boundary condition  $\mathbf{p}_{[N]}(t + \delta t) = \mathbf{p}_{[N+1]}(t + \delta t)$ , while being dynamically feasible, i.e.,  $f_{\min} \leq f \leq f_{\max}$  and  $\|\Omega\| \leq \Omega_{\max}$ .

Let  $\mathbb{T}$  be the free trajectory terminal time (which will henceforth be referred to as the *planning horizon*) and denote  $\mathbf{p}_0 = (\mathbf{x}_0^\top, \mathbf{v}_0^\top, \mathbf{a}_0^\top, \mathbf{j}_0^\top)^\top$ ,  $\mathbf{p}_\mathbb{T} = (\mathbf{x}_\mathbb{T}^\top, \mathbf{v}_\mathbb{T}^\top, \mathbf{a}_\mathbb{T}^\top, \mathbf{j}_\mathbb{T}^\top)^\top$  the trajectory boundary conditions at  $t = 0$  and  $t = \mathbb{T}$ , respectively. Then let  $\mathbf{u}(t) = \frac{d^4 \mathbf{x}(t)}{dt^4}$  denote snap, and treat it as the control input for the flat space dynamics. For a *relative weight parameter*  $k \in \mathbb{R}_+$ , the free terminal time optimal control problem the solution of which are the desired reference trajectories, is formulated as:

$$\begin{cases} \min \int_0^\mathbb{T} (k + \frac{1}{2} \|\mathbf{u}(t)\|^2) dt & \text{subject to} \\ \dot{\mathbf{x}}(t) = \mathbf{v}(t), \quad \dot{\mathbf{v}}(t) = \mathbf{a}(t), \quad \dot{\mathbf{a}}(t) = \mathbf{j}(t), \quad \dot{\mathbf{j}}(t) = \mathbf{u}(t) \\ \mathbf{p}(0) = \mathbf{p}_0, \quad \mathbf{p}(\mathbb{T}) = \mathbf{p}_\mathbb{T} . \end{cases} \quad (11)$$

The cost function of the optimal control problem (11) is a linear combination of two performance factors: (i) the incremental cost associated with the terminal time (time optimality), captured by the constant integrand term; and (ii) the incremental cost that penalizes the cumulative magnitude of snap  $\mathbf{u}$  along the trajectory. By tuning  $k$ , one can adjust how aggressive the reference trajectory obtained is. Smaller values for  $k$  place a smaller penalty on tracking time and therefore result to slower motion. GAZEBO simulation data at various  $k$  and speeds had been collected to fit a relation between the maximum speed along a candidate trajectory and parameter  $k$ . This relation comes out to be an interpolated curve of the form  $v_{\text{candidate}} = \alpha k^{1/\beta}$  and has been found to work well in practical scenarios to guide the selection of the cost weight  $k$  based on the maximum robot speed afforded for the mission at hand. The cost weight  $k$  is particular to the candidate trajectory and varies for the different candidate trajectories in the ensemble, since the maximum speed afforded along a candidate trajectory  $v_{\text{candidate}}$  itself varies in proportion to the ratio of the candidate trajectory's length to the length of the longest candidate trajectory. This leads to smaller trajectories having lower top speeds, making the MAV automatically slow down in the vicinity of dense obstacles.

Denoting  $t$  the time elapsed since the start of the whole planned maneuver,  $d$  the vehicle's remaining distance to its goal,  $r$  being the distance of the point in the ensemble from the camera focal point,  $v_{\max}$  the desired maximum MAV speed, and  $k_t$  and  $k_d$  being positive tuning parameters, we define the velocity profile along the reference trajectory

through the function

$$v_{\text{candidate}} = \text{erf}(k_t t) \text{erf}(k_d d) \left( \frac{r}{R_{\max}} \right) v_{\max} , \quad (12)$$

which now specifies  $v_{\text{candidate}}$  and allows the direct determination of the weighing factor  $k$  in (11). Compared to alternative trapezoid velocity profiles (Mellinger and Kumar 2011), (12) produces a velocity profile for the entire remaining quadrotor trajectory that is also effective in scenarios involving moving target interception, in which the vehicle needs to adjust its speed to match that of its target while at the vicinity of the latter. A more detailed discussion of the effect of this velocity profile on tracking dynamic targets follows in the section on target tracking.

To solve (11) one can utilize Pontryagin's Minimum principle (Athans and Falb 1966, Chapter 6) as follows. Let  $\lambda_x, \lambda_v, \lambda_a$  and  $\lambda_j$  be the *costate* vectors. Each such vector has three components, one for each spatial direction  $x, y$ , and  $z$ . Let  $\langle \cdot, \cdot \rangle$  denote the standard inner product between of two vectors, and express the Hamiltonian  $H$  for this problem as

$$H = k + \frac{1}{2} \|\mathbf{u}\|^2 + \langle \lambda_x, \mathbf{v} \rangle + \langle \lambda_v, \mathbf{a} \rangle + \langle \lambda_a, \mathbf{j} \rangle + \langle \lambda_j, \mathbf{u} \rangle . \quad (13)$$

The optimal solution is derived from the condition  $H(\mathbf{x}^*, \mathbf{u}^*, \lambda^*, t) \leq H(\mathbf{x}^*, \mathbf{u}, \lambda^*, t)$ , which, since the Hamiltonian is quadratic in the control inputs, leads to

$$\frac{\partial H}{\partial \mathbf{u}} = 0 \implies \begin{cases} u_x = -\lambda_{jx} \\ u_y = -\lambda_{jy} \\ u_z = -\lambda_{jz} . \end{cases} \quad (14)$$

The costate dynamics now is

$$\begin{aligned} \dot{\lambda}_x &= -\frac{\partial H}{\partial \mathbf{x}} = 0 & \dot{\lambda}_v &= -\frac{\partial H}{\partial \mathbf{v}} = -\lambda_x \\ \dot{\lambda}_a &= -\frac{\partial H}{\partial \mathbf{a}} = -\lambda_v & \dot{\lambda}_j &= -\frac{\partial H}{\partial \mathbf{j}} = -\lambda_a . \end{aligned}$$

Fortunately, this problem affords independent solution along each direction ( $x, y$  and  $z$  being independent flat outputs). Each costate vector consists of three costates corresponding to the  $x, y$  and  $z$  direction, respectively, i.e.,  $\lambda_x = \{\lambda_{px}, \lambda_{py}, \lambda_{pz}\}$ . If now  $\mathbf{c} = (c_{x0}, \dots, c_{x7})$  denotes an  $8 \times 1$  vector of constant coefficients for the  $x$  direction, the costate dynamics along with the system dynamics (11) in  $x$  direction can be resolved as follows:

$$\begin{aligned} \dot{\lambda}_{px} &= 0 \implies \lambda_{px} = c_{x7} \\ \dot{\lambda}_{vx} &= -c_{x7} \implies \lambda_{vx} = -c_{x7}t + c_{x6} . \end{aligned}$$

Recursively putting  $\lambda_{vx}$  into the equation of  $\lambda_{ax}$  and then  $\lambda_{ax}$  into the equation of  $\lambda_{jx}$  and using (14) along with the system dynamics in (11), yields the optimal solution along

in  $x$  direction:

$$\left. \begin{aligned} \lambda_{px} &= c_{x7} \\ \lambda_{vx} &= -c_{x7}t + c_{x6} \\ \lambda_{ax} &= c_{x7}\frac{t^2}{2} - c_{x6}t + c_{x5} \\ u_x &= c_{x7}\frac{t^3}{6} - c_{x6}\frac{t^2}{2} + c_{x5}t - c_{x4} \\ j_x &= c_{x7}\frac{t^4}{24} - c_{x6}\frac{t^3}{6} + c_{x5}\frac{t^2}{2} - c_{x4}t + c_{x3} \\ a_x &= c_{x7}\frac{t^5}{120} - c_{x6}\frac{t^4}{24} + c_{x5}\frac{t^3}{6} - c_{x4}\frac{t^2}{2} + c_{x3}t \\ &\quad + c_{x2} \\ v_x &= c_{x7}\frac{t^6}{720} - c_{x6}\frac{t^5}{120} + c_{x5}\frac{t^4}{24} - c_{x4}\frac{t^3}{6} + c_{x3}\frac{t^2}{2} \\ &\quad + c_{x2}t + c_{x1} \\ x_x &= c_{x7}\frac{t^7}{5040} - c_{x6}\frac{t^6}{720} + c_{x5}\frac{t^5}{120} - c_{x4}\frac{t^4}{24} \\ &\quad + c_{x3}\frac{t^3}{6} + c_{x2}\frac{t^2}{2} + c_{x1}t + c_{x0} . \end{aligned} \right\} (15)$$

The other two directions feature identical polynomials.

The optimal trajectory, therefore, is a 7<sup>th</sup> order polynomial in time. Enforcing the boundary conditions at  $t = 0$  gives  $c_{x0} = p_{x0}$ ,  $c_{x1} = v_{x0}$ ,  $c_{x2} = a_{x0}$  and  $c_{x3} = j_{x0}$ , while the remaining coefficients,  $c_{x4}$  through  $c_{x7}$ , are obtained from the boundary conditions at  $t = T$ :

$$\begin{bmatrix} c_{x7} \\ c_{x6} \\ c_{x5} \\ c_{x4} \end{bmatrix} = \begin{bmatrix} \frac{T^7}{5040} & -\frac{T^6}{720} & \frac{T^5}{120} & -\frac{T^4}{24} \\ \frac{T^6}{720} & -\frac{T^5}{120} & \frac{T^4}{24} & -\frac{T^3}{6} \\ \frac{T^5}{120} & -\frac{T^4}{24} & \frac{T^3}{6} & -\frac{T^2}{2} \\ \frac{T^4}{24} & -\frac{T^3}{6} & \frac{T^2}{2} & -T \end{bmatrix}^{-1} \begin{bmatrix} p_{xT} - (p_{x0} + v_{x0}T + \frac{1}{2}a_{x0}T^2) + \frac{1}{6}j_{x0}T^3 \\ v_{xT} - (v_{x0} + a_{x0}T + \frac{1}{2}j_{x0}T^2) \\ a_{xT} - (a_{x0} + j_{x0}T) \\ j_{xT} - j_{x0} \end{bmatrix}$$

expressing the optimal trajectory coefficients as a function of the (yet unknown) free terminal time  $T$ .

The free terminal time  $T$  can be determined as follows. With the  $c_x$  trajectory coefficients are explicitly expressed in terms of  $T$ , one substitutes and writes the control input at the terminal time as

$$u_{xT} = \frac{840(p_{x0} - p_{xT})}{T^4} + \frac{360v_{x0}}{T^3} + \frac{60a_{x0}}{T^2} + \frac{4j_{x0}}{T}$$

(control inputs at  $T$  in  $y$  and  $z$  are similarly obtained). Velocity, acceleration and jerk at time  $T$  are all set to zero, while the transversality condition (Athans and Falb 1966) at  $T$  requires  $H(T) = 0$ . Combining these conditions, with (14) results to

$$k + \frac{1}{2}\|u_T\|^2 + \langle \lambda_T, u_T \rangle = 0 \implies \|u_T\|^2 = 2k .$$

This is essentially an 8<sup>th</sup> degree polynomial equation which can now produce  $T$ . Indeed, if for the sake of brevity we set

$$l_1 \triangleq 840(p_{x0} - p_{xT}) \quad m_1 \triangleq 360v_{x0} \quad n_1 \triangleq 60a_{x0} \quad o_1 \triangleq 4j_{x0}$$

(work similarly for  $y$  and  $z$  using indices 2 and 3, respectively), and then substitute back we obtain the polynomial equation

$$\begin{aligned} & -2kT^8 + \sum_{i=0}^3 o_i^2 T^6 + 2 \sum_{i=0}^3 n_i o_i T^5 \\ & + \sum_{i=0}^3 (n_i^2 + 2m_i o_i) T^4 + 2 \sum_{i=0}^3 (l_i o_i + m_i n_i) T^3 \\ & + \sum_{i=0}^3 (m_i^2 + 2l_i n_i) T^2 + 2 \sum_{i=0}^3 l_i m_i T + \sum_{i=0}^3 l_i^2 = 0 . \end{aligned} \quad (16)$$

This equation can be efficiently solved numerically, returning three pairs of complex conjugate roots and a single pair of real roots, of which one is negative and the other is positive—the only acceptable root.

The vehicle's yaw angle  $\psi$  is dynamically decoupled from its position. A reference trajectory for yaw can be constructed in a similar way, given that the relative degree of flat output  $\psi$  is two. In the flat output space, therefore, the yaw dynamics can be expressed as a double integrator. Let us define the yaw state vector  $\Psi \triangleq (\psi, \dot{\psi})^T$ , treating  $w \triangleq \ddot{\psi}$  as the virtual control input for this dynamical subsystem. For the yaw trajectory we have boundary conditions on both sides, denoted  $\Psi(0) = \Psi_0 \triangleq (\psi_0, \dot{\psi}_0)$  and  $\Psi(T) = \Psi_T \triangleq (\psi_T, \dot{\psi}_T)$ . The reference yaw trajectory is obtained as Pontryagin's Minimum Principle solution of

$$\begin{cases} \min \frac{1}{2} \int_0^T w(s)^2 ds \\ \text{subject to} \\ \Psi = (\psi, \dot{\psi}), \quad \ddot{\psi}(t) = w(t) \\ \Psi(0) = \Psi_0, \quad \Psi(T) = \Psi_T , \end{cases} \quad (17)$$

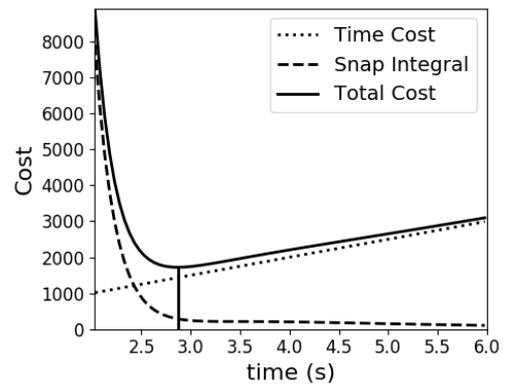
which is a 3<sup>rd</sup> order polynomial of the form

$$\psi(t) = \gamma_0 + \gamma_1 t + \gamma_2 t^2 + \gamma_3 t^3 ,$$

with coefficients given as

$$\begin{aligned} \gamma_0 &= \psi_0 & \gamma_2 &= \frac{1}{2} \left( \frac{6(\psi_T - \psi_0)}{T^2} - \frac{2(\dot{\psi}_T + 2\dot{\psi}_0)}{T} \right) \\ \gamma_1 &= \dot{\psi}_0 & \gamma_3 &= \frac{1}{6} \left( \frac{6(\dot{\psi}_T + \dot{\psi}_0)}{T^2} - \frac{12(\psi_T - \psi_0)}{T^3} \right) . \end{aligned}$$

Figure 3 shows a sample evolution of the cost functional in (11) for some particular set of boundary conditions. The cost associated with the terminal time (dotted curve) increases while the integral of the snap (dashed curve) reduces with time. The vertical line represents the positive real root  $T$  of (16) which minimizes the total cost (solid curve).

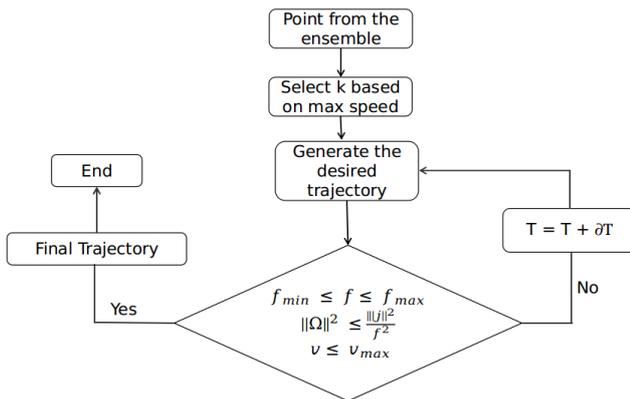


**Figure 3.** Temporal evolution of the optimization cost functional over time, showing how the variation of its two affine cost components contribute to its final value and yield a convex curve with respect to time.

In the context of the particular case study considered here, the best radiation detection performance has been shown to be achieved when the robot closes the distance between its sensor and the source as fast as possible (Sun and Tanner

2015; Yadav and Tanner 2019). This implies that aggressive reference trajectories are preferable, motivating us to select  $k$  based on the maximum speed limitations of the MAV. With this in mind, the methodology outlined above can thus provide motion plans that would be (a) feasible given the dynamic constraints of the vehicle, (b) as aggressive as possible, and (c) high-performance in terms of radiation detection accuracy.

While the reference trajectory is dynamically feasible by design, its conformity to actuation constraints is verified after its generation (Mueller et al. 2015) (see Fig. 4). Here, we have actuation (upper) bounds on the magnitude of the input thrust  $f = m\|a - g\|$  and on the square of the sum of roll and pitch angular velocities in terms of jerk and thrust,  $\|\Omega\|^2/f^2$ . These actuation constraints stem from (a) equipment limitations, and (b) the maximum deceleration that the vehicle can undergo during an emergency stopping maneuver. In the reactive navigation architecture described here, emergency stopping maneuvers are engaged when the MAV cannot find a safe path in its free workspace. In such a scenario, the vehicle traverses the collision-free portion of the reference trajectory from the previous planning cycle and then comes to a stop. Boundary conditions of zero velocity, acceleration and jerk for the candidate trajectories at time  $T$  ensure that any trajectory selected as reference is able to bring the quadrotor to a complete halt if the emergency stop is triggered. The linear velocity constraint is imposed in order to reduce motion blur which would otherwise affect the robot's visual inertial odometry (VIO). Contrary to posing the actuation constraints as hard constraints in the optimal control problem (11), computational efficiency reasons motivate an implementation where constraint satisfaction is verified a posteriori (Fig. 4) at various sampling points of a candidate trajectory. Candidate trajectories failing these tests are discarded.



**Figure 4.** The decision block diagram for the reference trajectory generation process. If the actuation limit check fails, the terminal time  $T$  is increased slightly by some  $\delta T$  and the trajectory is regenerated.

### Local Goal and Collision Costs

Once the ensemble of candidate reference trajectories is generated, the candidate trajectories that intersect with any obstacle is excluded from the ensemble. The methodology is motivated from earlier work of Yadav and Tanner (2020),

and has been modified to suit the new trajectory generation methodology.

The obstacle point cloud  $\mathcal{P}$  is first converted into a KD-Tree and each trajectory is discretized and represented in the form of a finite (size  $n$ ) sequence of points. Thereafter a query is run to find points of the KD-tree that lies within a ball of radius  $r$  from each of these points on the candidate trajectory. Parameter  $r$  is chosen so that it can fully enclose the quadrotor, with a suitable—based on how conservative with respect to the practical risk collision due to uncertainty or disturbances—safety margin. A candidate trajectory is said to be intersecting with an obstacle if any point (among the  $n$  points in which it is discretized) on it has any point of KD-tree within this ball of radius  $r$ ; such a trajectory is excluded from the ensemble.

Among the collision-free candidate trajectories, the optimal one should strike an acceptable balance between safety (against collision), and speed of convergence to the goal point. An illustration of the subsequent process of selecting this optimal reference trajectory is given in Fig. 5 (for details on the associated computational requirements, see Section 7). In lieu of a global planner which would otherwise sketch a complete path from initial to final configurations, the reactive local planner defines an *intermediate point* as the point in the (collision free) ensemble closest to the final goal (denoted IP in Fig. 5a). It then assigns a cost to each trajectory, in the form of a linear weighted sum of two cost components: the first component is based on the distance of each trajectory end-point to the intermediate point, normalized over the maximum distance; the second component is a normalized collision cost that captures how close the trajectory comes to  $\mathcal{P}$ .

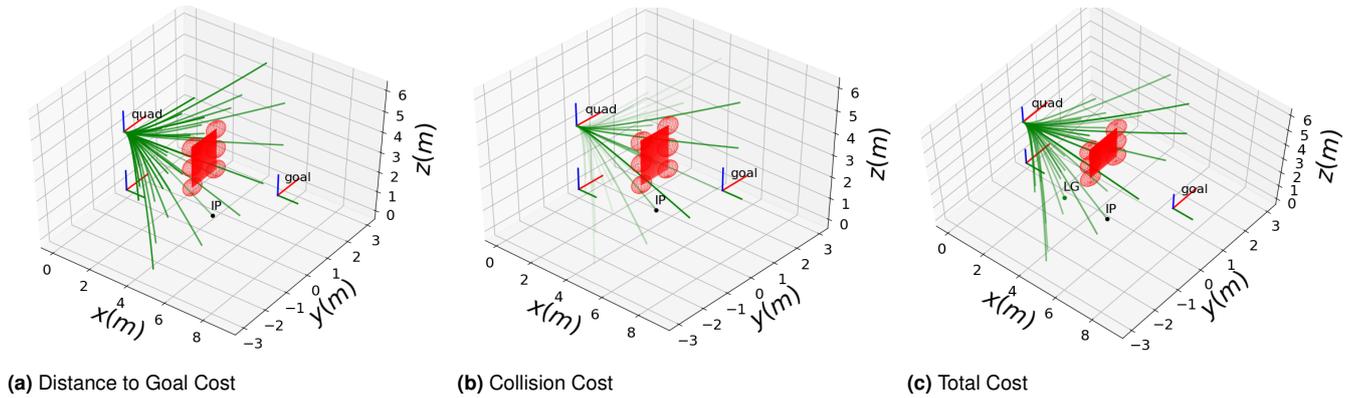
Denote the total number of collision-free trajectories and the Euclidean distance between the end point of the  $i^{\text{th}}$  trajectory and the intermediate point,  $p$  and  $d_i$ , respectively. Set  $d_{\max} \triangleq \max_i d_i$ , and let  $\hat{r} \geq r$  be an additional safety margin (over the radius around detected obstacles). The minimum distance  $\rho_i$  of trajectory  $i$  to obstacles is found by querying a KD-Tree over  $\mathcal{P}$  and minimizing over the query response. With these data at hand, the collision cost for trajectory  $i$  is found as

$$c_{\text{coll}_i} = \begin{cases} \frac{1+\hat{r}^4}{\hat{r}^4} \cdot \frac{[(\rho_i-r)^2-\hat{r}^2]^2}{1+[(\rho_i-r)^2-\hat{r}^2]^2} & \text{if } \rho_i - r \leq \hat{r} \\ 0 & \text{otherwise.} \end{cases}$$

The collision cost function normalizes the cost of each trajectory into the  $[0, 1]$  interval. Thus any trajectory that touches a ball around its nearest obstacle is assigned a collision cost of 1, while any trajectory that lies at least  $\hat{r}$ -away from every obstacle incurs zero cost. All other trajectories are assigned costs within the  $(0, 1)$  interval. The end point of the trajectory with the lowest total cost becomes the *local goal* (i.e., within the FOV) for the planner (marked with a green dot in Fig. 5c).

Selecting positive weights  $k_1, k_2 \in (0, 1)$ , the cost contributions of trajectory  $i \in \{0, \dots, p\}$  due to its proximity to the intermediate point and obstacles are combined into an aggregate cost expression

$$c_i = k_1 \frac{d_i}{d_{\max}} + k_2 c_{\text{coll}_i}.$$



**Figure 5.** Assigning costs to collision-free trajectories. The trajectories with high costs are plotted as darker green lines. The trajectory with minimum total cost is selected as the final trajectory through which the robot traverses. The corresponding end point is the *local goal*.

The trajectory with the least such cost is selected as the reference trajectory between the robot’s current position and the local goal.

The MAV tracks a small portion of the reference trajectory, hereafter referred to as the *control horizon*. The length of the control horizon is dependent on the vehicle’s speed, its computational capabilities, and its sensor update rate. Each time the MAV receives a sensor update, it generates a new reference trajectory (segment) and appends it to the end of the segment it is currently tracking. By design, the transition between subsequent reference trajectory segments is smooth. The process in between sensor updates constitutes a *replanning cycle*. This process is illustrated in Fig. 7 which shows different trajectories generated from the starting point until the end point. The complete implementation for planning, control and state estimation is open-source.<sup>†</sup>

## 6 Target Tracking

### *Object Detection based on SSD-MobileNetV2*

The motivating application for this work is the detection of weak radiological material on moving ground vehicles, and for this to be achieved via aerial means of sensing and measurement in a timely manner, the MAVs need to get close to their target as quickly as possible. To this end, the methodology of Section 5 is extended to dynamic target interception and tracking scenarios. Necessary extensions include the capability of the MAV to autonomously detect its target within its FOV, and estimate its target’s relative position and speed based on visual means. Contrary to Section 3 where the term detection referred to determining whether or not a given target is radioactive, the term “detection” in this section refers to detecting the given ground robot within the image. For this task an open source implementation of MobileNetV2 (Sandler et al. 2018) deep neural network in TensorFlow (Abadi et al. 2016) is utilized and combined with SSD (Liu et al. 2016) for object detection.

This implementation of SSD-MobileNetV2 has been trained on the standard COCO dataset. The desired neural network is then trained via transfer learning on a dataset of approximately 500 images of the Jackal mobile robot from Clearpath Robotics (Fig. 6a). These images were collected from an Intel RealSense D435 camera onboard the MAV in

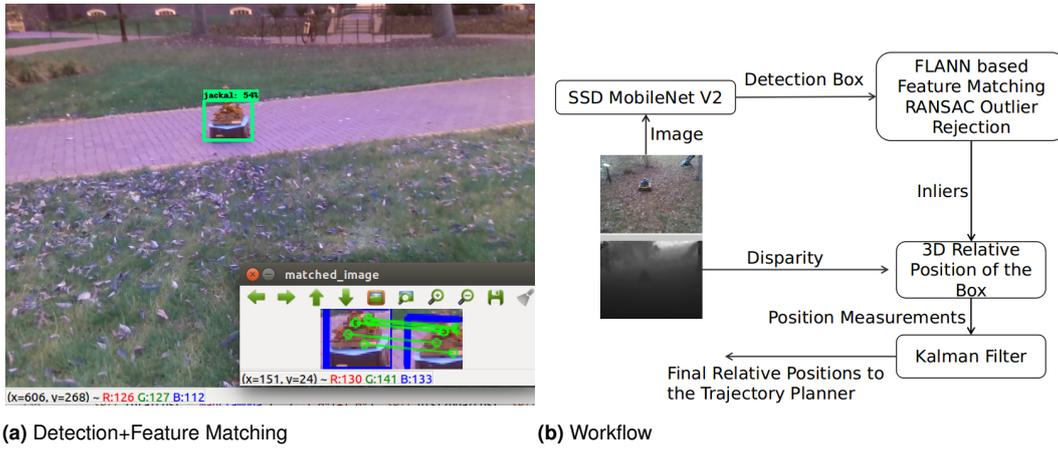
different indoor and outdoor environments, under different lightening conditions and varying background. The network utilizes a  $300 \times 300$  fixed image resizer to increase the inference speed. Adam optimizer was used with initial learning rate of  $2 \times 10^{-4}$  and subsequently  $10^{-4}$ ,  $8 \times 10^{-5}$  and  $4 \times 10^{-5}$  after 4500, 7000 and 10000 steps, respectively, for a total  $2 \times 10^4$  steps. All other parameters are kept to their default values. The network is trained on NVIDIA GTX 1060 GPU and the trained model is subsequently converted into a TensorRT model to run fast inference on NVIDIA Jetson-Nano GPU onboard the MAV.

### *Target 3D Position Estimation*

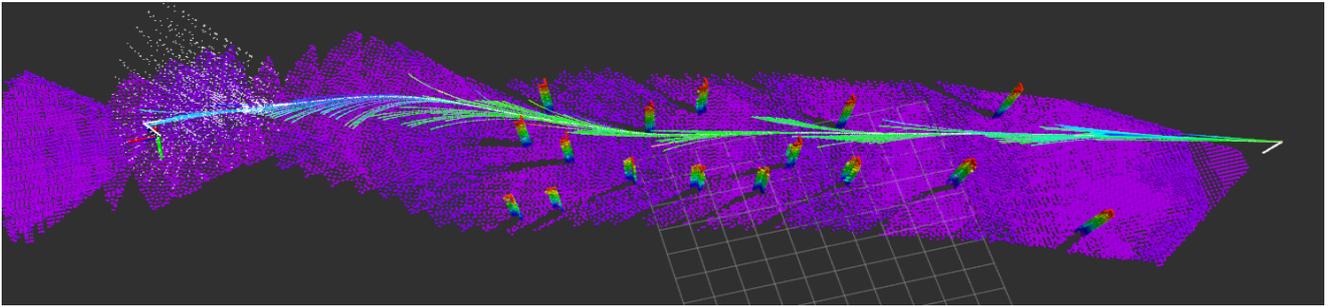
Given a bounding box obtained from the DNN, the onboard RGB-D sensor data allow direct measurement of relative position of the target with respect to the quadrotor. First, all SIFT features are extracted for two consecutive detection boxes and matched using a FLANN-based matcher utilizing a Lowe’s ratio test (Lowe 2004), and a ransac (RANSAC) based outlier rejection was implemented. Utilizing the disparity image from the RGB-D sensor, the  $(u, v)$  position of all the inliers (i.e. the features on the target) in the focal plane can be converted to 3D position of the feature with respect to the camera by utilizing the disparity map. The average of these 3D positions provides an aggregate measurement of relative position of the target with respect to the MAV. This aggregate relative position measurement is used by a Kalman filter, which based on a constant-acceleration motion model for the target, returns 3D position estimates of the target with respect to the MAV’s COG and feeds it to the motion planning algorithm. The entire target position estimation workflow is showcased in Fig. 6b.

The effectiveness of the velocity profile prescribed in (12) is pronounced in the case of intercepting a moving target. Note that in case of a static navigation end goal, the first factor (the time-dependent erf) increases slowly from 0 and converges to 1, while the second term (distance-dependent erf) starts from 1 and converges to 0; this steers the MAV so that it starts at lower speed, progresses to near maximum allowed speeds during most of the navigation maneuver, and

<sup>†</sup>[https://github.com/indsy123/quad\\_navigation\\_and\\_target\\_tracking](https://github.com/indsy123/quad_navigation_and_target_tracking)



**Figure 6.** a) The mobile target in the experimental study. The detection box is shown in green while the image matching between two successive images is shown in the inset. b) Workflow of the target tracker operating at 15 Hz.



**Figure 7.** Receding horizon planning in a cylindrical obstacle Poisson forest.

then smoothly slows down near its goal, thus mimicking a trapezoidal velocity profile. In contrast, in the case of moving target interception the distance-dependent erf in (12) converges instead to a fixed strictly positive value, allowing the MAV align its velocity with that of the target and maintain a fixed distance to it.

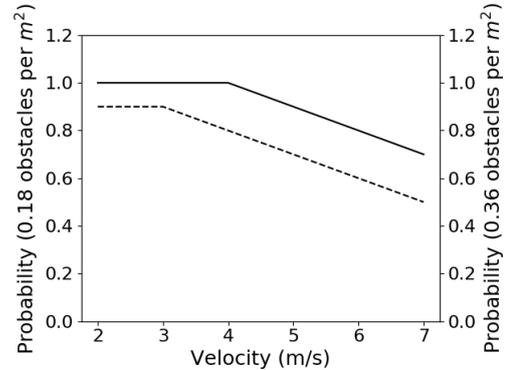
## 7 Implementation Results

### Numerical Testing

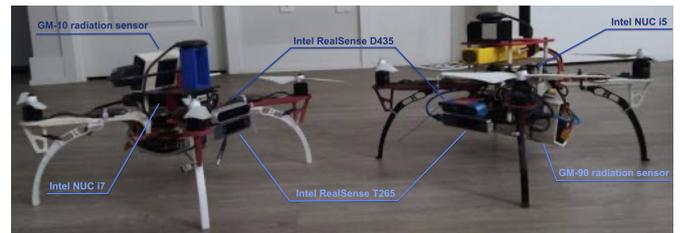
The reported planning and control framework was first tested in simulations based on the ROTORS package (Furrer et al. 2016), within a Poisson forest-like environment with obstacle densities of 18 and 36 obstacles in a 100 m<sup>2</sup> area. A typical RVIZ visualization of the resulting algorithm execution is shown in Fig. 7. As the MAV flies from right to left, the figure marks in green color the different trajectories generated within each planning cycle, while the FOV grid point ensemble is shown in white toward the end of the vehicle's path. Figure 8 shows the probability of mission success as a function of MAV speed at two different obstacle densities (cf. Karaman and Frazzoli (2012)).

### Experimental Testing

The custom-build MAVs used for experimental testing shown in Fig. 9 are based on a DJI Flamewheel F450 frame. Computationally, one of them features an on board Intel NUC Core i7-8650U quad core CPU@1.9 GHz×4 while the other has an Intel NUC Core i5-7300U dual core CPU@2.6 GHz×2. Both uses 16GB RAM, 128GB SSD and



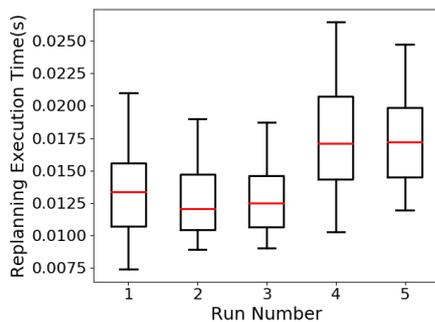
**Figure 8.** Mission (target interception with collision avoidance) success probability in Poisson forests as a function of obstacle density and MAV speed. Solid line corresponds to left axis (0.18 obstacle/m<sup>2</sup>) while dashed line corresponds to right.



**Figure 9.** A pair of quadrotors fitted with COTS GM counters. The one on the left carries a GM-10 counter (Blackcat Systems), while the one on the right features a GM-90 (more sensitive) GM counter from the same product line.

a Pixhawk flight controller that is given the desired thrust magnitude and rate commands, which are then tracked using an onboard body rate controller. A point cloud is generated by an Intel RealSense-D435 depth RGB-D camera (640×480 pixel, 30 Hz) while the RealSense-T265 v1-sensor (2 848×800 pixel 30 Hz cameras with a 200 Hz IMU) is used for inertial odometry. This lightweight sensor package provides reliable depth information for up to 5 m. A voxel filter reduces the density of the generated point cloud to a uniform size of 0.125 m, which is considered adequate for typical obstacle avoidance purposes. OpenVINS provides state estimation and real time odometry at 30 Hz. Ultimately, the pipeline’s computational bottleneck is the inference script that can only be run at 15 Hz on the CPU of the intel NUC, and therefore the entire receding horizon planning and target tracking software is constrained at 15 Hz.

Over five different runs, each of overall trajectory length of 25 m, in both indoor and outdoor environments the 95% quartile of execution time is well below 0.02 seconds (Fig. 10). The total number of candidate trajectories generated in each planning cycles can vary depending upon the hardware capabilities and replanning frequency and were varied between 50 to 300 in experimental setup. The input pointcloud after voxel filtering contained around 800-1400 points. The execution times are almost half of those obtained on those systems in earlier studies (Yadav and Tanner 2020) that uses an optimization based trajectory generation methodology. In this configuration, the MAV flies safely among moderately dense outdoor obstacle environments at speeds of 4.5–5 m/s. These speeds surpass or are comparable to those reported in recent literature (Gao and Shen 2016; Sikang Liu et al. 2016; Fragoso et al. 2018) without using any high-end and expensive sensors and their achievement is attributed to the ability to replan faster.



**Figure 10.** Replanning execution time for the quadrotor with NUC Core i7-8650U CPU@1.9 GHz×4 processor, averaged over five different experimental runs. The box corresponds to 5-95 quartile while the median is marked red.

The video attachment that accompanies this paper features a number of experiments, the first two of which show the MAV to navigate at speeds of 2 m/s and 4.5 m/s in an outdoor environment for a total of 40 m and 60 m respectively. The video showcases first and third person views of the cluttered scene traversed by the MAV, alongside an RVIZ representation of generated trajectories amongst the sensed obstacles. The octomap featured in the video is provided for illustration purposes only to mark obstacle locations along the path of the vehicle, and is not used for planning. The third

experiment included in the video attachment demonstrates the MAV’s capability to avoid moving obstacles in its field of view, while the fourth showcases receding horizon navigation with combined obstacle and target tracking abilities. The latter utilizes a neural network with feature matching between two subsequent images.

The approach presented here has some similarities with the motion primitive based planning approaches (Likhachev and Ferguson 2009) in which the feasible path is searched within the workspace using predefined motion primitives. Unlike such methods, the presented approach does not require a pre-computed lookup table to determine the primitives. Other motion primitive based approaches that do not require precomputation of lookup tables (Liu et al. 2017) are effective when the entire trajectory is searched a priori, but suffers from performance degradation when used in a receding horizon fashion (replanning at 3Hz at a maximum speed of up to 3m/s). None of these approaches allow the dynamic adjustment of speeds in the vicinity of the target that was made possible in our approach by formulation trajectory generation problem as multi-objective optimal control problem and dynamically adjusting the relative weight parameter  $k$  in (11).

Ultimately, the MAV’s speed will be limited primarily by the computational capabilities and the need to implement a safe emergency stop in the absence of an effective motion plan—a possibility which cannot be eliminated in purely reactive and local planning methods. Since purely reactive approaches are supposed to rely *exclusively* on local information, convergence to the navigation goal cannot be guaranteed for all possible scenarios. (The reported reactive planning architecture can however be integrated with an exact global planner, and overcome these limitations at the price of using global information (Yadav and Tanner 2021b).) This investigation naturally exposes the limits of purely reactive motion planning approaches. It is expected that knowledge of those limits, can guide the development of hybrid (local & global (Yadav and Tanner 2019)) MAV motion planning methodologies destined for deployment in environments where uncertainty is reasonably well characterized, in order to complement each other and operate robustly in real-world scenarios.

While on one hand the separation of sensing from trajectory generation (as in the presented approach) can induce latencies, their co-design using powerful on board computation and the incorporation of pre-trained neural networks for trajectory generation can further boost vehicle speeds to impressive levels (Loquercio et al. 2021). To reach speeds up to 10 m/s, however, one would also require a very favorable power-to-weight ratio (e.g.  $\sim 4.0$  (Loquercio et al. 2021)), which may be challenging to maintain depending on the *mission-mandated sensor payload*. In this paper, the MAVs featured a power-to-weight ratio in the 1.55–1.75 range. Besides an effort to achieve a more favorable power-to-weight ratio, we postulate that additional speed improvements can be achieved with the use of event cameras (Falanga et al. 2020).

### Radiation Detection

The radiation detection algorithm is based on a Neyman-Pearson based, fixed time interval binary hypothesis test

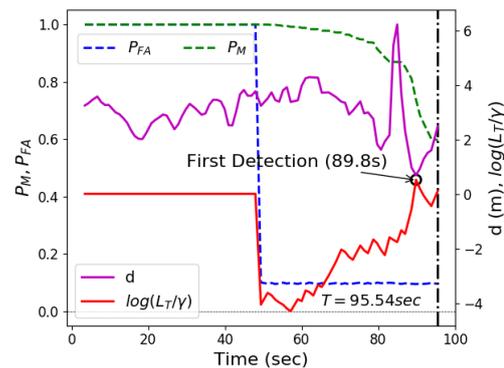
(refer to Pahlajani et al. (2014) for a more detailed exposition). At the heart of the test is a likelihood ratio of a statistic  $L_T$  (2) calculated based on the history of relative distance between the airborne radiation sensor and the hypothesized source in addition to the aggregated counts over the sensor's (predetermined) integration time interval  $T$ . This likelihood ratio is compared against a fixed threshold value  $\gamma$  that also depends on the relative distance and the acceptable bound on the probability of false alarm  $P_{FA}$ . The optimal value of  $p^*$  is obtained by solving (5) and then the threshold is calculated by evaluating (7) at  $p^*$ . The remaining parameters for this test, including background naturally occurring (background) radiation rate and sensor characteristics, are determined by computational and experimental calibration processes.

Our detection tests involve a sequence of controlled experiments in which the efficiency of the radiation sensors and the variation of its expected count rate as a function of its distance to the source is estimated. The MAVs featured in Fig. 9 were deployed in both indoor and outdoor experiments, where their task was to locate and intercept a ground target (the remotely controlled ClearPath Robotics Jackal in Fig. 1) moving along an unspecified path with unknown but bounded speed. The ground robot carried an approximately  $8 \mu\text{Ci}$  radioactivity source which the MAVs had to detect.

Each experiment involved the MAV tracking the ground vehicle for certain time  $T$ . To minimize radiation exposure risk, Monte Carlo simulations using GAZEBO have been performed to complement the experimental validation. The counts were generated using thinning algorithm in the simulations (Pasupathy 2009). These simulations were used for the estimation of the minimum length  $T$  of the time interval for radiation measurement collection that is more likely to yield a confident and accurate classification of the target relative to its radioactivity. This overall process suggested a sensor integration window set at  $T = 100$  seconds for the GM-10 counter (at the median of the distribution with 5% and 95% percentiles estimated at 71 and 136 seconds, respectively), for a radiation source of activity around  $8.2 \mu\text{Ci}$ . During that fixed time interval, the distance between sensor and source should be at most within 3–2.5 m, otherwise the source emissions completely blend into the background, rendering detection extremely unlikely given the MAV's operating time constraints imposed by on-board power. For that same source and range to sensor, the GM-90 counter appeared to need approximately  $T = 70$  seconds, with a 5% and 95% percentiles at 65 and 96 seconds, respectively.

The receding horizon planning and control strategy of Section 5 ensures that the MAV closes this distance as fast as possible, thus enabling the onboard Geiger counters to collect informative measurements. The MAV has to maintain a certain minimum distance from the target to keep it in its camera's limited FOV; as a result, the relative distance should not be arbitrarily reduced. Variations in the relative distance between sensor and source can be attributed to motion perturbations, as the (remotely operated) target performs avoidance maneuvers. Although generally robust, the neural network does not furnish guarantees against false target

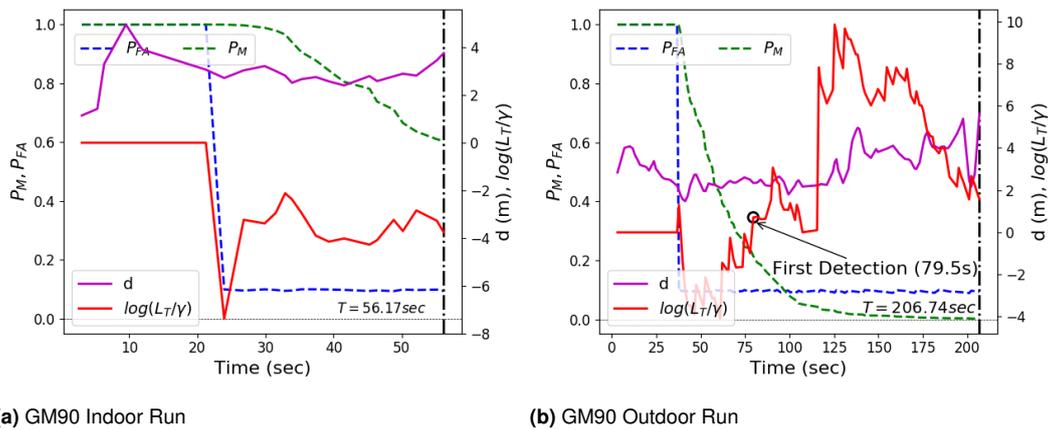
identification, and this can contribute to relative distance estimate outliers (see e.g. Fig. 11, around the 85<sup>th</sup> second).



**Figure 11.** Detection parameters for GM-10 sensor as a function of decision time  $T$ . Bound on probability of false alarm: dashed blue; bound on the probability of missed detection: dashed green; ratio  $\log L_T/\gamma$ : solid red; sensor-source distance: solid magenta.

Figure 11 presents the results of one radiation detection experiment conducted in an indoor environment (an abandoned warehouse) using the MAV that has GM-10 counter. It shows the evolution of the estimate of the relative distance  $d$  between the MAV and the ground robot as the latter moves with unknown and time-varying speed. The relative distance is estimated in real time via the target tracking pipeline described in the section on target tracking. The dashed curves in Fig. 11 indicate the evolution of Chernoff bounds on the probability of false alarm,  $P_{FA}$ , and probability of missed detection  $P_M$  (Pahlajani et al. 2014). The bound on the probability of false alarm appears to drop below the acceptable upper limit after approximately 50 seconds from the start of the experimental run, after which the bound on the probability of missed detection  $P_M$  also starts to slowly decrease monotonically—the latter is a decreasing function of the sensor integration time and distance between sensor and source (Pahlajani et al. 2014). The graph of the logarithm of the likelihood ratio  $L_T$  over the detection threshold  $\gamma$  over time is marked in red; this process is stochastic because it depends directly on the arrival time of gamma rays on the sensor. The initial segment of the red curve corresponds to the initial time period during which the constraint on  $P_{FA}$  has not been satisfied and  $\log L_T/\gamma$  has been kept at 0. The experiment is concluded at 95.54 seconds and the likelihood ratio  $L_T$  exceeds its threshold value at 89.8 seconds indicating the presence of the radiation source on the ground target (marked with a black circle in the plot). The likelihood ratio had actually crossed the threshold before that time, but the experiment was continued because that event was observed significantly earlier than the recommended sensor integration window.

Figures 12a and 12b showcase two different runs where the MAV featuring the GM-90 counter was utilized. The experimental run of Fig. 12a shows an instance where the MAV did not have enough time to detect the source. This experiment was performed in the same indoor facility as that used for the run of Fig. 11. Here, the radiation sensor integration window is 56 seconds. The bound on the probability of missed detection is then still around 0.6,



**Figure 12.** Detection parameters for GM-90 sensor as a function of decision time  $T$ . Bound on probability of false alarm: dashed blue; bound on the probability of missed detection: dashed green; ratio  $\log L_T/\gamma$ : solid red; sensor-source distance: solid magenta. a) Detection time interval 56 seconds. b) Detection time interval 206 seconds.

comparable to the conditions under which the detection of Fig. 11 was achieved, but this  $T$  is below the 5% percentile for the recommended exposure time.

Figure 12b depicts the results of a longer chase by the MAV carrying the GM-90 counter conducted outdoors. This time, the integration window was extended to more than 200 seconds. In addition to the effect of sensor integration window length on detection accuracy, Fig. 12b shows more clearly the evolution of the bounds on the decision test’s error probabilities  $P_{FA}$  and  $P_M$ . At the time of decision, the bound on the probability of miss,  $P_M$  is almost zero, indicating very high probability for accurate decision-making. Although the statistic  $\log L_T/\gamma$  crosses becomes positive for the first time shortly after 70 seconds, at that time the bound  $P_M$  is around 0.3. It is of interest that towards the end of the integration window, the statistic  $\log L_T/\gamma$  decreases, most likely due to the target being able to open up its distance with respect to the pursuing MAV—which by that time was experiencing a drop in its power reserves; same trend can be noticed in Fig. 12a.

The accompanying video attachment includes two experiments of target tracking with radiation detection (experiments #5 and #6). These are cases where the Jackal robot is steered manually while carrying an  $8.2 \mu\text{Ci}$  source. The plot at the bottom left of the video represents graphically the evolution of the detection parameters as depicted in Figs. 11 and 12b. To reduce the size of the video file, only selected portions of these experimental runs are included and the video is accelerated at four times the normal speed.

## 8 Conclusions

The challenges that a completely autonomous MAV equipped with short-ranged sensors faces when tasked to navigate in a completely unknown and cluttered environment are more pronounced when the complete data and signal processing (such as radiation detection) pipeline needs to be run onboard, and the motion of the vehicle can have adverse effect on the quality and quantity of the data. Under these conditions, a motion planner that aims at operating the robot within a reasonable safety envelope has to strike a balance between safety, platform limitations, and mission-informed

active sensing. In this context, an adaptive, purely reactive receding horizon motion planning and control strategy has been developed that co-designs the planning, safe navigation, target tracking and decision-making components. Not only can such a navigation strategy be remarkably effective even in the absence of global environment and platform information precludes formal completeness guarantees, but it can also be integrated with an exact global planner when prior knowledge permits, to furnish formal performance guarantees. The work thus pushes the envelope on what is achievable with *purely reactive* but deliberate MAV navigation in cluttered environments, particularly at the low-end of the technology and sensor sophistication spectrum.

## Acknowledgements

Special thanks to Paul Huang and his RPNG group, specifically to Kevin Eckenhoff and Patrick Geneva for his contributions on perception.

## Funding

This work was made possible in part by DTRA grant #HDTRA1-16-1-0039 and ARL grant #W911NF-20-2-0098.

## References

- Abadi M, Agarwal A, Barham P, Brevdo E, Chen Z, Citro C, Corrado GS, Davis A, Dean J, Devin M, Ghemawat S, Goodfellow I, Harp A, Irving G, Isard M, Jia Y, Jozefowicz R, Kaiser L, Kudlur M, Levenberg J, Mane D, Monga R, Moore S, Murray D, Olah C, Schuster M, Shlens J, Steiner B, Sutskever I, Talwar K, Tucker P, Vanhoucke V, Vasudevan V, Viegas F, Vinyals O, Warden P, Wattenberg M, Wicke M, Yu Y and Zheng X (2016) Tensorflow: Large-scale machine learning on heterogeneous distributed systems.
- Arslan O and Koditschek DE (2019) Sensor-based reactive navigation in unknown convex sphere worlds. *The International Journal of Robotics Research* 38(2-3): 196–223.
- Athans M and Falb PL (1966) *Optimal control : an introduction to the theory and its applications*. McGraw-Hill New York ; Sydney.
- Cadena C, Carlone L, Carrillo H, Latif Y, Scaramuzza D, Neira J, Reid I and Leonard JJ (2016) Past, present, and future of

- simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on Robotics* 32(6): 1309–1332.
- Chen J, Liu T and Shen S (2016) Tracking a moving target in cluttered environments using a quadrotor. In: *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*. pp. 446–453. DOI:10.1109/IROS.2016.7759092.
- Chojnacki M and Indelman V (2018) Vision-based dynamic target trajectory and ego-motion estimation using incremental light bundle adjustment. *International Journal of Micro Air Vehicles* 10(2): 157–170.
- Falanga D, Kleber K and Scaramuzza D (2020) Dynamic obstacle avoidance for quadrotors with event cameras. *Science Robotics* 5(40): eaaz9712.
- Fragoso AT, Cigla C, Brockers R and Matthies LH (2018) Dynamically feasible motion planning for micro air vehicles using an egocylinder. In: Hutter M and Siegwart R (eds.) *Field and Service Robotics*. Cham: Springer International Publishing. ISBN 978-3-319-67361-5, pp. 433–447.
- Furrer F, Burri M, Achtelik M and Siegwart R (2016) *Robot Operating System (ROS): The Complete Reference (Volume 1)*, chapter RotorS—A Modular Gazebo MAV Simulator Framework. Cham: Springer International Publishing. ISBN 978-3-319-26054-9, pp. 595–625.
- Gao F and Shen S (2016) Online quadrotor trajectory generation and autonomous navigation on point clouds. In: *proceedings of IEEE International Symposium on Safety, Security, and Rescue Robotics*. pp. 139–146.
- Geneva P, Eckenhoff K, Lee W, Yang Y and Huang G (2020) Openvins: A research platform for visual-inertial estimation. In: *Proceedings of the IEEE International Conference on Robotics and Automation*. Paris, France.
- Giusti A, Guzzi J, Cireşan DC, He F, Rodríguez JP, Fontana F, Faessler M, Forster C, Schmidhuber J, Caro GD, Scaramuzza D and Gambardella LM (2016) A machine learning approach to visual perception of forest trails for mobile robots. *IEEE Robotics and Automation Letters* 1(2): 661–667.
- Hornung A, Wurm KM, Bennewitz M, Stachniss C and Burgard W (2013) Octomap: an efficient probabilistic 3d mapping framework based on octrees. *Autonomous Robots* 34(3): 189–206.
- Karaman S and Frazzoli E (2012) High-speed flight in an ergodic forest. *CoRR* abs/1202.0253.
- Landry B, Deits R, Florence PR and Tedrake R (2016) Aggressive quadrotor flight through cluttered environments using mixed integer programming. In: *proceedings of IEEE International Conference on Robotics and Automation*. pp. 1469–1475.
- Lee T, Leoky M and McClamroch NH (2010) Geometric tracking control of a quadrotor uav on se(3). In: *Proceedings of 49th IEEE Conference on Decision and Control*. pp. 5420–5425.
- Likhachev M and Ferguson D (2009) Planning long dynamically feasible maneuvers for autonomous vehicles. *The International Journal of Robotics Research* 28(8): 933–945.
- Liu S, Atanasov N, Mohta K and Kumar V (2017) Search-based motion planning for quadrotors using linear quadratic minimum time control. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 2872–2879.
- Liu S, Watterson M, Mohta K, Sun K, Bhattacharya S, Taylor CJ and Kumar V (2017) Planning dynamically feasible trajectories for quadrotors using safe flight corridors in 3-d complex environments. *IEEE Robotics and Automation Letters* 2(3): 1688–1695.
- Liu W, Anguelov D, Erhan D, Szegedy C, Reed S, Fu CY and Berg AC (2016) Ssd: Single shot multibox detector. *Lecture Notes in Computer Science* : 21–37.
- Loquercio A, Kaufmann E, Ranftl R, Dosovitskiy A, Koltun V and Scaramuzza D (2020) Deep drone racing: From simulation to reality with domain randomization. *IEEE Transactions on Robotics* 36(1): 1–14.
- Loquercio A, Kaufmann E, Ranftl R, Müller M, Koltun V and Scaramuzza D (2021) Learning high-speed flight in the wild. *Science Robotics* 6(59): eabg5810.
- Loquercio A, Maqueda AI, del-Blanco CR and Scaramuzza D (2018) Dronet: Learning to fly by driving. *IEEE Robotics and Automation Letters* 3(2): 1088–1095.
- Lowe DG (2004) Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* 60(2): 91–110.
- Mellinger D and Kumar V (2011) Minimum snap trajectory generation and control for quadrotors. In: *Proceedings of IEEE International Conference on Robotics and Automation*. pp. 2520–2525.
- Mohta et al K (2018) Fast, autonomous flight in gps-denied and cluttered environments. *Journal of Field Robotics* 35(1): 101–120.
- Mueller MW, Hehn M and D’Andrea R (2015) A computationally efficient motion primitive for quadcopter trajectory generation. *IEEE Transactions on Robotics* 31(6): 1294–1310.
- Nemzek RJ, Dreicer JS, Torney DC and Warnock TT (2004) Distributed sensor networks for detection of mobile radioactive sources. *IEEE Transactions on Nuclear Science* 51(4): 1693–1700.
- Oleynikova H, Burri M, Taylor Z, Nieto J, Siegwart R and Galceran E (2016) Continuous-time trajectory optimization for online uav replanning. In: *proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*. pp. 5332–5339.
- Oleynikova H, Taylor Z, Fehr M, Siegwart R and Nieto J (2017) Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning. In: *proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*. pp. 1366–1373.
- Pahlajani CD, Sun J, Poulakakis I and Tanner HG (2014) Error probability bounds for nuclear detection: Improving accuracy through controlled mobility. *Automatica* 50(10): 2470–2481.
- Pasupathy R (2009) Generating nonhomogeneous Poisson processes. In: *Wiley Encyclopedia of Operations Research and Management Science*. Wiley.
- Penin B, Giordano PR and Chaumette F (2018) Vision-based reactive planning for aggressive target tracking while avoiding collisions and occlusions. *IEEE Robotics and Automation Letters* 3(4): 3725–3732.
- Redmon J, Divvala S, Girshick R and Farhadi A (2016) You only look once: Unified, real-time object detection. In: *proceedings of IEEE Conference on Computer Vision and Pattern Recognition*. pp. 779–788.
- Richter C, Bry A and Roy N (2016) *Polynomial Trajectory Planning for Aggressive Quadrotor Flight in Dense Indoor Environments*. Cham: Robotics Research: The 16th International Symposium ISRR. Springer International Publishing. ISBN

- 978-3-319-28872-7, pp. 649–666.
- Rimon E and Koditschek DE (1992) Exact robot navigation using artificial potential functions. *IEEE Transactions on Robotics and Automation* 8(5): 501–518. DOI:10.1109/70.163777.
- Ryll M, Ware J, Carter J and Roy N (2019) Efficient trajectory planning for high speed flight in unknown environments. In: *proceedings of International Conference on Robotics and Automation*. pp. 732–738.
- Sandler M, Howard A, Zhu M, Zhmoginov A and Chen L (2018) Mobilenetv2: Inverted residuals and linear bottlenecks. In: *proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4510–4520.
- Sikang Liu, Watterson M, Tang S and Kumar V (2016) High speed navigation for quadrotors with limited onboard sensing. In: *proceedings of IEEE International Conference on Robotics and Automation*. pp. 1484–1491.
- Smolyanskiy N, Kamenev A, Smith J and Birchfield S (2017) Toward low-flying autonomous mav trail navigation using deep neural networks for environmental awareness.
- Snyder DL (1975) *Random point processes*. Wiley New York. ISBN 0471810215.
- Sun J and Tanner HG (2015) Constrained decision making for low-count radiation detection by mobile sensors. *Autonomous Robots* 39(4): 519–536.
- Tanner HG and Piovesan JL (2010) Randomized receding horizon navigation. *IEEE Transactions on Automatic Control* 55(11): 2640–2644.
- Thomas J, Welde J, Loianno G, Daniilidis K and Kumar V (2017) Autonomous flight for detection, localization, and tracking of moving targets with a small quadrotor. *IEEE Robotics and Automation Letters* 2(3): 1762–1769. DOI:10.1109/LRA.2017.2702198.
- Vasilopoulos V, Pavlakos G, Bowman SL, Caporale JD, Daniilidis K, Pappas GJ and Koditschek DE (2020) Reactive semantic planning in unexplored semantic environments using deep perceptual feedback. *IEEE Robotics and Automation Letters* 5(3): 4455–4462.
- Wang CC, Thorpe C, Thrun S, Hebert M and Durrant-Whyte H (2007) Simultaneous localization, mapping and moving object tracking. *The International Journal of Robotics Research* 26(9): 889–916.
- Yadav I and Tanner HG (2017) Controlled mobile radiation detection under stochastic uncertainty. *IEEE Control Systems Letters* 1(1): 194–199.
- Yadav I and Tanner HG (2019) Mobile radiation source interception by aerial robot swarms. In: *2019 International Symposium on Multi-Robot and Multi-Agent Systems*. pp. 63–69.
- Yadav I and Tanner HG (2020) Reactive receding horizon planning and control for quadrotors with limited on-board sensing. In: *proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*. pp. 7058–7063.
- Yadav I and Tanner HG (2021a) Exact decentralized receding horizon planning for multiple aerial vehicles. In: *Proceedings of the IEEE Conference on Decision and Control*. (to appear).
- Yadav I and Tanner HG (2021b) Exact decentralized receding horizon planning for multiple aerial vehicles. In: *Conference on Decision and Control*. (Accepted).
- Yuan L, Reardon C, Warnell G and Loianno G (2019) Human gaze-driven spatial tasking of an autonomous mav. *IEEE Robotics and Automation Letters* 4(2): 1343–1350.