

Motion Planning and Visual-Inertial Target Tracking for UAV-based Radiation Detection

Indrajeet Yadav*, Kevin Eckenhoff*, Guoquan Huang, and Herbert G. Tanner

Abstract—This paper addresses the problem of detecting radioactive material in transit using an unmanned aerial vehicle (UAV) of minimal sensing capability, where the objective is to classify the target’s radioactivity as the vehicle plans its paths through the workspace while tracking the target for a short time interval. Assuming that the prior map of the workspace is available, this paper proposes a motion planning framework that integrates a navigation function based planner with a tightly-coupled visual-inertial localization and target tracking algorithm, and generates dynamically feasible trajectories that provably converge to a moving target while avoiding obstacles. The efficacy of the proposed approach is validated through Gazebo simulations.

I. INTRODUCTION

Consider a case where an (uncontrolled) mobile platform moves in a cluttered, GPS-denied environment while possibly carrying a package that is weakly radioactive. One would like to be able to send a robot after this target and be able to confidently ascertain (hereafter referred to as decision-making or detection) whether this platform is “hot” or “cold,” in a matter of seconds. The ability of this mobile robot to complete this task autonomously in a cluttered and possibly GPS-denied environment, relies on two critical capabilities: (i) estimating consistently and accurately its own state and that of its target within its workspace, and (ii) moving within that workspace in an agile yet safe way that can optimize the task’s performance metric. In this particular case, the performance metric can be formalized in terms of the probability of missed detection, given a certain upper bound on the probability of false alarm, thus framing the target classification problem as a *binary hypothesis testing problem* (whether the target is “hot” or “cold”). In *isolation*, both of these problems have received attention in literature, while existing techniques that offer combined solutions have limitations that stem from underlying assumptions in their constituent technologies for estimation and target tracking.

The problem that this paper is addressing is technically stated as follows: given a prior map of the robot’s workspace, and assuming that the robot’s maximum acceleration exceeds that of its target, one seeks to design (i) a robust and consistent estimator for the robot’s state and that of its moving target which is based solely on visual and inertial measurements, and (ii) a state feedback controller for the robot that uses those estimates to make the robot avoid

collisions with workspace boundaries, intercept the target as quickly as possible, and then keep tracking it.

The state of the art in terms of estimating the motion of a UAV using onboard sensors in GPS-denied environments, is primarily through the fusion of an inertial measurement unit (IMU) and a camera on the hardware side. Algorithmically, the lightweight multi-state constraint Kalman filter (MSCKF) [1] remains popular due to its low computational complexity. This method utilizes stochastic cloning to estimate a sliding window of past robot poses along with the current IMU state. Features are linearly marginalized to utilize their motion information without the need to store them in the state vector. The marginalization allows for the creation of constraints between the window poses while bounding the problem size, resulting in a computationally efficient estimation algorithm. This approach has been extended in many directions, e.g., including camera-to-IMU spatial and temporal calibration [2], handling degenerate motions [3], and enforcing the correct observability properties [4].

While the area of Simultaneous Localization, Mapping, and Moving Object Tracking (SLAMMOT) has seen research efforts in recent years ([5], [6] and references therein), few perform this task either using visual-inertial sensors or in combination with active following as in this work. One idea explored in the context of active tracking [7] is to choose robot motion in ways that minimize future target uncertainty. This methodology was developed for 2D problems and required that the state of the sensor be perfectly known. An impressive active visual-inertial target tracking using a quadrotor *was* performed in [8], yet, the size of tracking errors reported can be detrimental to detection of low-intensity mobile sources of radioactivity, because of the critical role that the distance between sensor and source plays in the context of nuclear measurement [9], [10].

Another work that is close to the one reported in this paper, is the aerial target tracking that has been recently demonstrated in the case of tracking a spherical rolling target [11]. That approach employed a geometric technique similar to visual servoing, a receding horizon strategy that penalizes velocity and position errors, and a UAV motion control scheme based on minimum-snap trajectory generation [12]. It is not clear to what degree this tracking algorithm depends on the target being spherical, but the overall estimation and motion control scheme was engineered to run fast.

The problem of detecting weak nuclear material with low-cost commercial-off-the-shelf (COTS) radiation counters has been looked at primarily in the context of *static* sensor networks [9], [13], [14], and the associated technical challenge that stems from the inverse square law decrease of signal-

*These authors contributed equally to this work.

This work was partially supported by DTRA (HDTRA1-16-1-0039).

The authors are with the Department of Mechanical Engineering, University of Delaware, Newark, DE 19716. Email: {indragt, keck, ghuang, btanner}@udel.edu

to-noise-ratio (SNR) with distance has been clearly identified, both at the analytical and the computational level [9], [13]. Moreover, the detection problem is exacerbated by its coupling with the source localization problem, where the source location is attempted to be ascertained based on counter readings alone. Even when the detection problem is considered in isolation, the inherent analytic complexity makes the Bayesian update of the posterior probabilities practically intractable, even for networks of modest size [9].

An alternative approach that circumvents the computational and analytical complexities associated with Bayesian updates is to treat analytically derived Chernoff bounds on the probability of missed detection (PM) (not being able to detect the present source) and probability of false alarm (FA) (detecting falsely) as surrogates for true probabilities [10], in a Neyman-Pearson framework. Derived in terms of relative distance between the source and the sensor, these bounds can be integrated in optimal motion control designs that steer sensors around mobile targets for the purpose of classifying the latter in terms of their radioactivity [15]. In essence, the optimal strategies appear to maximize over time the solid angle of the radiation sensor, boosting its SNR [9], thus improving the accuracy of decision-making.

Most of radiation detection literature treats mobility, either of sensor or of source, as a disturbance; it therefore does not provide sufficient insight on how to move the sensor to maximize decision-making accuracy. On the other hand, the body of robot navigation and motion planning literature cares about taking a platform from point A to point B , and has not incorporated objectives linked to the statistics of nuclear measurement, and cannot inform about trade-offs between collision avoidance and source interception in cluttered environments. In terms of radiation sensor motion control, the state of the art [15] establishes the validity of minimum-time interception motion planning heuristics for the case of detection of low-count mobile radiation sources and took a first step by integrating the detection optimization results in a navigation function-based motion planning and control framework to offer provably safe and convergent solutions for navigation and tracking in cluttered 2D environments—under the assumption that the sensor’s absolute location in the workspace, as well as the relative position between sensor and source, is *known*.

The work reported in this paper lifts this assumption and presents a motion planning approach that integrates platform model-based differential-geometric controllers [16] to time varying globally convergent 3D navigation function motion planners. This enables a UAV to intercept its target utilizing dynamically feasible and agile yet safe maneuvers. The potential field is integrated to a specifically designed tightly-coupled 3D visual-inertial localization and target tracking approach that provides the UAV and target states for motion planning and feedback control in a known map.

Tight coupling between the states of the UAV and that of the target ensures that the error on estimating the relative distance between them is minimized. Minimizing this error crucial in achieving high decision-making accuracy in radiation detection tasks, particularly when the sources are

weak and detectors are used. The assumption of a known map is not particularly restrictive—areas may have been mapped a priori; the critical element in this application is *time and accuracy*: the time available to detect the source and the accuracy of decision-making. The reported estimation navigation and control approach is integrated into a functional algorithmic architecture, and validated in ROS/Gazebo simulations.

II. UAV MOTION PLANNING AND CONTROL

The block diagram of the system is depicted in Fig. 1. Subsequent sections describe each component separately.

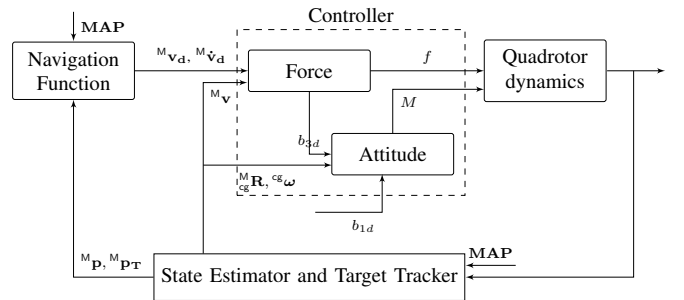


Fig. 1: The navigation function takes the state estimates of the UAV and the target and feeds the smooth velocity and acceleration trajectories to the force controller; which, using the velocity estimates, generates the required force and feeds the stabilizing direction b_{3d} [16] to the attitude controller. Given a suitable heading direction b_{1d} [16] and using estimates of angular velocity and orientation, the attitude controller generates required moments. The workspace map is assumed known.

A. Motion planning

This work extends the construction of navigation functions over squircle worlds (where the objects takes the form of squircles) [17] to 3D workspaces and integrates it with geometric velocity controllers [16]. A navigation function [18] is a real-valued map $\mathcal{V} : \mathcal{F} \rightarrow \mathbb{R}$, constructed on the UAV’s free configuration space \mathcal{F} that when tuned appropriately has a unique minimum at the desired goal configuration and is uniformly maximal over the boundary of \mathcal{F} .

Let ${}^M\mathbf{p}(t)$ denote the position of the UAV in the prior map’s inertial frame M , and ${}^M\mathbf{p}_T(t)$ that of the target at time t , both assumed in \mathcal{F} . For the purposes of control, ${}^M\mathbf{p}(t)$ and ${}^M\mathbf{p}_T(t)$ are assumed known; the process of estimating them is described in Section III. Take r to be the radius of the spherical bubble around the target and define the goal for navigation as minimizing the time varying goal function

$$J({}^M\mathbf{p}, {}^M\mathbf{p}_T) = \|{}^M\mathbf{p} - {}^M\mathbf{p}_T\|^2 - r^2 \quad (1)$$

It has been shown [15] that for a suitably selected obstacle function $\beta({}^M\mathbf{p})$ and a suitably large parameter $\lambda \in \mathbb{R}_+$, there exists a positive number N such that $\forall \kappa \geq N$, the function

$$\tilde{\varphi}({}^M\mathbf{p}, {}^M\mathbf{p}_T) = \frac{J({}^M\mathbf{p}, {}^M\mathbf{p}_T)}{[J({}^M\mathbf{p}, {}^M\mathbf{p}_T)^\kappa + \lambda \beta({}^M\mathbf{p})]^{1/\kappa}} \quad (2)$$

is a navigation function when the free configuration space of the robot is a sphere world \mathcal{S} . Note that the parameter $\lambda \neq 1$ [15] for better tuning.

If instead, the free configuration space of the robot is a forest of squircles then for a diffeomorphism $h_{\lambda_{sq}} : \mathcal{F} \rightarrow$

\mathcal{S} parameterized by a suitably chosen positive parameter $\lambda_{\text{sq}} \in \mathbb{R}_+$, the composition $\varphi = \tilde{\varphi} \circ h_{\lambda_{\text{sq}}}(\mathbf{M}\mathbf{p}, \mathbf{M}\mathbf{p}_T)$ can be shown to be a navigation function on \mathcal{F} [17]: for any position of the target satisfying some reasonable conditions, all (unstable) critical points outside the destination manifold are either nondegenerate with attraction region of measure zero, or inside the target's bubble.

B. UAV Control

Let m denote the mass of the UAV, $\mathbf{J} \in \mathbb{R}^{3 \times 3}$ its moment of inertia about a frame aligned with the principal axes and attached at the center of mass cg , and \mathbf{e}_3 the vector aligned with the standard gravity. The relative orientation between the inertial frame and the principal one at the UAV's center of mass is captured by the rotation matrix ${}^{\text{M}}\mathbf{R}_{\text{cg}} \in \text{SO}(3)$. The linear and angular velocity of the UAV relative to the inertial frame are denoted ${}^{\text{M}}\mathbf{v}$ and ${}^{\text{M}}\boldsymbol{\omega}$, respectively; the angular velocity vector relative to the body cg frame is denoted ${}^{\text{cg}}\boldsymbol{\omega}$. Let $\widehat{\cdot}$ denote the (wedge) operation that maps a vector in \mathbb{R}^3 to a member of the Lie algebra $\mathfrak{so}(3)$. With f expressing the magnitude of the total thrust produced by the thrusters, and \mathbf{M} the total moment relative to the body-fixed cg frame, the dynamics of the UAV is

$${}^{\text{M}}\dot{\mathbf{p}} = {}^{\text{M}}\mathbf{v}, \quad m {}^{\text{M}}\dot{\mathbf{v}} = -m {}^{\text{M}}\mathbf{g} + f {}^{\text{M}}\mathbf{R}_{\text{cg}} \mathbf{e}_3 \quad (3a)$$

$${}^{\text{M}}\dot{\mathbf{R}}_{\text{cg}} = {}^{\text{M}}\mathbf{R}_{\text{cg}} \widehat{{}^{\text{cg}}\boldsymbol{\omega}}, \quad \mathbf{J} {}^{\text{cg}}\boldsymbol{\omega} + {}^{\text{cg}}\boldsymbol{\omega} \times \mathbf{J} {}^{\text{cg}}\boldsymbol{\omega} = \mathbf{M} \quad (3b)$$

The *desired* velocity for the UAV is determined using the navigation function φ . Specifically, if v_{max} denotes the vehicle's maximum speed given the capabilities of its actuators or safety specifications, and k_φ is a positive control gain, then with ∇_x denoting the gradient with respect to variable x , the desired velocity relative to the inertial frame would be

$${}^{\text{M}}\mathbf{v}_d \triangleq -\text{erf}(k_\varphi(\|\mathbf{M}\mathbf{p} - \mathbf{M}\mathbf{p}_T\| - r)) \cdot \frac{\nabla_{\mathbf{p}}\varphi}{\|\nabla_{\mathbf{p}}\varphi\|} \cdot v_{\text{max}} \quad (4)$$

Consequently, the velocity error for the UAV would be $\mathbf{e}_v \triangleq {}^{\text{M}}\mathbf{v} - {}^{\text{M}}\mathbf{v}_d$. A suitable value of k_φ in erf function ensures that the velocity of the UAV approaches that of the target when they are within distance r to each other. The desired heading direction (b_{1d} in Fig. 1) can be selected such that target remains in the field of view, while the desired rotation matrix ${}^{\text{M}}\mathbf{R}_d$ is constructed so that the UAV achieves the desired velocity [16]. The orientation and angular velocity errors are defined accordingly (\top is for transpose), [16]

$$\widehat{\mathbf{e}}_R = \frac{1}{2} ({}^{\text{M}}\mathbf{R}_d \mathbf{R}_{\text{cg}}^\top - \mathbf{R}_{\text{cg}} \mathbf{R}_d) \quad (5a)$$

$$\mathbf{e}_\omega = {}^{\text{cg}}\boldsymbol{\omega} - \mathbf{R}_{\text{cg}} \widehat{\mathbf{e}}_R \mathbf{R}_{\text{cg}}^\top \quad (5b)$$

With these definitions in place, and with k_v , k_R and k_ω denoting positive control gains for velocity, orientation, and angular velocity, respectively, the UAV control inputs are

$$f = (-k_v \mathbf{e}_v + m {}^{\text{M}}\mathbf{g} + m {}^{\text{M}}\dot{\mathbf{v}}_d) \cdot ({}^{\text{M}}\mathbf{R}_{\text{cg}} \mathbf{e}_3) \quad (6a)$$

$$\mathbf{M} = -k_R \widehat{\mathbf{e}}_R - k_\omega \mathbf{e}_\omega + {}^{\text{cg}}\boldsymbol{\omega} \times \mathbf{J} {}^{\text{cg}}\boldsymbol{\omega} \quad (6b)$$

Proposition 1: For sufficiently large k_v , k_R , and k_ω , and given ${}^{\text{M}}\mathbf{p}(t)$ and ${}^{\text{M}}\mathbf{p}_T(t)$ for any $t > 0$, if ${}^{\text{M}}\mathbf{p}(0)$, ${}^{\text{M}}\mathbf{p}_T(0) \in \mathcal{F}$, then ${}^{\text{M}}\mathbf{p}$ converges asymptotically to a ball of radius r around ${}^{\text{M}}\mathbf{p}_T$, and \mathcal{F} is positively invariant.

Proof: Only a rough sketch of the proof is included here due to space limitations. The time-varying motion of the target makes $\tilde{\varphi}(\mathbf{M}\mathbf{p}, \mathbf{M}\mathbf{p}_T)$ a time varying Lyapunov function; invoking Barbalat's Lemma and showing $\lim_{t \rightarrow \infty} \dot{\tilde{\varphi}} = 0$ suffices to prove asymptotic convergence. For a UAV speed sufficiently higher than the target's, and with appropriately tuned κ , it can be shown that the $\tilde{\varphi}$ exists and bounded and that $\dot{\tilde{\varphi}}$ is also bounded. The conditions of the Lemma are subsequently fulfilled. ■

III. STATE ESTIMATION

The UAV is assumed to be equipped with a set of stereo cameras, and an IMU which provides linear acceleration and angular velocity readings. Cameras provide bearing measurements to three classes of objects: the target, known landmarks, and unknown features. To accommodate on-board computational constraints and need for low-latency, the MSCKF [1] is extended to estimate (i) the state of the UAV and its target, and (ii) the transformation between the UAV and map frames, through fusion of inertial (IMU) and visual (camera) data.

The state of the UAV is parameterized in terms of the unit quaternion ${}^{\text{G}}\bar{\mathbf{q}}$ expressing the rotation from the IMU's global frame G to the local IMU frame [19] I_k at measurement time step k , the IMU's gyro bias \mathbf{b}_{w_k} , the vehicle's velocity ${}^{\text{G}}\mathbf{v}_k$ in the IMU's global frame, the IMU's accelerometer bias \mathbf{b}_{a_k} , and the robot's current position ${}^{\text{G}}\mathbf{p}_k$ in the IMU's global frame:

$$\mathbf{x}_{\text{IMU},k} = \left[{}^{\text{I}_k}\bar{\mathbf{q}}^\top \quad \mathbf{b}_{w_k}^\top \quad {}^{\text{G}}\mathbf{v}_k^\top \quad \mathbf{b}_{a_k}^\top \quad {}^{\text{G}}\mathbf{p}_k^\top \right]^\top \quad (7)$$

The target's state for the estimator is a vector $\mathbf{T} \in \mathbb{R}^{3n+3}$ containing derivatives ${}^{\text{G}}\mathbf{p}_{T_k}^{(i)}$ of the target's current position ${}^{\text{G}}\mathbf{p}_{T_k}$ at the IMU's global frame up to order n .

$$\mathbf{T}_k = \left[{}^{\text{G}}\mathbf{p}_{T_k}^\top \quad {}^{\text{G}}\mathbf{p}_{T_k}^{(1)\top} \quad {}^{\text{G}}\mathbf{p}_{T_k}^{(2)\top} \quad \dots \quad {}^{\text{G}}\mathbf{p}_{T_k}^{(n)\top} \right]^\top$$

The estimator also needs to estimate a number of *static* (no dynamics involved) quantities, among them *clones* of its past IMU poses at measurement time step i as well as the pose of the frame M of the prior map

$$\mathbf{x}_{\text{clone},i} = \left[{}^{\text{I}_i}\bar{\mathbf{q}}^\top \quad {}^{\text{G}}\mathbf{p}_i^\top \right]^\top, \quad \mathbf{x}_{\text{MAP}} = \left[{}^{\text{M}}\bar{\mathbf{q}}^\top \quad {}^{\text{G}}\mathbf{p}_{\text{M}}^\top \right]^\top \quad (8)$$

relative to its own IMU's global frame G , which is needed as the estimator initializes its own frame during startup, *without* any knowledge of its pose relative to M . If the estimator maintains N past clones,

$$\mathbf{x}_{\text{clones},k} = \left[\mathbf{x}_{\text{clone},k}^\top \quad \mathbf{x}_{\text{clone},k-1}^\top \quad \dots \quad \mathbf{x}_{\text{clone},k-N+1}^\top \right]^\top$$

the static parameters that it has to estimate at step k can be grouped to a vector $\mathbf{x}_{\text{static},k} = \left[\mathbf{x}_{\text{clones},k}^\top \quad \mathbf{x}_{\text{MAP}}^\top \right]^\top$, allowing the combined state of the estimator to be written $\mathbf{x}_k = \left[\mathbf{x}_{\text{IMU},k}^\top \quad \mathbf{T}_k^\top \quad \mathbf{x}_{\text{static},k}^\top \right]^\top$. Note that \mathbf{x} does not live in a vector space, and so we employ indirect filtering [19]. In particular, the relationship between \mathbf{x} and its estimate $\hat{\mathbf{x}}$ is understood via the *error state* $\delta\mathbf{x}$ through the generalized update operation $\mathbf{x} = \hat{\mathbf{x}} \boxplus \delta\mathbf{x}$, where the operator \boxplus takes the form of simple addition for vector terms and quaternion

multiplication for some quaternion error $\delta\theta$, so that $\hat{\mathbf{q}} = \hat{\mathbf{q}} \boxplus \delta\theta \equiv \delta\hat{\mathbf{q}} \otimes \hat{\mathbf{q}}$ with $\delta\hat{\mathbf{q}} \approx \begin{bmatrix} \frac{\delta\theta^\top}{2} & 1 \end{bmatrix}^\top$.

A standard assumption is that the vertical (z) direction of M and G are aligned [20]. As a result, the transformation from G to M has only four degrees of freedom: a relative position between the respective origins and a relative yaw. Parameterizing orientation of M relative to G using a single degree-of-freedom quaternion (representing a rotation about the common gravity direction) allows one to write

$${}^M_G\hat{\mathbf{q}} = \begin{bmatrix} 0 & 0 & \frac{\theta}{|\theta|} \sin \frac{|\theta|}{2} & \cos \frac{|\theta|}{2} \end{bmatrix}^\top \approx \begin{bmatrix} 0 & 0 & \frac{\delta\theta}{2} & 1 \end{bmatrix}^\top \otimes {}^M_G\hat{\mathbf{q}}$$

and thus implement any updates to the estimate via rotations about the z -axis.

A. Filter Propagation

Let ${}^l\boldsymbol{\omega}$ and ${}^l\mathbf{a}$ be angular velocity and local linear acceleration of the local IMU frame l, while ${}^l\boldsymbol{\omega}_m$ and ${}^l\mathbf{a}_m$ denote the *measurements* of these quantities, respectively. Let also \mathbf{n}_w and \mathbf{n}_a represent continuous-time white Gaussian noise vectors that corrupt respective measurements, and express the measurement model as

$${}^l\boldsymbol{\omega}_m = {}^l\boldsymbol{\omega} + \mathbf{b}_w + \mathbf{n}_w, \quad {}^l\mathbf{a}_m = {}^l\mathbf{a} + {}^l\mathbf{R}^G \mathbf{g} + \mathbf{b}_a + \mathbf{n}_a$$

Measurement biases \mathbf{b}_w and \mathbf{b}_a are modeled as continuous-time random walks, driven by Gaussian white noises \mathbf{n}_{bw} and \mathbf{n}_{ba} . The estimator state components evolve continuously as

$$\begin{aligned} \dot{\mathbf{b}}_w &= \mathbf{n}_{bw} & \dot{\mathbf{b}}_a &= \mathbf{n}_{ba} & \dot{\mathbf{x}}_{\text{static}} &= \mathbf{0} & \dot{{}^l_G\hat{\mathbf{q}}} &= \frac{1}{2}\boldsymbol{\Omega}(\boldsymbol{\omega}) {}^l_G\hat{\mathbf{q}} \\ \dot{{}^G\hat{\mathbf{p}}} &= {}^G\hat{\mathbf{v}} & \dot{{}^G\hat{\mathbf{v}}} &= {}^l_G\mathbf{R}^\top {}^l\mathbf{a} & \boldsymbol{\Omega}(\boldsymbol{\omega}) &= \begin{bmatrix} -\hat{\boldsymbol{\omega}} & \mathbf{0} \\ \mathbf{0} & \hat{\boldsymbol{\omega}}^\top \end{bmatrix} \end{aligned}$$

The estimator assumes a random model for the target, driven by a Gaussian noise vector, \mathbf{n}_T , i.e., ${}^G\dot{\mathbf{p}}_T^{(i)} = {}^G\mathbf{p}_T^{(i+1)}$ and ${}^G\dot{\mathbf{p}}_T^{(n)} = \mathbf{n}_T$, giving rise to a (stochastic) linear model—involving a $(3n+3) \times 3$ evolution matrix \mathbf{H} having all zeros but for the bottom-right 3×3 block entry, which takes the form of an identity matrix, and a state Jacobian matrix \mathbf{B} , so that the total target state evolves continuously as $\dot{\mathbf{T}} = \mathbf{B}\mathbf{T} + \mathbf{H}\mathbf{n}_T$. Applying the expectation operator on the continuous evolution equations for the IMU and target states, yields the dynamics of the states' *estimates*,¹

$${}^l_G\dot{\hat{\mathbf{q}}} = \frac{1}{2}\boldsymbol{\Omega}({}^l\boldsymbol{\omega}_m - \hat{\mathbf{b}}_w) {}^l_G\hat{\mathbf{q}}, \quad {}^G\dot{\hat{\mathbf{v}}} = {}^l_G\hat{\mathbf{R}}^\top ({}^l\mathbf{a}_m - \hat{\mathbf{b}}_a) + {}^G\mathbf{g}$$

$${}^G\dot{\hat{\mathbf{p}}} = {}^G\hat{\mathbf{v}}, \quad \dot{\hat{\mathbf{b}}}_w = \mathbf{0}, \quad \dot{\hat{\mathbf{b}}}_a = \mathbf{0}, \quad \dot{\hat{\mathbf{x}}}_{\text{static}} = \mathbf{0}, \quad \dot{\hat{\mathbf{T}}} = \mathbf{B}\hat{\mathbf{T}}$$

These above equations allow efficient propagation of the state estimate by analytical solution of the differential equations across the measurement time-interval $[t_k, t_{k+1}]$ between measurements k and $k+1$.

Let \mathbf{F} and \mathbf{G} be the Jacobians of the IMU state with respect to itself and the IMU noises respectively. We say the target state is ℓ -dimensional while the static state is s -dimensional. Let $\mathbf{n}_{\text{IMU}} = [\mathbf{n}_w^\top \mathbf{n}_a^\top \mathbf{n}_{bw}^\top \mathbf{n}_{ba}^\top]^\top$ denote the stacked vector of IMU noise. The linearized dynamics now takes the form

$$\begin{bmatrix} \delta\dot{\mathbf{x}}_{\text{IMU}} \\ \delta\dot{\mathbf{T}} \\ \delta\dot{\mathbf{x}}_{\text{static}} \end{bmatrix} \approx \begin{bmatrix} \mathbf{F} & \mathbf{0}_{15 \times \ell} & \mathbf{0}_{15 \times s} \\ \mathbf{0}_{\ell \times 15} & \mathbf{B} & \mathbf{0}_{\ell \times s} \\ \mathbf{0}_{s \times 15} & \mathbf{0}_{s \times \ell} & \mathbf{0}_{s \times s} \end{bmatrix} \begin{bmatrix} \delta\mathbf{x}_{\text{IMU}} \\ \delta\mathbf{T} \\ \delta\mathbf{x}_{\text{static}} \end{bmatrix} + \begin{bmatrix} \mathbf{G} & \mathbf{0}_{12 \times 3} \\ \mathbf{0}_{\ell \times 12} & \mathbf{H} \\ \mathbf{0}_{s \times 12} & \mathbf{0}_{s \times 3} \end{bmatrix} \begin{bmatrix} \mathbf{n}_{\text{IMU}} \\ \mathbf{n}_T \end{bmatrix} \quad (9)$$

¹Notation $\hat{\cdot}$ for estimates should not be confused with the slightly wider differential geometric $\hat{\cdot}$ wedge mapping.

Let $\Phi(t_{k+1}, t_k)$ and $\mathbf{A}(t_{k+1}, t_k)$ be the discrete-time state transition matrices of the IMU and target based on (9). Matrices \mathbf{Q}_{IMU} and \mathbf{Q}_T henceforth denote the noise covariances for the IMU and target state evolution, associated with discrete-time noise characterization [19], [21]. Letting $\text{Diag}(\cdot, \cdot, \cdot)$ refer to the function which places the argument matrices on the block diagonal of an otherwise zero matrix, propagation of the covariance \mathbf{P} is given by:

$$\begin{aligned} \Psi &= \text{Diag}(\Phi(t_{k+1}, t_k), \mathbf{A}(t_{k+1}, t_k), \mathbf{I}_{s \times s}) \\ \mathbf{Q}_k &= \text{Diag}(\mathbf{Q}_{\text{IMU}}, \mathbf{Q}_T, \mathbf{0}_{s \times s}), \quad \mathbf{P}_{k+1} = \Psi \mathbf{P}_k \Psi^\top + \mathbf{Q}_k \end{aligned}$$

B. Filter Update

Define the projection operator $\Pi([x \ y \ z]^\top) \triangleq \begin{bmatrix} x/z \\ y/z \end{bmatrix}$, which is involved in mapping the global position vector ${}^G\mathbf{p}_f$ of an environment feature to the camera image plane. Considering the camera frame C at measurement step k , and letting ${}^C\mathbf{R}$ and ${}^C\mathbf{p}_l$ be the known transformation between the IMU and camera, the (noisy) camera measurement is

$$\mathbf{z}_{f_k} = \Pi\left({}^C\mathbf{R} {}^l_k\mathbf{R} ({}^G\mathbf{p}_f - {}^G\mathbf{p}_{l_k}) + {}^C\mathbf{p}_l\right) + \mathbf{n}_z$$

where \mathbf{n}_z is the imaging measurement noise vector. In addition, we can write the measurements with respect to the *target* by replacing ${}^G\mathbf{p}_f$ in the above equation with the position of the target, ${}^G\mathbf{p}_T$. Similarly, for a known landmark ${}^M\mathbf{L}$,² one would write

$$\mathbf{z}_{L_k} = \Pi\left({}^C\mathbf{R} {}^l_k\mathbf{R} ({}^G\mathbf{R}^\top {}^M\mathbf{L} + {}^G\mathbf{p}_M - {}^G\mathbf{p}_{l_k}) + {}^C\mathbf{p}_l\right) + \mathbf{n}_z$$

While the feature measurements with respect to known landmarks and the target are processed through the standard Extended Kalman Filter (EKF) update, linear marginalization through the MSCKF update step [1] is performed on the feature measurements, otherwise a standard EKF storing all the measured features in the state vector would lead to unbounded computation. Collect all camera measurements and corresponding noises in stacked vectors \mathbf{z} and \mathbf{n} respectively, and note that if instead of the true estimator state \mathbf{x} one uses its estimate $\hat{\mathbf{x}}$, the expressions above yield a camera measurement estimate, $\hat{\mathbf{z}}$. If the camera measurement model is linearized, then the camera measurement *residual* $\tilde{\mathbf{z}} \triangleq \mathbf{z} - \hat{\mathbf{z}}$ is expressed in terms the estimator state error $\delta\mathbf{x}$, the feature position error $\delta\mathbf{p}_f$, and the Jacobians \mathbf{H}_x , \mathbf{H}_f , with respect to the estimator state \mathbf{x} and feature position \mathbf{p}_f :

$$\tilde{\mathbf{z}} = \mathbf{H}_x \delta\mathbf{x} + \mathbf{H}_f \delta\mathbf{p}_f + \mathbf{n} \quad (10)$$

QR decomposition on \mathbf{H}_f gives a unitary matrix \mathbf{Q}_2 whose columns span the left nullspace of \mathbf{H}_f . If \mathbf{C} denotes the covariance of \mathbf{z} and one lets $\mathbf{C}' \triangleq \mathbf{Q}_2^\top \mathbf{C} \mathbf{Q}_2$ and $\mathbf{n}' \sim \mathcal{N}(\mathbf{0}, \mathbf{C}')$, this projection *eliminates* the dependency on feature errors of the new residual

$$\tilde{\mathbf{z}}' \triangleq \mathbf{Q}_2^\top \tilde{\mathbf{z}} = \mathbf{Q}_2^\top \mathbf{H}_x \delta\mathbf{x} + \mathbf{Q}_2^\top \mathbf{n} \quad (11)$$

²Assumed perfectly known without noise. Noise can be included by using the Cholesky-Schmidt-Kalman Filter proposed in [20].

C. Initialization

1) *Map Initialization:* To initialize a parameter not yet contained in the state, it is required to find an initial estimate of the new variable and compute Jacobians for a set of initializing measurements with respect to the current state and the new variable respectively [20]. In particular, bearing measurements to known landmarks are collected until the transformation between the global G and prior map M frames can be estimated. An estimate of the landmarks in the IMU's global frame, ${}^G\mathbf{L}_f$, can be recovered through multi-view triangulation. Using the positions for the landmarks expressed in both frames, the 6 DOF relative pose is estimated [22], from which the relative position and yaw are extracted. Variable initialization can be performed by separating the Jacobians of the collected landmark bearing measurements with respect to the current state and the new map parameters.

2) *Target Initialization:* Let ${}^G\mathbf{p}_{T_0}$ denote the target's position at the time t_0 that it is sighted by the UAV, and d , b denote the depth and bearing of the target in the camera's image. In general, d and ${}^G\mathbf{p}_{T_0}^{(i)}$ are unknown, but if one collects N consecutive measurements of the target's after the target's initial sighting and assumes a constant derivative model over the interval $t_j - t_0$ of time elapsed since the first sighting and a subsequent target measurement j , a constraint for each subsequent measurement can be stated about the position of the camera in the global frame for $j \in \{1, \dots, N\}$

$${}^G\mathbf{p}_{C_j} = -d_j b_j + \sum_{i=0}^n \frac{(t_j - t_0)^i}{i!} {}^G\mathbf{p}_{T_0}^{(i)} \quad (12)$$

Construct $\mathbf{p} = [{}^G\mathbf{p}_{C_1}^T, \dots, {}^G\mathbf{p}_{C_N}^T]^T$ by stacking all the relevant camera poses, and recall that \mathbf{T}_0 denotes the initial target state. Similarly, let \mathbf{d} to be the stack vector of all target depths d_j , and \mathbf{Y} be the matrix encoding the constraints (12). The least-squares estimates of $\hat{\mathbf{T}}_0$ and $\hat{\mathbf{d}}$ are: $\mathbf{p} = \mathbf{Y} \begin{bmatrix} \mathbf{T}_0 \\ \mathbf{d} \end{bmatrix} \implies \begin{bmatrix} \hat{\mathbf{T}}_0 \\ \hat{\mathbf{d}} \end{bmatrix} = (\mathbf{Y}^T \mathbf{Y})^{-1} \mathbf{Y}^T \mathbf{p}$. Propagating $\hat{\mathbf{T}}_0$ using the discrete-time target model $\mathbf{T}_{k+1} = \mathbf{A}(t_{k+1}, t_k) \mathbf{T}_k + \mathbf{n}_{TD}$ with discrete-time noise characterization $\mathbf{n}_{DT} \sim \mathbf{N}(\mathbf{0}, \mathbf{Q}_T)$, now yields estimates of the target state $\{\hat{\mathbf{T}}_0, \hat{\mathbf{T}}_1, \dots, \hat{\mathbf{T}}_N\}$ at each measuring time, as well as a constraint that, when used in combination with the collected bearing measurements, can be used as the initializing measurements for the target state.

IV. SIMULATIONS

The detection scenario consists of a simulated workspace of $20\text{ m} \times 20\text{ m} \times 5\text{ m}$ centered at $(0, 0, 2.5\text{ m})$ having three squircle shaped obstacles shown in Fig. 2a. A Firefly UAV equipped with a 20° downward facing VI-sensor intercepts a target (turtlebot, shown at the bottom of Fig. 2a). Figure 2b shows a section of the time-varying navigation function at an instance where the target is at $(7.5\text{ m}, 5\text{ m}, 0.5\text{ m})$, with parameters $\kappa = 6$, $\lambda = 10^4$ and $\lambda_{\text{sq}} = 10^4$.

A dense set of 1550 unknown features and a sparse set of 160 known landmarks are randomly placed in the environment. Noisy IMU measurements (with noise density 1×10^{-6} for accelerometer and gyroscope both) are simulated at 200 Hz while the bearing measurements are simulated at

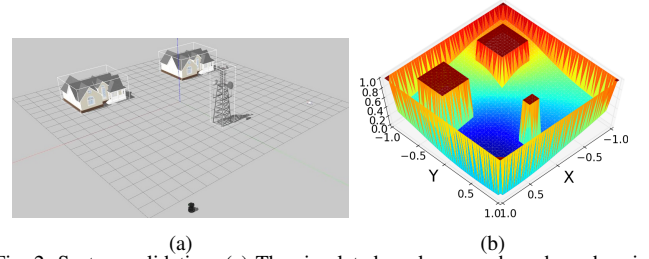


Fig. 2: System validation: (a) The simulated workspace whose boundary is not shown for clarity. (b) Navigation function constructed on unit squircle world in the plane containing the target at location $(7.5\text{ m}, 5\text{ m}, 0.5\text{ m})$. Variation across different heights can be seen in the attached video.

a rate of 20 Hz and are corrupted by Gaussian noise of one pixel standard deviation. The MSCKF window is of size 8. Occlusions are simulated by checking whether each bearing ray intersects with obstacle boundaries.

The target is assumed to carry an isotropic Cf-252 neutron source, and the UAV picks emitted neutrons using a Domino thermal neutron detector³. Radioactivity is simulated using the thinning algorithm [23]—benchmarked against a $5\ \mu\text{Ci}$ (micro Curie) source available—to guide the simulated realization of higher activity sources. (The activity of a $5\ \mu\text{Ci}$ source, 1.5 m away from the target essentially blends completely into background.)

The UAV behavior consists of three phases: (I) hovering temporarily to initialize tracking, (II) intercepting the target guided by the time-varying navigation function, and (III) return to hovering once the target exits the detection area. Figure 3 depicts two instances of interception phase (II). The target moves in a straight line with a velocity of 1 m/s and the UAV intercepts it with a maximum speed of 6 m/s (Fig. 3a). Once Intercepted, the UAV follows the target maintaining a constant distance of about 0.7 m (Fig. 3b).




Fig. 3: Stages of Target Interception for the setup shown in Fig. 2: (a) High speed Interception (b) Chase from a fix distance. Both maneuvers are guided by single time-varying navigation function.

Figure 4a shows the true and estimated relative distance between the UAV and the target. Starting high during tracking initializing phase, it eventually converges to approximately 15–20 mm. Figures 4b and 4c show the components of the relative position error between estimate and ground truth (simulated using Gazebo's odometry sensor), of the UAV (self position estimates) and the target, respectively. The estimation error in the target's position starts with a high value due to initialization at a distance, but eventually converges to acceptable values. The UAV state estimation error, due to informative known landmark bearing measurements, remains bounded within 10 cm over the whole path. Figure 5a presents the desired and achieved velocities of the UAV in all three phases, while Fig. 5b shows positions of the UAV

³Radiation Detection Technologies Inc <http://radectech.com/products/rtd-domino-v5-4>

and the target only in the interception phase (II) (as they are constant in I and III), during which synthetic radiation measurements are produced to simulate target radioactivity.




Fig. 4: (a) Estimated and Ground-truth values of Relative distance between the UAV and the target. (b) Error in VINS position estimates of the UAV. (c) Corresponding errors for the target.

Detection of weak radioactive material requires close proximity of sensor to target. To assess how close to the target the UAV should be, simulations were repeated with different safety bubble radii r . Figure 5c shows the variation of bound on PM with safety bubble radius. The bound on FA α is constrained at 0.001. The total time allocated to make the decision was capped at 20 sec while the experimentally obtained background radioactivity and detector's radiation cross-section used in the calculations were 0.005833 counts per second (CPS) and $2.12 \times 10^{-6} \text{m}^2$. Radiation sources of three different strengths 2.2×10^5 , 3.3×10^5 , and 4.4×10^5 CPS were simulated. The source activity is calculated assuming 4.4×10^9 neutrons/Curie/second from a Cf-252 source.⁴ A bound on probability of missed detection is set at 0.1 for when the UAV is within approximately 1-1.25m of the target—weak sources mandate close sensor proximity, justifying tightly coupling state estimation and target tracking.

V. CONCLUSIONS

An integrated motion-planning and estimation framework, which tightly couples visual-inertial navigation and target tracking implemented through a navigation function and a differential geometric controller, is shown to be appropriate for applications of UAV-based radiation detection. Such an integrated system offers a provable solution for guaranteed autonomous tracking (necessary for accurate detection) and collision avoidance, that can work without reliance on external infrastructure for positioning and tracking.

REFERENCES

- [1] A. I. Mourikis and S. I. Roumeliotis, "A multi-state constraint Kalman filter for vision-aided inertial navigation," in *Proceedings of the IEEE International Conference on Robotics and Automation*, Rome, Italy, Apr. 10–14, 2007, pp. 3565–3572.
- [2] M. Li and A. I. Mourikis, "Online temporal calibration for Camera-IMU systems: Theory and algorithms," *International Journal of Robotics Research*, vol. 33, no. 7, pp. 947–964, Jun. 2014.
- [3] D. Kottas, K. Wu, and S. Roumeliotis, "Detecting and dealing with hovering maneuverers in vision-aided inertial navigation systems," in *Proceedings of IEEE/RSSJ International Conference on Intelligent Robots and Systems*, Nov. 2013, pp. 3172–3179.
- [4] J. Hesch, D. Kottas, S. Bowman, and S. Roumeliotis, "Consistency analysis and improvement of vision-aided inertial navigation," *IEEE Transactions on Robotics*, vol. PP, no. 99, pp. 1–19, 2013.
- [5] C.-C. Wang, C. Thorpe, S. Thrun, M. Hebert, and H. Durrant-Whyte, "Simultaneous localization, mapping and moving object tracking," *The International Journal of Robotics Research*, vol. 26, no. 9, pp. 889–916, Sep. 2007.




Fig. 5: (a) Desired and achieved velocities in various phases of UAV motion. (b) Achieved positions of the UAV and the target in the chase phase (II). (c) Bound on Probability of missed detection with varying minimum distance between the robot and the target for a 20 seconds sensing window.

- [6] M. Chojnacki and V. Indelman, "Vision-based dynamic target trajectory and ego-motion estimation using incremental light bundle adjustment," *International Journal of Micro Air Vehicles*, vol. 10, no. 2, pp. 157–170, 2018.
- [7] K. Zhou and S. Roumeliotis, "Multirobot active target tracking with combinations of relative observations," *IEEE Transactions on Robotics*, vol. 27, no. 4, pp. 678–695, Aug. 2011.
- [8] J. Chen, T. Liu, and S. Shen, "Tracking a moving target in cluttered environments using a quadrotor," in *Proceedings of IEEE/RSSJ International Conference on Intelligent Robots and Systems*, Oct 2016, pp. 446–453.
- [9] R. J. Nemzek, J. S. Dreicer, D. C. Torney, and T. T. Warnock, "Distributed sensor networks for detection of mobile radioactive sources," *IEEE Transactions on Nuclear Science*, vol. 51, no. 4, pp. 1693–1700, Aug 2004.
- [10] C. D. Pahlajani, J. Sun, I. Poulakakis, and H. G. Tanner, "Error probability bounds for nuclear detection: Improving accuracy through controlled mobility," *Automatica*, vol. 50, no. 10, pp. 2470–2481, 2014.
- [11] J. Thomas, J. Welde, G. Loianno, K. Daniilidis, and V. Kumar, "Autonomous flight for detection, localization, and tracking of moving targets with a small quadrotor," *IEEE Robotics and Automation Letters*, vol. 2, no. 3, pp. 1762–1769, July 2017.
- [12] D. Mellinger and V. Kumar, "Minimum snap trajectory generation and control for quadrotors," in *Proceedings of IEEE International Conference on Robotics and Automation*, 2011, pp. 2520–2525.
- [13] S. Brennan, A. Mielke, and D. Torney, "Radioactive source detection by sensor networks," *IEEE Transactions on Nuclear Science*, vol. 52, no. 3, pp. 813–819, 2005.
- [14] J.-C. Chin, N. S. V. Rao, D. K. Y. Yau, M. Shankar, Y. Yang, J. C. Hou, S. Srivathsan, and S. Iyengar, "Identification of low-level point radioactive sources using a sensor network," *The ACM Transactions on Sensor Networks*, vol. 7, no. 3, pp. 21:1–21:35, 2010.
- [15] J. Sun and H. G. Tanner, "Constrained decision-making for low-count radiation detection by mobile sensors," *Autonomous Robots*, vol. 39, no. 4, pp. 519–536, 2015.
- [16] T. Lee, M. Leoky, and N. H. McClamroch, "Geometric tracking control of a quadrotor uav on se(3)," in *Proceedings of 49th IEEE Conference on Decision and Control*, 2010, pp. 5420–5425.
- [17] C. Li and H. Tanner, "Navigation functions with time-varying destination manifolds in star worlds," *IEEE Transactions on Robotics*, vol. PP, pp. 1–14, 10 2018.
- [18] E. Rimon and D. E. Koditschek, "Exact robot navigation using artificial potential functions," *IEEE Transactions on Robotics and Automation*, vol. 8, no. 5, pp. 501–518, Oct 1992.
- [19] N. Trawny and S. I. Roumeliotis, "Indirect Kalman filter for 3D attitude estimation," University of Minnesota, Department of Computer Science & Engineering, Tech. Rep., Mar. 2005.
- [20] R. C. DuToit, J. A. Hesch, E. D. Nerurkar, and S. I. Roumeliotis, "Consistent map-based 3d localization on mobile devices," in *Proceedings of IEEE International Conference on Robotics and Automation*, May 2017, pp. 6253–6260.
- [21] P. S. Maybeck, *Stochastic Models, Estimation and Control*, ser. Mathematics in Science and Engineering. London: Academic Press, 1982, vol. 141-2.
- [22] B. K. P. Horn, "Closed-form solution of the absolute orientation using quaternions," *Journal of Optical Society of America A*, vol. 4, pp. 629–642, Apr. 1987.
- [23] R. Pasupathy, "Generating nonhomogeneous Poisson processes," in *Wiley Encyclopedia of Operations Research and Management Science*. Wiley, 2009.

⁴<http://www.logwell.com/tech/nuclear/Californium-252.html>.