Propagation of Heterogeneous Receiver Self-Positioning Uncertainty to Target Localisation and Tracking in Airborne Multistatic Passive Radar

Bartlomiej M. Dudek[1]
Supervisor(s): Dr. Arash Asadi[1]**, Florian Kosterhon**[1]
[1]EEMCS, Delft University of Technology, The Netherlands

A Thesis Submitted to EEMCS Faculty Delft University of Technology,
In Partial Fulfilment of the Requirements
For the Bachelor of Computer Science and Engineering
June 21, 2026

Name of the student: Bartlomiej Mariusz Dudek
Final project course: CSE3000 Research Project
Thesis committee: Dr. Arash Asadi, Florian Kosterhon, Dr. Georgios Iosifidis

***Abstract*—This work investigates an airborne (UAV-based) multistatic passive radar using signals of opportunity in the context of a GNSS-denied environment. Surveyed works on multistatic passive radars either assume zero receiver positioning error, or otherwise ignore its heterogeneity across receivers, which we argue is expected in GNSS-denied navigation due to varying conditions. Rui and Ho's error model accommodates heterogeneous per-receiver errors but is evaluated only for the IID case. We adapt this model and study the heterogeneous regime it permits with a weighted least-squares (WLS) solver, and operationalise their deferred drifting-receiver case via an EKF tracker with per-receiver drift state. We run a Monte Carlo simulation with 10,000 trials per configuration varying the number of receivers and receiver positioning error parameters to compare a per-receiver-error-aware solver against a solver assuming IID error across receivers, and a solver assuming no self-positioning error at all. The findings show that, relative to the IID solver, the per-receiver-aware solver reduces the median target-position error by $\sim$13% in snapshot WLS (at N=10, $\sigma_{\mathbf{Rx}}$=5 m) and the median steady-state target-position RMSE by $\sim$28% in EKF tracking (at N=10, $\sigma_{\mathbf{drift}}$=0.1 m/$\sqrt{\mathbf{s}}$). This suggests that in a UAV-based multistatic passive radar, per-receiver self-positioning error is a non-negligible parameter to consider in design.**

***Index Terms*—airborne passive radar, UAV swarms, GNSS-denied navigation, multistatic passive radar, signals of opportunity, illuminators of opportunity, bistatic range**

I. Introduction

GNSS is the basis of modern navigation, enabling precise positioning for a wide range of applications, including autonomous UAVs. In recent years, GNSS jamming has become increasingly common, posing a significant threat to the reliability of GNSS-based navigation. Under GNSS-denied conditions, UAVs must rely on alternative methods to navigate effectively. Alternatives such as LiDAR or optical sensors are short-range, so they may not be sufficient in the open field. Utilising active radar or anchors could be an option, but they may be easily detected by adversaries, which may be expected since GNSS jamming is often used in an adversarial context.

To self-localise in such scenarios, UAVs may have to rely on passive sensing, which may involve dead reckoning from sensor readouts (IMU, gyroscope, compass) [1], or exploiting Illuminators of Opportunity (IoO): ambient radio emitters such as cellular networks, FM radio, or TV broadcasts [2], [3]. Such self-localisation methods are far from perfect, and the resulting position estimates are expected to have some error.

While Illuminators of Opportunity may already support navigation in terms of passive self-localisation, another important advantage they bring is that they can be used to detect, locate, and track targets. A system exploiting IoOs in this way is called a passive bistatic radar (PBR) [4]. Such a system works by listening to signals reflected off a target.

Just as the fixed distance between cameras in stereo vision enables depth perception, combined radar observations from multiple locations can synthesise spatial information. With multiple radar nodes, the position of a target can be resolved, unlike with a single node, which is the core idea behind a multistatic radar (Fig. 1) [5, Sec. 1.1]. Adding more radar nodes overdetermines the system, leading to improved target localisation accuracy. Existing works explore multistatic passive radars designed to detect both emitting targets [6] and silent targets via IoOs [7], [8].

Studies have experimentally demonstrated the feasibility of target tracking with airborne passive bistatic radars [4], [16], [17]. However, little work extends the concept of airborne passive bistatic radar to include multiple airborne receivers. Sagduyu et al. [9] have recently proposed a 5G-based multistatic passive radar system for ground-based receivers. UAVs, however, are not static objects and are prone to vibrations and drift, and their relative location in a multistatic configuration cannot be known perfectly, especially in GNSS-denied environments, where receivers must fall back on passive self-localisation. This error is expected to propagate through every receiver's measurement, ultimately affecting target localisation and tracking performance.

Rui and Ho [13] have proposed an estimator for target localisation in a multistatic passive radar system, which includes self-positioning error in the model. However, it has only been evaluated under the assumption that the receiver self-positioning error is independent and identically distributed (IID) across receivers. We claim that in challenging environments, self-positioning error may be heterogeneous across receivers, resulting from factors such as varying local conditions and sensor quality, and that incorporating per-receiver self-positioning awareness into the target estimator can improve localisation and tracking performance.

We explore tracking of an uncooperative drone in open, GNSS-denied airspace. A single, distant, uncooperative transmitter is the illuminator of opportunity. All receivers (N ranging from 3 to 10) are UAV-mounted, quasi-static, and experience some level of self-positioning error. We do not model the self-localisation process itself; instead, we model its output as a per-receiver position error, heterogeneous across receivers. Each receiver reports a bistatic range measurement, together with an estimate of its own position and the associated uncertainty. The measurements are taken at the same instant and combined centrally to estimate the target position. The target is in motion with fixed velocity, and we track it over time using a Kalman filter.

The main contributions of this work are as follows: modelling how receiver positioning uncertainty propagates to target localisation error in multistatic passive radar systems; and evaluating a per-receiver self-positioning awareness approach, based on the error model of [13] and tested under the heterogeneous and drifting errors that study leaves to future work, showing it reduces target positioning error in our system model.

The remainder of this paper is organised as follows. Sec. II surveys related work on multistatic and airborne passive radar. Sec. III formalises the system model. Sec. IV introduces the estimator and the proposed per-receiver weighting and tracking extension. Sec. V presents the experiments and results, and Sec. VI interprets them. Sec. VII addresses responsible research, and Sec. VIII concludes.

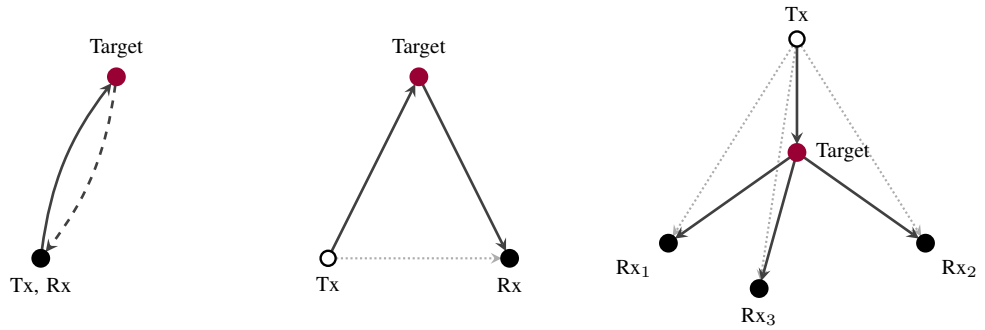

(a) Monostatic (b) Bistatic (c) Multistatic (M=1, N=3)

Fig. 1. Radar architectures. (a) Monostatic: transmitter and receiver co-located; (b) Bistatic and (c) multistatic ($M = 1$, $N = 3$): the transmitter (empty circle) is an Illuminator of Opportunity, and the receivers (filled) are spatially distributed. Solid arrows show the indirect path via the target, while dotted arrows show the direct path used as a reference.

TABLE I
LITERATURE TAXONOMY OF RELATED MULTISTATIC AND BISTATIC PASSIVE RADAR STUDIES.

Paper	**Modality**	**Architecture**	**Airborne Rx**	**Rx-pos unc.**	**IoO**	**Evidence**
Sagduyu et al. [9]	Bistatic Range, Doppler	multistatic	×	×	5G	Simulation
Malanowski & Kulpa [10]	Bistatic Range	bistatic, multistatic	×	×	generic	Simulation and Experiment
Noroozi & Sebt [11]	Bistatic Range	multistatic	×	×	generic	Simulation
Ahmed et al. [12]	Doppler	multistatic	✓	×	DVB-T	Simulation
Rui & Ho [13]	Bistatic Range	multistatic	×	✓(IID)	generic	Simulation
Brown et al. [14]	Bistatic Range, Doppler	bistatic	✓	×	FM	Experiment
Souli et al. [15]	Bistatic Range, Doppler	bistatic	✓	×	DVB-T	Experiment
This work	Bistatic Range	multistatic	✓	✓	generic	Simulation

II. RELATED WORK

Multistatic passive radar exploiting signals of opportunity is not a new concept: the Klein Heidelberg radar used an ambient signal as early as 1943 [18, §2.1.6]. This section focuses on recent works that bring the architecture to airborne context and works utilising modern illuminators of opportunity. The literature is grouped into two clusters (Table I): (A) multistatic passive radar localisation methods, which fuse bistatic measurements into a target position with fixed receivers; and (B) airborne implementations of passive radar, where the receivers are mounted on airborne platforms.

A. Multistatic passive radar localisation methods

There are multiple solutions to determine target position by synthesising measurements from several bistatic pairs. Being passive, these systems rely on a non-cooperative transmitter and differ in whether they treat its signal generically or exploit a specific waveform. Malanowski and Kulpa [10] introduce two algorithms for localisation based on known methods for time difference of arrival (TDOA) systems: spherical interpolation (SI) and spherical intersection (SX). Both rely on closed-form equations. The system enables 3D localisation in a configuration with multiple IoOs and a single receiver, which was verified both in simulation and, uniquely in this cluster, with a real-life FM-based experiment.

Noroozi and Sebt [11] introduce a closed-form weighted least squares (WLS) solver for a localisation setup with multiple IoOs and multiple receivers. The weighting matrix reduces propagation of noise by downweighting noisy bistatic range measurements, which may bring an advantage over the non-weighted LS approach found in [10]. The method does not exploit features of a particular signal of opportunity (SoO), in contrast to Sagduyu's 5G-oriented design.

Sagduyu et al. [9] introduce an Integrated Sensing and Communication (ISAC) configuration based on 5G which consists of a single IoO and multiple receivers, with an explicitly defined cross-ambiguity function. The estimator, in contrast to Noroozi and Sebt's WLS, is non-linear LS. The work shows that Extended Kalman Filter (EKF) tracking outperforms the Kalman Filter (KF) in filtering error by a factor of 13 for a circular trajectory, while KF performs marginally better on linear trajectories. The system is validated in simulation in 2D and 3D for both single target and multiple targets, for both localisation and tracking, under a favourable radar geometry.

Rui and Ho [13] demonstrate a multistatic target position estimator which incorporates Gaussian transmitter and receiver position error. While their model can represent heterogeneous receiver position error, it is evaluated only for the IID case.

The works listed demonstrate a number of solutions for multistatic passive radar localisation. Only Rui and Ho [13] explicitly model receiver position error, and they defer both non-IID errors and the drifting-receiver case to future work. The effect of non-IID receiver position error and of drift over

time on multistatic passive radar localisation performance is the focus of this work.

B. Airborne implementations of passive radar

Brown [14], [16] first demonstrated an airborne passive bistatic radar detecting aircraft using IoO (FM). The author explicitly identifies the single bistatic pair geometry as the key limitation and proposed multiple transmitters for future work.

Ahmed et al. [12] have compared two topologies for a UAV-based multistatic passive radar tracking ground targets in simulation. While this is a close architecture match to our problem, this work is more focused on optimising for reduced communication rather than receiver capabilities and assumes perfect receiver positioning.

Souli et al. [15] experimentally demonstrate a UAV-carried bistatic passive radar for detecting other UAVs which confines the target to an ellipse (detection) and utilises computer vision to further confine the visible target to a point (disambiguation). The radar component is based on software-defined radio (SDR) and is capable of employing various IoOs.

These works demonstrate a foundation for airborne PBRs detecting airborne targets including UAVs. For a single receiver, its measurements are useful in its own relative frame, which reduces the significance of receiver positioning error in that context. For a multi-receiver radar, fusion and disambiguation of measurements require a common geometric frame. Per-receiver position error propagates into fused estimates. We investigate this effect in detail.

No prior work addresses multiple airborne receivers whose positions are uncertain, with error that is heterogeneous across receivers and drifts over time, as we would expect in a non-uniform GNSS-denied environment when relying on passive self-localisation. This work studies how that self-positioning error propagates into target localisation and tracking, and whether per-receiver uncertainty awareness can mitigate it better than modelling the error as IID across receivers or ignoring it altogether.

III. System Model

A. Scenario and assumptions

The setup explored is a multistatic passive radar with one fixed, uncooperative IoO transmitter of known position and multiple UAV-based, quasi-static receivers. A single transmitter is chosen to keep complexity minimal. The selected receiver topology is a ring, so that every added receiver is geometrically equivalent, and the effect of N is not confounded with placement. Its radius $r = 1$ km is small relative to the 3 km target area and 5 km transmitter distance, modelling a compact swarm observing a more distant scene. A single target follows a constant-velocity linear trajectory at $10\,\mathrm{m/s}$ within a 3 km-radius disc centred on the receiver ring. The system model does not include any communication between the receivers. Each receiver reports its bistatic-range measurement, its self-position estimate, and the associated self-positioning uncertainty to a central node, which fuses them to localise the target. The system model does not restrict the type of IoO and does not define the self-localisation method; it only characterises the resulting error. The system is visualised in Fig. 2.

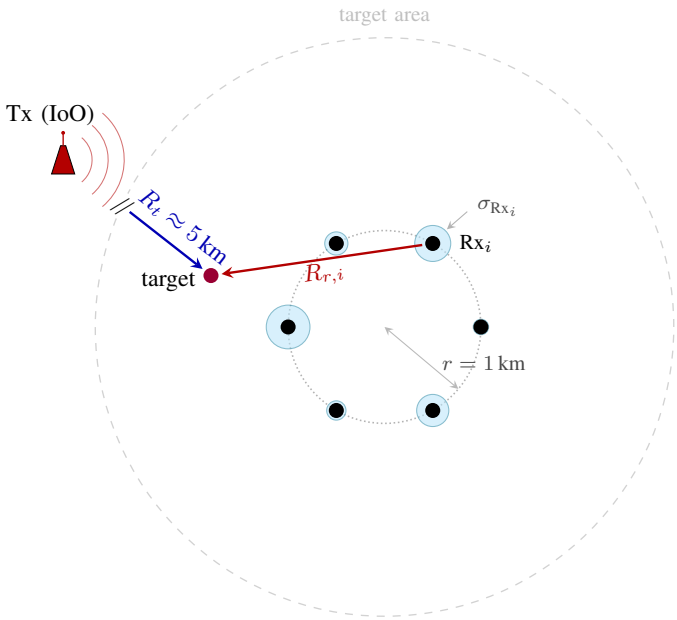

Fig. 2. Example setup explored in this work with receiver position uncertainty as cyan rings

B. Bistatic and multistatic radar

With a single source of signal of opportunity, the signal reaches the receiver through a direct path, and an indirect path, where the signal reflects off the target and only then reaches the receiver. The difference in time of arrival is used to confine the location of the target to a bistatic range. Such a setup (single transmitter of IoO, single receiver) will be referred to as a bistatic passive radar [18].

A generalisation of this setup is multistatic radar consisting of $M \geq 1$ transmitters and $N \geq 1$ receivers, all of which are spatially distributed (unlike in monostatic radar, where transmitter and receiver are co-located). For M transmitters and N receivers, there are MN bistatic pairs. The three architectures are compared in Fig. 1.

C. Bistatic range

The bistatic range B_i is the sum of the propagation distances from the transmitter to the target and from the target to receiver i:

$$B_i = \|\mathbf{x}_{\mathrm{Tx}} - \mathbf{x}_T\| + \|\mathbf{x}_T - \mathbf{x}_{\mathrm{Rx}_i}\|, \tag{1}$$

where B_i is a scalar bistatic range for receiver i, and $\mathbf{x}_{\mathrm{Tx}}$, $\mathbf{x}_T$, and $\mathbf{x}_{\mathrm{Rx}_i}$ are position vectors in $\mathbb{R}^2$ denoting the locations of the transmitter, target, and receiver i, respectively.

Given a bistatic range B_i, the target can be confined to an ellipse E_i (Fig. 3) with foci at Tx and Rx_i, such that Eq. (1) is satisfied. The target lies on the intersection of the

per-receiver ellipses, so in a 2D coordinate system three non-collinear bistatic pairs are required to locate it unambiguously. In 3D, four non-coplanar bistatic pairs are necessary and each E_i is an ellipsoid rather than a planar ellipse.

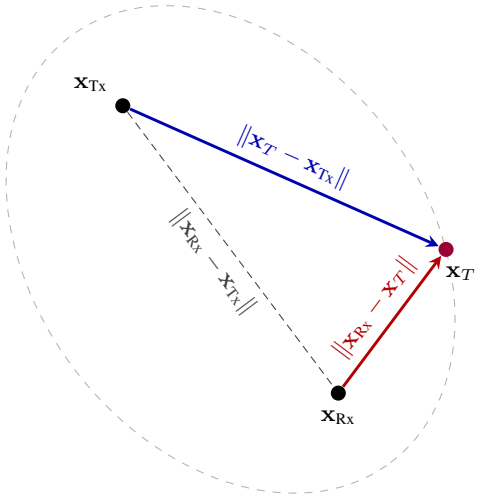

Fig. 3. Bistatic geometry. The bistatic range $B_i = \|\mathbf{x}_T - \mathbf{x}_{\mathrm{Tx}}\| + \|\mathbf{x}_{\mathrm{Rx}} - \mathbf{x}_T\|$ is the sum of Tx-target and Rx-target distances; the grey dashed ellipse with foci at $\mathbf{x}_{\mathrm{Tx}}$ and $\mathbf{x}_{\mathrm{Rx}}$ is the iso-range locus.

The output of a multistatic radar signal processing pipeline is a cross-ambiguity function (CAF), the correlation between the reference (direct-path) and surveillance (target-reflected) signals. The global peak of $\mathrm{CAF}(d, f)$ at (d_0, f_0) corresponds to the estimated bistatic range $\hat{B}_i = d_0\, c/f_s$, where c is the speed of light and f_s the sampling rate at the receiver. We adopt this CAF definition from [9]. The standard deviation of this measurement, $\sigma_{\hat{B}_i}$, is described next.

D. Bistatic range noise

One source of error is the bistatic range uncertainty $\sigma_{\hat{B}_i}$, the precision with which the CAF peak yielding $\hat{B}_i$ can be located. This precision is lower-bounded by the Cramér–Rao lower bound (CRLB) for range estimation. To keep the model simple, we derive a minimal distance-based expression for $\sigma_{\hat{B}_i}$ from the bistatic radar equation [18, Sec. 4.1]. The signal-to-noise ratio (SNR) of receiver i is

$$\mathrm{SNR}_i \;=\; \frac{P_t\, G_t\, G_r\, \lambda^2\, \sigma_{\mathrm{RCS}}\, T_{\mathrm{CPI}}}{(4\pi)^3\, R_t^2\, R_{r,i}^2\, N_0} \tag{2}$$

Here the numerator combines the standard radar-equation terms: transmit power P_t, antenna gains G_t and G_r, wavelength λ, radar cross-section σ_{RCS}, and coherent processing interval T_{CPI}; N_0 is the noise power spectral density. We assume that only the range factors R_t and $R_{r,i}$ vary across receivers and trials, giving the proportionality:

$$\mathrm{SNR}_i \;\propto\; \frac{1}{R_t^2\, R_{r,i}^2}. \tag{3}$$

The CRLB is related to the SNR as follows [18, Sec. 5.2]:

$$\sigma_{\hat{B}_i} \;\propto\; \frac{1}{\sqrt{\mathrm{SNR}_i}}, \tag{4}$$

We introduce reference uncertainty $\sigma_{B,\mathrm{ref}}$ scaling with SNR, together with a floor $\sigma_{B,\min}$ for SNR-independent receiver noise (e.g., clock and ADC jitter) not captured by the CAF peak precision. Combining equations (3) and (4) gives the following model for measurement uncertainty as a function of target and receiver range:

$$\sigma_{\hat{B}_i} \;=\; \max\!\left(\sigma_{B,\mathrm{ref}} \cdot \frac{R_t}{R_{t,\mathrm{ref}}} \cdot \frac{R_{r,i}}{R_{\mathrm{ref}}},\; \sigma_{B,\min}\right) \tag{5}$$

E. Receiver self-positioning error

In Section II we have found that most works treat receiver position as perfectly known, which is sensible for fixed radars. In works where Rx position was uncertain, such as [13], the error was assumed to be IID across receivers. This is a reasonable model for usual applications where the conditions for self-positioning are uniform, for example, GNSS reception in an open field. Since the receivers in this work rely on passive self-localisation, whose quality varies across the swarm, the self-positioning error distribution will differ per receiver. We define an additive Gaussian noise model on Rx position.

$$\hat{\mathbf{x}}_{\mathrm{Rx}_i} = \mathbf{x}_{\mathrm{Rx}_i} + \boldsymbol{\eta}_i, \qquad \boldsymbol{\eta}_i \sim \mathcal{N}(\mathbf{0}, \sigma^2_{\mathrm{Rx}_i}\mathbf{I}) \tag{6}$$

where $\boldsymbol{\eta}_i$ is receiver i's self-positioning error and σ_{Rx_i} its standard deviation, the per-receiver self-positioning uncertainty.

IoO-based self-positioning has been demonstrated experimentally, with operational σ_{Rx_i} magnitudes ranging from sub-metre to tens of metres [2], [3]. We treat receiver self-positioning as an upstream black box without the assumption of σ_{Rx_i} being IID as found in [13]. In our model, we assume that the error can be estimated from the onboard self-localisation filter output. The error is modelled as heterogeneous, since it would depend on the geometry of the SoO sources with respect to Rx_i, as well as different IMU drift characteristics [19] and SoO link quality.

We introduce a drift of receiver self-positioning over time. In a snapshot, receiver i's only self-positioning error is the random bias $\boldsymbol{\eta}_i$ with std σ_{Rx_i} from Eq. (6). In tracking scenarios, error additionally accumulates over time through a drifting offset $\boldsymbol{\Delta}_i(k)$, with $\boldsymbol{\Delta}_i(0) = \mathbf{0}$ and random-walk dynamics defined below. Unlike the static per-Rx bias $\boldsymbol{\eta}_i$, drift produces correlation in receiver i's noise across snapshots, which no single-snapshot weighting can compensate, so a tracker must carry the offset as state. The observed position of Rx_i at snapshot k is then given by:

$$\hat{\mathbf{x}}_{\mathrm{Rx}_i}(k) \;=\; \mathbf{x}_{\mathrm{Rx}_i} + \boldsymbol{\eta}_i + \boldsymbol{\Delta}_i(k) \tag{7}$$

The drift is modelled as a random walk.

$$\begin{aligned} \boldsymbol{\Delta}_i(k+1) &= \boldsymbol{\Delta}_i(k) + \mathbf{w}_i(k), \\ \mathbf{w}_i(k) &\sim \mathcal{N}\!\left(\mathbf{0},\; \sigma^2_{\mathrm{drift}_i}\,\Delta t\,\mathbf{I}\right) \end{aligned} \tag{8}$$

Here $k = 0, 1, \ldots, K-1$ indexes the K snapshots of a track, and Δt is the snapshot spacing. Fig. 4 illustrates an example drift trajectory.

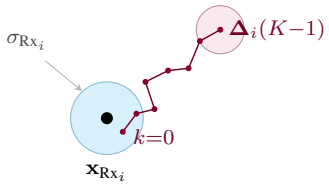

Fig. 4. Illustrative drift trajectory. The cyan circle indicates self-positioning error, the purple path represents drift. The purple circle is the cumulative drift uncertainty by the final snapshot.

F. Effective measurement variance

The measured bistatic range will be influenced by both an error from bistatic range estimation, and the error from receiver self-positioning. Time-varying quantities below are evaluated at snapshot k:

$$\begin{aligned} \hat{B}_i(k) &= B_i(\mathbf{x}_T, \mathbf{x}_{\text{Tx}}, \mathbf{x}_{\text{Rx}_i}) + \varepsilon_i(k), \\ \varepsilon_i(k) &\sim \mathcal{N}(0, \sigma^2_{\hat{B}_i}) \end{aligned} \tag{9}$$

The effective variance of the bistatic range measurement $\hat{B}_i(k)$ is estimated by summing the independent sources of error: range estimation noise, static bias, and accumulated drift.

$$\text{Var}(\hat{B}_i(k)) \approx \sigma^2_{\hat{B}_i} + \sigma^2_{\text{Rx}_i} + k\,\sigma^2_{\text{drift}_i}\,\Delta t \tag{10}$$

For the snapshot case, $k = 0$ and $\boldsymbol{\Delta}_i(0) = \mathbf{0}$, so the drift term vanishes and Eq. (10) reduces to $\sigma^2_{\hat{B}_i} + \sigma^2_{\text{Rx}_i}$.

Eq. (10) is an analytical approximation expressing the different components of measurement noise, used to define the weights in the solver. The simulated measurements and their noise are not drawn from that expression, but follow the true geometry, with self-positioning error entering through the use of the measured receiver position $\hat{\mathbf{x}}_{\text{Rx}_i}$.

IV. Methodology

The methodology of this work is a simulation study. The estimator is benchmarked across setups where (a) the receiver self-positioning error is assumed to be zero, as in most prior work, (b) the error is assumed to be IID, as in [13], and (c) the error is modelled as heterogeneous across receivers and, in tracking, drifting over time, breaking the IID assumption. The tracking functionality is implemented with a simple EKF on top of a snapshot estimator, and the same three error models are compared both in snapshot and tracking configurations.

A. Snapshot localisation solver

We adopt an iterative weighted least-squares (WLS) solver for snapshot localisation, an established approach for range-based localisation [9], [11]. We compare three weighting schemes: (a) "blind", which ignores self-positioning; (b) "iid", which uses the averaged variance $\bar{\sigma}^2_{\text{Rx}}$, matching the IID error model evaluated in [13]; (c) "per-Rx", our weighting that substitutes $\sigma^2_{\text{Rx}_i}$ per receiver, breaking the *identically-distributed* assumption when self-positioning quality is heterogeneous across receivers. Eq. (10) defines the effective variance of a single bistatic-range measurement, combining range estimation noise, static self-positioning error, and accumulated drift. Range estimation noise enters all three schemes identically and the drift term $k\,\sigma^2_{\text{drift}_i}\,\Delta t$ is deferred to the EKF tracking extension. The weighting matrix W then encodes how each scheme treats the self-positioning term:

$$W_{ii} = \begin{cases} 1/\sigma^2_{\hat{B}_i} & \text{(blind)} \\ 1/(\sigma^2_{\hat{B}_i} + \bar{\sigma}^2_{\text{Rx}}) & \text{(iid)} \\ 1/(\sigma^2_{\hat{B}_i} + \sigma^2_{\text{Rx}_i}) & \text{(per-Rx)} \end{cases} \tag{11}$$

The expectation is that per-Rx outperforms both iid and blind for growing $\sigma_{\text{Rx,mean}}$. By weighting each receiver according to its own self-positioning variance, it down-weights the least reliable ones, which matters increasingly as that term comes to dominate the measurement variance.

B. Tracking extension

Drift introduces temporal correlation within each receiver's noise sequence: unlike static per-Rx heterogeneity, it cannot be absorbed into a snapshot weighting and must be carried as estimator state. The EKF does this by augmenting the target state with per-Rx drift offsets, extending the static, IID error model evaluated in [13] to a drifting one. The state vector is

$$\mathbf{x}^{\text{state}} = [\mathbf{x}_T,\, \mathbf{v}_T,\, \boldsymbol{\Delta}_1,\, \ldots,\, \boldsymbol{\Delta}_N]^\top \in \mathbb{R}^{4+2N}. \tag{12}$$

where $\mathbf{v}_T$ is the target velocity. The input to the EKF is bistatic range and bistatic range-rate (related to velocity) measurements [9]. The EKF state starts warm with snapshot estimator output at $k = 0$ and then tracks the target for 50 steps. A tilde over a noise std (e.g., $\tilde{\sigma}_{\text{Rx},i}$, $\tilde{\sigma}_{\text{drift},i}$) denotes the solver's assumed value, distinct from the true σ. We initialise the $\boldsymbol{\Delta}_i$ states to zero with covariance $\tilde{\sigma}^2_{\text{Rx},i}\,\mathbf{I}_2$. The noise on drift increments $\boldsymbol{\Delta}_i$ is set by $\tilde{\sigma}_{\text{drift},i}$, the assumed value of the per-Rx drift rate σ_{drift_i} from the deferred drift term of Eq. (10):

$$\tilde{\sigma}_{\text{drift},i} = \begin{cases} 0 & \text{(blind)} \\ \bar{\sigma}_{\text{drift}} & \text{(iid)} \\ \sigma_{\text{drift}_i} & \text{(per-Rx)} \end{cases} \tag{13}$$

C. Monte Carlo procedure

The Monte Carlo simulation generates random configurations of target and receiver position within system model constraints. The constants are defined in Table II. Per-Rx self-positioning uncertainty σ_{Rx_i} is drawn IID per trial from a Uniform distribution centred on the cell's mean $\sigma_{\text{Rx,mean}}$, with spread factor $\sigma_{\text{Rx,high}}/\sigma_{\text{Rx,low}} = 10$ (Table II). The number of receivers N is swept from 3 to 10. The full sweep ranges are listed in Table III. The number of trials per cell is 10,000 for both snapshot (seed: `0xC5E3000`) and tracking (seed: `0x15EC75`) experiments.

D. Evaluation metrics

The main evaluation metric is the median target-position error with IQR across trials within each cell. For tracking, the per-trial error is the steady-state target-position RMSE over the last 30 snapshots, of which we likewise report the median (with IQR) across trials. Median is selected over mean or P90 to mitigate the heavy tail effect of outliers, which are expected from degenerate geometry cases.

V. Experiments and Results

We run four experiments: (1) snapshot WLS error vs. (N, σ_{Rx}) for the per-Rx solver, characterising the operating regime; (2) three-solver comparison in snapshot WLS; (3) three-solver comparison in EKF tracking; (4) per-Rx EKF vs. per-Rx snapshot WLS. The constants defining the system model in simulation are listed in Table II. The sweep variables are listed in Table III.

A. Snapshot target-position error vs. N and σ_{Rx}

In the first experiment, we demonstrate how the snapshot WLS target-position error varies with the number of receivers N and the per-Rx self-positioning uncertainty σ_{Rx} for the per-Rx solver. The results are shown in Fig. 5. Adding receivers reduces the error across the swept range: the benefit is significant for small σ_{Rx}, where the error is σ_B-dominated, and diminishes above the knee at $\sim$5 m, where σ_{Rx} dominates.

B. Solver comparison in snapshot WLS

Within this regime, we next ask whether solver choice matters. In this experiment, we compare the three solvers in the snapshot setting. The results are shown in Fig. 6. For small σ_{Rx}, the three solvers perform similarly. For higher σ_{Rx}, iid performs marginally better than blind, and per-Rx performs better than iid. We pick a particular value of $\sigma_{\text{Rx}} = 5$ m for analysis, which is a realistic range from [2]: at $N = 10$, the median target-position error is 4.85 m for per-Rx, 5.55 m for iid, and 6.55 m for blind, a 13 % reduction of per-Rx over iid and 26 % over blind. The ordering holds across the sweep: per-Rx achieves a lower median than iid in 64 of 72 (N, σ_{Rx}) cells. Intuitively, the closer a solver's weights match the true measurement variance, the lower its error: blind ignores self-positioning, iid accounts for it in aggregate, and per-Rx accounts for it per receiver, so the ordering follows.

C. Solver comparison in EKF tracking

In this experiment, we re-implement the three solvers in an EKF tracker and compare their steady-state tracking error. The drifting-receiver case is deferred to future work in [13], as its theoretical analysis requires the hybrid CRLB. We instead characterise it empirically via Monte Carlo, comparing the three solvers under drift. The results are displayed in Fig. 7. The difference between the solvers is more pronounced than in the snapshot case, with the same ordering visible across all σ_{drift}. At $N = 10$ and $\sigma_{\text{drift}} = 0.1\,\text{m}/\sqrt{\text{s}}$, the steady-state median is 2.90 m for per-Rx, 4.01 m for iid, and 4.80 m for blind, a 28 % reduction of per-Rx over iid and 40 % over blind.

D. EKF vs. snapshot WLS

In this experiment we compare how the per-Rx EKF tracker performs against the per-Rx snapshot WLS estimator. The results are shown in Fig. 8. This is an illustrative experiment to demonstrate the benefit of tracking over snapshot estimation. Across $N \in \{3, \dots, 10\}$ and $\sigma_{\text{drift}} \in \{0.01, \dots, 1\}\,\text{m}/\sqrt{\text{s}}$, the EKF tracker yields a 1.4–1.7$\times$ median improvement over snapshot WLS (median across cells: 1.6$\times$).

Symbol	Value	Description
Geometry		
r	1 km	Rx ring radius
target-area radius	3 km	max. target distance from ring centre
d_{Tx}	5 km	Tx distance from ring centre
Measurement noise (reference values)		
$\sigma_{B,\text{ref}}$	5 m	at $R_{t,\text{ref}}$, R_{ref}
R_{ref}	1 km	reference R_r
$R_{t,\text{ref}}$	15 km	reference R_t
$\sigma_{B,\min}$	0.5 m	non-SNR floor
$\sigma_{\text{Rx},\min}$	0.5 m	WLS regularisation floor
Tracking dynamics		
K	50	snapshots per track
Δt	0.5 s	snapshot spacing
$\lvert\mathbf{v}_T\rvert$	10 m/s	target speed (CV)
σ_{acc}	0.1 m/s^2	target process noise
n_{steady}	30	snapshots for steady-state
$\sigma_{\text{Rx}}^{\text{trk}}$	5 m	fixed (tracking sweeps σ_{drift})
Per-Rx heterogeneity (iid Uniform per trial)		
σ_{Rx} spread	10	$\sigma_{\text{Rx},\max}/\sigma_{\text{Rx},\min}$
σ_{drift} spread	10	$\sigma_{\text{drift},\max}/\sigma_{\text{drift},\min}$
Monte Carlo		
trials/cell (snapshot)	10 000	per (N, σ_{Rx})
trials/cell (tracking)	10 000	per $(N, \sigma_{\text{drift}})$

TABLE II
Fixed simulation parameters.

Variable	Range	Used in
N	$\{3, \dots, 10\}$	snapshot, tracking
σ_{Rx} [m]	$\{0.3, 0.5, 1, 3, 5, 10, 30, 50, 100\}$	snapshot
σ_{drift} [m/$\sqrt{\text{s}}$]	$\{0.01, 0.03, 0.1, 0.3, 1\}$	tracking

TABLE III
Sweep ranges for the experiments.

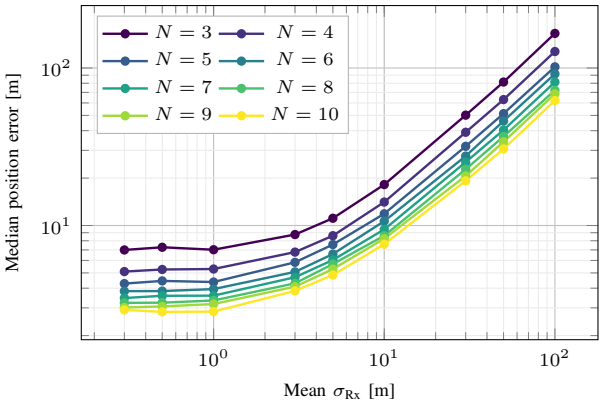

Fig. 5. Per-Rx median target-position error vs. mean self-positioning noise σ_{Rx}, one line per N (darker = higher N). Additional receivers reduce error across the swept range; the benefit is significant below the knee at $\sigma_{\text{Rx}} \approx 5$ m and diminishes above it, where σ_{Rx} dominates.

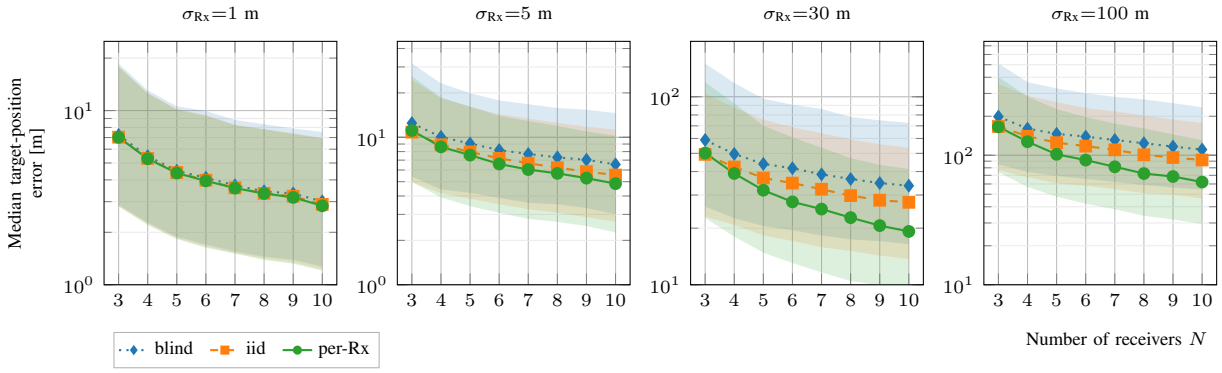

Fig. 6. Snapshot WLS target-position error vs. N across solvers (per-Rx, iid, blind); panels by mean σ_{Rx}. Lines: median; bands: P25–P75. The median ordering per-Rx $<$ iid $<$ blind is marginally visible.

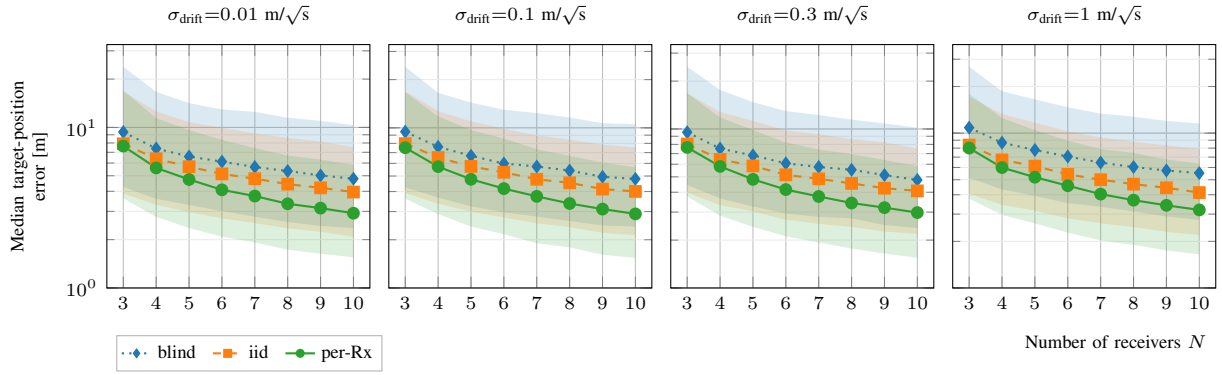

Fig. 7. EKF steady-state target-position error (last 30 of 50 snapshots) vs. N across solvers; panels by mean per-Rx drift rate σ_{drift}. Lines: median; bands: P25–P75. The per-Rx $<$ iid $<$ blind ordering is visible on the median across all σ_{drift} despite heavy IQR overlap. $\sigma_{\text{Rx,mean}} = 5\,\text{m}$ fixed.

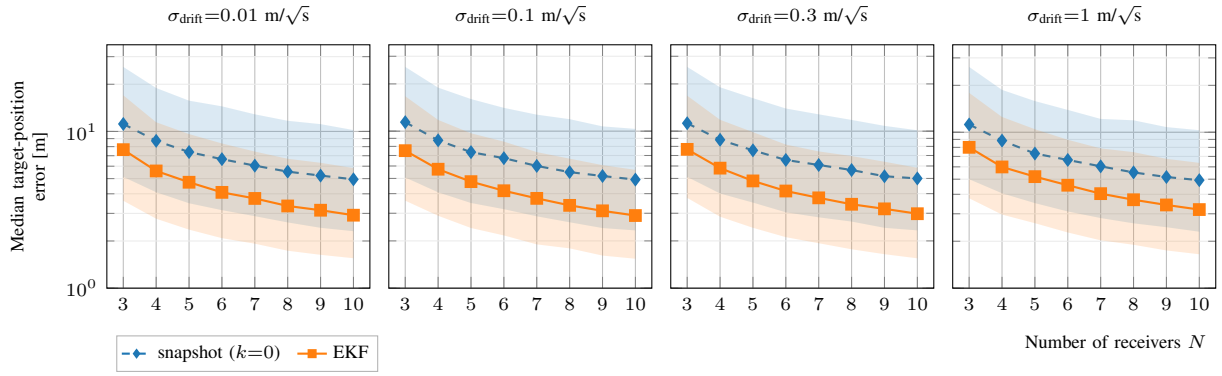

Fig. 8. Per-Rx solver only: EKF steady-state vs. snapshot WLS at k=0 (the EKF warm-start); panels by mean σ_{drift}. The tracker yields a 1.4–1.7$\times$ median improvement at every N across the swept σ_{drift} range. $\sigma_{\text{Rx,mean}} = 5\,\text{m}$ fixed.

VI. Discussion

We have shown that in our system model, per-receiver self-positioning awareness improves target-position accuracy over the iid [13] and blind baselines (Fig. 6, Fig. 7): the median gap is $\sim$13% in snapshot WLS, growing to $\sim$28% in EKF tracking. Snapshot per-Rx weighting addresses heterogeneity in space alone, the static positional uncertainty σ_{Rx_i} of each receiver. Tracking additionally addresses heterogeneity in time: the per-Rx drift rate σ_{drift_i} governs how each receiver's position uncertainty accumulates, and the EKF integrates this advantage across 50 snapshots.

This work adapts and extends the error model of [13] by empirically addressing the drifting-receiver case that study leaves to future work, and by evaluating it in the non-IID regime that study leaves unexamined. The practical implication is that in UAV-based multistatic passive radar, estimating per-

receiver self-positioning errors improves target localisation, particularly in tracking deployments. Although the per-Rx $>$ iid $>$ blind ordering is intuitive, surveyed multistatic passive radars (e.g., [9]) assume perfect or IID self-positioning in evaluation, so this empirical characterisation is novel. The 28% per-Rx improvement over the iid baseline in tracking suggests that per-receiver self-positioning awareness is a meaningful radar design parameter, with influence on accuracy non-negligible alongside other parameters such as receiver count or geometry.

The principal limitations of this work are: (a) the lack of experimental validation; (b) a simplified simulation model that treats receiver self-localisation noise as independent of bistatic range noise, though both may be correlated through the receiver's local signal conditions; (c) 2D simulation only; (d) a single transmitter; (e) quasi-stationary rather than mobile receivers, with no Doppler frequency modelling; (f) the assumption that per-Rx σ_{Rx_i} is reported exactly by the onboard self-localisation filter, which may be an estimation challenge in practice; (g) the IQR bands of all three solvers overlap heavily in both snapshot and tracking regimes, moderating the operational claim; (h) only a constant-velocity linear target trajectory, with non-linear paths (e.g., circular, as in [9]) untested.

VII. Responsible Research

A. Reproducibility

While the simulation and procedures are designed to be reproducible, simulation implementation details are not included in the paper. Code, seeds and data are available on GitHub under the MIT license at `https://github.com/doodek/multistatic-rx-uncertainty`. All parameters used, as well as dependencies with specific versions, are documented in the repository. Data from the simulations is archived in `data/*.npz`. Bit-level floating point reproducibility is not guaranteed but does not affect reported metrics significantly.

B. Ethical aspects and dual use

This work falls under the category of dual use research, and its position in that category is characterisation of a capability (the localisation and tracking of uncooperative drones) rather than enabling it. Explored passive radar technologies may be used for counter-drone kinetic countermeasures. False positives in this context may endanger civilian aircraft and other lawful airspace traffic.

Since passive radars do not emit signals, the tracked entity has no possibility of sensing that it is being tracked, enabling covert surveillance of UAV operators.

The simulation code is a Monte Carlo simulation framework which is not deployable to hardware. No proprietary or operator-specific technology was explored. There is no contribution to target classification or detection policy.

VIII. Conclusion

This work has surveyed the state of the art in multistatic passive radar for target localisation and tracking using signals of opportunity, with a focus on feasibility in airborne settings in GNSS-denied environments. We have identified receiver positioning uncertainty as a key error source in radar performance, and proposed a per-receiver self-positioning awareness approach to mitigate it, based on the error model of [13] and evaluated under the heterogeneous and drifting errors that study leaves to future work. We have demonstrated via Monte Carlo simulation that per-receiver awareness improves target localisation and tracking accuracy over IID and blind baselines, with a $\sim$13% reduction in median target-position error over iid (snapshot WLS) and $\sim$28% in median steady-state RMSE (tracking).

Future work could include implementing the proposed approach in a real-world airborne multistatic passive radar system, as well as extending the simulation to 3D, including correction for moving receivers, evaluating non-linear target trajectories, and extending to multiple transmitters. Establishing theoretical limits such as the hybrid CRLB for the drifting-receiver case deferred in [13] would be a valuable complement to the empirical characterisation presented here. Furthermore, error distributions could be characterised beyond median and IQR, in particular by isolating geometry-degenerate trials which dominate distribution tails.

Appendix

AI Tooling (Claude Code) was used in work on this paper. All content is under complete ownership and responsibility of the author. The tooling was used for language editing, formatting, drawing figures, generating feedback on the text, and co-developing the simulation software. Prompt engineering was used to guide AI on these tasks, specifically to prevent AI from taking over the intellectual contribution of the work, and to restrict it to a supportive role. The author reviewed and edited all AI-generated content.

References

[1] S. Alghamdi, S. Alahmari, S. Yonbawi, K. Alsaleem, F. Ateeq, and F. Almushir, "Autonomous navigation systems in GPS-denied environments: A review of techniques and applications," in *2025 11th International Conference on Automation, Robotics, and Applications (ICARA)*, 2025, pp. 290–299.

[2] Z. M. Kassas, J. Khalife, A. A. Abdallah, and C. Lee, "I am not afraid of the GPS jammer: Resilient navigation via signals of opportunity in GPS-denied environments," *IEEE Aerospace and Electronic Systems Magazine*, vol. 37, no. 7, pp. 4–19, 2022.

[3] J. Khalife and Z. M. Kassas, "On the achievability of submeter-accurate UAV navigation with cellular signals exploiting loose network synchronization," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 58, no. 5, pp. 4261–4278, Oct. 2022.

[4] K. Kulpa, M. Malanowski, P. Samczynski, and B. Dawidowicz, "The concept of airborne passive radar," in *2011 Microwaves, Radar and Remote Sensing Symposium (MRRS)*, Kiev, Ukraine, 2011, pp. 267–270.

[5] V. S. Chernyak, *Fundamentals of Multisite Radar Systems*. Gordon and Breach Science Publishers, 1998.

[6] U. Bhattacherjee, E. Ozturk, O. Ozdemir, I. Guvenc, M. L. Sichitiu, and H. Dai, "Experimental study of outdoor UAV localization and tracking using passive RF sensing," in *Proceedings of the 15th ACM Workshop on Wireless Network Testbeds, Experimental evaluation & CHaracterization (WiNTECH '21)*. New Orleans, LA, USA: ACM, 2022, pp. 1–8.

[7] G. Fang, J. Yi, X. Wan, Y. Liu, and H. Ke, "Experimental research of multistatic passive radar with a single antenna for drone detection," *IEEE Access*, vol. 6, pp. 33 542–33 551, 2018.

[8] S. Rzewuski, K. Kulpa, M. Pachwicewicz, M. Malanowski, and B. Salski, "Drone detectability feasibility study using passive radars operating in WiFi and DVB-T band," NATO STO MP-MSG-SET-183 paper 13, 2018.

[9] Y. E. Sagduyu, K. Davaslioglu, T. Erpek, S. Kompella, G. Anderson, and J. Ashdown, "MULTI-SCOUT: Multistatic integrated sensing and communications in 5G and beyond for moving target detection, positioning, and tracking," in *MILCOM 2025 — 2025 IEEE Military Communications Conference*, Los Angeles, CA, USA, 2025, pp. 693–698.

[10] M. Malanowski and K. Kulpa, "Two methods for target localization in multistatic passive radar," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 48, no. 1, pp. 572–580, 2012.

[11] A. Noroozi and M. A. Sebt, "Target localization from bistatic range measurements in multi-transmitter multi-receiver passive radar," *IEEE Signal Processing Letters*, vol. 22, no. 12, pp. 2445–2449, 2015.

[12] A. Ahmed, S. Zhang, and Y. D. Zhang, "Multi-target motion parameter estimation exploiting collaborative UAV network," in *ICASSP 2019 — 2019 IEEE International Conference on Acoustics, Speech and Signal Processing*, Brighton, UK, 2019.

[13] L. Rui and K. C. Ho, "Elliptic localization: Performance study and optimum receiver placement," *IEEE Transactions on Signal Processing*, vol. 62, no. 18, pp. 4673–4688, 2014.

[14] J. Brown, K. Woodbridge, A. Stove, and S. Watts, "Air target detection using airborne passive bistatic radar," *Electronics Letters*, vol. 46, no. 20, pp. 1396–1397, Sep. 2010.

[15] N. Souli, P. Kardaras, P. Kolios, and G. Ellinas, "Onboard passive radar system implementation for detection and tracking of rogue UAVs," in *2023 International Conference on Unmanned Aircraft Systems (ICUAS)*, 2023, pp. 820–826.

[16] J. W. A. Brown, "FM airborne passive radar," Ph.D. dissertation, University College London, 2013.

[17] B. Gabard, V. Wasik, O. Rabaste, T. Deloues, D. Poullin, and H. Jeuland, "Airborne targets detection by UAV-embedded passive radar," in *2020 17th European Radar Conference (EuRAD)*, 2021.

[18] N. J. Willis, *Bistatic Radar*, 2nd ed. SciTech Publishing, 2005.

[19] VectorNav Technologies, "VN-100 SMD IMU/AHRS Sensor Datasheet (Hardware v7.0)," https://www.vectornav.com/products/detail/vn-100, 2024, document ID: DS100-SMD-70-R1.