engineering proceedings

MDPI

Proceeding Paper

Low-Rate Cross-View Map Updates for Drift-Bounded Fixed-Wing VIO in GNSS-Denied Missions †

Arif Hajiyev [1,*], Zeki Yetgin [1] and Sabina Aghazada [2]

[1] Department of Computer Engineering, Faculty of Engineering, Mersin University, Ciftlikkoy Campus, 33343 Yenisehir, Mersin, Türkiye; zyetgin@mersin.edu.tr
[2] School of Information Technology, National Aviation Academy, Mardakan Avenue 30, AZ1045 Baku, Azerbaijan; sabina.aghazada@naa.edu.az
* Correspondence: 24023010008@mersin.edu.tr
† Presented at the 13th International Electronic Conference on Sensors and Applications, 18–20 November 2026; Available online: https://sciforum.net/event/ECSA-13.

Abstract

Fixed-wing UAVs operating in GNSS-denied conditions require stable state estimates over long distances, yet standalone visual–inertial odometry (VIO) still accumulates drift during sustained forward flight because loop closures are rare and viewpoint variation is limited. This paper presents a drift-bounded navigation architecture with five components: a camera–IMU VIO front end, a motion-aware keyframe selector, a cross-view aerial-map retrieval module, an image-to-map alignment stage, and a confidence-gated conservative correction back end. The VIO front end runs continuously to provide high-rate local pose, while map anchoring is invoked only at a low rate on selected keyframes. A map update is accepted only when the retrieval score is high, the geometric inlier ratio is sufficient, the reprojection error is low, and the estimated heading is consistent with the onboard solution; in the present implementation, the corresponding gates are 0.80, 0.35, 3.0 px, and 8°, respectively. Accepted updates correct global position and yaw without resetting the local VIO trajectory. In representative high-fidelity simulation with staged GNSS loss over a 24 km route, standalone VIO reached 0.92% final drift, whereas the proposed method reduced drift to 0.28%, with an 81% update success rate and 22.4 m absolute trajectory error. In shorter fixed-wing flight tests over 8.6 km, drift decreased from 1.34% to 0.41%, with 74% accepted updates and a 34.7 m absolute trajectory error. Compared with continuous map matching, the low-rate policy preserved similar drift behaviour while reducing map-matching calls by 75% and sustaining 18 FPS on an Intel NUC 8i7BEH-class onboard computer. The main contribution is a fixed-wing-oriented fusion policy that uses geo-referenced maps as occasional, verified corrections rather than as a continuously active localisation source. Because complete GNSS denial is only the endpoint of a wider range of GNSS fault conditions, the paper additionally specifies an innovation-consistency layer that admits, down-weights, or rejects GNSS measurements when they are available but degraded; this layer is presented as an architectural extension with a defined validation protocol and is not included in the experimental results reported above.

Keywords: fixed-wing UAV; GNSS-denied navigation; visual-inertial odometry; cross-view geo-localisation; map matching; drift-bounded navigation; robust multi-sensor integration; GNSS fault detection; real-time onboard estimation

Check for updates

Academic Editor: Francisco Falcone

Published: 28 October 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Long-endurance fixed-wing unmanned aerial vehicles are used for corridor inspection, wide-area surveys, and search tasks in which a single sortie may cover tens of kilometres. Almost all such platforms rely on GNSSs for the absolute component of their navigation solution. That dependence is convenient but fragile: signal availability degrades under jamming, spoofing, and interference, and it can be lost entirely for extended periods [1]. When this happens, the aircraft must continue to produce a usable global position from onboard sensing alone, and it must do so for long enough to complete or safely abort the mission.

Visual–inertial odometry is the standard onboard fallback. Filter-based [2] and optimisation-based [3–6] estimators fuse a camera with an inertial measurement unit and deliver locally accurate, high-rate motion estimates at modest computational cost. Their weakness is well known and structural: the estimate is dead-reckoned, so position and yaw errors grow without bound in the absence of an absolute reference. On a fixed-wing platform the growth is worse than on the rotorcraft benchmarks that dominate the VIO literature. Cruise flight is close to a straight line at nearly constant altitude, so the aircraft rarely revisits an earlier location and loop closure is effectively unavailable. Viewpoint diversity is small, the observed scene is far away, and the resulting parallax is weak, which leaves the scale and heading poorly constrained. In our own experience, this produces a slow, monotonic heading bias whose contribution to horizontal error grows roughly linearly with distance flown.

A geo-referenced aerial or satellite orthophoto of the operating area is an attractive way to break that growth, because it is an absolute reference that requires no radio link. Absolute visual localisation for UAVs has been surveyed comprehensively by Couturier and Akhloufi [7], and the underlying image-retrieval problem—matching an onboard image against a georeferenced map tile—has advanced considerably through learned cross-view descriptors [8–10] and through direct registration of UAV imagery against pre-existing satellite imagery [11]. Combined with modern local features and matchers [12,13] and a robust geometric estimator [14], these components can, in favourable conditions, recover a metric global position and heading from a single frame.

The difficulty is not whether such a match can be computed, but whether it should be trusted. Cross-view matching over farmland, forest, water, or repetitive suburban texture fails in a characteristic way: it returns a confident but wrong tile. Injecting one such match into the navigation filter produces a position jump of hundreds of metres, and on a fixed-wing aircraft flying a pre-planned route that jump can propagate into the guidance loop before any subsequent measurement can correct it. Running the matcher continuously, as is common in the localisation literature, both maximises the exposure to this failure mode and consumes a large share of a small onboard computing budget. The design question for a real mission is therefore a policy question—how often to consult the map, and under what evidence to act on the answer—rather than a question of retrieval accuracy alone.

There is a further point that we did not treat adequately in the first version of this work and that we address explicitly here. "GNSS-denied" is not a separate operating mode so much as the extreme end of a continuum. In practice, a receiver passes through healthy operation, degraded operation with inflated noise or multipath, intermittent and outlier-contaminated operation, slowly growing bias, and only then complete loss. The integrated-navigation community has studied precisely this continuum for many years under the heading of robust and adaptive filtering: Mahalanobis-distance criteria for identifying abnormal observations and inflating their covariance [15,16], hypothesis testing and sequential probability ratio tests for separating abrupt outliers from slow-growing faults [17,18], robust unscented and cubature formulations with measurement-error detec-

tion [19,20], anomaly detection embedded in windowed factor-graph optimisation [21], and, more recently, fuzzy, kernel-density, kernel-weighted, and neuro-fuzzy schemes that replace a hard accept/reject decision with a graded adjustment of measurement influence [22–26]. A navigation architecture that claims to operate in GNSS-denied missions should state where it sits with respect to that body of work and should say what it does when a GNSS is present but untrustworthy rather than absent.

Against this background, the contributions of this paper are as follows. First, we describe a five-component architecture in which a continuously running camera–IMU VIO front end is anchored, at a low rate and only on motion-aware keyframes, by cross-view retrieval and image-to-map alignment. Second, we define a conservative acceptance policy in which a candidate map update must simultaneously satisfy a retrieval-score, inlier-ratio, reprojection-error, and heading-consistency gate before it is allowed to modify the global position and yaw, and in which an accepted update never resets the local VIO trajectory. Third, we report quantitative drift, accuracy, acceptance-rate, and throughput results for this configuration in a high-fidelity simulation over 24 km and in fixed-wing flight over 8.6 km, together with a comparison against continuous map matching. Fourth, and new in this version, we place the method within the robust multi-sensor integrated-navigation literature and specify an innovation-consistency layer that governs how the same architecture should treat GNSS measurements that are available but degraded, together with the fault-injection protocol and detector-level metrics required to validate it. We are explicit that this fourth contribution is an architectural specification: it has not yet been exercised experimentally, and none of the numbers reported in Section 5 depend on it.

The remainder of the paper is organised as follows. Section 2 reviews visual–inertial odometry, absolute visual localisation, and robust integrated navigation. Section 3 describes the architecture and the acceptance policy, including the GNSS-validation extension. Section 4 describes the experimental setup and states the scope of the evidence. Section 5 reports results, Section 6 discusses them and sets out limitations and the planned validation, and Section 7 concludes.

2. Related Work

2.1. Visual–Inertial Odometry on Fixed-Wing Platforms

Vision-aided inertial navigation has converged on two families. The multi-state constraint Kalman filter [2] and its descendants marginalise feature states to keep computation bounded, which suits embedded targets. Optimisation-based estimators built on IMU preintegration [3] achieve higher accuracy by re-linearising over a sliding window; VINS-Mono [4], OpenVINS [5], and ORB-SLAM3 [6] are representative and widely used as baselines. Reported benchmark performance is, however, dominated by handheld and multirotor datasets characterised by rich rotational excitation and frequent revisits of previously mapped areas.

Fixed-wing flight violates several of the assumptions behind those results. Rotational excitation is small, so the IMU biases and the visual scale are only weakly observable; the scene is at long range, so triangulation is ill-conditioned; and the flight path rarely closes a loop, so the mechanism that keeps drift bounded in indoor and urban SLAM is unavailable. The practical consequence is that even a well-tuned VIO stack on a fixed-wing aircraft produces error that grows with distance travelled rather than saturating, which is the behaviour our own standalone baseline exhibits in Section 5.

2.2. Absolute Visual Localisation and Cross-View Geo-Localisation

Because loop closure is unavailable, the absolute reference has to come from outside the flight itself. Couturier and Akhloufi [7] survey the alternatives for UAVs and note the

recurring trade-off between map-based methods that are accurate but brittle and odometric methods that are stable but unbounded. Cross-view geo-localisation addresses the appearance gap between an onboard view and an overhead map: CVM-Net [8] introduced a learned cross-view embedding, spatial-aware feature aggregation [9] improved the handling of orientation ambiguity, and VIGOR [10] relaxed the one-to-one retrieval assumption that earlier benchmarks imposed. On the registration side, Goforth and Lucey [11] demonstrated GPS-denied UAV localisation by aligning onboard imagery with pre-existing satellite imagery, and learned local features and matchers [12,13] with RANSAC-family geometric verification [14] provide the pixel-level correspondence needed to turn a retrieved tile into a metric pose.

What this literature does not settle is the fusion policy. Most systems query the map on every frame or at a fixed high rate and treat the result as a measurement of known quality. For a small onboard computer carrying a full guidance stack, that is expensive; more importantly, it exposes the estimator to the confident-but-wrong retrieval failure that dominates the practical risk. The present work is aimed squarely at that gap.

2.3. Robust Multi-Sensor Filtering and GNSS Fault Handling

The problem of admitting or rejecting an unreliable measurement is not specific to map matching. It is the central concern of robust integrated navigation, and the methods developed there transfer directly to the architecture described in Section 3.

The oldest and most widely used mechanism is a consistency test on the innovation. Chang [27] formalised Mahalanobis distance as an outlier-judging criterion for Kalman filtering, and Gao et al. [15] embedded the same criterion in a cubature Kalman filter for vehicular INS/GNSS integration, using it to compute a scaling factor that inflates the observation-noise covariance and thereby reduces the filter gain when an abnormal observation is identified. This is an important point of principle for our own design: the response to a suspicious measurement need not be binary. Hu et al. [19] reached a similar conclusion from the direction of innovation orthogonality, detecting measurement errors by hypothesis test and rescaling the Kalman gain so that erroneous measurements are down-weighted without discarding the information they carry when they are in fact valid. Gao et al. [16] went further and combined an adaptive CKF for kinematic-model error with a robust CKF for observation-model error inside an interacting-multiple-model structure, so that a single estimator can respond to either error source.

A single-epoch consistency test is, however, blind to faults that grow slowly. Gao et al. [17] addressed this with a double-channel sequential probability ratio test that runs a horizon-based channel for slow-growing failures alongside a current-sample channel for abrupt ones and demonstrated that the combination detects both classes where either alone does not. The same reasoning underlies the hypothesis-test-constrained robust Kalman filter of Gao et al. [18], in which a chi-square test characterises abrupt outliers and an improved sequential probability ratio test characterises slow-growing ones, with both tests then constraining the maximum-likelihood estimate of the measurement-noise covariance. Hu et al. [20] extended robust UKF processing to decentralised fusion of INS, GNSS, and celestial measurements, using a Mahalanobis-based hypothesis test at the local filters before global fusion—an arrangement conceptually close to the source arbitration we describe in Section 3.7. Robustification is also not confined to Kalman back ends: Dong et al. [21] embedded Mahalanobis-distance anomaly detection as a node inside a sliding-window factor graph for underwater INS/DVL/USBL integration, retaining real-time performance while resisting observation anomalies.

More recent work replaces fixed thresholds with adaptive weighting. Hu et al. [22] proposed an indirect fuzzy robust CKF in which normalised innovation quantities drive a

fuzzy inference system that produces a scaling matrix for filter correction. Gao et al. [23] introduced a two-level reliable measurement selection mechanism that screens aiding measurements before they enter an extended Kalman filter, which is precisely the "check before you accept" principle that our map-update gates implement in a different modality. Gao et al. [24] modelled the unknown probability density of abrupt-change measurement-noise covariance using kernel density estimation, and Gao et al. [25] constructed Gaussian-kernel weights with adaptive bandwidth to handle both slightly and abruptly varying noise statistics. Gao et al. [26] embedded a dual-event-triggering adaptive neuro-fuzzy inference system inside a UKF for cluster cooperative navigation, activating the adaptive machinery only when triggering conditions indicate anomalous measurements and thereby reducing computational load in nominal operation—an event-triggered philosophy that closely parallels our own decision to invoke expensive map matching only on selected keyframes. Kernel-based and correntropy-type criteria have likewise been applied to INS/GNSS and INS/GNSS/BPNS integration to limit the influence of non-Gaussian noise and outliers [28].

Table 1 summarises how each of these lines of work relates to the mechanism used in this paper. The purpose of the table is not to claim that we implement these estimators; it is to make explicit which established principle each element of our design corresponds to and where a stronger alternative exists for future work.

Table 1. Robust multi-sensor integrated-navigation prior art and its relation to the present architecture.

Work	Year	Mechanism	Relation to this Paper
Chang [27]	2014	Mahalanobis distance as outlier-judging criterion	Theoretical basis for the innovation-consistency gate of Section 3.6
Gao et al. [15]	2019	Robust CKF; Mahalanobis-derived scaling factor inflates observation covariance	Direct precedent for graded down-weighting rather than hard rejection
Hu et al. [19]	2019	Innovation-orthogonality robust UKF with measurement-error detection	Supports gain rescaling instead of discarding suspect measurements
Gao et al. [16]	2021	Adaptive-plus-robust CKF fused by interacting multiple models	Alternative to the fixed two-threshold policy adopted here
Gao et al. [17]	2021	Double-channel SPRT for abrupt and slow-growing failures	Basis for the persistent-fault monitor of Section 3.6
Gao et al. [18]	2023	Chi-square-plus-improved SPRT constraining ML noise-covariance estimation	Justifies combining a per-epoch gate with a sequential monitor
Hu et al. [20]	2023	Robust UKF decentralised multi-sensor information fusion	Precedent for local validation prior to global source arbitration
Dong et al. [21]	2024	Windowed factor graph with Mahalanobis anomaly detection	Candidate back end if the estimator moves from filtering to windowed FGO
Hu et al. [22]	2024	Indirect fuzzy robust CKF with normalised input parameters	Graded alternative to the discrete accept/down-weight/reject policy
Gao et al. [23]	2024	Two-level reliable measurement selection before EKF update	Same principle as the multi-gate map-update admission test
Gao et al. [24]	2025	KDE-based robust KF for abrupt-change noise covariance	Option when GNSS error is heavy-tailed rather than sparsely faulty
Gao et al. [25]	2026	Kernel-function-weighted robust KF, adaptive bandwidth	Adaptive replacement for the fixed inflation factor of Section 3.6
Gao et al. [26]	2026	Dual-event-triggering ANFIS-UKF for measurement anomalies	Event-triggered computation parallels low-rate map invocation

2.4. Positioning of the Present Work

We do not propose a new robust estimator, and the paper should not be read as competing with the filters listed in Table 1. Our contribution is a fusion policy for a specific and, we would argue, under-served platform class: a fixed-wing aircraft that must survive a long GNSS outage using a geo-referenced map that it can consult only occasionally and cannot fully trust. The robust-filtering literature supplies the statistical machinery for deciding whether a measurement deserves to be believed; what we add is the argument that on this platform the map should be treated as an occasional, heavily verified correction source rather than a continuously active localisation source, together with an evaluation of what that policy costs and saves. Section 3.6 then shows that the same admission logic extends naturally to GNSSs when GNSSs are present but degraded.

3. Materials and Methods

3.1. System Overview

The architecture comprises five components, shown in Figure 1: a camera–IMU visual–inertial front end, a motion-aware keyframe selector, a cross-view retrieval module operating against a georeferenced orthophoto database, an image-to-map alignment stage, and a confidence-gated correction back end that maintains the global anchor.

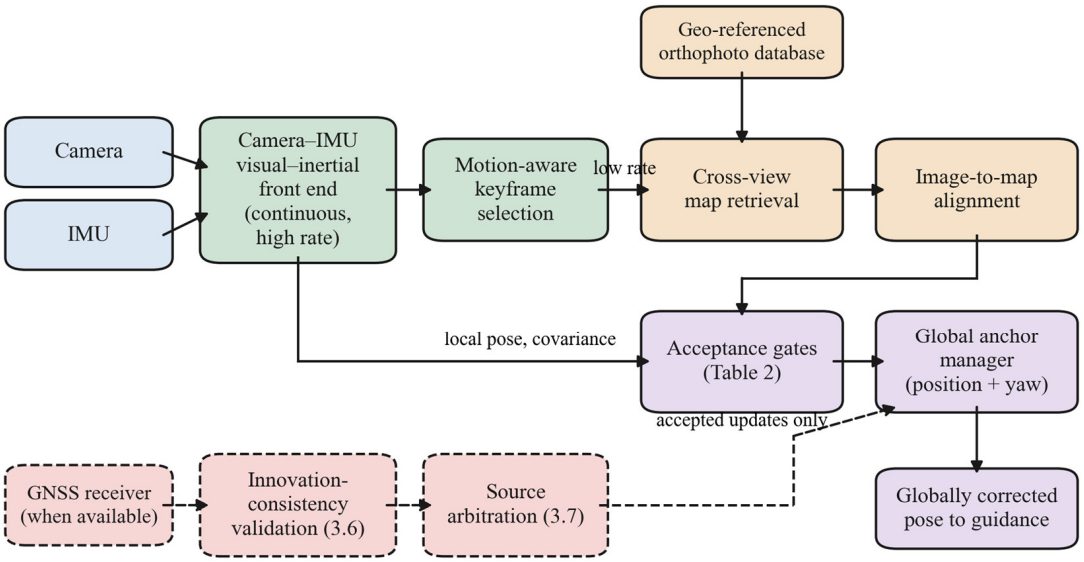


Figure 1. System architecture. The camera–IMU front end runs continuously; cross-view retrieval and alignment are invoked at a low rate on selected keyframes; and the correction back end applies a global position and yaw update only when all acceptance gates are satisfied. The GNSS validation branch (dashed) is the architectural extension described in Section 3.6 and was not active during the experiments of Section 5.

The separation of concerns is deliberate. The VIO front end runs continuously and owns the local, high-rate trajectory. The map branch runs intermittently and owns nothing except a candidate global correction, which it must justify before the back end applies it. The global anchor is the only object that both branches touch, and it is updated by a bounded, gated operation rather than by resetting the estimator. This means that a failure in the map branch—including a completely wrong retrieval that somehow survives the

gates—degrades global accuracy but cannot corrupt local odometry, attitude, or the velocity estimate on which the flight controller depends.

3.2. Visual–Inertial Front End

The front end is a tightly coupled keyframe-based visual–inertial estimator using IMU preintegration between image frames [3] and a sliding window over recent keyframes. It outputs a pose at IMU rate and a marginalised covariance for the current state. We deliberately treat this component as replaceable: the policy contributed by this paper is agnostic to whether the underlying estimator is filter-based [2,5] or optimisation-based [4,6], and any estimator that exposes a pose, a covariance, and a keyframe stream can be substituted.

The front end operates in a local frame that is initialised at take-off and is never reset during flight. All global corrections are expressed as an updated transformation between this local frame and the georeferenced map frame, not as a modification of the local trajectory itself. This is the mechanism that keeps map errors from contaminating the local solution.

3.3. Motion-Aware Keyframe Selection

Querying the map on every frame is wasteful, because consecutive frames from a fixed-wing aircraft in cruise are highly redundant, and it is also counter-productive, because most of those frames are no more informative than the one before. The selector therefore proposes a keyframe for map anchoring only when the accumulated motion since the last accepted anchor justifies it. The criteria are the along-track distance travelled, the change in heading, the change in altitude, and the elapsed time since the last successful anchor, evaluated against the current estimated uncertainty: as the global covariance grows, the selector becomes more willing to spend a query.

Frames are additionally rejected before retrieval on the basis of cheap image-quality and geometry checks—excessive bank angle, motion blur, and cloud or haze occlusion—since a query issued on an unusable frame consumes the same computation as a useful one and can only add risk. In cruise, this policy yields map queries at intervals of the order of tens of seconds rather than at the frame rate, which is the source of the workload reduction reported in Section 5.3.

3.4. Cross-View Retrieval and Image-to-Map Alignment

For a selected keyframe, a search region in the georeferenced orthophoto database is defined from the current global position estimate and its covariance, which keeps the retrieval problem local rather than global and substantially reduces the chance of a distant false positive. A cross-view descriptor is computed for the query image and matched against descriptors of the candidate tiles, following the learned cross-view formulation established in [8–10]; the best candidate and its similarity score are passed to alignment.

Alignment estimates the geometric relationship between the query image and the retrieved tile using local features and a learned matcher [12,13], followed by robust model estimation with RANSAC [14]. From the resulting correspondence set and the known ground sample distance and geo-referencing of the tile, a global horizontal position and a map-frame heading are recovered, together with the inlier ratio and the mean reprojection error of the inlier set. These two quantities, rather than the retrieval score alone, carry most of the discriminative power in the acceptance test that follows, because a wrong tile can occasionally produce a high similarity score but will rarely produce a geometrically consistent inlier set.

3.5. Confidence-Gated Conservative Correction

A candidate update is accepted only if it satisfies all four gates in Table 2 simultaneously. The gates are deliberately conjunctive: each addresses a different failure mode, and a candidate that fails any one of them is discarded without further consideration.

Table 2. Map-update acceptance gates used in the present implementation.

Quantity	Condition	Failure Mode Addressed
Cross-view retrieval score	≥0.80	Wrong tile retrieved from the search region
Geometric inlier ratio	≥0.35	High-scoring tile without consistent geometry
Mean reprojection error of inliers	≤3.0 px	Degenerate or poorly conditioned alignment
Heading discrepancy vs. onboard estimate	≤8°	Rotationally ambiguous match over repetitive terrain

The thresholds in Table 2 were set from the development sequences and held fixed for all results reported in Section 5. They are conservative by design. The asymmetry of the cost function motivates this: a rejected good update costs nothing except a slightly longer interval of unaided drift, whereas an accepted bad update injects a large, unmodelled position and heading error into the global anchor. Under that asymmetry, it is rational to accept a lower true-positive rate in exchange for a low false-positive rate, and the acceptance rates of 74–81% reported in Section 5 should be read in that light rather than as a detection-performance deficit.

When a candidate passes all four gates, the back end updates the global anchor—the transformation between the local VIO frame and the map frame—in position and yaw. The correction is applied as a weighted update rather than a hard assignment, with the weight determined by the current global covariance relative to the covariance attributed to the map observation, so that a single accepted update cannot produce a discontinuity larger than the outstanding uncertainty justifies. The local VIO trajectory, its velocity, and its attitude are left untouched. Rejected candidates are logged with the score, inlier ratio, reprojection error, and heading discrepancy that caused the rejection; this log is what makes the acceptance statistics in Section 5 auditable after flight.

3.6. Extension to Intermittent, Uncertain, and Faulty GNSSs

The architecture described so far assumes that a GNSS is absent. As argued in Section 1, this is the endpoint of a continuum, and a system intended for real missions must also define its behaviour when a GNSS is present but of uncertain quality. This subsection specifies that behaviour. It is an architectural extension: it was not active during the experiments reported in Section 5, and no result in this paper depends on it.

The principle follows the robust integrated-navigation literature reviewed in Section 2.3. A GNSS measurement is not admitted to the estimator on the strength of its availability; it is first tested for consistency against the prediction produced by the VIO and map-anchored state. For a measurement z at epoch k with observation model h, the innovation and its covariance are

$$\nu_k = z_k - h\left(\hat{x}_k^-\right) \tag{1}$$

$$S_k = H_k P_k^- H_k^T + R_k \tag{2}$$

and the normalised innovation squared, equivalently the squared Mahalanobis distance of the innovation, is

$$d_k^2 = \nu_k^T S_k^{-1} \nu_k \tag{3}$$

Under the nominal Gaussian assumption, the statistic in Equation (3) follows a chi-square distribution with degrees of freedom equal to the measurement dimension, which makes threshold selection interpretable rather than arbitrary [27,29]. The decision is three-way. Below a soft threshold, the measurement is accepted with its nominal covariance. Between the soft and hard thresholds, it is neither accepted at face value nor discarded; instead, its covariance is inflated,

$$R_{k,\mathrm{eff}} = \kappa_k R_k, \quad \kappa_k = \min\left(\kappa_{\max}, \frac{d_k^2}{\tau_{\mathrm{warn}}}\right) \tag{4}$$

which reduces its weight in the update in the manner established for robust CKF and UKF formulations [15,19]. Above the hard threshold, the measurement is rejected for that epoch and the persistent-fault monitor is notified.

Table 3 gives chi-square thresholds at the 99% and 99.9% levels for the measurement dimensions most likely to be used. We emphasise that these are statistically interpretable starting points, not tuned optima: they have not been optimised on flight data in this work, and any deployment should calibrate them against logged innovations from the specific receiver, antenna installation, and flight profile.

Table 3. Recommended initial chi-square thresholds for the GNSS innovation-consistency test. These are starting values, not experimentally optimised parameters.

GNSS Update	Degrees of Freedom	Soft Threshold, 99%	Hard Threshold, 99.9%
Scalar test	1	6.63	10.83
Horizontal position	2	9.21	13.82
Three-dimensional position	3	11.34	16.27
Position and velocity	6	16.81	22.46

A per-epoch test cannot detect a fault that grows slowly enough to keep every individual innovation inside the gate, which is the case for a ramping bias or a slowly developing multipath condition. Following the double-channel and hypothesis-test-constrained approaches of Gao et al. [17,18], the per-epoch gate is therefore complemented by a sequential monitor that accumulates evidence across epochs, so that repeated marginal innovations or a consistent sign in the innovation sequence eventually declares the source faulty even though no single epoch triggers the hard gate. When that declaration occurs, GNSS updates are suspended and the system reverts to exactly the VIO-plus-verified-map configuration that Section 5 evaluates.

Recovery is deliberately asymmetric with respect to failure. Returning a source to service immediately after one good epoch invites oscillation between modes, so re-entry requires a run of consecutive nominal epochs—three is a reasonable starting value, again to be calibrated rather than assumed—after which the global anchor is corrected gradually rather than by a step. Every rejection is logged with the timestamp, the raw measurement, the innovation, the innovation covariance, the test statistic, the receiver status flags, and the active navigation mode, which is the minimum record needed to compute the detector-level metrics defined in Section 6.4. The complete decision flow is shown in Figure 2.

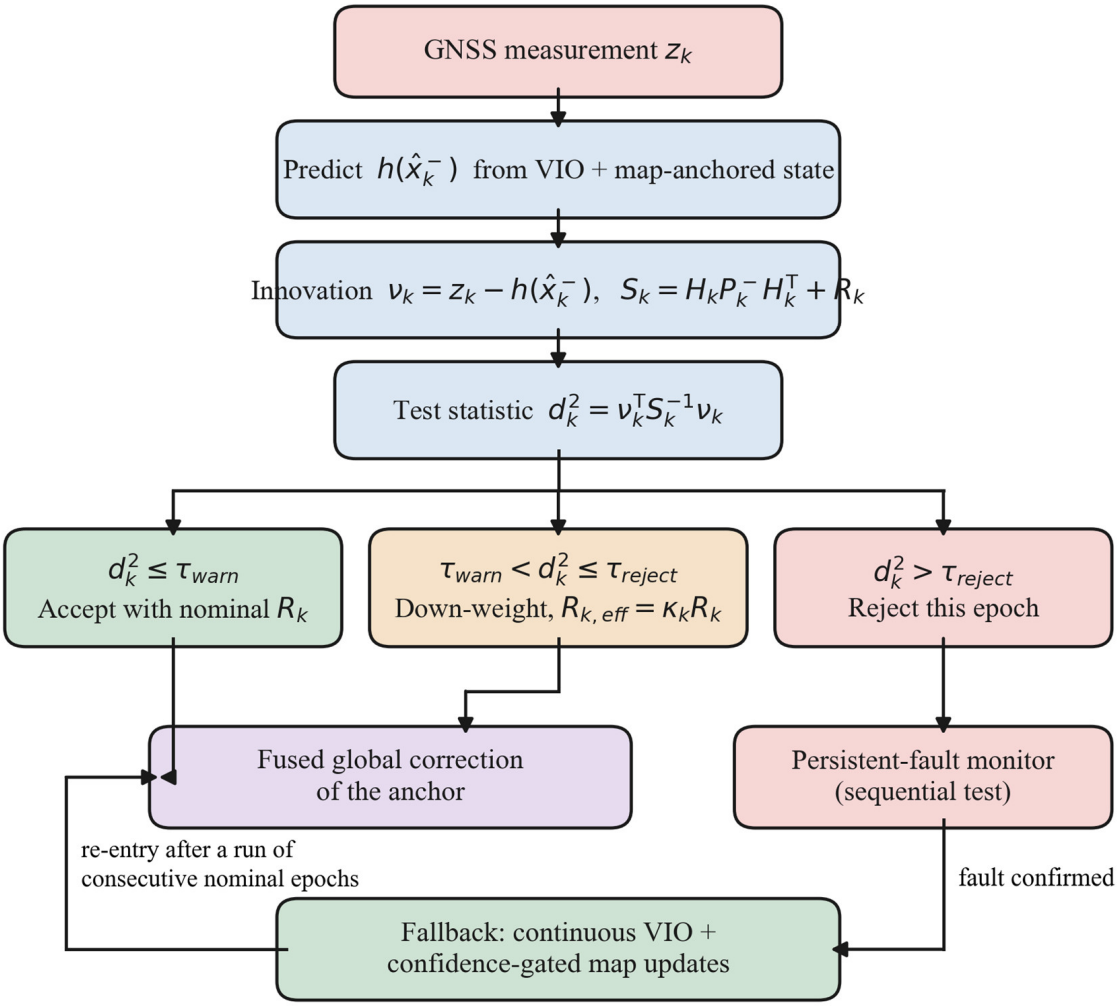


Figure 2. GNSS measurement validation. Each measurement is tested for innovation consistency before it is allowed to influence the state. Consistent measurements are accepted nominally, marginal measurements are down-weighted by covariance inflation, and inconsistent measurements are rejected and reported to a persistent-fault monitor. On rejection, the system continues in the VIO-plus-verified-map mode evaluated in Section 5; the local VIO trajectory is never reset.

3.7. Source Arbitration and Global Anchor Management

Once GNSS and map anchoring can both be active, the back end needs a rule for the case in which both offer a global correction at the same time. Averaging them is unwise, because the two sources fail in uncorrelated ways and an average of a good estimate and a badly wrong one is simply a less obviously wrong estimate. Instead, the two candidates are first tested against each other. Given the two global position estimates and their covariances, the difference and its normalised squared magnitude

$$r = p_{\mathrm{GNSS}} - p_{\mathrm{map}} \tag{5}$$

$$d^2 = r^T \left(P_{\mathrm{GNSS}} + P_{\mathrm{map}} \right)^{-1} r \tag{6}$$

indicate whether the two sources are mutually consistent. If they are, both are fused. If they are not, the estimator does not choose arbitrarily: it retains the source with the stronger independent evidence—a GNSS solution with good receiver integrity indication, or a map update that passed all four gates of Table 2 with a margin—holds the other as provisional, and requires corroboration from a subsequent epoch before acting on it. This mirrors the local-validation-before-global-fusion structure used in decentralised robust fusion [20].

The resulting operating envelope has four regimes. With a healthy GNSS, the receiver supplies the routine absolute correction and map matching remains at a low rate as a redundant check. With a degraded GNSS, the receiver's influence is reduced by covariance inflation, which correspondingly increases the relative authority of accepted map anchors. With a rejected or faulty GNSS, the source is disabled and the system operates in the VIO-plus-map mode. With complete denial, no GNSS measurement arrives at all and the behaviour is identical to the configuration evaluated in Section 5. In every regime, the local VIO trajectory is preserved and only the global anchor is modified.

4. Experimental Setup

4.1. Simulation Environment

The primary evaluation uses a high-fidelity fixed-wing simulation with a 24 km route flown at cruise, in which the GNSS is removed in a staged manner during the flight rather than at the outset, so that the estimator begins from a well-initialised state and the drift observed thereafter is attributable to the denial period. Imagery is rendered from a georeferenced source, and the map database used for retrieval is drawn from the same geographic area but is not the rendering source, so that query and map images differ in season, illumination, and resolution in a manner representative of an operational deployment.

4.2. Fixed-Wing Flight Sequence

The flight evaluation uses a fixed-wing sequence of 8.6 km with a camera–IMU payload and a recorded GNSS solution retained only as ground truth for evaluation; it is not supplied to the estimator during the GNSS-denied portion. The shorter distance relative to the simulation, together with real illumination, real vibration, and real map–image appearance differences, is the reason the flight numbers in Section 5 are consistently worse than the simulation numbers.

4.3. Onboard Computing and Baselines

All timing figures were obtained on an Intel NUC 8i7BEH-class onboard computer (Intel Corporation, Santa Clara, CA, USA) with the full pipeline running. Two baselines are used. The first is standalone VIO with the map branch disabled, which isolates the drift the map updates are intended to bound. The second is continuous map matching, in which retrieval and alignment are invoked at a fixed high rate rather than on motion-aware keyframes, which isolates the cost and benefit of the low-rate policy itself.

4.4. Evaluation Metrics

Final drift is reported as the terminal position error as a percentage of distance travelled, which is the convention that allows the 24 km and 8.6 km sequences to be compared. Absolute trajectory error is reported after alignment of the estimated and reference trajectories, following standard quantitative trajectory-evaluation practice for visual–inertial odometry [30]. The update success rate is the fraction of issued map queries whose candidate satisfied all four gates in Table 2. Throughput is reported as the sustained end-to-end frame rate of the complete pipeline on the target computer and map-matching workload as the number of retrieval-plus-alignment invocations over the sequence.

4.5. Scope of the Experimental Evidence

To avoid any ambiguity: the results in Section 5 characterise the VIO-plus-low-rate-map configuration under GNSS denial. The GNSS-validation branch of Section 3.6 and the source arbitration of Section 3.7 were not active in any reported run, no GNSS measurement was fused after the denial point, and no claim is made in this paper about their measured

performance. Section 6.4 specifies the experiments that would be required to make such a claim.

5. Results

5.1. Drift in Simulation

Over the 24 km simulated route, standalone VIO reached a final drift of 0.92%, corresponding to a terminal error of roughly 220 m. Enabling low-rate map anchoring reduced the final drift to 0.28%, with an absolute trajectory error of 22.4 m and 81% of issued map queries accepted. The qualitative difference matters as much as the ratio: the standalone error grows monotonically with distance, whereas the anchored error is bounded and resets toward the map's own accuracy each time an update is accepted, so that the error between anchors is set by the anchoring interval rather than by the total route length.

5.2. Drift in Flight

On the 8.6 km fixed-wing sequence, standalone VIO drifted 1.34% and the proposed method 0.41%, with 74% of queries accepted and an absolute trajectory error of 34.7 m. Both the higher drift and the lower acceptance rate relative to simulation are to be expected. Real imagery differs from the georeferenced map in season, illumination, and shadow direction; airframe vibration degrades feature quality; and the terrain along the route includes stretches of low-texture agricultural land where cross-view retrieval is genuinely ambiguous. The relevant observation is that the gates behaved as intended under those harder conditions: acceptance fell, but the accepted updates remained sound enough that the bounded-drift behaviour was preserved.

Figure 3 compares the final drift of standalone VIO and the proposed method in both the simulation and flight experiments.

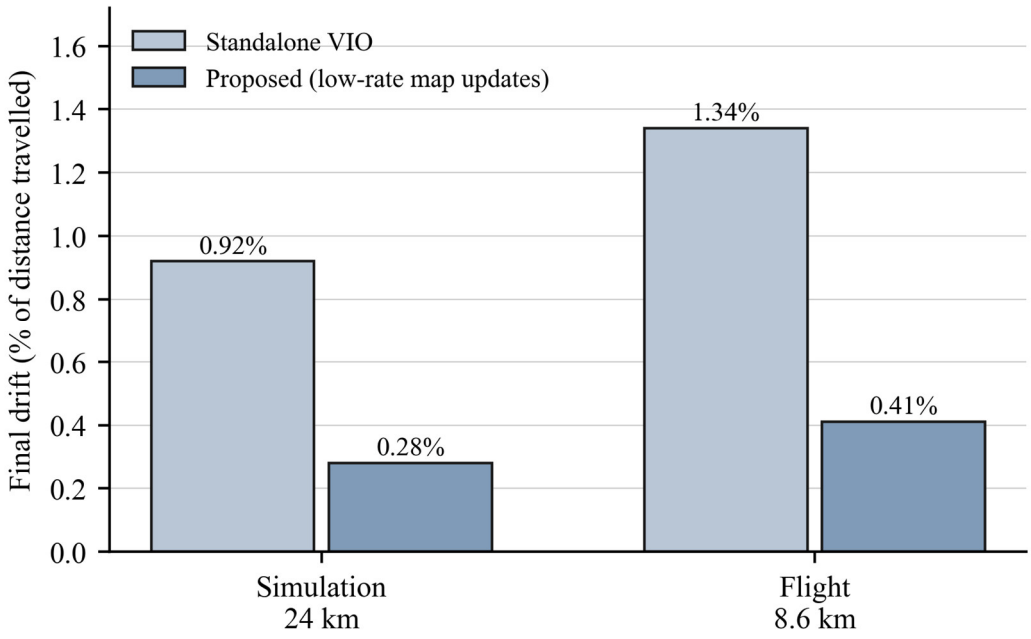


Figure 3. Final drift as a percentage of distance travelled, for standalone VIO and the proposed low-rate map-anchored method, on the 24 km simulated route and the 8.6 km flight sequence.

5.3. Low-Rate Versus Continuous Map Matching

Compared with continuous map matching, the low-rate policy reduced the number of map-matching invocations by 75% while preserving similar drift behaviour. This is the

central practical result of the paper. It indicates that on a fixed-wing platform in cruise the marginal value of an additional map query falls off sharply once an anchor has been established, because the intervening VIO solution is locally accurate and the aircraft's motion between anchors is smooth and well modelled. Consulting the map four times as often buys little accuracy, and it costs both computation and additional exposure to false-positive retrieval.

5.4. Onboard Computational Behaviour

The complete pipeline sustained 18 FPS on the Intel NUC 8i7BEH-class computer. Because map matching is invoked only on selected keyframes, the load profile is strongly bimodal: the continuous cost is that of the VIO front end alone, with retrieval and alignment appearing as infrequent bursts. This matters for integration on a real airframe, where the navigation stack shares the processor with guidance, control, and payload tasks, and where a predictable baseline load with occasional bursts is considerably easier to schedule than a uniformly high load.

Figure 4 shows the reduction in map-matching workload and the sustained throughput of the proposed system.

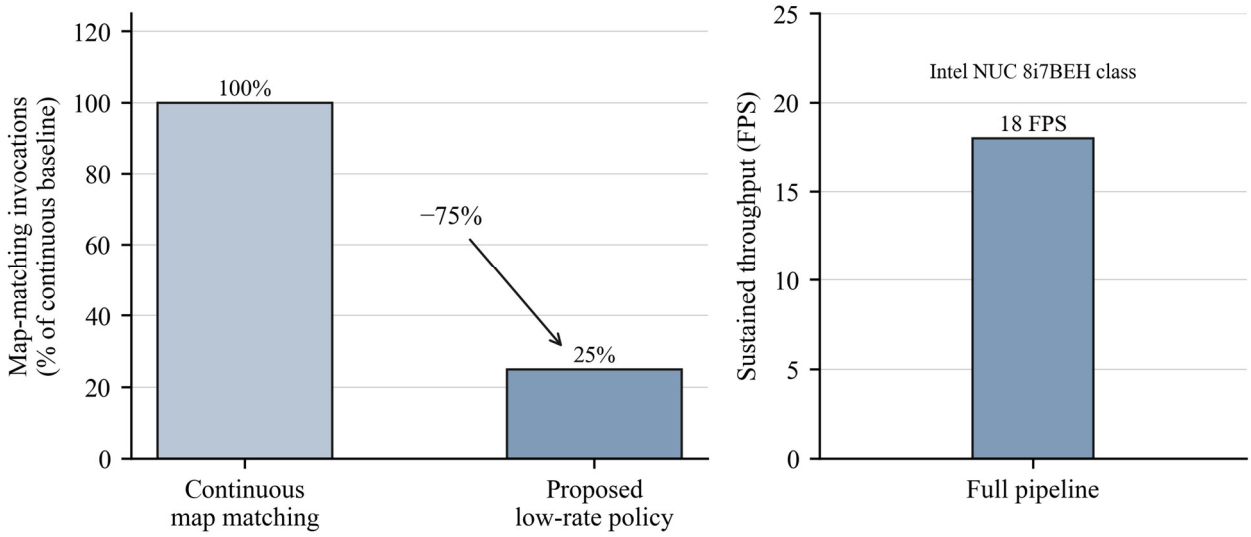


Figure 4. Relative map-matching workload for continuous matching and for the proposed low-rate policy, together with the sustained end-to-end throughput achieved by the full pipeline on the target onboard computer.

6. Discussion

6.1. Why Conservative Gating Matters

The acceptance rates in Table 4 are, on their face, modest: roughly one query in four or five is discarded. Interpreted as a detector, this looks like a limitation. Interpreted as a navigation policy, it is the intended behaviour, because the two error types have radically different costs. A discarded good update leaves the aircraft on its VIO solution for another anchoring interval, during which the error grows by a bounded and predictable amount. An accepted bad update writes an error of hundreds of metres into the global anchor and, on a fixed-wing aircraft following a planned route, that error reaches the guidance loop immediately. The conjunctive gate structure exists precisely to make the second event rare, and the price paid in acceptance rate is the correct price to pay.

The same asymmetry explains why the heading gate is worth its cost. Cross-view retrieval over regular field patterns, orchards, and grid-planned suburbs is prone to rota-

tionally ambiguous matches that are geometrically self-consistent and therefore pass an inlier-ratio test. Checking the recovered heading against the onboard attitude solution, which is independently observable from the IMU, catches this class of failure at negligible computational cost.

Table 4. Quantitative results for the GNSS-denied evaluation. Values are as previously reported and are unchanged in this revision.

Sequence	Method	Final Drift	Update Success Rate	Absolute Trajectory Error
Simulation, 24 km	Standalone VIO	0.92%	—	—
Simulation, 24 km	Proposed	0.28%	81%	22.4 m
Flight, 8.6 km	Standalone VIO	1.34%	—	—
Flight, 8.6 km	Proposed	0.41%	74%	34.7 m

6.2. Relation to Robust Integrated-Navigation Filtering

Viewed through the lens of Section 2.3, the map-update gates of Table 2 are a domain-specific measurement admission test, and the mechanism they implement is the same one that reliable measurement selection [23] applies to aiding sensors and that Mahalanobis-based robust filtering [15,19,27] applies to GNSS observations: establish whether a measurement is consistent with what the estimator already believes and modulate its influence accordingly. Our gates are cruder than a covariance-scaling rule, in the sense that they are conjunctive thresholds rather than a continuous weighting. The graded alternatives—fuzzy scaling [22], kernel-weighted covariance adaptation [24,25], and event-triggered adaptation [26]—are the natural direction in which to develop this component, and doing so would allow marginal map updates to contribute partial information instead of being discarded outright.

The low-rate invocation policy also has a counterpart in that literature. The dual-event-triggering scheme of Gao et al. [26] activates expensive adaptive machinery only when triggering conditions indicate anomalous measurements, in order to avoid paying its cost during nominal operation. Our motion-aware keyframe selector applies the same reasoning to a different expense: cross-view retrieval is invoked only when accumulated motion and grown uncertainty indicate that an anchor is worth its cost. We did not arrive at the design from that literature, but the convergence is worth noting, and it suggests that a formal event-triggering criterion derived from the covariance evolution would be a more principled replacement for the current motion-and-time heuristics.

6.3. Limitations

Several limitations should be stated plainly.

The GNSS-validation branch of Section 3.6 and the source arbitration of Section 3.7 are specified but not experimentally validated. They are consistent with established practice, but consistency with the literature is not evidence, and we make no performance claim for them.

The thresholds in Table 2 were tuned on development sequences from the same environments as the evaluation. Their transferability to substantially different terrain—desert, dense forest canopy, snow cover, or coastline—is untested, and we would expect the retrieval-score gate in particular to require recalibration.

The evaluation covers one simulated route and one flight sequence. This is adequate for a proceedings paper, but it is not a statistically meaningful sample, and single-sequence

results cannot separate the contribution of the method from the characteristics of the particular terrain flown.

The method presumes an up-to-date georeferenced orthophoto of the operating area. Where the map is stale relative to the ground—new construction, seasonal agricultural change, flooding, or post-disaster damage—retrieval and alignment will degrade, and the gates will respond by rejecting updates and allowing drift to grow. This is the safe failure mode, but it is still a failure mode.

Finally, the reported results were obtained under daylight conditions in reasonable visibility. Night operation, low sun angles, and cloud cover between the aircraft and the ground are outside the tested envelope.

6.4. Planned Validation Protocol for the GNSS Branch

To convert Section 3.6 from a specification into a validated component, we intend to run controlled fault injection on the recorded sequences. Table 5 lists the fault classes and what each is intended to exercise.

Table 5. Planned fault-injection protocol for validating the GNSS measurement-validation branch.

Injected Condition	Parameterisation	Property Exercised
Abrupt position outlier	5, 10, 20, 50 m step	Hard-gate detection of sudden faults
Slowly growing bias	Ramp equivalent to 0.1–1.0 m/s	Sequential monitor sensitivity to slow faults
Noise degradation	Covariance inflated 2×, 5×, 10×	Benefit of graded down-weighting
Burst anomaly	2–10 consecutive bad epochs	Separation of isolated outliers from persistent faults
Dropout	1, 5, 10, 30 s	Continuity of the VIO-plus-map fallback
Measurement latency	0.1–2.0 s	Sensitivity to stale or asynchronous measurements
Recovery	Clean measurements after a fault	Safety of hysteretic re-entry
GNSS–map conflict	Artificial offset between the two sources	Behaviour of source arbitration

Trajectory error alone is insufficient to evaluate a detector. The protocol therefore also records detector-level quantities: the false-alarm rate, defined as the fraction of healthy measurements rejected; the missed-detection rate, defined as the fraction of faulty measurements accepted; the detection delay between fault onset and alarm; the recovery time; and the magnitude of any position discontinuity at a source transition. Where the estimator covariance is trustworthy, the normalised estimation error squared should also be reported, since it evaluates the calibration of the uncertainty estimate and not merely the accuracy of the mean. Only after these quantities have been measured would it be appropriate to claim robustness to GNSS faults rather than to specify a mechanism for handling them.

7. Conclusions

This paper presented a drift-bounded navigation architecture for fixed-wing UAVs in GNSS-denied missions, built around a continuously running camera–IMU VIO front end anchored at a low rate by confidence-gated cross-view map updates. Over a 24 km simulated route, final drift fell from 0.92% for standalone VIO to 0.28%, with 81% of map queries accepted and a 22.4 m absolute trajectory error; over an 8.6 km flight sequence, it fell from 1.34% to 0.41%, with 74% acceptance and a 34.7 m absolute trajectory error. Relative to continuous map matching, the low-rate policy reduced map-matching invocations by 75% while preserving comparable drift behaviour and sustaining 18 FPS on an Intel NUC

8i7BEH-class onboard computer. The result supports the central claim: on a fixed-wing platform, a georeferenced map is most usefully treated as an occasional, verified correction rather than as a continuously active localisation source.

Responding to the observation that complete GNSS denial is only the extreme case of a broader set of GNSS fault conditions, the paper additionally specified an innovation-consistency layer that accepts, down-weights, or rejects GNSS measurements according to a normalised-innovation test, complemented by a sequential monitor for slowly growing faults and by an arbitration rule for the case in which GNSS and map updates disagree. This layer is an architectural extension and is not experimentally validated here; Section 6.4 defines the fault-injection experiments and detector-level metrics required to validate it, and that validation is the immediate next step of this work. Beyond it, replacing the conjunctive acceptance gates with a graded, covariance-based weighting and deriving the map-query trigger formally from the covariance evolution rather than from motion heuristics are the two changes we expect to yield the largest further improvement.

Author Contributions: Conceptualisation, A.H. and Z.Y.; methodology, A.H.; software, A.H.; validation, A.H., Z.Y. and S.A.; formal analysis, A.H. and Z.Y.; investigation, A.H.; resources, Z.Y. and S.A.; data curation, A.H.; writing—original draft preparation, A.H.; writing—review and editing, Z.Y. and S.A.; visualisation, A.H.; supervision, Z.Y.; project administration, Z.Y. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are available on request from the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Psiaki, M.L.; Humphreys, T.E. GNSS spoofing and detection. *Proc. IEEE* **2016**, *104*, 1258–1270. [CrossRef]
2. Mourikis, A.I.; Roumeliotis, S.I. A multi-state constraint Kalman filter for vision-aided inertial navigation. In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Rome, Italy, 10–14 April 2007; pp. 3565–3572.
3. Forster, C.; Carlone, L.; Dellaert, F.; Scaramuzza, D. On-manifold preintegration for real-time visual–inertial odometry. *IEEE Trans. Robot.* **2017**, *33*, 1–21. [CrossRef]
4. Qin, T.; Li, P.; Shen, S. VINS-Mono: A robust and versatile monocular visual-inertial state estimator. *IEEE Trans. Robot.* **2018**, *34*, 1004–1020. [CrossRef]
5. Geneva, P.; Eckenhoff, K.; Lee, W.; Yang, Y.; Huang, G. OpenVINS: A research platform for visual-inertial estimation. In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Paris, France, 31 May–31 August 2020; pp. 4666–4672.
6. Campos, C.; Elvira, R.; Gómez Rodríguez, J.J.; Montiel, J.M.M.; Tardós, J.D. ORB-SLAM3: An accurate open-source library for visual, visual–inertial, and multimap SLAM. *IEEE Trans. Robot.* **2021**, *37*, 1874–1890. [CrossRef]
7. Couturier, A.; Akhloufi, M.A. A review on absolute visual localization for UAV. *Robot. Auton. Syst.* **2021**, *135*, 103666. [CrossRef]
8. Hu, S.; Feng, M.; Nguyen, R.M.H.; Lee, G.H. CVM-Net: Cross-view matching network for image-based ground-to-aerial geo-localization. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Salt Lake City, UT, USA, 18–23 June 2018; pp. 7258–7267.
9. Shi, Y.; Liu, L.; Yu, X.; Li, H. Spatial-aware feature aggregation for image based cross-view geo-localization. In Proceedings of the Advances in Neural Information Processing Systems (NeurIPS 32), Vancouver, BC, Canada, 8–14 December 2019; pp. 10090–10100.
10. Zhu, S.; Yang, T.; Chen, C. VIGOR: Cross-view image geo-localization beyond one-to-one retrieval. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Nashville, TN, USA, 20–25 June 2021; pp. 3640–3649.
11. Goforth, H.; Lucey, S. GPS-denied UAV localization using pre-existing satellite imagery. In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Montreal, QC, Canada, 20–24 May 2019; pp. 2974–2980.

12. DeTone, D.; Malisiewicz, T.; Rabinovich, A. SuperPoint: Self-supervised interest point detection and description. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), Salt Lake City, UT, USA, 18–22 June 2018; pp. 224–236.
13. Sarlin, P.-E.; DeTone, D.; Malisiewicz, T.; Rabinovich, A. SuperGlue: Learning feature matching with graph neural networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA, 13–19 June 2020; pp. 4938–4947.
14. Fischler, M.A.; Bolles, R.C. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM* **1981**, *24*, 381–395.
15. Gao, B.; Hu, G.; Zhu, X.; Zhong, Y. A robust cubature Kalman filter with abnormal observations identification using the Mahalanobis distance criterion for vehicular INS/GNSS integration. *Sensors* **2019**, *19*, 5149. [CrossRef] [PubMed]
16. Gao, B.; Hu, G.; Zhong, Y.; Zhu, X. Cubature Kalman filter with both adaptability and robustness for tightly-coupled GNSS/INS integration. *IEEE Sens. J.* **2021**, *21*, 14997–15011. [CrossRef]
17. Gao, G.; Zhong, Y.; Gao, S.; Gao, B. Double-channel sequential probability ratio test for failure detection in multisensor integrated systems. *IEEE Trans. Instrum. Meas.* **2021**, *70*, 1–14. [CrossRef]
18. Gao, G.; Gao, B.; Gao, S.; Hu, G.; Zhong, Y. A hypothesis test-constrained robust Kalman filter for INS/GNSS integration with abnormal measurement. *IEEE Trans. Veh. Technol.* **2023**, *72*, 1662–1673. [CrossRef]
19. Hu, G.; Gao, B.; Zhong, Y.; Ni, L.; Gu, C. Robust unscented Kalman filtering with measurement error detection for tightly coupled INS/GNSS integration in hypersonic vehicle navigation. *IEEE Access* **2019**, *7*, 151409–151421. [CrossRef]
20. Hu, G.; Xu, L.; Gao, B.; Chang, L.; Zhong, Y. Robust unscented Kalman filter-based decentralized multisensor information fusion for INS/GNSS/CNS integration in hypersonic vehicle navigation. *IEEE Trans. Instrum. Meas.* **2023**, *72*, 1–11. [CrossRef]
21. Dong, X.; Hu, G.; Gao, B.; Zhong, Y.; Ruan, W. Windowing-based factor graph optimization with anomaly detection using Mahalanobis distance for underwater INS/DVL/USBL integration. *IEEE Trans. Instrum. Meas.* **2024**, *73*, 1–13. [CrossRef]
22. Hu, G.; Xu, L.; Yang, Z.; Gao, B.; Zhong, Y. Indirect fuzzy robust cubature Kalman filter with normalized input parameters. *IEEE Trans. Aerosp. Electron. Syst.* **2024**, *60*, 5880–5890. [CrossRef]
23. Gao, B.; Yin, H.; Hu, G.; Zhong, Y. Reliable measurement selection mechanism-based tightly coupled inertial/bionic polarization integration with position correction. *Aerosp. Sci. Technol.* **2024**, *154*, 1–14. [CrossRef]
24. Gao, G.; Yi, Y.; Zhong, Y.; Liu, Z. A robust Kalman filter based on kernel density estimation for system state estimation against measurement outliers. *IEEE Trans. Instrum. Meas.* **2025**, *74*, 1–12. [CrossRef]
25. Gao, G.; Zhong, Y.; Gao, B.; Hu, G.; Peng, X. Kernel function-weighted robust Kalman filter for integrated navigation system. *IEEE Trans. Veh. Technol.* **2026**, *75*, 7544–7556. [CrossRef]
26. Gao, B.; Ma, P.; Hu, G.; Zhong, Y.; Liu, Z. Dual-event-triggering ANFIS-based unscented Kalman filter for cluster cooperative navigation with measurement anomalies. *Chin. J. Aeronaut.* **2026**, *39*, 103968. [CrossRef]
27. Chang, G. Robust Kalman filtering based on Mahalanobis distance as outlier judging criterion. *J. Geod.* **2014**, *88*, 391–401. [CrossRef]
28. Yin, H.; Gao, B.; Hu, G.; Zhong, Y.; Liu, Z. Robust navigation for INS/GNSS/BPNS integration using adaptive weighted q-Rényi kernel mixture correntropy filter. *IEEE Trans. Instrum. Meas.* **2026**, *in press*.
29. Bar-Shalom, Y.; Li, X.-R.; Kirubarajan, T. *Estimation with Applications to Tracking and Navigation: Theory, Algorithms and Software*; Wiley: New York, NY, USA, 2001.
30. Zhang, Z.; Scaramuzza, D. A tutorial on quantitative trajectory evaluation for visual(-inertial) odometry. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Madrid, Spain, 1–5 October 2018; pp. 7244–7251.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.