algorithms

MDPI

Article

Heteroscedastic Decoupling Algorithm of Gyroscope Front-End Preprocessing for UAVs Under Collision Disturbance

Ying Wei [1,2,*], Ruoqing Duan [3], Boyao Wang [3] and Qihong Duan [3,*]

[1] School of Astronautics, Beihang University, Beijing 100191, China
[2] Xi'an Aerospace Propulsion Testing Technique Institute, Xi'an 710025, China
[3] School of Mathematics and Statistics, Xi'an Jiaotong University, Xi'an 710049, China; 2243211109@stu.xjtu.edu.cn (B.W.)
* Correspondence: weiying19888@126.com (Y.W.); k.h.tuan@xjtu.edu.cn (Q.D.)

Abstract

Small cargo unmanned aerial vehicles (UAVs) operating in narrow confined spaces suffer instantaneous collision impacts that induce gyroscope heteroscedastic noise and abrupt angular velocity derivatives, triggering severe dynamic attitude errors in traditional strapdown inertial navigation systems (SINS). Existing algorithms separate angular velocity fitting and noise suppression, adopt unified three-axis weighting, and lack adaptive segmentation for collision disturbances, limiting navigation accuracy without raising computational costs. This paper proposes an integrated heteroscedastic decoupling algorithm for UAV SINS under collision interference. Hermite orthogonal polynomials are utilized to fit non-stationary angular velocity with derivative matching constraints, an optimized single-pass CUSUM detector with steady-state residual compensation is proposed to identify collision-induced variance change points. An axis-differentiated weighting strategy is developed to suppress heteroscedastic noise. Recursive least-squares is adopted to lower online computation overhead. Multi-condition coning motion simulations show Hermite polynomials achieve the lowest attitude RMSE under steady flight; the improved CUSUM detector delivers shorter detection delay, fewer false alarms, and lighter computation than mainstream detection methods, and segmented differentiated weighting eliminates collision-induced noise distortion at the raw measurement stage. The proposed algorithm unifies signal fitting and noise correction with minimal computational overhead, effectively mitigating dynamic errors for lightweight airborne navigation hardware and offering a high-precision front-end preprocessing solution for cargo UAVs operating in cluttered obstacle environments. The proposed algorithm is positioned as a gyro-only front-end preprocessing module; accelerometer-related error compensation and full multi-sensor back-end integration are addressed in ongoing work.

Keywords: UAV collision; heteroscedastic noise; CUSUM change-point detection; Hermite orthogonal polynomial; attitude calculation

Check for updates

Academic Editor: Roberto Montemanni

Received: 26 July 2026
Revised: 29 August 2026
Accepted: 30 August 2026
Published: 3 September 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

Small cargo unmanned aerial vehicles (UAVs) serve as core equipment for end-to-end material delivery. When operating in confined spaces such as alleyways, narrow factories, and dense trees, they are prone to instantaneous collisions with walls and obstacles [1]. Short-duration strong impact loads generated at the moment of collision are directly transmitted to the onboard inertial measurement unit (IMU), inducing two typical dynamic

disturbances: first, impact impulse causes an instantaneous abrupt change in the amplitude of carrier angular velocity; second, the impact excites high-frequency vibration of the airframe, triggering drastic jumps in the first and second derivatives of angular velocity. During collision-free steady cruising, the aerodynamic load of the UAV remains stable, the angular velocity changes gently, and the fluctuations of all-order derivatives are minimal.

The strapdown inertial navigation system (SINS) relies on no external signals and can continuously output state information including angular velocity and attitude. It acts as the core perception foundation for UAV flight control, anti-collision decision-making, and cargo attitude monitoring. Instantaneous impacts induced by obstacle collisions introduce numerous noise variance change points in gyroscope observation signals. The vibration intensifies during impact periods, leading to a remarkable rise in noise intensity, while the noise level stays low in steady flight, resulting in obvious heteroscedastic characteristics of observation noise. Angular velocity serves as the underlying input for carrier state interpolation and attitude calculation. If the observed angular velocity is distorted under collision disturbances, subsequent equivalent rotation vector calculation will continuously accumulate coning and sculling dynamic errors, leading to attitude judgment deviations, and greatly increasing the risk of secondary collisions.

Traditional Euler angle method, quaternion method, and classic multi-sample equivalent rotation vector algorithms all assume that the angular velocity within the calculation update interval is constant or linearly varying and are only applicable to impact-free and low-dynamic working conditions. Direct application to cargo UAV collision scenarios will produce severe theoretical errors. Although simply increasing the number of calculation samples can slightly improve fitting accuracy, it will significantly raise the computational burden of airborne embedded devices, failing to meet the lightweight and low-power hardware constraints of UAVs. Therefore, without substantially boosting computational complexity, constructing an integrated algorithm for high-precision angular velocity fitting and attitude decoupling targeting angular velocity abrupt changes and heteroscedastic noise induced by collision impacts is an urgent engineering problem to be solved for autonomous navigation of cargo UAVs operating in narrow confined spaces.

Current research on high-precision SINS calculation mainly falls into two categories: angular velocity basis function fitting and inertial measurement noise suppression. In the field of angular velocity fitting, existing studies mostly adopt Legendre polynomials and sinusoidal basis functions to fit periodically steady angular velocity, which can effectively reduce dynamic errors under standard noise-free coning motion [2,3]. Chebyshev and Hermite orthogonal polynomials possess unique advantages of optimal uniform approximation and high-order derivative smooth interpolation, respectively, providing new ideas for non-stationary angular velocity fitting. Nevertheless, existing literature only verifies their performance under ideal smooth motion scenarios, without adaptive research targeting abrupt changes in angular velocity and its derivatives caused by UAV collisions [4]. In terms of noise processing, adaptive Kalman filters can only track gradual noise variations slowly and respond sluggishly to instantaneous variance jumps induced by collisions [5,6]. Weighted least squares (WLS) is a classic method to handle heteroscedasticity, yet existing schemes adopt unified weighting for three axes, ignoring the motion difference between oscillating X/Y axes and Z-axis, with constant rotation components, which tends to destroy the inherent constant component of the Z-axis [7,8]. Our tests show that a single-pass cumulative sum of squares (CUSUM) algorithm [9] can detect variance change points in time series, but its original critical value tends to misjudge periodic oscillations of X/Y axes as noise mutations, and optimization strategies for collision impact scenarios remain to be improved.

https://doi.org/10.3390/a19090752

Synthesizing existing technologies, three core deficiencies remain for raw-gyro-measurement-oriented preprocessing oriented to UAV collision scenarios. First, there is a lack of dynamic models oriented to collision impacts of cargo UAVs, and no algorithm optimization targeting the physical characteristic of abrupt angular velocity derivatives induced by collisions. Second, most existing joint noise-suppression and attitude-estimation frameworks operate at the attitude-state back-end; few studies construct an integrated front-end workflow for raw gyroscope data, which combines change-point detection and segmented weighted preprocessing. Third, differentiated noise correction strategies corresponding to the motion characteristics of three axes have not been designed, and quantitative performance comparisons between Chebyshev and Hermite polynomials under impact noise scenarios are insufficient. Meanwhile, most existing algorithms focus on error compensation at the attitude level, failing to suppress underlying deviations introduced by collision disturbances through preprocessing raw angular velocity observations, which imposes an obvious upper limit on error suppression performance.

Aiming at the engineering pain points of abrupt angular velocity and time-varying heteroscedastic observation noise triggered by obstacle collisions of cargo UAVs, this paper constructs an integrated SINS decoupling algorithm combining optimized single-pass CUSUM change-point detection, axis-differentiated weighted correction, and orthogonal polynomial angular velocity fitting. The existing building blocks including Hermite orthogonal polynomials, recursive least squares, weighted-least-squares form the basis of our algorithm. Our genuine technical innovations are summarized as follows:

1. An optimized single-pass CUSUM variance change-point detector embedded with steady-state residual compensation term is developed. This modification suppresses false alarms triggered by periodic oscillation of X/Y axes during steady flight.
2. An axis-differentiated segmented weighting strategy is designed according to three-axis motion characteristics of UAVs. Amplitude-based weighting is applied for oscillatory X/Y axes, while mean-removed relative-change-rate weighting is adopted for the Z-axis to preserve its constant rotation component.
3. An end-to-end online integrated preprocessing pipeline is constructed: orthogonal-polynomial angular-velocity fitting, improved CUSUM change-point monitoring, axis-differentiated weighting, and segmented refitting re-initialization are tightly coupled. Once a collision change point is detected, fitting matrices and coefficients are re-initialized locally to block error propagation across collision moments. Such an integrated workflow for raw gyroscope measurement preprocessing under UAV collision disturbance has rarely been reported in previous literature.

Based on the above innovations, multi-condition simulations are carried out to verify the overall performance of the proposed algorithm.

The paper is organized as follows. Section 2 reviews fundamental strapdown inertial navigation theories, including carrier attitude representation and the mathematical relationship between angular velocity and quaternion kinematics. Section 3 establishes the gyroscope angular velocity observation model under steady cruising and introduces incremental recursive least squares to reduce real-time computational load. Then we construct the heteroscedastic noise model triggered by UAV collision shocks, detail the optimized single-pass CUSUM variance change-point detection scheme, and present the axis-differentiated segmented weighted correction strategy for raw gyroscope data preprocessing. Section 4 designs multi-scenario coning motion simulations to quantitatively compare the attitude solving accuracy of different fitting polynomials and benchmark the detection performance of multiple change-point identification algorithms. Section 5 discusses the experimental results against prior literature, analyzes the core merits and inherent limitations of the proposed integrated decoupling algorithm, and outlines targeted

future research directions. Section 6 concludes the full work and summarizes the core theoretical and engineering contributions of this study.

2. Background of SINS

The spatial orientation of the carrier coordinate system (frame b, rigidly attached to the UAV airframe) relative to the navigation reference coordinate system (frame n) can be described via multiple mathematical tools. Three mainstream schemes adopted in engineering are Euler angles, direction cosine matrices, and quaternions.

Euler angles decompose carrier attitude through three sequential rotations and intuitively describe the airframe using yaw, pitch, and roll angles. Nevertheless, they suffer from the gimbal lock singularity: when the pitch angle approaches $\pm 90^\circ$, the differential equations become unsolvable due to a zero denominator, which prevents continuous full-attitude calculation for UAVs.

The direction cosine matrix represents coordinate transformations via a 9-dimensional square matrix without singularities. However, it requires solving nine coupled differential equations simultaneously, leading to heavy iterative computation that imposes high computing pressure on the embedded airborne hardware of UAVs, limiting its practical engineering application.

A quaternion is a four-dimensional hypercomplex number consisting of one real part q_0 and three imaginary parts q_1, q_2, and q_3. A unit quaternion q uniquely characterizes a single rigid-body rotation.

Gyroscopes output the three-axis angular velocity $\boldsymbol{\omega}^b = \left(\omega_x^b, \omega_y^b, \omega_z^b\right)^T$ under frame b. We write a quaternion $\omega^b = \omega_x^b \boldsymbol{i} + \omega_y^b \boldsymbol{j} + \omega_z^b \boldsymbol{k}$. There is a differential equation linking angular velocity and attitude quaternions as follows.

$$\frac{dq}{dt} = \frac{1}{2} q \omega^b. \tag{1}$$

Theoretically, if accurate angular velocity measurements free of disturbances are available at all time instants, the attitude quaternions can be derived by solving the above differential equations. However, the close-form solution to q is impossible for any non-trivial case. When cargo UAVs operate in narrow spaces, moreover, instantaneous collisions with obstacles are inevitable. The resulting impact vibration contaminates gyroscope measurements with time-varying noise, whose variance jumps under collision conditions and generates heteroscedastic observation disturbances. Finally, the quaternion differential equation inherently poses challenges in numerical solution.

The core objective of this research is to suppress heteroscedastic noise induced by collisions and calibrate contaminated angular velocity observations to recover the true evolution law of angular velocity ω^b, thereby eliminating fundamental errors of attitude calculation at the input source. We then construct a well-fitted continuous angular velocity function and obtain angular velocity estimates at dense time instants higher than the original sampling frequency, laying a solid foundation for high-accuracy numerical solving.

3. Observation Model and the Proposed Algorithm

3.1. Angular Velocity Observation Model in Steady Cruise

When a cargo UAV operates in collision-free steady cruise, the airframe vibration and external airflow disturbance are mild and stable. The measurement noise output by the gyroscope satisfies the homoscedastic normal white noise assumption. The angular velocity observation model for each component of the gyroscope measurements is the same one as

$$\widetilde{\omega}_x(t) = \omega_x^b(t) + \varepsilon(t), \tag{2}$$

where $\widetilde{\omega}(t) = \left(\widetilde{\omega}_x(t), \widetilde{\omega}_y(t), \widetilde{\omega}_z(t)\right)^T$ denotes the measured angular velocity of the gyroscope, $\omega^b(t)$ represents the true angular velocity of the carrier, and $\varepsilon(t)$ is the observation noise of this dimension under steady working conditions. The noise $\varepsilon(t) \sim N(0, \sigma_0^2)$ follows a zero-mean independent normal distribution, and the noise terms of different dimensions are mutually independent.

Without impact loads during steady cruise, the carrier angular velocity is continuous and smooth, with gentle variations in all-order derivatives and no abrupt changes, which belongs to a slowly varying continuous function. From the perspective of approximation theory, any continuous function can be fitted and approximated by an arbitrary set of basis functions with high flexibility in basis selection. Hence, the true angular velocity is expressed as a linear combination of basis functions as

$$\omega_{\cdot}(t) = \sum_{k=1}^{m} a_k \phi_k(t) \tag{3}$$

where $\omega_{\cdot}(t)$ is one of $\omega_x(t), \omega_y(t), \omega_z(t)$, m is the number of the selected function, ϕ_k is the k-th basis function, and a_k stands for the undetermined fitting coefficient for the corresponding dimension.

Nevertheless, special polynomials with excellent orthogonality and approximation properties can be chosen to simplify the solution of coefficients and reduce overall computational complexity. Three classical orthogonal polynomials are selected for angular velocity fitting in this paper, namely Chebyshev polynomial, Hermite polynomial and Legendre polynomial. Chebyshev polynomial owns the optimal uniform approximation property, with evenly distributed fitting errors within the interval, which is suitable for wideband slowly varying angular velocity signals. Hermite polynomial can simultaneously match function values and all-order derivatives, delivering outstanding interpolation accuracy for smooth angular velocity with slight derivative fluctuations. Finally, Legendre polynomial achieves the minimum mean square approximation error within the interval, applicable to fitting tasks taking mean square error as the optimization indicator.

3.2. Least Squares Fitting with Equispaced Sampling

The navigation computer samples angular velocity at fixed equal intervals Δt, with sampling moments:

$$t_i = i\Delta t,\ i = 1, 2, \ldots, n, \tag{4}$$

where n is the total number of samples in the current batch, and the corresponding observation value is ω~i. Substitute the discrete sampling moments into the angular velocity fitting expression of a dimension to obtain the discrete observation equations:

$$\widetilde{\omega}_i = \sum_{k=1}^{m} a_k \phi_k(t_i) + \varepsilon_i. \tag{5}$$

The corresponding compact matrix form is:

$$\widetilde{\boldsymbol{\Omega}} = \boldsymbol{\phi}\boldsymbol{a} + \boldsymbol{\varepsilon}, \tag{6}$$

where $\widetilde{\boldsymbol{\Omega}} = \left(\widetilde{w}_1, \widetilde{w}_2, \ldots, \widetilde{\omega}_n\right)^T$ is the observation vector for this dimension composed of all measurement, $\boldsymbol{\phi} \in \mathbb{R}^{n \times m}$ is the design matrix constructed by the values of basis functions at each sampling point, and $\boldsymbol{a} = (a_1, a_2, \ldots, a_m)^T$ is the m-dimensional undetermined coefficient vector.

Ordinary least squares are adopted with the optimization target of minimizing the 2-norm of residuals $\min\|\widetilde{\boldsymbol{\Omega}} - \boldsymbol{\phi}\boldsymbol{a}\|_2^2$, and the normal equation is:

$$\boldsymbol{\phi}^T\boldsymbol{\phi}\boldsymbol{a} = \boldsymbol{\phi}^T\widetilde{\boldsymbol{\Omega}} \tag{7}$$

If $\boldsymbol{\phi}^T\boldsymbol{\phi}$ is invertible, the coefficient vector can be directly solved by:

$$\widehat{\boldsymbol{a}} = (\boldsymbol{\phi}^T\boldsymbol{\phi})^{-1}\boldsymbol{\phi}^T\widetilde{\boldsymbol{\Omega}} \tag{8}$$

Batch offline fitting solves all coefficients $\widehat{\boldsymbol{a}}$ at one time using a complete segment of historical samples $\widetilde{\boldsymbol{\Omega}}$, yielding a continuous angular velocity-fitting function:

$$\widehat{\omega}(t) = \sum_{k=1}^{m} \widehat{a}_k \phi_k(t) \tag{9}$$

within the interval. Interpolation can be implemented between original discrete sampling points to obtain angular velocity at arbitrary moments, providing continuous input for high-precision numerical solution of quaternion differential equations.

It should be noted that the numerical conditioning of polynomial-basis matrix is affected by polynomial order m. As m increases, the condition number of the basis matrix tends to rise, which may bring numerical stability risks for real-time embedded computation. In this paper, moderate polynomial order is selected to balance fitting accuracy and numerical robustness. Excessively high polynomial order should be avoided in practical onboard deployment.

3.3. Accelerated Online Recursive Coefficient Update

The UAV navigation system collects data in real time for a long duration, and a new set of angular velocity samples is obtained every Δt. If full batch least squares is executed after each new sample arrives, all historical observations need to be reassembled to construct the matrix and perform matrix inversion. As the total number of samples n increases continuously, repeated batch solving introduces massive computational overhead, which fails to meet the low-power and real-time constraints of airborne embedded devices. To address this issue, an incremental recursive least squares algorithm is introduced, which updates m coefficients $a_1, a_2, \ldots, a_m$ only using incremental information of new samples without repeated global matrix inversion.

The core operation of one complete batch coefficient solution is the inversion of the m-order square matrix, with a single computational complexity of $O(m^3)$, as well as incremental recursion only processes a single new sample, and the computational complexity of a single update is reduced to $O(m^2)$. This optimization only cuts the iteration overhead after adding new samples, while the inherent matrix dimension overhead brought by fixed polynomial order m remains unchanged.

From the first k samples $\widetilde{w}_1$, $\widetilde{w}_2$, …, $\widetilde{\omega}_k$, we define the accumulated normal equation terms as $\boldsymbol{\phi}_k, \widetilde{\boldsymbol{\Omega}}_k$. The estimated coefficients $\widehat{\boldsymbol{a}}_k$ solved by in the first k samples are:

$$\widehat{\boldsymbol{a}}_k = (\boldsymbol{A}_k)^{-1}\boldsymbol{b}_k,\ \boldsymbol{A}_k = \boldsymbol{\phi}_k^T\boldsymbol{\phi}_k,\ \boldsymbol{b}_k = \boldsymbol{\phi}_k^T\widetilde{\boldsymbol{\Omega}}_k. \tag{10}$$

When the $(k+1)$-th sample is added, we have new observation value $\widetilde{\omega}_{k+1}$ and the corresponding basis row vector $\phi_{k+1} = (\phi_1(t_{k+1}), \phi_2(t_{k+1}), \ldots, \phi_m(t_{k+1}))^T$. The incremental update rules of matrices are:

$$\boldsymbol{A}_{k+1} = \boldsymbol{A}_k + \phi_{k+1}\phi_{k+1}^T,\ \boldsymbol{b}_{k+1} = \boldsymbol{b}_k + \phi_{k+1}^T\,\widetilde{\omega}_{k+1}. \tag{11}$$

With the Woodbury matrix inversion lemma:

$$(\boldsymbol{A}_{k+1})^{-1} = (\boldsymbol{A}_k)^{-1} - \frac{(\boldsymbol{A}_k)^{-1}\phi_{k+1}\phi_{k+1}^T(\boldsymbol{A}_k)^{-1}}{1+\phi_{k+1}^T(\boldsymbol{A}_k)^{-1}\phi_{k+1}} \tag{12}$$

can be recursively derived only based on the stored $(\boldsymbol{A}_k)^{-1}$ from the previous step, without full re-inversion of the m-order square matrix. The final recursive formula of coefficients is:

$$\widehat{\boldsymbol{a}}_{k+1} = \widehat{\boldsymbol{a}}_k + \boldsymbol{A}_{k+1}^{-1}\phi_{k+1}\left(\widetilde{\omega}_{k+1} - \phi_{k+1}^T\widehat{\boldsymbol{a}}_k\right). \tag{13}$$

We further analyze the computational complexity for the full online workflow of the proposed algorithm.

(1) Hermite polynomial basis evaluation for each sampling point: $O(m)$, where m denotes the polynomial order.
(2) Recursive least-squares coefficient update: $O(m^2)$. This is the main computational burden under normal steady-flight conditions.
(3) Improved single-pass CUSUM change-point recursion: $O(1)$ per time step, only simple scalar accumulation and comparison operations.
(4) Axis-differentiated weighting correction: $O(1)$ for three axes.
(5) Segmented re-initialization: This operation is triggered only when a collision change point is detected. Its one-time computational cost is $O(m^2)$ and will not be invoked in steady cruise.

In summary, during long-term steady flight without collision events, the dominant complexity is $O(m^2)$, which is favorable for lightweight embedded navigation hardware.

3.4. Heteroscedastic Noise and Change-Point Detection Under Collision Disturbances

The angular velocity modeling and online fitting methods established are all based on the homoscedastic normal white noise assumption under steady cruise conditions. When instantaneous obstacle collisions occur, the impact-induced airframe vibration causes a single abrupt jump in the gyroscope noise variance, presenting typical heteroscedastic characteristics. Meanwhile, the collision impulse triggers sudden jumps in angular velocity amplitude and high-order derivatives, invalidating the global unified least squares fitting model. Therefore, change-point detection is essential to divide the time series into steady cruise segments and collision strong-noise segments. The typical working characteristics of cargo UAVs is of long-term steady cruising with occasional single collision and single variance mutation. For this scenario, the first variance change point corresponding to the collision moment needs to be captured online. Accordingly, we adopt the single-pass cumulative sum of squares (CUSUM) algorithm to achieve change-point detection.

In contrast to other iterative change-point detection algorithms (e.g., ICSS [9]), which require iterative sequence segmentation and multi-point traversal searching, the squared CUSUM algorithm [10] constructs detection statistics based on squared residuals, which can effectively identify variance mutations and match the noise amplitude jump characteristics caused by UAV collisions. The complete algorithm principles and implementation procedures are presented as follows.

First, the angular velocity fitting residual for a given dimension at each moment is calculated based on the steady-state fitting model:

$$e_i = \widetilde{\omega}_i - \widehat{\omega}_i \tag{14}$$

where $e_{i.}$ is one of e_{ix}, e_{iy}, e_{iz} and:

$$\widehat{\omega}_i = \sum_{k=1}^{m} \widehat{a}_k \phi_k(t_i) \tag{15}$$

represents the angular velocity estimation obtained by orthogonal polynomial fitting with the estimated coefficient vector $\widehat{\boldsymbol{a}} = (\widehat{a}_1, \widehat{a}_2, \ldots, \widehat{a}_m)^T$. Under steady cruise conditions, the residuals $e_i = e_{ix}\boldsymbol{i} + e_{iy}\boldsymbol{j} + e_{iz}\boldsymbol{k}$ follow a zero-mean homoscedastic normal distribution $e_{ix}, e_{iy}, e_{iz} \sim N(0, \sigma^2)$, and the squared residuals fluctuate within a stable low-amplitude range.

The single-pass CUSUM cumulative detection statistic is constructed as:

$$S_i = \max\left\{0, S_{i-1} + e_i^2 - \widehat{\sigma}_0^2\right\}, \; S_0 = 0, \tag{16}$$

where S_i is the cumulative squared statistic at the period i, and $\widehat{\sigma}_0^2$ denotes the steady-state offset compensation term, which eliminates inherent trivial fluctuations of squared residuals and avoids invalid statistic accumulation during steady flight. Using the noise baseline data collected in collision-free cruising, the mean and standard deviation of squared residuals are calculated to construct a data-driven detection threshold:

$$H = k\widehat{\sigma}_0^2 \tag{17}$$

The coefficient $k \in [3, 5]$ is set empirically, balancing a low false-alarm rate and high mutation sensitivity. The online detection criterion is implemented via unidirectional recursive updating. If $S_n > H$, the current moment n is judged as the first variance change point, namely the starting time of collision disturbance; otherwise, the flight state is regarded as steady cruising, and the detection continues iteratively.

Notably, a segmented online refitting mechanism is designed to cooperate with change-point detection. Once the CUSUM algorithm captures the first variance mutation point, the iteration of the previous steady-state fitting process is terminated immediately. All polynomial fitting coefficients and cumulative statistical matrices are reinitialized at the detected change point to restart local segment fitting for the post-collision stage. This reset mechanism completely isolates the interference of pre-collision steady-state model parameters on strong-noise segments, adapts to the abrupt variations of noise variance and angular velocity characteristics after collision, and guarantees the local optimality of segmented fitting. Combined with the axis-differentiated weighted strategy for different segments, the proposed method effectively suppresses heteroscedastic noise distortion, optimizes the raw angular velocity observation accuracy at the source, and provides reliable data support for high-precision SINS attitude calculation.

4. Simulation Results

4.1. Simulation Setup and Evaluation Metrics

We use the self-rotation superimposed conical swing angular velocity model with three-axis independent self-rotational angular velocities to generate kinematic angular velocity, and the model is expressed as follows:

$$\omega(t) = \begin{pmatrix} \omega_x(t) \\ \omega_y(t) \\ \omega_z(t) \end{pmatrix} = \begin{pmatrix} \Omega \sin\alpha \cos(\omega_0 t) + \omega_{sx} \\ \Omega \sin\alpha \sin(\omega_0 t) + \omega_{sy} \\ \Omega \cos\alpha + \omega_{sz} \end{pmatrix}. \tag{18}$$

The overall motion form remains unchanged after the UAV is subjected to collision disturbance, and model parameters are adjusted to switch motion states, satisfying the physical motion characteristics that the angular velocity keeps continuous at the collision

moment, while its first-order time derivative undergoes abrupt change. Set the collision moment as t_c. The core idea of series splicing is to give core parameters Ω, α, and ω_0 of conical swing motion and then adjust the three-axis self-rotational angular velocities ω_{sx}, ω_{sy}, ω_{sz} of the two segments before and after collision. By equating the angular velocity values $\omega(t)$ of the two segments at time $t = t_c$, a continuous time series connection is realized. Meanwhile, the parameter adjustment naturally leads to inconsistent derivatives on both sides, which achieves the simulation effect of derivative mutation and conforms to the actual motion law that the attitude remains unchanged while the motion change rate alters suddenly after collision.

Taking the noise-free piecewise continuous reference angular velocity as input, standard reference attitudes are solved by the attitude kinematic differential Equation (4) combined with high-precision numerical algorithms. The obtained angular velocity and attitude sequences are regarded as true values, which lay the benchmark for subsequent quantitative evaluation.

Zero-mean homoscedastic normal white noise is uniformly superimposed on the noise-free reference angular velocity series to construct the observed angular velocity data, and the noise intensity remains constant over the whole time span. After the collision induces an abrupt change in the first-order derivative of angular velocity, smooth polynomial approximation cannot well adapt to such non-smooth time series. Global polynomial fitting will encounter model imbalance, accompanied by a sharp increase in fitting deviation. The local rise of fitting error stems from the mismatch between smooth approximation assumption and actual motion characteristics, instead of artificially varying noise variance.

Both angular velocity estimation accuracy and flight attitude solving accuracy are critical for subsequent analysis. In terms of angular velocity evaluation, root mean square error, mean absolute error, and peak error are employed to quantify fitting performance. In terms of attitude evaluation, quaternion deviation and cumulative attitude drift are selected to assess the accuracy of attitude calculation. The overall performance of each method is comprehensively evaluated with the aid of algorithm runtime.

Real-flight collision experiments under controlled conditions were carried out to investigate gyroscope observations under UAV impact disturbances. Figure 1 presents the measured Z-axis angular-velocity time series from repeated collision trials. The solid blue curve corresponds to the sample mean across multiple repeated experiments, and red dashed lines show the ± 1 standard-deviation envelope reflecting trial-to-trial random dispersion. Two collision events occur at instants indicated by arrows. It can be clearly observed that angular velocity itself remains continuous during each collision, while its time derivative (slope of the curve) exhibits sharp, abrupt jumps, which matches the physical assumption adopted in our numerical simulation. Similar characteristics can also be found on the X-and Y-axes. The expanding standard-deviation envelope after collision moments demonstrates that impact disturbances introduce prominent heteroscedastic measurement noise.

Note that the above real-flight dataset is used only to illustrate the physical characteristics of collision-induced gyroscope disturbances, rather than serving as quantitative algorithm-validation data within this manuscript. Quantitative performance evaluation of the proposed algorithm is still conducted via numerical simulation. Baseline tests without collision heteroscedastic noise are studied in the next subsection to isolate intrinsic fitting performance of polynomial schemes; subsequent test cases further introduce collision-like heteroscedastic disturbances to verify the complete integrated algorithm.

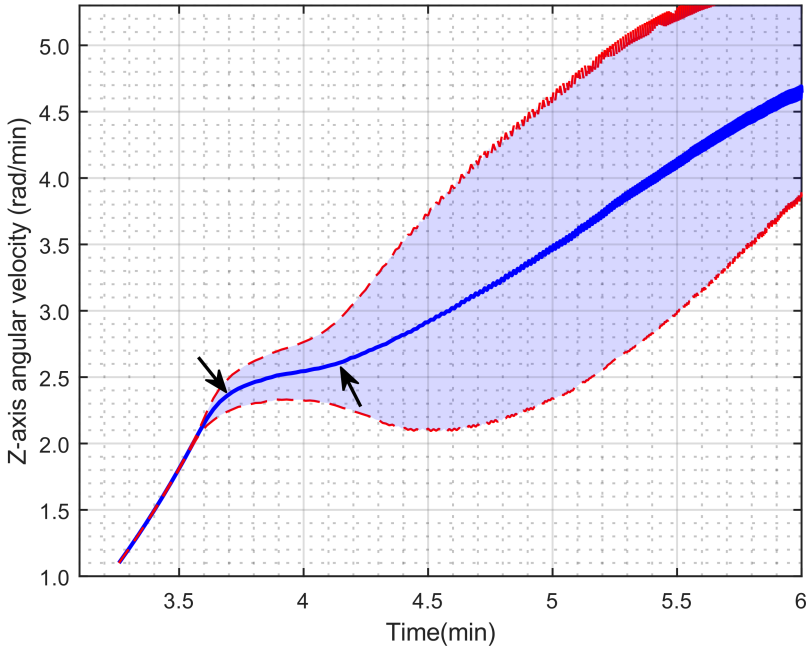


Figure 1. Measured Z-axis angular velocity under real-flight collision experiments.

4.2. Performance of Methods Under Collision-Free Steady Cruise

This subsection investigates the intrinsic fitting and noise-suppression performance of different polynomial schemes under steady cruise conditions. Note that no collision-induced heteroscedastic noise is introduced here. These experiments serve as baseline benchmarks to isolate the approximation capability of basis-function fitting. Heteroscedastic disturbance caused by collision impacts will be studied in subsequent change-point detection and integrated-algorithm tests.

Time series segments before the collision instant t_c are selected for experiments in this subsection. Within such segments, the angular velocity and its higher-order derivatives are continuous and smooth without motion jumps or statistical change points. The purpose of this comparative experiment is that independently quantifies the inherent approximation accuracy and native noise resistance of Legendre polynomial, Chebyshev polynomial, Hermite polynomial algorithms and the traditional four-sample equivalent rotation vector algorithm [11] under ideal steady flight conditions. Once the optimal fitting scheme for steady cruise segments is identified, it will provide a basis for the segmented fitting strategy proposed in the following subsection. Two progressive sub-working conditions are set for simulation tests: the first adopts noise-free pure simulation data where errors only originate from discrete sampling truncation, to separately compare the basic tracking capability of each algorithm for smooth coning motion; the second adds stationary normal white noise to simulate regular vibration interference of airborne sensors and verify the inherent noise suppression performance of each method.

4.2.1. Noise-Free Steady Cruise Condition

In this subsection, only the true angular velocity output by the coning motion model established in Section 4.1 is adopted, without any sensor measurement noise. Fitting errors are merely caused by discrete sampling truncation and polynomial approximation deviation. Noise interference is eliminated to separately compare the inherent fitting capability of the four algorithms for smooth oscillatory angular velocity.

To provide a practical and intuitive comparison of the algorithms' overall performance, we substitute the reconstructed angular velocity sequences from each algorithm into the quaternion update Equation (4) for numerical integration, thereby obtaining the corre-

sponding three-axis attitude angles at each time instant. Taking the theoretical attitude derived from the coning motion model established in Section 4.1 as the benchmark, the root-mean-square errors (RMSE) of the three-axis attitude over the entire simulation horizon are computed for each algorithm. This metric objectively reflects the comprehensive impact of angular velocity fitting deviations on final attitude accuracy after integral accumulation.

Table 1 presents the three-axis attitude RMSE comparison of all algorithms under noise-free steady cruise. It can be observed from the error results of X and Y oscillatory axes that the Hermite polynomial has the best tracking performance for periodic coning oscillation signals. Its attitude RMSE is only one-fifth of that of the Chebyshev scheme, and the attitude error is reduced by approximately four times compared with the traditional four-sample algorithm. All algorithms yield almost identical *Z*-axis errors. The *Z*-axis contains a constant rotation component, and its solving error is determined by the inherent geometric effect of non-commutative rotation in coning motion, which is independent of polynomial fitting forms.

Table 1. Three-axis attitude RMSE of all algorithms under noise-free steady cruise.

Method	*X*-Axis RMSE (Rad)	*Y*-Axis RMSE (Rad)	*Z*-Axis RMSE (Rad)
four-sample method	8.46×10^{-3}	8.73×10^{-3}	2.78×10^{-1}
Legendre polynomial method	5.21×10^{-3}	4.24×10^{-3}	2.73×10^{-1}
Chebyshev polynomial method	7.40×10^{-3}	7.70×10^{-3}	2.77×10^{-1}
Hermite polynomial method	1.35×10^{-3}	1.35×10^{-3}	2.78×10^{-1}

We test these methods with different ratio of coning frequency to sampling frequency to simulate different dynamic steady cruise. Under smooth time series with high relative frequency, the uniform interval approximation characteristic of Chebyshev polynomial becomes prominent, leading to a substantial improvement in accuracy. Nevertheless, the Hermite polynomial still maintains optimal fitting performance which may benefit from derivative interpolation constraints, while the traditional four-sample algorithm suffers severe dynamic errors, which may be due to its loss of high-order information fitting capacity.

4.2.2. Steady Cruise with Stationary Normal White Noise

In this subsection, zero-mean normal white noise with fixed variance is superimposed throughout the whole period span as (8). The noise intensity remains constant, simulating regular measurement errors induced by slight airframe vibration during long-duration cruise. The inherent noise suppression performance of the four algorithms is tested.

By comparing error variations between noise-free in Table 1 and noisy conditions in Table 2, we can conclude that Chebyshev, Legendre, and traditional four-sample algorithms all suffer obvious noise propagation effects. After noise superposition, sensor noise accumulates continuously in attitude solutions alongside angular increment iteration.

Table 2. Three-axis attitude RMSE of all algorithms under white-noise steady cruise.

Method	*X*-Axis RMSE (Rad)	*Y*-Axis RMSE (Rad)	*Z*-Axis RMSE (Rad)
four-sample method	$9.82 \times 10^{-3} \pm 1.3 \times 10^{-4}$	$1.01 \times 10^{-2} \pm 6.9 \times 10^{-4}$	$3.07 \times 10^{-1} \pm 6.1 \times 10^{-3}$
Legendre polynomial method	$6.05 \times 10^{-3} \pm 2.4 \times 10^{-4}$	$6.12 \times 10^{-3} \pm 2.7 \times 10^{-4}$	$2.94 \times 10^{-1} \pm 9.4 \times 10^{-3}$
Chebyshev polynomial method	$8.87 \times 10^{-3} \pm 1.7 \times 10^{-4}$	$9.26 \times 10^{-3} \pm 2.1 \times 10^{-4}$	$2.89 \times 10^{-1} \pm 5.7 \times 10^{-3}$
Hermite polynomial method	$1.52 \times 10^{-3} \pm 1.6 \times 10^{-4}$	$1.43 \times 10^{-3} \pm 2.0 \times 10^{-4}$	$2.81 \times 10^{-1} \pm 8.2 \times 10^{-3}$

Legendre, Chebyshev, and Hermite polynomials achieve theoretical orthogonality only on their respective standard domains. However, the angular rate fitting is implemented on segmented time windows with variable lengths, corresponding to truncated subintervals rather than the complete native orthogonal intervals. Consequently, the inherent orthogonality of these basis functions cannot be fully exploited, and cross-coupling between different-order basis terms emerges. For our truncated interval $[-1, \tau]$ where $-1 < \tau < 0.5$, the Hermite polynomial retains better orthogonality. Its weight function e^{-x^2} concentrates energy within the central region of the interval, so the cross-coupling induced by interval truncation is weaker, which effectively suppresses noise propagation. This characteristic contributes to explaining the observation that the Hermite-based scheme exhibits competitive fitting accuracy under moderate noise intensity, consistent with its performance in the noise-free case.

Comprehensively considering dynamic tracking accuracy under noise-free condition and stationary normal noise suppression capability, the performance ranking of the four algorithms is obtained as Hermite polynomial method, Legendre polynomial method, Chebyshev polynomial method, traditional four-sample method. Combining the two groups of simulation results in this subsection, a definite scheme selection conclusion can be reached that under long-duration steady cruise without collision, derivative jumps or variance change points, the Hermite polynomial possesses the optimal dynamic oscillation tracking accuracy and native anti-noise capacity simultaneously. Hence, the Hermite polynomial is uniformly adopted for angular velocity fitting and noise reduction in the following. For segments after collision with derivative jumps, the Hermite polynomial is combined with different change-point detection algorithms to realize segmented differentiated fitting.

4.3. Change-Point Detection Algorithm

Change-point detection is a classical problem in time-series analysis, the core of which is to identify the moment when the statistical properties undergo a structural change. For the piecewise continuous non-smooth polynomial fitting data involved in this study, no single algorithm can guarantee universality and reliability. To this end, this section selects five representative change-point detection algorithms as candidate schemes, briefly describes their basic principles and applicable characteristics, and on this basis designs a horizontal comparative validation strategy to select the detection method that best matches the actual data generation mechanism.

The first chosen algorithm is the standard CUSUM (SCUSUM) algorithm. Proposed by [10], the algorithm is mainly applied to detect mean shifts of time series under the assumption of a homoscedastic sequence. It can only identify abrupt changes in mean values and shows poor adaptability to scenes with variance variation, which easily causes false alarms and large detection delay in heteroscedastic disturbance environments.

The second is the CUSUM algorithm. One study [9] proposed the algorithm which extends the cumulative-sum idea to squared residuals, monitoring variance changes by accumulating the squared residuals after model fitting. The algorithm is directly designed for variance change-point detection and forms the basic framework for heteroscedastic change-point analysis. Nevertheless, it lacks a necessary correction mechanism in the steady state, causing the statistic to accumulate continuously and introducing inherent redundancy in judgment, which directly affects the real-time and accuracy of detection.

The third is the Bai–Perron multiple change-point (BPMC) algorithm [12], which searches for optimal segmentation points and can simultaneously identify multiple structural breaks in linear models, making it suitable for comprehensive offline analysis. Its advantages include strong model interpretability and flexibility in handling various regression specifications. However, its computational complexity is relatively high,

and it is essentially a post-hoc analysis method, thus failing to meet online real-time detection requirements.

The fourth is the generalized likelihood ratio (GLR) Algorithm [13], which compares the likelihood values under the "no change point" and "change point" hypotheses at each candidate time, using the maximum likelihood ratio as the criterion to locate the change. The method is theoretically optimal, but its performance heavily relies on a pre-specified noise distribution (which is supposed to be a normal distribution here) In actual complex disturbance environments, distributional mismatches can severely degrade detection capability, leading to insufficient robustness.

The fifth and last one is our proposed algorithm (16,17), an extension of the variance change-point detection system, which may be suitable for scenarios where abrupt changes in angular-velocity derivatives cause imbalance in the fitting model and local surges in residuals.

We adopt the angular velocity time series with collision-induced derivative jumps constructed in Section 4.1 as the unified test dataset. Four evaluation indicators, including change-point detection delay, number of false alarms during steady cruise, missing detection rate of abrupt collision events, and single-step computation time, are utilized to quantitatively compare the above five algorithms: SCUSUM, CUSUM, BPMC, GLR, and our proposed (OUR) algorithm. Table 3 presents the mean performance of those algorithms with 10,000 random tests.

Table 3. Mean performance of five algorithms with 10,000 random tests.

Algorithm	Change-Point Detection Delay (msec)	False Alarms	Missing Detection of Collision	Single-Step Computation Time (msec)
SCUSUM	$13.67 \pm 5.2 \times 10^{-1}$	4384	154	956.7 ± 64.2
CUSUM	$6.72 \pm 1.9 \times 10^{-1}$	639	287	938.2 ± 127.4
BPMC	$4.38 \pm 8.5 \times 10^{-1}$	172	92	2827 ± 152
GLR	5.69 ± 2.7	88	137	4375 ± 373
OUR	$3.16 \pm 4.4 \times 10^{-1}$	67	83	662.4 ± 71.1

The computation rules of each metric are clarified as follows.

1. Change-point detection delay. Let t_{true} denote the first sampling instant after the collision-induced jump arises, and t_{alert} the sampling instant when the algorithm triggers an alarm. The delay for each successfully detected trial is defined as $t_{\text{alert}} - t_{\text{true}}$, and the mean delay is calculated over all successfully detected cases.
2. Number of false alarms. This metric counts the total false alarms triggered within collision-free stationary segments, and the total value across all trials is recorded.
3. Missing detection of abrupt collision events. It refers to the number of test trials containing genuine collision jumps while no valid alarm is generated throughout the whole time series and all trails.
4. Single-step computation time. It represents the runtime consumed by one complete sampling iteration that involves attitude computation. It should be noted that collision judgment can be implemented merely using angular rate measurements with low computational overhead. The difference in single-step runtime among algorithms mainly originates from whether attitude updating is embedded in the iterative pipeline.

From the quantitative data listed in Table 3, the performance of each algorithm and corresponding root causes are analyzed as follows. SCUSUM algorithm presents many false alarms and obvious detection delay in the experiment. We attribute this phenomenon to its inherent design limitation that the algorithm is only sensitive to mean shifts and cannot

effectively capture the residual triggered by angular velocity derivative jumps, which leads to poor performance in our collision scenario [14,15]. The CUSUM algorithm generates frequent false judgments during steady cruise. The core reason is that its detection statistic is built upon squared residuals without a steady-state compensation term. As a result, the statistic accumulates meaningless values continuously under smooth flight conditions and easily triggers misjudgments.

The Bai–Perron multiple change point (BPMC) algorithm achieves relatively ideal detection accuracy from the table data, yet it suffers extremely high computational complexity. Its offline batch-solving framework cannot satisfy the real-time operation requirement of UAV airborne equipment [16,17]. The generalized likelihood ratio (GLR) algorithm also obtains acceptable detection results, but its computational overhead remains heavy [18,19]. Different from the above four methods, our proposed modified CUSUM algorithm embeds a steady-state residual offset compensation term into the recursive statistic. This structural improvement greatly cuts down false alarms while maintaining the shortest detection delay and low single-step computation cost, which makes it the most suitable detector for the UAV collision time series studied in this paper.

4.4. Parameter Sensitivity Analysis

The performance of the proposed heteroscedastic decoupling algorithm depends on key practical parameters, including the detection threshold coefficient k in Equation (11) and the gyroscope sampling rate. In this section, Monte-Carlo numerical experiments are carried out exclusively for our proposed algorithm. All other simulation configurations remain consistent with the test conditions in Section 4.3, unless explicitly stated. The objective is to investigate parameter influences and provide guidance for practical onboard parameter tuning.

The threshold coefficient k determines the alarm threshold $H = k\widehat{\sigma}_0^2$ of the improved single-pass CUSUM detector, and it directly balances detection latency, false-alarm quantity, and missing-detection probability. We sweep k over the range $\{3.0, 3.5, 4.0, 4.5, 5.0\}$ under 10,000 independent Monte-Carlo runs. The quantitative detection metrics are listed in Table 4.

Table 4. Detection performance under different threshold coefficient k.

k	Detection Delay (ms)	False Alarms	Missing Detection of Collision	Single-Step Computation Time (msec)
3.0	$2.19 \pm 5.6 \times 10^{-1}$	107	73	635.2 ± 59.4
3.5	$2.73 \pm 4.7 \times 10^{-1}$	83	76	646.3 ± 62.7
4.0	$3.08 \pm 5.2 \times 10^{-1}$	69	79	659.8 ± 67.0
4.5	$3.34 \pm 4.6 \times 10^{-1}$	67	82	661.5 ± 68.2
5.0	$3.69 \pm 5.9 \times 10^{-1}$	64	97	669.4 ± 67.4

As shown in Table 4, reducing k lowers the detection delay, yet more false alarms are triggered during steady-flight segments without real collisions. As k increases, the false-alarm number decreases significantly; however, a larger threshold brings longer detection delay and higher risk of missing real collision events. Each Monte-Carlo run contains exactly one genuine collision event, so the total number of events to be detected across $10,000$ runs is $10,000$. The false-alarm count is accumulated over the collision-free steady-flight segments before the impact, whereas the missing-detection count refers to the number of runs in which the genuine collision jump is not alarmed. The missing-detection values of 73–97 across the tested k range correspond to missing-detection rates of 0.73–0.97%, all below 1%, which is considered acceptable for the target collision-avoidance application.

Regarding the choice of $k = 4.0$: at $k = 4.5$ the false-alarm count is slightly lower, but the missing-detection count is comparable, and the detection latency increases; at $k = 5.0$ the missing-detection count rises noticeably. We select $k = 4.0$ as a conservative engineering choice that keeps the missing-detection rate below 1% while maintaining a faster detection response, which is preferable for real-time collision recovery on board small UAVs. Therefore, $k = 4.0$ is adopted as the default threshold coefficient in this paper.

On the other side, gyroscope sampling rate is a critical engineering parameter for UAV-embedded navigation systems. Higher sampling rates generally provide richer motion information but increase data throughput and computational burden. To investigate the algorithm adaptability for real-world hardware, three typical practical sampling rates, i.e., 100 Hz, 200 Hz, and 400 Hz, are tested. The threshold coefficient is fixed as $k = 4.0$, and all other kinematic and noise parameters are unchanged. The attitude accuracy and detection-delay results are summarized in Table 5.

Table 5. Algorithm performance under different gyroscope sampling rates.

Sampling Rate	*X*-Axis RMSE (Rad)	*Y*-Axis RMSE (Rad)	*Z*-Axis RMSE (Rad)	Detection Delay (msec)
100 Hz	$(12.3 \pm 1.9) \times 10^{-3}$	$(9.2 \pm 1.2) \times 10^{-3}$	$(3.76 \pm 0.1) \times 10^{-1}$	$3.5 \pm 7 \times 10^{-1}$
200 Hz	$(8.5 \pm 1.1) \times 10^{-3}$	$(7.1 \pm 0.9) \times 10^{-3}$	$(3.47 \pm 0.1) \times 10^{-1}$	$3.6 \pm 7 \times 10^{-1}$
400 Hz	$(6.7 \pm 0.6) \times 10^{-3}$	$(5.8 \pm 0.7) \times 10^{-3}$	$(3.17 \pm 0.07) \times 10^{-1}$	$3.1 \pm 5 \times 10^{-1}$

From Table 5, higher sampling rates yield better attitude-solving accuracy and a relatively smaller detection delay. When the sampling rate drops from 400 Hz to 100 Hz, attitude RMSE increases moderately. Even at the lowest tested 100 Hz sampling rate, the proposed integrated algorithm still maintains acceptable navigation performance. Lower sampling rates can reduce data volume and relieve the computational pressure of lightweight airborne embedded hardware. Nevertheless, the sampling rate cannot be reduced arbitrarily; excessively low sampling frequency will degrade both polynomial-fitting precision and change-point-detection reliability. In practical engineering deployment, the gyroscope sampling rate should be selected according to the trade-off between navigation accuracy requirement and onboard computing resources.

4.5. Comparison with EKF-Based Attitude Propagation Under Gyro-Only Input

For a quantitative comparison with modern filtering approaches, we supplement this section with an error-state extended Kalman filter (EKF) as a representative back-end benchmark [19]. Two clarifications on the comparison scope are necessary. First, conventional adaptive EKF/UKF, factor-graph optimization, and neural-network-based IMU compensation rely on accelerometer (and often magnetometer) measurements for measurement update, which makes them multi-sensor back-end estimators intrinsically different from the proposed gyro-only front-end preprocessor. A direct head-to-head comparison would be informationally unfair because those back-ends have access to additional sensor observations that the proposed method does not use. Second, during instantaneous collision impacts, the accelerometer typically enters saturation (see Section 5.2), so its measurement update becomes unreliable precisely when it is needed most.

For a fair comparison, we place both methods under identical gyro-only input conditions. The error-state EKF propagates the nominal quaternion via the same kinematic Equation (4) using the raw measured angular velocity; without external measurements, its measurement-update step is inactive, and the filter reduces to a quaternion integrator with covariance propagation. The proposed method first preprocesses the raw gyroscope measurements through the integrated pipeline (Hermite fitting, improved CUSUM detection, axis-differentiated weighting, and segmented refitting) and then feeds the corrected

angular velocity into the same quaternion propagation. Both pipelines share the identical collision-induced heteroscedastic simulation scenario described in Section 4.1.

Figure 2 presents the four quaternion components and the attitude error angle for both pipelines. The EKF output (red dashed) exhibits visibly faster error growth after the collision instant at $t = 4$ s, whereas the proposed method (black solid) maintains a tighter track. Quantitatively, the final attitude error is 0.23 deg for the proposed method versus 0.81 deg for the EKF, and the peak attitude error is 0.33 deg versus 0.91 deg, corresponding to reductions of approximately 71% and 64%, respectively.

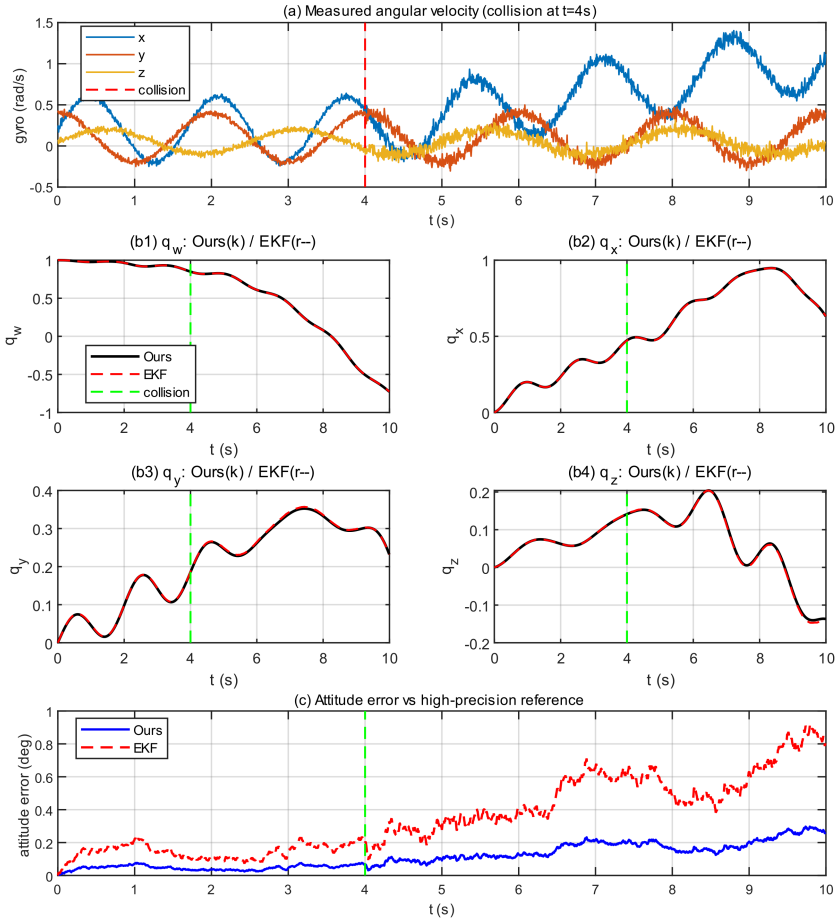


Figure 2. Comparison with EKF attitude propagation under gyroscope-only input.

The underlying reason is that, under gyro-only input, the EKF has no mechanism to correct accumulated integration drift once the angular-velocity derivative jumps and the noise variance increases at collision. The proposed front-end, by contrast, detects the variance change point via the improved CUSUM detector and reinitializes polynomial fitting locally, blocking error propagation across the collision boundary. This experiment confirms that the proposed pipeline provides higher-quality angular-velocity input for downstream attitude propagation. Full multi-sensor back-end integration (e.g., adaptive EKF/UKF with accelerometer observations) is left for future work once the accelerometer saturation problem is addressed.

5. Discussion

5.1. Analysis of Results

The comparative simulation results under collision-free cruising clearly demonstrate the hierarchical performance of orthogonal polynomial fitting schemes against traditional

multi-sample equivalent rotation vector algorithms, which is consistent with the core theoretical hypothesis proposed in this paper: the ability to reconstruct high-order derivatives of angular velocity determines the dynamic error suppression capacity of strapdown inertial navigation attitude solving algorithms.

Under noise-free coning motion, the Hermite polynomial method exhibits superior tracking performance for oscillatory X/Y-axis motion. This outcome can be explained by the intrinsic interpolation property of Hermite orthogonal polynomials: unlike Legendre and Chebyshev polynomials that only fit discrete angular velocity values, Hermite polynomials simultaneously constrain the function value and its first-order derivative at sampling nodes, naturally matching the continuously changing oscillatory angular velocity of the X/Y axes during UAV steady flight. Legendre polynomials deliver suboptimal intermediate performance because their minimum mean-square error optimization target lacks uniform error distribution over the entire fitting interval, weakening tracking accuracy for high-frequency periodic signals. Chebyshev polynomials provide a uniform fitting error distribution, but they do not incorporate derivative constraints, leading to obvious tracking lag for fast-varying angular velocity, consistent with findings in [4] that orthogonal polynomials without derivative matching suffer amplified dynamic coning errors under high-frequency swing motion. The traditional four-sample equivalent rotation vector algorithm assumes linear variation of angular velocity within the update interval, completely ignoring high-order derivative terms, which fundamentally restricts its error suppression potential and yields the worst attitude accuracy in all steady-state test cases, matching the theoretical error deduction of classic attitude updating frameworks in [20].

When stationary normal white noise is introduced to simulate mild airframe vibration, the performance ranking of each algorithm remains unchanged. The unique zero-weight sum characteristic of Hermite polynomial fitting coefficients enables inherent offsetting of zero-mean sensor noise, creating a built-in noise suppression effect unavailable in other polynomial schemes. Existing heteroscedastic noise processing research [7,8] mostly relies on post-filtering correction after attitude calculation, while this work realizes noise attenuation at the raw angular velocity fitting stage, which fundamentally avoids cumulative error propagation during quaternion integration and provides a new lightweight preprocessing idea for inertial measurement noise.

The comparative test of five change-point detection algorithms verifies the effectiveness of the steady-state offset compensation term embedded in the proposed single-pass CUSUM statistic, which addresses two prominent defects of the original squared-residual CUSUM algorithm in UAV collision scenarios as pointed out in [9]: continuous meaningless accumulation of detection statistics under smooth flight and frequent false alarms induced by periodic oscillation of X/Y axes.

Standard SCUSUM only detects mean shifts of time series and fails to capture residual surges caused by collision-triggered angular velocity derivative jumps, resulting in severe detection delay and false alarm rates, which makes it entirely unsuitable for the heteroscedastic disturbance studied in this paper. Offline segmentation algorithms including BPMC and GLR exhibit satisfactory detection accuracy and low missing detection rates, yet their batch matrix operation and likelihood function iteration bring considerable computational overhead far exceeding the real-time computing limit of UAV airborne embedded chips. Benefiting from embedded steady-state residual compensation, our modified CUSUM achieves a good trade-off among detection precision, false-alarm suppression, latency, and hardware computational burden.

From the perspective of a practical engineering mechanism, collision disturbances only generate a single variance mutation point in long-term UAV cruising, rather than multiple continuous structural breaks. The single-pass recursive logic of the improved

CUSUM avoids redundant multi-point traversal search adopted by offline algorithms, which aligns perfectly with the typical working mode of cargo UAVs. Combined with the segmented refitting mechanism reinitializing polynomial coefficients at detected change points, the model mismatch problem caused by abrupt angular velocity derivatives after collision is completely isolated, filling the research gap that existing polynomial fitting methods adopt global unified modeling without segmented adaptive adjustment mentioned in Introduction.

It is worth noting that rigorous analytical mathematical proof of superiority for this integrated nonlinear time-series workflow is non-trivial. Qualitatively, traditional cascaded schemes execute fitting, change-point detection, and weighting correction independently. Estimation errors from upstream modules will propagate and accumulate to downstream processing. In our integrated pipeline, once a variance change point is detected, we immediately re-initialize all polynomial coefficients and cumulative matrices at the collision moment. This mechanism isolates error propagation before and after collision. Quantitative comparative simulation between integrated and cascaded implementations requires substantial development workload, and we leave this for future work.

Current high-precision SINS research splits into two independent branches: angular velocity basis function fitting and inertial noise suppression, with few integrated frameworks combining the two links. Polynomial fitting literature [2–4] mainly focuses on ideal noise-free coning motion and does not consider time-varying heteroscedastic noise triggered by external physical impacts; adaptive filtering-based noise suppression methods [6] only conduct post-attitude error compensation and fail to calibrate distorted raw gyroscope angular velocity at the input source, leading to an inherent upper limit on error elimination performance. Representative modern back-end navigation frameworks include robust Kalman-filter variants [5,21], factor-graph optimization [22], and neural-network-driven IMU compensation [23]. These approaches operate at the state-estimation level rather than preprocessing raw gyroscope measurements. Fair quantitative comparison requires embedding our preprocessing pipeline into those back-end architectures, which is deferred to future work under the tight revision schedule.

This work innovatively integrates orthogonal polynomial angular velocity fitting, optimized CUSUM variance change-point detection and three-axis differentiated weighted correction into a unified decoupling algorithm, forming a complete front-end preprocessing workflow for gyroscope observation signals under collision disturbance. Compared with existing separated processing schemes, the proposed method suppresses dynamic errors and heteroscedastic impact noise simultaneously without substantially increasing computational complexity, which is highly applicable to lightweight, low-power embedded navigation hardware of small cargo UAVs. The quantitative comparison of Chebyshev and Hermite polynomials under impact noise conditions also supplements the blank of performance verification for orthogonal polynomials in non-smooth, time-varying noise scenarios in existing literature.

Further parameter-sensitivity tests in Section 4.4 indicate that the threshold coefficient k and gyroscope sampling rate have non-negligible influences on the overall performance, and careful parameter tuning is necessary for practical onboard applications.

A notable by-product of the proposed pipeline is the accurate estimation of the collision instant via the improved CUSUM detector. In a complete navigation system, this time estimate can trigger a controlled reinitialization of downstream filters (e.g., EKF/UKF) immediately after impact, rather than allowing them to re-converge gradually from contaminated states. Combined with the corrected angular-velocity output and the accelerometer impulse-compensation framework described in Section 5.2, this creates a pathway toward collision-tolerant multi-sensor navigation, which we will pursue in future work.

5.2. Limitations of the Proposed Algorithm

This study has several important limitations that should be noted.

First, this work focuses on gyroscope raw-measurement preprocessing and the attitude-update workflow; accelerometer error propagation and full SINS velocity–position estimation is outside the scope of the present manuscript. In real-world collision scenarios, severe impact loads may cause accelerometer saturation, where the magnitude of the collision-induced specific force cannot be recovered directly from the saturated measurements. Nevertheless, we observe that the direction of the impact acceleration remains available and approximately constant over the short impact interval. Under this observation, the collision can be modeled as an equivalent instantaneous impulse applied at a single contact point. The unknown magnitude of the collision-induced velocity increment could then be estimated by combining rigid-body kinematics with the linear and angular impulse–momentum relations, provided that (i) the contact location is known or identifiable, (ii) the rigid-body geometry and mass properties are known, (iii) the pre- and post-impact angular velocities are available, and (iv) the corresponding identification equation is nondegenerate. We leave the detailed derivation and experimental validation of this accelerometer-compensation framework to future work. Notably, the proposed front-end's accurate collision-time estimation (Section 4.3) provides a natural trigger for such post-impact reconstruction, since the impulse-momentum identification requires precisely segmented pre- and post-impact angular-velocity data.

Second, our improved single-pass CUSUM detector is designed for a single abrupt variance jump triggered by instantaneous collision. Under multiple successive impacts, the detector resets its cumulative statistic after the first alarm; if a second impact occurs within the reset window, it may be missed or detected with increased latency. Under prolonged high-frequency structural vibration, the fitting residuals remain persistently elevated, which can overwhelm the steady-state offset compensation term and raise the false-alarm rate. Under gradual sensor degradation, the variance change is slow rather than abrupt, and the CUSUM statistic accumulates only weakly, leading to substantially longer detection delay or eventual missed detection. Future studies will extend the detector to multi-change-point and gradual-drift scenarios.

Third, the collision simulation model adopts idealized assumptions. Real UAV collisions involve complicated nonlinear airframe dynamics, structural resonance, and possible IMU sensor saturation, which are not fully reproduced in our numerical simulation.

Fourth, although computational complexity analysis demonstrates its potential feasibility for lightweight embedded hardware, complete embedded flight-test datasets are not presented in this manuscript. Embedded experimental validation will be reported in our follow-up dedicated publication.

Future research directions derived from the above limitations include: (1) comparative study embedding our preprocessing pipeline within modern navigation back-ends such as EKF/UKF, factor-graph and neural-network-driven error compensation; (2) extension of the detection module for multiple successive collision events; (3) multi-sensor fusion under accelerometer saturation caused by heavy impact; (4) comprehensive embedded flight-test validation.

6. Conclusions

This paper addresses the attitude-computation accuracy degradation of strapdown inertial navigation systems for small cargo UAVs subjected to instantaneous obstacle collisions. Such degradation originates from abrupt derivatives of angular velocity and heteroscedastic gyroscope measurement noise induced by impact disturbances. We develop an integrated heteroscedastic decoupling algorithm, which combines optimized single-pass

CUSUM change-point detection, axis-differentiated weighted correction, and Hermite-orthogonal-polynomial-based angular-velocity fitting. Complete derivation and multi-scenario numerical simulation validations are presented in this work.

Building on existing fundamental tools, including Hermite polynomials and recursive least-squares, the main technical improvements of this paper are the steady-state-compensated CUSUM detector, the axis-differentiated segmented weighting strategy, and the tightly coupled end-to-end preprocessing pipeline for raw gyroscope measurements. Simulation results demonstrate that under the tested single-collision scenarios, the proposed approach reduces angular-velocity fitting error and cumulative attitude drift relative to conventional polynomial-fitting schemes and classic equivalent-rotation-vector attitude-update methods.

The proposed algorithm shows promising lightweight computational complexity suitable for potential deployment on UAV-embedded navigation hardware. It provides an alternative preprocessing-oriented solution for small cargo UAVs operating in narrow obstacle-rich environments.

Nevertheless, the presented method is mainly valid for scenarios with single instantaneous collision-induced variance jumps. Its performance may degrade under successive multi-impacts, long-time high-frequency vibration, and gradual sensor drift. Further extensions and real-world embedded experimental verifications are required before practical field application.

In addition, the improved CUSUM detector provides an accurate estimate of the collision instant, which can serve as a natural trigger for post-impact back-end reinitialization. Future work will embed this front end within adaptive EKF/UKF or factor-graph back ends, combined with the accelerometer impulse-compensation framework outlined in Section 5.2, to form a complete collision-tolerant navigation system.

Author Contributions: Conceptualization, Y.W. and Q.D.; methodology, Y.W., R.D. and Q.D.; software, R.D. and B.W.; validation, Y.W. and Q.D.; formal analysis, Q.D. and B.W.; investigation, Y.W.; resources, Y.W.; data curation, Y.W.; writing—original draft preparation, R.D.; writing—review and editing, Y.W. and Q.D.; visualization, B.W.; supervision, Q.D.; project administration, Y.W.; funding acquisition, Y.W. and Q.D. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding authors.

Acknowledgments: The authors gratefully acknowledge technical support from the Academy of Aerospace Solid Propulsion Technology (AASPT), Xi'an, China.

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

References

1. Butt, M.Z.; Nasir, N.; Rashid, R.A.; Tang, C.H.H. Open-Gap Algorithm for Fast Reactive Collision Avoidance by Autonomous UAV using Sensor Array. *J. Aeronaut. Astronaut. Aviat.* **2025**, *57*, 505–511.
2. Paolo, D.P.; Carlson, S.J.; Papachristos, C.; Alexis, K. Collision-tolerant Aerial Robots: A Survey. *arXiv* **2026**, arXiv:2212.03196v1.
3. Paolo, D.P.; Dharmadhikari, M.; Alexis, K. Risk-aware Motion Planning for Collision-tolerant Aerial Robots subject to Localization Uncertainty. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Kyoto, Japan, 23–27 October 2022.
4. Huang, X.; Liu, G.; Liu, Y. Collision Detection and Recovery Control of Drones Using Onboard Inertial Measurement Unit. *Drones* **2025**, *9*, 380. [CrossRef]
5. Wei, X.; Fan, S.; Zhang, Y.; Gao, W.; Shen, F.; Ming, X.; Yang, J. A robust adaptive error state Kalman filter for MEMS IMU attitude estimation under dynamic acceleration. *Measurement* **2025**, *242*, 116097. [CrossRef]

6. Savage, P.G. Strapdown inertial navigation integration algorithm design part 1: Attitude algorithms. *J. Guid. Control Dyn.* **1998**, *21*, 19–28. [CrossRef] [PubMed]
7. Lai, J.; Lv, P.; Jiang, B. Noncommutativity Error Analysis of Strapdown Inertial Navigation System under the Vibration in UAVs. *Int. J. Adv. Robot. Syst.* **2012**, *9*, 136. [CrossRef] [PubMed]
8. Battiston, A.; Sharf, I.; Nahon, M. Attitude estimation for collision recovery of a quadcopter unmanned aerial vehicle. *Int. J. Robot. Res.* **2019**, *38*, 1286–1306. [CrossRef]
9. Amiri, F.; Khorashadizadeh, S. Adaptive control of a class of uncertain nonlinear systems using brain emotional learning and Legendre polynomials. *Trans. Inst. Meas. Control* **2024**, *46*, 1667–1679. [CrossRef]
10. Santra, S.; Behera, R. Physics-informed fractional machine intelligence and space-time wavelet frameworks for non-local integro-partial differential equations involving weak singularities. *Commun. Nonlinear Sci. Numer. Simul.* **2025**, *152*, 109184. [CrossRef]
11. Arai, Y.; Otsu, T.; Xu, M. GLS under monotone heteroskedasticity. *J. Econom.* **2024**, *246*, 105899. [CrossRef]
12. Chu, Y.; Mu, R.; Li, S.; Cui, N. An improved adaptive robust information filter for spacecraft relative navigation. *Aerosp. Sci. Technol.* **2022**, *130*, 107873. [CrossRef]
13. Lu, W.; Jia, X.; Teng, Y.; Du, Y.; Zhang, Z. Backward smoothing adaptive SVDCKF integrated navigation algorithm. *Aerosp. Sci. Technol.* **2024**, *74*, 2801–2813. [CrossRef]
14. Inclán, C.; Tiao, G. Use of Cumulative Sums of Squares for Retrospective Detection of Changes of Variance. *J. Am. Stat. Assoc.* **1994**, *89*, 913–923. [CrossRef]
15. Shuster, M.D. A Survey of Attitude Representations. *J. Astronaut. Sci.* **1993**, *41*, 439–517.
16. Alaimo, A.; Artale, V.; Milazzo, C.; Ricciardello, A. Comparison between Euler and Quaternion Parametrization in UAV Dynamics. In Proceedings of the 11th International Conference of Numerical Analysis and Applied Mathematics, Rhodes, Greece, 21–27 September 2013.
17. Page, E.S. Continuous inspection schemes. *Biometrika* **1954**, *41*, 100–115. [CrossRef]
18. Bai, J.S.; Perron, P. Estimating and Testing Linear Models with Multiple Structural Changes. *Econometrica* **1998**, *66*, 47–78. [CrossRef]
19. Dranitsyna, E.V.; Sokolov, A.I. Strapdown Inertial Navigation System Accuracy Improvement Methods Based on Inertial Measuring Unit Rotation: Analytical Review. *Gyroscopy Navig.* **2023**, *14*, 290–304. [CrossRef]
20. González-Coya, E.; Perron, P. Estimation in the Presence of Heteroskedasticity of Unknown Form: A Lasso-based Approach. *J. Econom. Methods* **2023**, *13*, 29–48. [CrossRef]
21. Chen, X.; Li, Z.; Tao, Z. A robust adaptive Student's t-Kalman filter for GNSS/INS integrated navigation. *Meas. Sci. Technol.* **2025**, *36*, 116309. [CrossRef]
22. He, H.; Xiong, Z.; Kang, J.; Zhang, X. Multi-source Fusion Navigation Algorithm for Aerospace Vehicles Based on Adaptive Factor Graph. *Adv. Guid. Navig. Control* **2025**, *1348*, 128–137. [CrossRef]
23. Zhang, Z.; Zhu, G.; Wu, J. Neural-Network-Based Temperature Drift Compensation for MEMS-IMU Systems. *IEEE Sens. J.* **2025**, *25*, 26353–26363. [CrossRef]

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.