

Article

An Evidence-Weighted Multi-Rate Inertial Fusion Method for Industrial Load Displacement Reconstruction

Peiyi Zhou ¹, Weige Liang ^{1,*}, Qizheng Zhou ¹, Chi Li ¹, Shiyan Sun ¹ and Jixin Song ²¹ Naval University of Engineering, Wuhan 430033, China² No. 208 Research Institute of China Ordnance Industries, Beijing 102202, China

* Correspondence: 1312021010@nue.edu.cn

Abstract

An evidence-weighted multi-rate inertial fusion method is proposed for load displacement reconstruction in high-vibration industrial transportation scenarios, where inertial displacement reconstruction is readily affected by acceleration bias, attitude error, vibration disturbance, and false zero-velocity decisions. Multi-rate inertial data and high-bandwidth vibration/impact observations are used. The main inertial channel is used for motion-dynamics modeling, low-frequency inertial and attitude information is used to constrain the trajectory trend, and high-bandwidth vibration and impact responses are converted into motion-veto evidence so that unreliable stationary decisions under high-vibration conditions can be weakened. A Dempster–Shafer-style evidence fusion model is adopted to estimate motion, stationary, and unknown confidence, and the stationary confidence is propagated to ZUPT gating, covariance weighting, multi-branch displacement fusion, and keyframe factor-graph optimization. The experimental results show that stable load displacement reconstruction can be achieved by the proposed method, while interpretable diagnostic information, including displacement, velocity, stationary probability, ZUPT triggering, residuals, and drift risk, is also output. A diagnosable and auditable framework for inertial displacement reconstruction is therefore provided for high-vibration industrial scenarios where external references are difficult to deploy.

Keywords: displacement reconstruction; inertial sensor; multi-rate fusion; zero-velocity update; Dempster–Shafer evidence; factor graph

MSC: 94A12



Academic Editor: Ping Lang

Received: 15 June 2026

Revised: 27 July 2026

Accepted: 27 July 2026

Published: 4 August 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 \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Smart manufacturing, robotic load handling, and condition-aware maintenance driven by digital twins are driving the transformation of industrial systems from offline monitoring to real-time or quasi-real-time status updates [1–5]. In such scenarios, the displacement trajectories of loads, fixtures, or moving parts not only describe geometric motion, but also serve for path verification, transport repeatability assessment, anomaly detection, and maintenance decisions. Existing displacement measurements often rely on external reference tools, such as optical motion capture, laser displacement meters, encoders, vision systems, or vision-acceleration fusion. The occlusion, limited space, lighting changes, metal structure interference, and strong mechanical vibration in industrial sites make it difficult to deploy these external benchmarks stably in the long term. Inertial sensors are self-contained, miniaturized, and easy to install, and they are suitable for installation on

moving loads; however, the displacement needs to be obtained by double integration of acceleration, and small offsets, attitude errors, gravity leakage, synchronization errors, and vibration interference will be rapidly amplified [6–10]. The issue that this study focuses on is thus formed: under industrial transportation conditions with multiple sampling rates, strong vibrations, and limited independent external measurements, how to convert observations in different frequency bands into three types of credible evidence, namely motion, stationary, and unknown, and transfer this uncertainty to the ZUPT, covariance weighting, and factor-graph displacement reconstruction processes.

Inertial displacement reconstruction is mainly limited by two types of factors. First, the inertial sensor directly measures acceleration and angular velocity. Displacement is obtained by double integration of acceleration; therefore, accelerometer bias, attitude errors, and gravity leakage accumulate rapidly in the reconstructed trajectory. Second, the effectiveness of ZUPT depends on whether the stationary detection is reliable. During industrial load transportation, low-speed rigid body motion, mechanical shaking, continuous vibration, and local impact may occur alternately. It is difficult for a single acceleration threshold to stably distinguish motion from rest. If the moving segment is misjudged as stationary, the incorrect zero-velocity constraint will suppress the actual velocity; if the stationary segment is missed, double-integral drift will continue to accumulate. A more appropriate treatment is to describe stationary states as quantities of evidence with uncertainty. For industrial loads, this uncertainty comes both from low-frequency inertial changes and from the violation of the zero-velocity assumption by vibration and shock [11–15].

Multi-rate data in industrial load displacement reconstruction have different information roles. The main inertial channel provides motion-dynamic information with higher time resolution, but it is susceptible to high-frequency noise and integral drift in strong-vibration environments; the low-frequency inertial and attitude channels help provide trajectory trends, attitude assistance, and stationary segment stability; the high-bandwidth vibration channels mainly represent continuous mechanical disturbances and local structural vibrations; and the high-bandwidth impact channels are more sensitive to short-term impacts, collisions, or sudden disturbances. High-bandwidth vibration and shock observations mainly describe local dynamic responses and cannot directly represent rigid body translational displacements. If double-integrated directly, local vibration and impact can easily be misinterpreted as low-frequency displacement drift. This study therefore converts vibration and impact information into motion-veto evidence: when high-frequency vibration or impact activity is significant, the stationary confidence at the current time is reduced; when high-frequency activity is weak, it only indicates the lack of obvious motion-veto evidence, and the unknown confidence level is still retained. In this way, multi-rate data are transformed from simple signal superposition to evidence collaboration oriented to uncertainty management. Dempster–Shafer evidence theory can describe uncertainty, conflicting evidence, and epistemic uncertainty in multi-source information, and it has been used in multi-source information fusion, fault diagnosis, and risk assessment [16–21]. Manifold pre-integration provides a strict model for inertial keyframe propagation [22]; Meanwhile, iSAM2/Bayes-tree methods and factor-graph optimization frameworks have become relatively mature back-end optimization tools in robot perception, visual-inertial odometry, and multi-sensor navigation [23–25]. These methods provide a methodological basis for integrating low-frequency trajectory priors, inertial dynamic constraints, and zero-velocity constraints into the keyframe smoothing model used in this study. Factor-graph-based fusion has subsequently been extended to GPS/VIO integration, UWB/IMU cooperative localization, and robust radar-inertial odometry [26–28], while adaptive multi-sensor and multi-inertial fusion strategies have also been investigated for heterogeneous and dynamically varying environments [29,30]. Against this background, the present study focuses

specifically on the unified propagation of vibration-informed stationary-state uncertainty throughout industrial load displacement reconstruction.

To address the uncertainty in stationary-state identification, the heterogeneous information roles of multi-rate observations, and the difficulty of deploying external reference systems during high-vibration industrial load transportation, an evidence-weighted multi-rate inertial displacement reconstruction framework is developed. The methodological contribution lies in the scenario-oriented integration of heterogeneous multi-rate observations and the consistent propagation of stationary-state uncertainty throughout the reconstruction process. Unlike conventional pipelines that treat stationary detection as an independent binary preprocessing step, the proposed framework represents motion, stationary, and unknown states as continuous evidence masses and propagates the resulting confidence through ZUPT activation, covariance adaptation, multi-branch displacement fusion, and keyframe factor-graph optimization. The main contributions are as follows: (1) a role-aware evidence model is constructed for multi-rate inertial, vibration, and impact observations, in which high-frequency vibration and impact information serves as motion-veto evidence rather than being directly integrated for displacement, thereby reducing the risk of false zero-velocity constraints under high-frequency disturbances; (2) a unified confidence-propagation mechanism is developed, allowing stationary confidence to continuously regulate ZUPT activation and covariance and to be further incorporated into displacement fusion and factor-graph optimization; (3) an adaptive displacement fusion strategy based on branch consistency, residual roughness, and stationary-interval drift is developed to balance high-dynamic responsiveness and low-frequency stability; and (4) a diagnosable evaluation framework is established for industrial environments with limited external references, providing displacement, velocity, ZUPT activation, residual, and reliability outputs for result interpretation and subsequent digital-twin applications.

2. Methods

2.1. Multi-Rate Observation Framework and Signal Preprocessing

A composite sensor package was installed on the industrial load to acquire multi-rate inertial measurements and high-bandwidth vibration and impact signals. The sensor package contains main inertial acceleration channel, low-frequency inertial/attitude channel, vibration-monitoring channel, and shock-monitoring channel. The main inertial channel provides motion-dynamic information with higher time resolution; the low-frequency inertial/attitude channel provides attitude reference and low-frequency trajectory trends; the vibration-monitoring channel reflects continuous mechanical disturbances and local structural vibrations; and the impact-monitoring channel captures short-term collisions, impacts, or sudden disturbances. Since different channels have different sampling frequencies and physical meanings, each channel was assigned a specific role according to its bandwidth and physical meaning. The main inertial channel participates in dynamic constraints, the low-frequency branch provides trajectory soft prior, and the high-bandwidth channel provides motion-veto evidence [31–33].

In the experimental system, the four types of observations are provided by the 1 kHz main inertial channel, the 200 Hz low-frequency inertial/attitude channel, the 30 kHz vibration-monitoring channel, and the 50 kHz impact-monitoring channel. The main inertial branch refers to the processing path driven by the 1 kHz acceleration measurements. This branch provides navigation-frame acceleration for inertial state propagation and the dynamic constraints in the factor graph. All channels are acquired by the same composite sensor package and unified to a 1 ms master timeline in offline processing. Let the main timeline be

$$t_i = \frac{i}{f_0}, f_0 = 1000 \text{ Hz} \quad (1)$$

where i is the sampling sequence number of the main timeline. For channels with different sampling rates, low-sampling rate inertia/attitude information is mapped to the main time axis through interpolation, and high-sampling rate vibration and impact signals are aggregated at the block level according to the main time axis window. This not only ensures that subsequent evidence calculations and factor-graph optimization are performed under a unified time base, but also avoids direct resampling of high-bandwidth vibration and impact signals for displacement double integration.

For any sensing channel, s , the original observation is first calibrated through scale, bias, and axial-direction correction:

$$\tilde{a}_s(t) = D_s S_s (a_s^{\text{raw}}(t) - b_s) \quad (2)$$

where $a_s^{\text{raw}}(t)$ is the original acceleration observation, b_s is the offset term, S_s is the scale correction matrix, and D_s is the axial sign and installation direction correction matrix. The initial offset, b_s , of each acceleration axis was estimated as the median of the samples identified as stationary and was subtracted before filtering, coordinate transformation, and inertial reconstruction. The median estimator was selected to reduce the influence of vibration peaks and transient disturbances. For the main inertial acceleration channel, the calibrated body coordinate system acceleration is further converted to the navigation coordinate system using low-frequency attitude information, and the gravity term is deducted:

$$a^n(t) = R_{nb}(t) \tilde{a}_{1k}(t) - g^n \quad (3)$$

where $R_{nb}(t)$ represents the rotation matrix from the body coordinate system to the navigation coordinate system, and g^n is the gravity vector in the navigation coordinate system. The resulting $a^n(t)$ serves as the primary dynamic input for subsequent inertial propagation constraints and keyframe factor-graph optimization.

The low-frequency inertial/attitude channel is also used to construct the low-frequency trajectory prior. This branch is filtered, attitude-compensated, and zero-velocity-corrected to obtain a low-frequency displacement estimate, which is then mapped to the 1 ms main time axis through linear interpolation, keeping time consistent with the main inertial channel and back-end keyframe status. The corresponding velocity prior is obtained by differencing the interpolated low-frequency displacement sequence. The low-frequency displacements and velocities here are only used as soft priors in the back-end factor graph, and the weights are adjusted by indicators such as branch consistency, residual roughness, and stationary segment drift.

For vibration-monitoring and impact-monitoring channels, this study does not perform displacement integration. High-bandwidth responses mainly reflect local vibrations, structural disturbances, and transient impacts, and are valuable for judging whether the zero-velocity assumption is reliable; there is no direct equivalent relationship between them and the translational displacement of the load rigid body. Direct double integration of such signals can easily misinterpret the local high-frequency response as low-frequency displacement drift. This study extracts activity features such as RMS, peak, and band energy according to the main timeline window and converts them into motion-veto evidence for subsequent motion, stationary, and unknown confidence estimation.

After the above preprocessing, this study obtains three types of inputs on the unified main timeline: navigation system main acceleration, $a^n(t_i)$; low-frequency displacement/velocity prior, $p_0(t_i)$, $v_0(t_i)$; and high-bandwidth activity characteristics generated by vibration and shock channels. Subsequent methods will estimate the confidence of three types of states, namely motion, stationary, and uncertain, on this basis and further transfer

the uncertainty of the stationary state to ZUPT gating, multi-branch displacement fusion, and keyframe factor-graph reconstruction.

2.2. High-Frequency-Driven Motion-Veto Evidence Modeling

During the transportation of high-vibration industrial loads, the reliability of ZUPT depends both on whether the low-frequency acceleration or angular velocity is close to the stationary state, and is also affected by sustained vibration, local structural disturbances, and transient impacts. If a binary stationary determination is made based only on low-frequency inertial signals, ZUPT may be triggered incorrectly when the load still exhibits significant vibration or impact response, thus suppressing the actual velocity and introducing displacement deviations. This study transforms vibration-monitoring channels and shock-monitoring channels into motion-veto evidence to weaken unreliable stationary judgments.

For high-bandwidth channels, h , this study divides them into several short time windows according to the main timeline. Assuming that the k -th window is $\mathcal{B}_{h,k}$, then the RMS activity intensity and peak activity intensity within the window are respectively defined as follows [34]:

$$e_{h,k}^{\text{rms}} = \sqrt{\frac{1}{|\mathcal{B}_{h,k}|} \sum_{i \in \mathcal{B}_{h,k}} \|\tilde{a}_h(i)\|_2^2}, \quad e_{h,k}^{\text{pk}} = \max_{i \in \mathcal{B}_{h,k}} \|\tilde{a}_h(i)\|_2 \quad (4)$$

where $\tilde{a}_h(i)$ is the calibrated high-bandwidth acceleration observation, and $\mathcal{B}_{h,k}$ is the number of samples in the window. The RMS feature mainly reflects the sustained vibration energy in the window, whereas the peak feature is more sensitive to transient impacts and sudden disturbances. In practice, the two features can be weighted to obtain a composite high-frequency activity score [34]:

$$e_h(t) = \rho e_{h,k}^{\text{rms}} + (1 - \rho) e_{h,k}^{\text{pk}} \quad (5)$$

where $\rho \in [0, 1]$ is the weighting coefficient. To reduce the influence of amplitude-scale variation among different tests, installation states, and load conditions, robust within-record statistics are used to normalize the activity score [35]:

$$z_h(t) = \frac{e_{h,k} - Q_{0.50}(e_h)}{Q_{0.90}(e_h) - Q_{0.10}(e_h) + \epsilon_{\text{norm}}} \quad (6)$$

where $Q_{0.50}(\cdot)$, $Q_{0.90}(\cdot)$, and $Q_{0.10}(\cdot)$ denote the median, the 90th percentile, and the 10th percentile, respectively, and ϵ_{norm} is a stabilization term used to prevent the denominator from becoming too small. The normalized $z_h(t)$ represents the high-frequency activity intensity of the current window relative to the background level of the record.

This study uses three types of states, namely motion, stationary, and unknown, to describe the evidence distribution at each moment. The three types of states are recorded as M , S , and U , where M represents movement, S represents stationary, and U represents unknown or insufficient evidence. For vibration and shock channels, high-frequency activity enhances the evidence of motion, but weaker high-frequency activity does not directly prove that the load is at rest. Therefore, high-bandwidth channels only provide moving and unknown masses, not strong stationary masses [36,37]:

$$m_h(M, t) = \frac{\eta_h}{1 + \exp[-\kappa_h(z_h(t) - \tau_h)]} \quad (7)$$

$$m_h(S, t) = 0, \quad m_h(U, t) = 1 - m_h(M, t) \quad (8)$$

where $m_h(M, t)$, $m_h(S, t)$, and $m_h(U, t)$ are the evidence masses assigned by the high-bandwidth channel to the motion, stationary, and unknown states, respectively; τ_h is the high-frequency activity threshold; κ_h is the mapping slope; and η_h is the upper bound of the high-frequency motion evidence mass.

The low-frequency inertial channel provides both motion evidence and stationary evidence. Let the low-frequency activity score be z_l . It can be mapped to low-frequency evidence masses as follows [37]:

$$m_l(M, t) = \frac{1}{1 + \exp[-\kappa_M(z_l(t) - \tau_M)]} \quad (9)$$

$$m_l(S, t) = \frac{1}{1 + \exp[\kappa_S(z_l(t) - \tau_S)]} \quad (10)$$

$$m_l(U, t) = 1 - m_l(M, t) - m_l(S, t) \quad (11)$$

where $m_l(M, t)$, $m_l(S, t)$, and $m_l(U, t)$ denote the motion, stationary, and unknown evidence masses provided by the low-frequency inertial channel, respectively. The tolerances in Equations (7)–(10) are determined from robust statistics calculated independently for each measurement record. Specifically, the median and selected percentiles of the activity indicators are used to reduce sensitivity to transient outliers and variations in background noise. The percentile settings and mapping parameters are selected through preliminary sensitivity analysis and are subsequently kept fixed during displacement reconstruction. If negative values occur in Equation (11), the three evidence masses are non-negatively truncated and normalized so that their sum is one.

Multi-source evidence is fused by using a Dempster–Shafer-style combination rule [16–21]. For two evidence sources, their combined result is

$$m_{12}(A, t) = \frac{\sum_{B \cap C = A} m_1(B, t) m_2(C, t)}{1 - K(t)} \quad (12)$$

where $A, B, C \in \{M, S, U\}$, and the conflict term, $K(t)$, is defined as

$$K(t) = \sum_{B \cap C = \emptyset} m_1(B, t) m_2(C, t) \quad (13)$$

By sequentially fusing the low-frequency inertial evidence, vibration evidence, and impact evidence, the motion confidence, $P_m(t)$; stationary confidence, $P_s(t)$; and unknown confidence, $P_u(t)$, can be obtained at each sample on the main time axis. Here, $P_m(t)$ characterizes the possibility of motion or high-frequency disturbance at the current time, $P_s(t)$ characterizes the credibility that a zero-velocity constraint is suitable, and $P_u(t)$ characterizes insufficient evidence or multi-source conflict.

In order to pass the evidence results to the subsequent ZUPT and factor-graph reconstruction, this study further constructs continuous stationary confidence [38]:

$$\pi_s(t) = \text{clip}\{P_s(t)[1 - P_m(t)][1 - P_u(t)], 0, 1\}. \quad (14)$$

Equation (14) shows that the stationary confidence is high only when the stationary evidence is strong, the motion evidence is weak, and the unknown degree is low. When vibration or impact activity is significant, $P_m(t)$ increases and $\pi_s(t)$ decreases accordingly. When multi-source evidence is conflicting or insufficient, $P_u(t)$ increases, and the stationary confidence is also weakened.

2.3. Confidence-Weighted ZUPT and Multi-Branch Displacement Priors

In inertial displacement reconstruction, ZUPT is an important constraint to suppress double-integral drift. Traditional ZUPT is usually based on fixed threshold stationary detection: when a certain moment is determined to be stationary, a pseudo-measurement of zero-velocity is directly applied. However, during the transportation of high-vibration industrial loads, the stationary state has obvious uncertainties: low-frequency inertial signals may appear smooth, but high-frequency vibration or shock may still indicate that the load is not suitable for imposing strong zero-velocity constraints. Therefore, this study does not use a single binary stationary label as the basis for ZUPT triggering; instead, it uses the continuous stationary confidence obtained in the previous section to adjust the activation strength and covariance weight of ZUPT.

For the navigation-frame velocity, $v(t)$, the ZUPT pseudo-measurement can be written as

$$z_{\text{zupt}}(t) = 0 \quad (15)$$

The corresponding residual is

$$r_{\text{zupt}}(t) = v(t) - z_{\text{zupt}}(t) = v(t) \quad (16)$$

In order to avoid imposing strong constraints on motion or uncertainty intervals, this study adaptively adjusts the ZUPT covariance according to the stationary confidence:

$$R_{\text{zupt}}(t) = \frac{R_0}{\pi_s(t) + \epsilon} \quad (17)$$

where R_0 is the basic zero-velocity covariance, and ϵ is a stabilization term used to prevent the denominator from becoming too small. When $\pi_s(t)$ is high, strong stationary evidence is available at the current time, the ZUPT covariance is reduced, and the zero-velocity constraint is strengthened. When motion confidence, unknown confidence, or multi-source conflict is high, $\pi_s(t)$ is reduced, the ZUPT covariance is increased, and the zero-velocity constraint is weakened. In practice, a candidate stationary threshold, τ_{zupt} , can be further set so that ZUPT is enabled only in the candidate interval of $\pi_s(t) > \tau_{\text{zupt}}$, while the constraint strength is still determined by the continuous weight.

In the front-end reconstruction, the low-frequency inertial/attitude branch is used to generate soft priors for displacement and velocity. After filtering, attitude compensation, zero-velocity correction, and drift suppression, a low-frequency displacement sequence, $p_{200}(t)$, is obtained by this branch. Because the back-end reconstruction is performed on the 1 ms main time axis, the low-frequency displacement sequence is interpolated onto the main time axis to obtain $p_0(t)$. The corresponding velocity prior is obtained by differencing $p_0(t)$ and is denoted as $v_0(t)$. $p_0(t)$ and $v_0(t)$ only represent the trajectory trend provided by the low-frequency branch and are used as soft priors in the subsequent factor graph.

Meanwhile, the main inertial channel provides the navigation-frame acceleration, $a^n(t)$, which is used to describe motion dynamics with higher temporal resolution. Compared with the low-frequency displacement prior, the main inertial acceleration can preserve richer transient motion information, but direct double integration is susceptible to bias and vibration disturbance. Therefore, a complementary strategy is adopted: the low-frequency branch provides a relatively stable trajectory scale and slowly varying trend, the main inertial channel provides dynamic propagation constraints, and the high-bandwidth vibration/impact channels affect the ZUPT strength through stationary confidence.

To prevent the low-frequency displacement prior from excessively dominating the back-end optimization under local abnormalities, its reliability is adjusted according to branch consistency, residual roughness, and stationary-segment drift. Let $p_{1k}(t)$ be the

auxiliary displacement estimate obtained from the main inertial branch through basic integration and drift suppression. The branch difference in the q -th direction can then be expressed as

$$d_p(t) = |p_{0,q}(t) - p_{1k,q}(t)| \quad (18)$$

where $q \in \{x, y, z\}$. The roughness of the low-frequency trajectory in a local window can be characterized by the second-order difference:

$$\rho_q(t) = |p_{0,q}(t + \Delta t) - 2p_{0,q}(t) + p_{0,q}(t - \Delta t)| \quad (19)$$

In addition, within platform intervals with high stationary confidence, if the low-frequency displacement still changes significantly, a stationary-segment drift risk is indicated for this branch. The stationary-segment drift index is denoted as $\delta_q^{\text{stat}}(t)$. Considering the above factors, the direction-dependent standard deviation of the low-frequency position prior can be written as

$$\sigma_{p,q}(t) = \sigma_{p,0} + c_d d_q(t) + c_p \rho_q(t) + c_s \delta_q^{\text{stat}}(t) + c_u P_u(t) \quad (20)$$

where $\sigma_{p,0}$ is the basic position-prior standard deviation, and c_d , c_p , c_s , and c_u are the weighting coefficients for branch difference, trajectory roughness, stationary-segment drift, and unknown confidence, respectively. Equation (20) indicates that when the low-frequency displacement differs greatly from the main inertial branch, the local trajectory is not smooth, drift still exists in the stationary segment, or the unknown confidence is high, the position-prior standard deviation in this direction is increased and the corresponding constraint weight is reduced.

Similarly, the standard deviation of the velocity prior can be adjusted according to the low-frequency velocity stability and stationary confidence [39]:

$$\sigma_{v,q}(t) = \sigma_{v,0} + c_v |v_{0,q}(t)| (1 - \pi_s(t)) + c_u^v P_u(t) \quad (21)$$

where $\sigma_{v,0}$ is the basic velocity-prior standard deviation, and c_v and c_u^v are adjustment coefficients. c_v controls the covariance inflation caused by instability in the low-frequency velocity prior, while c_u^v controls the inflation caused by decreasing stationary confidence. In high-confidence stationary intervals, the velocity prior and ZUPT jointly constrain the velocity to be close to zero. In motion or unknown intervals, the velocity-prior weight is reduced to avoid incorrect suppression of true motion.

After the above processing, three types of information are provided by the front end for back-end reconstruction: first, low-frequency position and velocity soft priors, $p_0(t)$ and $v_0(t)$; second, the main acceleration in the navigation frame, $a^n(t)$; and third, the continuous stationary confidence, $\pi_s(t)$, obtained from motion, stationary, and unknown evidence. Here, $p_0(t)$ and $v_0(t)$ are used to constrain the low-frequency trend of the trajectory, $a^n(t)$ is used to establish dynamic consistency between adjacent keyframes, and $\pi_s(t)$ is used to adjust the strength of ZUPT and stationary-related factors.

2.4. Evidence-Weighted Keyframe Factor-Graph Reconstruction

After the front-end obtains the low-frequency position/velocity soft prior, the main acceleration of the navigation system, and the continuous stationary confidence, this section further constructs a keyframe factor graph to jointly smooth the load displacement, velocity, and acceleration offset. Different from direct point-by-point double integration of acceleration, factor graphs can integrate low-frequency trajectory priors, inertial dynamic consistency, zero-velocity constraints, and bias-smoothing constraints in a unified optimiza-

tion framework, thereby reducing the impact of local noise, acceleration bias, and false triggering of zero-velocity updates on the final displacement.

Assume that K keyframes are selected on the main time axis and that the keyframe times are t_k . The corresponding state is defined as

$$x_k = [p_k^\top, v_k^\top, b_{a,k}^\top]^\top, k = 1, 2, \dots, K \quad (22)$$

where p_k is the keyframe position, v_k is the keyframe velocity, and $b_{a,k}$ is the acceleration bias. The keyframe state is required to remain close to the low-frequency trajectory trend provided by the front end and to satisfy the adjacent-frame dynamic constraints provided by the main inertial acceleration.

The low-frequency branch provides soft priors for position and velocity. Let the low-frequency position and velocity priors at the k -th keyframe be $p_{0,k}$ and $v_{0,k}$, respectively. The corresponding residuals are

$$r_{p,k}^0 = p_k - p_{0,k} \quad (23)$$

$$r_{v,k}^0 = v_k - v_{0,k} \quad (24)$$

Here, $p_{0,k}$ and $v_{0,k}$ are the soft priors generated by the low-frequency branch. Their covariance is jointly adjusted by branch consistency, trajectory roughness, stationary-segment drift, and unknown confidence, as described in the previous section. When the low-frequency prior is stable and consistent with the main inertial branch, higher weights are assigned to the position and velocity priors. When the branch difference is large or the unknown confidence is high, the corresponding constraint weight is reduced.

Dynamic propagation constraints are established between adjacent keyframes by the main inertial acceleration. Let the interval between adjacent keyframes be

$$\Delta t_k = t_{k+1} - t_k \quad (25)$$

and let the mean main acceleration in the navigation frame over this interval be $\bar{a}_k^n$. Over a short time interval, the position and velocity propagation relationships can be approximated as

$$p_{k+1} = p_k + v_k \Delta t_k + \frac{1}{2}(\bar{a}_k^n - b_{a,k}) \Delta t_k^2 \quad (26)$$

$$v_{k+1} = v_k + (\bar{a}_k^n - b_{a,k}) \Delta t_k \quad (27)$$

The corresponding position and velocity dynamic residuals are

$$r_{p,k}^{\text{dyn}} = p_{k+1} - p_k - v_k \Delta t_k - \frac{1}{2}(\bar{a}_k^n - b_{a,k}) \Delta t_k^2 \quad (28)$$

$$r_{v,k}^{\text{dyn}} = v_{k+1} - v_k - (\bar{a}_k^n - b_{a,k}) \Delta t_k \quad (29)$$

Equations (28) and (29) belong to weak dynamic consistency factors for low-speed industrial load transportation. This factor is used to constrain the keyframe displacement, velocity, and main inertial acceleration to be consistent in a low-frequency sense. For scenarios with strong angular vibrations, rapid attitude changes, or significant non-rigid body shaking, attitude errors may be converted into equivalent acceleration biases through gravity leakage, thereby affecting velocity and displacement estimates. This study jointly suppresses such errors through low-frequency trajectory soft prior, evidence-weighted ZUPT, stationary segment position constraints, and acceleration bias estimation.

The temporal evolution of the acceleration-bias state is modeled as a Gaussian random-walk process:

$$b_{a,k+1} = b_{a,k} + w_{b,k}, \quad w_{b,k} \sim \mathcal{N}(0, Q_{b,k}) \quad (30)$$

The corresponding bias-increment residual incorporated into the factor-graph objective is

$$r_{b,k} = b_{a,k+1} - b_{a,k} \quad (31)$$

Here, $b_{a,k}$ denotes the three-axis acceleration-bias state at keyframe, k ; $Q_{b,k} = \sigma_b^2 \Delta t_k \mathbf{I}_3$ is the covariance matrix of the inter-keyframe bias increment; σ_b denotes the corresponding acceleration-bias random-walk standard deviation; and Δt_k denotes the time interval between keyframes k and $k + 1$. This factor penalizes unlikely abrupt bias changes while allowing the bias to evolve gradually over time. The bias state is jointly estimated with position and velocity in the factor-graph optimization.

For keyframes with high stationary confidence, an evidence-weighted ZUPT factor is introduced. Let the continuous stationary confidence at the k -th keyframe be $\pi_{s,k} = \pi_s(t_k)$. The ZUPT weight is then defined as follows:

$$w_k = \text{clip}(\pi_{s,k}, w_{\min}, w_{\max}) \quad (32)$$

The corresponding zero-velocity residual is as follows [39]:

$$r_{\text{zupt},k} = \frac{\sqrt{w_k}}{\sigma_v} v_k \quad (33)$$

where σ_v is the basic velocity-constraint standard deviation. Unlike fixed-strength ZUPT, Equation (32) allows the zero-velocity constraint strength to vary continuously with stationary confidence. When vibration or impact evidence is enhanced, $\pi_{s,k}$ decreases and the ZUPT weight is reduced. When low-frequency inertial evidence and high-frequency motion-veto evidence jointly support stationarity, the ZUPT weight is increased, so velocity drift in stationary segments can be more effectively suppressed.

To reduce displacement overshoot caused by residual velocity in platform–motion–platform transition segments, a weak main-direction kinematic constraint is introduced. Motion segments are first extracted from low-stationary-confidence intervals, and the main displacement direction of the r -th segment is estimated from the low-frequency trajectory prior as

$$d_r = \frac{p_0(t_r^1) - p_0(t_r^0)}{\|p_0(t_r^1) - p_0(t_r^0)\|_2 + \epsilon} \quad (34)$$

where t_r^0 and t_r^1 denote the start and end times of the motion segment, respectively. For a keyframe, $t_k \in \mathcal{M}_r$, the velocity component along the estimated main direction is first projected, and only its negative component is converted into a soft reverse-velocity penalty:

$$v_{r,k}^{\text{rev}} = \max(0, -d_r^\top v_k), t_k \in \mathcal{M}_r \quad (35)$$

This constraint suppresses only obvious reverse velocity along the estimated transport direction; it does not impose monotonicity on each coordinate axis and does not force any motion segment to reach a preset displacement. Therefore, it is a weak kinematic consistency term rather than a semantic stage prior.

$$r_{\text{mono},k} = \frac{v_{r,k}^{\text{rev}}}{\sigma_{\text{mono}}} \quad (36)$$

where σ_{mono} denotes the velocity-scale standard deviation of the reverse-velocity soft constraint, which controls the penalty strength of the monotonic consistency factor. For

continuous platform segments with high stationary reliability, the stationary segment position locking residual is further introduced:

$$r_{\text{lock},k} = \frac{p_{k+1} - p_k}{\sigma_{\text{lock}}}, \quad \pi_{s,k} > \tau_{\text{lock}}, \quad \pi_{s,k+1} > \tau_{\text{lock}} \quad (37)$$

This item is used to suppress the phenomenon of “velocity close to zero but displacement slowly drifting” in the stationary section. Different from the endpoint closure constraint, Equation (37) only constrains the positions of adjacent keyframes inside the stationary platform to be consistent, and it does not use external references or theoretical trajectories.

Based on the above factors, the complete keyframe factor-graph optimization objective is formulated as

$$\begin{aligned} \min_{\{x_k\}} & \sum_k \|r_{p,k}^0\|_{\Sigma_{p,k}^{-1}}^2 + \sum_k \|r_{v,k}^0\|_{\Sigma_{v,k}^{-1}}^2 \\ & + \sum_k \|r_{p,k}^{\text{dyn}}\|_{\Sigma_{p,\text{dyn}}^{-1}}^2 + \sum_k \|r_{v,k}^{\text{dyn}}\|_{\Sigma_{v,\text{dyn}}^{-1}}^2 \\ & + \sum_k \|r_{b,k}\|_{Q_{b,k}^{-1}}^2 + \sum_{k \in \mathcal{S}} \|r_{\text{ZUPT},k}\|_2^2 \\ & + \sum_{k \in \mathcal{M}} \|r_{\text{mono},k}\|_2^2 + \sum_{k \in \mathcal{P}} \|r_{\text{lock},k}\|_2^2 \end{aligned} \quad (38)$$

Equation (38) defines the complete factor-graph optimization objective by combining all the residual terms described above. Minimizing this objective yields the joint estimates of keyframe position, velocity, and acceleration bias, from which the final displacement trajectory is reconstructed, where $\Sigma_{p,k}$ and $\Sigma_{v,k}$ denote the covariance matrices of the position and velocity soft priors, respectively, and vary according to branch consistency, trajectory roughness, stationary-segment drift, and unknown confidence; and $\Sigma_{p,\text{dyn}}$ and $\Sigma_{v,\text{dyn}}$ denote the covariance matrices of the position and velocity dynamic-propagation residuals, respectively. The matrix, $Q_{b,k}$, is the covariance of the inter-keyframe acceleration-bias increment defined in Equation (30); and $\mathcal{S}$, $\mathcal{M}$, and $\mathcal{P}$ denote the set of high-confidence stationary keyframes, the set of motion-segment keyframes, and the set of stationary-platform keyframes, respectively.

3. Experimental and Analysis

In order to verify the effectiveness of the proposed evidence-weighted multi-rate inertial fusion method in industrial load transportation scenarios, this study conducted an offline reconstruction experiment around a cylindrical load. The experiment focuses on investigating the effects of multi-rate inertial observations, vibration/impact motion-veto evidence, confidence-weighted ZUPT, and keyframe factor-graph smoothing on three-axis displacement reconstruction.

The experimental object is a cylindrical load, and the sensor assembly is rigidly installed inside the cylindrical load to synchronously collect the inertia, vibration, and impact response during load transportation. The experimental system employed a composite sensor package rigidly installed inside the cylindrical payload. Before data acquisition, the payload was maintained stationary during the built-in initialization and zero-calibration procedure. The package integrated four functional sensing channels: a 1 kHz main triaxial acceleration channel, a 200 Hz inertial/attitude channel, a 30 kHz triaxial vibration-monitoring channel, and a 50 kHz triaxial impact-monitoring channel. Their principal measurement specifications and functional roles in the proposed reconstruction framework are summarized in Table 1. The 1 kHz channel provides dynamic acceleration information with higher time resolution, the 200 Hz channel provides low-frequency trend and attitude assistance, the 30 kHz channel mainly reflects continuous vibration information, and the 50 kHz channel mainly captures transient impact information. The 30 kHz and 50 kHz

channels do not participate in the double integration of displacement and are only used as evidence to refute vibration/shock-driven motion to assist in judging whether the current zero-velocity assumption is reliable.

Table 1. Multi-rate sensing channels and their functions.

Channel	Measured Quantity and Range	Sampling Rate	Resolution/Accuracy	Main Function
Main inertial channel	Triaxial acceleration, ± 5 g	1 kHz	12 bit; $\pm 1\%$ amplitude linearity	High-time-resolution dynamic acceleration information is provided
Low-frequency inertial/attitude auxiliary channel	Triaxial angular rate, attitude angle and body-frame acceleration; accelerometer ± 16 g, gyroscope $\pm 2000^\circ/\text{s}$	200 Hz	Dynamic roll/pitch accuracy: 0.1° RMS; heading accuracy: 0.5° RMS	Low-frequency trends, attitude assistance, and displacement priors are provided
Vibration-monitoring channel	Triaxial vibration acceleration, ± 10 g	30 kHz	12 bit; $\pm 1\%$ amplitude linearity	Continuous vibration activity is characterized and used in motion-veto evidence modeling
Impact-monitoring channel	Triaxial impact acceleration, ± 100 g	50 kHz	12 bit; $\pm 1\%$ amplitude linearity	Transient impacts are captured and used in motion-veto evidence modeling

The sensor assembly was rigidly fixed inside the cylindrical load, and its orientation relative to the load remained unchanged throughout the experiment. However, the initial horizontal sensor axes could not be guaranteed to coincide with the actual horizontal motion directions. Therefore, two independent motion intervals along mutually perpendicular horizontal directions were used to establish a fixed horizontal coordinate system for result presentation and evaluation. The same fixed coordinate transformation was applied to all evaluated methods.

During the experiment, the cylindrical load underwent multiple repeated transport stages with motion and platform-holding intervals, and the total recording time was approximately 125 s. The raw signals collected during the experiment are shown in Figure 1. In offline processing, all data are uniformly mapped to the 1 ms main time axis: the 200 Hz low-frequency channel is aligned to the main time axis through interpolation, and the high-frequency vibration and shock channels are used to extract activity features through short-window aggregation. Subsequently, the algorithm successively completed signal calibration, navigation system acceleration construction, motion-stationary-unknown evidence estimation, confidence-weighted ZUPT, multi-branch displacement prior construction, and keyframe factor-graph reconstruction. The final displacement reconstruction results are presented in Figure 2.

Figure 1 presents the raw multi-rate observations before displacement reconstruction. The low-frequency angular velocity and acceleration channels show intermittent motion intervals separated by long quasi-stationary platforms, whereas the 30 kHz vibration and 50 kHz impact channels exhibit dense high-bandwidth fluctuations during transport and transition periods; this difference supports using high-bandwidth signals as motion-veto evidence rather than direct displacement inputs.

The cylindrical load was transported by a transfer tooling system according to a factory-defined operating procedure. The complete operation consisted of a sequence of

prescribed transport and holding operations. For each operation, the motion direction, target displacement, operating velocity, and holding state were specified by the transfer tooling system. These predefined operating parameters were used to construct the reference trajectory shown in Figure 2. The reference trajectory was used exclusively for offline accuracy evaluation and was not introduced into stationary detection, inertial integration, evidence fusion, or factor-graph optimization. The shaded regions in Figure 2 denote the ZUPT intervals, during which the three-dimensional zero-velocity factors were activated. Figure 2 compares the reconstructed displacement with the reference trajectory in the X, Y, and Z directions. The reconstructed curves generally follow the staged displacement changes and platform-holding periods in all three coordinate directions, indicating that the proposed method recovers the overall three-axis displacement evolution of the load.




Figure 1. Raw signals collected by the sensor. (a) 200 Hz angular velocity; (b) 200 Hz angular acceleration; (c) 200 Hz acceleration; (d) 1 kHz acceleration; (e) 30 kHz vibration acceleration; (f) 50 kHz impact acceleration.




Figure 2. Displacement reconstruction results. The shaded regions indicate the ZUPT intervals, during which the three-dimensional zero-velocity constraint is activated. (a) Displacement in the x-direction; (b) Displacement in the y-direction; (c) Displacement in the z-direction.

3.1. Comparison of Displacement Reconstruction Results

To evaluate the proposed method against representative recent inertial displacement reconstruction approaches under controlled and reproducible conditions, two literature-informed baseline pipelines were implemented for comparison. All three methods used the same sensor records, calibrated inputs, master time base, and output-coordinate transformation; only their fusion and constraint strategies were different. Baseline 1 represents the multi-rate filtering and weighted-fusion route commonly used in recent multi-sensor and multi-inertial fusion studies [6,10,29,30]. It exploits the complementary characteristics of inertial branches with different sampling rates to reconstruct the displacement trend, but it does not introduce vibration/impact-based motion-veto evidence or factor-graph smoothing. Baseline 2 combines adaptive stationary detection, ZUPT, and keyframe factor-graph smoothing, representing the adaptive ZUPT-FGO route used for inertial drift suppression [11–15,22–28]. However, it does not use high-frequency vibration and impact observations to evaluate the reliability of stationary decisions. These baselines were selected from recent methodological routes that can operate under the same inertial sensing configuration. Approaches requiring additional external positioning sensors or task-specific training data were not included, because they would change the sensing conditions and prevent a controlled comparison. The reference trajectory was used only for quantitative evaluation and was not introduced into any reconstruction method. This controlled comparison enables the effects of motion-veto evidence and continuous stationary-confidence propagation to be evaluated without introducing differences in sensor configuration, input data, or coordinate definition.

Figure 3 compares the reconstructed displacements of the three methods along the three navigation axes. All methods reproduce the main stage-wise variation in the load displacement, but differences remain in reconstruction accuracy and trajectory shape. Baseline 1 recovers the approximate displacement trend through multi-rate inertial fusion,

but noticeable drift, transient reversal, and amplitude deviation remain during holding and motion-transition intervals. Baseline 2 improves trajectory continuity and suppresses part of the holding-interval drift by introducing ZUPT and factor-graph smoothing. Nevertheless, because the reliability of stationary detection is not adjusted using vibration and impact observations, false zero-velocity constraints and local trajectory distortion may still occur during high-vibration intervals. In comparison, the proposed method is generally closer to the reference trajectory along the three axes and better preserves the displacement increments, holding plateaus, and motion-transition characteristics. These results demonstrate the contribution of vibration/impact-based motion-veto evidence and continuous confidence propagation to ZUPT weighting, displacement fusion, and keyframe factor-graph optimization.




Figure 3. Comparison of displacement reconstruction results of different methods in the three coordinate directions. (a) Comparison of x-directional displacement reconstruction results; (b) Comparison of y-directional displacement reconstruction results; (c) Comparison of z-directional displacement reconstruction results.

In order to quantitatively evaluate the reconstruction accuracy of the three methods, this study uses RMSE in the three directions of X, Y, and Z and 3D RMSE as evaluation indicators. Among them, the single-axis RMSE is used to reflect the average reconstruction deviation in each coordinate direction, and the 3D RMSE is used to comprehensively evaluate the consistency between the entire three-dimensional trajectory and the reference displacement. Table 2 gives the quantitative comparison results of the three methods.

Table 2. Quantitative comparison of different baseline.

Method	X RMSE	Y RMSE	Z RMSE	3D RMSE
Baseline 1	1.281	1.362	0.442	1.921
Baseline 2	0.478	1.085	0.325	1.229
Proposed algorithm	0.274	0.522	0.280	0.652

As shown in Table 2, the proposed method achieves the lowest RMSE values in all three coordinate directions and in the 3D RMSE. Its 3D RMSE is 0.652 m, which is lower than 1.921 m for Baseline 1 and 1.229 m for Baseline 2, corresponding to reductions of approximately 66.1% and 46.9%, respectively. In the X, Y, and Z directions, the RMSE values of the proposed method are 0.274 m, 0.522 m, and 0.280 m, respectively, all lower than those of the two comparison methods. This demonstrates that the proposed method improves not only single-axis reconstruction accuracy but also the consistency of the overall three-dimensional displacement trajectory. The error of Baseline 1 mainly comes from integral drift and platform-segment offsets. Although Baseline 2 reduces part of the drift through ZUPT and graph optimization, its stationary-reliability judgment in high-vibration intervals remains insufficient, so the overall error is still higher than that of the proposed method.

Figure 3 and Table 2 further show that the proposed method outperforms the two comparison baseline in both three-dimensional trajectory consistency and single-axis reconstruction accuracy. Its improvement mainly comes from three aspects: first, the vibration/shock motion-veto evidence reduces the risk of falsely triggering ZUPT in high-vibration intervals; second, the confidence-weighted ZUPT can adjust the zero-velocity constraint strength according to the stationary confidence to avoid over-suppressing uncertain motion segments as stationary; finally, the keyframe factor graph puts the low-frequency trajectory prior, main inertial dynamic constraints, and bias-smoothing constraints into a unified optimization framework, improving trajectory continuity and platform stability. The above results show that the proposed method can obtain more stable three-axis displacement reconstruction results in cylindrical load transport experiments.

3.2. Parameter Sensitivity Analysis

To evaluate the influence of the main parameters on displacement reconstruction, a one-at-a-time sensitivity analysis was performed around the nominal configuration. Five representative fixed parameters were selected: the stabilization term, ϵ ; the basic zero-velocity covariance, R_0 ; the candidate stationary threshold, τ_{zupt} ; the basic position-prior noise density, $\sigma_{p,0}$; and the acceleration-bias random-walk noise density, σ_b . Their nominal values were 0.03, 0.30, 1.6×10^{-4} , 0.385, and 0.40, respectively. The nominal configuration was selected through preliminary bounded parameter calibration and was subsequently fixed for the sensitivity analysis.

Each parameter was independently varied to 60%, 80%, 100%, 120%, and 140% of its nominal value, while the remaining parameters were kept unchanged. Because the five parameter groups shared the same nominal configuration, a total of 21 distinct reconstruction configurations were evaluated. The horizontal axis in Figure 4 is the normalized parameter value, θ/θ_0 , where θ_0 denotes the nominal value.

As shown in Figure 4, the nominal configuration produced X-, Y-, Z-, and three-dimensional RMSE values of 0.2741, 0.5217, 0.2798, and 0.6524 m, respectively. Within the investigated range, the maximum relative changes in three-dimensional RMSE caused by ϵ , R_0 , $\sigma_{p,0}$, and σ_b were 0.416%, 1.721%, 1.036%, and 0.305%, respectively. The unchanged result for τ_{zupt} indicates that the same candidate stationary intervals were retained throughout the tested range. The Z-direction RMSE remained nearly unchanged, whereas the main variations occurred in the horizontal reconstruction. The nominal configuration was

located at or close to the lower envelope of the three-dimensional RMSE curves. Although $R_0 = 1.4R_{0,0}$ produced a numerically lower RMSE than the nominal value, the difference was only approximately 5.9×10^{-6} m and was negligible. These results indicate that moderate perturbations of the selected parameters do not materially affect the reconstruction accuracy in the present experiment.




Figure 4. One-at-a-time sensitivity of displacement reconstruction RMSE to the normalized parameter value. ϵ , stabilization term in the confidence-adaptive ZUPT covariance; τ_{zupt} , candidate stationary threshold; R_0 , basic zero-velocity covariance; $\sigma_{p,0}$, basic position-prior noise density; and σ_b , acceleration-bias random-walk noise density. The vertical dashed line denotes the nominal configuration, $\theta/\theta_0 = 1$. (a) 3D RMSE; (b) X RMSE; (c) Y RMSE; (d) Z RMSE.

4. Discussion

Experimental results show that the proposed evidence-weighted multi-rate inertial fusion method can achieve more stable displacement reconstruction in high-vibration cylindrical load transport scenarios. Compared with the basic multi-rate fusion and ZUPT-factor-graph smoothing methods, the algorithm in this study achieves lower errors in the three directions of X, Y, Z, and 3D RMSE. This shows that it is necessary to divide the functions of low-frequency inertial/attitude information, main inertial dynamic information, and vibration/shock observation information. The low-frequency channel provides trajectory trends and scale constraints, the main inertial channel preserves motion dynamics, and the vibration and shock channels are used to determine whether the zero-velocity assumption is reliable. Converting high-bandwidth observations into motion-veto evidence helps reduce the suppression of true motion by false ZUPT in high-vibration intervals.

From the perspective of error distribution, the proposed method reduces reconstruction deviations in the X, Y, and Z directions and improves the overall three-dimensional trajectory consistency. The current results show that the reconstruction error is not only determined by a single zero-velocity detection link, but is also related to the acceleration integral area, attitude compensation error, low-frequency trajectory prior stability, and

velocity tail term in the motion transition section. ZUPT can effectively suppress the velocity drift in the stationary section, but it is difficult to handle the state uncertainty under high-vibration conditions only by relying on fixed-strength zero-velocity constraints. This study alleviates this problem through motion-veto evidence and confidence weighting mechanisms, and further improves trajectory continuity through factor-graph smoothing.

Overall, this study experimentally verifies the complementary value of multi-rate inertial data and high-bandwidth vibration/shock observations in industrial load displacement reconstruction. By converting vibration and shock information into motion-veto evidence and passing stationary confidence to the ZUPT and factor-graph optimization processes, the proposed method outperforms the two contrasting methods in both 3D RMSE and single-axis RMSE values. This result shows that inertial displacement reconstruction for high-vibration industrial scenarios should not rely solely on direct double integration of acceleration or fixed-threshold zero-velocity detection; instead, it should be completed in combination with multi-source evidence, uncertainty propagation, and back-end smoothing optimization. This provides a diagnosable, auditable displacement status input for industrial digital twins, load status awareness, and maintenance decisions.

5. Conclusions

This study proposes an evidence-weighted multi-rate inertial fusion method for industrial load displacement reconstruction under high-vibration and limited-reference conditions. High-bandwidth vibration and impact responses are converted into motion-veto evidence, fused with low-frequency inertial evidence to estimate stationary confidence, and propagated to ZUPT weighting, multi-branch priors, and keyframe factor-graph smoothing. Cylindrical-load experiments show that the proposed method achieves lower 3D RMSE and lower single-axis RMSE values in the X, Y, and Z directions than the two comparison methods; improves platform stability and staged displacement reconstruction; and outputs interpretable diagnostics such as stationary confidence, ZUPT activation, residuals, and drift risk. Remaining direction-dependent and transition-stage errors indicate that attitude compensation and vibration-induced bias modeling require further improvement, while validation under a broader range of motion profiles and vibration conditions remains an important direction for future work.

Author Contributions: Conceptualization, P.Z., Q.Z. and S.S.; methodology, P.Z., W.L. and S.S.; software, P.Z. and Q.Z.; validation, P.Z., W.L. and Q.Z.; formal analysis, P.Z., Q.Z. and J.S.; investigation, P.Z. and Q.Z.; resources, P.Z., Q.Z. and J.S.; data curation, P.Z. and C.L.; writing—original draft, P.Z.; writing—review and editing, P.Z., W.L., Q.Z., C.L. and S.S.; visualization, C.L.; supervision, P.Z., W.L. and S.S.; project administration, W.L. and S.S.; funding acquisition, W.L. and S.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Natural Science Foundation of Hubei Province (2023AFB900).

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

Conflicts of Interest: Author Jixin Song was employed by No. 208 Research Institute of China Ordnance Industries. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

DS	Dempster–Shafer
FGO	Factor-graph optimization
GNSS	Global Navigation Satellite System
IMU	Inertial Measurement Unit
INS	Inertial Navigation System
iSAM2	Incremental Smoothing and Mapping 2
RMS	Root Mean Square
RMSE	Root Mean Square Error
RTK	Real-Time Kinematic
UWB	Ultra-Wideband
VIO	Visual-inertial odometry
ZUPT	Zero-velocity update

References

- Renard, D.; Saddem, R.; Annebicque, D.; Riera, B. From Sensors to Digital Twins toward an Iterative Approach for Existing Manufacturing Systems. *Sensors* **2024**, *24*, 1434. [\[CrossRef\]](#)
- Dadfarnia, M.; Sharp, M.; Herrmann, J. Comprehensive evaluations of condition monitoring-based technologies in industrial maintenance: A systematic review. *J. Manuf. Syst.* **2025**, *82*, 449–477. [\[CrossRef\]](#)
- Cao, Y. Multisensor Data Fusion-Driven Digital Twins in Computer Numerical Control Machining: A Review. *Machines* **2025**, *13*, 921. [\[CrossRef\]](#)
- Alfaro-Viquez, D.; Zamora-Hernandez, M.; Fernandez-Vega, M.; Garcia-Rodriguez, J.; Azorin-Lopez, J. A Comprehensive Review of AI-Based Digital Twin Applications in Manufacturing: Integration Across Operator, Product, and Process Dimensions. *Electronics* **2025**, *14*, 646. [\[CrossRef\]](#)
- Lapiente, G.A.; Avinzano, P.M.; Castillo, M.P.L.; Seneviratne, D. Methodologies in digital twin for manufacturing industry: A systematic literature review. *Future Gener. Comput. Syst.* **2025**, *174*, 107997. [\[CrossRef\]](#)
- Jin, Z.; Chen, G.; Niu, Y.; Zhang, C.; Zhang, X.; Shu, J. Variational mode decomposition-based multirate data-fusion framework for estimating structural dynamic displacement by integrating vision- and acceleration-based measurements. *Mech. Syst. Signal Process.* **2024**, *211*, 111252. [\[CrossRef\]](#)
- Niu, Y.; Li, Z.; Li, J.; Sun, B. Accelerometer-assisted computer vision data fusion framework for structural dynamic displacement reconstruction. *Measurement* **2025**, *242*, 116021. [\[CrossRef\]](#)
- Gao, S.; Chen, X.; Pan, Z.; Tian, Z.; Liu, F. A method for reconstructing the dynamic displacement of floating structures based on acceleration measurements and comparison with data-reconstruction techniques. *Commun. Eng.* **2025**, *4*, 68. [\[CrossRef\]](#) [\[PubMed\]](#)
- Huang, X.; Zhang, Z.; Cheng, T.; Li, J.; Ma, J. Distributed strain monitoring method for structural vibration based on multi-point acceleration measurement. *Measurement* **2024**, *230*, 114520. [\[CrossRef\]](#)
- Wang, J.; Wang, L.; Lin, J.; Xiao, R.; Chen, J.; Jin, P. Sensor fusion noise suppression method based on finite impulse response complementary filters. *Measurement* **2024**, *232*, 114680. [\[CrossRef\]](#)
- Li, M.; Nie, W.; Suvorkin, V.; Rovira-Garcia, A.; Zhang, W.; Xu, T.; Xu, G. Stationary Detection for Zero Velocity Update of IMU Based on the Vibrational FFT Feature of Land Vehicle. *Remote Sens.* **2024**, *16*, 902. [\[CrossRef\]](#)
- Zhang, W.; Wang, L.; Zu, Y. The Zero-Velocity Correction Method for Pipe Jacking Automatic Guidance System Based on Fiber Optic Gyroscope. *Sensors* **2024**, *24*, 5911. [\[CrossRef\]](#) [\[PubMed\]](#)
- Li, Q.; Li, K.; Liang, W. A zero-velocity update method based on neural network and Kalman filter for vehicle-mounted inertial navigation system. *Meas. Sci. Technol.* **2023**, *34*, 045110. [\[CrossRef\]](#)
- Zhang, S.; Gao, H.; Yang, F. Research on Wearable Devices for Pedestrian Navigation Based on the Informer Model Zero-Velocity Update Architecture. *Sensors* **2025**, *25*, 2587. [\[CrossRef\]](#) [\[PubMed\]](#)
- Skog, I.; Handel, P.; Nilsson, J.O.; Rantakokko, J. Zero-Velocity Detection-An Algorithm Evaluation. *IEEE Trans. Biomed. Eng.* **2010**, *57*, 2657–2666. [\[CrossRef\]](#) [\[PubMed\]](#)
- Liu, N.; Zhou, Z.; Zhu, L.; He, Y.; Huang, F. Fault Diagnosis of Unmanned Aerial Systems Using the Dempster-Shafer Evidence Theory. *Actuators* **2024**, *13*, 264. [\[CrossRef\]](#)
- Wang, X.; Zhang, P.; Zhang, Q.; Liu, J.; Wang, D.; Ye, H.; Li, T. Information fusion and feature selection for multi-source data utilizing Dempster-Shafer evidence theory and K-nearest neighbors. *Inf. Sci.* **2025**, *718*, 122408. [\[CrossRef\]](#)

18. An, Y.; Li, L.; Gao, H.; Luo, Z.; He, Y. Fusing multiple source data for foundation excavation risk assessment based on cloud model and Dempster-Shafer evidence theory. *Sci. Rep.* **2025**, *15*, 9894. [[CrossRef](#)] [[PubMed](#)]
19. Wang, K.; Wang, W.; Zhao, Y.; Yuan, B.; Xiang, Z. Multisensor fault diagnosis via Markov chain and Evidence theory. *Eng. Appl. Artif. Intell.* **2023**, *126*, 106851. [[CrossRef](#)]
20. Dempster, A.P. Upper and Lower Probabilities Induced by a Multivalued Mapping. *Ann. Math. Stat.* **1967**, *38*, 325–339. [[CrossRef](#)]
21. Shafer, G. *A Mathematical Theory of Evidence*; Princeton University Press: Princeton, NJ, USA, 1976.
22. 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](#)]
23. Kaess, M.; Johannsson, H.; Roberts, R.; Ila, V.; Leonard, J.J.; Dellaert, F. ISAM2: Incremental Smoothing and Mapping Using the Bayes Tree. *Int. J. Robot. Res.* **2012**, *31*, 216–235. [[CrossRef](#)]
24. Lyu, P.; Wang, B.; Lai, J.; Bai, S.; Liu, M.; Yu, W. A Factor Graph Optimization Method for High-Precision IMU-Based Navigation System. *IEEE Trans. Instrum. Meas.* **2023**, *72*, 9509712. [[CrossRef](#)]
25. Fan, G.; Wang, Q.; Yang, G.; Liu, P. RFG-TVIU: Robust Factor Graph for Tightly Coupled Vision/IMU/UWB Integration. *Front. Neurobotics* **2024**, *18*, 1343644. [[CrossRef](#)] [[PubMed](#)]
26. Karimi, M.M.; Mosavi, M.R. GPS/VIO integrated navigation system based on factor graph and fuzzy logic. *Sci. Rep.* **2024**, *14*, 30937. [[CrossRef](#)] [[PubMed](#)]
27. Wang, R.; Xu, C.; Li, R.; Duan, S.; Zhang, X. Cooperative Localization and Mapping Based on UWB-IMU Fusion Using Factor Graphs. *IEEE Sens. J.* **2024**, *24*, 14. [[CrossRef](#)]
28. Lyu, Y.; Hua, L.; Wu, J.; Liang, X.; Zhao, C. Robust Radar Inertial Odometry in Dynamic 3D Environments. *Drones* **2024**, *8*, 197. [[CrossRef](#)]
29. Luo, L.; Peng, F.; Dong, L. Improved Multi-Sensor Fusion Dynamic Odometry Based on Neural Networks. *Sensors* **2024**, *24*, 6193. [[CrossRef](#)] [[PubMed](#)]
30. Dong, W.; Lu, C.; Bao, L.; Li, W.; Shin, K.; Han, C. A Planar Multi-Inertial Navigation Strategy for Autonomous Systems for Signal-Variable Environments. *Sensors* **2024**, *24*, 1064. [[CrossRef](#)] [[PubMed](#)]
31. He, Q.; Yu, H.; Liang, D.; Yang, X. Enhancing Pure Inertial Navigation Accuracy through a Redundant High-Precision Accelerometer-Based Method Utilizing Neural Networks. *Sensors* **2024**, *24*, 2566. [[CrossRef](#)] [[PubMed](#)]
32. Sun, X.; Lai, J.; Lyu, P.; Liu, R.; Gao, W. A Fast Self-Calibration Method for Dual-Axis Rotational Inertial Navigation Systems Based on Invariant Errors. *Sensors* **2024**, *24*, 597. [[CrossRef](#)] [[PubMed](#)]
33. Wu, J.; Jiang, J.; Zhang, C.; Li, Y.; Yan, P.; Meng, X. A Novel Optimal Robust Adaptive Scheme for Accurate GNSS RTK/INS Tightly Coupled Integration in Urban Environments. *Remote Sens.* **2023**, *15*, 3725. [[CrossRef](#)]
34. Ge, J.; Yin, G.; Wang, Y.; Xu, D.; Wei, F. Rolling-Bearing Fault-Diagnosis Method Based on Multimeasurement Hybrid-Feature Evaluation. *Information* **2019**, *10*, 359. [[CrossRef](#)]
35. Hall, D.L.; Llinas, J. An Introduction to Multisensor Data Fusion. *Proc. IEEE* **1997**, *85*, 6–23. [[CrossRef](#)]
36. Huber, P.J.; Ronchetti, E.M. *Robust Statistics*, 2nd ed.; John Wiley & Sons: Hoboken, NJ, USA, 2009. [[CrossRef](#)]
37. Denceux, T. A k-Nearest Neighbor Classification Rule Based on Dempster-Shafer Theory. *IEEE Trans. Syst. Man. Cybern.* **1995**, *25*, 804–813. [[CrossRef](#)]
38. Smets, P.; Kennes, R. The Transferable Belief Model. *Artif. Intell.* **1994**, *66*, 191–234. [[CrossRef](#)]
39. Jao, C.-S.; Shkel, A.M. ZUPT-Aided INS Bypassing Stance Phase Detection by Using Foot-Instability-Based Adaptive Covariance. *IEEE Sens. J.* **2021**, *21*, 24338–24348. [[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.