remote sensing

MDPI

Article

Causal Subspace Detection and Reconstruction with Protected Local Deformation Modes in GNSS Time Series

Junpeng Shi [1,2,*], Junjie Chen [1], Weitao Yan [1], Mingyuan Wang [1,3], Haofei Ban [1], Sen Li [1] and Di Zhao [1]

[1] School of Geomatics and Aerospace Information, Henan Polytechnic University, Jiaozuo 454000, China
[2] Highway of Institute, Henan College of Transportation, Zhengzhou 451460, China
[3] The 2nd Geological Brigade of Jiangsu Geological Bureau, Changzhou 213300, China
* Correspondence: sjp@home.hpu.edu.cn

Highlights

What are the main findings?

- Joint integrates causal anomaly detection and reconstruction with explicit protection of local offsets and linear rates.
- Historical contamination control reduces mean fallback RMSE by 0.232–0.307 mm relative to NoGuard across three synthetic scenarios.

What are the implications of the main findings?

- Predefined motion protection supports interpretable GNSS quality control and displacement tracking.
- Calibration and transient response require assessment alongside final signal-amplitude retention.

Abstract

Preserving weak deformation while detecting observation faults remains challenging in Global Navigation Satellite System (GNSS) monitoring. We develop Joint, a causal subspace framework combining protected local offsets and linear rates, dual-timescale background learning, covariance whitening, historical contamination control, matched-subspace detection, and masked generalized least-squares reconstruction. Only observations available at the decision epoch are used, and missing observations are selected before whitening. Theoretical analysis establishes residual invariance and preservation of protected components under fixed-model, fixed-support, and full-rank reconstruction conditions. Evaluation uses three synthetic scenarios with 20 random seeds each and nine receivers at six sites. Historical contamination control increases F1 by 0.093–0.190 and reduces reconstruction root-mean-square error (RMSE) by 0.232–0.307 mm relative to its removal. Joint achieves event recall of 92.81–98.13%, exceeding HampelAxis with more false positives. Field injections transmit a 1 mm plateau as 0.999855–1.000000 mm, although maximum transition RMSE reaches 1.8276 mm. These results support explicit protection of prescribed motion components and suppression of historical contamination, while identifying calibration and transient response as limitations. Joint provides an interpretable workflow for GNSS quality control and displacement tracking, with potential applications in landslide, mining-subsidence, and structural monitoring.

Keywords: GNSS time series; deformation monitoring; subspace detection; causal estimation; anomaly detection; generalized least squares

Check for updates

Academic Editor: Antonio Miguel Ruiz Armenteros

Received: 3 October 2026
Revised: 13 November 2026
Accepted: 20 November 2026
Published: 25 November 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

Continuous Global Navigation Satellite System (GNSS) observations support surface and structural deformation monitoring. Applications include landslide displacement prediction [1], combined GNSS and interferometric synthetic aperture radar characterization of tectonic, hydrological, and anthropogenic motion [2], slow-slip detection [3], and offshore-platform monitoring [4]. In these applications, weak motion can be superimposed on multipath, environmental effects, reference instability, equipment changes, data gaps, and observation faults. Removing these disturbances is difficult because their amplitudes and time scales may overlap those of genuine motion. A method that removes large residuals without an explicit signal model can also suppress the deformation being monitored.

The distinction between signal and noise depends on both the functional model and the stochastic model. Santamaría-Gómez and Ray [5] showed that inferred GNSS noise characteristics change with model specification, while Gobron et al. [6] examined the influence of unmodeled offsets on low-frequency stochastic properties. Existing approaches include Bayesian piecewise-trend models [7], machine-learning discontinuity detection [8], and statistical assessment of wavelet features [9]. Spatial filtering and missing-data reconstruction have been addressed using weighted expectation-maximization principal component analysis [10], extended principal component analysis [11], and temporal–spatial Hankel matrix factorization [12]. These methods motivate careful treatment of incomplete observations and model dependence when interpreting filtered coordinates.

Learning-based approaches offer additional representations of nonlinear temporal and spatial dependence. Examples include comparisons of gradient boosting, long short-term memory (LSTM), and support-vector models [13], spatiotemporal graph networks for slow-slip extraction [14], and near-real-time neural denoising [15]. Multi-station consistency [16] and trend filtering [17] also provide ways to detect gradual motion. Related multivariate anomaly-detection methods use temporal convolution and attention [18], singular-value decomposition and subspace reconstruction [19], or LSTM variational autoencoders [20]. These approaches address different objectives and information settings; their results do not establish that arbitrary weak deformation will survive an adaptive detection and reconstruction pipeline.

Classical matched-subspace detection provides the projection and likelihood-ratio foundations for detecting signals in structured interference. Subspace constraints, SVD, covariance whitening, group-sparse estimation, and GLS likewise provide established methodological foundations. Building on these principles, this study formulates Joint as a sequential GNSS workflow that explicitly connects prescribed motion protection, historical contamination handling, current-epoch detection, and displacement reconstruction. The algorithmic contribution lies in the coupling between the protected background model, estimated historical contamination support, and reconstruction mask. Background learning uses available recent and phase-aligned historical blocks after removing the prescribed motion components. Within the current window, epoch-grouped sparse estimation identifies potentially contaminated earlier observations. Their selected support augments the nuisance model for current-epoch detection, and the corresponding observations are excluded from the subsequent reconstruction fit. Observed-coordinate selection precedes covariance whitening to accommodate incomplete records. These connections specify how the components interact within the sequential procedure. Motion preservation requires fixed covariance, model components, masks, and support, together with full-rank reconstruction; it does not automatically extend to adaptive reselection. The resulting framework therefore integrates background learning, anomaly identification, and signal reconstruction within a single causal estimation procedure, while providing an explicit mathematical guarantee that deformation components contained in the prescribed protection subspace do not con-

tribute to the anomaly residual. Its performance is evaluated in terms of anomaly-detection accuracy, preservation of weak deformation, reconstruction error, and uncertainty behavior under different noise, missing-data, and contamination conditions.

2. Deformation-Preserving Causal Dual-Timescale Detection and Reconstruction

Long-term GNSS displacement series contain deformation signals, temporally correlated background variations, stochastic noise, data gaps, and episodic observation faults. A central difficulty is that slowly varying deformation may occupy temporal scales similar to those of environmental and common-mode variations. Consequently, unconstrained smoothing or low-rank approximation may absorb part of the deformation signal together with the background.

To avoid ambiguity between weak deformation and structured disturbance, the proposed method explicitly defines a deformation subspace that is excluded from anomaly testing. Figure 1 illustrates the resulting causal dual-timescale representation of the GNSS time series. In this representation, neighboring observations within the current local window are arranged along the circumferential direction, while historical reference windows are arranged along the axial direction. The highlighted cells indicate the target and reference observations used for background learning and anomaly assessment. Normal background variability is estimated exclusively from past observations at two complementary temporal scales, and anomalies are detected in the residual subspace after removing the protected deformation, the learned background, and previously identified historical contamination. Observations classified as admissible are then used for structural displacement reconstruction. Since only information available up to epoch is used, the entire procedure is strictly causal.

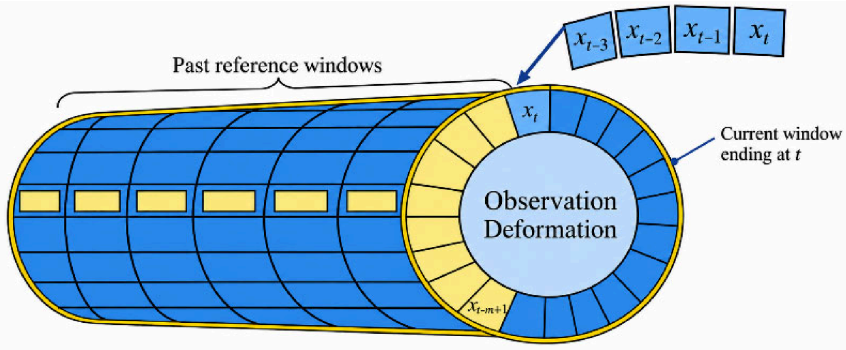

Figure 1. Schematic illustration of the deformation-preserving causal dual-timescale framework. The circumferential direction represents neighboring observations within the current analysis window, and the axial direction represents historical reference windows used for background learning. The highlighted cells indicate the target and reference observations involved in detection. A deformation subspace is explicitly protected from anomaly testing, while normal background variability is learned from past observations only. Future observations do not participate in the decision at the current epoch, thereby ensuring causality.

2.1. Observation Model and Deformation-Preserving Subspace

Let $\mathcal{W}_t = \{\tau_1, \ldots, \tau_m\}$, $\tau_m = t$ define the causal window of m epochs up to the current epoch t. Stack the three coordinate components epoch-wise into $y_t \in R^{3m}$. The observation model is

$$\mathbf{y}_t = \boldsymbol{\mu}_t + \mathbf{F}_t\boldsymbol{\beta}_t + \mathbf{U}_t\boldsymbol{\alpha}_t + \mathbf{E}_t\mathbf{a}_t + \mathbf{e}_t + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{C}_t) \tag{1}$$

where μ_t is a conditional deterministic term estimated from training data, $\mathbf{F}_t\boldsymbol{\beta}_t$ is the protected deformation component, $\mathbf{U}_t\boldsymbol{\alpha}_t$ is the recurrent background, and $\mathbf{E}_t\mathbf{a}_t$ represents

sparse observation anomalies. The discrepancy $\mathbf{e}_t$ collects deterministic variation outside the specified subspaces, and ε_t denotes correlated stochastic noise.

External explanatory variables may contribute to $\boldsymbol{\mu}_t$ when their relationship to displacement is independently justified. No external forcing model is used in the experiments, so this term is set to zero. The reconstruction target includes the protected component and the modeled background; it is not a uniquely identified geophysical deformation component.

Within a short window, a slowly varying displacement can be approximated by an offset and a linear rate. The protection basis is

$$\mathbf{F}_t = \begin{bmatrix} 1 & (\tau_1 - \tau_0)/h \\ \vdots & \vdots \\ 1 & (\tau_m - \tau_0)/h \end{bmatrix} \otimes \mathbf{I}_3, \quad \tau_0 = \frac{1}{m}\sum_{i=1}^{m} \tau_i \tag{2}$$

where h is the characteristic window duration and $\otimes$ denotes the Kronecker product. The time origin is the mean epoch within the window.

Within a local window, slowly varying displacement is approximated by an offset and a linear rate in each coordinate. Equation (2) defines the corresponding six-dimensional protection space when at least two distinct epochs are available. Time is centered at the mean window epoch and scaled by the characteristic window duration to improve numerical conditioning without changing the column space. The preservation property applies to motion within this space under fixed-model and observation-selection conditions. Curved trajectories, abrupt steps, and changes in rate are generally outside this space. Window duration therefore involves a trade-off between local approximation accuracy and sample support. Higher-order protection can better represent curved motion but consumes additional residual degrees of freedom and does not necessarily improve reconstruction accuracy.

2.2. Causal Dual-Timescale Background Learning

Two historical libraries provide complementary support. The recent library $\mathcal{L}_t$ contains nearby past windows; the phase-aligned library $\mathcal{P}_t$ contains windows from previous cycles at approximately the same phase. Both libraries contain only complete reference blocks.

For historical windows represented on the current relative-time grid, removal of the protected component gives

$$\mathbf{x}_{t,\ell} = \mathbf{P}_{F,t}^{\perp}(\mathbf{y}_{t-\ell} - \boldsymbol{\mu}_{t-\ell}), \quad \mathbf{X}_t = \left[\mathbf{x}_{t,\ell}\right]_{\ell \in \mathcal{L}_t \cup \mathcal{P}_t}, \quad \mathbf{P}_{F,t}^{\perp} = \mathbf{I} - \mathbf{F}_t\mathbf{F}_t^{\dagger} \tag{3}$$

with $\max\mathcal{W}_{t-\ell} < \min\mathcal{W}_t$. The symbol $(\cdot)^{\dagger}$ denotes the Moore–Penrose pseudoinverse. Each historical block has the same coordinate ordering and aligned relative epochs as the current full window. Matching uses actual timestamps. Blocks outside the matching tolerance or containing missing coordinates are excluded rather than treated as zero-filled or interpolated observations.

2.3. Noise Covariance and Background Basis

A white-noise-plus-exponential covariance approximates correlation over the finite analysis window:

$$\mathbf{C}_t^{(ik)} = \left[(1-\nu)\exp\left(-\frac{|\tau_i - \tau_k|}{\ell_c}\right) + \nu\delta_{ik}\right]\boldsymbol{\Sigma}, \quad 0 < \nu \leq 1 \tag{4}$$

where $\boldsymbol{\Sigma}$ is the 3×3 cross-component covariance matrix, ℓ_c is the correlation time, and δ_{ik} is the Kronecker delta. Positive definiteness is required for whitening.

This finite-window covariance is a working model. It does not describe the full range of long-term GNSS power-law noise or remove the need to assess distribution shifts.

Let $\mathbf{C}_t = \mathbf{L}_t\mathbf{L}_t^T$ be a Cholesky factorization. The historical matrix is whitened and projected away from the whitened protection space before weighted singular-value decomposition (SVD):

$$\begin{aligned}\mathbf{Q}_{F,t} &= \mathrm{orth}\left(\mathbf{L}_t^{-1}\mathbf{F}_t\right)\\ \mathbf{Z}_t &= \left(\mathbf{I} - \mathbf{Q}_{F,t}\mathbf{Q}_{F,t}^T\right)\mathbf{L}_t^{-1}\mathbf{X}_t\mathbf{\Omega}_t^{1/2}\\ \mathbf{Z}_t &= \mathbf{V}_t\mathbf{S}_t\mathbf{H}_t^T, \quad \mathbf{U}_t = \mathbf{L}_t\mathbf{V}_t(:, 1:r)\end{aligned} \tag{5}$$

Here, $\mathbf{\Omega}_t$ contains historical-block weights, and r is the retained background rank. The weights combine the two library weights with robust down-weighting of unusually large historical columns.

The projection prevents the retained background basis from overlapping the protected space in the full-window whitened metric. The historical-block weights combine the two library weights with robust down-weighting of unusually large historical columns, and denote the effective retained background rank. Projection prevents the background basis from overlapping the protected space in the full-window whitened metric. The requested maximum rank controls model complexity, whereas the effective rank is additionally limited by the numerical rank of the historical matrix and residual-dimension constraints. Increasing the rank can represent additional recurrent variation but can also absorb directions informative for anomaly detection and reduce the degrees of freedom available for estimation. The requested maximum rank, window duration, and weighting parameters are fixed before testing. When fewer than five complete historical blocks are available, the background basis is omitted and this fallback is recorded separately. Sensitivity analyses over requested ranks zero through four show that additional background dimensions do not consistently improve detection or reconstruction. The default upper limit of two should therefore be interpreted as an evaluated configuration, not an established optimum across receivers or noise conditions.

2.4. Missing Observations and Residual Construction

Let $\mathbf{M}_t$ select available coordinate observations. Missing-data selection is applied to the covariance before whitening:

$$\mathbf{C}_t^o = \mathbf{M}_t\mathbf{C}_t\mathbf{M}_t^T, \quad \mathbf{W}_t\mathbf{C}_t^o\mathbf{W}_t^T = \mathbf{I} \tag{6}$$

The initial nuisance space, orthogonal residual projector, and residual are

$$\begin{aligned}\mathbf{Q}_{0,t} &= \mathrm{orth}(\mathbf{W}_t\mathbf{M}_t[\mathbf{F}_t, \mathbf{U}_t])\\ \mathbf{R}_{0,t} &= \mathbf{I} - \mathbf{Q}_{0,t}\mathbf{Q}_{0,t}^T\\ \mathbf{v}_{0,t} &= \mathbf{R}_{0,t}\mathbf{W}_t\mathbf{M}_t(\mathbf{y}_t - \boldsymbol{\mu}_t)\end{aligned} \tag{7}$$

The orth operator returns an orthonormal basis for the numerical column space. Under correlated noise, whitening the full vector and then discarding coordinates generally differs from whitening the selected covariance. Using the latter ordering keeps the data, design, and covariance on the same observation domain.

2.5. Historical Contamination Control

A fault earlier in the current window can influence several subsequent decisions. An epoch-grouped sparse fit identifies potential historical interference:

$$\hat{\mathbf{o}}_t = \mathrm{argmin}_{\mathbf{o}}\left\{\frac{1}{2}\left\|\mathbf{v}_{0,t} - \mathbf{R}_{0,t}\mathbf{W}_t\mathbf{M}_t\mathbf{D}_t\mathbf{o}\right\|_2^2 + \lambda\sum_{g\in\mathcal{G}_t}\left\|\mathbf{o}_g\right\|_2\right\} \tag{8}$$

The group-sparse problem is solved independently in each window using the fast iterative shrinkage-thresholding algorithm (FISTA). Both the coefficient and extrapolated

vectors are initialized to zero, with an initial momentum parameter of one and no warm start. The dictionary incorporates coordinate scales estimated from the training data, observation-domain whitening, and projection away from the protected and background spaces. Each iteration applies a gradient step followed by epoch-group soft thresholding. The step size is the reciprocal of the squared spectral norm of this standardized, projected dictionary.

The historical support is then

$$\mathcal{S}_t = \left\{ g < t : \left\| \hat{\mathbf{o}}_g \right\|_2 > \gamma \right\} \tag{9}$$

where $\mathbf{D}_t$ accounts for coordinate-dependent scales, λ controls sparsity, and γ is the support threshold. Each group contains the three coordinate components associated with one epoch.

The sparsity penalty λ is fixed at 2.8, and the historical-support threshold γ is fixed at 0.8 in standardized coefficient units. These are implementation constants used throughout the reported experiments, rather than parameters established as statistically optimal. The support threshold is distinct from the empirically calibrated anomaly-detection threshold. Iteration stops when the Euclidean norm of the coefficient change is less than 10^{-6} times one plus the previous coefficient norm, or when 60 iterations have been completed. If the tolerance is not met within this limit, the last iterate is retained; reaching the limit is not interpreted as convergence.

Groups are defined by epoch using only observed coordinate components. A complete epoch contains three coefficients, a partially observed epoch contains one or two, and a completely missing epoch contributes no group. Missing components are neither imputed nor included as zero-valued observations. No group-size normalization is applied, so equal selection behavior across different missingness patterns is not guaranteed. A historical epoch is selected when its final standardized coefficient-group norm exceeds 0.8; all observed coordinates of that epoch then enter the historical support. The current epoch participates in the sparse fit but is excluded from this support, preventing its own anomaly template from being explicitly absorbed into the historical nuisance space before testing.

The current epoch participates in the sparse fit but is excluded from the historical support $\mathcal{S}_t$. This prevents its own anomaly template from being explicitly absorbed into the historical nuisance space before testing. The support remains data dependent, so this construction does not make the operational score an exact fixed-design test.

Augmenting the nuisance space with the selected historical columns gives

$$\begin{aligned} \mathbf{Q}_t &= \operatorname{orth}\left(\mathbf{W}_t \mathbf{M}_t [\mathbf{F}_t, \mathbf{U}_t, \mathbf{D}_{\mathcal{S}_t}]\right) \\ \mathbf{R}_t &= \mathbf{I} - \mathbf{Q}_t \mathbf{Q}_t^T \\ \mathbf{v}_t &= \mathbf{R}_t \mathbf{W}_t \mathbf{M}_t (\mathbf{y}_t - \boldsymbol{\mu}_t) \end{aligned} \tag{10}$$

Because the nuisance space contains the selected and whitened protection basis, $\mathbf{R}_t \mathbf{W}_t \mathbf{M}_t \mathbf{F}_t = \mathbf{0}$; hence $\mathbf{v}_t(\mathbf{y}_t + \mathbf{F}_t \mathbf{b}) = \mathbf{v}_t(\mathbf{y}_t), \forall \mathbf{b}$. This algebraic identity holds when the covariance, selection, background, and historical support are held fixed.

The identity excludes additions in the column space of $\mathbf{F}_t$ from the residual. It neither distinguishes physical motion from a fault in the same space nor guarantees invariance after retraining, remasking, or reselection of the nuisance model. The protection space creates a deliberate blind spot for faults with the same local signature.

2.6. Matched-Subspace Detection and Calibration

For the current epoch j, let $\mathbf{E}_{j,t}$ embed a possible three-component anomaly at that epoch j into the full window. Applying the same transformations to this template gives

$$\mathbf{A}_t = \mathbf{R}_t \mathbf{W}_t \mathbf{M}_t \mathbf{E}_{j,t}, \quad q_t = \operatorname{rank}(\mathbf{A}_t) \tag{11}$$

Following the matched-subspace formulation [21], the projection is $\mathbf{P}_{A,t} = \mathbf{A}_t(\mathbf{A}_t^T\mathbf{A}_t)^{\dagger}\mathbf{A}_t^T$ and the test statistic is $T_t = \mathbf{v}_t^T\mathbf{P}_{A,t}\mathbf{v}_t$. We use the rank-adjusted reference score $s_t = -\log_{10}\left[1 - F_{\chi^2_{q_t}}(T_t)\right]$. A disjoint calibration interval sets the threshold $\xi_{\text{cal}} = Q_{1-\alpha_0}\{s_\tau : \tau \in \mathcal{T}_{\text{cal}}\}$, with the decision rule

$$s_t > \xi_{\text{cal}} \Rightarrow H_1 \tag{12}$$

where H_1 denotes a candidate anomalous epoch. A candidate is a model-based flag rather than an independently verified observation fault.

For a fixed nuisance model, correct whitening, Gaussian noise, and zero deterministic discrepancy, $T_t|H_0 \sim \chi^2_{q_t}$ is the reference distribution. Estimated covariance and data-dependent background and sparse-support selection prevent an automatic exact-probability interpretation of the reference score. The operational threshold is therefore determined empirically from the 99th percentile of a separate calibration interval and remains fixed during the primary test evaluation. Its transfer to subsequent observations depends on the stability of the score distribution. Changes in noise amplitude, temporal correlation, missingness, or background variation can alter this distribution and invalidate the nominal false-positive target. We assess these effects through paired distribution-shift experiments and separately examine recalibration using preceding observations. Neither the fixed threshold nor the evaluated updating procedures provide a general guarantee of 1% test false-positive control. If $q_t = 0$ or too few observations remain, no anomaly score is reported. Rank and conditioning checks distinguish an unscorable epoch from an epoch classified as normal.

2.7. Masked Generalized Least-Squares Reconstruction

Detection and reconstruction have different targets. Let $\mathbf{G}_t$ retain available observations that are neither flagged nor included in the selected historical contamination support.

The reconstruction design is $\mathbf{B}_t = [\mathbf{F}_t, \mathbf{U}_t]$, with selected covariance and design $\mathbf{C}_{G,t} = \mathbf{G}_t\mathbf{C}_t\mathbf{G}_t^T$, $\mathbf{B}_{G,t} = \mathbf{G}_t\mathbf{B}_t$. Generalized least squares (GLS) estimates the coefficients as

$$\hat{\boldsymbol{\theta}}_t = \left(\mathbf{B}_{G,t}^T\mathbf{C}_{G,t}^{-1}\mathbf{B}_{G,t}\right)^{-1}\mathbf{B}_{G,t}^T\mathbf{C}_{G,t}^{-1}\mathbf{G}_t(\mathbf{y}_t - \boldsymbol{\mu}_t) \tag{13}$$

The current structural coordinate estimate is

$$\hat{\mathbf{y}}_{t,\text{str}} = \boldsymbol{\mu}_{t,c} + \mathbf{B}_{t,c}\hat{\boldsymbol{\theta}}_t \tag{14}$$

with conditional plug-in covariance

$$\hat{\mathbf{V}}_{t,\text{str}} = \mathbf{B}_{t,c}\left(\mathbf{B}_{G,t}^T\mathbf{C}_{G,t}^{-1}\mathbf{B}_{G,t}\right)^{-1}\mathbf{B}_{t,c}^T \tag{15}$$

where $\mathbf{B}_{t,c}$ contains the three current-epoch design rows. The estimate in Equation (14) retains modeled background as well as protected deformation.

The displayed inverses specify the estimator; numerical implementation uses factorizations or linear solves. Reconstruction is unavailable when the retained design lacks full column rank, residual degrees of freedom are insufficient, or conditioning fails. A raw-observation fallback is used only for the synthetic error metric described below and is not counted as a successful reconstruction.

For a fixed design, mask, and covariance, adding a design-space signal shifts the full-rank GLS estimate by the same signal. This follows because the GLS operator is a left inverse of the retained design. The result is conditional on preserving the design and observation support. The covariance in Equation (15) treats the estimated covariance parameters, learned background, selected historical support, and reconstruction mask as fixed. It does not propagate uncertainty from estimating or selecting these quantities, or account for bias

caused by model discrepancy. The resulting intervals are conditional precision summaries whose empirical coverage must be assessed separately for the complete adaptive procedure.

The overall workflow of the proposed deformation-preserving causal dual-timescale framework is illustrated in Figure 2. The method consists of four stages: background modelling, anomaly detection, displacement reconstruction, and result output. GNSS observations are first processed under deformation protection and dual-timescale background learning, then screened for outliers through residual-based matched-subspace detection, and finally reconstructed by masked generalized least squares to obtain structural displacement and uncertainty.

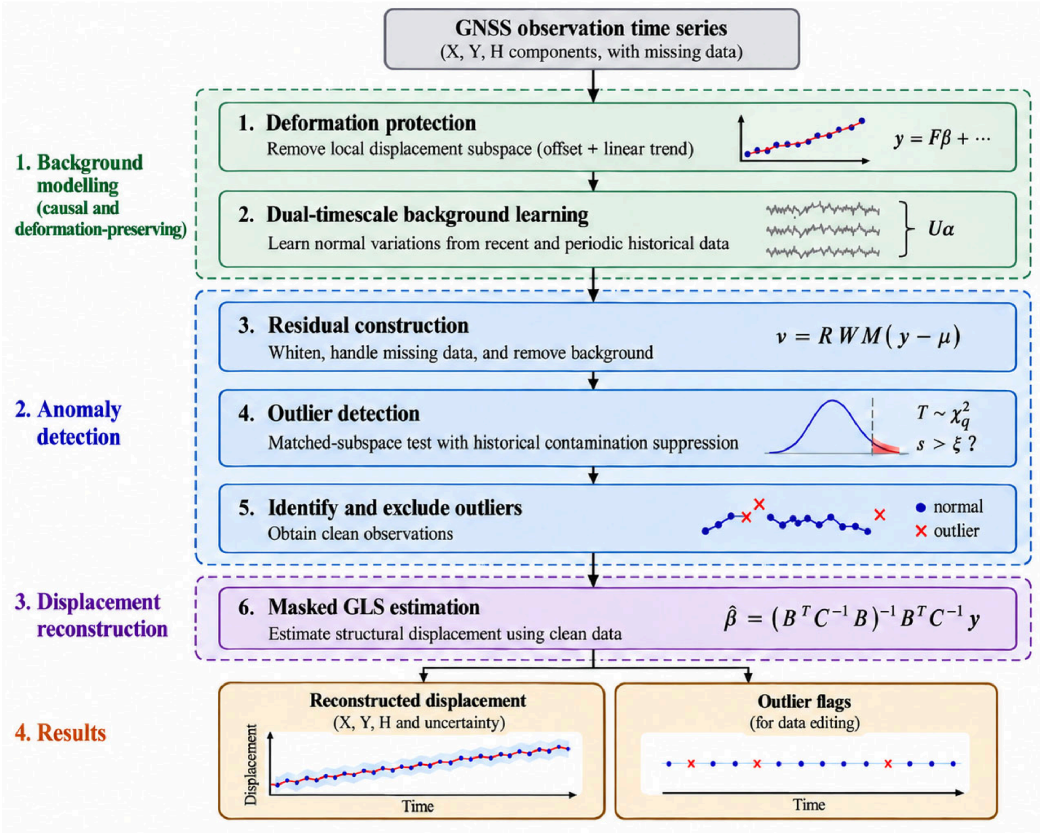

Figure 2. Overall workflow of the proposed deformation-preserving causal dual-timescale framework. The method consists of four stages: background modeling, anomaly detection, displacement reconstruction, and result output. GNSS observations are first processed under deformation protection and dual-timescale background learning, then screened for outliers through residual-based matched-subspace detection, and finally reconstructed by masked generalized least squares to obtain structural displacement and uncertainty.

2.8. Implementation and Controlled Comparisons

Joint uses linear protection and a background rank of at most two. The default window is 12 h for hourly data. S03 uses a 48 h window because its 6 h sampling interval provides only two scheduled epochs within a 12 h window, below the minimum requirement of five window-grid epochs. Under the implemented open-left, closed-right convention, the selected windows contain twelve scheduled epochs for hourly records and eight for S03. This provides a sample-support rationale for extending the S03 window, although it does not establish 48 h as optimal. Recent reference endpoints are spaced 6 h apart over the preceding 7 d, while phase-aligned references use the preceding fourteen 24 h cycles. Both libraries have an initial weight of 0.5, and timestamp matching allows a tolerance of 0.10 h. Incomplete historical blocks are excluded; at least five complete blocks are required to

construct a learned background. When this requirement is unmet, a background-free fallback is used, subject to the remaining support and rank conditions. Sensitivity analyses compare window durations of 6, 12, 24, and 48 h at hourly sampling with linear protection; constant, linear, and quadratic protection at the default 12 h setting; and sampling intervals of 1, 2, 3, and 6 h at a fixed 48 h window. Each comparison uses twenty paired realizations per scenario, with configuration-specific calibration and thresholds fixed during testing. The additional analyses report reconstruction accuracy, detection performance, output availability, and historical-background diagnostics. These analyses assess configuration dependence rather than establish a universally optimal setting. All calculations were performed in MATLAB R2024a.

The training interval estimates cross-component covariance and the correlation time. Fixed parameters are a white-noise mixing proportion of 0.20, robust truncation coefficient of 2.5, group-sparsity penalty of 2.8, maximum of 60 iterations, and support threshold of 0.8. Reconstruction requires a full-rank design, at least two residual degrees of freedom, and a normal-matrix condition number no greater than 10^{11}. These values define the evaluated configuration; they do not establish optimal settings for other monitoring environments.

Table 1 defines eight single-module ablations and a simple baseline. NoMask reuses Joint detection scores, thresholds, scoreability, and flags, changing only the reconstruction mask. Consequently, all detection metrics must match Joint on the same epochs. This invariant holds for all 60 paired cases in the result records. HampelAxis is evaluated as a separate component-wise baseline combining Hampel screening with local linear GLS reconstruction. It does not include Joint's learned multivariate background or epoch-grouped historical contamination control. This comparison evaluates whether the complete combination improves performance relative to a simpler procedure, while the single-module ablations assess the contribution of individual components within Joint.

Table 1. Controlled variants and comparison baseline.

Method	Change Relative to Joint
Joint	Complete joint detector with masked GLS
NoBackground	Set background rank to zero
NoGuard	Disable historical group-sparse contamination control
NoSlope	Protect constants only; remove linear-rate columns
WhiteNoise	Use an identity temporal kernel; retain cross-component covariance
LocalOnly	Use recent historical support only
PeriodicOnly	Use phase-aligned historical support only
NoRobust	Disable robust amplitude down-weighting; retain library weights
NoMask	Keep Joint detection; reconstruct from all originally valid observations
HampelAxis	Component-wise Hampel detection and local linear GLS; no learned joint background or historical sparse guard

2.9. Synthetic Design and Statistical Analysis

Three scenarios represent single-axis faults, coupled three-axis faults, and stress conditions. Each hourly series lasts 60 d: 20 d for training, 15 d for calibration, and 25 d for testing. Twenty random seeds per scenario, identified as 20261201–20261220, yield 60 cases and 600 method-level records. These numbers are seed identifiers rather than acquisition dates. They provide repeated realizations of the same generators, not independent physical monitoring environments.

The comprehensive observation combines real deformation, periodic background, time- and cross-component correlated noise, independent white noise, and injected observation faults. The reconstruction reference includes deformation and background. Matching single-axis and coupled scenarios share signal, noise implementation, event time, and fault-vector specifications. Pressure conditions can also introduce noise scale drift, missing observations, and timestamp jitter. These disturbances and true deformations are not marked as observed faults. The complete generation rules, event schedule, and implementation details are listed in Table 2.

Table 2. Synthetic data generation parameters.

Component	Generation Settings
Linear trend	Rates of 0.0012, −0.0007, and 0.0004 mm/h for X, Y, and H, respectively.
Smooth nonlinear deformation	Logistic transition centered at day 44, with a time-scale parameter of 36 h and amplitudes of 2.0, −1.0, and 1.5 mm.
Acceleration segment	Quadratic growth begins at day 48 with zero initial displacement and rate for this component. Its contributions reach 0.40, −0.20, and 0.30 mm at day 53 and continue quadratically without clipping.
Genuine displacement step	A permanent 1 mm step in X at day 52; treated as genuine deformation rather than an observation fault.
Lagged periodic background	Sinusoidal forcing with periods of 24 and 216 h and amplitudes of 8 and 2. A first-order lag uses a 5 h time constant and an initial state of zero. Component gains are 0.14, −0.09, and 0.19.
Weekly background	A zero-phase sinusoid with a 168 h period and component amplitudes of 0.25, 0.12, and −0.20 mm.
Colored noise	Zero-mean Gaussian noise with a 1 h temporal correlation of 0.65, adjusted exponentially for actual time separations. Component variances are 0.1225, 0.2025, and 0.4900 mm^2; XY, XH, and YH covariances are 0.025, 0.020, and −0.020 mm^2, respectively.
White noise	Independent Gaussian noise with a standard deviation of 0.15 mm in each component.
Injected faults	Three isolated training faults with uniformly distributed vector norms of 8–12 mm; no calibration faults. Testing contains eleven isolated faults and five blocks of 3, 5, 8, 4, and 6 epochs. Test-event norms are uniformly distributed over 3–10 mm, with equiprobable signs and a constant vector within each block.
Fault directions	Single-axis and stress faults cycle through X, Y, and H in the specified event order. Coupled faults use the normalized direction proportional to 1, −0.8, and 1.15, preserving the corresponding single-axis fault norm.
Stress noise-scale drift	Multipliers increase linearly from 1 to 1.6, 1.25, and 2.0 between days 39 and 45, then remain constant. Scaling applies to the combined colored and white noise.
Stress missing observations	Independent missingness probabilities of 2.5%, 3.0%, and 4.0% throughout training, calibration, and testing. An additional H-component gap occupies one-based epochs 1270–1281.
Stress timestamp jitter	Deterministic sinusoidal jitter bounded by five minutes, applied every 37 epochs starting at epoch 31. Perturbed times are used in signal generation and noise propagation.

The empirical 99% quantile on common valid calibration epochs determines each threshold, which remains fixed during the primary test evaluation. A separate analysis evaluates updates every 24 or 168 h, starting one update interval after the test begins. Each update recomputes the empirical 99% threshold using that method's complete, scoreable

epochs from the preceding 360 h, excluding training data and the current epoch. At least 100 eligible scores are required; otherwise, the previous threshold is retained. Updates use no fault labels and do not exclude previously flagged observations. Detection is evaluated on the original common test domain, and reconstruction is not recomputed using the updated flags. Epoch precision, recall, F1, and false-positive rate use complete three-component test epochs scoreable by all compared methods. Evaluable events contain a valid, scoreable faulty epoch, detected or not. Event recall is the fraction of evaluable events with a detected faulty epoch. Reconstruction truth includes the generated background and deformation. Additional metrics include RMSE on each method's successfully reconstructed coordinates and on their intersection across all ten methods, with a separate Joint–HampelAxis intersection. Failure rate is the proportion of originally observed coordinates without a finite reconstruction; coordinate counts and domain shares accompany these conditional comparisons. Successful-coordinate errors are stratified by injected-fault status and current detector flags. Flags do not imply removal for NoMask, and these strata do not capture all historical exclusions. Raw-observation errors at reconstruction failures are reported separately. Metrics are calculated per realization and summarized across twenty seeds, with empty strata treated as unavailable rather than assigned zero error. Means and sample standard deviations summarize the twenty seeds. Ablation differences are paired by seed and defined as variant minus Joint. Nominal 95% percentile bootstrap intervals use 2000 resamples of the twenty seed pairs; the aggregation bootstrap seed is 20261003. No multiple-comparison adjustment is applied, so these intervals provide exploratory module comparisons. Consecutive epochs are not treated as independent replicates. The error bars describing seed variability and the intervals describing paired effects have different interpretations.

2.10. Field Data and Paired Signal Injection

The field evaluation uses nine receivers at six sites (Table 3), comprising 39,031 grid epochs and 104,783 valid coordinate values; 82,098 valid values fall in the test intervals. The first 30 d of each receiver record are used for training, the next 15 d for calibration, and the remainder for testing. S03 retains a 6 h cadence and a 48 h current window; the other receivers use a 1 h grid and a 12 h window. The original X, Y, and H labels are retained because their mapping to north, east, and up has not been established.

The field inputs are exported coordinates, not raw satellite observations. Components are sorted, duplicate timestamps with identical coordinates are collapsed, and X/Y/H are outer-joined on exact timestamps. For this retrospective evaluation, campaigns are segmented at gaps >168 h or adjacent three-coordinate jumps >1000 mm. The latest-ending segment spanning at least 90 d is selected and capped at its final 270 d. This displacement-dependent selection can exclude physical or instrumental discontinuities and conditions the scope of inference. A 6 h grid is used when the median native interval is ≥3 h; otherwise, a 1 h grid is used. The hourly :00 or :30 phase is selected by its frequency in the first 30 d, with ties assigned to :00. Exact reindexing applies no averaging, interpolation or nearest-neighbor reassignment. Both S08 and S09 use :30 phase: 6184 of 12,543 native epochs and 1960 of 4962 native epochs are retained, excluding 6359 and 3002 off-grid epochs. Each axis is expressed in millimeters relative to its first valid grid value.

Two structural comparisons supplement Joint and HampelAxis. SplitEpoch and SplitAxis use component-wise detection with a calibrated maximum-axis score and identical candidate decisions. SplitEpoch discards an entire epoch when any component is flagged, whereas SplitAxis excludes only flagged components. These variants also change background capacity, noise parameterization, and grouping relative to Joint. They are structural comparisons rather than isolated ablations.

Table 3. Selected field records and observation availability.

Receiver	Site	Start to End Date	Cadence (h)	Grid Epochs	Valid Values (%)
S01	Qianfengshan	8 December 2021 16 April 2022	1	3099	89.1
S02	Qianfengshan	27 July 2021 29 November 2021	1	3006	89.3
S03	Buyiqiao coal-mine goaf	27 March 2026 17 September 2026	6	697	97.0
S04	Qifengshan	16 October 2024 13 July 2025	1	6481	79.7
S05	Shiyan Village	8 February 2025 13 July 2025	1	3728	91.1
S06	Shunguoshan	16 October 2024 13 July 2025	1	6481	89.7
S07	Shunguoshan	18 October 2024 13 July 2025	1	6448	97.5
S08	Huaqiangu	16 October 2024 13 July 2025	1	6481	95.4
S09	Huaqiangu	2 June 2026 18 September 2026	1	2610	75.1

The initial experiment simultaneously adds a linear ramp to X, a quadratic ramp to Y, and a permanent step to H. To remove this fixed pairing between deformation pattern and coordinate component, an additional crossed experiment injects each pattern separately into each of X, Y, and H at all nine receivers. Each run modifies only one component, giving 81 receiver–pattern–component conditions per method. All additions have a terminal amplitude of 1 mm and begin on test day 7. The ramps reach 1 mm on day 14 and subsequently remain constant, whereas the step reaches its full amplitude immediately. Original and injected records use identical training data and frozen thresholds. The difference between paired estimates measures the transmitted signal. Transition RMSE is evaluated over test days 7–21 and plateau displacement over days 21–35; paired response coverage is reported alongside error. Baseline and injected runs retain identical training and calibration data, original missingness patterns, and frozen method-specific thresholds. The paired response is the difference between their reconstructed series. Transition RMSE compares this response with the known addition over test days 7–21, while plateau response is summarized by its median over days 21–35. Evaluation requires an originally observed coordinate and finite estimates in both runs. Results are presented by receiver, deformation pattern, and coordinate component (Figure 3). The crossed design enables comparisons without fixing one pattern to one component, while retaining interactions with the observed noise and adaptive processing. These experiments measure responses to known additions to field records, rather than absolute reconstruction accuracy against independently measured displacement.

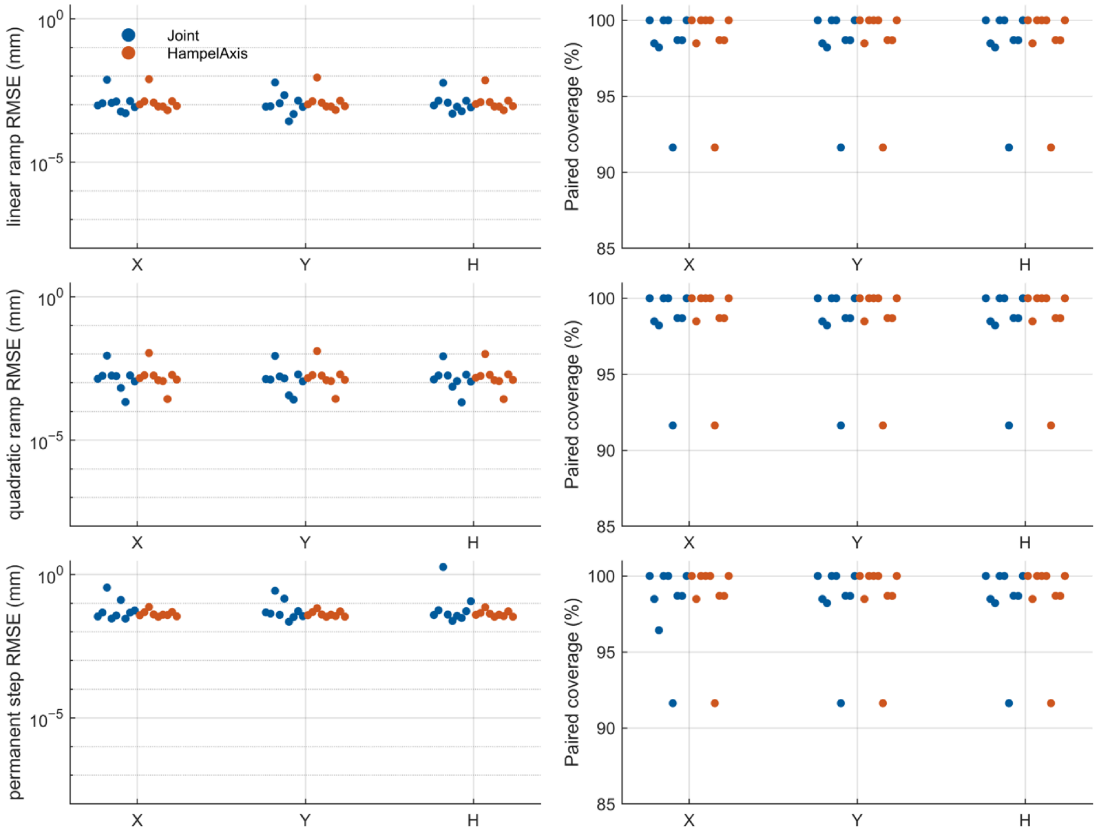

Figure 3. Separate injections of linear ramps, quadratic ramps and permanent steps into X, Y and H. Each marker is a receiver–method condition; horizontal jitter separates overlapping values. Left: transition RMSE on a logarithmic scale. Right: paired transition coverage. The 81 conditions per method are repeated observations on nine receivers, not independent site replicates.

3. Results

3.1. Synthetic Detection and Reconstruction

Table 4 and Figure 4 show that HampelAxis outperforms Joint in epoch-wise F1 and the main reconstruction RMSE for all three scenarios. Joint has F1 scores of 0.568, 0.535, and 0.359 and RMSEs of 0.928, 1.008, and 1.233 mm. The corresponding HampelAxis values are 0.845, 0.798, and 0.758 for F1 and 0.560, 0.704, and 0.833 mm for RMSE. The results do not support a general performance advantage for the complete Joint configuration. Joint mean reconstruction coverage is 98.79%, 99.40%, and 97.50%, respectively. The empirical coverage of Joint's nominal 95% plug-in intervals is 93.10%, 90.46%, and 79.30%, showing undercoverage, especially under stress. In the stress scenario, doubling the reported standard errors increases mean coverage to 96.70%, while increasing mean full interval width from 1.989 to 3.978 mm (Table 4; Figure 5). This exploratory analysis demonstrates a coverage–width trade-off, but does not establish a calibrated multiplier for new data.

Figure 5 evaluates standard-error factors 1, 1.25, 1.5 and 2 on each method's successful reconstruction domain. Joint stress coverage increases from 79.30% to 87.05%, 92.15% and 96.70%; mean width increases from 1.989 to 2.486, 2.983 and 3.978 mm. Inflation improves average synthetic coverage at a clear width cost.

Table 4. Synthetic F1 and fallback RMSE across twenty seeds.

Method	Single-Axis		Coupled		Stress	
	F1	RMSE (mm)	F1	RMSE (mm)	F1	RMSE (mm)
Joint	0.568 ± 0.056	0.928 ± 0.112	0.535 ± 0.052	1.008 ± 0.108	0.359 ± 0.042	1.233 ± 0.161
NoBack-ground	0.594 ± 0.035	0.939 ± 0.113	0.596 ± 0.048	0.946 ± 0.102	0.377 ± 0.048	1.192 ± 0.116
NoGuard	0.378 ± 0.033	1.235 ± 0.221	0.361 ± 0.041	1.240 ± 0.236	0.266 ± 0.040	1.525 ± 0.189
NoSlope	0.695 ± 0.047	0.686 ± 0.073	0.612 ± 0.052	0.780 ± 0.115	0.420 ± 0.053	0.866 ± 0.092
WhiteNoise	0.563 ± 0.042	0.974 ± 0.102	0.528 ± 0.047	1.043 ± 0.101	0.342 ± 0.055	1.272 ± 0.119
LocalOnly	0.587 ± 0.037	0.927 ± 0.104	0.546 ± 0.045	1.003 ± 0.098	0.368 ± 0.053	1.200 ± 0.106
PeriodicOnly	0.566 ± 0.041	0.923 ± 0.094	0.533 ± 0.046	1.005 ± 0.096	0.368 ± 0.043	1.198 ± 0.126
NoRobust	0.573 ± 0.040	0.926 ± 0.108	0.544 ± 0.047	0.985 ± 0.113	0.361 ± 0.043	1.217 ± 0.149
NoMask	0.568 ± 0.056	0.974 ± 0.093	0.535 ± 0.052	1.011 ± 0.079	0.359 ± 0.042	1.143 ± 0.088
HampelAxis	0.845 ± 0.071	0.560 ± 0.032	0.798 ± 0.101	0.704 ± 0.064	0.758 ± 0.082	0.833 ± 0.038

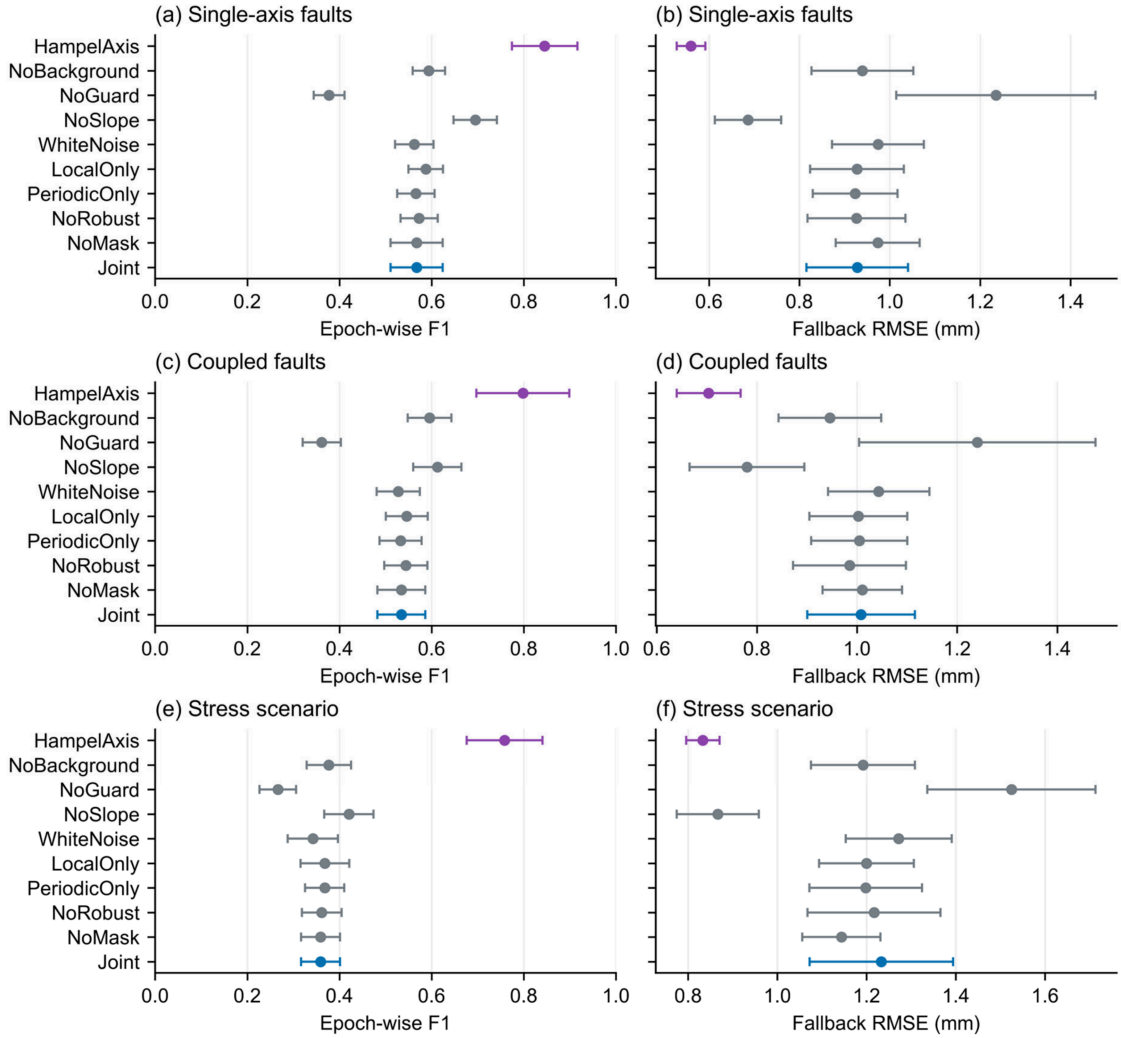

Figure 4. Epoch-wise detection F1 and reconstruction RMSE for single-axis, coupled, and stress scenarios. Points and error bars denote the mean and sample standard deviation over twenty seeds. Fallback RMSE includes all originally valid test coordinates, using raw observations when reconstruction is unavailable.

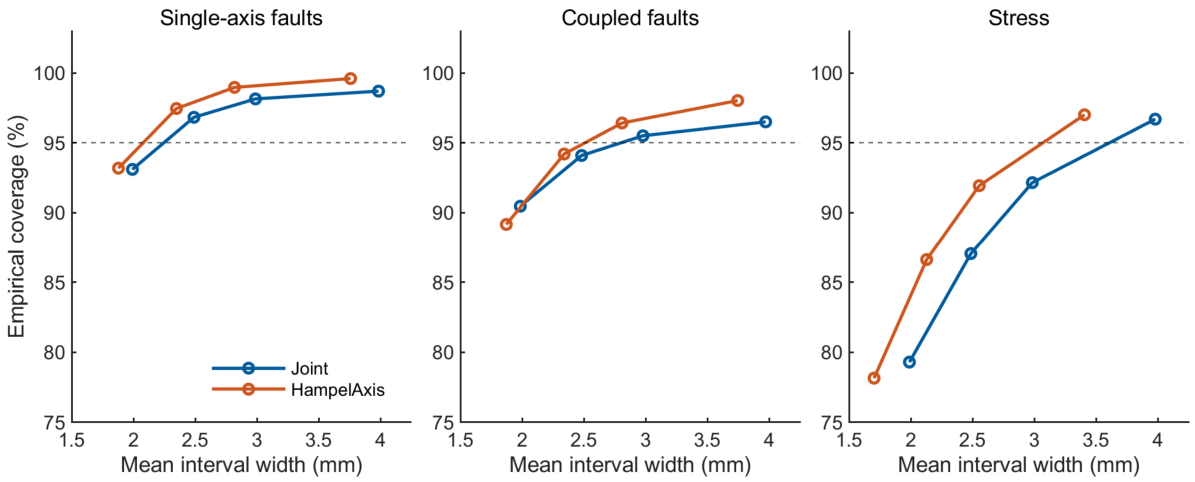

Figure 5. Mean empirical interval coverage versus mean width. Along each curve, points from left to right multiply plug-in standard errors by 1, 1.25, 1.5 and 2. Dashed line denotes 95%. Results use twenty seeds and each method's successful reconstruction domain.

3.2. Effects of Individual Modules

The most consistent benefit comes from historical contamination control (Figure 6). Removing it reduces F1 by 0.1901, 0.1731, and 0.0926 and increases RMSE by 0.3065, 0.2322, and 0.2919 mm in the single-axis, coupled, and stress cases, respectively. All six paired intervals exclude zero. This pattern is consistent with historical faults affecting later windows, although the ablation alone does not uniquely identify the mechanism.

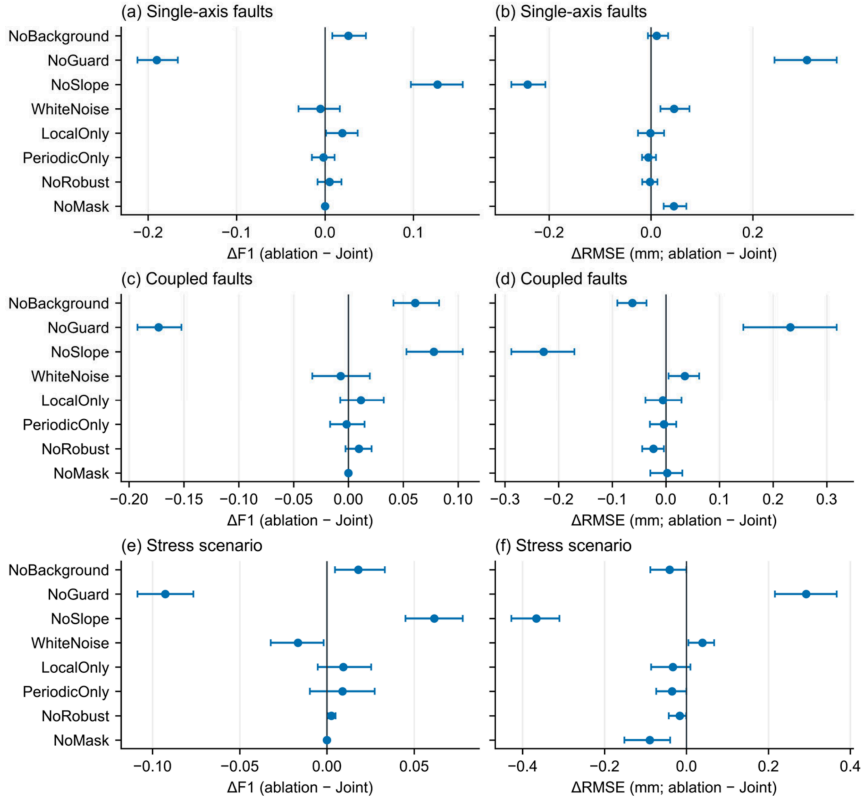

Figure 6. Paired ablation differences, defined as variant minus Joint. Error bars are nominal 95% percentile bootstrap intervals from 2000 resamples of twenty seed pairs. Positive ΔF1 favors the ablation; positive ΔRMSE favors Joint. Intervals are exploratory and are not adjusted for multiple comparisons.

Other components show mixed or unfavorable effects on aggregate performance. Removing the background improves F1 in all three scenarios, and the complete dual-library background provides no consistent reconstruction benefit. LocalOnly and PeriodicOnly remain close to Joint, while robust historical weighting does not consistently improve performance. NoSlope achieves both higher F1 and lower fallback RMSE than Joint in all three scenarios (Table 3). The inclusion of these components therefore cannot be justified by aggregate F1 and RMSE alone. A narrower difference occurs near the genuine displacement step in the stress scenario, where the non-fault flag proportion is 20.00% for Joint and 24.94% for NoSlope. This descriptive local comparison does not establish protection of arbitrary steps or nonlinear deformation, and it does not reverse the overall performance comparison. The correlated temporal model reduces RMSE relative to WhiteNoise by 0.0455, 0.0352, and 0.0391 mm, with paired intervals excluding zero. Masking is less consistent: NoMask increases RMSE by 0.0451 mm in single-axis cases but reduces it by 0.0895 mm in stress cases. The stress result indicates that erroneous exclusions can outweigh the benefit of removing observation faults. More model components therefore do not necessarily yield better reconstruction.

3.3. Event Recall and False-Positive Burden

At the original fixed calibration thresholds, Joint event recall is 98.13%, 92.81%, and 97.37%, versus 89.38%, 81.25%, and 81.09% for HampelAxis. These rates do not control the false-positive burden. Figure 7 therefore also compares threshold-dependent event recall at matched false-positive budgets and interpolated precision–recall curves. At a retrospective 1% budget, Joint recall is 69.06%, 55.63%, and 69.93%, versus 89.69%, 85.00%, and 79.28%. Paired differences are −20.63, −29.38, and −9.34 percentage points; exploratory 95% seed-bootstrap intervals are [−28.44, −12.81], [−36.26, −22.19], and [−19.43, 0.66]. Thus, higher original event recall does not establish a detection advantage at a low matched false-positive budget. These retrospective operating points use test labels for evaluation and are not deployable calibration rules.

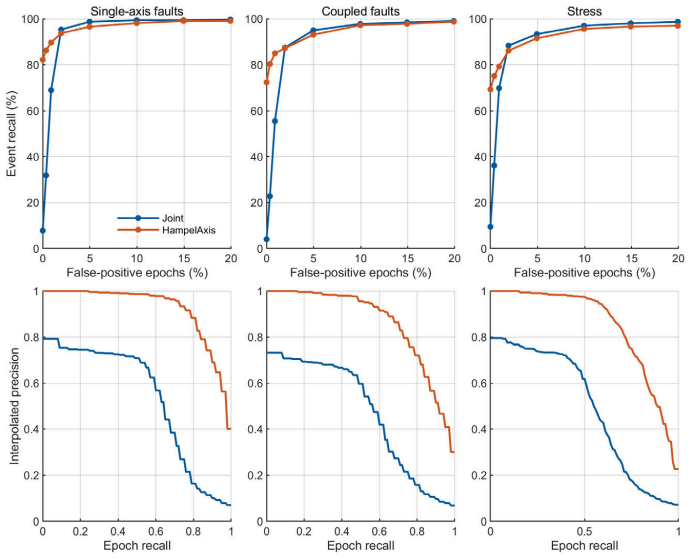

Figure 7. Threshold-dependent comparison of Joint and HampelAxis. Top: mean event recall versus achieved epoch-wise false-positive rate. Bottom: mean interpolated epoch-level precision–recall curves. Results summarize twenty paired realizations per scenario on the common evaluation domain. Test labels are used for retrospective comparison; these curves do not represent operational threshold calibration.

At the original thresholds, Joint false-positive rates are 3.97%, 3.73%, and 13.90%, versus 0.97%, 0.77%, and 1.36% for HampelAxis. Under stress, Joint produces 2.774 falsely flagged epochs and 2.202 false-positive episodes per monitoring day on the stated evaluation domain. An episode is a run of consecutive eligible falsely flagged grid epochs; an unevaluable or non-false-positive epoch breaks a run. Counts are divided by the full 25 d test duration, so missing/unevaluable periods do not contribute false positives.

Periodic recalibration reduces this burden but also reduces event sensitivity (Table 5). Weekly updating lowers Joint's stress-scenario FPR to 1.73%, while event recall decreases from 97.37% to 44.38%. Daily updating yields an FPR of 0.86% and event recall of 39.59%. Thus, the reduction in false positives does not establish improved detection performance independently of the accompanying loss of event recall.

Table 5. Joint performance under fixed thresholds and periodic recalibration in the stress scenario.

Calibration Strategy	FPR (%)	Event Recall (%)	False-Positive Epochs/Day	False-Positive Episodes/Day
Fixed threshold	13.90	97.37	2.774	2.202
Weekly recalibration	1.73	44.38	0.344	0.240
Daily recalibration	0.86	39.59	0.172	0.136

Values are means across twenty realizations evaluated on the same common test domain. FPR uses eligible non-faulty epochs, and event recall uses evaluable fault events. Daily counts use the full 25 d test duration. Recalibration uses preceding observations only.

3.4. Field Availability and Candidate Behavior

Of the 39,031 selected grid epochs, 34,895 are complete, 52 contain partial observations, and 4084 have no observations (Figure 8). Partial missingness occurs only at S02. Complete-epoch percentages range from 75.10% to 97.49%. The distinction between available coordinates and complete epochs matters because joint detection and component-wise reconstruction may use different domains.

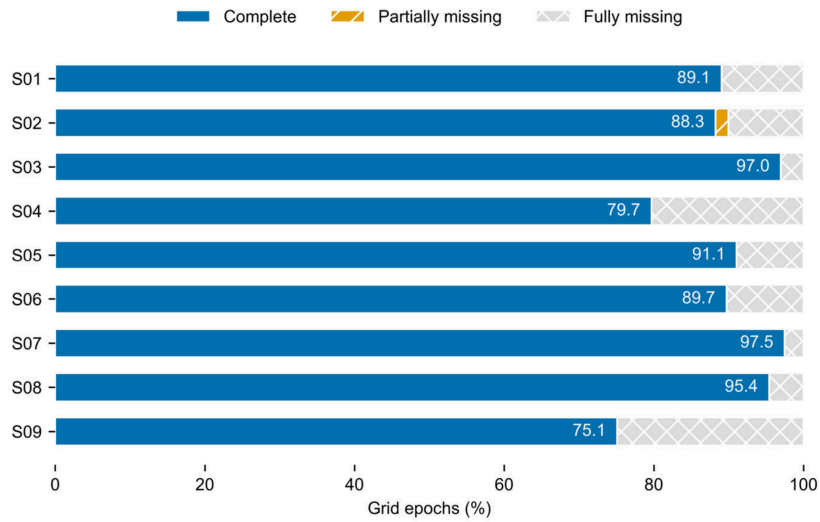

Figure 8. Complete, partially missing, and fully missing epochs on the selected receiver grids. Categories sum to 100% for each receiver. Labels within the complete-epoch bars give their percentages. S03 has a 6 h cadence; the remaining receivers have a 1 h cadence.

Joint candidate rates range from zero at S02 to 8.76% at S05 (Table 6). The absence of S02 candidates does not establish the absence of faults; its calibrated threshold is high relative to

its test scores. At S04, the learned background is active for only 57.89% of scoreable epochs. The remaining epochs use the fallback without a background basis. Accordingly, the field evaluation does not imply that every Joint output uses the complete dual-history model.

Table 6. Field candidate rates and coordinate reconstruction coverage.

Receiver	Joint Candidates (%)	Split Candidates (%)	Hampel Candidates (%)	Joint Coverage (%)	SplitEpoch Coverage (%)	SplitAxis Coverage (%)	Hampel Coverage (%)
S01	3.19	4.13	2.52	98.81	97.86	97.86	98.81
S02	0.00	0.00	0.36	99.29	98.93	98.93	99.29
S03	6.95	9.39	4.07	97.99	88.80	94.50	97.79
S04	0.48	0.84	1.77	99.89	99.52	99.68	99.77
S05	8.76	14.95	2.33	97.81	92.41	96.54	98.16
S06	2.18	3.89	0.61	99.03	98.14	98.45	99.01
S07	1.57	1.07	2.07	99.69	99.48	99.54	99.69
S08	1.09	1.81	1.72	99.50	99.15	99.15	99.50
S09	0.30	0.84	1.35	98.58	98.06	98.06	98.81

SplitEpoch and SplitAxis share candidate decisions. Candidate rates are the proportions of method-specific scoreable test epochs flagged by the detector. They describe model-generated candidates, not independently verified observation faults or false positives. Reconstruction coverage is the proportion of originally observed test coordinates for which a finite estimate is returned. It measures output availability and does not establish coordinate accuracy, uncertainty calibration, or the absence of undetected faults. Accordingly, a zero candidate rate does not imply a fault-free record. Hampel denotes HampelAxis. These field candidate rates are not false-positive rates.

Selective component masking improves coverage for some receivers. S03 coverage increases from 88.80% with SplitEpoch to 94.50% with SplitAxis; S05 increases from 92.41% to 96.54%. Joint coverage at S03 is 97.99%. These differences describe output availability, not fault classification accuracy. Additional internal consistency checks assess reconstruction-minus-observation residuals, adjacent coordinate increments, and absolute reconstruction changes for each receiver and component (Table 7). Across the 27 receiver–component series, the median scaled residual MAD on unflagged successful coordinates is 0.673 mm for Joint and 1.051 mm for HampelAxis. The corresponding median reconstructed-increment MADs are 1.141 and 0.890 mm. These statistics use each method's available observations and reconstructions; increment comparisons use adjacent pairs valid in both the original and reconstructed series. Increments refer to each receiver's analysis-grid interval and are not normalized to hourly rates.

Table 7. Internal field consistency by receiver and method. Each row gives the median across its three coordinate-specific statistics, in mm.

Receiver	Method	Residual MAD	Raw Increment MAD	Output Increment MAD	Median Axis Max Change
S01	HampelAxis	0.800	0.890	0.574	4.809
S01	Joint	0.488	0.890	0.780	6.518
S02	HampelAxis	0.757	0.741	0.448	7.889
S02	Joint	0.461	0.741	0.570	4.492
S03	HampelAxis	1.036	1.631	0.979	9.388
S03	Joint	0.673	1.631	1.234	8.095
S04	HampelAxis	1.200	1.038	0.853	17.239
S04	Joint	0.852	1.038	0.855	32.148
S05	HampelAxis	0.709	0.741	0.701	4.797
S05	Joint	0.356	0.741	0.679	22.427

Table 7. *Cont.*

Receiver	Method	Residual MAD	Raw Increment MAD	Output Increment MAD	Median Axis Max Change
S06	HampelAxis	0.467	0.445	0.458	4.132
S06	Joint	0.237	0.445	0.484	7.915
S07	HampelAxis	0.355	0.445	0.339	4.091
S07	Joint	0.247	0.445	0.384	5.112
S08	HampelAxis	2.878	4.893	2.152	194.513
S08	Joint	2.269	4.893	3.141	192.075
S09	HampelAxis	1.385	3.262	1.702	11.844
S09	Joint	0.875	3.262	2.579	10.602

Residuals compare reconstruction and raw input on successful unflagged coordinates. Increment MAD uses adjacent common valid pairs. The last column is a median of three axis-specific maximum absolute changes, not a station maximum.

3.5. Paired Deformation Transfer

All four methods retain the final injected plateau closely (Table 8; Figure 9). For Joint, transmitted plateau displacement ranges from 0.999855 to 1.000000 mm across the 27 receiver–axis combinations. Its median transition RMSE is 0.0108 mm, but the maximum is 1.8276 mm. The final plateau therefore provides an incomplete description of signal preservation.

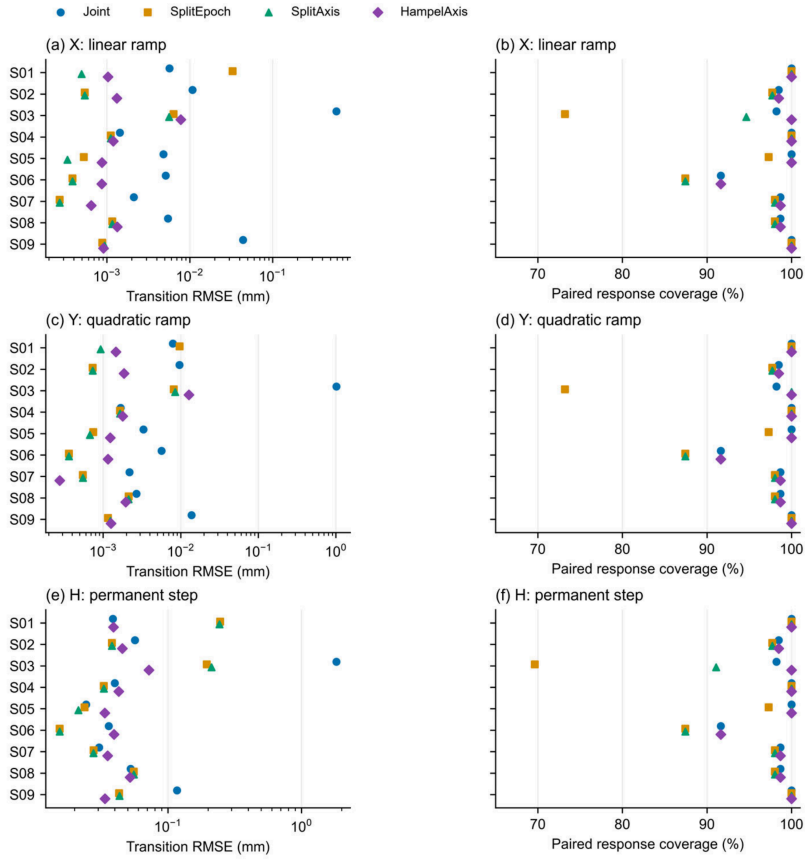

Figure 9. Transition error and paired response coverage for the field injections. Rows correspond to the X linear ramp, Y quadratic ramp, and H permanent step. Each marker represents one receiver and method; the horizontal error scale is logarithmic. Joint has its largest transition error for the S03 H step. Shape is fixed to component in this experiment, so these points are not independent replicated comparisons of shape.

Table 8. Paired 1 mm injection responses across 27 receiver–axis combinations.

Method	Median Transition RMSE (mm)	Maximum Transition RMSE (mm)	Transmitted Plateau Range (mm)	Minimum Paired Coverage (%)
Joint	0.0108	1.8276	0.999855–1.000000	91.63
SplitEpoch	0.0021	0.2463	1.000000–1.000000	69.64
SplitAxis	0.0012	0.2435	1.000000–1.000000	87.45
HampelAxis	0.0018	0.0721	1.000000–1.000000	91.63

The median and maximum summarize heterogeneous receiver–axis combinations rather than independent random replicates. Transition errors use available paired responses; paired coverage is required to interpret selective availability. Plateau values are transmitted displacement, not absolute positioning accuracy.

Detailed S03 diagnostics reproduce the original combined-injection transition RMSE of 1.8276 mm. An H-only injection yields 1.8280 mm over 55 of 56 originally observed H coordinates, indicating that simultaneous X/Y ramps are not required for the large error. At test day 7.25, six hours after step onset, the H observation changes from 3.7 to 4.7 mm, whereas the reconstructed value changes from 7.395 to 21.817 mm. The resulting paired response is 14.423 mm for a 1 mm addition. The detection score rises from 1.3625 to 2.2703, crossing the frozen threshold of 1.5075 (Figure 10).

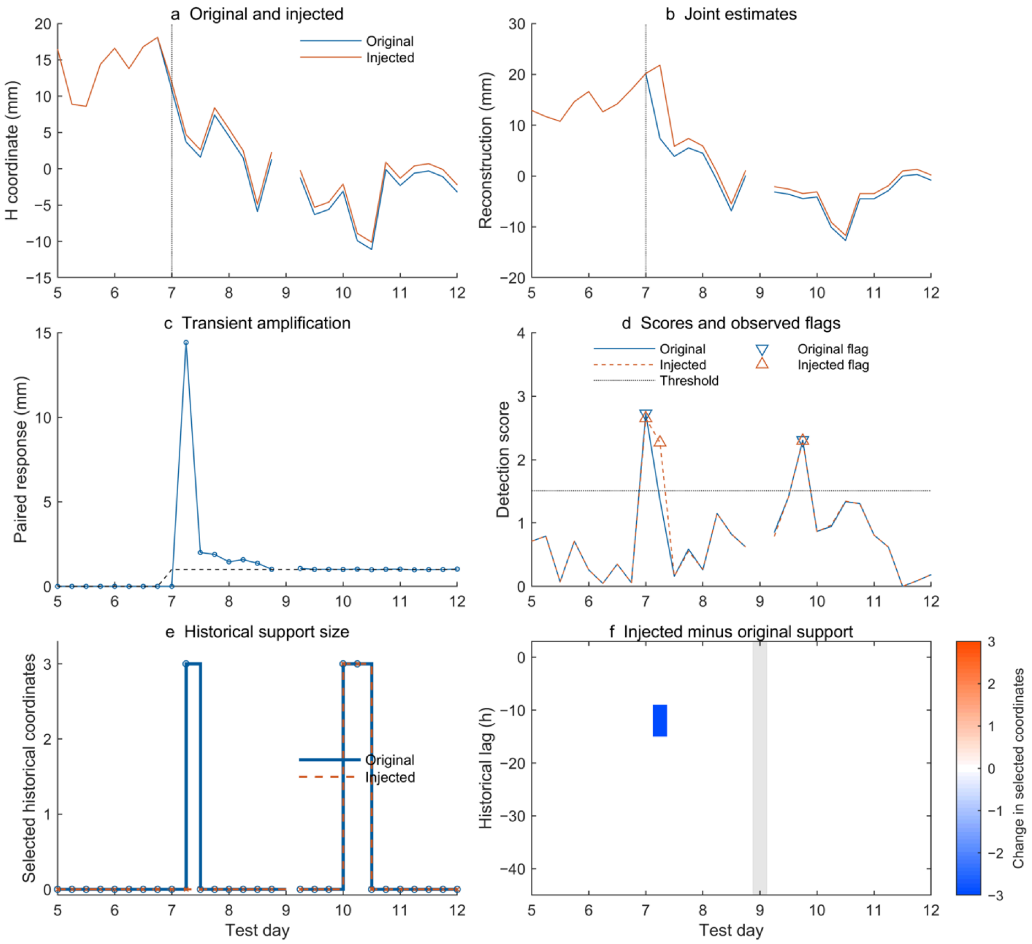

Figure 10. S03 H-only step response. Panels show original/injected observations, reconstructed coordinates, paired signal response, detection scores with actual flags and frozen threshold, selected historical coordinate counts, and injected-minus-original support by relative historical lag. Gray cells indicate no diagnostic information.

At this epoch, the selected historical support changes from three coordinate rows to zero. These rows belong to one historical epoch. That previously excluded epoch becomes available to the reconstruction, while the newly flagged current epoch is excluded. Both fits retain 18 coordinate observations, but their composition differs. The whitened normal-matrix condition estimates increase from approximately 154 to 372, remaining below the rejection threshold of 10^{11}. Thus, the peak error occurs despite successful numerical checks.

4. Discussion

4.1. Trade-Offs Between Deformation Protection and Anomaly Detection

The protection basis specifies which motion components are excluded from anomaly testing and retained during reconstruction. Under fixed covariance, model components, masks, and observation support, additions within the protected constant-plus-linear space do not affect the anomaly residual and are preserved by full-rank reconstruction. These conditional properties do not imply higher F1 or lower reconstruction error. Fault components occupying the same protected space are also excluded from the detection evidence, so the constraint cannot by itself distinguish genuine deformation from faults with identical temporal signatures. The existing results quantify this cost: removing slope protection reduces mean fallback RMSE from 0.928 to 0.686 mm in the single-axis scenario and from 1.233 to 0.866 mm under stress (Table 3). The protected space must therefore be selected according to the deformation components of interest, with its effects on fault sensitivity and reconstruction accuracy evaluated explicitly. Historical contamination control provides a clearer benefit: removing this component increased the single-axis RMSE from 0.928 to 1.235 mm, with deterioration also observed in the other scenarios. In contrast, the dual-history background and robust weighting produced no consistent improvement. Calibration and masking introduce another limitation. Under stress, Joint reached a false-positive rate of 13.90%, while retaining observations without the reconstruction mask reduced RMSE from 1.233 to 1.143 mm. This combination is consistent with useful observations being excluded after calibration deteriorates, although the aggregate results do not isolate that mechanism. The present evidence therefore supports historical contamination control more strongly than the complete combination of background learning, protection, and masking. Periodic recalibration reduces false positives under stress but substantially reduces event recall (Table 5). Using preceding observations preserves causality without ensuring a fault-free calibration pool: unidentified faults or genuine motion can increase the upper score quantile and suppress detections. These detection results do establish improved reconstruction or guaranteed 1% false-positive control.

4.2. Field Signal Transfer and Monitoring Reliability

The field injections reveal a substantial difference between preserving final displacement and reproducing its transition. Joint transmitted the 1 mm plateau as 0.999855–1.000000 mm across the 27 receiver–component combinations, but its paired-response transition RMSE reached 1.8276 mm for the S03 H-component step (Table 5). HampelAxis also preserved the plateau and had a smaller maximum transition RMSE of 0.0721 mm. Plateau retention alone therefore does not establish an advantage for Joint. S03 uses a 6 h sampling interval and a 48 h window. When this window spans the step, the signal contains a discontinuity that cannot be represented exactly by the protected constant-plus-linear basis. Changes in the selected nuisance support and reconstruction mask may further affect the response; their individual contributions cannot be separated by the present experiment.

These injection errors quantify the transfer of an added signal through the processing procedure. They do not measure absolute field displacement accuracy, because differencing the original and injected outputs partly cancels their shared background and noise.

Reliability also depends on uncertainty calibration. The 79.30% coverage of nominal 95% intervals under stress indicates that reconstructed displacement can appear more certain than warranted, potentially overstating the significance of weak changes. Interval widening improves coverage in these simulations but requires independent calibration and validation before operational use. More complete assessment could use independent displacement references or temporally structured resampling that repeats all estimation and selection stages. For monitoring applications, the evaluation must therefore consider transition error, unavailable outputs, false alarms, and interval coverage together. Independent displacement observations are needed to determine whether naturally occurring changes are retained accurately and whether candidate flags correspond to actual observation faults.

5. Conclusions

This study develops Joint, a causal subspace framework for anomaly detection and displacement reconstruction in GNSS coordinate time series. The framework combines explicit protection of local offsets and linear rates with dual-timescale background learning, correlated-noise whitening, and historical contamination control. Matched-subspace detection and masked generalized least-squares reconstruction connect anomaly screening with displacement estimation. Its causal structure supports sequential processing, while observation selection before whitening accommodates incomplete records.

The theoretical analysis establishes that, under a fixed model and observation support, protected components do not affect the anomaly residual and are preserved by full-rank reconstruction. This property provides an explicit, interpretable constraint on which motion components are retained during data cleaning. It also clarifies the conditions under which deformation preservation can be guaranteed. Validation includes three synthetic scenarios with 20 random seeds each and field records from nine receivers at six sites. Controlled ablations identify historical contamination control as the most consistently beneficial module. Field injections demonstrate near-complete retention of the final 1 mm plateau and quantify the limitations of transient recovery. The framework is therefore relevant to GNSS data quality control for landslide displacement, mining subsidence, and structural deformation monitoring. Its practical value is the coordinated treatment of prescribed deformation modes, historical contamination, and missing observations, with site-specific calibration and transient-response validation required for deployment.

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

Funding: This research was funded by the National Natural Science Foundation of China (42207534, U21A20108), the Natural Science Foundation of Henan Province (242300421363, 262300422738), Henan Province Science and Technology R&D Program (262102210234), and the Excellent Youth Foundation of Henan Polytechnic University (J2024-4).

Data Availability Statement: The algorithm validation data is provided by the 2nd Geological Brigade of Jiangsu Geological Bureau and can be obtained from the corresponding author Shi Junpeng (sjp@home.hpu.edu.cn). The algorithm was developed on MATLAB 2024a and executed on a workstation (Intel i7-8700K, 32 GB RAM, 2 TB SSD) to ensure computational efficiency.

Acknowledgments: The authors sincerely thank Kezhao Li of Henan Polytechnic University for his valuable theoretical guidance and support in experimental instrumentation throughout this study. His insightful academic suggestions greatly contributed to this research. Li passed away on 4 January 2026. The authors respectfully dedicate this work to his memory and outstanding scientific contributions.

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

References

1. Wang, J.; Nie, G.; Gao, S.; Wu, S.; Li, H.; Ren, X. Landslide Deformation Prediction Based on a GNSS Time Series Analysis and Recurrent Neural Network Model. *Remote Sens.* **2021**, *13*, 1055. [CrossRef]
2. Hu, X.; Bürgmann, R.; Xu, X.; Fielding, E.; Liu, Z. Machine-Learning Characterization of Tectonic, Hydrological and Anthropogenic Sources of Active Ground Deformation in California. *J. Geophys. Res. Solid Earth* **2021**, *126*, e2021JB022373. [CrossRef]
3. Donoso, F.; Yáñez, V.; Ortega-Culaciati, F.; Moreno, M. A machine learning approach for slow slip event detection using GNSS time-series. *J. S. Am. Earth Sci.* **2023**, *132*, 104680. [CrossRef]
4. Xiong, K.; Liu, Z.; Niu, Y. GNSS-RTK time series denoising based on deep learning and mode decomposition techniques for offshore platform. *GPS Solut.* **2025**, *29*, 132. [CrossRef]
5. Santamaría-Gómez, A.; Ray, J. Chameleonic Noise in GPS Position Time Series. *J. Geophys. Res. Solid Earth* **2021**, *126*, e2020JB019541. [CrossRef]
6. Gobron, K.; Rebischung, P.; de Viron, O.; Demoulin, A.; Van Camp, M. Impact of offsets on assessing the low-frequency stochastic properties of geodetic time series. *J. Geod.* **2022**, *96*, 46. [CrossRef]
7. Oelsmann, J.; Passaro, M.; Sánchez, L.; Dettmering, D.; Schwatke, C.; Seitz, F. Bayesian modelling of piecewise trends and discontinuities to improve the estimation of coastal vertical land motion. *J. Geod.* **2022**, *96*, 62. [CrossRef]
8. Crocetti, L.; Schartner, M.; Soja, B. Discontinuity Detection in GNSS Station Coordinate Time Series Using Machine Learning. *Remote Sens.* **2021**, *13*, 3906. [CrossRef]
9. Ghaderpour, E.; Pagiatakis, S.D.; Scarascia Mugnozza, G.; Mazzanti, P. On the stochastic significance of peaks in the least-squares wavelet spectrogram and an application in GNSS time series analysis. *Signal Process.* **2024**, *223*, 109581. [CrossRef]
10. Li, W.; Li, Z.; Jiang, W.; Chen, Q.; Zhu, G.; Wang, J. A New Spatial Filtering Algorithm for Noisy and Missing GNSS Position Time Series Using Weighted Expectation Maximization Principal Component Analysis: A Case Study for Regional GNSS Network in Xinjiang Province. *Remote Sens.* **2022**, *14*, 1295. [CrossRef]
11. Ji, K.; Shen, Y.; Chen, Q.; Feng, T. Extended Principal Component Analysis for Spatiotemporal Filtering of Incomplete Heterogeneous GNSS Position Time Series. *IEEE Trans. Geosci. Remote Sens.* **2023**, *61*, 1–19. [CrossRef]
12. Liu, H.; Li, L. Missing Data Imputation in GNSS Monitoring Time Series Using Temporal and Spatial Hankel Matrix Factorization. *Remote Sens.* **2022**, *14*, 1500. [CrossRef]
13. Gao, W.; Li, Z.; Chen, Q.; Jiang, W.; Feng, Y. Modelling and prediction of GNSS time series using GBDT, LSTM and SVM machine learning approaches. *J. Geod.* **2022**, *96*, 71. [CrossRef]
14. Costantino, G.; Giffard-Roisin, S.; Dalla Mura, M.; Socquet, A. Denoising of Geodetic Time Series Using Spatiotemporal Graph Neural Networks: Application to Slow Slip Event Extraction. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2024**, *17*, 17567–17579. [CrossRef]
15. Mastella, G.; Bedford, J.; Corbi, F.; Funiciello, F. Denoising daily displacement GNSS time series using deep neural networks in a near real-time framing: A single-station method. *Geophys. J. Int.* **2025**, *242*, ggaf207. [CrossRef]
16. Okada, Y.; Nishimura, T.; Tabei, T.; Matsushima, T.; Hirose, H. Development of a detection method for short-term slow slip events using GNSS data and its application to the Nankai subduction zone. *Earth Planets Space* **2022**, *74*, 18. [CrossRef]
17. Yano, K.; Kano, M. L1 Trend Filtering-Based Detection of Short-Term Slow Slip Events: Application to a GNSS Array in Southwest Japan. *J. Geophys. Res. Solid Earth* **2022**, *127*, e2021JB023258. [CrossRef]
18. Yu, L.-R.; Lu, Q.-H.; Xue, Y. DTAAD: Dual TCN-attention networks for anomaly detection in multivariate time series data. *Knowl.-Based Syst.* **2024**, *295*, 111849. [CrossRef]
19. Ge, N.; Weng, X.; Yang, Q. Unsupervised anomaly detection via two-dimensional singular value decomposition and subspace reconstruction for multivariate time series. *Appl. Intell.* **2023**, *53*, 16813–16829. [CrossRef]
20. Fährmann, D.; Damer, N.; Kirchbuchner, F.; Kuijper, A. Lightweight Long Short-Term Memory Variational Auto-Encoder for Multivariate Time Series Anomaly Detection in Industrial Control Systems. *Sensors* **2022**, *22*, 2886. [CrossRef]
21. Scharf, L.L.; Friedlander, B. Matched subspace detectors. *IEEE Trans. Signal Process.* **1994**, *42*, 2146–2157. [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.