



20th NKG General Assembly — Vilnius, Lithuania — September 2026

Anatomy of Offset-Induced Noise Disturbances GNSS Time Series

Nordic–Baltic regional experiment


Component-Dependence, Magnitude-Scaling, and localization within the residual structure

Alejandro Arias-Gallegos · Hannu Koivula · Maaria Nordman

Department of Geodesy and Geodynamics, FGI — National Land Survey of Finland / Aalto University



Offsets are not just position jumps



- **An abrupt discontinuity**
instrumental, geophysical, or processing-related — documented or not.
- **Poorly modeled offsets bias velocities**
and their uncertainties, and can mimic spurious low-frequency noise.
- **But the effect is not only positional**
an offset retained in the residual also changes how the local noise behaves.

Central question: how, where, and with what intensity do offsets leave a measurable imprint on GNSS residuals?

1 Memory
of Noise

2 Experimental
design

3 Existence of
the Event




4 Impact of
the
magnitude

5 Residual
structure

6 Conclusions

Concept

Local noise memory — and how an offset fakes it



Memory = how much a residual's past values remain informative about its present state — quantified via autocorrelation. Not global and fixed: here it's estimated **locally, in moving windows**, as a function $M(t)$.

Key idea

A retained offset introduces a step-like structure inside what should be purely stochastic. Memory metrics can interpret that structure as temporal persistence — an apparent memory, not necessarily a real change in the physical noise-generating process.

Nordic–Baltic network, 146 unified offsets

13 Nordic–Baltic stations

N, E, U components · JPL + BIFROST solutions

146

unified offset
events (vs. 105
from a single source)

137

valid events
Met coverage and criteria for offset vs.
random)

Moving-window noise-memory metrics

1.5 years

Window

30 days

Step

≥ 70%

Coverage

- **Lag-1 autocorrelation** — *day-to-day persistence*
- **CAR(1) correlation time** — *AR(1)-derived timescale*
- **ACF decorrelation time** — *most discriminant metric*
- **Local noise amplitude** — *control, not memory*


- *Residuals: linear trend + annual + semi-annual signal removed by least squares; offsets deliberately retained.*
- *Local metrics*
- *Disturbance*

The study focuses on the statistical response of the residuals around documented offsets, not on detecting them.

Step 1

The phenomenon is real: offsets vs. random epochs

137 documented offsets vs. 146 random control epochs. One-sided Mann-Whitney U + Cliff's delta — all four metrics significant ($p < 0.01$).

**p-values**

local_amplitude_mm 1.7×10^{-10}

acf_decorr_days 2.4×10^{-9}

lag1_acf 1.2×10^{-5}



car1_days 9.2×10^{-3}

*All four metrics distinguish offsets from random epochs
— a methodological step, not the headline result.*

Central result A

Perturbation intensity scales with offset magnitude

Spearman correlation between |magnitude| and MAD-normalized perturbation, using magnitudes reported by catalog.



12 / 12

metric × component combinations statistically significant

Pattern most clear in North, more scattered in East and Up for this dataset — consistent with PELT's finding that Up is the noisiest, hardest-to-detect component.

Estimated magnitudes included → pattern weakens (regression dilution)

Does result A survive scrutiny?

Multiple-comparison correction

12 / 12 combinations remain significant after Benjamini-Hochberg and Holm correction ($\alpha = 0.05$).

Station-level cluster bootstrap

5000 resamples — 95% CI excludes zero for all four metrics (ρ point estimates 0.63–0.74).

Leave-one-station-out


Correlation stays significant with any single station removed — no station drives the result.

Caveat: catalog-reported magnitudes come from only 4 of the 13 stations (JPL coverage). Robust within that subset; the pattern holds — more weakly — across all 13 with estimated magnitudes.

Central result B

Where the perturbation lives in the residual

An edge-preserving bilateral filter separates the residual into a smoothed/correlated layer and a high-frequency remainder. **Two-sided control:** *does the offset signature survive in each layer?*



ACF decorrelation time

27 d

Full residual

92 d

Filtered

1 d

High-freq. remainder

Effect collapses in the remainder, is preserved in the filtered layer — confirmed independently by each layer's own autocorrelation structure.

Things to keep in mind

No physical noise claim

We describe apparent memory / residual structure, not a genuine change in the noise-generating process.

Nearby offsets

Events within the same ± 150 -day window can partially contaminate each other's background estimate.

Overlapping windows

~95% shared data between consecutive windows — deliberate smoothing, not independence.

Structural, not spectral

The bilateral filter is an edge-preserving separation, not a classical Fourier/wavelet decomposition.

Why this matters - not just an offset detector

This is not an offset detector — dedicated methods like PELT already do that, and do it faster. Its value is diagnostic: once an offset is known, it tells us how it interacts with the surrounding noise structure.

Offset-catalog validation

Magnitude / perturbation mismatches -- like the PTBB case -- flag processing artifacts worth a second look.

Triage for undocumented events

A strong local perturbation with no cataloged offset is a concrete candidate for manual review.

Stochastic modeling

The perturbation lives in the correlated structure -- informs covariance models that allow local, transient changes around known discontinuities, for more realistic velocity uncertainties.

Regional extension

First application of this diagnostic lens to the Nordic-Baltic / BIFROST network, distinct from OMP's global scope.

The goal isn't finding offsets faster -- it's understanding what they do to the noise once they're there.

Offsets leave a measurable, structured imprint

Offsets retained in GNSS residuals generate measurable local perturbations in noise-memory and amplitude metrics, which **increase with the catalog-reported offset magnitude** — a relation that survives multiple-comparison correction, station bootstrap, and leave-one-station-out.

The response is **component-dependent** — clearest in North, more scattered in East and Up for this dataset.

The perturbation concentrates in the **preserved/smoothed structure** of the residual and nearly vanishes in the weakly-correlated high-frequency remainder.

A diagnostic framework, not just a detector — *see previous slide for what this means for catalog validation, stochastic modeling, and the Nordic–Baltic network.*

Thank you

Alejandro Arias-Gallegos

julio.ariasgallegos@maanmittauslaitos.fi

FGI — National Land Survey of Finland