NERVOUS MACHINE

Space Weather Workshop 2026 · Boulder, CO · Poster #18 · Wed April 29

# SPACE ATLAS

Driver-resolved attribution and forecasting of voxel-level thermospheric density, with per-edge uncertainty and learned time-lag.

Heidi Bennett · Nervous Machine · space-atlas.nervousmachine.com

## Problem

Operational thermospheric models (JB2008, NRLMSISE-00) estimate bulk density but do not resolve *which* driver — or at what lag — is responsible for a given perturbation. Without per-driver attribution and learned time-lag structure, there is no mechanism to forecast how a driver impulse propagates across orbital shells over hours to days.

## Method

A causal-discovery pipeline ingests five space-weather indices — `Dst`, `IMF Bz`, `AE`, solar wind density, solar wind speed — and decomposes voxel-level density into per-driver contributions. Each causal edge carries a learned multiplier `m`, a time-lag distribution, and an evidence-weighted certainty `Z`. Activity-weighted attribution and quiet-time baselines prevent confidence inflation during geomagnetic silence.

*24 voxels · 104 causal edges · 460–510 km altitude band.*

## Validation

GRACE-FO accelerometer density (TU Delft TOLEOS v02c, 10-second raw averaged to hourly, 2024-12 → 2025-12): 3.41 M raw / 89.8 k processed observations.

Every driver recovered from near-noise TLE baseline. Four of five drivers cleared `Z > 0.91`; all five cleared `Z > 0.80`. Largest recovery: solar wind density `Δ +0.74`. Learned lag structure produced validated multi-day forecasts using held-out GRACE-FO as error signal.

## Operational path forward

Live LEO GPS drag residuals already provide GRACE-FO-comparable cadence across the operational fleet but remain unexploited for thermospheric characterization. Ingesting these signals yields certainty-quantified, driver-resolved nowcasts and spatiotemporal perturbation forecasts without dedicated science missions. The same attribution + lag mechanism extends to seasonal-latitudinal variation and Joule-heating-driven density blooms.


## FIVE-DRIVER LEARNING TRAJECTORIES

One representative causal edge per driver · GRACE-FO TOLEOS v02c · 2024-12 → 2025-12

Z curve (cyan) · driver multiplier m (warm, dashed) · Δ certainty recovered vs TLE baseline


`η(Z) = 1 / (1 + e^{10(Z − 0.5)})` learning rate contracts as certainty earns itself · obs count, 0 → 5,200 on shared scale 24 voxels · 104 edges · 460–510 km · 89.8 k hourly obs

## Driver attribution · GRACE-FO vs TLE baseline

| DRIVER | $Z_{TLE}$ | $Z_{GRACE\text{-}FO}$ | Δ | M | EDGES |
|---|---|---|---|---|---|
| **Dst** *geomagnetic storm index* | 0.49 | 0.940 | +0.45 | 1.44 | 24 |
| **IMF Bz** *interplanetary magnetic field, z* | 0.45 | 0.912 | +0.46 | 1.72 | 24 |
| **AE index** *auroral electrojet · Joule-heating proxy* | 0.39 | 0.949 | +0.56 | 1.71 | 8 |
| **Solar wind ρ** *number density · dynamic-pressure proxy* | 0.20 | 0.937 | +0.74 | 1.28 | 24 |
| **Solar wind v** *bulk speed* | 0.13 | 0.854 | +0.72 | 1.08 | 24 |

*TLE baseline certainties at top of pipeline metadata. GRACE-FO values are per-driver averages across all edges for that driver. Δ is the certainty recovered when cadence tightens from ~daily to sub-minute — the signal was there; the sampling wasn't.*

## PER-DRIVER LAG DISTRIBUTION

fitted τ across all 104 causal edges · 0–24 h search window · ordered by learned lag

*Each edge scans τ and picks the peak of its own r(τ); dot opacity ∝ |r|, median marked in cyan.*

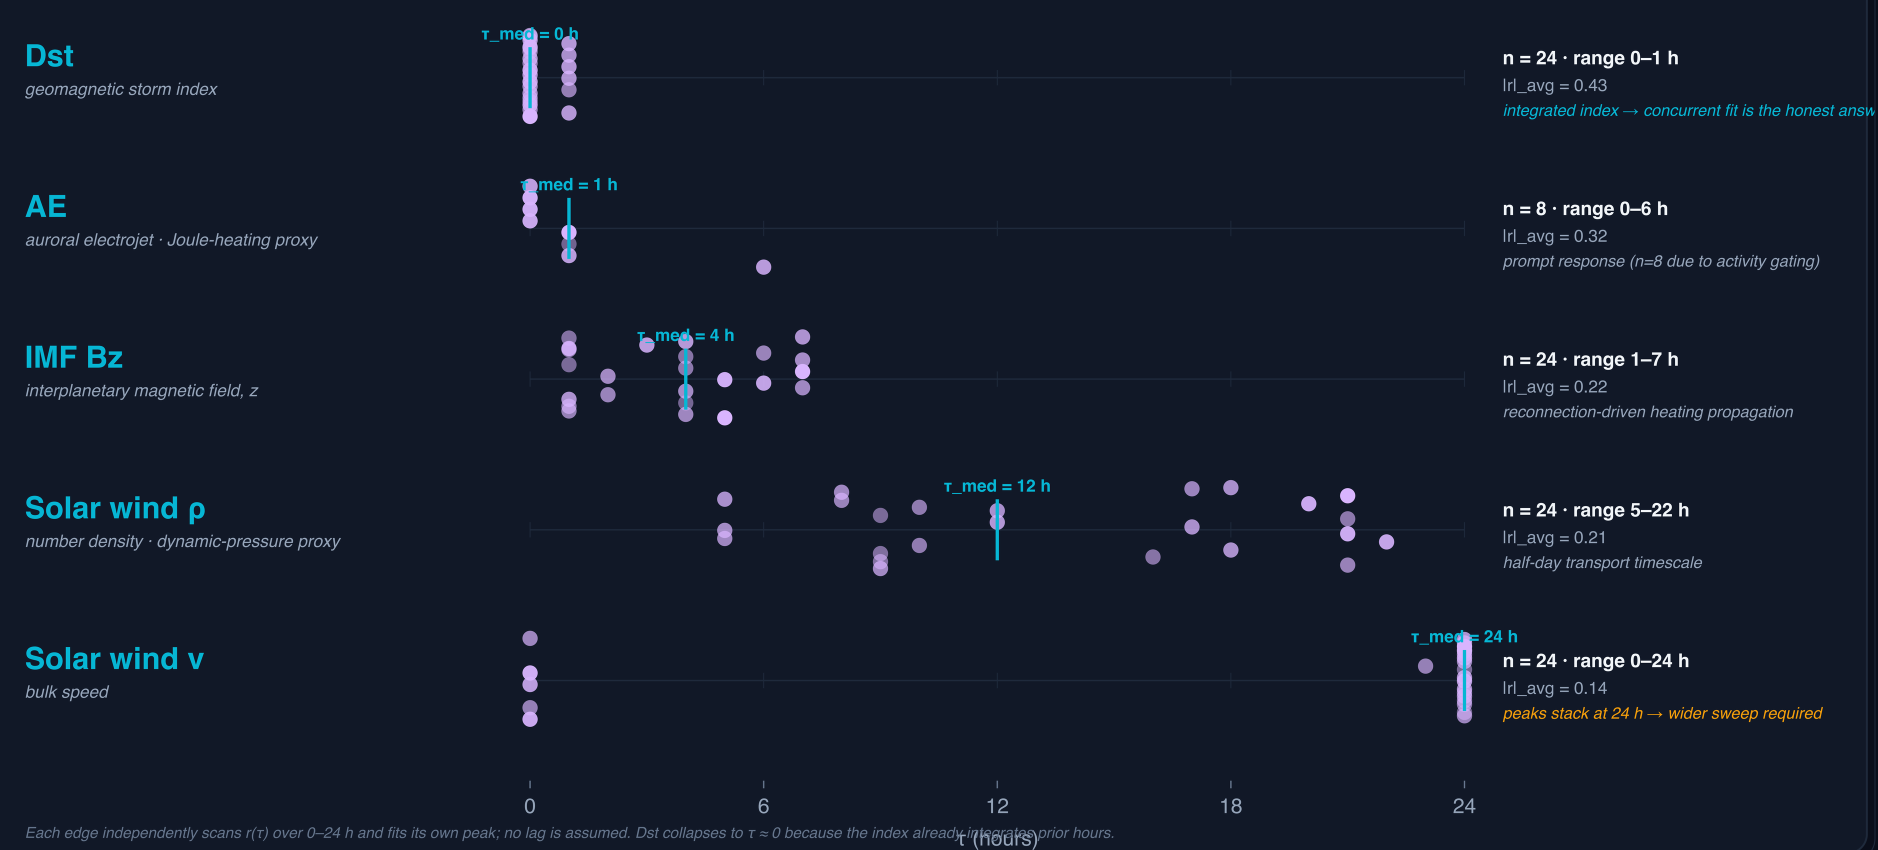


*Each edge independently scans r(τ) over 0–24 h and fits its own peak; no lag is assumed. Dst collapses to τ ≈ 0 because the index already integrates prior hours.*