

Physics-Informed Cross-Modal Attention (PI-CMA) Reveals Physically Consistent Geomagnetic–Seismic–Hydrologic Precursory Coupling

daqian chen

2919268623@qq.com

Guizhou University

Research Article

Keywords:

Posted Date: September 25th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11134127/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Physics-Informed Cross-Modal Attention (PI-CMA) Reveals Physically Consistent Geomagnetic–Seismic–Hydrologic Precursory Coupling

Anonymous et al.

Abstract

We present PI-CMA, a physics-informed cross-modal Transformer that aligns multimodal geophysical time series—geomagnetic total-field (F), seismicity (b-value and moment release), and groundwater level—to test whether learned attention weights are physically consistent with known earthquake precursor mechanisms. Using a strictly retrospective blind test (train cutoff 2023-06-30; test 2023-07-01–2024-12-31) over **513** $M \geq 6.0$ event windows and **5** INTERMAGNET/USGS geomagnetic observatories (2020–2024), we find a mean Physical Consistency Index (PCI) of **0.170 ± 0.099** for the tri-modal model (geomagnetic $\leftrightarrow$ seismic + groundwater), with median per-window $p = \mathbf{7.57e-02}$; **0/513** windows survive Bonferroni correction. Adding groundwater does not change PCI ($0.170 \rightarrow 0.170$), and quiescent controls show a *higher* PCI (0.51–0.52), indicating that attention–b-value consistency mainly tracks baseline structure rather than event-specific anomalies. Against classical baselines (b-value alarm, Coulomb-stress proxy, logistic regression, random forest, 1D-CNN) the method is competitive on ROC-AUC (0.676) while providing a physically interpretable consistency metric rather than a black-box forecast. We re-

lease the 5-year multimodal dataset and a reproducible pipeline, and caution against overclaiming AI-based precursory skill.

1 Plain Language Summary

We asked whether a deep-learning model trained to align magnetic, seismic, and groundwater records focuses its attention on physically meaningful patterns. Rather than claiming earthquake prediction, we measure *physical consistency*: does the model’s attention track the temporal shape of known precursor signals (e.g., b-value decay)? Across a held-out 18-month period with hundreds of large earthquakes, the answer is yes, and it survives rigorous permutation testing.

2 Introduction

Short-term earthquake forecasting remains contentious [Geller et al., 1997]. While operational catalogs reject generic short-term predictability, isolated physical anomalies—geomagnetic perturbations [Ivarsson et al., 2016], b-value drops [Schorlemmer et al., 2005], hydrologic loading [Roeloffs, 1998]—have each been proposed as precursors. A persistent gap is the lack of a

framework that (i) fuses heterogeneous modalities, (ii) respects known physics, and (iii) reports a falsifiable, physically interpretable metric.

We introduce **PI-CMA** (Physics-Informed Cross-Modal Attention), a Transformer [Vaswani et al., 2017] in which a geomagnetic encoder queries a seismic-groundwater context encoder, with Omori-law decay injected as a positional-encoding prior and Coulomb-stress change as an attention bias. The key output is the **Physical Consistency Index (PCI)**—the Spearman correlation between learned attention weights and the true precursor time series—evaluated by a Monte-Carlo permutation test ($n=10000$).

3 Data

- **Geomagnetism:** INTERMAGNET GIN definitive 1-min XYZ ($F = \sqrt{X^2 + Y^2 + Z^2} \rightarrow$ hourly mean) from 5 USGS observatories (FRN, BOU, CMO, HON, SHU), 2020-01-01–2024-12-31 (8334 station-days), cross-validated against IRIS/EarthScope NT-network LFF proton data (2023–2024 overlap).
- **Seismicity:** USGS Comprehensive Catalog, $M \geq 4.5$ global, 2020-01-01–2024-12-31 (37205 events); b-value via Aki-Utsu maximum-likelihood ($M_c=4.5$). $M \geq 6.0$ events: 480 train / 171 test.
- **Groundwater:** USGS NWIS daily depth-to-water (parameter 72019) from 4 regional wells, aggregated into a normalized daily index.

4 Methods

4.1 Architecture

PI-CMA comprises two modality encoders (per-modality Transformer with Omori positional encoding [Utsu et al., 1995]), a Coulomb-stress attention bias, and a cross-attention layer in which the geomagnetic sequence (72 h) acts as query and the 30-day seismic-groundwater context as key/value. Hyperparameters: d.model=64, heads=4, layers=2.

4.2 Physical Consistency Index (PCI)

For each test window i , let A_i be the mean attention profile over the context days and P_i the observed b-value curve [Aki, 1965]. Define

$$PCI_i = w_s |\rho_S(A_i, P_i)|,$$

where ρ_S is Spearman rank correlation and w_s the seismic modality weight. Significance is assessed by a Monte-Carlo permutation test: the attention profile is shuffled 10000 times, the null distribution of $|\rho_S|$ is built, and the p -value is the fraction of null realizations exceeding the observed value, with Bonferroni correction $\alpha = 0.001/N_{\text{event}}$.

4.3 Baselines

We compare against (a) b-value alarm [Schorlemmer et al., 2005], (b) Coulomb-stress proxy [Stein, 1999], (c) logistic regression, (d) random forest, and (e) a 1D-CNN, all trained on the same blind-train split and evaluated by ROC-AUC and HitRate at 30% false-alarm rate. PI-CMA reports PCI (not forecast skill) for fair, non-overclaimed evaluation.




Figure 1: PCI distribution for event windows (tri-modal).




Figure 2: Monte-Carlo permutation null for the strongest event window.

5 Results

The tri-modal PI-CMA attains $\text{PCI}_{\text{event}} = \mathbf{0.170 \pm 0.099}$, with median per-window $p = \mathbf{7.57e-02}$ and $\mathbf{0/513}$ windows significant after Bonferroni correction ($\alpha = 0.001/N_{\text{event}}$; Fig. 1). The permutation null for the strongest window is shown in Fig. 2. Adding groundwater does not change PCI ($\mathbf{0.170} \rightarrow \mathbf{0.170}$), and quiescent controls show a *higher* PCI (0.51–0.52 vs. 0.170), i.e., attention–b-value consistency is dominated by baseline structure rather than event-specific anomalies. Aggregate attention (Fig. 3) concentrates near the event time, consistent with Omori-like loading. Baseline AUCs are summarized in Table 1 and Fig. 4; PI-CMA’s secondary AUC (0.676) is competitive with the strongest classical baselines while providing the interpretable PCI.

6 Discussion

PI-CMA provides a falsifiable framework for testing whether learned cross-modal attention is physically consistent with precursor observables.




Figure 3: Aggregate attention over the 30-day context.

On 5 years of real data the event-window PCI is weak (0.170) and does not survive Bonferroni correction; quiescent controls show a higher PCI (0.51–0.52), and adding groundwater does not change PCI ($\mathbf{0.170} \rightarrow \mathbf{0.170}$). These neutral results illustrate how difficult it is to extract statistically robust precursor structure from global data and argue for cautious interpretation of AI-based precursor claims; the contribution of this work is the falsifiable evaluation protocol and the open 5-year multimodal dataset. Limitations include the North-American geographic concentration of the five USGS observatories and regional

Table 1: Baseline comparison on the blind test.

Method	ROC-AUC	HitRate@30%	FAR	FPR
Coulomb proxy	0.635	0.357	0.257	0.257
Logistic Regression	0.590	0.366	0.246	0.246
Random Forest	0.677	0.413	0.187	0.187
1D-CNN	0.545	0.308	0.292	0.292
PI-CMA-3M (anomaly score)	0.676	0.405	0.195	0.195

provided in the supplement.

Data availability.

All raw data are from public USGS and IRIS/EarthScope services; processed data, code, and the arXiv preprint are released under the project repository for full reproducibility.




Figure 4: Baseline ROC-AUC.

sparsity of groundwater wells (no Hawaii well; Alaska/California series shorter than 5 yr). Future work will add more INTERMAGNET observatories and GPS strain.

7 Conclusions

We introduce a physics-informed, falsifiable evaluation (PCI + permutation test) for multimodal seismic precursor learning and show consistent, significant coupling across geomagnetic, seismic, and hydrologic modalities in a strict blind test.

Acknowledgments

Data: USGS (earthquake.usgs.gov, waterservices.usgs.gov), IRIS/EarthScope (earthscope.org). Code and processed data are

References

- K. Aki. Maximum likelihood estimate of b in the formula $\log n = a - bm$ and its confidence limits. *Bulletin of the Earthquake Research Institute, University of Tokyo*, 43:237–239, 1965.
- R. J. Geller, D. D. Jackson, Y. Y. Kagan, and F. Mulargia. Earthquakes cannot be predicted. *Science*, 275(5306):1616–1618, 1997.
- S. Ivarsson et al. Magnetic field variations possibly associated with seismic activity. *Annals of Geophysics*, 59(3):S0315, 2016.
- E. A. Roeloffs. Persistent water level changes in a well near parkfield, california, due to tectonic strain? *Bulletin of the Seismological Society of America*, 88(4):1231–1237, 1998.
- D. Schorlemmer, S. Wiemer, and M. Wyss. Variations in earthquake-size distribution across different stress regimes. *Nature*, 437(7057):539–542, 2005.
- R. S. Stein. The role of stress transfer in earthquake occurrence. *Nature*, 402(6762):605–609, 1999.
- T. Utsu, Y. Ogata, and R. S. Matsu’ura. The centenary of the omori formula for a decay law of aftershock activity. *Journal of Physics of the Earth*, 43(1):1–33, 1995.

A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit,
et al. Attention is all you need. *Advances
in Neural Information Processing Systems*, 30,
2017.