PISCES: Physics-Informed Convolutional Autoencoder for Solar-Wind Anomaly Detection and Space-Weather Early Warning

Alison J. March[1,*] (alison.march@colorado.edu), Kevin Lee[2,3,*] (kevinlee69720@g.ucla.edu)

[1]University of Colorado, Boulder [2]University of California, Los Angeles [3]NASA Jet Propulsion Laboratory

*Equal Contribution

NASA 5th Eddy Cross-Disciplinary Symposium, Boulder, Colorado, 2026

Introduction

Problem: Data-driven methods for solar-wind transients anomaly detection at L1 Lagrange point typically only provide a single opaque anomaly score. They are useful as an alarm, but lack the physical attribution for operational diagnostic.

Model: Physics-Informed Solar-wind Convolutional autoEncoder for Space-weather (PISCES) is a physics-informed convolutional autoencoder model trained on NASA OMNI data.

Physics awareness: Embeds seven solar wind physical constraints: (1) magnetic field consistency, (2) Parker-spiral angle, (3) temperature–velocity relationship, and smoothness priors on (4) proton entropy, (5) total pressure, (6) plasma beta, and (7) dynamic pressure.

Explainable Outputs: Decomposes anomaly score into interpretable physical components rather than a scalar metric; provides event-type hints (interplanetary shocks, magnetic ejecta, compression).

Validation: Evaluated on a 2018–2024 held-out set (2.97 million windows; 863 catalog events), PISCES achieves competitive aggregate detection across unsupervised models and delivers interpretable physical decomposition scores.

Data, Methods, and Model

Data and Input Specification

Category	Details
Training Data	NASA OMNI high-resolution solar wind, 2005–2024 (1-min cadence at L1)
Real-time Input	NOAA SWPC RTSW JSON feeds
Evaluation Catalogs	DONKI HSS & IPS; Richardson–Cane near-Earth ICME catalog
Train Split	2005–2015
Validation Split	2016–2017
Test Split	2018–2024
Training Labels	None (unsupervised)
Evaluation Labels	DONKI IPS/HSS and Richardson–Cane ICME

Input Channels (7 Features)

Domain	Channels
Magnetic Field	B_x, B_y, B_z, B_t [nT]
Plasma	n_p [cm^{-3}], V [km s^{-1}], T_p [K]

Preprocessing Pipeline

Step	Method
Missing Data	Forward-fill $\leq$ 2 min; interpolate $\leq$ 15 min; NaN: longer gaps
Transform	$\log(n_p, V, T_p)$ and z-score standardization
Windowing	7 × 60-min windows; 15-min step for training; 1-min step for inference

Physical Scoring Criteria

Class	Signatures
Quiet Solar Wind	Smooth entropy, thermal pressure, magnetic pressure, plasma β
Shock/Compression	Abrupt jumps in total pressure, dynamic pressure, and mass flux
ICME	$\beta < 0.5$; magnetic field rotation; Parker-spiral and T–V deviation

Data, Model, and Training Summary.

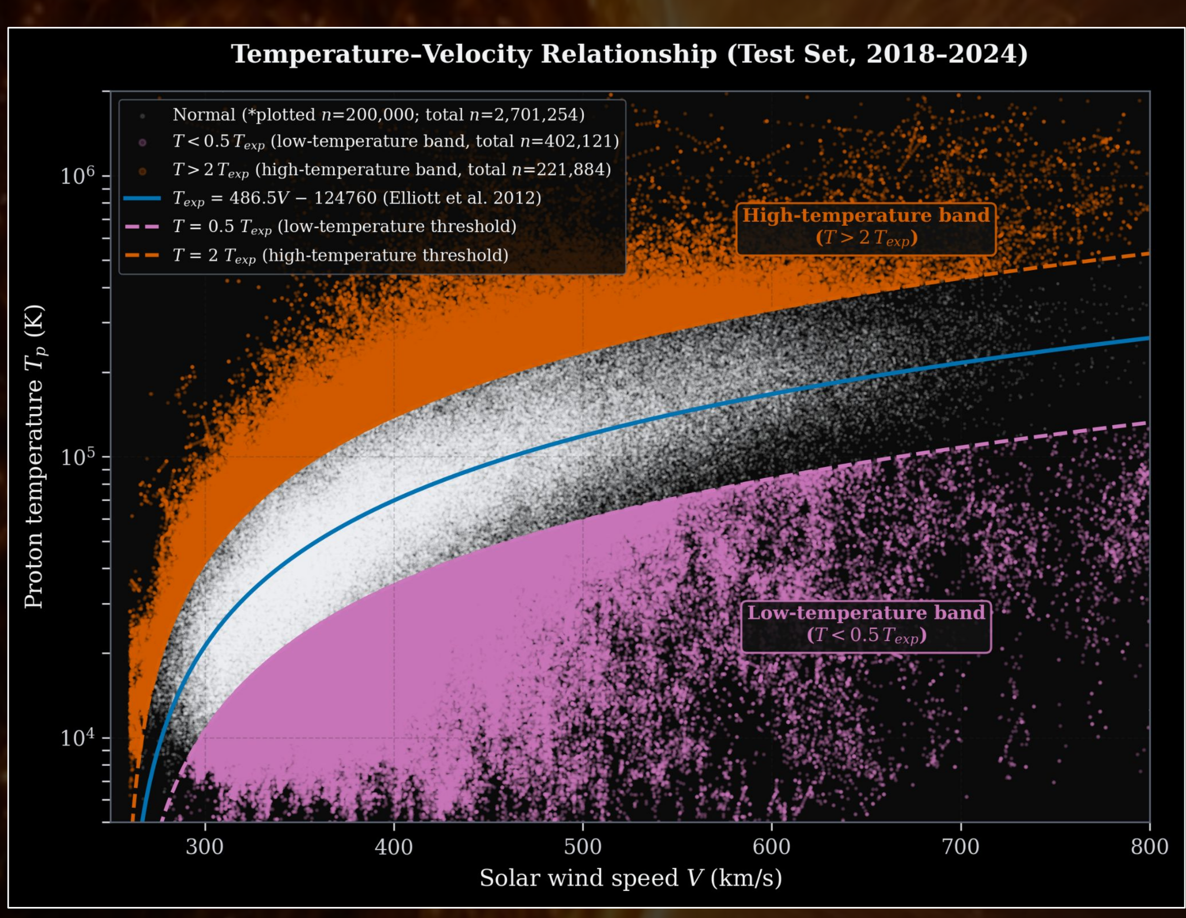

Example Constraint Used: T–V Relationship on Filtered 2018–2024 Test Set (3.3 Million 1-Minute Sample).

The solid blue shows Texp = 486.5V - 124760 K. Dashed lines mark low-temperature and high-temperature thresholds. 81% of minutes lie between thresholds (200K of selected classes plotted), 12% below, and 7% above. The normal points are subsampled for visibility.

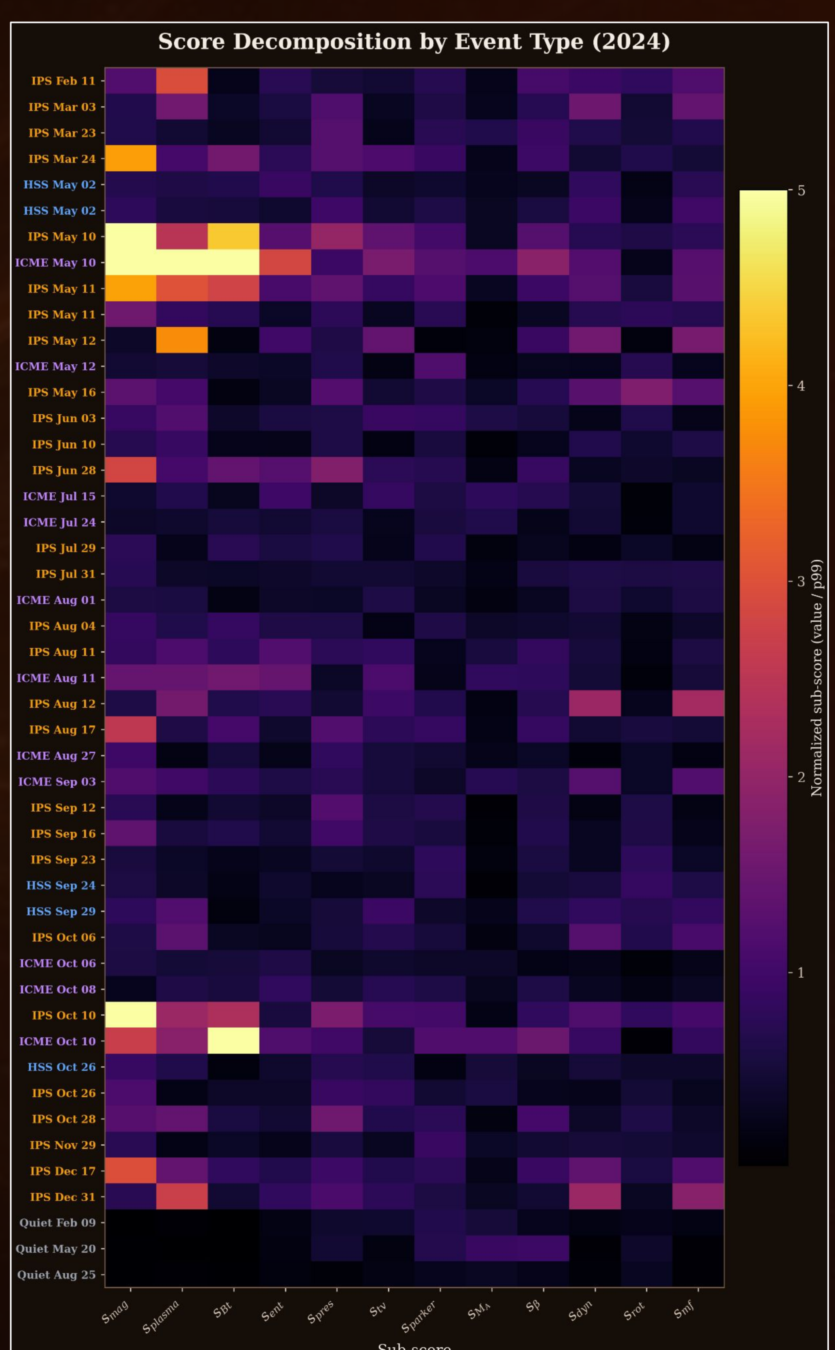

Score Decomposition Heatmap For 2024 Event Types.

Columns show 12 non-residual sub-scores: s_mag, s_plasma, s_Bt, s_ent, s_pres, s_tv, s_parker, s_MA, s_β, s_dyn, s_rot, and s_mf. Differences across event types are shown and consistent with expected physical behavior.

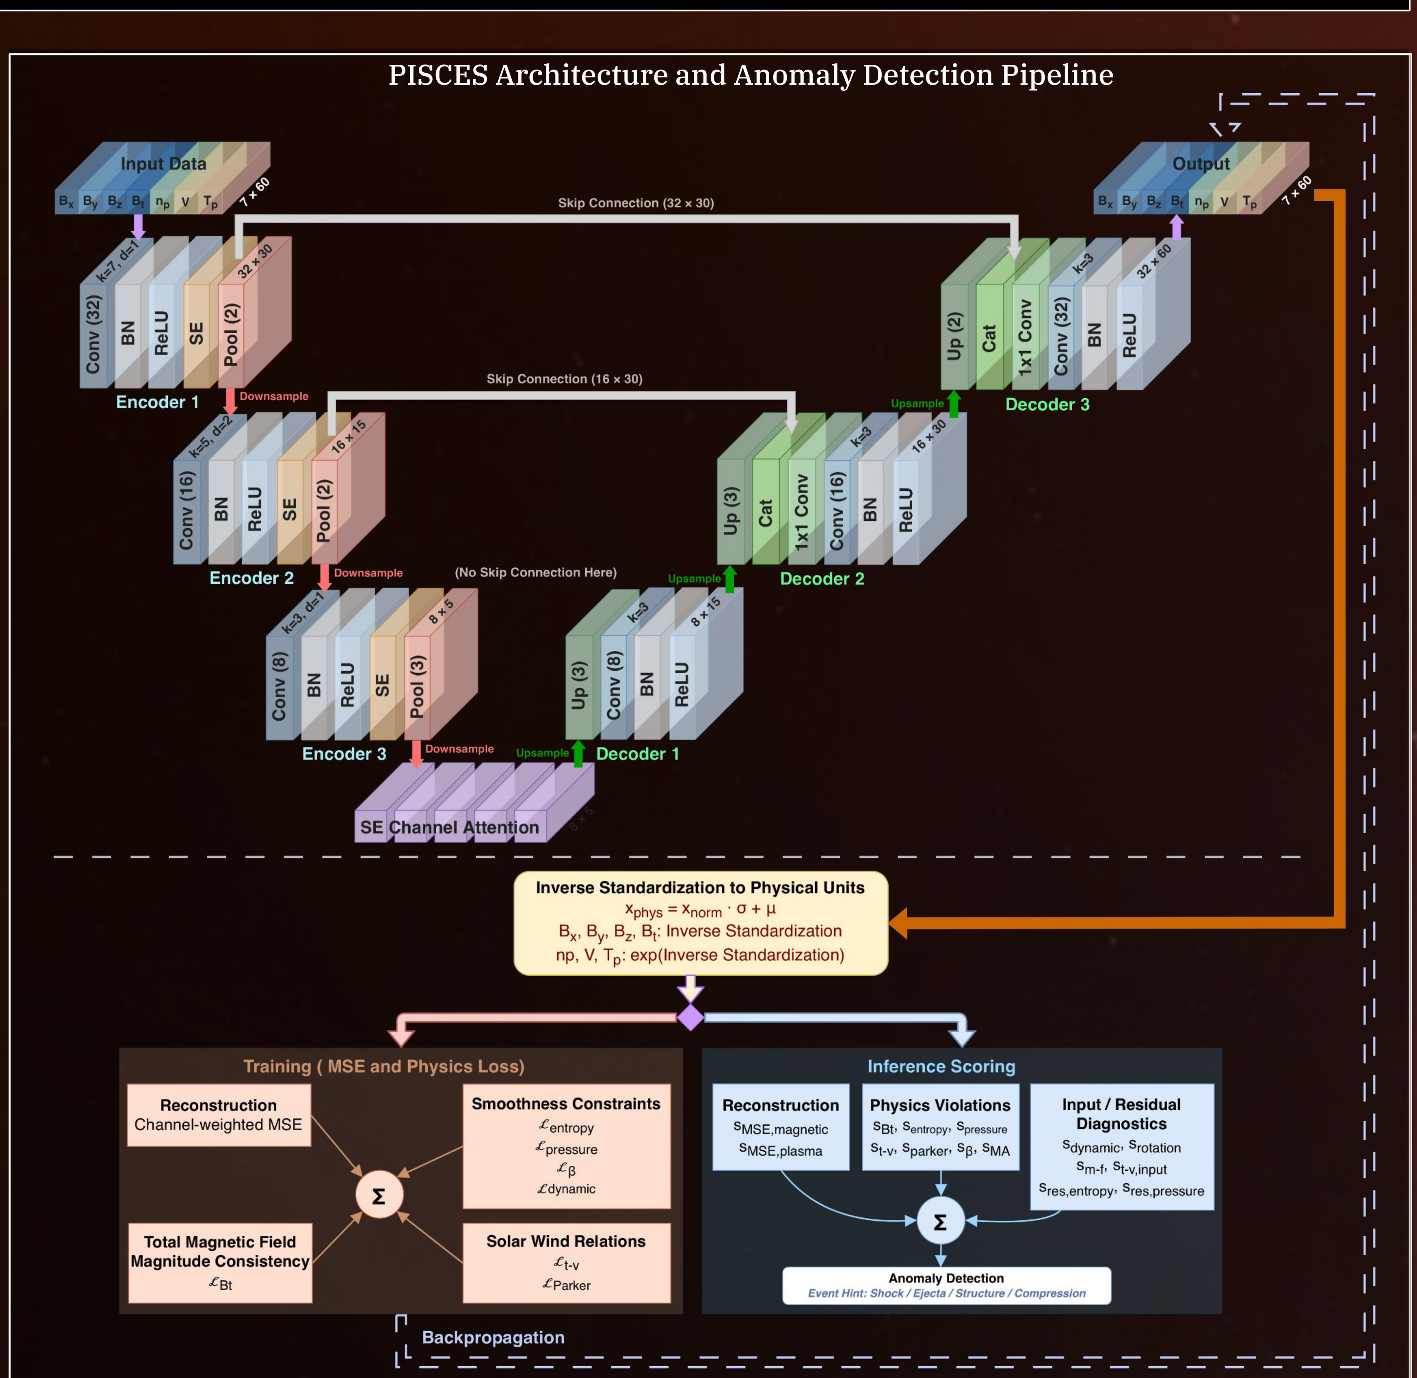

PISCES: Physics-Informed Convolutional Autoencoder For 60-min Windows × 7 Solar-wind Channels.

Top: The convolutional autoencoder (U-Net) uses the encoder (Conv–BN–ReLU–SE–Pool) to compress the input window to the bottleneck, and the decoder reconstructs the window using upsampling and skip fusion.

Bottom: training combines channel-weighted MSE in normalized space with physics losses, while inference combines normalized sub-scores into a smoothed composite anomaly score with event type hints.

Results and Case Studies

Method	PR-AUC	POD	FAR	Precision	HSS
PISCES	**0.349**	0.059	**0.305**	**0.695**	0.092
OC-SVM	0.256	**0.082**	0.328	0.672	**0.124**
PCA	0.228	0.039	0.437	0.563	0.058
Conv-VAE	0.222	0.035	0.440	0.560	0.053
LSTM-AE	0.214	0.038	0.445	0.555	0.056
Max-Jump	0.190	0.054	0.614	0.386	0.068

Performance Benchmark Between Different Models and Methodologies.

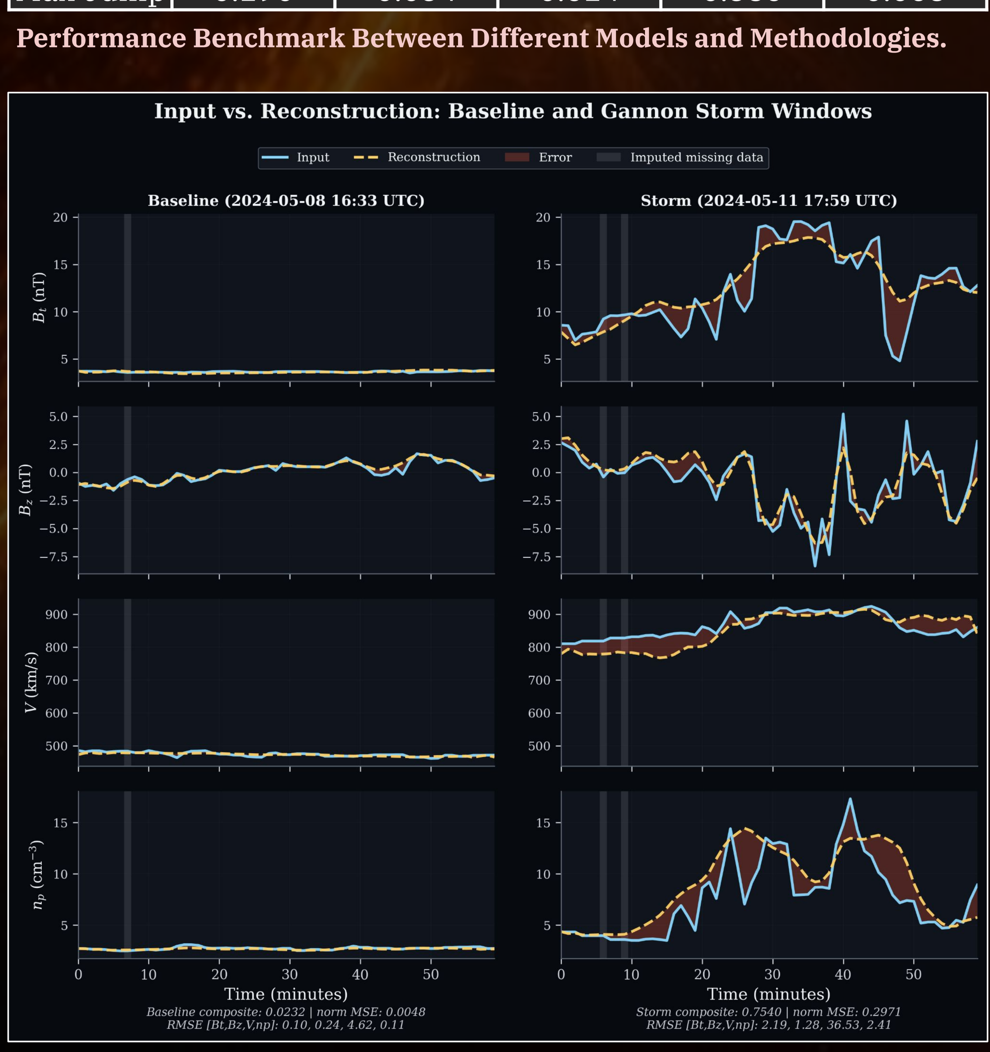

Input vs. Reconstruction for Two Fixed 60-minute Solar-Wind Windows.

Left: Low-anomaly baseline (2024-05-08 16:33 UTC). Right: Gannon storm (2024-05-11 17:59 UTC). Channels signals of B_t, B_z, V, and n_p.

The model reconstructs a low-anomaly reference closely while diverging during the Gannon storm, especially in magnetic and plasma channels with composite score (0.023 vs. 0.754) and normalized MSE (0.0048 vs. 0.297).

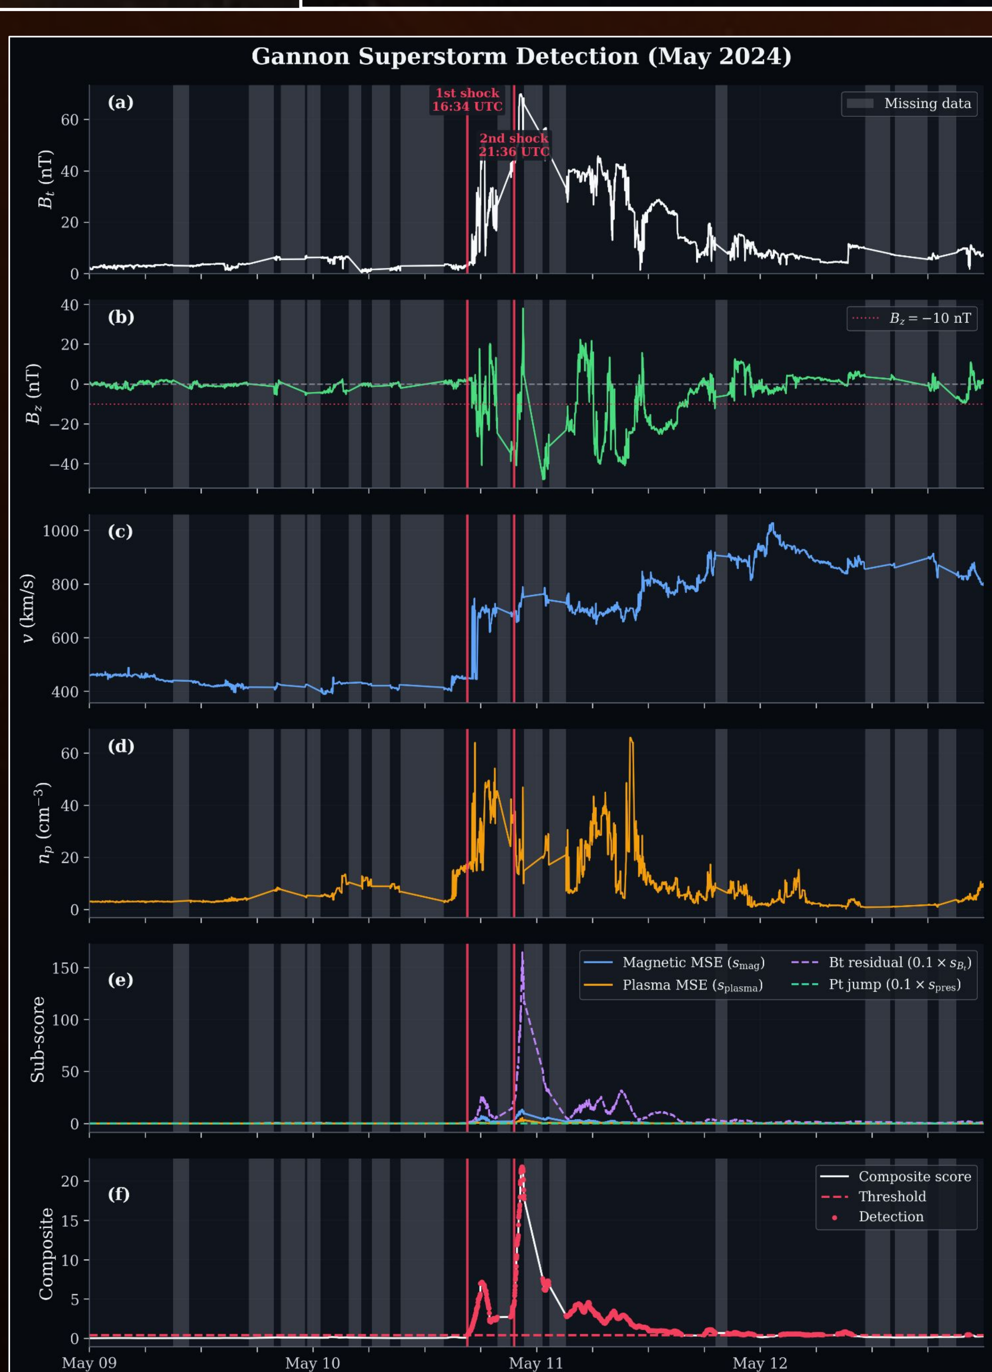

PISCES Response During the May 2024 Gannon Superstorm.

(a) - (d): Preprocessed OMNI solar-wind signals (B_t, B_z, V, and n_p).

(e) - (f): Selected sub-scores and the composite score against the detection threshold value. The anomaly response keeps elevating in the storm, with the strongest peak around the shock interval.

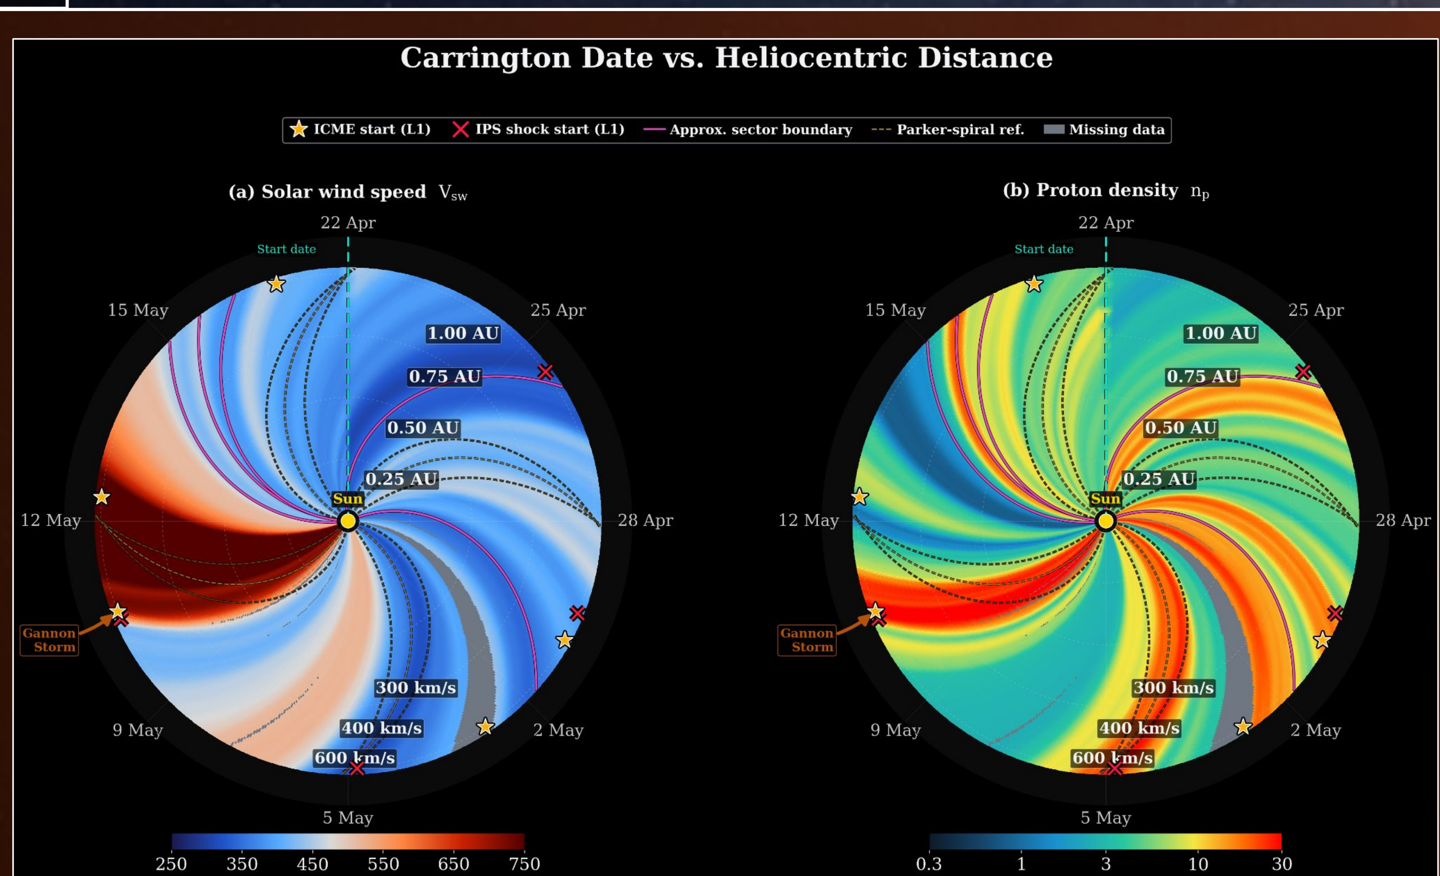

Ballistic Back-mapping of May 2024 Gannon Storm From L1.

Left: Solar Wind Speed back-mapping from L1; Right: Proton Density back-mapping from L1.

Conclusion and Future Work

Conclusion: PISCES decomposes unsupervised anomaly detection from a scalar anomaly score into physical components of shock, ejecta, and compression features. The skip-connection dropout forces the bottleneck to learn from the physical features; the physics constraints anchor the latent space for physical meaning.

Future Work: Multi-scale parallel inference (30, 60, 120 min) for gradual structures. Deployment for continuous monitoring of spacecraft real-time feeds.

PISCES Repo