Machine Learning Concentration Proposal

Predicting Relativistic Electron Flux at LEO from Solar Wind Data using Neural Networks

Student: Parth Kheni (Computer Engineering)
Supervisor: Dr. Luisa Capannolo • luisacap@bu.edu
Team: Individual Project

11/01/2025

The outer radiation belt is a toroidal region surrounding Earth, populated by very energetic charged particles. Due to changes in the solar wind flow impacting Earth's magnetic field, particles can fall into the upper atmosphere and deposit energy there, leading to chemical changes. Additionally, this radiation environment is detrimental to satellites from low-Earth orbit to geostationary orbit and beyond. This project aims to model, through machine learning, the relativistic electron population of the outer belt, with a focus on Relativistic Electron Precipitation (REP).

Relativistic Electron Precipitation (REP) can damage satellites and affect the upper atmosphere. I will build a neural network that predicts > 700 keV electron flux at low-Earth orbit (LEO) using only upstream solar-wind measurements. The model will output both precipitating ($0°$) and trapped ($90°$) flux from POES/MetOp, enabling event-level insight and a potential 30–60 minute heads-up for space-weather operations. Concretely, I will assemble a time-aligned dataset linking 1-minute OMNI solar-wind/IMF(interplanetary magnetic field) parameters to POES/MetOp E4_0 (precipitating) and E4_90 (trapped) with L and MLT; train and compare architectures (a baseline feed-forward network plus LSTM/GRU or 1D-CNN) to predict $\log_{10}$-flux at T_0 using multi-hour lookback windows; evaluate skill across L-shell, MLT, and storm phase with emphasis on peak capture critical to REP; and deliver an interpretable, reproducible pipeline with sensitivity analyses over inputs, history window length, and downsampling choices.

Data & preprocessing. I will use OMNI (flow speed, density, pressure, IMF components) at 1-minute cadence; POES/MetOp > 700 keV (E4) at 2-second cadence; and satellite position data. I will downsample POES to 1 minute with peak-preserving statistics (comparing mean, median, trimmed mean, and quantiles) to avoid smearing REP bursts, align a 4-hour history (5-minute stride) for each solar-wind feature, and append instantaneous L and MLT encoded as $\sin / \cos$. Targets are $\log_{10}$(E4_0) and $\log_{10}$(E4_90).

Modeling. Inputs will be the solar-wind time series plus L and MLT encodings, with optional geomagnetic indices (SYM-H, AL). I will train a feed-forward model on flattened windows and temporal models (LSTM/GRU or 1D-CNN), using the **Adam optimizer** with a log-MSE loss, early stopping, and time-aware splits (rolling 30-day or month-by-month blocks) to prevent leakage. Interpretability will rely on permutation importance across channels and lags, along with **input group sensitivity tests** (e.g., removing driver subsets) and history-window length studies (2–48 hr).

Timeline to May 2026

- **Nov–Dec 2025:** Finalize supervisor sign-off; complete dataset build; choose POES downsampling.
- **Jan–Feb 2026:** Train/tune baseline FFN; fix time-segmented train/val/test splits; add MLT encodings.
- **Mar 2026:** Implement temporal models (LSTM/GRU or 1D-CNN); run hyper-parameter search.
- **Apr 2026:** History-window length checks; calibration; event-peak analysis; L-binned error.
- **May 2026:** Finalize results; prepare submission (PDF report, slides, approvals).

This work directly addresses the concentration outcomes. On the data side, I will make and synchronize multi-source time series, perform resampling and windowing, and apply strict quality control. Algorithmically, I will select and compare neural architectures, design temporal features, and optimize and assess performance. Tooling will center on PyTorch, NumPy, and Pandas with a reproducible pipeline and optional GPU training. For communication, I will read and explain relevant ML literature and present results with clear metrics and figures. I will also have opportunities to present progress in our research group (from undergrads to faculty) and, schedule permitting, attend a scientific conference (e.g., GEM in June 2026).

Validation and Assessment. The primary metric is RMSE in $\log_{10}$-space, also reported as multiplicative error (10^{RMSE}). Secondary metrics include Pearson r; an event-peak capture rate that checks timing (±1–5 minutes) and amplitude (±20–50%); and skill relative to persistence and climatology, broken down by by L, MLT, storm phase, and 0° vs. 90°. Success means achieving test RMSE $\leq 0.40 \log_{10}$ (approximately $\times 2.5$), $r \geq 0.85$, capturing at least 60% of major peaks within tolerance, and outperforming persistence/climatology. If results are partial—e.g., timing is good but peaks are underestimated—I will adjust inputs, architecture, or loss. If the model fails—no improvement over persistence or strong L-dependent bias—I will revise inputs, splits, and models.

References

Selected neural-network radiation-belt studies and internal group materials inform the lookback design, metrics, and evaluation; full citations will be included in the final report and slides.