Identification and Classification of Relativistic Electron Precipitation Events at Earth Using Supervised Deep Learning

Luisa Capannolo[1], Wen Li[1†], and Sheng Huang[1†]

[1] Center for Space Physics, Boston University, Boston, MA, USA

† These authors have contributed equally to this work and share last authorship

*** Correspondence:**
Luisa Capannolo
luisacap@bu.edu

Keywords: electron precipitation, wave-particle interactions, current sheet scattering, space sciences, supervised classification, LSTM, deep learning, supervised deep learning

Abstract

We show an application of supervised deep learning in space sciences. We focus on the relativistic electron precipitation into Earth's atmosphere that occurs when magnetospheric processes (wave-particle interactions or current sheet scattering, CSS) violate adiabatic invariants of trapped radiation belt electrons leading to electron loss. Electron precipitation is a key mechanism of radiation belt loss and can lead to several space weather effects due to its interaction with the Earth's atmosphere. However, the detailed properties and drivers of electron precipitation are currently not fully understood yet. Here, we aim to build a deep learning model that identifies relativistic precipitation events and their associated driver (waves or CSS). We use a list of precipitation events visually categorized into wave-driven events (REPs, showing spatially isolated precipitation) and CSS-driven events (CSSs, showing an energy-dependent precipitation pattern). We elaborate the ensemble of events to obtain a dataset of randomly stacked events made of a fixed window of data points that includes the precipitation interval. We assign a label to each data point: 0 is for no-events, 1 is for REPs and 2 is for CSSs. Only the data points during the precipitation are labeled as 1 or 2. By adopting a long short-term memory (LSTM) deep learning architecture, we developed a model that acceptably identifies the events and appropriately categorizes them into REPs or CSSs. The advantage of using deep learning for this task is meaningful given that classifying precipitation events by its drivers is rather time-expensive and typically must involve a human. After post-processing, this model is helpful to obtain statistically large datasets of REP and CSS events that will reveal the location and properties of the precipitation driven by these two processes at all *L* shells and MLT sectors as well as their relative role, thus is useful to improve radiation belt models. Additionally, the datasets of REPs and CSSs can provide a quantification of the energy input into the atmosphere due to relativistic electron precipitation, thus providing valuable information to space weather and atmospheric communities.

1 Introduction

The radiation belt environment is highly dynamic and it is governed by acceleration, transport and loss processes (e.g., Li and Hudson, 2019; Reeves et al., 2003). One of the loss mechanisms is electron precipitation (EP), which occurs when the conservation of the first adiabatic invariant is violated (e.g., Horne and Thorne, 1998; Shulz and Lanzerotti, 1974): electrons are no longer trapped by the Earth's magnetic field and fall into the upper atmosphere. Not only electron depletion is important in the radiation belt evolution in time and flux, but electron precipitation is also known to drive many atmospheric effects related to space weather. Multiple studies have indeed associated conductivity variations and atmospheric chemistry changes (potentially leading to ozone reduction) with electron precipitation (Duderstadt et al., 2021; Fytterer et al., 2015; Khazanov et al., 2018; Meraner & Shmidt, 2018; Mironova et al., 2015; Robinson et al., 1987; Sinnhuber et al., 2021; Tyssøy et al., 2021; Yu et al., 2018).

It is well understood that electron precipitation can occur as a result of interactions between plasma waves existing in the magnetosphere and the trapped electron population in the radiation belts (e.g., Millan and Thorne, 2007; Thorne, 2010). Electrons can also be lost if the magnetic field line around which they gyrate is stretched away from Earth or undergoes a significant geometry variation such that the curvature radius of the field line is comparable to the gyroradius of the electrons (e.g., Buchner & Zelenyi, 1989; Dubyagin et al., 2021; Sergeev et al., 1983, 1993). This process is called field line curvature scattering or current sheet scattering (CSS). Under these conditions, the field line no longer traps the electrons, and these electrons can precipitate into the atmosphere. The location where precipitation occurs (called isotropic boundary, IB) depends on electron energy (Capannolo et al., 2022; Yahnin et al., 2016; 2017). This phenomenon has also been widely studied for protons (Dubyagin et al., 2018; Ganushkina et al., 2005; Gilson et al., 2012; Liang et al., 2014).

A comprehensive understanding of which mechanism (waves or CSS) dominates the electron precipitation and thus the energy input into the Earth's atmosphere is still under active research. Given the Earth's magnetic field geometry, one would expect that on the dayside and at low L shells CSS does not contribute much, but more quantitative studies are still needed. Overall, while wave-driven precipitation can occur at all MLT (magnetic local time) sectors, CSS-driven precipitation is indeed primarily observed over 20–04 MLT (Yahnin et al., 2016; 2017), and overlaps with precipitation driven by waves (for the most part, electromagnetic ion cyclotron waves, EMIC) in the midnight sector (Capannolo et al., 2022).

These studies use data from the constellation of satellites called POES (Polar Orbiting Environmental Satellites) and MetOp (Meteorological Operational), described in Section 2. An example of a wave-driven (REP, relativistic electron precipitation) event is shown in Figure 1a, together with an example of a CSS-driven (CSS) event (Figure 1b). REP events show enhancements in the relativistic (>700 keV) precipitating electron flux (solid red line) and the precipitation is rather isolated (gray region) in space (L shell) with little/no precipitation around the main event. This region generally matches the location where the wave-particle interaction is efficient to violate an adiabatic invariant. CSS events, instead, show an energy-dependent precipitation with higher energy electrons precipitating at lower L shells than lower energy electrons (Figure 1b; green, black, and blue solid lines). This is a direct result from the fact that the electron gyroradius depends on electron energy: higher energy electrons have a larger gyroradius, thus are lost by a stretched magnetic field line at distances closer to Earth (smaller L shells) than lower energy electrons. Given such a distinct pattern of precipitation, we can distinguish the precipitation drivers.

So far, existing analyses aiming to distinguish the precipitation drivers have either focused on a limited time span (Yahnin et al., 2016; 2017) or on a limited MLT sector (Capannolo et al., 2022). Identifying precipitation events and visually inspecting their precipitation patterns to categorize their driver (waves or CSS) is a rather time-expensive task. Algorithms that find relativistic electron precipitation events (based on count rate or flux thresholds) exist in literature (e.g., Capannolo et al., 2022; Gasque et al., 2021; Shekhar et al, 2017), but they do not include the distinction between wave-driven precipitation and CSS-driven precipitation, which is a much more complex task to perform using algorithms. The goal of this work is to take advantage of deep learning techniques not only to find precipitation events, but also to categorize them into wave-driven (REP) and CSS-driven (CSS) events. We use the dataset of precipitation events analyzed in Capannolo et al. (2022), which were visually classified between wave-driven (REPs) and CSS-driven (CSSs) precipitation events (details in Capannolo et al., 2022). This work is an example of an application of supervised deep learning classification in space sciences that is able to provide a large dataset of precipitation events classified by driver (waves or CSS) after an initial manual classification of events.

2 Satellite Data Description

We use data from the POES and MetOp network of sun-synchronous satellites in polar orbits at ~800–850 km of altitude (Evans and Greer, 2004). The Medium Energy Proton and Electron Detector (MEPED) provides electron (and proton) flux in 3 integral channels with cutoff energies of >30 keV (E1), >100 keV (E2), and >300 keV (E3) (Rodger et al., 2010). The P6 proton channel is designed to measure >6.9 MeV protons, however, it is also sensitive to electrons at >700 keV (Yando et al., 2011) in absence of high energy protons. Thus, we use the P6 channel as a fourth virtual electron channel, E4 (Green, 2013). Additionally, each satellite is equipped with two telescopes: one oriented along zenith (0° telescope) and one perpendicular to it (90° telescope), both with full field-of-view angle of 30°. At mid-to-high latitudes, the 0° telescope provides measurements of electrons precipitating deep into the loss cone and the 90° telescope provides observations of trapped electrons. Strong precipitation typically occurs when the flux observed by the 0° telescope approaches the flux observed by the 90° telescope, indicating that a large percentage of trapped electrons are precipitating. Precipitation events are marked in gray in Figure 1, 2 and 3, and highlighted in brown (REP) and blue (CSS) in Figure 4. The resolution of the electron flux is 2 seconds, and the constellation of satellite covers a rather broad *L*-shell range and MLT sectors. Typical observations of POES/MetOp are shown in the Supplementary Figure 1. Each panel shows ¼ orbit of a POES/MetOp satellites (one pass through the radiation belts) and highlights the significant variability of flux during the satellite trajectory.

3 Methods

In this section, we describe how we prepared the dataset of precipitation events in order to obtain a well-performing model. We also describe the model architecture and how it was decided, as well as how we trained the deep learning model.

3.1 Dataset Preparation

Capannolo et al. (2022) analyzed relativistic electron precipitation events observed by POES/MetOp from 2012 to 2020 over 22–02 MLT and classified these events between those driven by waves (called REP events in this work) from those driven by CSS (CSSs hereafter) using their characteristic precipitation profile (Figure 1). Note that this dataset was obtained after careful event classification: only events that clearly belonged to either category (REP or CSS) were considered, while ambiguous precipitation events were carefully discarded. More details on the classification are provided in Capannolo et al. (2022). In this work, we use this dataset of precipitation events classified over 22–02 MLT with additional preprocessing to improve the model performance as explained below.

Our goal is to build a dataset of precipitation events randomly stacked one after the other. We consider all four POES/MetOp electron channels and the two look directions (0° and 90°) for a total of 8 inputs at a given time. The model output (or target) is the data point class (or label, used interchangeably hereafter): 0 is for no-event, 1 is for REP, and 2 is for CSS. Given one event, the data points are labeled as 1 or 2 during the precipitation (gray regions of Figure 1) and the adjacent data points (to the left and right of the event) are labeled with 0. Fluxes ≤ 0 for all channels are set to 0.01 (100) $s^{-1}cm^{-2}sr^{-1}$ for the 0° (90°) telescope measurements (negative values in POES/MetOp data indicate unreliable flux measurements). We apply the natural logarithm to the fluxes and normalize the train dataset.

As shown in Supplementary Figure 1, each pass through the radiation belts highlights a significant flux variability observed by POES/MetOp, while the precipitation events are rather short-lived (< 30–60 seconds). As a result, if we use the full day of data when a given REP/CSS event occurs, we will obtain a label of mostly zeroes (no-event) and only a few data points at 1 or 2 (indicating the REP/CSS). This would make the full dataset of stacked events extremely imbalanced, where only a few percent of the labels are non-zero. With such dataset, the deep learning model is unable to perform well and identify only the no-events correctly. In order to overcome this obstacle, we consider a much shorter window of data for each event: given one event, we label the data points during precipitation with 1 or 2, but label with 0 only the data points adjacent to the left and right of the event such that the total number of data points is 50. In this way, we have windows of 50-point-long for each event which we stack one after the other in a random order. Additionally, we ensure that no other nearby events were occurring within the 50-point-long window such that in this window there is only one type of non-zero label (either 1 or 2). Note that if two events of different classes are adjacent to each other, we rule out both. Instead, if two REP events are adjacent to each other within the 50-point-long window, we widen the label of 1 to include both to ensure that in each 50-point-long window, there is only one continuous non-zero label. For the CSS events, we also manually extended the boundary of the precipitation events to include the full energy dispersion observed by POES/MetOp because we do not limit ourselves to the E4 precipitation alone (as done in Capannolo et al., 2022). This ensures that the full precipitation pattern (from low to high electron energy) is identified as a CSS event and used to train the model. Using the boundaries as in Capannolo et al. (2022) worsens the model performance because the full extent of the energy-dependent pattern is not correctly learned by the model. We show a portion of the dataset in Figure 2: panel a) indicates the label and panel b) shows the electron flux for all energy channels and look directions, where the precipitation events are highlighted in gray.

In order to augment our dataset and provide the model with a wider variety of precipitation patterns, we also mirror each precipitation event about its main axis. This does not introduce data redundancy since each precipitation event (either mirrored or not) carries a meaningful information. In other words, a REP/CSS event can be directly observed by a POES/MetOp satellite following its actual trajectory (e.g., from low to high L shells), but the precipitation pattern would still be observed (though symmetrically) if the same POES/MetOp satellite was travelling along its opposite orbit (e.g., from high to low L shells) through the precipitation region at the same time. Note that this is possible since we are only interested in the profile of the precipitation (i.e., flux evolution as a function of dataset index) and not its temporal evolution. By using this methodology, we obtain a dataset of 460 REPs and 348 CSSs for a total dataset length of 40,400 data points. Although only ~20% of the data points are labeled with 1 or 2 (making this dataset still imbalanced with respect to the 0 class), the REP and CSS classes are approximately balanced (~10% data points are REPs and ~8% data points are CSSs) and the model is able to identify correctly no-events, REPs and CSSs as we show in the following sub-sections.

3.2 Model Structure and Training

We adapt a long short-term memory (LSTM; Hochreiter and Schmidhuber, 1997) architecture (a type of artificial recurrent neural network, RNN; Rumelhart et al., 1986a) for the deep learning model because it retains input information at much earlier time steps, making it more efficiently than RNNs for problems that treat time series. As a matter of fact, the problem of our work is a time series classification. Although the time variable is not explicitly used, it is instead intrinsically represented by the shape of the precipitation. It is indeed the evolution of the precipitation pattern (isolated *vs.* energy-dependent) that differentiates between the two drivers of precipitation, as mentioned in Section 1.

The input format required by LSTM is a tensor, which is composed of a stack of snapshots of the dataset identified by a sliding window with stride 1 and length 7. The label in each snapshot is assigned as the most probable one (*i.e.,* if the majority of data points have label of 0, the label assigned to that snapshot is also 0) and is *one-hot* encoded. The length of 7 is set after trying different sliding window lengths and choosing the one that provided the best model performance.

The metrics we use are those of a standard classification problem and we focus on the F1 score (calculated as the weighted average of the precision and recall; it expresses how many events the classifier identifies correctly quantifying also how many are missed or mislabeled), the AUC (area under the ROC (Receiver Operating Characteristic) recall *vs.* false-positive-rate curve) and the AUPRC (area under the precision *vs* recall curve). We perform a k-fold cross validation with k=10: the whole dataset is split into 10 portions of which one is used as a test set and the remaining 9 are used as training set. We also consider a validation set that is 15% of the training set in each k-fold. The k-fold cross validation consists in training the model on k different datasets (described above) and estimating the model performance for each of the k iterations. The final model performance is the average of the k performances and the final model weights are obtained by training the model on the whole dataset (with the exception of 15% of the dataset used for testing purposes). During training, we use early stopping (with patience of 10 epochs) on the AUC calculated for the validation dataset.

4 Model Performance

We tried different model configurations, all made of a LSTM layer followed by a fully connected (*i.e.,* dense) layer, ending with a dense layer of 3 neurons that outputs one predicted class. There are two dropout layers (with 0.5 dropout rate) after the LSTM layer and after the first dense layer. We validated each model configuration using the k-fold cross-validation (mentioned above) and we selected the model configuration with the highest F1 score, AUC and AUPRC. Out of all the configurations we tried (64 LSTM cells + 256 dense cells; 128 LSTM cells + 128 dense cells; 128 LSTM cells + 256 dense cells; 64 bidirectional LSTM cells + 256 dense cells; 64 bidirectional LSTM cells + 64 bidirectional LSTM cells + 128 dense cells) the model with the best performance is the one with a layer of 64 bidirectional LSTM cells followed by a fully connected layer of 256 cells (total number of free parameters is 71,171). The metrics resulting from the k-fold cross-validation for this model are: F1 ~ 0.948, AUC ~ 0.995, and AUPRC ~ 0.990. Note that the performance among the different model configurations is similar and differs only on the second or third decimal figure. Table 1 in the Supplementary Material shows the performance scores (F1, AUC, AUPRC) resulting from the k-fold cross-validation for each architecture tested. As an example, Supplementary Figure 2 (panels a–e) shows the metrics as a function of epoch for the k=3 fold. Panel f) shows the confusion matrix averaged from all the confusion matrices of each k-fold: the highest values are focused along the diagonal, indicating that the model performs well in assigning the correct class to each snapshot.

To highlight that the model appropriately identifies and classifies precipitation events, we show in Figure 3 three examples of how the model performs on three portions of the test dataset. Panels a, c, and e present the model (solid) and original (dashed) labels and panels b, d, f show the electron fluxes in a similar format as Figure 2. The precipitation events (originally assigned) are highlighted in gray and their associated class is reported in the panels a, c, e. Not only the model identifies all precipitation events, but each event is categorized in the class originally assigned. Note that the indices where the labels are non-zero only indicate that nearby that region the probability of finding an event is higher than the probability of a no-event, but these indices do not necessarily represent the exact precipitation event boundaries (as the original class does). Nevertheless, the labels predicted by the model are in good agreement with the original location and class of the events highlighted in gray. The model labels seem to be shifted to the left by a few data points compared to the original classes, due to the fact that we assign a class to each snapshot of length 7 (described in Section 3.1). In other words, the very first snapshot is classified with the most probable label in the first 7 data points. As the sliding window progresses with stride 1, each label is associated with the following 7 data points resulting in anticipating the snapshot classification.

4.1 Model Application on Several Days of POES/MetOp data: Preliminary Results

As we showed in Section 3.1, the dataset used for training has been significantly shrunk to only 50 data points for each precipitation event observed by POES/MetOp. In this section, we explore the model performance on longer time periods (full day of POES/MetOp data, the significant flux variability of which is shown in Supplementary Figure 1) to test its generalization ability.

We apply the model to several POES/MetOp days and show the results in Figure 4 and Supplementary Figure 3. Each panel in these figures is from a different date and none of the events shown belong to the dataset prepared in Section 3.1 (they are all out-of-sample). Here, we are only considering events occurring in the outer radiation belt, thus we filter out any events occurring at $L < 2.5$ or $L > 8.5$ (L is expressed using the International Geomagnetic Reference Field, IGRF, model in POES/MetOp data). The panels on the left column of Figure 4 show REP events (highlighted in brown), whereas the events on the right column are CSSs (highlighted in blue). This classification is accurate because the classified REPs indeed show isolated E4 precipitation, while the classified CSSs display an energy-dependent precipitation. During REP events (Figures 2, 3, 4), although the low-energy electrons (E1, E2 and E3 channels) appear to precipitate as well, their flux is likely the result of proton contamination, which is known to affect the electron measurements onboard POES/MetOp satellites (e.g., Capannolo et al. 2019, 2021; Evans and Greer, 2000; Yando et al., 2011). Note again that the location where these events are identified by the model differs from the exact event location by a few data points. This is not a major concern as this shift appears to be systematic and can be corrected in the post-processing by shifting the predicted model class by a few data points.

On the contrary, Supplementary Figure 3 shows examples when the model does not perform very well and identifies two adjacent precipitation events belonging to different classes (panels a and b), mislabeled events (panel c) or false positive events (panel d). The cases in panel a) only last one data point and could be potentially disregarded since the model does not identify a long enough non-zero label. The event in panel d) shows a precipitating E4 flux that is higher than the others, which could indicate a potential issue in the recorded POES/MetOp data. Events in panels b) and c) instead must be appropriately ruled out or inspected further (e.g., what is the probability of each class? Is the probability of the CSS class comparable to that of the REP?). Handling false positives is beyond the scope of this work and we are aware that post-processing on the model output is needed before using these results for scientific research. The post-processing should rule out events lasting only one data point, adjacent events belonging to different non-zero classes, and events in the South Atlantic Anomaly, as well as improving the *L* shell calculation for each event (using Tsyganenko models such as the T89 (Tsyganenko, 1989) or T05 (Tsyganenko and Sitnov, 2005)) used to consider events occurring only in the outer radiation belt.

5 Conclusions and Discussion

In this work, we showed an example of an application of supervised deep learning to space sciences. Understanding when, where and why relativistic electrons precipitate into the Earth's atmosphere has a longstanding relevance for a variety of reasons (from improving our knowledge on plasma dynamics to study the space weather impacts of electron precipitation). In this work, we focused specifically on relativistic electron precipitation. Our goal was to classify the relativistic electron precipitation events depending on their spatial precipitation pattern, which in turn corresponds to their magnetospheric driver (waves or current sheet scattering). We used data from the POES/MetOp constellation of low-Earth-orbit satellites. Our task was supervised because we used the list of events studied by Capannolo et al. (2022), which were visually classified. Note that these events were classified only in a limited MLT sector (22–02); however, their MLT value was not used as input in the model, and in fact, our model is able to identify precipitation events at any MLT.

The dataset preparation was key to obtain a satisfying model performance. By considering only a short time window around each event instead of the full day of POES/MetOp data, using non-zero labels to indicate REPs (class of 1) or CSSs (class of 2) and labels at 0 to indicate the no-event, and including electron fluxes observed at different energies and look directions, we were able to obtain an appropriate dataset to use for training. We found that the LSTM architecture is suitable for identifying precipitation events and classifying them by precipitation pattern given its ability to consider the data history (in our case the precipitation pattern profile evolution along the satellite trajectory).

Our model is composed of one layer of 64 bidirectional LSTM cells, one layer of 256 fully connected neurons, and one layer of 3 dense cells. The inputs are the electron fluxes at different energies and look directions, and the output is the class of each data point. We obtained the model metrics (F1 ~ 0.948, AUC ~ 0.995, and AUPRC ~ 0.990) by conducting a k-fold cross-validation (k=10). Our model is able to learn the dataset properties correctly. The model is not only able to identify the electron precipitation events, but it also appropriately classifies them by their drivers.

Since the dataset used for training and testing purposes has been specifically designed to obtain a good model performance, it shows less variability than that typically observed by POES/MetOp over an entire orbit. Nevertheless, our model is still able to identify and classify the precipitation events when applied to a full day of data (Figure 4), though some false positives might still be identified (Supplementary Figure 3). Post-processing of these results is needed before being able to use the model outputs for scientific research; however, this is beyond the scope of this paper and left for future investigation. Once the post-processing routine is developed, this model could be easily used as a tool to produce lists of relativistic electron precipitation events in a very short amount of time, overcoming the complex task of developing deterministic algorithms based on flux thresholds to delineate the precipitation patterns and the time-expensive task of visually classifying these events by driver. In this way, we would be able to extend the study conducted in Capannolo et al. (2022) to the whole MLT range and statistically investigate on where the CSS effects should be considered for radiation belt and precipitation modeling, as well as compare them with the precipitation driven by waves. Such event dataset would also potentially open additional avenues of machine learning applications to space sciences; for example, from a space weather point of view, we could investigate if the electron precipitation events can be predicted by using solar images, solar wind data and/or geomagnetic indices.

6 Conflict of Interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

7 Author Contributions

L. C. conducted the core of this work (dataset preparation, model development, model training, etc.). W. L. and S. H. contributed equally by offering feedback during the preparation of this work, and they share the last authorship.

8 Funding

This research is supported by the NSF grants AGS-1723588 and AGS-2019950, the NASA grants 80NSSC20K0698 and 80NSSC20K1270, and the Alfred P. Sloan Research Fellowship FG-2018-10936. SH would like to acknowledge the NASA FINESST Award 80NSSC21K1385.

9 Acknowledgments

L. C. would like to acknowledge M. Capannolo for the insightful conversations on machine and deep learning. L. C. also acknowledges her teammates in the 2020 Heliophysics hackweek (B. Tremblay, A. K. Tiwari, A. Hu, B. Thompson, S. Shekhar, M. Shumko, S. Forsyth, J. Li, and D. Linko).

10 References

Büchner, J., and Zelenyi, L. M. (1989), Regular and chaotic charged particle motion in magnetotaillike field reversals: 1. Basic theory of trapped motion, J. Geophys. Res., 94(A9), 11821–11842, doi:10.1029/JA094iA09p11821.

Capannolo, L., Li, W., Ma, Q., Shen, X.-C., Zhang, X.-J., Redmon, R. J., et al. (2019). Energetic electron precipitation: Multievent analysis of its spatial extent during EMIC wave activity. *Journal of Geophysical Research: Space Physics*, 124, 2466– 2483. https://doi.org/10.1029/2018JA026291

Capannolo, L., Li, W., Millan, R., Smith, D., Sivadas, N., Sample, J., & Shekhar, S. (2022). Relativistic Electron Precipitation Near Midnight: Drivers, Distribution, and Properties. *Journal of Geophysical Research: Space Physics*, 127, e2021JA030111. https://doi.org/10.1029/2021JA030111

Capannolo, L., Li, W., Spence, H., Johnson, A. T., Shumko, M., Sample, J., & Klumpar, D. (2021). Energetic electron precipitation observed by FIREBIRD-II potentially driven by EMIC waves: Location, extent, and energy range from a multievent analysis. *Geophysical Research Letters*, 48, e2020GL091564. https://doi.org/10.1029/2020GL091564

Dubyagin, S., Apatenkov, S., Gordeev, E., Ganushkina, N., & Zheng, Y. (2021). Conditions of loss cone filling by scattering on the curved field lines for 30 keV protons during geomagnetic storm as inferred from numerical trajectory tracing. Journal of Geophysical Research: Space Physics, 126, e2020JA028490. https://doi.org/10.1029/2020JA028490

Dubyagin, S., Ganushkina, N. Y., & Sergeev, V. A. (2018). Formation of 30 keV proton isotropic boundaries during geomagnetic storms. Journal of Geophysical Research: Space Physics, 123, 3436–3459. https://doi.org/10.1002/2017JA024587

Duderstadt, K. A., Huang, C.-L., Spence, H. E., Smith, S., Blake, J. B., Crew, A. B., et al. (2021). Estimating the impacts of radiation belt electrons on atmospheric chemistry using FIREBIRD II and Van Allen Probes observations. Journal of Geophysical Research: Atmospheres, 126, e2020JD033098. https://doi.org/10.1029/2020JD033098

Evans, D. S., and M. S. Greer (2004), Polar Orbiting Environmental Satellite Space Environment Monitor-2: Instrument Descriptions and Archive Data Documentation, NOAA Tech. Mem. 93, version 1.4, Space Weather Predict. Cent., Boulder, Colo.

Fytterer, T., Mlynczak, M. G., Nieder, H., Pérot, K., Sinnhuber, M., Stiller, G., and Urban, J.: Energetic particle induced intra-seasonal variability of ozone inside the Antarctic polar vortex observed in satellite data, Atmos. Chem. Phys., 15, 3327–3338, https://doi.org/10.5194/acp-15-3327-2015, 2015.

Ganushkina, N. Yu., Pulkkinen, T. I., Kubyshkina, M. V., Sergeev, V. A., Lvova, E. A., Yahnina, T. A., Yahnin, A. G., and Fritz, T.: Proton isotropy boundaries as measured on mid- and low-altitude satellites, Ann. Geophys., 23, 1839–1847, https://doi.org/10.5194/angeo-23-1839-2005, 2005.

Gasque, L. C., Millan, R. M., & Shekhar, S. (2021). Statistically determining the spatial extent of relativistic electron precipitation events using 2-s polar-orbiting satellite data. Journal of Geophysical Research: Space Physics, 126, e2020JA028675. https://doi.org/10.1029/2020JA028675

Gilson, M. L., Raeder, J., Donovan, E., Ge, Y. S., and Kepko, L. (2012), Global simulation of proton precipitation due to field line curvature during substorms, J. Geophys. Res., 117, A05216, doi:10.1029/2012JA017562.

Green, J. C. (2013), MEPED Telescope Data Processing Algorithm Theoretical Basis Document, Natl. Oceanic and Atmos. Admin. National Geophysical Data Center, Boulder, Colorado

Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780

Horne, R. B., & Thorne, R. M. (1998). Potential waves for relativistic electron scattering and stochastic acceleration during magnetic storms. Geophysical Research Letters, 25(15), 3011–3014. https://doi.org/10.1029/98GL01002

Khazanov, G. V., Robinson, R. M., Zesta, E., Sibeck, D. G., Chu, M., & Grubbs, G. A. (2018). Impact of precipitating electrons and magnetosphere-ionosphere coupling processes on ionospheric conductance. Space Weather, 16, 829– 837. https://doi.org/10.1029/2018SW001837

Li, W., & Hudson, M. K. (2019). Earth's Van Allen Radiation Belts: From Discovery to the Van Allen Probes Era. Journal of Geophysical Research: Space Physics, 124, 8319–8351. https://doi.org/10.1029/2018JA025940

Liang, J., Donovan, E., Ni, B., Yue, C., Jiang, F., and Angelopoulos, V. (2014), On an energy-latitude dispersion pattern of ion precipitation potentially associated with magnetospheric EMIC waves, *J. Geophys. Res. Space Physics*, 119, 8137– 8160, doi:10.1002/2014JA020226.

Meraner, K. and Schmidt, H. (2018): Climate impact of idealized winter polar mesospheric and stratospheric ozone losses as caused by energetic particle precipitation, Atmospheric Chemistry and Physics, 18, 1079-1089, https://doi.org/10.5194/acp-18-1079-2018.

Millan, R. M., & Thorne, R. M. (2007). Review of radiation belt relativistic electron loss. Journal of Atmospheric and Solar - Terrestrial Physics, 69, 362–377. https://doi.org/10.1016/j.jastp.2006.06.019

Mironova, Irina A.; Aplin, Karen L.; Arnold, Frank; Bazilevskaya, Galina A.; Harrison, R. Giles; Krivolutsky, Alexei A.; Nicoll, Keri A.; Rozanov, Eugene V.; Turunen, Esa; Usoskin, Ilya G. (2015) Energetic Particle Influence on the Earth's Atmosphere, Space Science Reviews, Volume 194, Issue 1-4, pp. 1-96, doi 10.1007/s11214-015-0185-4

Robinson, R. M., Vondrak, R. R., Miller, K., Dabbs, T., and Hardy, D. (1987), On calculating ionospheric conductances from the flux and energy of precipitating electrons, J. Geophys. Res., 92(A3), 2565– 2569, doi:10.1029/JA092iA03p02565.

Rodger, C. J., M. A. Clilverd, J. C. Green, and M. M. Lam (2010), Use of POES SEM-2 observations to examine radiation belt dynamics and energetic electron precipitation into the atmosphere, Journal of Geophysics Research Space Physics, 115, A04202, doi:10.1029/2008JA014023.

Rumelhart, D., Hinton, G., and Williams, R. (1986a). Learning representations by back-propagating errors. Nature, 323, 533–536.

Schulz, M., & Lanzerotti, L. J. (1974). Particle diffusion in the radiation belts. In Physics and Chemistry in Space (Vol. 7). Berlin: Springer.

Sergeev, V. A., Malkov, M., and Mursula, K. (1993), Testing the isotropic boundary algorithm method to evaluate the magnetic field configuration in the tail, J. Geophys. Res., 98(A5), 7609– 7620, doi:10.1029/92JA02587.

Sergeev, V.A., E.M. Sazhina, N.A. Tsyganenko, J.Å. Lundblad, F. Søraas, Pitch-angle scattering of energetic protons in the magnetotail current sheet as the dominant source of their isotropic precipitation into the nightside ionosphere, Planetary and Space Science, Volume 31, Issue 10, 1983, Pages 1147-1155, ISSN 0032-0633, https://doi.org/10.1016/0032-0633(83)90103-4.

Shekhar, S., Millan, R., & Smith, D. (2017), A statistical study of the spatial extent of relativistic electron precipitation with polar orbiting environmental satellites. Journal of Geophysical Research: Space Physics, 122, 11,274–11,284. https://doi.org/10.1002/2017JA024716

Sinnhuber, M., Tyssoy, H. N., Asikainen, T., Bender, S., Funke, B., Hendrickx, K., et al. (2021). Heppa III intercomparison experiment on electron precipitation impacts, part II: Model-measurement intercomparison of nitric oxide (NO) during a geomagnetic storm in April 2010. Journal of Geophysical Research: Space Physics, 126, e2021JA029466. https://doi.org/10.1029/2021JA029466

Tsyganenko, N. A. (1989). A solution of the Chapman-Ferraro problem for an ellipsoidal magnetopause. Planetary and Space Science, 37(9), 1037–1046. https://doi.org/10.1016/0032-0633(89)90076-7

Tsyganenko, N. A., & Sitnov, M. I. (2005). Modeling the dynamics of the inner magnetosphere during strong geomagnetic storms. Journal of Geophysics Research, 110, A03208. https://doi.org/10.1029/2004JA010798

Tyssøy, H. N., Sinnhuber, M., Asikainen, T., Bender, S., Clilverd, M. A., Funke, B., et al. (2021). HEPPA III intercomparison experiment on electron precipitation impacts, part I: Estimated ionization rates during a geomagnetic active period in April 2010. Journal of Geophysical Research: Space Physics, 126, e2021JA029128. https://doi.org/10.1029/2021JA029128

Yahnin, A. G., T. A. Yahnina, N. V. Semenova, B. B. Gvozdevsky, and A. B. Pashin (2016), Relativistic electron precipitation as seen by NOAA POES, Journal of Geophysics Research Space Physics, 121, doi:10.1002/2016JA022765.

Yahnin, A. G., T. A. Yahnina, T. Raita, and J. Manninen (2017), Ground pulsation magnetometer observations conjugated with relativistic electron precipitation, Journal of Geophysics Research Space Physics, 122, 9169-9182, doi:10.1002/2017JA024249.

Yando, K., Millan, R. M., Green, J. C., & Evans, D. S. (2011). A Monte Carlo simulation of the NOAA POES medium energy proton and Electron detector instrument. Journal of Geophysics Research, 116, A10231. https://doi.org/10.1029/2011JA016671

Yu, Y., Jordanova, V. K., McGranaghan, R. M., & Solomon, S. C. (2018). Self-consistent modeling of electron precipitation and responses in the ionosphere: Application to low-altitude energization during substorms. Geophysical Research Letters, 45, 6371– 6381. https://doi.org/10.1029/2018GL078828

11 Supplementary Material

The Supplementary Material shows three additional figures (Supplementary Figure 1, 2, and 3) and one Table (Table 1).

12 Data Availability Statement

The POES/MetOp 2s data used in this work are available at https://satdat.ngdc.noaa.gov/sem/poes/data/processed/ngdc/uncorrected/full/. The dataset preparation and model training are done on a Linux OS (version 3.10.0-1160.49.1.el7.x86_64) machine (Shared Computer Cluster at Boston University) in Python (version 3.8.6), using the TensorFlow library (version 2.5.0, https://www.tensorflow.org) and the Python packages: Matplotlib (https://matplotlib.org), Scikit Learn (https://scikit-learn.org/stable/), Xarray (https://xarray.pydata.org/en/stable/), Joblib (https://joblib.readthedocs.io/en/latest/), Seaborn (https://seaborn.pydata.org/), Numpy (https://numpy.org), and Pandas (https://pandas.pydata.org). Figures 1 and 4 and Supplementary Figures 1 and 3 have been produced in IDL (version 8.6.0). The trained model and a sample Python script to apply the model on any POES/MetOp date can be found in the GitHub repository here: https://github.com/luisacap/REPs_classifier_codes_for_paper.git.

Figure 1. Examples of a) a wave-driven (REP) precipitation event and b) a CSS-driven (CSS) precipitation event. Electron flux observed by POES n19 (a) and MetOp m02 (b) satellites is color-coded by energy channel (as indicated in panel b), and shown as a function of time and satellite trajectory expressed in *L* and MLT. Dashed (solid) lines are relative to the 90° (0°) telescope, indicating the trapped (precipitating) electrons. The precipitation events are highlighted by the gray rectangles.

Figure 2. Portion of the training dataset: a) class of each data point and b) electron flux for different energies. Dashed and solid lines in panel b) indicate the 90° and 0° telescope observations, respectively, as in Figure 1. Precipitation events are highlighted in gray in panel b) and their relative class is shown in panel a), where class 0 indicates “no event”, class 1 indicates “REP event” and class 2 indicates “CSS event”.

Figure 3. Three different portions of the test dataset in a similar format as Figure 2. Panels a), c) and e) show the original class of each event in the dashed gray line and the class of each event predicted by the model in solid black. Panels b), d) and f) show the electron flux as Figure 2b, where each event (originally identified) is highlighted in gray.

Figure 4. Identification and classification of precipitation events on 6 days of POES/MetOp data. Each panel shows the electron flux color-coded in energy (legend in panel a) as a function of *L*, MLT, and time. Dashed (solid) lines indicate observations of trapped (precipitating) electrons from the 90° (0°) telescope. REP events identified by the model are highlighted in brown, while CSS events identified by the model are marked in blue.