

JGR Space Physics

METHOD

10.1029/2022JA030683

Key Points:

- We explore a transfer learning approach to detecting lightning-induced electron precipitation events
- We automatically 25961 with a small apparent error rate
- We analyze statistics of LEP events, finding, for example, that >30% of LEP events are only detectable in VLF phase

Correspondence to:

N. Pailoor,
npailoor3@gatech.edu

Citation:

Pailoor, N., & Cohen, M. (2022). A large database of lightning-induced electron precipitation events. *Journal of Geophysical Research: Space Physics*, 127, e2022JA030683. <https://doi.org/10.1029/2022JA030683>

Received 26 MAY 2022

Accepted 26 OCT 2022

A Large Database of Lightning-Induced Electron Precipitation Events

Nikhil Pailoor¹  and Morris Cohen¹ 

¹School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, USA

Abstract Lightning-induced Electron Precipitation (LEP) events represent a known loss mechanism for radiation belt electrons. The quantitative role LEP events play in the radiation belt dynamics, however, is not fully understood. Direct measurements of the D region of the ionosphere, where LEP electrons are deposited, are very limited. However, Very Low Frequency (VLF) remote sensing has for a long time been used to detect perturbations in the D region, in particular using the narrow-frequency signals of VLF transmitting beacons. We present an automated neural-network based search for LEP events using our network of VLF receivers in the continental US detecting, seeded by National Lightning Detection Network (NLDN) lightning strokes. We arrive at a database of nearly 26000 LEP events over a span of 18 months within the continental USA. We find several notable patterns in the occurrence and behavior of events, such as 30.1% of events only containing disturbances in phase, and a substantially greater probability of perturbations in phase and major axis amplitude being negative.

1. Introduction

The Van Allen radiation belts are layers of charged particles trapped by the Earth's magnetic field, captured from solar wind. The inner radiation belt, parameterized by the McIlwain L-parameter, or L-shell, exists in the range of $L = 1.1$ to 2.5 , and the outer radiation belt spans $L = 3$ to 6 . At L shells above $L = 1.25$, electrons are sensitive to the effects of whistler-mode electromagnetic waves. Electrons which encounter whistler-mode waves under the gyroresonance condition (Bell, 1984), experience changes in their parallel velocities (defined relative to the magnetic field lines). If their parallel velocity is accelerated sufficiently, altering the pitch angle, the electrons trajectory leads them to collide with thermospheric neutral particles in the 60–90 km range, which is also the D region of the ionosphere. The predominant sources of whistler-mode radiation in different regions of the ionosphere remains the subject of investigation and research. These sources include the plasmaspheric hiss, artificial VLF transmission from manmade transmitters, and VLF emissions from high intensity lightning strokes.

As lightning is an intense and global source of VLF waves, via impulsive radio atmospherics, it is also a contributor to electron loss from the radiation belts. This process of electron precipitation, known as lightning-induced electron precipitation (LEP), represents a key loss mechanism of radiation belt electrons (Abel & Thorne, 1998). An example of direct observation of LEP events made by a satellite can be found in Inan et al. (2007).

However, LEP events are more commonly observed by VLF remote sensing. As the precipitating electrons re-enter the atmosphere, they deposit their energy and create additional ionization. If a VLF signal propagating below the ionosphere crosses through this region, the signal will be scattered by the disturbed ionization. LEP events are observed exclusively at night, when D-region ionization levels are fairly low. LEP events exist within a family of general lightning-generated VLF perturbations, which includes similar looking disturbances such as early/fast events. Early/fast events and LEP events both fit a similar waveform of an initial perturbation shortly after the lightning stroke, followed by a gradual recovery. However, high resolution data can distinguish these events: LEP events delay between the time of the lightning stroke and the onset of the perturbation, usually between 0.1 and 1 s.

Despite the work of Abel and Thorne (1998), however, it remains unknown when and where, and to what extent, lightning may be a significant or even dominant contributor to radiation belt losses, via LEP events. Previous approaches to quantifying the specific role of LEP events has relied on either specific case studies, or making broad assumption of LEP occurrence based on global lightning patterns, not on an actual database of LEP events. An example of the former approach is Peter (2007), which applied advances in LEP modeling from Bortnik et al. (2006) to quantify the electron precipitation from two LEP events detected in Colorado in March 2001.

This approach relied on VLF measurements of signals perturbed by the disturbed section of the ionosphere, and successfully derived a linear relationship between signal changes and the total number of precipitated electrons.

An example of the latter approach is Sousa (2018), which utilized ray tracing to compute the average power density for a stencil array of flashes across latitude, for both dayside and nightside, which was then applied to the full GLD360 database (Demetriades et al., 2010) of global lightning flashes to obtain a global estimate of the impact LEPs have on the radiation belts. But the presence and impact of LEP events is highly dependent on geomagnetic conditions, in a way that is not understood, so calculating LEP events strictly from lightning occurrences is incomplete.

Analysis from the DEMETER satellite over 2006–2008 shows a correlation between changes in flux among 126 keV electrons and average nighttime lightning flash intensity over the year (Gemelos et al., 2009). This occurred in the range of $2 < L < 3$, consistent with the assumption that whistler waves from LEP events are a dominant mechanism in electron flux changes in that region of the radiation belts. However, this study was also only able to correlate changes in electron flux with the occurrence of high intensity lightning, as sourced from the National Lightning Detection Network (NLDN) (Cummins et al., 1998), rather than with actual occurrence of LEP events. Recent work by Martinez-Calderon et al. (2020), using data from the Van Allen Probes in comparison with NLDN's data set, has supported this correlation between lightning occurrence and measured changes in electron flux.

Many previous studies of LEP events have been constrained to specific time durations or geographical searches. For instance, (Peter, 2004) detected a total of 161 LEP events from two specific thunderstorms, using the HAIL array of VLF receivers positioned across the mountain region of the United States. (Peter, 2005) examined data corresponding to Hurricane Isabel in 2003, specifically on the days on September 8, 11, 14, 16, and 18. This approach used a strict definition of LEP events, excluding cases with an amplitude perturbation less than 0.5 dB. This approach detected 139 LEP events on the days of September 11, 14, and 16, but no events on the 8th or 18th. (Clilverd et al., 2002) examined 1000 LEP events in the Antarctic peninsula during 1994 and 1995, using a perturbation detector that examined the VLF data from 10 transmitter-receiver paths and searched for statistically significant perturbations in amplitude and phase relative to the background noise.

In this paper, we attempt to bridge the gap between these approaches by designing a classifier to detect LEP events from their signature on narrowband VLF transmitters. The classifier will build on previous work using neural network classifiers to detect Early-Fast events in narrowband VLF data. We deploy that classifier to existing VLF data in order to build a large scale database of events.

We will examine this large database of LEP events to build a more comprehensive understanding of the causes of LEP events, as well as the behavior of LEP events in the VLF waveforms.

2. Methodology

2.1. Data Collection

The first step to building a database of LEP Events is obtaining a data set of candidate lightning strokes, since we presume that all or nearly all LEP events will originate from lightning strokes, particularly with high peak current. For the purposes of being able to detect perturbations in the D region through VLF remote sensing, we further narrow this data set to strokes occurring at nighttime (defined at 80 km altitude), at locations such that the potentially disturbed region of the ionosphere intersects with a transmitter-receiver path in our VLF receiver network, and that the entire transmitter-receiver path is under a nighttime ionosphere. We build our data set of events from the VLF data collected at each of those receivers.

This database serves two key purposes: allowing us to quantify the event occurrence rate for a given lightning stroke on a given day, and providing a training data set for a machine learning classifier to automatically detect LEP events from VLF data.

The data set of lightning strokes we use is Vaisala's National Lightning Detection Network (Cummins et al., 1998). This database contains detailed and precise information about the time, location, intensity, and polarity of each stroke in the continental United States, with location accuracy in the 100-m range, and timing accuracy in the μ s

Table 1
VLF Transmitters Operating in the United States, by Location and Frequency

Call sign	Location	Frequency (kHz)
NLK	Jim Creek, Washington	24.8
NML	LaMoure, North Dakota	25.2
NAA	Cutler, Maine	24.0
NPM	Lualualei, Hawaii	21.4
NAU	Aguada, Puerto Rico	40.75

range. We select only strokes with a peak current greater in magnitude than 100 kA, as observable LEP events are thought to occur primarily due to high peak current lightning.

Georgia Tech's LF Receiver network, which comprises a group of LF receivers placed at sites around North America, serves as a starting point for VLF data collection. Each receiver uses two cross-looped air-core antennas, which capture the incoming magnetic field in two perpendicular directions, North/South and East/West. The receivers use anti-aliasing filters to achieve a flat 1–470 kHz passband. The noise levels remain flat at $-10 \text{ dB } fT/\sqrt{Hz}$ over the 18–30 kHz frequency band. At roughly 7.3 nT magnetic fields, the receiver measurements "clip" due to the $\pm 5 \text{ V}$ range of the ADC. This is well above the fields generated by distant VLF radio transmitters and most

global lightning strokes, making this network of receivers suitable for measuring LEP events without saturating (Cohen, 2018).

We focus in particular on five US-Navy operated transmitters which are sources of VLF signals, as seen in Table 1. These transmitters operate over bandwidths of a couple hundred Hz around a carrier, and use minimum shift keying (MSK) modulation at 200 baud (Watt, 1967).

Due to trans-ionospheric propagation of lightning-generated whistler waves causing the waves to rise in L shell number, the region of disturbance is polewards-displaced from the site of the stroke. At 39 degrees geographic latitude, the oblique propagation of the whistler wave results in the precipitation occurring at a point polewardly displaced 6–8 degrees relative to the lightning stroke (Lauben et al., 1999), which we approximated as a standard 700 km northwards. As these models predict a range of possible displacements at different latitudes, we approximated the maximum radius of the possible disturbance to be 400 km, drawn as a circle around the polewards location.

Figure 1 provides a summary of the collection process thus far. Here, we begin with a lightning stroke off the coast of North Carolina in October 2018. We assume a region of disturbance 700 km north, shown in the gray circle. This circle intersects the path from the NAA transmitter to the Dover receiver. The graph shown is the amplitude data collected at the Dover receiver, centered such that $t = 0$ marks the time of the lightning stroke.




Figure 1. An LEP event displayed over a map of North America. This event was generated by a -217 kA stroke that occurred on 27 October 2018 at 07:36:18.217 UT.




Figure 2. Adapted from Gross et al. (2018). Geometric setup for a VLF receiver's antenna orientation (black axes), wave propagation axes (red axes), and polarization ellipse (blue ellipse) of the magnetic flux density. The angle θ_{az} gives the nominal arrival direction from the source. The green markings show the major axis and minor axis of the ellipse, τ is the tilt angle of the ellipse from ϕ , and χ is the ellipticity angle. The red vector on the ellipse shows the start phase and rotation sense of the ellipse.

given the amplitudes and phases of the two components, which is typically an ellipse. The four components of the ellipse are then extracted: (a) Major axis, or the longest diameter of the ellipse, (b) Minor axis, of the shortest diameter of the ellipse, (c) Tilt angle, which captures the rotation of the major axis counterclockwise from the ϕ direction, and (d) Start phase, which captures where along the ellipse is the magnetic field at $t = 0$, and which direction it rotates.

The AWESOME receivers produce VLF data in two resolutions, a high resolution sampling of 50 Hz, and a low resolution sampling of 1 Hz. We show a comparison of these resolution formats for the minor axis variable in




Figure 3. An LEP event, shown through the minor axis of the polarization ellipse. Note the delay between the lightning sferic and the onset of the perturbation. This event occurred on 3 April 2019 detected at the Delaware receiver from the NAA signal.

We see from this data that an LEP event did indeed occur, as there is a sharp reduction in amplitude immediately after the stroke, a result of the ionosphere over the transmitter-receiver path seeing a rapid increase in electron density. This increase in electron density causes the VLF signal from NAA to attenuate more rapidly, resulting in a weaker signal detected at Dover. Over the next 50s, chemical processes in the ionosphere cause the electron density to return to normal, allowing the signal to return to its previous strength.

The transmitters make use of a 200 baud Minimum Shift Key (MSK) modulation scheme. Here, frequency variances of ± 50 Hz (one fourth of the baud rate) from the center frequency define a bit of "1" or "0" being communicated. In MSK communications with this frequency variance, the transmitter uses a bit period of 5 ms, within which the phase rises or falls by 90 degrees. This serves to limit the overall bandwidth by ensuring continuous phase during the bit sequence. By decoding this modulation and removing the modulated phase shifts, we can recover an effective continuous wave (CW) phase, as if the transmitter were sending a monotone single frequency signal. After removing ambiguities in the phase, we can describe the signal using the four components of an elliptically polarized wave. The process of converting amplitude and phase measurements from the N/S and E/W aligned receivers into the polarization ellipse format is described in detail in Gross et al. (2018). Figure 2 shows the structure of the polarization ellipse, where the vertical and horizontal axes represent the orientation of the receiver's loop antennas, typically in the north-south and east-west directions, and the red axes shows the direction away from ($\hat{r}$) and orthogonal to ($\hat{\phi}$) the source.

The blue curve shows an example of how the magnetic field might oscillate

given the amplitudes and phases of the two components, which is typically an ellipse. The four components of the ellipse are then extracted: (a) Major axis, or the longest diameter of the ellipse, (b) Minor axis, of the shortest diameter of the ellipse, (c) Tilt angle, which captures the rotation of the major axis counterclockwise from the ϕ direction, and (d) Start phase, which captures where along the ellipse is the magnetic field at $t = 0$, and which direction it rotates.

The AWESOME receivers produce VLF data in two resolutions, a high resolution sampling of 50 Hz, and a low resolution sampling of 1 Hz. We show a comparison of these resolution formats for the minor axis variable in Figure 3. In the high resolution format, we can see that the signal does not significantly deviate from its pre-stroke average immediately after the stroke, but instead begins decreasing roughly 300 ms after. By roughly 1000 ms after the lightning stroke, the signal is at a level significantly below the pre-stroke average, and remains at this level for the next few seconds. Due to the noise, the precise time of the onset is difficult to determine clearly. Furthermore, other lightning sferics and short-term ionospheric disturbances are visible and shape the data. Because the full event lasts over a period of tens of seconds, up to several minutes, we choose to primarily use the low resolution data in the interest of minimizing noise, but do use the high resolution data to check that we have not included many events with small onset delay (which might indicate an Early VLF event).

2.2. Data Set of Candidates

Our goal is to identify LEP events in mass based on the data collection described in the previous section. We have built this database on data from 4 November 2017 to 13 August 2019.

We begin with NLDN-identified lightning strokes above 100 kA peak current (absolute value), which we refer to as a candidate stroke. There were 157,780 candidate strokes. However, many of these strokes have multiple

transmitter-receiver paths that may have been disturbed, which we refer to as a candidate sample. There were 755,524 candidate samples.

Our classifier will examine each candidate sample, and identify whether or not an LEP event occurred. The ones which pass this classifier will be referred to as event samples. In addition, since the same stroke may have multiple event samples disturbed, the event samples are then grouped into unique events.

Not all lightning strokes, even intense strokes, generate a noticeable LEP event. Due to the noisy nature of VLF data, the perturbations associated with large-scale changes in the ionosphere are difficult to analytically separate from those caused by sferics, background static, or instrumentation artifacts. Previous case studies have relied on manual classification, which results in a limited sample size with narrow geographic and temporal scope. It would be too time consuming to rely on manual classification over the entire set of candidates within the continental US over even a period of a single year, so we aim to create an automated process for event detection.

2.3. LEP Classifier

We employed a transfer learning approach toward building our LEP Classifier. We began by deploying a previous classifier trained in Pailoor and Cohen (2022) on Early-Fast events. At the lower signal resolution, LEP Events and Early-Fast events have similar appearance in their waveforms, allowing us to begin with the same classifier. The inputs were the four channels of signal data (major axis, minor axis, tilt angle, and start phase), with a time sample spanning 10 s before the lightning stroke and 30 s after. This network was structured with three fully connected hidden layers. The outputs were limited to two cases: event and non-event. We deployed the Early-Fast classifier over our database of candidate samples. We extracted 500 samples classified as "Events", 500 samples classified as "non-Events", and 500 randomly selected samples. We refer to each of these 500 sample bins as the Predicted Positive Data set, the Predicted Negative Data set, and the Random Data set respectively. We manually examined and labeled each of these samples. In the Predicted Positive Data set, 298 of the samples were actual events, while 202 were false positives. In the Predicted Negative Data set, 493 of the samples were accurately classified as "non-Events", while 7 were false negatives (actual events classified as non-events). In the Random Data set, 54 of the samples were events, while 446 were non-events.

In total, this gives us a data set of 359 events and 1141 non-events. We then use this data set to train a new neural network, which will ultimately be our final LEP classifier.

One potential source of False Negatives comes from the differences between the waveforms of Early-Fast events and LEP events. While the two types of events have similar perturbation and recovery patterns, Early-Fast events are more likely to result in a positive perturbation (signal strength increased), while LEP events can result in either a positive or negative perturbation. To correct for this, we take each data channel after the preprocessing, and "invert" the data, resulting in a new sample with values of "1-x", where x is the original sample after normalizing. This inverted sample is then added as an additional sample in the data set, effectively doubling the data set size and training the classifier to recognize both positive and negative polarity events.

The structure of the classifier network was something we investigated. The Early-Fast classifier was built using a traditional artificial neural network (ANN) structure, using multi-layer perceptrons, also known as "fully connected layers". A fully connected layer acts upon its inputs through a linear transformation, and a nonlinear activation function. In this case, we chose our activation function to be a Rectified Linear Unit (ReLU), equivalent to $f(x) = \max(0, x)$. ReLU has long been a popular activation function in deep learning due to sparse activation, scale invariance, and efficient computation of gradients (Glorot et al., 2011).

Recent research has indicated that 1D Convolutional layers can improve classifier accuracy (Kiranyaz et al., 2021). The mechanism for this improvement comes from the use of convolutional kernels to match the shape of specific waveforms in a 1D signal. The use of 1D convolution has been growing in recent years in deep learning applications, beginning with efforts to classify and detect Electrocardiogram (ECG) signals (Kiranyaz et al., 2014), and later to applications such as structural damage detection using wireless sensors (Avci et al., 2018) and fault detection in modular multilevel converters (Abdeljaber et al., 2018).

To compare the effectiveness of different network structures for our purposes, we trained and tested three different networks: a predominantly convolution-based network, a fully connected network, and a "hybrid" network




Figure 4. Test accuracy over each training batch for the three networks compared.

including both convolutional and fully connected layers. Note that even in the convolution-based network, a fully connected layer is still used to sort the data into one of the two result classes.

The classifier's performance can be measured using a metric called the loss function, which is a function that assigns a penalty to incorrect classification outputs. The classifier "learns" by carrying out a gradient descent optimization in order to minimize the loss function output for the training set. Initially, the loss function outputs for both the training and the test set will decrease as the classifier attempts to minimize loss for the training set. However, because the classifier is not learning from errors from the test, it will eventually "overfit" (Tetko et al., 1995) to the specific characteristics of the training set, and the loss function output for the test set (test loss) may increase. We use a cross entropy loss function, which will be described in more comprehensive detail further down.

We present the comparison of the test performance over the course of classifier training in Figure 4, which shows the trajectories of the loss functions of each of the three classifiers, as we increase the number of training batches. As a lower loss value corresponds to better performance, the largely convolutional network (CNN) performed significantly better than either the fully connected network (FC) or the hybrid network. Note that for the fully connected network and the hybrid network, the test loss begins overfitting a certain number of training batches (2200 and 2050) respectively. Also note that while the CNN performed better than its rivals, it converged significantly slower, and does not appear to have reached its peak after 2500 training batches.

The precise time samples chosen also have an effect on the classifier performance. Although our database stored low resolution data ranging from 40 s before the stroke to 120 s after, this 160 s long sample is impractical for use in the classifier, as the main event behavior occurs very close to the $t = 0$ region. Additionally, because the data preprocessing normalizes each channel between 0 and 1, other anomalies and events that occur before and after the stroke can greatly distort the sample.

The original classifier used for Early-Fast event used a 40 s window, ranging from 10 s before the stroke to 30 s after. For the LEP event classifier, we tested variations of the classifier that used 40 s, 60 s, and 80 s sample widths, ranging from $t = -10:30$, $t = -20:40$, and $t = -30:50$ respectively. We found that a 60 s sample with produced the best results, as shown in Figure 5.

Finally, we note that we can vary the classifier's loss function to achieve different goals. In this classifier, we have used a Cross Entropy loss measurement. The classifier provides, as its output, a pair of values representing one of the two classes (Non-Event and Event). The Cross Entropy of an output x corresponding to a sample with a label y is as follows:

$$l(x, y) = -w_y \log \frac{\exp(x_y)}{\exp(x_0) + \exp(x_1)}$$

where w represents the weights assigned to each of the two output values.




Figure 5. Test accuracy over training batch for different time widths.

For example, if we have an LEP Event ($y = 1$) and the classifier produces the output $[0,1]$, the loss function should return the value $-w_y \log \frac{e}{1+e} = -w_y (1 - \log(1 + e)) \approx 0.31w_y$. If the classifier had instead returned the output $[1,0]$, the loss function would return $-w_y \log \frac{1}{e+1} = -w_y (0 - \log(1 + e)) \approx 1.31w_y$.

Changing w_y changes the magnitude of the error the classifier assigns to getting each class wrong. By setting a relatively higher weight for Events than Non-Events, we can trade off the True Positive Rate (TPR, defined by $\frac{TP}{P}$) and the Positive Predictive Value (PPV, defined by $\frac{TP}{PP}$).

Figure 6 shows the trade between TPR and PPV. Each dot represents a classifier trained with a different set of weights. We can see, particularly with the green set of dots, that attempting to achieve a high TPR results in a lower PPV, and vice versa. Dots that are further to the top left correspond to classifiers which may miss a large fraction of events, but has predicted events which reliably are real events. Dots further to the lower right will rarely miss an event, but will have a high fraction of false positives.

If the relative weight assigned to Events is low, we can achieve a high Positive Predictive Value (most samples classified as events will really be events), but this comes at a cost of a low TPR (most events will be missed). Likewise, by increasing the weight assigned to events, we can capture most events and achieve a high TPR, but at the cost of reducing our PPV (a sample classified as an event is less likely to really be an event).




Figure 6. Positive Predictive Value (PPV) versus True Positive Rate (TPR). Each dot represents a classifier trained with a different set of weights.

For our final network, we set our weight value equal to 2.0 to balance the TPR and PPV effectively, using the convolutional network described previously as well as preprocessing our data samples to have a 60s with. Combining these choices, we achieved a network that lowered test loss to 0.335, and achieved a TPR of 0.92 and a PPV of 0.79. Compared to the original classi this is an improvement of the PPV by 32.6 from the original classifier, and a lowering of the False Negative Rate (FNR) by 38.5%.

2.4. Comparison to Alternative Approaches

The process described so far has involved a transfer learning approach, using a classifier originally trained on Early/Fast events to detect candidate samples, which are then added to the training data set to train a new classifier. Because of the low occurrence of LEP events within the broader population of candidate samples, if we were to start from scratch it would require considerably more manual classification in order to get a statistically significant number of events.

To illustrate this, we manually labeled an additional 1000 samples randomly selected from the population of candidate samples. Adding to our existing pool of 500 random samples, this gave us a data set of 1500 samples, equal in size to the one produced with the assistance of the Early/Fast classifier. As shown in Figure 6, this classifier had a significantly worse performance over a wide range of weight values.

We also compare to one of the most simple types of classifier: perturbation detection. Since a defining characteristic of LEP Events is an initial sharp perturbation following a lightning stroke, we can examine the data preceding and proceeding the stroke to see if such a perturbation has occurred. We can define this using the variance and average of the data points immediately prior and immediately after the stroke occurrence. For our perturbation detector, we define the incident field average using the average of data points ranging from 12 to 2 s before the stroke, and define the post-stroke average using the average of data points 2–12 s after the stroke. We exclude the time points including the stroke and directly surrounding it to avoid the waveform of the sferic distorting the average or variance.

We can then define the scattered field using the difference between the incident field and the post-stroke field. If the scattered field is larger than the sum of the variances of these fields, we consider a perturbation to have happened. There are four channels, and if a perturbation happens in any of these, we consider the sample to have passed our perturbation detector.

However, just because a sample contains a perturbation does not mean it contains an LEP Event. When we run this perturbation detector over a randomly selected data set of events, we find a PPV value of only 0.153. It achieved a perfect TPR value of 1.0, but this just means that, even if every LEP event contains a perturbation, most perturbations do not result in events.

Our final classifier used a combination of the perturbation detector and the CNN. Of the 755,524 nighttime candidate samples, our classifier detected 25,961 event samples. When categorized by unique lightning strokes, there are 157,780 candidate strokes, out of which 18,109 were classified as unique events. This occurrence rate (11.5%) is similar to the 10.8% occurrence found in the Random Data set.

2.5. High Resolution Analysis

We note that the survey parameters used to search for LEP events in some cases allow for cases where both the stroke, and the assumed region of perturbation, are in close proximity to the transmitter-receiver path. This is particularly more likely in paths that span over a greater North-South range. In this cases, there is a chance that detected perturbations may be the result of Early/Fast events, a phenomena that at low resolutions appears similar to LEP events.

To gauge the incidence of Early/Fast events in our database, we examined the corresponding high resolution samples of 145 events. The high resolution data is sampled at 50 Hz. At this resolution, it is possible to detect the onset delay between the initial lightning sferic, and the perturbation that occurs as a result of electron precipitation. This onset delay is the result of whistler-mode VLF waves having to propagate large distances along the magnetic field lines in order to scatter energetic electrons in the radiation belts, and typically ranges from 200 ms to 2.5 s (Peter, 2004).




Figure 7. High resolution image of an LEP Event detected on 26 March 2019, at the Burden receiver from the NAA transmitter.

In our survey of high resolution samples, 59 out of 145 had a clear, identifiable onset delay, while 8 had an identifiable perturbation within 200 ms of the lightning stroke. 78 of the samples were too noisy to clearly distinguish the precise time of perturbation. Figure 7 is an example of an event, with the major axis shown, that has a clearly identifiable onset delay.

The ability to accurately classify a data set of this size allows us to map in detail the occurrence of events across different points of time, allowing us to find key patterns and relationships to direct observations of precipitating electrons from satellites.

3. Results

3.1. Distribution of Events

Our database allows us to look at the distribution of LEP events by parameters such as lightning current and geography in relation to the transmitter-receiver path.

Figure 8 shows the distribution of events by lightning current. Because we only examined candidate strokes with $|I| < 100$ kA, there is a "gap" in the center. The errorbars reflect the True Positive Rate (as an upper bound error) and the Positive Predictive Value (as a lower bound error) of the classifier.

We see a positive correlation between peak current magnitude and the likelihood of an event to occur. There is an asymmetry in the event occurrence, with positive current strokes generating an event 16.2% of the time while negative current strokes generate an event 10.4% of the time. While this difference is larger than the error range, it is not as significant as the variance by stroke magnitude.

3.2. Event Geometry Analysis

Our initial search for LEP event candidates assumed a polewards displacement of 700 km and a maximum distance from this displacement point to a transmitter-receiver path of 400 km. Now that we have a classifier to distinguish candidate event samples and actual events, we can re-examine these initial assumptions.

If we knew with certainty that the disturbed region of the ionosphere was displaced by a fixed distance from the lightning stroke, we would expect samples from paths further away from this displacement points to be progressively less likely to yield a signal perturbation. In other words, there should be an inverse correlation between the distance from the displacement point to the path and the likelihood of a sample to be an event. We see this inverse relationship in early-fast events (Pailoor & Cohen, 2022), where the disturbance is inherently centered at the point of the stroke.




Figure 8. Left: histogram of candidate strokes by peak current Right: Event occurrence rate for each peak current bin.




Figure 9. Top: distribution of total samples for various distances from displacement point to tx-rx path. Bottom: event likelihood as a function of distance to path.

We chose three different values for the displacement point (600 km, 700 km, and 800 km), and for each of these values examine the relationship between event likelihood and distance to path.

Figure 9 shows these relationships. We can see that samples tend to be concentrated closer to the displacement point. For 600 and 700 km displacement, we see an inverse relationship between distance to path and likelihood, where paths further away from the displacement are less likely to yield an event. However, the reduction is not particularly drastic, with the occurrence rate dropping about half a percent. For an 800 km displacement, there appears to be no correlation between distance to path and occurrence rate, which suggests that 800 km is not an accurate choice for centering our assumed disturbance location.

This suggests that a 600–700 km polewards displacement is consistent with the observed distribution of where LEP events tend to be.

3.3. Amplitude and Phase Disturbance Analysis

LEP events have long been reported using the amplitude disturbance. It is also known, however, that the phase may be disturbed, either in addition to, or instead of, the amplitude. As such, some analyses that consider only amplitude may be missing some LEP events. In our case, we are converting our amplitude and phase measurements into polarization ellipse, following the method described by (Gross et al., 2018). With the polarization ellipse technique, the amplitude and phase on each of two horizontal magnetic field components are together converted into a single ellipse that has four components: A major axis, a minor axis, a tilt angle, and a start phase.

We examine the disturbance using the “perturbation detector” method described in Section 2.4, where we examine the average values for each of the four polarization variables in the 10 s preceding the stroke, and compare it to the post-stroke average 2–12 s after. For each variable this value is greater than the sum of the variances of the data before and after the stroke, we consider the disturbance to be nonzero. We note that while this method is resilient to noise, it is sensitive to changes in the event recovery time; if the event recovers within 12 s of the lightning stroke, the size of the scattered field can be underestimated.

Table 2 shows the distribution of events according to which components of the polarization ellipse is disturbed. Each entry in the table shows the fraction of events that have measurable scattered field components corresponding to the combination described in the row and column name. Each row and column represents a pair of scattered field components, with the rows referring to the amplitudes (major axis and minor axis), and the columns

Table 2

Distribution of LEP Events According to Components of the Polarization Ellipse That Are Disturbed

		Row total	Angle/phase disturbance			
			Start phase	Tilt angle	Both	Neither
Amplitude Disturbance	Neither axis	35.3%	30.1%	0.96%	4.2%	0%
	Major axis only	16.5%	8.4%	1.6%	4.6%	18.9%
	Minor axis only	22.7%	11.1%	3.1%	6.3%	2.2%
	Both axes	25.5%	9.9%	3.2%	10.3%	3.1%
Column total		100%	59.5%	8.9%	25.4%	7.2%

Note. There are 25961 events in total.

referring to the tilt angle and starting phase. The "Row total" and "Column total" sums up the rows, and columns, respectively.

The top-right entry in the table shows a count of 0, as all identified LEP events must be disturbed in at least one of the four components. On the other hand, 10.3% of events had measurable disturbances on all four ellipse components ('both' row and 'both' column). Looking at the distribution of amplitude disturbances (row total), we can see that both the major and minor axes are disturbed in roughly equal proportion: 16.5% and 22.7%, with 25.5% being disturbed on both.

However, 35.3% of events had no amplitude disturbance, meaning that many of these events would be missed if looking only at amplitude to identify LEP events. The event shown in Figure 10 is an example of one of these "phase-only" events. Of these, most had a disturbance exclusively in the start phase. In fact, this category (disturbance only in start phase) is the largest one, representing 30.1% of the LEP events in our database. For a single-mode TEM wave, the start phase corresponds to the phase velocity of the wave. To the extent that VLF waves may sometimes be dominated by a single mode, we may interpret this result, very roughly, to mean that some LEP events cause a change in group velocity of the propagating VLF waves, but do not change the attenuation rate by much. Gross et al. (2018) identified these types of LEP events, and suggested that these are due to the wide region of disturbance causing the phase velocities of the predominant signal propagation mode to be affected in a roughly equal way, without changing the signal strength of the combined modes.

More generally, it can be seen that 84.9% of events disturb the start phase (59.5% + 25.4%). The start phase can be thought of as a common-mode phase disturbance between the two components of the polarization ellipse.

The tilt angle, however, was disturbed only 34.3% of the time. This may be a consequence of the fact that LEP ionospheric disturbances are relatively smooth spatially compared to wavelength. As such, the effect on VLF propagation is primarily a change in the attenuation rate and/or wavenumber of each mode, and less so on mode

conversion on a sharp boundary as is often seen in the day/night terminator (Westerlund, 1968), or some Early/Fast events (Poulsen et al., 1993). In addition, we will show later that the tilt angle disturbance, when it does occur, does not appear to be related to the peak current of the lightning stroke triggering the LEP event.

Now that we have analyzed the occurrence of disturbances, we can examine their distribution. Figure 11 shows histograms of nonzero disturbances in all four polarization variables. Note that we present the major axis disturbances in log scale, while the minor axis is in absolute scale. This is because while the major axis amplitude is normalized in calculations to always have a positive value, the minor axis can take on negative or zero values, and as such is easier to examine in linear units.

The major axis is more likely to see negative perturbations (meaning a reduction) than positive perturbations (meaning an increase). This may be a consequence of the precipitation of electrons decreasing the ionosphere's effective height, which when applied to a single-mode propagation theory




Figure 10. An LEP event displayed in low resolution format. This event occurred on 4 December 2018. Data displayed is from the Briarwood receiver, filtered to the narrowband NML signal.




Figure 11. Histogram of nonzero perturbations for each polarization ellipse component, from left to right, major axis, minor axis, tilt angle, and start phase.

results in an increased absorption of the signal and, on balance, a decreased amplitude at the receiver (Inan & Carpenter, 1987), since at lower altitudes there are more collisions and thus higher absorption. However, in reality VLF propagation is complicated by multi-mode propagation so depending on the makeup of waveguide modes, may result in an increase in signal strength depending on how the constructive and destructive interference of the modes works. But this single-mode picture may explain why a reduction in amplitude is more common.

The starting phase more commonly decreases but sometimes increases. Using a similar single-mode simplified picture, the reduction in reflection height increases the incidence angle of propagating modes since the waveguide cutoff increases. This would cause an increase in the phase velocity along the waveguide, and thus a lower phase. But, as in the previous case, the single-mode picture may only help explain the imbalance, and the multi-mode propagation muddies this simplified picture.

The minor axis and tilt angle appear to have a more symmetric distribution between positive and negative changes.

Our database also allows us to better understand the role of peak current in determining the disturbances seen the VLF signals. Figure 12 shows our analysis of perturbations as a result of lightning strokes of different peak current ranges. Within each panel, each bar represents a different percentile value of the respective (absolute) field component perturbations. The database from which this is drawn is event samples corresponding to lightning strokes with a peak current (magnitude) within the range indicated on the horizontal axis. There are correlations between the magnitude of perturbations within the major axis, minor axis, and starting phase, particularly for the 50th and 75th percentiles of the disturbances. This is consistent with (Inan & Carpenter, 1987), which relates




Figure 12. Percentile magnitudes of field component perturbations for different peak current ranges.

Table 3
Likelihood of Individual Field Components to Be Perturbed in an Event at Different Peak Current Ranges

Peak current magnitude	Likelihood of field component perturbation			
	Major axis	Minor axis	Tilt angle	Start phase
100–150 kA	40.89%	41.68%	31.77%	84.46%
150–200 kA	49.71%	46.16%	35.12%	84.66%
200–250 kA	57.90%	41.08%	33.20%	85.61%
250–300 kA	59.90%	46.36%	36.75%	85.55%

the changes in signal amplitude and phase to the quantity of the precipitated electron flux, and (Bortnik et al., 2003), which estimates the power density of lightning-generated whistler waves as proportional to the square of the peak current. However, we do not see a significant variation in tilt angle change with respect to peak current. This logically follows from the assumption that changes to the tilt angle, which represent the effects of field modes arriving at different scattering angles, may be largely determined by geographical factors rather than the intensity of the electron precipitation.

The role of peak current in determining the likelihood of individual field component perturbations appears to largely be confined to the major axis. Table 3 shows, for each field component, the likelihood of an event to have a perturbation in that polarization variable. We see a 19% increase in likelihood

for the major axis to perturb between events formed by strokes with $I/I < 150$ kA and those from strokes with 250 kA $< I/I < 300$ kA. The other three field components do not appear to have this relationship, however, even if the likelihood of an event to occur in the first place is more strongly determined by peak current.

4. Conclusion

Our transfer-learning based machine learning approach to detecting LEP events has allowed us to assemble large database of events across North America. This database provides insights into the occurrence and behavior of events. We show that lightning peak current has a positive correlation with event likelihood, and also positively correlates with larger changes in amplitude and phase perturbation. We observe a quantified trend of amplitude and phase components having a negative perturbation. We also observe that a large share, over 30%, of LEP events only see a perturbation in the start phase.

Our approach has allowed us to assemble a larger database than previous studies have obtained, with 18,109 events observed over 567 days, a far larger number over a longer period of time than would be feasible with a manual survey. This provides us with a higher confidence in the global behavior and occurrence of LEP events than if we were limited to a smaller and more localized data set of events. Furthermore, by combining radial and azimuthal field components of the VLF signals, we show that LEP events produce distinct patterns of effects on the major and minor axes of the elliptically polarized waves.

Further work will involve using trends and patterns in LEP occurrence to provide a clearer picture of the larger impact that electron precipitation has as a loss mechanism for the radiation belts.

Data Availability Statement

All raw data is archived at the Worldwide Archive of Low-frequency Data and Observations (WALDO), described by (Cohen, 2020), and openly available at <http://waldo.world>. Once going to waldo.world, click on “Browse data” in the menubar, and then select “Broadband VLF data” for the raw VLF data at 100 kHz sampling, and “narrow-band data” for post-processed transmitter amplitude and phase data. Once there, select the year of interest, then open the folder for the site of interest, inside that select the subfolder of interest to refer to the date, then inside of that are individual files. While navigating, if you hover the mouse on the bar right of any folder or file, there is a little checkbox that you can click on to select it. You can select many files or folder, and then at the top of that navigation section, click the three dots and then pick “Download selected” to down everything that you have selected. If any issues are encountered using WALDO, or if larger scale access is needed that can't be met by the WALDO interface, contact morris.cohen@waldo.world, who can work to solve any problem or provide whatever you need. The “Format” section of the menubar will describe the file labeling and header format. The “Scripts” section gives some sample Matlab and python scripts to analyze the data.

References

- Abdeljaber, O., Avci, O., Kiranyaz, M. S., Boashash, B., Sodano, H., & Inman, D. J. (2018). 1-d CNNs for structural damage detection: Verification on a structural health monitoring benchmark data. *Neurocomputing*, 275, 1308–1317. <https://doi.org/10.1016/j.neucom.2017.09.069>
- Abel, B., & Thorne, R. M. (1998). Electron scattering loss in earth's inner magnetosphere: 1. dominant physical processes. *Journal of Geophysical Research*, 103(A2), 2385–2396. <https://doi.org/10.1029/97ja02919>

Acknowledgments

The authors thank Mark Golkowski and Vijay Harid for helpful discussions.

- Avci, O., Abdeljaber, O., Kiranyaz, S., Hussein, M., & Inman, D. J. (2018). Wireless and real-time structural damage detection: A novel decentralized method for wireless sensor networks. *Journal of Sound and Vibration*, 424, 158–172. <https://doi.org/10.1016/j.jsv.2018.03.008>
- Bell, T. F. (1984). The nonlinear gyroresonance interaction between energetic electrons and coherent VLF waves propagating at an arbitrary angle with respect to the earth's magnetic field. *Journal of Geophysical Research*, 89(A2), 905. <https://doi.org/10.1029/ja089ia02p00905>
- Bortnik, J., Inan, U. S., & Bell, T. F. (2003). Energy distribution and lifetime of magnetospherically reflecting whistlers in the plasmasphere. *Journal of Geophysical Research*, 108(A5). <https://doi.org/10.1029/2002ja009316>
- Bortnik, J., Inan, U. S., & Bell, T. F. (2006). Temporal signatures of radiation belt electron precipitation induced by lightning-generated MR whistler waves: 1. Methodology. *Journal of Geophysical Research*, 111(A2). <https://doi.org/10.1029/2005ja011182>
- Clilverd, M. A., Nunn, D., Lev-Tov, S. J., Inan, U. S., Dowden, R. L., Rodger, C. J., & Smith, A. J. (2002). Determining the size of lightning-induced electron precipitation patches. *Journal of Geophysical Research*, 107(A8). <https://doi.org/10.1029/2001ja000301>
- Cohen, M. (2018). Broadband VLF/LF/MF Radio Reception and Remote Sensing with the AWESOME Instrument.
- Cohen, M. (2020). Returning lightning data to the cloud. *Eos*, 101. <https://doi.org/10.1029/2020eo142801>
- Cummins, K. L., Murphy, M. J., Bardo, E. A., Hiscok, W. L., Pyle, R. B., & Pifer, A. E. (1998). A combined TOA/MDF technology upgrade of the u.s. national lightning detection network. *Journal of Geophysical Research*, 103(D8), 9035–9044. <https://doi.org/10.1029/98jd00153>
- Demetriades, N., Murphy, M., & Cramer, J. (2010). Validation of vaisala's global lightning dataset (gld360) over the continental United States.
- Gemelos, E. S., Inan, U. S., Walt, M., Parrot, M., & Sauvaud, J. A. (2009). Seasonal dependence of energetic electron precipitation: Evidence for a global role of lightning. *Geophysical Research Letters*, 36(21). <https://doi.org/10.1029/2009gl040396>
- Glorot, X., Bordes, A., & Bengio, Y. (2011). Deep sparse rectifier neural networks. In G. Gordon, D. Dunson, & M. Dudík (Eds.), *Proceedings of the fourteenth international conference on artificial intelligence and statistics* (Vol. 15, pp. 315–323). PMLR. Retrieved from <https://proceedings.mlr.press/v15/glorot11a.html>
- Gross, N. C., Cohen, M. B., Said, R. K., & Gokowski, M. (2018). Polarization of narrowband VLF transmitter signals as an ionospheric diagnostic. *Journal of Geophysical Research: Space Physics*, 123(1), 901–917. <https://doi.org/10.1002/2017ja024907>
- Inan, U. S., & Carpenter, D. L. (1987). Lightning-induced electron precipitation events observed at $f_i \sim 2.4$ as phase and amplitude perturbations on subionospheric VLF signals. *Journal of Geophysical Research*, 92(A4), 3293. <https://doi.org/10.1029/ja092ia04p03293>
- Inan, U. S., Piddiyachiy, D., Peter, W. B., Sauvaud, J. A., & Parrot, M. (2007). DEMETER satellite observations of lightning-induced electron precipitation. *Geophysical Research Letters*, 34(7). <https://doi.org/10.1029/2006gl029238>
- Kiranyaz, S., Avci, O., Abdeljaber, O., Ince, T., Gabbouj, M., & Inman, D. J. (2021). 1d convolutional neural networks and applications: A survey. *Mechanical Systems and Signal Processing*, 151, 107398. <https://doi.org/10.1016/j.ymssp.2020.107398>
- Kiranyaz, S., Ince, T., & Gabbouj, M. (2014). Multidimensional particle swarm optimization for machine learning and pattern recognition. In *Multidimensional particle swarm optimization for machine learning and pattern recognition*. Springer Berlin Heidelberg. <https://doi.org/10.1007/978-3-642-37846-1>
- Lauben, D. S., Inan, U. S., & Bell, T. F. (1999). Poleward-displaced electron precipitation from lightning-generated oblique whistlers. *Geophysical Research Letters*, 26(16), 2633–2636. <https://doi.org/10.1029/1999gl900374>
- Martinez-Calderon, C., Bortnik, J., Li, W., Spence, H., Claudepierre, S. G., Douma, E., & Rodger, C. J. (2020). Comparison of long-term lightning activity and inner radiation belt electron flux perturbations. *Journal of Geophysical Research: Space Physics*, 125(3). <https://doi.org/10.1029/2019ja027763>
- Pailoor, N. A., & Cohen, M. B. (2022). Mass statistical analysis of early VLF events. *Journal of Geophysical Research: Space Physics*, 127(5). <https://doi.org/10.1029/2021ja030204>
- Peter, W. B. (2004). On the occurrence and spatial extent of electron precipitation induced by oblique nonducted whistler waves. *Journal of Geophysical Research*, 109(A12). <https://doi.org/10.1029/2004ja010412>
- Peter, W. B. (2005). Electron precipitation events driven by lightning in hurricanes. *Journal of Geophysical Research*, 110(A5). <https://doi.org/10.1029/2004ja010899>
- Peter, W. B. (2007). *Quantitative measurement of lightning-induced electron precipitation using vlf remote sensing* (Unpublished doctoral dissertation). Stanford University.
- Poulsen, W. L., Inan, U. S., & Bell, T. F. (1993). A multiple-mode three-dimensional model of VLF propagation in the earth-ionosphere waveguide in the presence of localized/region disturbances. *Journal of Geophysical Research*, 98(A2), 1705–1717. <https://doi.org/10.1029/92ja01529>
- Sousa, A. (2018). *Global and seasonal effects of lightning-induced electron precipitation* (Unpublished doctoral dissertation). Stanford University.
- Tetko, I. V., Livingstone, D. J., & Lui, A. I. (1995). sep). Neural network studies. 1. Comparison of overfitting and overtraining. *Journal of Chemical Information and Computer Sciences*, 35(5), 826–833. <https://doi.org/10.1021/ci00027a006>
- Watt, A. D. (1967). *VLF radio engineering*. Elsevier. <https://doi.org/10.1016/c2013-0-02069-5>
- Westerlund, S. (1968). Periodic fading of VLF signals received over a polar path during sunset and sunrise. In *Ionospheric radio communications* (pp. 108–116). Springer US. https://doi.org/10.1007/978-1-4899-5511-1_10