AGU PUBLICATIONS

JGR

Journal of Geophysical Research: Space Physics

RESEARCH ARTICLE

10.1002/2015JA021916

Global occurrence rate of elves and ionospheric heating due to cloud-to-ground lightning

P. R. Blaes[1], R. A. Marshall[2,3], and U. S. Inan[1,4]

[1]Department of Electrical Engineering, Stanford University, Stanford, California, USA, [2]Department of Aerospace Engineering Sciences, University of Colorado Boulder, Boulder, Colorado, USA, [3]Department of Aeronautics and Astronautics, Stanford University, Stanford, California, USA, [4]Department of Electrical Engineering, Koç University, Istanbul, Turkey

Key Points:
- We find that elves can be predicted with 90% accuracy based on peak current
- We extrapolate to find the global occurrence rate of elves
- We calculate the global contribution of CG lightning to ionospheric heating

Correspondence to:
R. A. Marshall,
robert.marshall@colorado.edu

Citation:
Blaes, P. R., R. A. Marshall, and U. S. Inan (2016), Global occurrence rate of elves and ionospheric heating due to cloud-to-ground lightning, *J. Geophys. Res. Space Physics*, *121*, 699–712, doi:10.1002/2015JA021916.

Received 14 SEP 2015
Accepted 19 DEC 2015
Accepted article online 22 DEC 2015
Published online 19 JAN 2016

Abstract We present ground-based observations of elves made using an optical free-running photometer along with VLF/LF observations of the lightning electromagnetic pulse (EMP) magnetic field. We use these experimental observations to investigate the properties of the lightning return stroke that control the production of optical elve emissions. Two summers of data containing observations of over 600 elves along with the LF magnetic field of the associated lightning are analyzed. By training a classifier on features of the EMP ground wave, we find that we are able to accurately predict whether or not a stroke produced an elve. We find that the peak current of the causative discharge predicts elve production with 90% accuracy. Further, we find that the production probability of elves as a function of peak current fits a linear regression, with a 50% elve probability for peak currents of 88 kA. We use this finding along with global data from the GLD360 lightning geolocation network to extrapolate the global elve production rate; we show that ~0.8% of all cloud-to-ground lightning discharges produce elves. Finally, using GLD360 data and a numerical model of the lightning EMP, we estimate the total amount of ionospheric heating due to lightning, amounting to approximately 2 MW of continuous power dissipated globally in the lower ionosphere.

1. Introduction

Lightning is one of the most frequently occurring natural phenomena on Earth and forms an integral component of the global electric circuit. At any given moment, there are on average 2000 active thunderstorms covering about 10 percent of the Earth's surface [*Rakov and Uman*, 2003, p. 10] with a global average flash rate of about 45 lightning discharges per second [*Christian et al.*, 2003]. Lightning activity is not uniformly distributed; particular hotspots such as the Caribbean, Central Africa, and Southeast Asia account for much of the global activity, with much of it occurring during summer months. Lightning can be very destructive and significantly impacts human society. It is estimated that lightning kills 24,000 people globally and causes and results in $300 million in insured property damage in the U.S. each year [*Holle*, 2008].

In addition to the direct impact of strikes on the ground, lightning also radiates an intense electromagnetic pulse (EMP). This EMP propagates long distances and may interfere with electronics and radio communications [*Wait and Spies*, 1964]. A portion of the EMP energy propagates upwards where it interacts with plasma in a region of the upper atmosphere called the ionosphere. This EMP-ionosphere interaction results in energy transfer to free electrons and, subsequently, electron-neutral collisions which may excite optical emissions known as elves [*Inan et al.*, 1997]. The heating and ionization associated with elves further affects the propagation of lightning-generated sferics into the magnetosphere by modifying the propagation medium [*Marshall*, 2014]. Elves belong to the broad class known as transient luminous events (TLE) and are a relatively recent discovery, having only been first scientifically observed in 1990 [*Fukunishi et al.*, 1996; *Inan et al.*, 1991].

Despite the frequency of lightning and the danger to humans posed by ground strikes, there is still a great deal unknown about the phenomenon. In particular, there is uncertainty in the properties of the current propagation known as the return stroke. Furthermore, there are outstanding questions regarding the conditions necessary in the return stroke and ionosphere for elve production. In this paper, we utilize radio remote sensing and photometric imaging techniques to observe elves and to answer questions pertaining to elve production and ionospheric heating associated with elves.

Several prior studies have attempted to quantify the production rate and probability of elves as a function of lightning parameters, in particular the National Lightning Detection Network (NLDN)-reported peak current. The first such study was conducted by *Barrington-Leigh et al.* [1999] using the "Fly's Eye" photometer which recorded 34 manually triggered events that were identified as elves. All of the elves had NLDN peak currents greater than 56 kA, suggesting a peak current threshold for elve production. Measurements from the satellite borne Imager of Sprites and Upper Atmospheric Lightning (ISUAL) instrument, another triggered photometer, suggested that all Cloud-to-Ground (CG) lightning with peak current magnitudes greater than 60 kA produce elves [*Chen et al.*, 2008b]. A later study by [*Newsome*, 2010; *Newsome and Inan*, 2010] conducted over multiple summers using the nontriggered Photometric Imaging of Precipitation of Electron Radiation (PIPER) photometer in which ~1000 elves were observed suggested, however, that only about 10% of 60 kA CG strokes produce elves. Puzzlingly, the results of [*Newsome*, 2010] show a surprising number of large CG strokes, many greater than 100 kA, for which an elve was not observed.

A global distribution and production rate for elves was also given by [*Chen et al.*, 2008b] using satellite data from the ISUAL photometer. These data, however, are biased by the fact that ISUAL used triggered recording and the fact that it was in a polar orbit with a narrow field of view. No prior studies have shown a global elve rate inferred from nontriggered ground based measurements.

2. Data

In this study, we use optical elve observations collected using the PIPER instrument [*Marshall et al.*, 2008]. The PIPER photometer was originally designed at Stanford in an attempt to observe the optical signatures of lightning-induced electron precipitation (LEP) events; however, it was found to be particularly useful for observing elves [*Newsome*, 2010; *Newsome and Inan*, 2010]. Due to the high data rates, many previous instruments operated in modes based on triggered recordings of lightning, introducing a possible sampling bias. PIPER, on the other hand, is a free-running instrument and does not suffer from such bias. Data are sampled continuously at a high rate, and all data are saved. Event detection and labeling is then performed in postprocessing. The PIPER instrument is outfitted with a longpass red filter with a lower wavelength cutoff of 650 nm, and the photomultiplier tube (Hamamatsu R5900U-20-L16) has a flat response up to about 850 nm. Hence, the instrument is responsive to photons in the range 650–850 nm, which is dominated by the N_2 1P emission in elves.

An experimental field campaign was run in the western United States from mid-June until September of 2013 in which we observed 671 elves along with magnetic field recordings of the coincident VLF/LF (1 MHz sample rate) radio atmospherics, or sferics, that were emitted by the causative lightning return strokes. This data set was used for the bulk of the analysis carried out in this paper. Two PIPER photometers were deployed: one at McDonald Observatory, Texas and one at Langmuir Laboratory, New Mexico. Both were pointed over roughly the same region of northern Texas and western Oklahoma. Both photometers recorded simultaneously and could be pointed independently to track storms. The major advantage of the multisite photometer setup was that though both sites pointed over the same region, it dramatically increased the size of our data set due to the redundancy. The probability that one of the sites would have its viewing obscured due to local weather is quite high, but the probability that both sites have cloud cover on the same night is much lower. Several example PIPER observations which were recorded during the 2013 summer campaign are shown in Figure 1, including a single elve in Figures 1a and 1b; a PIPER observation containing an elve, halo, and sprite in Figures 1c and 1d; a "cloud flash" followed by an elve in Figures 1e and 1f; and an "elve doublet" event in Figures 1g and 1h. These doublet events have been recently linked to compact intracloud discharges (CIDs) as their source [*Marshall et al.*, 2015]. For the analysis presented in this paper, we do not include elve double events since they are quite rare, and we are primarily concerned with CG lightning. Also, care was taken to only include the contributions of a single elve once when it was observed from both stations simultaneously.

In addition to the photometers, an array of six AWESOME VLF/LF magnetic field receivers was installed in Oklahoma, Kansas, and Texas, roughly located under the fields of view of the two PIPERs. The intended purpose of these receivers is to record the radio emissions from nearby lightning, especially the causative lightning that produced elves so as to infer the return stroke properties that are important for elve production. These receivers are based on the design presented in *Cohen et al.* [2010] but modified to encompass the LF band up to 400 kHz and sampled at 1 MS/s on each of two orthogonal channels.

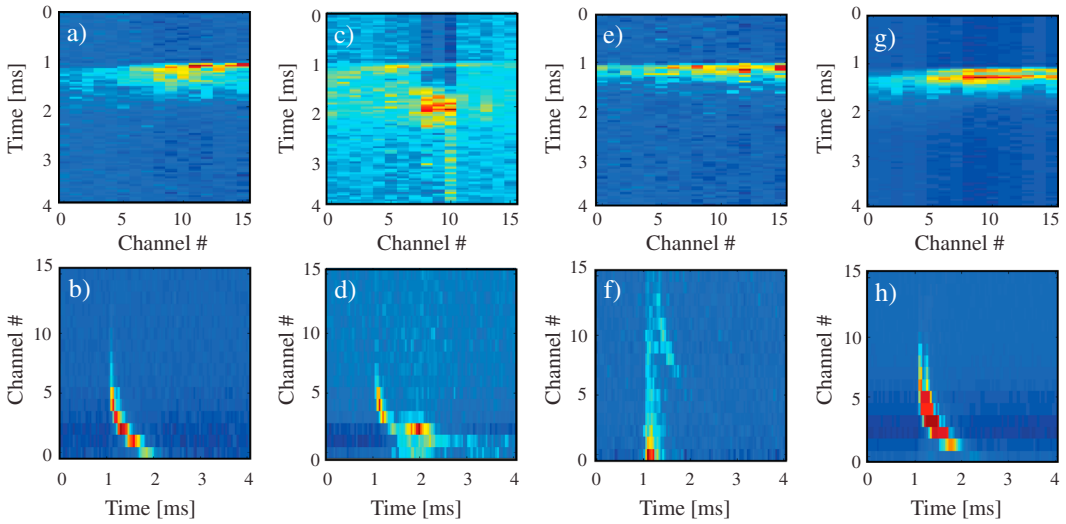

Figure 1. Example of elves observed by PIPER in 2013. (a, c, e, and g) The horizontal photometer channels, and (b, d, f, and h) shows the vertical channels. Figures 1a and 1b show a single elve. Figures 1c and 1d show an elve followed by a halo and sprite. Figures 1e and 1f show an elve preceded by a cloud flash. Figures 1g and 1h show an elve doublet event.

3. Properties of Elve-Producing Lightning

In this section, we use the data set of elve observations from PIPER in conjunction with VLF/LF sferics from the Oklahoma LF array to answer several questions about elve-producing lightning. We begin by observing the qualitative features of sferics from elve-producing lighting. Then, we use the features of the LF sferics to determine the important factors that discriminate between elve-producing and nonelve-producing lightning. Finally, we estimate the likelihood of elve production based on features that we observe in the lightning sferic waveform.

3.1. Determining Elve Causation

Before sferics from elve causative lightning strokes can be compared with those from nonelve causative strokes, a reliable method of labeling the two classes of sferics must be determined. For the analysis presented in this paper, each PIPER elve observation is matched with the causative NLDN CG event based on its time stamp and whether or not the discharge occurred within the PIPER photometer field of view. The sferic from the elve causative lightning stroke is then extracted from the broadband VLF/LF recordings based on the timestamp after accounting for the propagation delay. Determining that a lightning discharge did *not* produce an elve, however, is more difficult and great care is taken. We first create a list of all large (>50 kA) NLDN CG events that occurred during the same storm as a known elve causative discharge. Based on the latitude and longitude reported by NLDN, all events that occurred outside of the PIPER field of view are discarded. Then, we use coaligned CCD images to ensure that there was no cloud cover during these events (i.e., stars are visible in the images). Strokes reported by NLDN during times of cloud cover are ignored. Finally, for the remaining events, we review the PIPER data set to ensure that there was not actually an elve at that time that was previously missed. After performing these steps, we can be reasonably confident that we have obtained a set of sferics from large lightning discharges that did not produce elves—if they did occur, they would have been visible in the PIPER data.

Once sferics have been classified into elve and nonelve causative events, they still cannot be directly compared because they occur at different distances from the receiver. As a sferic propagates a distance r along the Earth's surface, it is attenuated by a factor of $\sim 1/r$ in the first ~1000 km. Furthermore, dispersion distorts the sferics due to the finite conductivity of the ground. Therefore, the shape of a sferic waveform changes significantly as it propagates, and we cannot reliably compare the features of sferics from different distances. Hence, we bin sferics based on their distance from the VLF/LF receiver before computing average waveforms and extracting features.

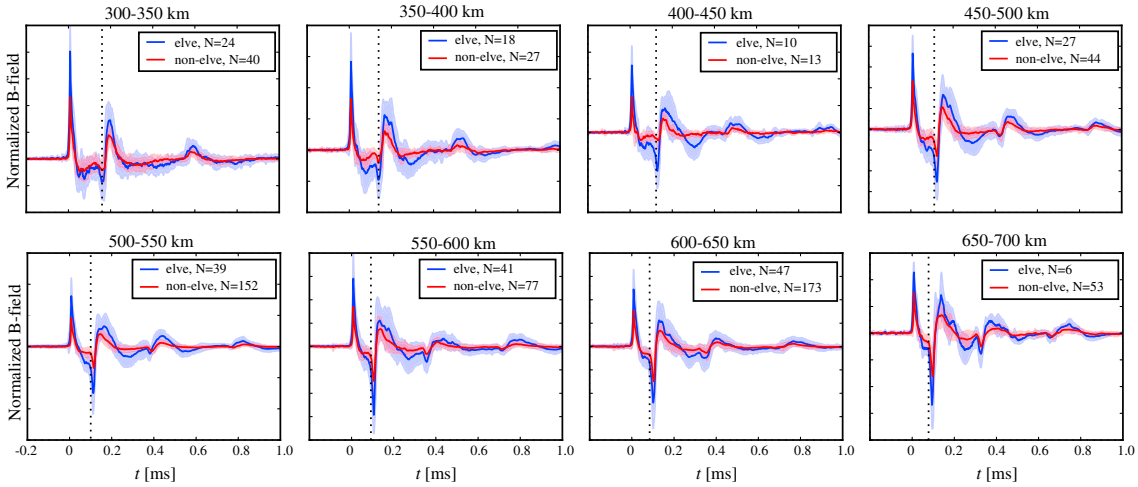

Figure 2. Distance binned average sferics from elve-producing and nonelve-producing lightning from multiple storms with ±1 standard deviations shown as shaded regions.

3.2. General Features of Sferics From Elve-Producing Lightning

The first sferic feature that we examine is the arrival time of the first ionospheric reflection, also called the first sky wave. Some authors have reported variations in the nighttime ionospheric density above thunderstorms, with the hypothesis that the height of the *D* region is measurably lowered over the course of a storm [*Shao et al.*, 2012; *Lay et al.*, 2014]. When this happens, the VLF reflection height is lower in altitude, so the first sky wave should arrive at the receiver earlier in time due to the shorter propagation path. The question naturally arises as to whether the ionosphere height may play a role in elve production—perhaps a lower ionosphere makes it more difficult to produce an elve since E/N will be lower. In Figure 2, we show average distanced-binned sferics from both +CG and −CG events for the two classes of lightning, with the ±1 standard deviations shown as shaded regions. A vertical, black dashed line is plotted where we would expect the sky wave to start if the ionosphere was a perfectly conducting flat plane at 88 km altitude. The first sky wave is the negative peak following the ground wave. We see that there is not a significant difference in the average arrival times of the sky waves for elve causative and nonelve causative lightnings, suggesting that these sferics propagated through similar ionospheres. Several caveats exist, however. The sferics in Figure 2 are averaged from events in multiple storms and hence do not convey any information about the temporal variation of the ionospheric conditions over the course of a storm or daily variations. Also, the ionosphere is not a perfect reflector, and it is difficult to pin down an exact reflection height since the lightning EMP is a broadband signal and different frequencies will reflect at different heights.

Next, we turn our attention to the portion of the EMP which propagates along the line of sight from the lightning channel to the receiver on the ground, known as the sferic ground wave. As it does not interact with the ionospheric plasma, this wave is directly related to the return stroke current propagation in the lightning channel. Assuming a flat earth and ignoring ground conductivity, the ground wave magnetic field is given by

$$\mathbf{B}(r,t) = \frac{\mu_0}{2\pi c r}\frac{\partial}{\partial t}\int_0^{\ell_{\text{chan}}} I\left(z', t - \frac{r}{c}\right) \mathrm{d}z'\,\hat{\phi} \tag{1}$$

where r is range along the ground, t is time, c is the speed of light, z' is elevation above the ground, and μ_0 is the permeability of free space.

When we compare the distance-binned ground waves from elve and nonelve causative lightnings, normalized by dividing by the peak of the ground wave, we find that across all of the distance bins, sferics from elve causative lightning exhibit a slightly longer rise-time.

In Figure 3, we bin the ground waves by their NLDN reported peak current rather than distance from the receiver. Propagation is accounted for by multiplying by r, the distance from the receiver. This effectively cancels out the factor of $\frac{1}{r}$ in equation (1). It should be noted that this is only an approximation; dispersion and attenuation due to finite ground conductivity and earth curvature also have a small effect on the shape of the

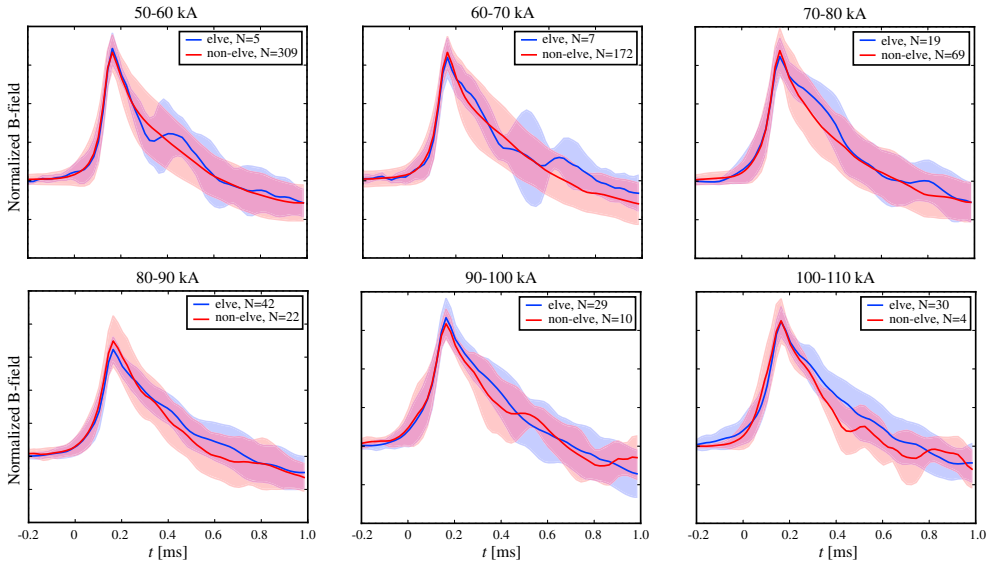

Figure 3. Range normalized average ground waves binned by NLDN peak current. We find that within each peak current bin, the two classes of lightning have approximately the same peaks and rise times.

ground wave. We observe that the difference in ground wave rise time disappears—the slope and rise times in Figure 3 match almost perfectly between the two classes. This observation is interesting because it suggests that the ground wave rise time is linearly correlated with the NLDN reported peak current, and hence with the peak magnitude of the radiated field.

4. Elve Production Probability Modeling

In this section, we develop a more rigorous statistical model for analyzing the differences between elve and nonelve causative lightnings. In particular, we train a logistic regression classifier model which is capable of estimating the probability that an elve was produced for a given observed sferic waveform.

4.1. Features of the Sferic Ground Wave

Five hand-picked waveform features are used as features in the logistic regression. These features are illustrated with a typical LF sferic in Figure 4.

First, we use the range normalized ground wave peak, also referred to as the range normalized signal strength (RNSS). This feature is measured by first multiplying the waveform by the range from the receiver to the stroke to account for the $1/r$ propagation decay and then searching for the peak. A parabola is fit to the three maximum points and the maximum computed.

The second feature used is the 10–90% rise time, τ_r. This is measured by first detecting the ground wave peak, and then computing the time it takes the magnetic field to go from 10% to 90% of that value in the preceding time segment. The 10–90% fall time is computed in a similar manner on the time segment immediately following the ground wave peak.

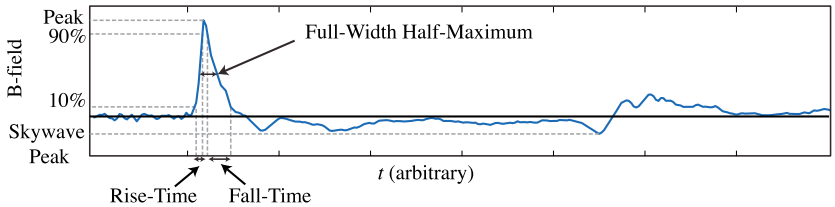

Figure 4. Example ground wave with the features labeled.

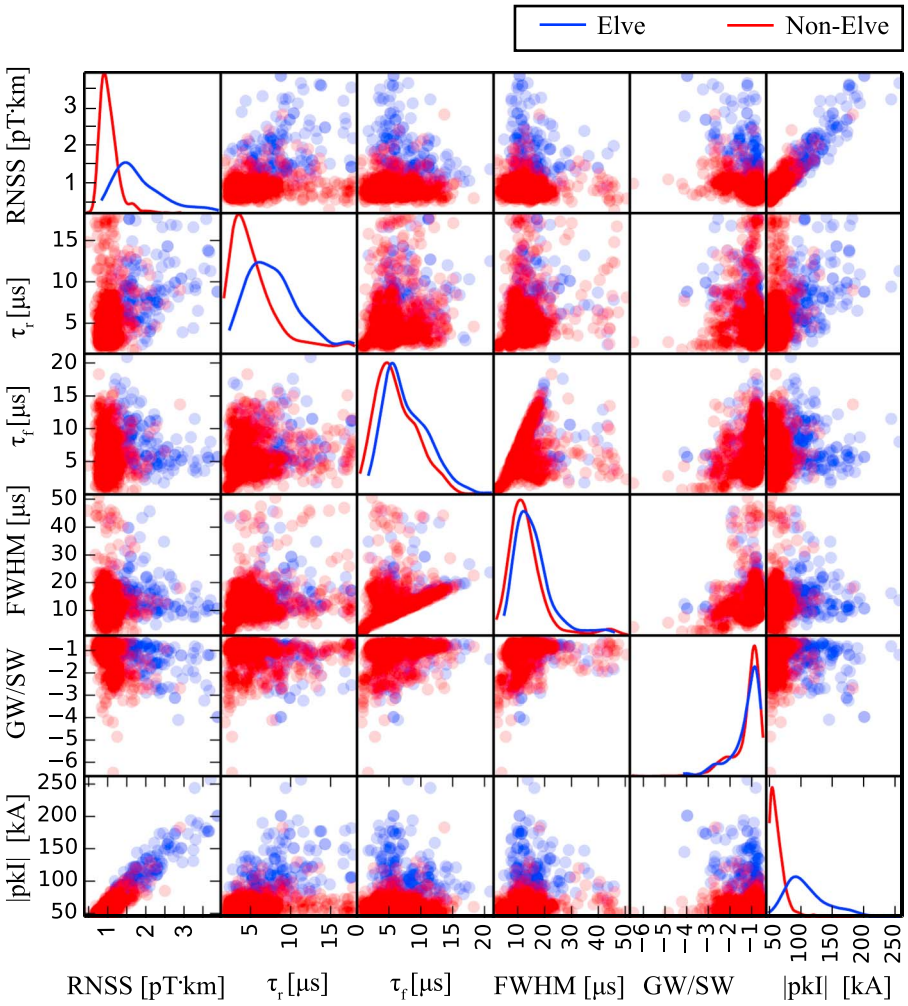

Figure 5. Feature scatter plot matrix. The diagonal plots show the KDE estimated probability distributions for each feature, while the off-diagonal plots show pairwise scatter plots for each pair of features.

The full-width half-maximum (FWHM) is determined by computing the time it takes for the waveform to go from 50% of the peak value, up to the peak, and then back down to 50%. This gives a sense of the "width" of the ground wave.

The ground wave to sky wave ratio (GW/SW) is computed by dividing the range normalized ground wave peak by the value of the first sky wave peak. The sky wave peak is found by searching for the smallest value in a time window determined by the range from the stroke and an assumed ionosphere height. The motivation for using this feature is to capture some information about the radiation pattern of the EMP. However, this feature is likely not very informative because the sky wave peak is heavily influenced by the ionospheric profile midway between the lightning stroke and the receiver, which we do not control for.

The NLDN peak current associated with the detected waveform is used as the final feature.

In Figure 5, we show a scatter plot matrix of the features extracted from our data set for both classes of lightning. The off-diagonals show pairwise scatter plots for each pair of features, while the diagonals show the probability distributions of each individual features. The distributions are estimated using a kernel density estimator (KDE) with a Gaussian kernel [*Parzen*, 1962; *Rosenblatt*, 1956]. This technique estimates probability density by smoothing a histogram with a kernel function, in this case a Gaussian. From this figure, we can gain a sense as to which features are important. For example, RNSS, τ_r, and peak current all have some discriminatory power, while the distributions for τ_f, FWHM and GW/SW are nearly identical for elve and nonelve causative lightning.

4.1.1. Logistic Regression Model

In order to model the production probability of elves given the observed ground wave features, we use a logistic regression. This is a binary classifier which maps our feature vector to a class probability of the form

$$p(\text{elve} \mid x) = \left[1 + \exp(-\theta^T x)\right]^{-1} \tag{2}$$

where x is the vector of sferic features and θ is a vector of model weights. For a detailed derivation of logistic regression, see, for example, *Bishop* [2006, chap. 4].

4.2. Results

Our logistic regression model is trained on the sferic features, and the classifier performance is evaluated using a k-fold cross-validation procedure with $k=3$. This involves randomly partitioning the data set into three subsets, training the model on two of the subsets, and evaluating the accuracy on the remaining 1/3. This is performed for each permutation of the subsets, and the accuracies from each of the three folds is averaged to provide a single estimate of the classifier accuracy. Using this approach provides a more robust estimate for how the classifier will generalize to new data.

When trained using the five waveform features discussed in section 4.1 in addition to the NLDN peak current, the logistic regression is able to correctly classify an event as being elve causative or nonelve causative with 91.7% accuracy. Remarkably, when trained on a single feature, NLDN peak current alone, the model is able to classify an event with 90.2% accuracy. The reasons for this are apparent after careful examination of Figure 5. Fall time, full-width half-maximum, and ground wave to sky wave ratio are distributed nearly identically for the two classes of lightning and thus do not provide much discriminatory information. All of the other features—RNSS, rise time, and peak current—are highly correlated. In fact, NLDN uses a variant of the RNSS, albeit using electric field measurements, along with transmission line return stroke assumptions to calculate the peak current estimate.

This peak current only logistic model is of the form

$$p\left(\text{elve} \mid I_p\right) = \frac{1}{1 + \exp(-\theta_1 |I_p| - \theta_0)} \tag{3}$$

where I_p is the peak current and θ_1 and θ_0 are learned model parameters.

In Figure 6a, we show the probability that an elve is produced given the peak current, as modeled by the logistic regression. The red circles show the empirical fraction of events that produced elves in peak current bins 15 kA wide. We see that the modeled probability provides a very reasonable fit with the empirical point estimates. This model predicts that CG return strokes with $|I_p| = 88$ kA have a 50% probability of producing elves, while those with $|I_p| = 106$ kA have a 90% probability. These results differ significantly from those of *Newsome* [2010], who found a significantly worse fit at larger peak currents and estimated that elves are produced with 50% probability at approximately 160 kA. The reasons for the differing results are unclear, but we suspect that more care was taken in producing our data set, in particular by excluding events outside of the PIPER field of view and during cloud cover. It is important to note that the peak currents provided by NLDN are only estimates based on a simple regression model and the peak observed fields and that there does not exists a ground truth to compare against and determine the absolute accuracy of the peak currents for naturally occurring lightning [*Cummins et al.*, 1998]. Hence, some of that uncertainty is transferred to our elve prediction model.

We examine a possible distance bias in Figure 6b by plotting each event's peak current as a function of distance from the PIPER photometer. The mean peak current is shown as a solid curve and the $\pm 1\sigma$ intervals are shown as shaded regions. We observe that the elve causative strokes have statistically significant higher peak current across all distances. This discounts the alternative hypothesis that PIPER was preferentially detecting large peak current events far from the photometer, thus biasing the results.

The conclusion that elves can be reliably predicted using only the NLDN peak current is extremely powerful, as it enables us to classify events when LF sferics and elve photometer measurements are not available. In particular, this result can be applied to the peak current estimate from global lightning detection networks, such as GLD360, to predict elve production outside of NLDN's coverage area. Further validation of this result is discussed in section 5.2 in which we find good agreement with satellite based elve observations from ISUAL.

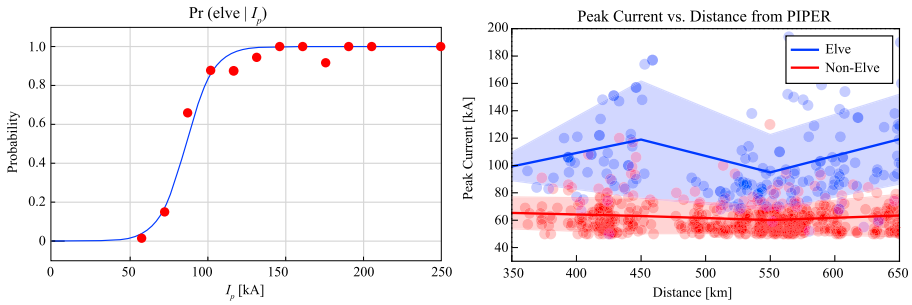

Figure 6. (a) Elve production probability as a function of NLDN peak current. In blue, we show the analytical probability provided by the logistic regression, while in red we show the empirical fraction of events which produced elves in several peak current bins. (b) NLDN peak currents of all elve and nonelve events in our data set, plotted as a function of distance. We see that elve causative strokes consistently have higher peak currents across all distances, suggesting that there is not a bias resulting in only observing the most intense events at long distances.

5. Global Elve Production Rates

The global production rate and geographic distribution of elves are not well understood. It is of particular importance for determining the amount of energy coupled between the lower troposphere, thunderstorms, into the ionosphere through heating by the return stroke EMP and injection of wave energy into the magnetosphere. In the following sections, we review prior studies on global elve production using satellite observations, and then we introduce our results using ground based measurements.

5.1. Prior Satellite Observations

Elves have been observed many times from Earth orbit, including their first observation in 1991 on board the Space Shuttle [*Boeck et al.*, 1992]. In particular, the Imager of Sprites and Upper Atmospheric Lightning (ISUAL) instrument [*Chern et al.*, 2003], a scientific payload on the Taiwanese FORMOSAT-2 satellite, observed thousands of elves (in addition to sprites and jets) over the course of multiple years. Over a 3 year period from July 2004 to June 2007, ISUAL observed 5434 elves, 633 sprites, 657 halos, and 13 gigantic jets, from which global distributions and occurrence rates for the various TLEs were estimated [*Chen et al.*, 2008a]. A map of the elve distribution and occurrence rates is shown in Figure 7a.

There are, however, a number of issues with the ISUAL data set which may bias its predicted elve occurrence rate. The FORMOSAT-2 satellite was in a Sun-synchronous orbit with an altitude of 890 km, and the ISUAL payload had a relatively narrow field of view with an eastward pointing direction. These limitations can be seen in the coverage and accumulative observation time map shown in Figure 7b. Furthermore, ISUAL had to be powered off during its passage through the South Atlantic Anomaly (SAA) to prevent radiation damage to the instruments. These factors introduce a significant observation bias which may affect the global rate estimate.

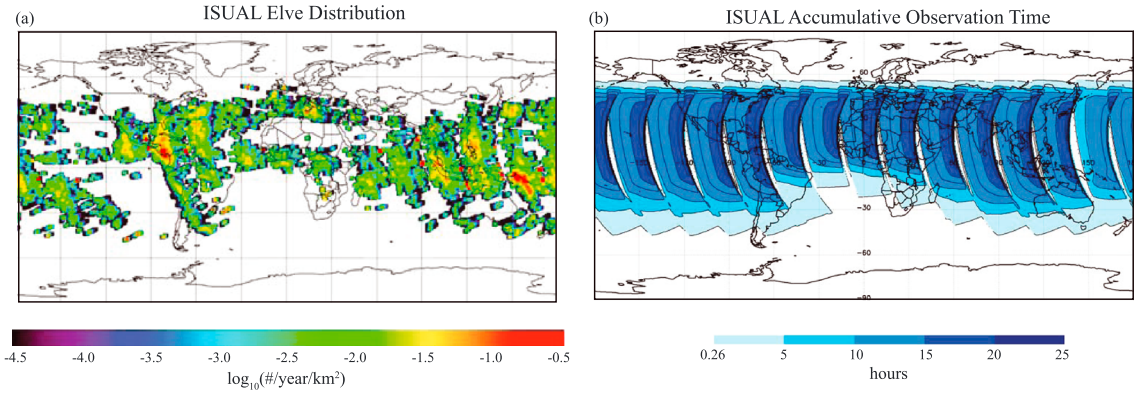

Figure 7. (a) The global elve distribution estimated from ISUAL satellite observations. (b) The accumulative observation time of ISUAL. We notice that the polar orbit produces significant gaps in coverage. Reproduced from *Chen et al.* [2008a].

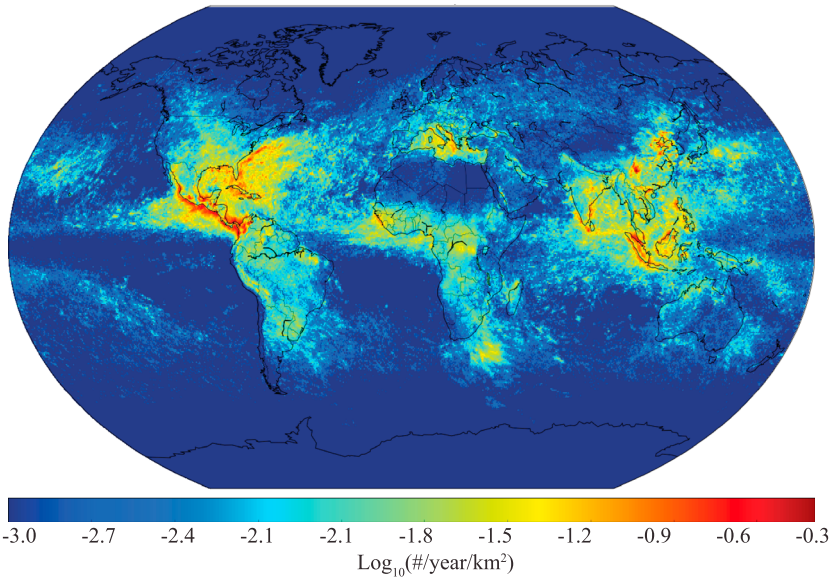

Figure 8. The estimated global distribution of elve production estimated using GLD360 peak currents. The probability that a return stroke with a given peak current will produce an elve is estimated using PIPER elve observations and NLDN peak current estimates.

ISUAL contained three sensors: a CCD camera, a six-channel spectrophotometer, and a 16-anode photometer array (similar to the one used in PIPER) which measured light variation at various altitudes. Due to limited onboard data storage and downlink bandwidth, the instruments could not be operated in a continuous acquisition, free-running manner. Instead, the spectrophotometer was used to trigger events based on emissions from lightning and TLEs. The triggering required an empirically tuned threshold so as to maximize collection of true TLEs while minimizing false positives. This triggering threshold inherently introduces a triggering bias into the data set: there may be many dim TLEs that ISUAL did not register as events, hence lowering the estimated global occurrence rates.

Due to these potential limitations and potential biases, it is desirable to develop a method for predicting global elve production using only continuous ground-based observations. A method for doing this is presented in the following section.

5.2. Production Rate Estimated From Ground Observations

In section 4, we compared the properties of elve and nonelve causative lightnings and find that NLDN peak current provides an accurate method for estimating the probability that a lightning stroke produced an elve. In particular, we trained a logistic model of the form shown in equation (3) which estimates the probability that an elve was produced given the NLDN reported peak current. In this section, we apply this model to the peak current estimates from another lightning geolocation network, GLD360.

GLD360 is a global network of VLF receivers maintained by Vaisala, Inc. which provides accurate global geolocation of CG lightning strokes. The network uses a combination of time of arrival along with correlation with waveform banks of average sferics to perform the geolocation [*Said et al.*, 2010]. In addition to location, GLD360 provides an estimate for the return stroke peak current by calibrating the received VLF peak magnetic fields with NLDN reported peak currents. We input GLD360 peak current estimates into equation (3) to predict whether a stroke produced an elve. This allows us to determine global elve production density because we are not limited by the NLDN coverage area.

The procedure we use to estimate the global elve production rate is as follows: (1) begin with a data set consisting of 1 year of GLD360 events which occurred during local night; (2) for each event, label it as elve-producing with probability $p\left(\text{elve} \mid I_p\right)$; (3) bin each elve causative event into a two-dimensional latitude-longitude histogram with $0.5° \times 0.5°$ bin sizes; and (4) normalize each bin by its surface area to obtain units of #/km^2/year.

The resulting distribution of elve production rates is shown in Figure 8. We estimate that 4,658,579 elves were produced in 2013, representing 0.82% of all GLD-recorded CG lightning events that year. When compared to

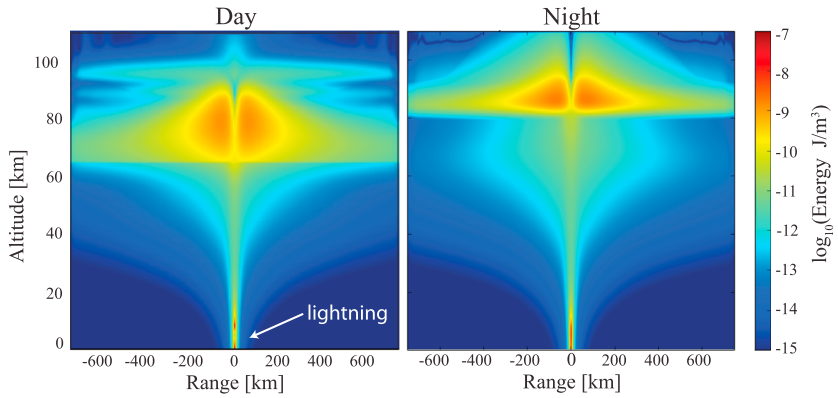

Figure 9. The simulated Joule heating of the *D* region due to a lightning EMP during the (left) day and at (right) night.

the empirical elve distribution from ISUAL observations shown in Figure 7, we see remarkable similarities. The occurrence rates are roughly on the same order of magnitude—the peak production rates occur at roughly 0.5 per year per km^2. Additionally, the locations of peak elve production occur in similar locations in Central America, the Caribbean, and Southeast Asia. The fact that our statistical model, trained only on events over North America within NLDN coverage, extrapolates to a global scale and is in such good agreement with satellite observations provides a good deal of confidence about the global validity of the model.

Several sources of uncertainty exist in this elve occurrence estimate. The peak current estimates from GLD360 exhibit a mean error of 21% when compared to NLDN and a ground flash detection efficiency of 67% for strokes with a peak current >15 kA [*Said et al.*, 2013]. Also, there exist slight variations in the height of the *D* region globally which may influence elve production. Furthermore, there exists uncertainty not captured in the simple logistic model due to global meteorological variation. For example, regions with higher average thunderstorm altitudes could potentially produce elves at lower peak currents, since the longer lightning channel would result in a larger current moment, and hence more intense EMP, for a given peak current. However, the overall good agreement with satellite observations suggests that the effects of these uncertainties may be minimal.

6. Heating of the *D* Region by Lightning

The upward propagating portion of the lightning EMP interacts with the ionospheric plasma, transferring energy through a collisional heating process, also called Joule heating [*Inan et al.*, 1997]. In this section we perform numerical simulations to estimate the amount of heating due to lightning EMPs from all CGs, not just elve causative strokes, and then extrapolate the results to find the global contribution of elve causative lightning to *D* region heating.

6.1. Numerical Simulations

To investigate the relationship between the return stroke and *D* region heating, we use the finite difference time domain EMP model from *Marshall* [2012] to simulate return strokes with a variety of parameters and to calculate the total heating from each one. The return stroke parameters that were varied include the peak current I_p, the return stroke speed v_{RS}, the channel length ℓ_{chan}, and the current rise time τ_r. These simulations are performed with IRI ionosphere profiles typical for day and night. Additionally, the daytime simulations were run three times with the IRI electron density profiles multiplied by $\frac{1}{3}$, 1, and 3 to account for ionospheric variability.

At each time step, the model computes the time integrated Joule heating $\mathbf{J}\cdot\mathbf{E}$ in every grid cell. Once the elve has dissipated, we are left with the spatial distribution of the total Joule heating associated with the EMP, $S(r, h)$, which is a function of range r and altitude h. Example heating distributions for both day and night are shown in Figure 9. The total energy transferred to the *D* region can then be computed by integrating vertically and radially

$$\text{Total Energy} = U = 2\pi \int_{h_{iono}}^{\infty} \int_{0}^{r_{max}} S(r,h) r dr dh \tag{4}$$

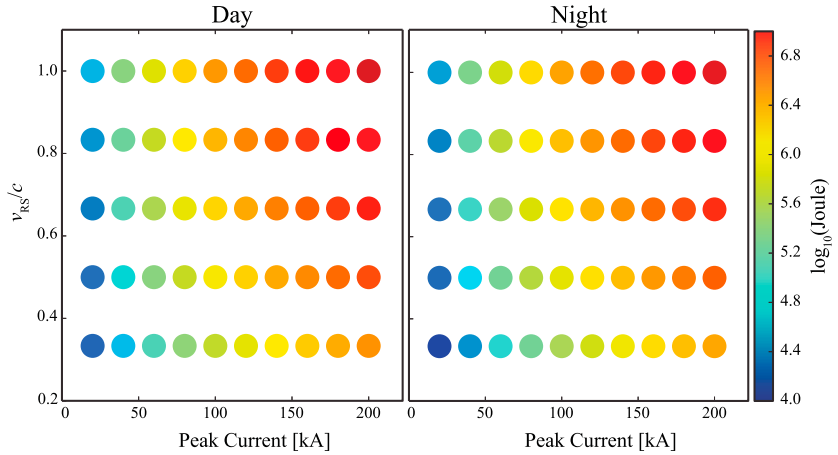

Figure 10. Simulated *D* region Joule heating as a function of peak current and return stroke speed for daytime and nighttime ionospheres.

where h_{iono} is taken to be 60 km during the day and 75 km at night and r_{max} is the maximum radial extent of the elve and is taken to be 750 km. We start integrating at these higher altitudes so as to avoid contributions due to the quasi-electrostatic (QES) field and the lightning channel itself.

Of the return stroke parameters that were varied in these simulations, the peak current and return stroke speed have the greatest impact on the resulting *D* region heating. The total energy as a function of these two parameters is shown in Figure 10. Larger peak currents result in greater heating, as intuitively expected. Return stroke speed plays a significantly smaller role, with faster return stroke speeds resulting in slightly more heating.

Using the result from *Blaes et al.* [2014] in which a maximum likelihood estimate for return stroke speed of 0.64*c* was found for elve causative lightning, we can extract a horizontal slice through the data in Figure 10 at 2/3*c* to obtain heating as a function of peak current only. This result is plotted for both day and night in Figure 11a along with quadratic best fit lines. The shaded red region shows the range in heating as the daytime electron density is varied from IRI/3 to IRI ×3, while the solid red line shows the heating for the unmodified IRI profile. Since this ionospheric variation leads to only a small variation in the total heating, using the quadratic fit for the typical IRI daytime profile is sufficient. Using this analytical quadratic fit, we can now estimate the total Joule heating in the *D* region due to a CG return stroke using only the return stroke peak current.

The question naturally arises as to what fraction of the total energy dissipated by a lightning stroke gets transfered to the *D* region ionosphere through heating. *Quick and Krider* [2013] used calibrated optical measurements of the visible and near-infrared radiation produced by CG lightning to find average optical energies

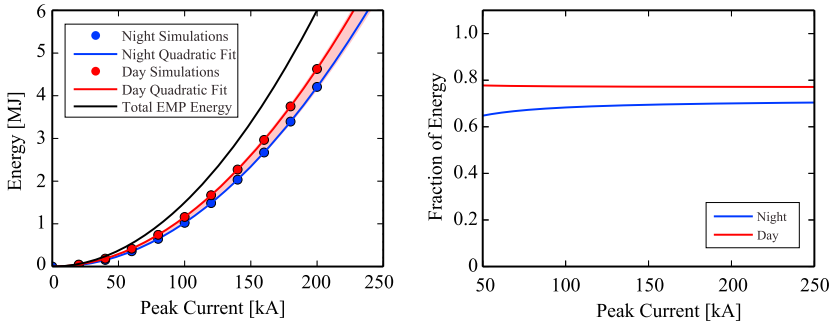

Figure 11. (a) Simulated *D* region Joule heating as a function of peak current. We find that this follows a simple quadratic relationship of the form $ax^2 + bx + c$ where x is the peak current. The points show the simulated heating results for day (red) and night (blue), while the solid lines show the quadratic best fit. The shaded region shows the range in daytime heating as the electron density profile is varied from IRI/3 to IRI ×3. (b) The fraction of total EMP energy transferred to the *D* region as a function of peak current. We find this to be relatively constant across peak currents with roughly 77% dissipated in the *D* region during the day and 70% at night.

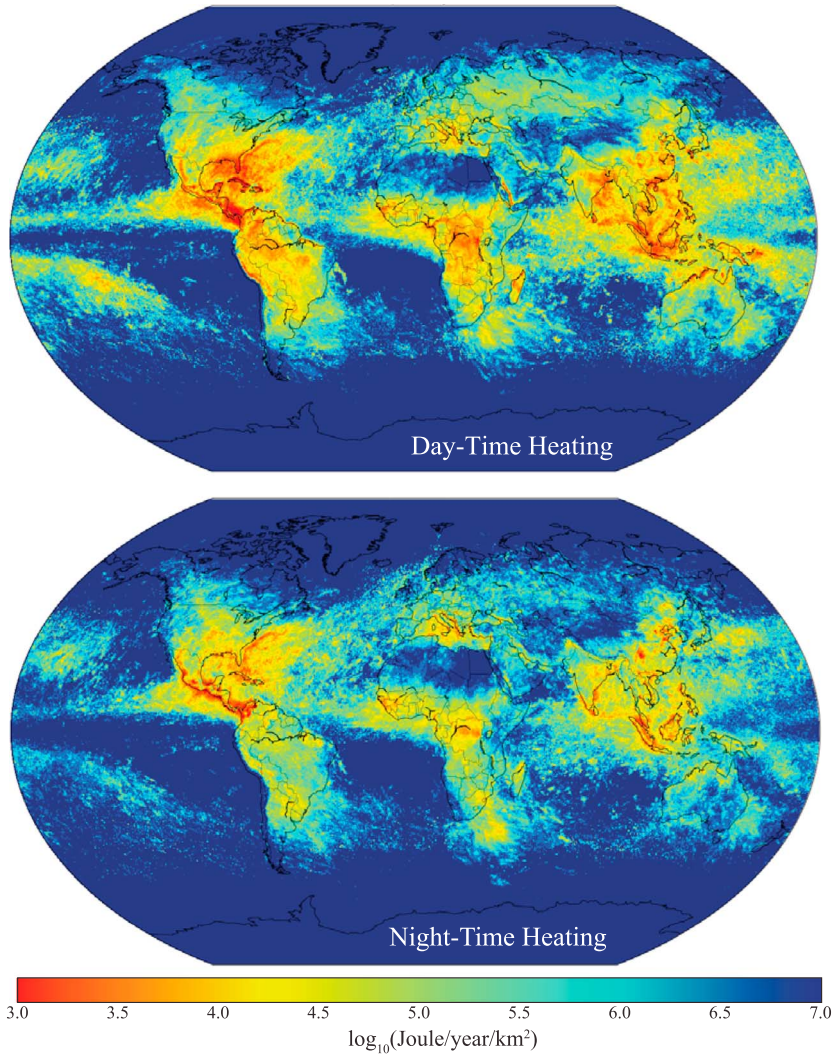

Figure 12. Here we show the estimated geographic distribution of one year's worth of *D* region heating at local (top) daytime and (bottom) nighttime.

of 3.6 MJ and 1.2 MJ for first strokes and subsequent strokes, respectively. *Krider and Guo* [1983] estimated the peak EMP power by assuming a small vertical dipole channel and then computing the surface integral of the Poynting flux through an enclosing hemisphere, finding

$$P_{\text{peak}} = \int_S (\mathbf{E} \times \mathbf{H}) \cdot \hat{\eta} \mathrm{d}S = \frac{4\pi}{3} R^2 \frac{E_p^2}{\mu_0 c} \tag{5}$$

where R is range along the ground and $E_p \simeq I_p/3$ is the peak electric field at a range of 100 km. If we approximate the return stroke current as a triangular current pulse, the width of which is the fall time τ_f, then the total EMP energy is given by

$$U_{\text{EMP}} = \frac{\tau_f P_{\text{peak}}}{2} \tag{6}$$

This total EMP energy is computed for a range of peak currents and is plotted as the black line in Figure 11a. As expected, we find that this curve lies slightly above our *D* region heating estimates for both daytime and nighttime ionosphere profiles. The ratio of the heating energy to the total EMP energy is shown in Figure 11b. We find this fraction to be approximately constant across the range of peak currents, with approximately 77%

of the EMP energy being dissipated in the ionosphere during the day and 70% at night. This result suggests that the majority of the total lightning-radiated energy is actually transferred to the *D* region plasma by the causative lightning stroke.

6.2. Global Heating Estimate

Using an approach similar to the one in section 5.2, we can utilize the peak current estimates from GLD360 to estimate the global contribution of lightning EMPs to *D* region heating. Using one year's worth of GLD360 CG events, we label each stroke as occurring in either day or night based on whether the Sun was visible at *D* region altitudes directly above the lightning. Subsequently, using the peak current and the quadratic fits shown in Figure 11, the amount of energy transferred to the *D* region plasma is estimated. These heating results are then integrated into latitude-longitude bins and normalized by area to give units of J/km^2/year. In Figure 12, we show the geographic distribution of the EMP-driven heating. Notably, the locations of the nighttime "hot spots" are highly correlated with the regions of maximum elve production shown in Figure 8, as expected.

Next, we integrate the power density distributions shown in Figure 12 over the surface area of the Earth to obtain the total energy deposition due to lightning EMP-driven *D* region heating. Doing so, we estimate that CG lightning transfers a total of 61.1 TJ of energy into the *D* region ionosphere due to Joule heating each year, with 39.9 TJ transferred during the day and 21.2 TJ at night. This deposition corresponds to an average power of 1.94 MW.

It is worth comparing this energy dissipation to other forms of upper atmospheric heating, through lightning or other means. *Mareev et al.* [2006] estimate that sprites are responsible for dissipating a maximum of 10 MW on average in the global electric circuit. However, both of these estimates of the energy transfer by the lightning EMP and QE fields (in the case of sprites) are several orders of magnitude lower than reported auroral energy depositions rates. Observations using the Ultraviolet Imager on board the Polar spacecraft showed rates of auroral energy deposition by the solar wind as low as 1 GW during the night and varying between 4 and 10 GW during the day, even during times of relatively quiet solar activity ($0 \leq K_p \leq 1$) [*Brittnacher et al.*, 1997]. However, while auroral energy is deposited primarily at higher latitudes, the energy transfer due to lightning is concentrated at equatorial and low latitudes, and in particular, geographic regions where lightning is most frequent.

It should also be noted that the 1.94 MW power estimate only includes the effects of CG lightning, due to the limitations of the two-dimensional numerical model used. This model is only capable of simulating vertical lightning channels as a result of the spherical symmetry of the simulation grid. CG lightning has significantly higher peak currents, and thus more intense EMPs, than IC lightning but only comprises approximately 25% of total lightning activity. Hence, the total lightning EMP driven *D* region heating may be higher if the contribution from IC strokes is included.

7. Summary

In this paper we have used PIPER observations of elves and VLF/LF sferic data to examine the detailed properties of elve causative lightning. We compare average waveforms from the two classes of lightning, finding that elve causative ground waves tend to exhibit a longer 10–90% rise time for a given distance from the receiver. However, when we control for peak current, the difference in the rise times disappears. Next we describe how features are extracted from the sferics and how a logistic regression is used to model elve production probability given those features. Doing this, we find that we are able to classify the events in our data set with 90.2% accuracy using NLDN peak current alone.

We use this logistic regression classifier to classify a year's worth of GLD360 events. This allows us to estimate global elve occurrence rates based on our PIPER observations, despite the fact that the PIPER data were collected exclusively in North America. Our resulting global occurrence map agrees well with satellite-based measurements from the ISUAL instrument, providing confidence that our elve production probability model using peak current is correct and applicable outside of North America. We find that 0.8% of all lightning located by GLD360 produces elves, amounting to 4.7×10^6 total elves per year. Next, we perform numerical simulations using the EMP model to find the total Joule heating in the *D* region due to the lightning EMP as a function of peak current. Joule heating in our model follows a quadratic relationship with peak current, and

we find that 70% of the total radiated EMP energy is transferred to the *D* region at night and 77% during the day. These results are then applied to the GLD360 data set to provide a global estimate and spatial distribution for *D* region heating. We find that lightning EMPs transfer approximately 60 TJ of energy per year into the *D* region through collisional heating, or 2 MW on average.

Acknowledgments
This work was supported by DARPA grant HR0011-10-1-0058 under the NIMBUS program and by NSF CEDAR grant AGS-1243176. We thank Vaisala Inc. and Ryan Said for access to GLD360 lightning location data. The data used for this study are available upon request from the authors.

References

Barrington-Leigh, C. P., U. S. Inan, M. Stanley, and S. A. Cummer (1999), Sprites triggered by negative lightning discharges, *Geophys. Res. Lett.*, *26*(24), 3605–3608.

Bishop, C. M. (2006), *Pattern Recognition and Machine Learning*, Springer, New York.

Blaes, P. R., R. A. Marshall, and U. S. Inan (2014), Return stroke speed of cloud-to-ground lightning estimated from elve hole radii, *Geophys. Res. Lett.*, *41*, 9182–9187, doi:10.1002/2014GL062392.

Boeck, W. L., O. H. Vaughan Jr., R. Blakeslee, B. Vonnegut, and M. Brook (1992), Lightning induced brightening in the airglow layer, *Geophys. Res. Lett.*, *19*(2), 99–102.

Brittnacher, M. J., R. Elsen, G. K. Parks, L. Chen, G. A. Germany, and J. F. Spann (1997), A dayside auroral energy deposition case study using the Polar ultraviolet imager, *Geophys. Res. Lett.*, *24*(8), 991–994.

Chen, A. B., C. L. Kuo, Y. J. Lee, and H. T. Su (2008a), Global distributions and occurrence rates of transient luminous events, *J. Geophys. Res.*, *113*, A08306, doi:10.1029/2008JA013101.

Chen, A. B., et al. (2008b), Global distributions and occurrence rates of transient luminous events, *J. Geophys. Res.*, *113*, A08306, doi:10.1029/2008JA013101.

Chern, J. L., R. R. Hsu, H. T. Su, and S. B. Mende (2003), Global survey of upper atmospheric transient luminous events on the ROCSAT-2 satellite, *J. Atmos. Sol. Terr. Phys.*, *65*, 647–659.

Christian, H. J., et al. (2003), Global frequency and distribution of lightning as observed from space by the optical transient detector, *J. Geophys. Res.*, *108*(D1), 4005, doi:10.1019/2002JD002347.

Cohen, M. B., U. S. Inan, and E. W. Paschal (2010), Sensitive broadband ELF/VLF radio reception with the AWESOME instrument, *IEEE Trans. Geosci. Remote Sens.*, *48*, 3–17.

Cummins, K. L., M. J. Murphy, E. A. Bardo, W. L. Hiscox, R. B. Pyle, and A. E. Pifer (1998), A combined TOA/MDF technology upgrade of the U.S. National lightning detection network, *J. Geophys. Res.*, *103*, 9035–9044, doi:10.1029/98JD00153.

Fukunishi, H., Y. Takahashi, M. Kubota, K. Sakanoi, U. S. Inan, and W. A. Lyons (1996), Elves: Lightning-induced transient luminous events in the lower ionosphere, *Geophys. Res. Lett.*, *23*, 2157–2160, doi:10.1029/96GL01979.

Holle, R. L. (2008), Annual rates of lightning fatalities by country, paper presented at 20th International Lightning Detection Conference, pp. 1–14, Tuczon, Ariz., 21–23 Apr.

Inan, U. S., T. F. Bell, and J. V. Rodriguez (1991), Heating and ionization of the lower ionosphere by lightning, *Geophys. Res. Lett.*, *18*(4), 705–708.

Inan, U. S., C. P. Barrington-Leigh, S. Hansen, V. S. Glukhov, T. F. Bell, and R. Rairden (1997), Rapid lateral expansion of optical luminosity in lightning-induced ionospheric flashes referred to as 'elves', *Geophys. Res. Lett.*, *24*(5), 583–586.

Krider, E. P., and C. Guo (1983), The peak electromagnetic power radiated by lightning return strokes, *J. Geophys. Res.*, *88*(3), 8471–8474.

Lay, E. H., X.-M. Shao, and A. R. Jacobson (2014), D region electron profiles observed with substantial spatial and temporal change near thunderstorms, *J. Geophys. Res. Space Physics*, *119*, 4916–4928, doi:10.1002/2013JA019430.

Mareev, E. A., A. A. Evtushenko, and S. A. Yashunin (2006), On the modeling of sprites and sprite-producing clouds in the global electric circuit, in *Sprites, Elves and Intense Lightning Discharges, NATO Sci. Ser.*, vol. 225, edited by M. Fullekrug, E. Mareev, and M. Rycroft, pp. 313–340, Springer, Netherlands.

Marshall, R. A. (2012), An improved model of the lightning electromagnetic field interaction with the D-region ionosphere, *J. Geophys. Res.*, *117*, A03316, doi:10.1029/2011JA017408.

Marshall, R. A. (2014), Effect of self-absorption on attenuation of lightning and transmitter signals in the lower ionosphere, *J. Geophys. Res. Space Physics*, *119*, 4062–4076, doi:10.1002/2014JA019921.

Marshall, R. A., R. T. Newsome, and U. S. Inan (2008), Fast photometric imaging using orthogonal linear arrays, *IEEE Trans. Geosci. Remote Sens.*, *46*(11), 3885–3893.

Marshall, R. A., C. L. da Silva, and V. P. Pasko (2015), Elve doublets and compact intracloud discharges, *Geophys. Res. Lett.*, *42*, 6112–6119, doi:10.1002/2015GL064862.

Newsome, R. (2010), Ground-based photometric imaging of lightning EMP-induced transient luminous events, Thesis, Stanford Univ., Stanford, Calif.

Newsome, R. T., and U. S. Inan (2010), Free-running ground-based photometric array imaging of transient luminous events, *J. Geophys. Res.*, *115*, A00E41, doi:10.1029/2009JA014834.

Parzen, E. (1962), On estimation of a probability density function and mode, *Ann. Math. Stat.*, *33*(3), 1065–1076.

Quick, M. G., and E. P. Krider (2013), Optical power and energy radiated by natural lightning, *J. Geophys. Res.*, *118*(4), 1868–1879.

Rakov, V. A., and M. A. Uman (2003), *Lightning: Physics and Effects*, Cambridge Univ. Press, Cambridge, U. K.

Rosenblatt, M. (1956), Remarks on some nonparametric estimates of a density function, *Ann. Math. Stat.*, *27*(3), 832–837.

Said, R. K., U. S. Inan, and K. L. Cummins (2010), Long-range lightning geolocation using a VLF radio atmospheric waveform bank, *J. Geophys. Res.*, *115*, D23108, doi:10.1029/2010JD013863.

Said, R. K., M. B. Cohen, and U. S. Inan (2013), Highly intense lightning over the oceans: Estimated peak currents from global GLD360 observations, *J. Geophys. Res. Atmos.*, *118*, 6905–6915, doi:10.1002/jgrd.50508.

Shao, X.-M., E. H. Lay, and A. R. Jacobson (2012), Reduction of electron density in the night-time lower ionosphere in response to a thunderstorm, *Nat. Geosci.*, *6*(1), 29–33, doi:10.1038/ngeo1668.

Wait, J. R., and K. P. Spies (1964), Characteristics of the Earth-ionosphere waveguide for VLF radio waves, *Tech. Note 300*, National Bureau of Standards, Boulder, Colo.