Posted on 24 Nov 2022 — The copyright holder is the author/funder. All rights reserved. No reuse without permission. — https://doi.org/10.1002/essoar.10506273.1 — This a preprint and has not been peer reviewed. Data may be preliminary.

QuakeCast, an Earthquake Forecasting System Using Ionospheric Anomalies and Machine Learning

Jessica Reid[1], Jeffrey Liu[1], and Bhavani Ananthabhotla[1]

[1]MIT Lincoln Laboratory,Massachusetts Institute of Technology

November 24, 2022

Abstract

Distribution: A. Approved for public release – distribution is unlimited QuakeCast is a novel system for short-term earthquake prediction using global ionosphere Total Electron Content (TEC) data. In the last 20 years, earthquakes have caused over 800,000 deaths and $650 billion in economic damage. While current seismic early warning systems provide up to a minute's notice by detecting an earthquake's non-damaging P-waves prior to the damaging S-waves, earlier warnings could allow more significant emergency preparations. Electromagnetic ionospheric phenomena have recently been observed many days before major earthquakes, notably in 2015 in Nepal. QuakeCast explores whether such signals could be used to predict earthquakes in a global dataset of ionosphere and earthquake data. The ionosphere data consists of global TEC data from NASA's Crustal Dynamics Data Information System (CDDIS), for the years 2005-2015 at 15 minute intervals. The earthquake data consists of over 10,000 events of at least M5 from the International Seismological Centre (ISC-GEM) Global Instrumental Earthquake Catalogue over the same time period. We used the dataset to explore the following classification problem: given a 24-hour sequence of ionosphere TEC data in a 30° latitude by 30° longitude window, is the sequence preseismic (an earthquake occurs in the next time step) or nominal (no earthquake)? We built two models to address this. The first is a classical logistic regression model trained on radially binned data. Analysis of the classifier weights indicate a localized effect in the ionosphere: TEC data spatially and temporally closer to the event contributes more significantly to the prediction. The second model is a deep learning ConvLSTM autoencoder trained to reproduce nominal TEC data sequences. Since autoencoders reproduce data which resemble their training data better than data which does not, we used the reconstruction error to classify sequences as anomalous (preseismic) or nominal. Both methods were found to perform significantly better than a random null classifier, indicating that the ionosphere contains information useful for predicting earthquakes at least 15 minutes in advance. While further research is needed to incorporate additional data features and address noise, we believe that these results are a promising development toward forecasting earthquakes using geoelectric signals.

QuakeCast: Earthquake Forecasting Using Machine Learning to Identify and Classify Preseismic Ionospheric Anomalies as Earthquake Signals

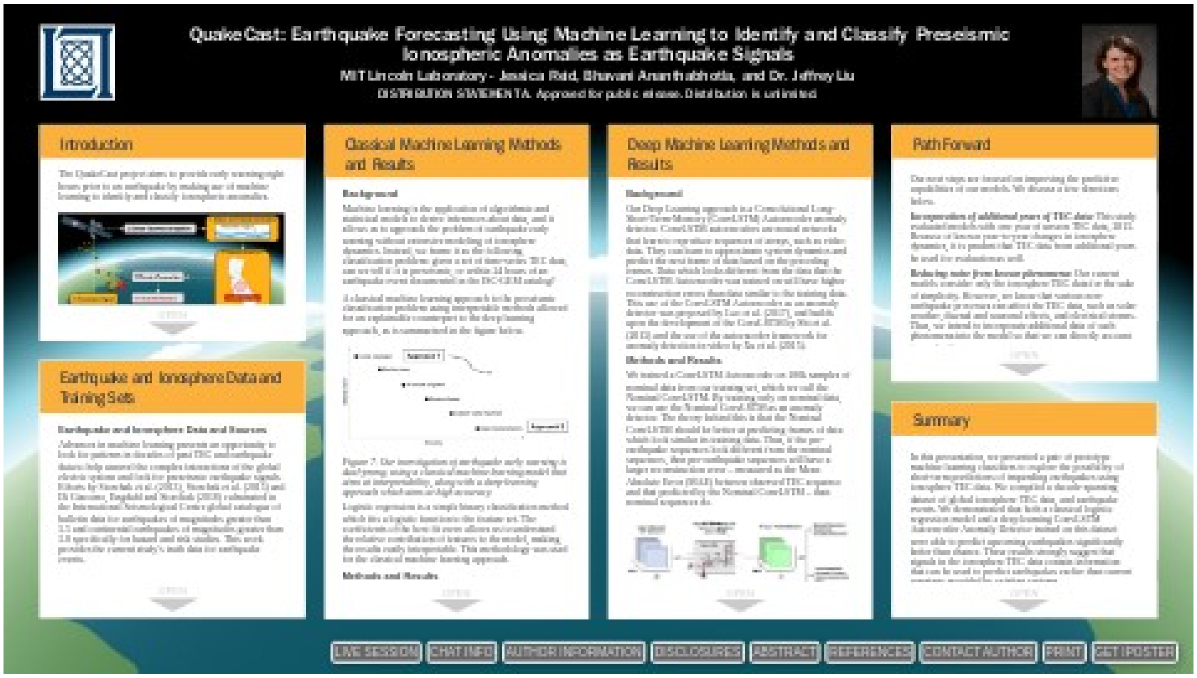

MIT Lincoln Laboratory - Jessica Reid, Bhavani Ananthabhotla, and Dr. Jeffrey Liu

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

PRESENTED AT:

INTRODUCTION

The QuakeCast project aims to provide early warning eight hours prior to an earthquake by making use of machine learning to identify and classify ionospheric anomalies.

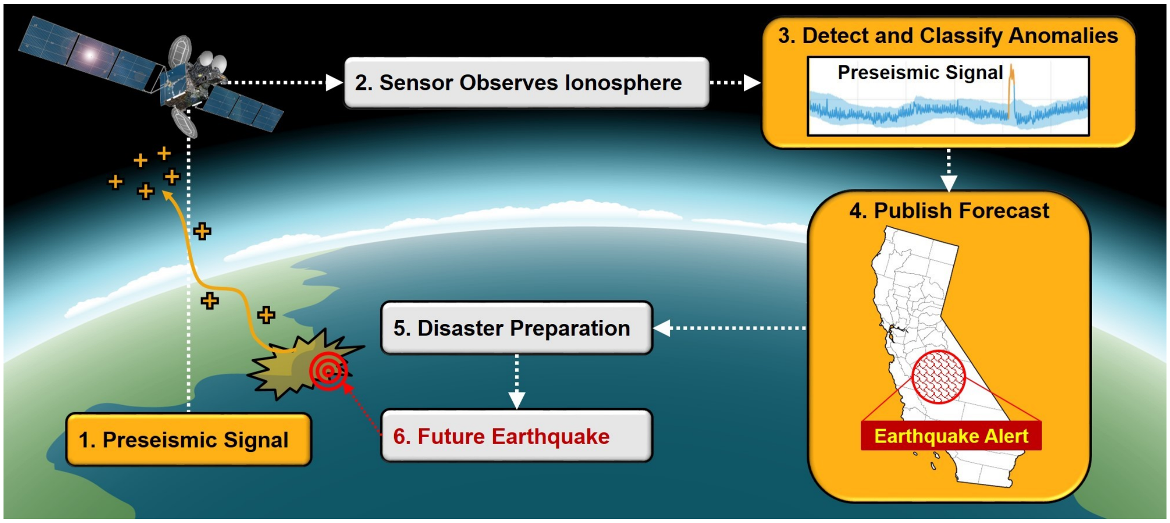

Figure 1: QuakeCast forecasting steps.

(1) Preseismic Signal

Anomalous disturbances in total electron content (TEC) in the ionosphere have been observed preceding strong earthquakes greater than magnitude 5 (Liu 2018, Zhou 2017, Zolotov 2012). Some theories suggest that microcracking in the lithosphere during preseismic processes generates stress-activated electrical currents due to lattice structure defects in the igneous rocks.

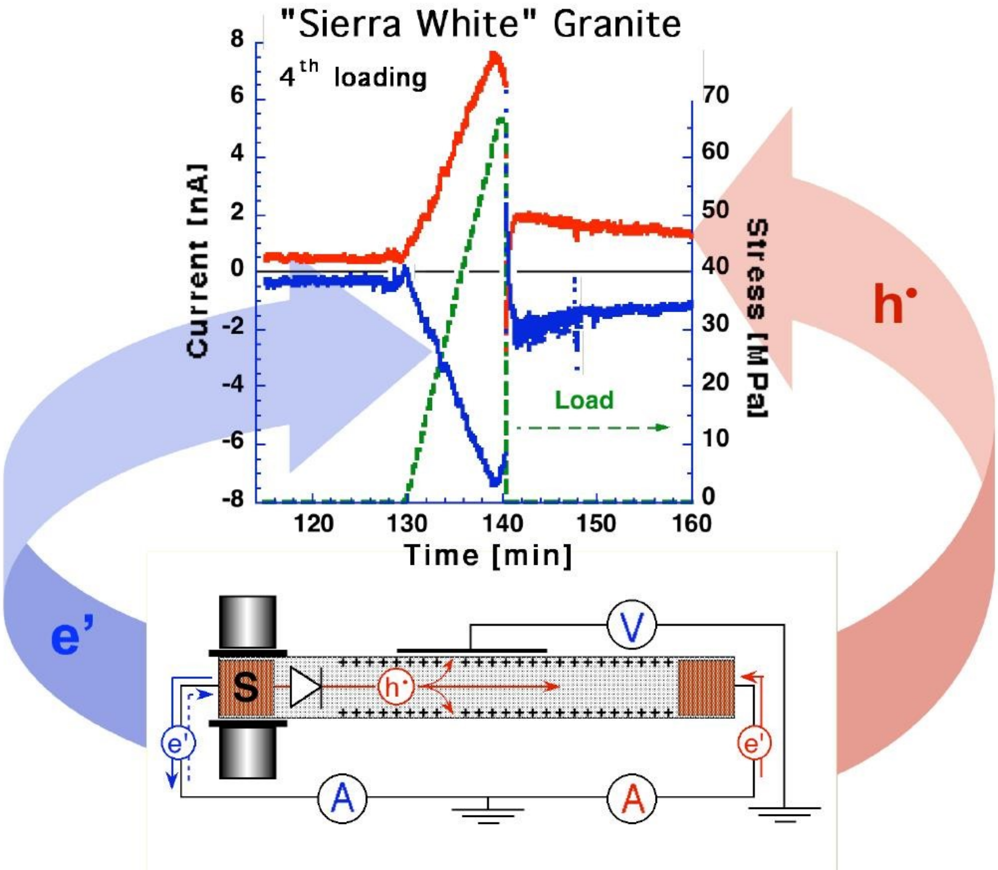

Figure 2: Laboratory testing of granite blocks, such as those found in the lithosphere, demonstrated a current flow when stressed (Freund 2005 - Figure 2)

The ionosphere is coupled to these electrical currents through the global electric circuit and results in a rapid change in TEC above the affected region (Freund 2005, Kuo 2018, Molchanov 2004).

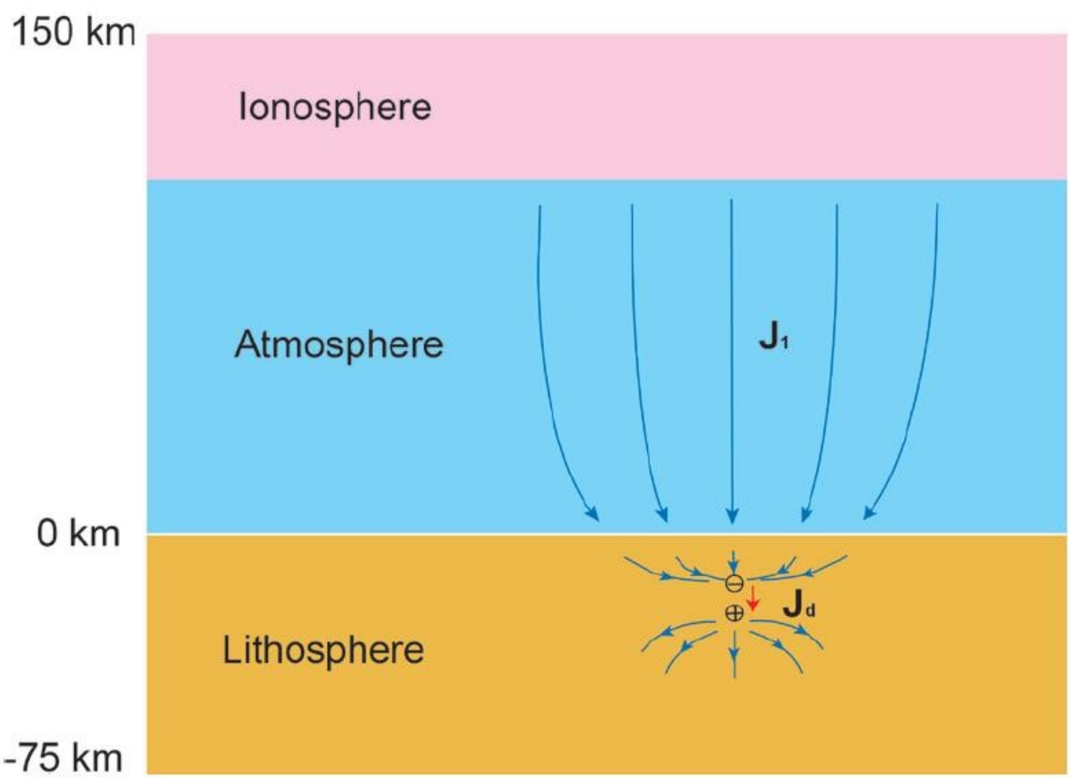

Figure 3: Current flow between the ionosphere, atmosphere, and lithosphere due to the electric current generated in the lithosphere, Jd, by preseismic microcracking (Kuo 2018 - Figure 1).

(2) Sensor Observes Ionosphere

The Global Positioning System (GPS) monitors the ionosphere TEC globally by measuring the time delay between the L1 and L2 frequencies between satellites and ground receivers. TEC data collected by GPS can be monitored for anomalous changes due to preseismic activity.

(3) Detect and Classify Anomalies

QuakeCast is developing machine learning algorithms to identify these TEC anomalies and classify them as a preseismic earthquake signal.

(4) Publish Forecast

By continuously monitoring the TEC above earthquake prone regions, QuakeCast aims to provide earthquake forecasts that will warn the area of an impending event and provide a projected risk region, event window, and event magnitude.

(5) Disaster Preparation

Based on the earthquake forecast, the region can prepare by staging supplies, evacuating people, and putting some infrastructure into safe modes. This will save lives and reduce financial losses.

(6) Future Earthquake

The earthquake occurs in the projected region and window at the projected magnitude.

In the following sections, we will present two types of prototype machine learning classifiers to explore the possibility of short-term predictions of impending earthquakes and introduce the earthquake and ionosphere data set used to train these classifiers.

EARTHQUAKE AND IONOSPHERE DATA AND TRAINING SETS

Earthquake and Ionosphere Data and Sources

Advances in machine learning presents an opportunity to look for patterns in decades of past TEC and earthquake data to help unravel the complex interactions of the global electric system and look for preseismic earthquake signals. Efforts by Storchak et al. (2013), Storchak et al. (2015) and Di Giacomo, Engdahl and Storchak (2018) culminated in the International Seismological Centre global catalogue of bulletin data for earthquakes of magnitudes greater than 5.5 and continental earthquakes of magnitudes greater than 5.0 specifically for hazard and risk studies. This work provides the current study's truth data for earthquake events.

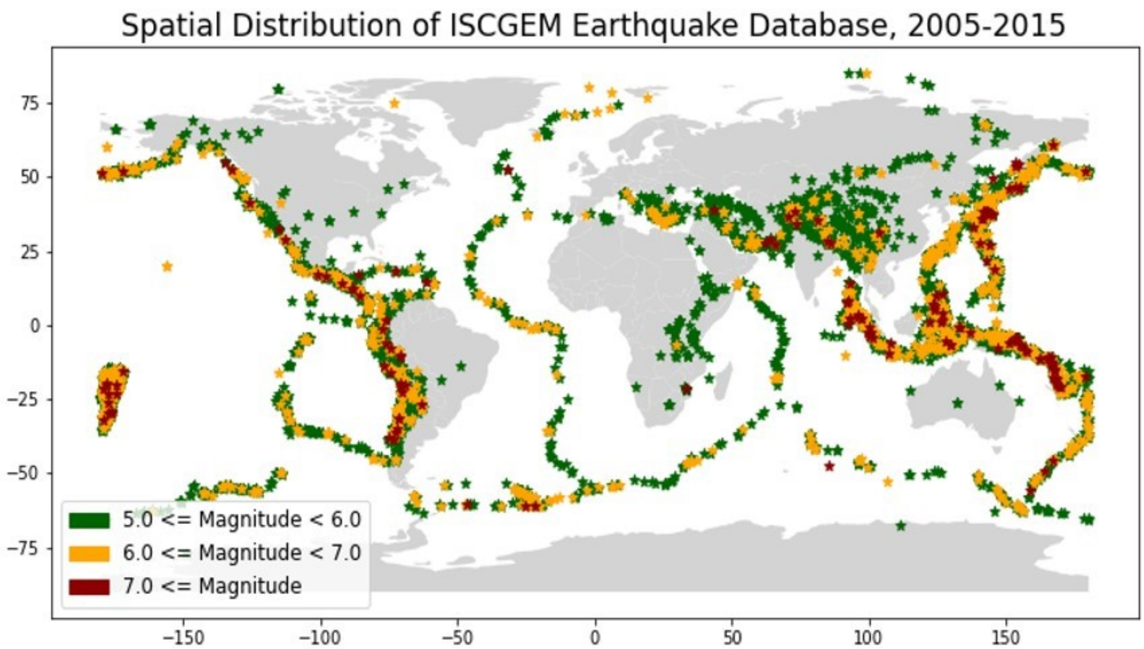

Figure 4: Spatial Distribution of earthquakes in the ISC-GEM Global Instrumental Earthquake Catalogue between 2005 and 2015.

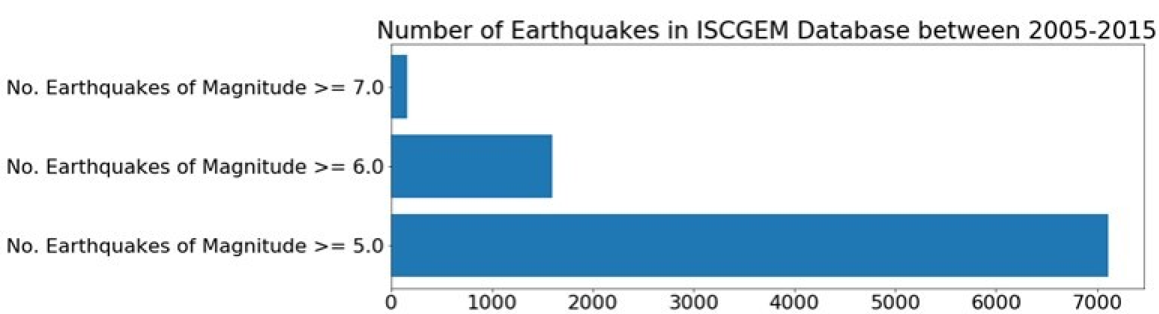

Figure 5: Distributions of magnitude of earthquakes between 2005 and 2015 in ISCGEM.

Global time-series vertical TEC data is available at temporal resolution of every 15-minutes, and a spatial resolution of 2.5-degree latitude, 5-degree longitude, as a GNSS Atmospheric Data Product from NASA CDDIS (Noll 2010). Data between 2005 and 2015 was used for this study.

The ionosphere is not a static medium and TEC dynamics are caused by known conditions; time of day, time of year, sun activity, and even location. For example, as the sun rises the ionosphere is energized and the TEC rises, after the sun sets the ionosphere is de-energized and the TEC falls. Below is a visualization of the ionosphere TEC on a nominal day where no > 5.0 magnitude earthquakes occurred, the TEC rises as the sun rise terminator travels over the globe.

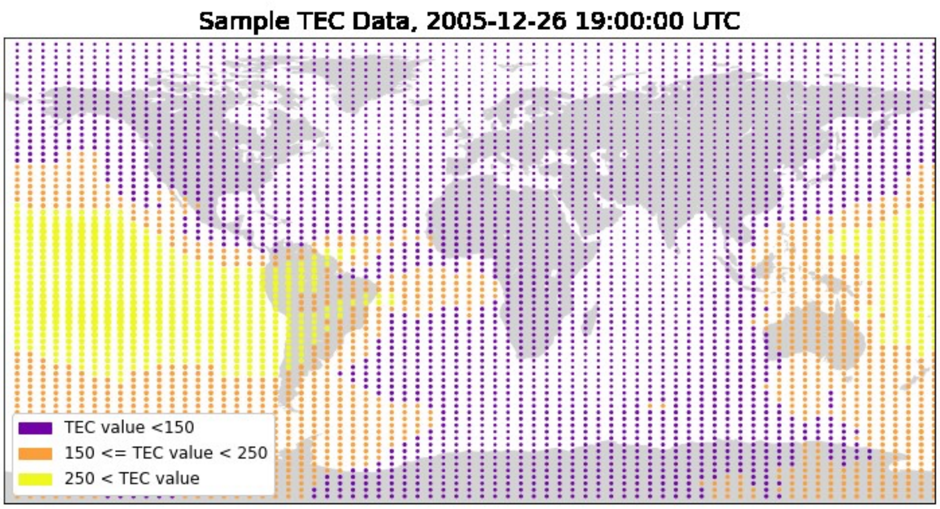

Figure 6: Ionosphere TEC rising as sun rise terminator travels over the globe.

Data Management and Training Sets

We considered 24-hour sequences of 30°×30° latitude/longitude windows of TEC data. The TEC data occasionally had missing observations and time steps; we considered only sequences that had a full 24-hours of observations.

We classified sequences into two categories: nominal, and pre-earthquake. Nominal data is defined as sequences which had no earthquakes within the spatial and temporal window, nor in the subsequent 24-hours. For the pre-earthquake sequences, we further defined sequences as X-lookahead pre-earthquake sequence, where X is a time interval. An X-lookahead pre-earthquake sequence means that an earthquake event occurs within the spatial window between X to X+15 minutes after the end of the sequence. In this poster, we discuss results for 15 minute- and 8 hour-lookahead pre-earthquake sequences.

The data were normalized with mean 0 and variance 1, estimated from a random sample of 10k nominal sequences. We partition the data into a train/test split based on the timestamp of the data. We used the data from the years 2005-2014 as the training set, and those from 2015 as the test set. This is done to simulate a forward-looking prediction: had we made the predictions based only on the data from 2005-2014, how would it have performed in 2015? The following sections show the methods and results from two different machine learning approaches, classical and deep learning.

CLASSICAL MACHINE LEARNING METHODS AND RESULTS

Background

Machine learning is the application of algorithmic and statistical models to derive inferences about data, and it allows us to approach the problem of earthquake early warning without extensive modeling of ionosphere dynamics. Instead, we frame it as the following classification problem: given a set of time-series TEC data, can we tell if it is preseismic, or within 24 hours of an earthquake event documented in the ISC-GEM catalog?

A classical machine learning approach to the preseismic classification problem using interpretable methods allowed for an explainable counterpart to the deep learning approach, as is summarized in the figure below.

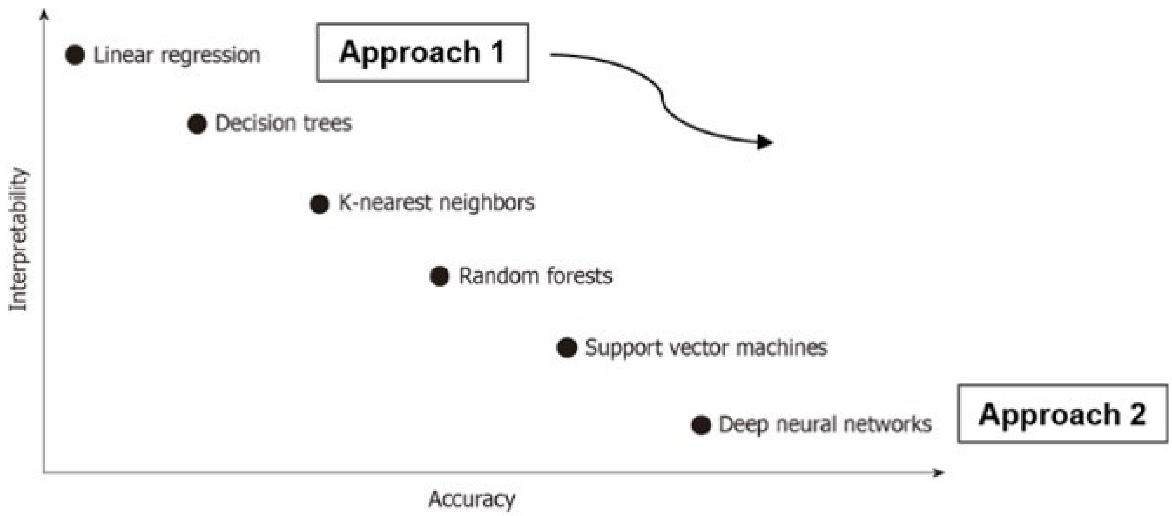

Figure 7. Our investigation of earthquake early warning is dual-prong; using a classical machine learning model that aims at interpretability, along with a deep-learning approach which aims at high accuracy.

Logistic regression is a simple binary classification method which fits a logistic function to the feature set. The coefficients of the best fit curve allows us to understand the relative contribution of features to the model, making the results easily interpretable. This methodology was used for the classical machine learning approach.

Methods and Results

The following feature reduction strategy was used in this initial approach: first, the TEC data was binned radially within each timestep of the 30-degree by 30-degree pre-event window, as can be seen in the figure below. The maximum value of each bin at each timestep was used as input to the model, for a total of 192 features, rather than 8736 features.

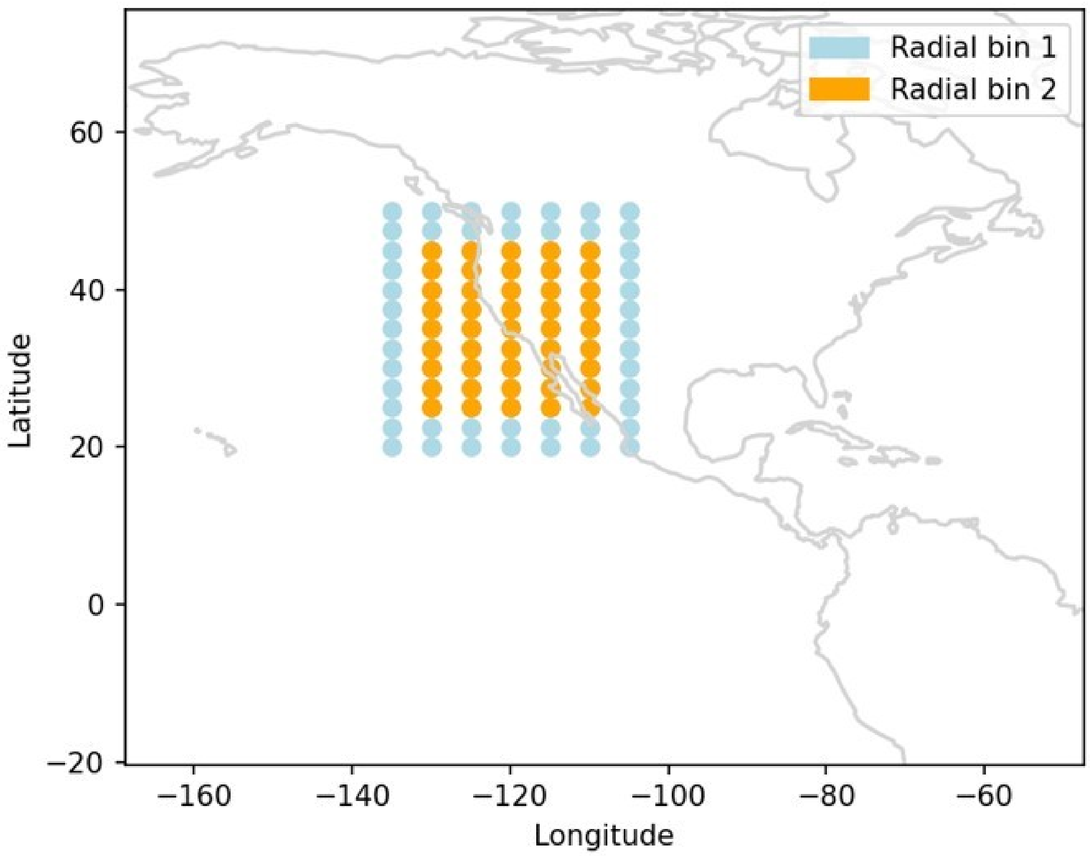

Figure 8. At each of the 96 timesteps within the 24-hour sample, the maximum of two pseudo-radial bins was computed. The resulting 192 values were the feature set for the logistic regression.

Using a methodology called bootstrapping, 500 logistic regression models trained over class-balanced subsets of data from 2005-2014 were generated. They were each evaluated over the same set as well as an unseen class-balanced subset from 2015. The precision and recall of each model as a function of the threshold used to determine class labels is shown below, as well as the average precision and recall of the models over 2005-2014 and 2015 data. For comparison, the results of a baseline classifier, which generates a random number to determine class label, is also shown.

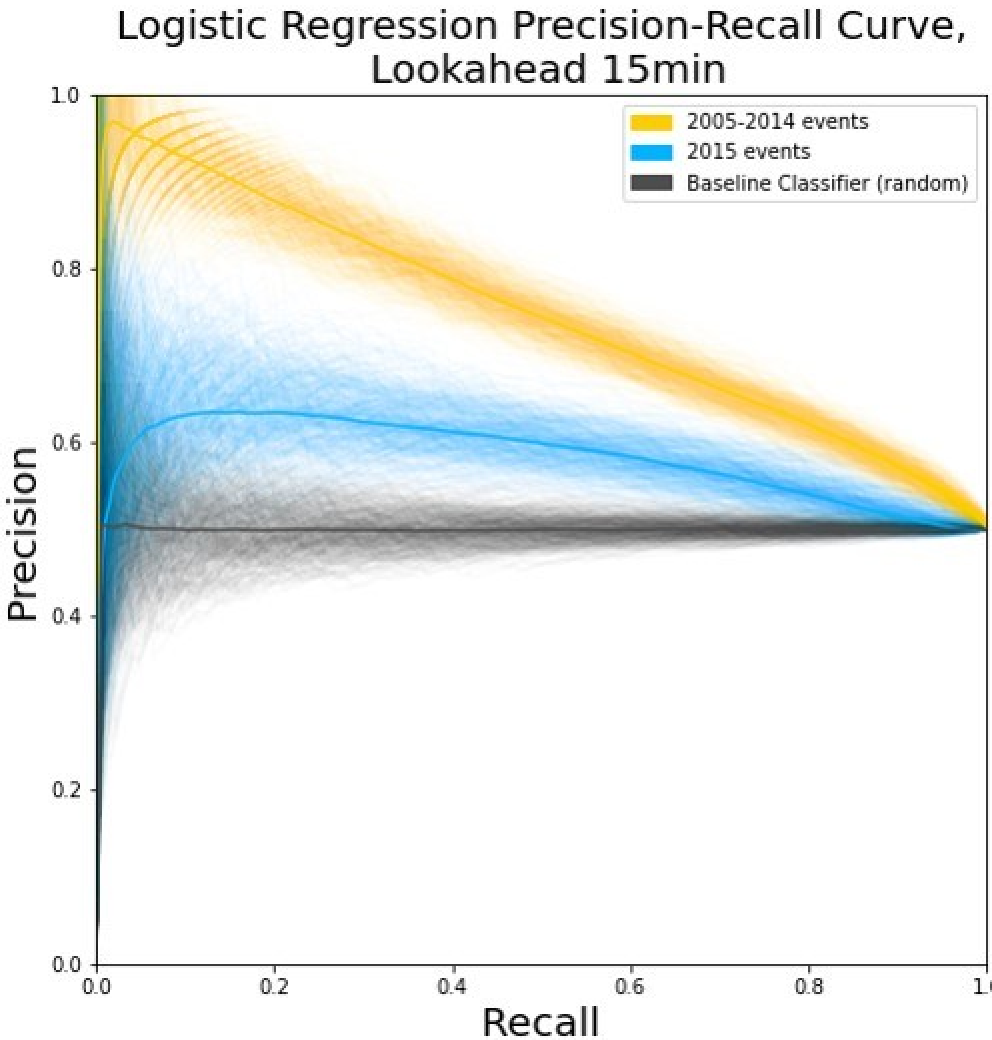

Figure 9. Precision-Recall curve for a log regression classifier of preseismic data, using bootstrapping (sample size = 1000, number of bootstraps = 500, sampled data balanced and with replacement). The feature set was the max TEC value with radial bins at each timestep over a 24-hour period and within a 30-degree by 30-degree spatial window (192 features).

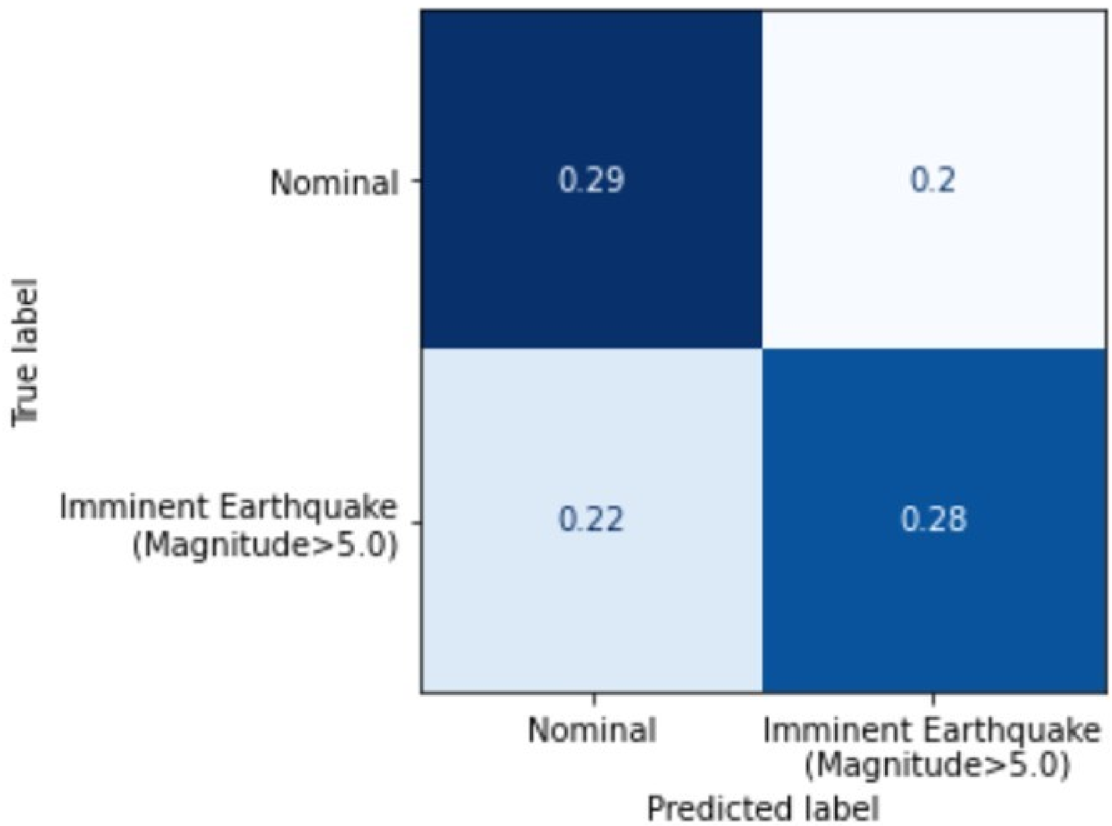

Figure 10. Confusion matrix for one representative classifier of preseismic data from the bootstrapping study using a threshold of 0.5, trained on 500 nominal and 500 imminent samples and applied to a test set of 500 nominal and 500 imminent samples. This plot shows the percent of all data that were (reading across, then down) true negatives, false positives, false negatives, and true positves when the model was applied to unseen test data. Accuracy = 0.578.

The results of the simple model provide compelling evidence for information content in TEC data for earthquake early warning. Analysis of the results and why the feature reduction strategy was successful in retaining information relevant to preseismic classification would provide paths forward for improving classical machine learning models for preseismic classification in the future.

DEEP MACHINE LEARNING METHODS AND RESULTS

Background

Our Deep Learning approach is a Convolutional Long-Short-Term-Memory (ConvLSTM) Autoencoder anomaly detector. ConvLSTM autoencoders are neural networks that learn to reproduce sequences of arrays, such as video data. They can learn to approximate system dynamics and predict the next frame of data based on the preceding frames. Data which looks different from the data that the ConvLSTM Autoencoder was trained on will have higher reconstruction errors than data similar to the training data. This use of the ConvLSTM Autoencoder as an anomaly detector was proposed by Luo et al. (2017), and builds upon the development of the ConvLSTM by Shi et al. (2015) and the use of the autoencoder framework for anomaly detection in video by Xu et al. (2015).

Methods and Results

We trained a ConvLSTM Autoencoder on 100k samples of nominal data from our training set, which we call the Nominal ConvLSTM. By training only on nominal data, we can use the Nominal ConvLSTM as an anomaly detector. The theory behind this is that the Nominal ConvLSTM should be better at predicting frames of data which look similar its training data. Thus, if the pre-earthquake sequences look different from the nominal sequences, then pre-earthquake sequences will have a larger reconstruction error – measured as the Mean Absolute Error (MAE) between observed TEC sequence and that predicted by the Nominal ConvLSTM – than nominal sequences do.

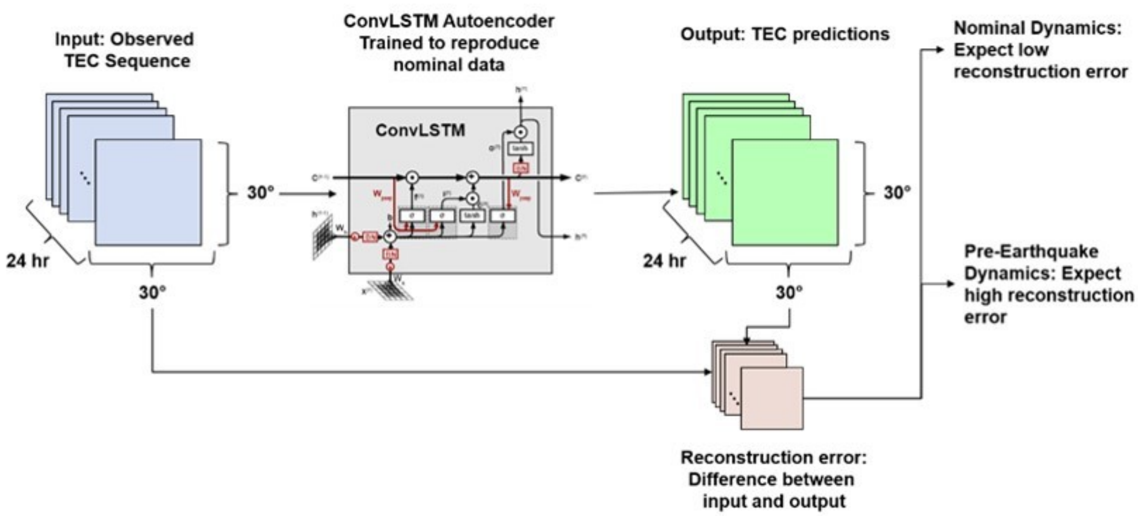

Figure 11: Deep learning approach using a ConvLSTM Autoencoder trained on nominal data to detect preseismic earthquake signals by looking for sequences with high reconstruction errors.

We test if the pre-earthquake sequences are different from the nominal sequences by posing it as a hypothesis test:

H0: The distributions of reconstruction errors for nominal sequences and pre-earthquake sequences from the Nominal ConvLSTM are the same

H1: The reconstruction errors from the Nominal ConvLSTM for pre-earthquake sequences are larger than those for nominal sequences

Under the null hypothesis, the Nominal ConvLSTM is equally accurate in reproducing both nominal and pre-earthquake TEC sequences. This would suggest that pre-earthquake sequences look sufficiently similar to nominal sequences that a ConvLSTM trained only on nominal data can predict the dynamics of pre-earthquake data. The alternative hypothesis states that the Nominal ConvLSTM has larger errors when trying to predict pre-earthquake dynamics. This would suggest that pre-earthquake sequences indeed look different from nominal sequences.

Note that since the Nominal ConvLSTM is trained exclusively on nominal data, we can use it to evaluate any 24-hour TEC sequence and look for high reconstruction errors. Thus, we are able to evaluate sequences with various lookahead values without any additional effort. In this presentation, we evaluate both 15 minute- and 8 hour- lookaheads pre-earthquake sequences, where a 15 minute lookahead would provide 15 minutes of warning and an 8 hour lookahead would provide 8 hours of warning.

Hypothesis Test Results

We test the hypothesis using a one-sided Mann-Whitney U Test, which is a non-parametric test of whether two samples come from the same distribution. We chose a non-parametric test since we do not know the distribution of the reconstruction errors, and thus do not want to assume normality. We use the one-sided version because we are testing specifically if the reconstruction errors for pre-earthquake sequences are greater than those for nominal sequences.

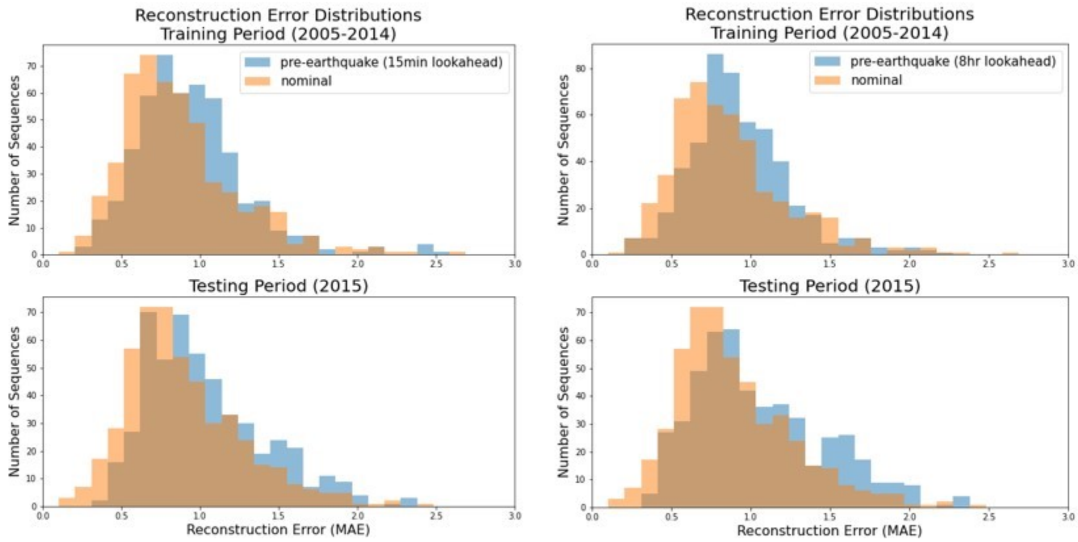

Figure 12: Results of a one-sided Mann-Whitney U test which show that the pre-earthquake data is different than nominal data for both a 15 minute lookahead (left) and 8 hour lookahead (right).

We generated 1000 bootstrap samples, each with 500 nominal and 500 pre-earthquake sequences. For each sample, we computed the reconstruction error for every nominal and pre-earthquake sequence, and computed the Mann-Whitney U Test on the reconstruction errors for the two populations. Figure 12 shows an example of the distributions of reconstruction error for each of the 15 min- and 8 hr-lookahead pre-earthquake and nominal sequence populations for both the Training and Testing data splits. We can see that the distribution of reconstruction errors for the pre-earthquake sequences are skewed rightward compared to that of the nominal sequences. We report the median bootstrap p-value in the table below. In both the Train and Test splits, we are able to reject the null hypothesis for all cases at the 99.99% confidence level. This means that the reconstruction errors for the 15 min and 8 hr-lookahead pre-earthquake sequences are indeed larger than those for the nominal sequences, and that the pre-earthquake sequences look different from the nominal sequences.

Split	Lookahead	p-value	Significance level
Train (2005-2014)	15 minute	1.8e-5	p<0.0001
	8 hour	4.5e-5	p<0.0001
Test (2015)	15 minute	8.9e-10	p<0.0001
	8 hour	8.7e-11	p<0.0001

Table 1: Both the train and test splits are able to reject the null hypothesis at the 99.99% confidence level, confirming the pre-earthquake data is different than the nominal data. These results hold for both 15-minute and 8-hour lookaheads

Precision-Recall Curve

In addition to the hypothesis test, we present the precision recall (PR) curve for the anomaly detector below. The PR curve demonstrates the tradeoff between precision (fraction of alarms which are correct) and recall (fraction of events which are detected). The baseline null classifier (randomly guessing) is shown in grey as a comparison. Confidence intervals are generated by the same bootstrap resampling process described in the Hypothesis Testing section. We see that the anomaly detector outperforms the null classifier across the board. However, the precision is still relatively low, and additional work is necessary to bring the performance up to operationally-useful standards. Nevertheless, these results give us confidence that there are indeed some signals in ionosphere TEC data that can be used to predict some earthquakes before they happen.

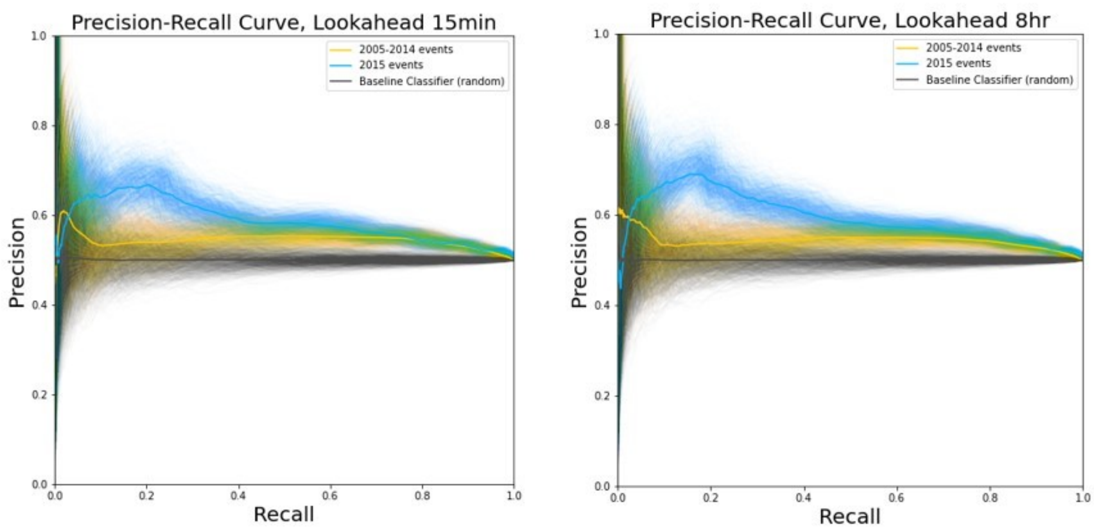

Figure 13: The precision-recall curves for both the 15 minute and 8 hour lookahead anomaly detectors outperform the null classifier, more work is needed to improve the performance for operational use.

In summary, we trained a ConvLSTM Autoencoder on nominal TEC sequences to create a model of nominal TEC dynamics. This was used as an anomaly detector by comparing the reconstruction error of TEC sequences to the predicted dynamics from the Nominal ConvLSTM. We found that the reconstruction error of TEC sequences prior to an earthquake event were significantly higher than the reconstruction error for nominal sequences, suggesting that pre-earthquake sequences indeed look different than nominal sequences. We also show that these differences are apparent at least 8 hours prior to earthquake events.

PATH FORWARD

Our next steps are focused on improving the predictive capabilities of our models. We discuss a few directions below.

Incorporation of additional years of TEC data: This study evaluated models with one year of unseen TEC data, 2015. Because of known year-to-year changes in ionosphere dynamics, it is prudent that TEC data from additional years be used for evaluation as well.

Reducing noise from known phenomena: Our current models consider only the ionosphere TEC data for the sake of simplicity. However, we know that various non-earthquake processes can affect the TEC data, such as solar weather, diurnal and seasonal effects, and electrical storms. Thus, we intend to incorporate additional data of such phenomena into the model so that we can directly account for such effects.

Model refinement: In addition to incorporating additional data sources, we will explore additional classifier types, data representations and dimensionality reduction, and model architectures. We aim to improve accuracy, as well as predict additional characteristics of the earthquake events, such as onset time, magnitude, and location.

Expand model domain and resolution: Our current models consider windows of 30°×30° lat/lon, for 24 hour sequences; we plan to expand the spatial window to see if considering regional and global effects improve accuracy beyond the local data. In addition, we also intend to incorporate higher-resolution data, where available, to try and improve accuracy. Obviously, this will increase the dimensionality of the data, and thus, this will have to be explored in parallel with the work on dimensionality reduction above.

Characterization of signal and physical models: The deductive, statistical nature of our approach means that while we are quite confident that a signal exists in the ionosphere TEC data which can help predict earthquakes, we do not know about the physical characteristics of the signal. A characterization of the physical characteristics and processes related to this signal would provide valuable insight for improving an earthquake prediction system.

SUMMARY

In this presentation, we presented a pair of prototype machine learning classifiers to explore the possibility of short-term predictions of impending earthquakes using ionosphere TEC data. We compiled a decade-spanning dataset of global ionosphere TEC data, and earthquake events. We demonstrated that both a classical logistic regression model and a deep learning ConvLSTM Autoencoder Anomaly Detector trained on this dataset were able to predict upcoming earthquakes significantly better than chance. These results strongly suggest that signals in the ionosphere TEC data contain information that can be used to predict earthquakes earlier than current warnings provided by existing systems.

While these initial results are promising, significant work is still required to bring this approach to operational capacity. We will address several of these challenges in our future work, but also invite others to lend their expertise and examine this problem. To facilitate this, we intend to release data, code and reference models open source to the community.

DISCLOSURES

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

This material is based upon work supported by the United States Air Force under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force.

Copyright 2020 Massachusetts Institute of Technology.

Delivered to the U.S. Government with Unlimited Rights, as defined in DFARS Part 252.227-7013 or 7014 (Feb 2014). Not withstanding any copyright notice, U.S. Government rights in this work are defined by DFARS 252.227-7013 or DFARS 252.227-7014 as detailed above. Use of this work other than as specifically authorized by the U.S. Government may violate any copyrights that exist in this work.

AUTHOR INFORMATION

Jessica Reid is a Technical Staff member at MIT Lincoln Laboratory in the Advanced Sensor Systems and Test Beds group. She specializes in finding the intersections of new technology to solve next generation challenges. She applies her background in thermal engineering and design of complex systems to developing new technology in air, missile, and maritime defense. She received her BSE and MSE in Aerospace Engineering from the University of Michigan and also holds a MBA in Technology and Space Policy from the University of Colorado. Email: jreid@ll.mit.edu

Bhavani Ananthabhotla is an Assistant Staff member at MIT Lincoln Laboratory in the Humanitarian Assistance and Disaster Relief group, where her research interests include the application of machine learning and data analysis to the domain. She received a B.S. in Computer Science from Yale University. Email: bhavani.ananthabhotla@ll.mit.edu

Dr. Jeffrey Liu is a Technical Staff member at MIT Lincoln Laboratory in the Humanitarian Assistance and Disaster Relief group. He works on applications of artificial intelligence and machine learning for disaster prediction, characterization, and response. In particular, he is interested in anomaly detection and computer vision applications. He received his PhD from MIT in Civil Engineering and Computation, and also holds a BSE in Engineering Physics from the University of Michigan. Email: jeffrey.liu@ll.mi.edu

ABSTRACT

QuakeCast is a novel system for short-term earthquake prediction using global ionosphere Total Electron Content (TEC) data. In the last 20 years, earthquakes have caused over 800,000 deaths and $650 billion in economic damage. While current seismic early warning systems provide up to a minute's notice by detecting an earthquake's non-damaging P-waves prior to the damaging S-waves, earlier warnings could allow more significant emergency preparations.

Electromagnetic ionospheric phenomena have recently been observed many days before major earthquakes, notably in 2015 in Nepal. QuakeCast explores whether such signals could be used to predict earthquakes in a global dataset of ionosphere and earthquake data. The ionosphere data consists of global TEC data from NASA's Crustal Dynamics Data Information System (CDDIS), for the years 2005-2015 at 15 minute intervals. The earthquake data consists of over 10,000 events of at least M5 from the International Seismological Centre (ISC-GEM) Global Instrumental Earthquake Catalogue over the same time period.

We used the dataset to explore the following classification problem: given a 24-hour sequence of ionosphere TEC data in a 30° latitude by 30° longitude window, is the sequence preseismic (an earthquake occurs in the next time step) or nominal (no earthquake)? We built two models to address this. The first is a classical logistic regression model trained on radially binned data. Analysis of the classifier weights indicate a localized effect in the ionosphere: TEC data spatially and temporally closer to the event contributes more significantly to the prediction. The second model is a deep learning ConvLSTM autoencoder trained to reproduce nominal TEC data sequences. Since autoencoders reproduce data which resemble their training data better than data which does not, we used the reconstruction error to classify sequences as anomalous (preseismic) or nominal.

Both methods were found to perform significantly better than a random null classifier, indicating that the ionosphere contains information useful for predicting earthquakes at least 15 minutes in advance. While further research is needed to incorporate additional data features and address noise, we believe that these results are a promising development toward forecasting earthquakes using geoelectric signals.

REFERENCES

Bondár, I., et al. “ISC-GEM: Global Instrumental Earthquake Catalogue (1900–2009), II. Location and Seismicity Patterns.” Physics of the Earth and Planetary Interiors, vol. 239, Feb. 2015, pp. 2–13, doi:10.1016/j.pepi.2014.06.002.

Di Giacomo, Domenico, James Harris, et al. “ISC-GEM: Global Instrumental Earthquake Catalogue (1900–2009), I. Data Collection from Early Instrumental Seismological Bulletins.” Physics of the Earth and Planetary Interiors, vol. 239, Feb. 2015, pp. 14–24, doi:10.1016/j.pepi.2014.06.003.

Di Giacomo, Domenico, István Bondár, et al. “ISC-GEM: Global Instrumental Earthquake Catalogue (1900–2009), III. Re-Computed MS and Mb, Proxy MW, Final Magnitude Composition and Completeness Assessment.” Physics of the Earth and Planetary Interiors, vol. 239, Feb. 2015, pp. 33–47, doi:10.1016/j.pepi.2014.06.005.

Di Giacomo, Domenico, E. Robert Engdahl, et al. “The ISC-GEM Earthquake Catalogue (1904–2014): Status after the Extension Project.” Earth System Science Data, vol. 10, no. 4, Oct. 2018, pp. 1877–99, doi:https://doi.org/10.5194/essd-10-1877-2018.

Freund, et al. "Cracking the Code of Pre-Earthquake Low Frequency EM Emissions." 2005

Kuo, Cheng-Ling, et al. “Electrical Coupling Between the Ionosphere and Surface Charges in the Earthquake Fault Zone.” Geophysical Monograph Series, edited by Dimitar Ouzounov et al., John Wiley & Sons, Inc., 2018, pp. 99–124, doi:10.1002/9781119156949.ch7.

Liu, Cheng-Yan, et al. “Statistical Analyses on the Ionospheric Total Electron Content Related to M ≥ 6.0 Earthquakes in China during 1998 - 2015.” Terrestrial, Atmospheric and Oceanic Sciences, vol. 29, no. 5, 2018, pp. 485–98, doi:10.3319/TAO.2018.03.11.01.

Molchanov, O., et al. “Lithosphere-Atmosphere-Ionosphere Coupling as Governing Mechanism for Preseismic Short-Term Events in Atmosphere and Ionosphere.” Natural Hazards and Earth System Sciences, vol. 4, no. 5/6, Nov. 2004, pp. 757–67, doi:10.5194/nhess-4-757-2004.

Noll, Carey E. “The Crustal Dynamics Data Information System: A Resource to Support Scientific Analysis Using Space Geodesy.” Advances in Space Research, vol. 45, no. 12, June 2010, pp. 1421–40, doi:10.1016/j.asr.2010.01.018.

Storchak, D. A., et al. “The ISC-GEM Global Instrumental Earthquake Catalogue (1900–2009): Introduction.” Physics of the Earth and Planetary Interiors, vol. 239, Feb. 2015, pp. 48–63, doi:10.1016/j.pepi.2014.06.009.

Storchak, Dmitry A., et al. “Public Release of the ISC–GEM Global Instrumental Earthquake Catalogue (1900–2009).” Seismological Research Letters, vol. 84, no. 5, Sept. 2013, pp. 810–15, doi:10.1785/0220130034.

Zhou, Yiyan, et al. “Ionospheric Disturbances Associated with the 2015 M7.8 Nepal Earthquake.” Geodesy and Geodynamics, vol. 8, no. 4, 2017, pp. 221–28, doi:10.1016/j.geog.2017.04.004.

Zolotov, et al. "TEC Disturbances Prior to the Great Tohoku March 11, 2011 and October 23, 2011 Turkey Van Earthquakes and Their Physical Interpretation." Proceedings of the MSTU, Vol. 15, No. 3, 2012