PREDICTING EARTHQUAKES USING LONG SHORT-TERM MEMORY NETWORKS

A Research Paper submitted to the Department of Computer Science
In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science in Computer Science

By

Keerat Singh

April 28, 2023

On my honor as a University student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments.

ADVISOR
Geoffrey E. Fox, Department of Computer Science

PREDICTING EARTHQUAKES USING LONG SHORT-TERM MEMORY NETWORKS

CS4991 Capstone Report, 2023

Keerat Singh
Computer Science
The University of Virginia
School of Engineering and Applied Science
Charlottesville, Virginia USA
ks5hrx@virginia.edu

ABSTRACT

Past attempts to predict earthquakes have been largely unsuccessful. Professor Geoffrey Fox and I, on behalf of the University of Virginia Biocomplexity Institute, tried feeding earthquake data into machine learning models to predict events. We collected time-series data on earthquake events from the United States Geological Survey (USGS) for select regions in the United States; and we built a Long Short-Term Memory (LSTM) model to predict earthquakes months in advance. Our results show that earthquakes appear to be predictable in the long term, although more so for some locations than others. Inconsistency in the results of this study suggest more work is needed before application into an Earthquake Early Warning system.

1. INTRODUCTION

Predicting earthquakes is a significantly challenging problem with a long history of largely unsuccessful attempts. As defined by USGS, an acceptable earthquake prediction must include a time, location and magnitude (USGS, n.d.a). There is a distinction between earthquake prediction and earthquake forecasts. Earthquake predictions are specifically predictions months or years in advance of the event, while forecasting is on the scale of hours. There has never been an accurate prediction of any major earthquake (USGS, n.d.b). The technical research is focused on predicting earthquakes by feeding large time-series data sets of earthquake events into machine learning (ML) models.

It would be very useful to know when and where large earthquakes will occur. Many lives could be saved. Even for communities that cannot access transportation to geographically relocate during an event, moving to open areas or holding onto shelter and covering one's head is fairly effective in preventing casualties (USGS, n.d.c). Furthermore, mitigation efforts would be able to take place far in the future, at least as far as the prediction window of the model. Resources such as buildings, food, and livestock would also be able to be moved and saved. Pre-designated dangerous areas, such as areas with buildings that are not earthquake-resistant, could be closed off and evacuated in advance. Simply put, given an accurate prediction of a large event, the public would have more time to avoid potential harm.

2. RELATED WORKS

While there are known physical equations that govern when and where an earthquake will happen, humans are unable to obtain the parameters needed to perform such a calculation. One traditional field of attempts for predicting earthquakes has been monitoring for suspected local precursors to large events (Hayakawa, 2016; Korepanov, 2016). This is challenging because of the high complexity in correlation and the relatively short prediction range this allows. Therefore, at best it could only be used for earthquake forecasting.

Another avenue has been mathematically analyzing long-term trends in geophysical-related patterns (Boucouvalas et. al, 2015; Kannan, 2014). No consistent pattern has emerged as it relates to earthquakes. This technical research is on predicting earthquakes using ML models rather than physical models. ML methods are possible because of the large amount of historical earthquake data available. With this data and high computational power, computers can be trained to fit a certain dataset. Furthermore, the built-in ability of time-series-based methods to correlate complex spatio-temporal state spaces makes them promising for this problem. The goal is to accurately predict the time, location and magnitude of major earthquakes months or years into the future.

3. PROJECT DESIGN

USGS has developed the most extensive, active database on earthquakes events worldwide since around 1950 (USGS, n.d.d) This public data includes events characterized by their magnitude, depth, date, and latitude/longitude coordinates. An example of data from Southern California is shown in Figure 1.

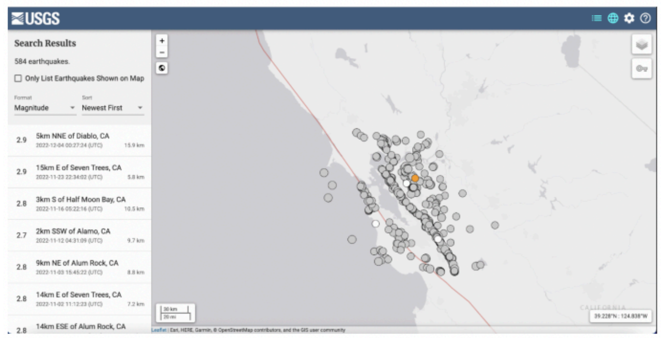

Fig. 1. Earthquake data in Southern California from USGS

The data was transformed into a space-time matrix by binning the events based on bin size. For example, an event at (34.9, 116.4) and (34.84, 116.2) would be processed as events at the same location. The time unit was chosen to be one day. Therefore, if the two events from the previous example occurred on different days, they would be stored with the same location index but a different time index.

This was used to train a single LSTM model. The ML model is made for problems that deal with predicting certain values over time, which is exactly the framework of the earthquake prediction model. The model was trained on 80% of the locations and tested on the rest. If the model had substantially learned some general pattern, then its predictions on the locations it had never seen before should resemble the actual earthquake activity that happened there. Other parameters were tweaked in an effort to increase the similarity between the prediction and the actual output.

4. RESULTS

Initially, data was gathered from all areas with significant earthquake activity, including Southern California, Japan, Mexico and others. After further inspection, it was found that only data from the United States contained a large enough record of small events with a magnitude less than 0.5. It proved to be challenging to accurately predict future events in regions that did not record small events, which suggested that these small events act as necessary predictors for big events. Therefore, results were collected for regions in the United States such as Southern and Northern California, Hawaii, Washington, and Alaska.

Figure 2 shows the results for a model asked to predict 6 months in the future on earthquakes from Hawaii. Within each graph, the actual cumulative magnitude within the entire location's space is plotted over time in blue. At every point in time, the model is given a history of actual events up to that point, and is asked to predict the magnitudes across the location space six months into the future. Therefore, the predicted value on January 1, 2000 in orange was made by the model with only an understanding of what had happened up to June 1, 1999.

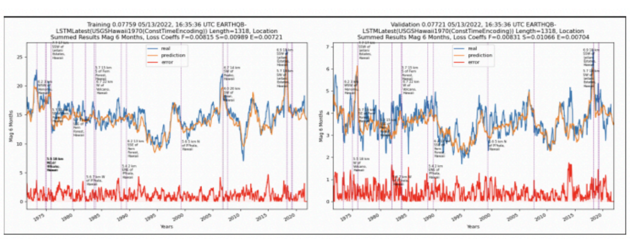

Fig. 2. 6-month prediction on Hawaii

On both graphs, the orange line shows these predictions. The left graph contains locations included in the training set from which the model learned a pattern in input parameters and output prediction. The right graph serves as the test, as it shows the model's predictions on locations within Hawaii it had never seen before. The results for Hawaii are promising. The 1976 7.7 M earthquake and 2018 6.9 M earthquake, both near Leilani Estates, Hawaii, were accurately predicted with notable spikes in the test graph.

The model is not always accurate, however. As shown in Figure 3, for a model trained on data from Northern California, a major spike appears unnecessarily in 1980 and does not appear for the 6.9 M Loma Prieta earthquake or the 6.1 Carter Springs earthquake. Therefore, there is a potential for false positives and negatives.

Fig. 3. 6-month prediction on Northern California

While false positives and negatives are inherent in any prediction model, they especially need to be minimized for an earthquake prediction model. We do not want to waste resources in a long-term preparation effort for an earthquake that does not happen.

Upon closer investigation into Figure 3, one can see that the reason for the predicted event in 1980 most likely was because the training data contained the 6.1 M earthquake in Aspen Springs in 1980. It was suspected that the model was overfitting to its training data. An overfitted model is not useful, since it will not help us predict earthquakes it has not seen before.

At this point, one could argue that the results on Hawaii from Figure 2 could have also been due to overfitting. This is especially likely since the major earthquakes that affected locations in the test set, such as the 1976 7.7 M earthquake and the 2018 6.9 M earthquake near Leilani Estates, also affected locations in the training set. To test whether the model was actually learning a generalizable pattern of when earthquakes would happen in the long term, we used a model trained on a given location to predict earthquakes in a completely different location. This effectively removed the potential for overlap in the training and test sets that might have prevailed in the results from Figures 2 and 3. As shown in Figure 4, the model trained on Northern California's data was asked to predict earthquake activity in Southern California six months in advance.

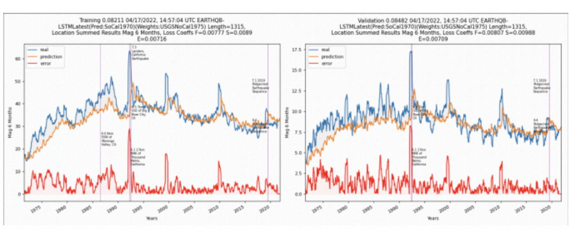

Fig. 4. 6-month prediction on Southern California using the Northern California model

In this case, both the left and right graphs are test data, since the model never saw any location from Southern California during its training. While there is a discrepancy in the absolute values of the spikes, the relative values appear to line up. The highest prediction peak in 1993 correctly corresponds with the three 6+ M events that occurred that year. Similarly, other local spikes also match, such as the 2019 7.1 M earthquake in Ridgecrest. This is especially promising, as it suggests that earthquakes have predictable properties in the long term, and LSTM models have the ability to learn and encode them.

The performance of the model is largely location-specific and not yet consistent. For example, Figure 5 shows a model trained on data from Washington predicting earthquakes in Northern California. The predictions are largely sporadic, although the actual data also appears sporadic. Nevertheless, actionable insight could not be gained from such a prediction, showing the need for greater consistency.

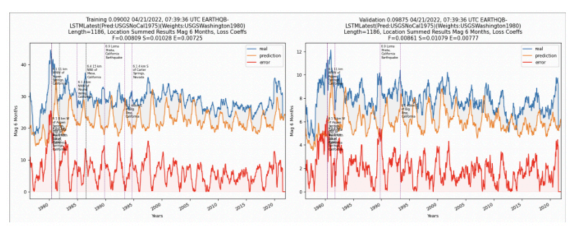

Fig. 5. 6-month prediction on Northern California using the Washington model

5. CONCLUSION

The problem of earthquake prediction appears to be solvable. One potential solution explored in this study combined large time-series earthquake datasets with LSTM models. While the results were largely dependent on location, there were cases where such a model was able to predict significant earthquakes six months in advance.

6. FUTURE WORK

Future work should aim to reduce the number of false positives and negatives of the model, as that will be key if it is ever to be deployed into an Earthquake Early Warning system. Additionally, as with most ML problems, expanding data collection efforts to ensure large, complete datasets are available is extremely helpful. These efforts are most necessary outside of the United States, as their current datasets are not thorough enough to be used with this solution.

7. ACKNOWLEDGMENTS

This work was started and led by Geoffrey C. Fox of UVA's Biocomplexity Institute. The entire study was outlined in a paper written by Geoffrey C. Fox, John B. Rundle, Andrea Donnellan, and Bo Feng (Fox et al., 2022).

REFERENCES

Boucouvalas, A. C., Gkasios, M., Tselikas, N. T., & Drakatos, G. (2015). Modified-fibonacci-dual-lucas method for earthquake prediction. *SPIE Proceedings*, 95–96. https://doi.org/10.1117/12.2192683

Fox, G. C., Rundle, J. B., Donnellan, A., & Feng, B. (2022). Earthquake nowcasting with deep learning. *GeoHazards*, *3*(2), 199–226. https://doi.org/10.3390/geohazards3020011

Hayakawa, M. (2016). Earthquake prediction with electromagnetic phenomena. *AIP Conference Proceedings*. https://doi.org/10.1063/1.4941201

Kannan, S. (2014). Innovative mathematical model for earthquake prediction. *Engineering Failure Analysis*, *41*, 89–95. https://doi.org/10.1016/j.engfailanal.2013.10.016

Korepanov, V. (2016). Possibility to detect earthquake precursors using cubesats. *Acta Astronautica*, *128*, 203–209. https://doi.org/10.1016/j.actaastro.2016.07.031

Psarosa, A. F., Menga, X., Zoua, Z., Guob L., & Karniadakis G. E. (2022). Uncertainty quantification in scientific machine learning: methods, metrics, comparisons. *Cornell Unversity*, *10*(2), 1-104. https://doi.org/10.48550/arXiv.2201.0776

Reksten J., Salberg A., & Solberg R. (2019). Flood detection in Norway based on Sentinel SAR-1 imagery. The International Archives of the Photogrammetry, *Remote Sensing and Spatial Information Sciences*, *42*(3), 349-355. https://doi.org/10.5194/isprs-archives-XLII-3-W8-349-2019

The United States Geological Survey. (n.d.a). *What is the difference between earthquake early warning, earthquake forecasts, earthquake probabilities, and earthquake prediction?* Retrieved October 26, 2022, from https://www.usgs.gov/faqs/what-difference-between-earthquake-earlywarning-earthquake-forecasts-earthquake-probabilities

The United States Geological Survey. (n.d.b). *Can you predict earthquakes?* Retrieved October 26, 2022, from https://www.usgs.gov/faqs/can-you-predictearthquakes

The United States Geological Survey. (n.d.c). *What should I do during an earthquake?*. Retrieved October 26, 2022, from https://www.usgs.gov/faqs/what-should-i-doduring-earthquake

The United States Geological Survey. (n.d.d). *Lists, Maps, Statistics*. Retrieved October 26, 2022, from https://www.usgs.gov/programs/earthquake-hazards/listsmaps-and-statistics