PREDICTING EARTHQUAKES USING LONG SHORT-TERM MEMORY NETWORKS

ANALYZING THE NECESSARY CONFIGURATIONS OF A LONG-TERM EARTHQUAKE EARLY WARNING SYSTEM

A Thesis Prospectus
In STS 4500
Presented to
The Faculty of the
School of Engineering and Applied Science
University of Virginia
In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science in Computer Science

By
Keerat Singh

October 27, 2022

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.

ADVISORS

Catherine D. Baritaud, Department of Engineering and Society

Geoffrey C. Fox, UVA Biocomplexity Institute

Predicting earthquakes is a significantly challenging problem with a long history of largely unsuccessful attempts. As defined by the United States Geological Survey (USGS), an acceptable earthquake prediction must include its date and time, location, and magnitude (*What is*, n.d.). In addition, earthquake predictions deal with larger-scale time windows on the order of months and years. There has never been an accurate prediction of any major earthquake (*Can*, n.d.). The technical research is focused on predicting earthquakes by feeding large historical datasets of events into machine learning (ML) models.

The ability to predict naturally leads to the deployment of earthquake early warning (EEW) systems in given regions. Earthquake forecasting models, predicting the characteristics of an earthquake hours or days in advance, are becoming more successful. As a result, there are many short-term EEWs, built on these forecasting models, dispersed in high-impact regions of the world. While potentially a powerful tool in any field, early warning (EW) systems must be carefully configured due to the uncertainty in any prediction. Especially if longer-range prediction from the technical project proves to be successful, EEWs built on prediction models need to be standardized before implemented. As the overall power of EEWs will increase, the chance for mispredictions increases with it. The STS research is comparing and contrasting EW systems from many different fields to see how such tradeoffs in predictive power and uncertainty are generally made. This is tightly coupled with the technical research, as this knowledge ultimately assists in imagining how a larger-window EEW would be best configured and deployed.

The technical work was started at the beginning of 2022 with Geoffrey C. Fox. Certain models can accurately predict earthquakes in multiple regions throughout the US around six months in advance. Present work is related to statistically analyzing the reliability of these results, particularly as the model makes predictions further out into the future. Once these results are gathered, a paper will be written. It will hopefully communicate the success of using LSTM networks for earthquake prediction in certain regions, as well as improvements that need to be made before deployment to the public. The paper will most likely be completed by the end of 2022. For the STS research, the main sources for comparison have been collected. Currently, efforts are being made to read, analyze, and compare these papers to develop a knowledge base for the STS question. This work will be carried out through 2022 and into 2023. A comprehensive argument answering the question will hopefully be able to be made at the end of January 2023.

PREDICTING EARTHQUAKES USING LONG SHORT-TERM MEMORY NETWORKS

It would be very useful to know when and where large earthquakes will occur. Countless lives would be able to 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 (*What should*, n.d.). 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 the downtown of cities, could be closed off during the predicted time of the event. Simply put, given an accurate prediction of a large event, the public would have more time to prepare which directly relates to less chance for harm.

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 into the future 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 machine learning models rather than physical models. Machine learning 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 date, location, and magnitude of major earthquakes months or years into the future.

USGS has developed the most extensive, up-to-date database on earthquakes events worldwide since around 1950 (*Lists*, n.d.). This publicly-available data includes events characterized by their magnitude, depth, date, and latitude/longitude coordinates. All of the computer processing was done using Google Colab, and the subscription fee was paid for by UVA's Biocomplexity Institute. No additional resources were needed.

The data was transformed into a space-time matrix by binning the events based on bin size. A bin size of 0.1, for example, included all events within a 0.1 by 0.1 degree box into the same location. Therefore, an event at (34.9, 116.4) and (34.84, 116.2) would be processed as events at the same location. A final matrix was constructed such that for every time unit, a location space at that time was stored. The time unit was chosen to be one day. Therefore, if the two events from the previous example occurred on the same day, the value in the array on that day in that box of space would be the sum of the two events. Multiple major events in the same region on the same day would rarely happen, and usually spots in the array either held the event that happened there at that time or a 0, signifying no earthquake activity.

This was used to train a single Long-Short-Term Memory (LSTM) model. This ML model is made for problems that deal with predicting certain values over time, which is exactly the framework of the earthquake prediction model. Its architecture is two stacked LSTM layers with single fully-connected networks at the start and end. The architecture was experimented with and it ultimately dictates how the input will be processed to get the desired output. Training the model consisted of giving it one “slice” of the time series data, say from 1950 to 2000, to find patterns from. Once it had “learned” some pattern, it was asked to predict earthquake activity from 2000 to 2022. If the model had substantially learned something, then its predictions should have matched the actual earthquake activity from 2000 to 2022. Other parameters were tweaked in an effort to increase the similarity between the prediction and the actual output.

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 enough small events with magnitude less than 0.5. It was impossible 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 1 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 locations 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. 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 a test, as it shows the model's predictions on locations within Hawaii it had never seen before. The results are promising, as it seems earthquake prediction is feasible with such a model. A scholarly paper will be written on the findings once more statistical analysis is done to ensure confidence.

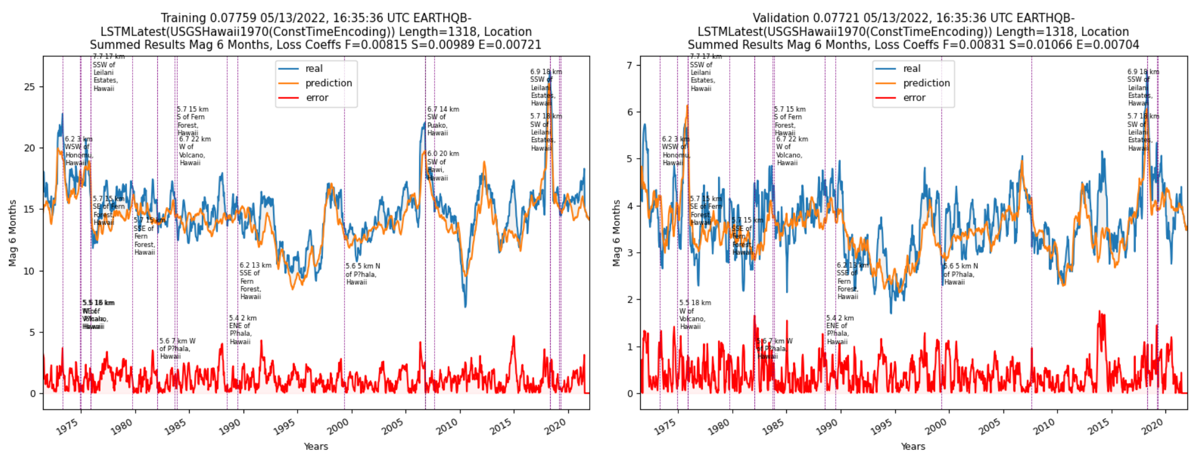

Figure 1: 6-month prediction window LSTM output on Hawaii (Fox & Singh, 2022)

ANALYZING THE NECESSARY CONFIGRUATIONS OF A LONG-TERM EARTHQUAKE EARLY WARNING SYSTEM

For any technology built on a prediction algorithm, there is generally an exponential relationship between how far into the future it is predicting and a resulting uncertainty as shown in Figure 2. This is because of compounding errors, and specifically the compounding error of predicting 50 days in advance that builds on predicting 25 days in advance, for example. Most of the literature on EW systems revolves around systems that are contained in the rectangular black box on the figure. Therefore, to imagine a long-term EEW system that would be far outside the box requires both technological and STS research-backed support.

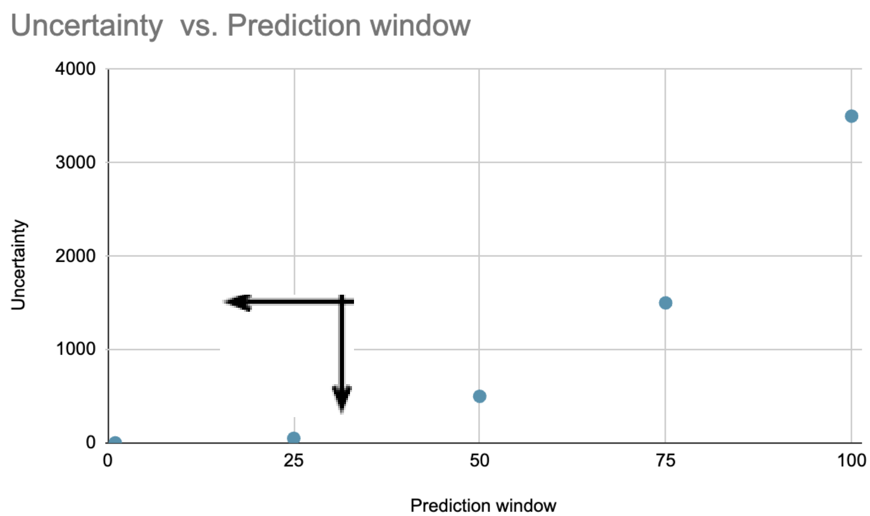

Singh, K (2022). *Relationship between uncertainty and prediction window.* [2]. *Prospectus* (Unpublished undergraduate thesis). School of Engineering and Applied Science, University of Virginia. Charlottesville, VA.

Most current EEW systems are incredibly short-term in their alerts, as they report the presence of an event based on real-time, event-driven data. With an EEW system predicting months and potentially years in advance, the uncertainty naturally rises, and thus an appropriate threshold of confidence before prediction must be established. Establishing this threshold is necessary for any technology that calls for action based on the prediction of a future event. This is a very common problem when preparing for not only earthquakes, but also floods, landslides, and other natural disasters (Reksten et. al, 2019; Guzetti et. al, 2020). For example, floods within Norway have the potential to cause significant damage to towns on the edge of fjords. Reksten, Salberg, and Solberg have researched the development of a technology that uses satellite images to detect the floods right before they happen. Subsequently, alerts are delivered so people can mobilize to a safe altitude above the flood. These authors mention that the challenge for preventing false positives still remains, but efforts such as identifying critical scores between precision and recall and select filtering of the output maps slightly improve the problem.

While removing false positives and false negatives entirely is the ultimate goal, it is largely unrealistic. More importantly, EW systems that are known to potentially be wrong yet are still deployed need to have a set level of trust established, which is a discussion that leads to many ethical questions of how to distribute resources with variable uncertainty. For example, Alam, Hobbelink, and other authors presented an overview of EW systems for patient outcomes in medical settings (Alam et. al, 2014). They analyzed seven studies performed in different hospitals that compare outcomes before and after the introduction of an EW system, providing details on each implementation, the following course of action, and the outcomes. What they found is largely inconclusive in itself, but when combined with other similar analyses from other domains it could be helpful.

The STS question is how should this tradeoff in predictive power and uncertainty be made for a larger window EEW. Insight can be gained by comparing and contrasting other successful and unsuccessful EW systems in different domains. In terms of natural disaster preparedness, there have been multiple successful landslide EW systems (Guzetti et. al, 2020). These fields will not directly map on to each other. For example, the negative effects of a false positive from an EW system for patients in a hospital are much less than those for an earthquake (Alam et. al, 2014). Nevertheless, insight can still be gained to imagine the ethical, yet useful configuration of a long-term EEW system.

TAKING A SOCIAL CONSTRUCTION THEORY OF TECHNOLOGY APPROACH

Allen and Melgar conducted several case studies that look at how EEW systems around the world are implemented, configured, and deployed (Allen & Melgar, 2019). They generally found that, among other things, studies on the public's preferences are used to assist the configuration of the EEW system, an otherwise tricky task. This suggests that the STS research done here will take a Social Construction Theory of Technology (SCOT) view on EW systems. Figure 3 shows the general mapping of SCOT to any EW system. Once a warning is triggered, related scientists must assess and explain specifics of the warning, local officials or decision makers must accordingly outline necessary actions, mitigation teams must carry out those actions, and locals will be directly affected by actions taken or not taken. There may be other supporting groups depending on the specifics of the EW system. Nevertheless, all groups must be involved in the process of recieving a warning to taking action.

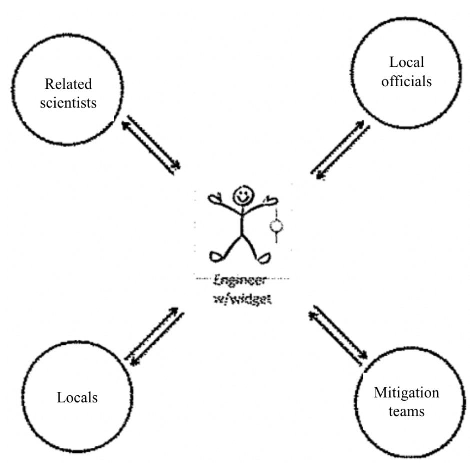

Singh, K (2022). *Mapping of SCOT to problem space.* [3]. *Prospectus* (Unpublished undergraduate thesis). School of Engineering and Applied Science, University of Virginia. Charlottesville, VA.

COMPARING AND CONTRASTING USING THE SCOT APPROACH

Using case studies of other EW systems, the interaction of these groups will be studied to understand how to negotiate configurations for a successful EW system. It may be challenging to extract general patterns when decisions are made on a specific basis within each case study. But, given that the specifics are known through the technical, there is greater potential for useful analysis. In each case, it is necessary to understand who defines each group, how their preferences and needs are assessed, and how those are translated into decisions about the implementation and/or deployment of the technology. Additionally, it will be helpful to observe how conflicts are resolved, as the error-prone nature of EW systems can create disagreement in action that can't easily be settled with confidence. It will be useful to compare which groups tend to take precedence over others within both successful and unsuccessful EW systems. With this understanding of both the makeup of relevant groups and their impact on the EW system, hopefully a well-supported theory can be put forth that answers these same questions related to handling a long-term EEW system. Figure 4 shows the general outline of the study.

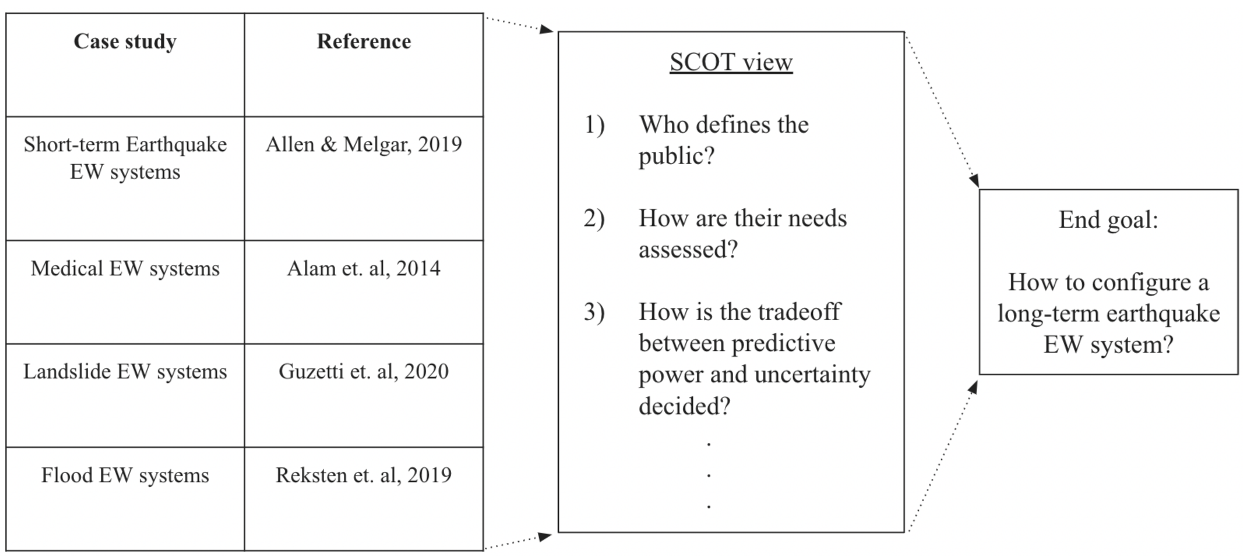

Singh, K (2022). *Outline of the proposed STS study.* [4]. *Prospectus* (Unpublished undergraduate thesis). School of Engineering and Applied Science, University of Virginia. Charlottesville, VA.

Each paper will be used to answer each relevant question as it relates to the specific case study. The knowledge will be combined to hypothesize the answers to the same questions revolving around the negotiations between relevant social groups for a long-term EEW system.

EFFECTIVELY CREATING AND DEPLOYING A LONG-TERM EEW SYSTEM

The technical research aims to create a reliable long-term EEW system for select regions in the United States. With an appropriately comprehensive dataset, it appears that using LSTM networks are promising prediction models anywhere in the world, although more work needs to be done to support this claim. Nevertheless, the advent of long-term EEW systems built on such models is nearing, and necessary questions suggested by SCOT need to be addressed. This will be done by comparing and contrasting how these questions were answered with other EW technologies. Although there is not necessarily a clear mapping from each domain to the other, the similarities and differences can be contextually related to each other to imagine a successful implementation of a long-term EEW system.

REFERENCES

Alam, N., Hobbelink, E. L., van Tienhoven, A. J., van de Ven, P. M., Jansma, E. P., & Nanayakkara, P. W. B. (2014). The impact of the use of the early warning score (EWS) on patient outcomes: A systematic review. *Resuscitation*, *85*(5), 587–594. https://doi.org/10.1016/j.resuscitation.2014.01.013

Allen R., & Melgar D. (2019). Earthquake Early Warning: Advances, Scientific Challenges, and Societal Needs. *Annual review of Earth and Planetary Sciences*, *47*(1), 361-388. https://doi.org/10.1146/annurev-earth-053018-060457

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

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

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

Guzzetti, F., Gariano, S. L., Peruccacci, S., Brunetti, M. T., Marchesini, I., Rossi, M., & Melillo, M. (2020). Geographical landslide early warning systems. *Earth-Science Reviews*, *200*, 102973. https://doi.org/10.1016/j.earscirev.2019.102973

Joffe, H., Rossetto, T., Bradley, C., & O'Connor, C. (2017). Stigma in science: The case of earthquake prediction. *Disasters*, *42*(1), 81–100. https://doi.org/10.1111/disa.12237

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

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

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

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

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