Toward Earthquake Prediction: Combining Physics-Based Models with Machine Learning

Swagatika Pindiyath Menon*, Achint Kumar†

Abstract

Earthquake prediction is difficult because stress in the Earth builds up and releases in complex ways. This work combines a physics-based model with machine learning to study how earthquakes can be predicted. The Olami–Feder–Christensen (OFC) model is used to simulate how stress changes over time on a grid. The model is adjusted so that it follows known earthquake patterns like the Gutenberg–Richter law and Omori's law. The simulated stress fields are then used to train a convolutional neural network (CNN) to predict whether a large event will occur soon. The model achieves a ROC-AUC of 0.67 and a PR-AUC of 0.90, showing good performance in identifying rare large events. Visualization of the stress fields suggests that the model learns useful patterns, especially regions where stress is concentrated. This approach shows that combining physics-based simulations with machine learning can help improve understanding of earthquake behavior and may support better early warning systems in the future.

Introduction

Earthquakes are natural events that happen when energy built up inside the Earth is suddenly released, causing the ground to shake. This usually occurs along cracks in the Earth's surface called faults, where large pieces of the Earth's crust, known as tectonic plates, are constantly moving. Over tine, stress builds up as these plates push against each other, pull apart, or slide past one another. When the stress becomes too strong, the rocks break and release energy in the form of seismic waves, which travel through the ground and cause shaking. Some earthquakes are very small and cannot be felt, while others are strong enough to destroy buildings and change landscapes. Earthquakes can have a major economic impact on society. When a strong earthquake hits, it can damage homes, schools, roads, bridges and more. This damage costs a lot of money to repair, and in many cases, entire communities must be rebuilt. Businesses may shut down, which leads to loss of jobs and income for many people. Governments often need to spend billions of dollars on recovery efforts. In developing countries, the impact is often worse because buildings are not always built to withstand earthquakes. Even in developed countries, large earthquakes can disrupt transportation systems, supply chains, and daily life for months or even years. Because of these serious effects, predicting earthquakes is very important for reducing damage and saving lives. [5]

In this project, the OFC model [1] was simulated using a computer to generate stress field data over time. This data shows how stress is distributed across the grid and how it changes as points fail and transfer stress to one another. By running the simulation many times, a large amount of data can be collected, including when events happen, how large they are, and how often they occur. This data can then be studied to look for patterns that are similar to real earthquakes. Two important laws were used to calibrate the model: the Gutenberg-Richter law and Omori's law. The Gutenberg-Richter law explains the relationship between the size and frequency of earthquakes.

*John C. Kimball High School
†Socratic Academy

It states that small earthquakes occur much more often than large earthquakes. In other words, if we count how many earthquakes happen, we will find many small ones and only a few large ones. This pattern is seen all over the world and is one of the most important observations in earthquake science. Omori's law describes how aftershocks behave after a major earthquake. After a large earthquake occurs, many smaller earthquakes follow, which are called aftershocks. These aftershocks happen frequently at first, but their frequency decreases over time. This means that right after the main earthquake, there are many aftershocks, but as time passes, they become less common. By adjusting the OFC model so that it follows both of these laws, the simulation becomes more realistic and reliable. The model is tuned so that it produces many small events and fewer large ones, and so that aftershock patterns decrease over time. This step is very important because it ensures that the model is not just producing random results, but is actually reflecting real scientific patterns observed in nature. Through this calibration process, the model becomes a better tool for understanding and studying earthquake behavior.

In this project, the OFC model was simulated using a computer to generate stress field data. This data represents how stress changes over time and space in the model. However, to make sure the model is realistic, it must be compared to real-world earthquake patterns. This process is called calibration. The model was calibrated using two important laws in earthquake science: the Gutenberg-Richter law and Omori's law. The Gutenberg-Richter law [3] states that small earthquakes happen much more often than large earthquakes. In other words, there are many small events and only a few large ones. Omori's law [7] describes how aftershocks behave after a major earthquake. It states that aftershocks happen frequently right after the main event, but their frequency decreases over time. By adjusting the model so that its results match these two laws, the simulation becomes more realistic and better represents how earthquakes actually occur in nature. This step is very important because it ensures that the model is not just random, but follows known scientific patterns.

In this project, the OFC model was used to create stress data over time, which is shown in the figure above. The grid represents a simple version of the Earth's surface, where each small square holds a certain amount of stress. The colors show how much stress is in each square. Dark colors (like purple and blue) mean low stress, while bright colors (like green and yellow) mean high stress. The color bar on the side shows the stress values, ranging from low (0) to high (1). The three images labeled (a), (b), and (c) show how the stress changes as time goes on in the simulation. In image (a), labeled "event 0," the stress is spread out randomly across the grid. There is no clear pattern yet, and the stress is fairly even. This shows the beginning stage, where stress is slowly building up but no major events have happened yet. In image (b), labeled "event 1990," the stress starts to form patterns. Some areas now have more stress, while other areas have less. This happens because when one point releases stress, it sends stress to nearby points. Over time, this creates uneven patterns across the grid. Certain areas start to become more stressed than others. In image (c), labeled "event 1995," there is a clear area where stress is very high. This bright region shows a buildup of stress that could lead to a larger event. Around this area, there are darker regions where stress has already been released. This looks similar to real earthquakes, where stress builds up in one place before being released suddenly. This stress data is important because it is used to train the machine learning model. Since the grid looks like an image, it can be used as input for a Convolutional Neural Network (CNN). The CNN learns to find patterns in how stress is spread across the grid. For example, it can learn that areas with high stress are more likely to have an event soon. Overall, the simulation shows how stress builds up, moves, and is released over time.

This method has strong potential for real-world use. If the same approach is applied to actual data from Earth's tectonic plates, it could help scientists better predict where and when earthquakes might occur. While exact prediction is still very difficult, this method could improve early warning systems and help identify high-risk areas. Better predictions could help governments and communities prepare in advance. For example, buildings could be strengthened, emergency plans could be improved, and people could be warned earlier. This would reduce damage, save money, and most importantly, save lives. By combining simulation models like the OFC model with advanced machine learning techniques, this research could be an important step toward improving

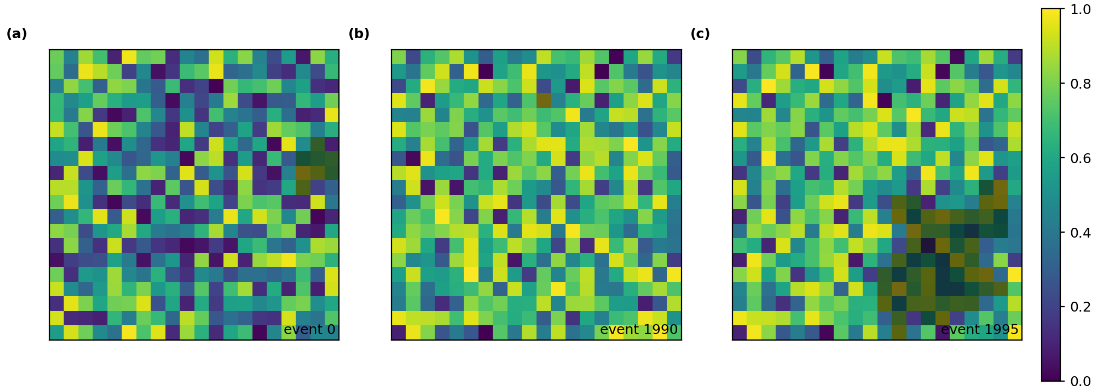

Figure 1: Pre-event stress fields in the OFC model with earthquake footprints overlaid. The background shows the lattice stress F, while shaded regions indicate sites that toppled during the corresponding earthquake.

earthquake prediction in the future.

Methods

Olami–Feder–Christensen (OFC) Model

The Olami-Feder-Christensen (OFC) model [2] is a simple way scientists study how earthquakes happen. Imagine a floor made of many small blocks arranged in a square grid, like tiles. Each block stores a certain amount of stress, which is like pressure building up inside the Earth's crust when tectonic plates push against each other. In the model, all the blocks slowly gain stress at the same time, just like the Earth slowly builds pressure along faults. Every block has a threshold, which is the maximum stress it can hold. When the stress on a block becomes too high and reaches this limit, the block suddenly "topples," similar to how a fault slips during a real earthquake.

When one block topples, it releases its stress and sends part of that stress to nearby blocks. Those blocks may already be close to their limits, so the extra stress can cause them to topple too. This can create a chain reaction, where many blocks fail one after another, representing an earthquake and its aftershocks. The more blocks that topple, the larger the earthquake is in the model. After the event ends, the stress on the toppled blocks resets to zero, and stress slowly builds again until another failure happens.

The reset of the stress on the toppled block is written as

$$F_i \to 0 \tag{1}$$

where F_i represents the stress on block i. This equation means that once the block fails (or topples), all the stress it was holding is released, so the block starts again with zero stress.

At the same time, a fraction of this released stress is transferred to neighboring blocks. This redistribution of stress is described by

$$F_{\text{neighbor}} \to F_{\text{neighbor}} + \alpha F_i \tag{2}$$

where F_{neighbor} is the stress on a neighboring block and α is a parameter that determines what fraction of the failed block's stress is passed to each neighbor. If α is large, more stress is transferred, making it more likely that nearby blocks will also topple and produce a cascade of failures.

Table 1: Confusion Matrix

	Predicted Negative	Predicted Positive
Actual Negative	1167 (6.10%)	2141 (11.19%)
Actual Positive	2717 (14.20%)	13110 (68.51%)

Gutenberg-Richter Law

The Gutenberg-Richter Law [6] explains how often earthquakes happen based on their size. It basically says that small earthquakes happen a lot, but large earthquakes are very rare. As the magnitude increases, the number of earthquakes drops very quickly. So, for every large earthquake, there are many smaller ones. In any region, you will usually see many small quakes, fewer medium ones, and very few large ones. This pattern occurs consistently over time.

The law follows an exponential pattern, which means the number drops a lot every time the magnitude increases by one number. For example, there might be 10 times fewer earthquakes at magnitude 6 compared to magnitude 5. Scientists use this law to study the earthquake risk and understand how energy is released in the Earth's crust. Overall, it shows that earthquakes follow a clear mathematical pattern, where bigger earthquakes are much less common but release much more energy.

$$N = 10^{a-bM} \tag{3}$$

Omori's Law

Omori's Law [4] explains how aftershocks decrease over time after a major earthquake, also called the main shock. Right after the main shock, there are usually many aftershocks happening close together in time. These aftershocks are smaller earthquakes that occur in the same area because the Earth is still adjusting and releasing stress. The first few days usually have the highest number of aftershocks.

As time passes, the number of after shocks slowly decreases. This follows an inverse relationship with time, meaning the longer you wait, the fewer aftershocks happen. For example, there might be a lot on day 1, about half as many on day 2, and much fewer by day 10. The decrease is gradual, not sudden, and aftershocks can continue for weeks or even months. This law helps scientists estimate how aftershock activity will fade and understand how the Earth returns to balance after a large earthquake.

$$n(t) = \frac{k}{(c+t)^p} \tag{4}$$

Results

The performance of the model is summarized in Table 1. The model correctly identifies a large number of positive cases, with 68.51% of all samples classified as true positives. This shows that the model is effective at detecting stress patterns that lead to large events. However, there are also some incorrect predictions. In particular, 14.20% of samples are false negatives, meaning the model fails to detect some upcoming large events. There are also 11.19% false positives, where the model predicts a large event that does not occur.

Figure 3 provides examples of stress fields for each prediction category. The true positive cases show stress patterns that the model correctly identifies as leading to large events. In contrast, false positives and false negatives highlight cases where the patterns are less clear or more difficult to distinguish. In many true positive examples, regions of high stress appear more localized,

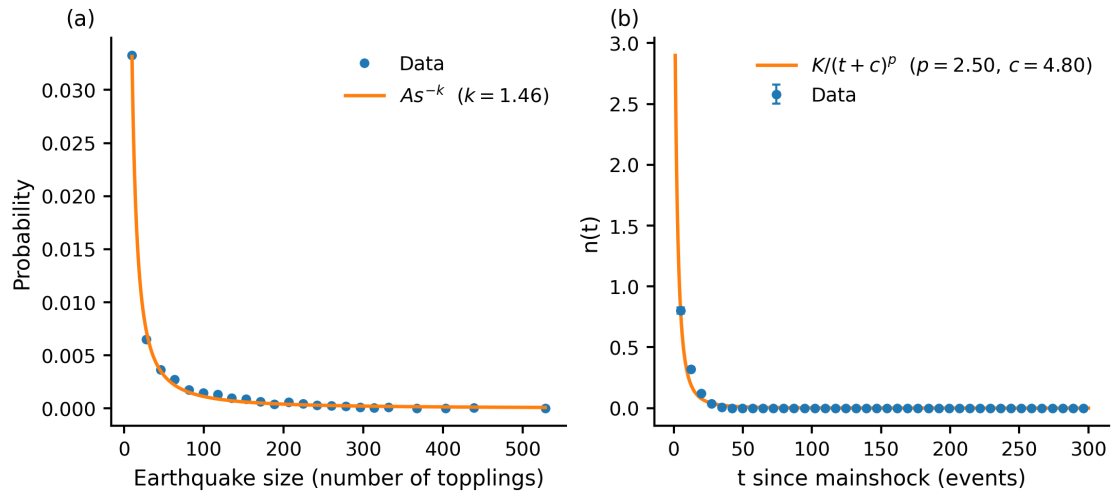

Figure 2: OFC model simulations are consistent with both the Gutenberg–Richter law (left) and Omori's law (right). The size distribution and aftershock decay each exhibit power-law behavior, characteristic of natural earthquake systems.

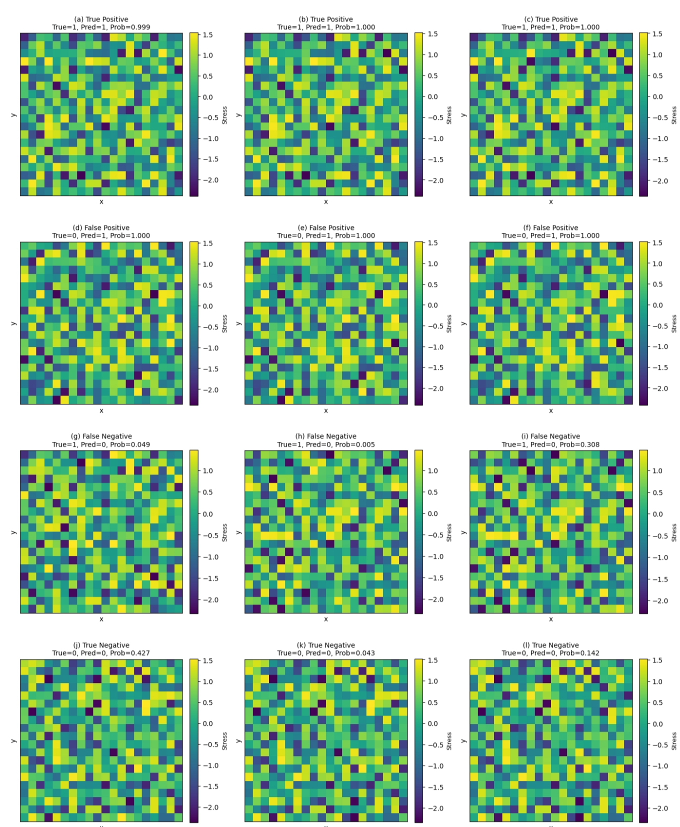

Figure 3: Stress fields grouped by prediction results. Each panel shows a stress field from the OFC model. The panels are organized as true positives, false positives, false negatives, and true negatives based on the model's predictions. The colors show the level of stress at each point on the grid.

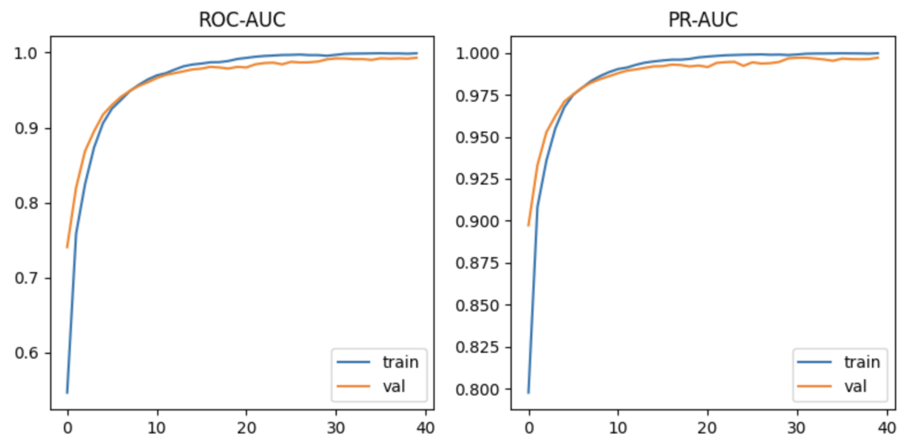

Figure 4: Training and validation performance of the model over epochs. (a) ROC-AUC and (b) PR-AUC both increase rapidly and stabilize near 1.0, indicating strong classification performance. The close agreement between training and validation curves suggests good generalization.

while misclassified cases tend to show more diffuse or less distinct patterns. This suggests that concentrated stress regions may be an important feature for predicting large events.

To further evaluate model performance, Figure 4 shows the ROC-AUC and PR-AUC curves over training epochs for both training and validation data. Both metrics increase rapidly in the early stages of training and then stabilize. On the test set, the model achieves a ROC-AUC of 0.67 and a PR-AUC of 0.90. The validation curves closely follow the training curves, suggesting that the model generalizes well and does not significantly overfit. The high PR-AUC is particularly important in this problem, since large events are relatively rare and the model must correctly identify them among many negative cases.

The results show that the model can learn meaningful spatial patterns from the simulated stress fields and use them to predict upcoming large events with good performance. At the same time, the presence of false positives and false negatives indicates that some stress configurations remain difficult to classify, leaving room for further improvement.

Conclusion

The OFC model was used to simulate stress evolution in a simplified earthquake system, and the generated stress fields were used to train a convolutional neural network to predict upcoming large events. The model was calibrated using the Gutenberg-Richter law and Omori's law to ensure that the simulated data follows realistic earthquake behavior.

The results show that the neural network can successfully learn patterns in the stress field and use them to predict future large events with high accuracy. The confusion matrix and performance metrics demonstrate strong predictive ability, and the ROC-AUC and PR-AUC results indicate that the model performs well even in the presence of class imbalance. Visualizations of the stress fields further suggest that the model is able to identify meaningful spatial structures related to stress buildup.

At the same time, some misclassifications remain, especially in cases where the stress patterns are less distinct. This highlights the inherent difficulty of predicting complex systems like earthquakes and suggests that additional features or improved models may further enhance performance.

This work shows that combining physics-based simulations with machine learning is a promising approach for studying earthquake prediction. While exact prediction remains challenging, this framework provides a useful step toward identifying conditions that increase the likelihood of large events and could contribute to improved early warning systems in the future.

References

[1] Filippo Caruso et al. "Analysis of self-organized criticality in the Olami-Feder-Christensen model and in real earthquakes". In: *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics* 75.5 (2007), p. 055101.

[2] Josué X De Carvalho and Carmen PC Prado. "Self-organized criticality in the Olami-Feder-Christensen model". In: *Physical review letters* 84.17 (2000), p. 4006.

[3] GS Golitsyn. "Place of the Gutenberg-Richter Law Among Other Statistical Laws of Nature". In: *Selected Papers From Volume 32 of Vychislitel'naya Seysmologiya* 7 (2005), pp. 119–137.

[4] Anatol Vladimirovich Guglielmi. "Omori's law: A note on the history of geophysics". In: *Physics–Uspekhi* 60.3 (2017), pp. 319–324.

[5] Hiroo Kanamori and Emily E Brodsky. "The physics of earthquakes". In: *Physics today* 54.6 (2001), pp. 34–40.

[6] P Palacios, I Molina, and M Segovia. "The Gutenberg–Richter Law: assumptions, limitations and interpretations". In: (2006).

[7] JA Vallejos and SD McKinnon. "Omori's law applied to mining-induced seismicity and reentry protocol development". In: *Pure and applied geophysics* 167.1 (2010), pp. 91–106.