
DEEP VS. SHALLOW LEARNING: A BENCHMARK STUDY IN LOW MAGNITUDE EARTHQUAKE DETECTION

A PREPRINT

Akshat Goel

Department of Computer Science
University College London

Denise Gorse

Department of Computer Science
University College London

May 3, 2022

ABSTRACT

While deep learning models have seen recent high uptake in the geosciences, and are appealing in their ability to learn from minimally processed input data, as 'black box' models they do not provide an easy means to understand how a decision is reached, which in safety-critical tasks especially can be problematical. An alternative route is to use simpler, more transparent 'white box' models, in which task-specific feature construction replaces the more opaque feature discovery process performed automatically within deep learning models. Using data from the Groningen Gas Field in the Netherlands, we build on an existing logistic regression model by the addition of four further features discovered using elastic net driven data mining within the catch22 time series analysis package. We then evaluate the performance of the augmented logistic regression model relative to a deep (CNN) model, pre-trained on the Groningen data, on progressively increasing noise-to-signal ratios. We discover that, for each ratio, our logistic regression model correctly detects every earthquake, while the deep model fails to detect nearly 20 % of seismic events, thus justifying at least a degree of caution in the application of deep models, especially to data with higher noise-to-signal ratios.

Keywords Seismology · Machine learning · Feature selection · Benchmark study

1 Introduction

The geosciences have been transformed in recent years by a substantial growth in the quantity and quality of available data. This has spurred interest in machine learning methods for seismological tasks, and especially in deep learning models, as powerful feature extractors (see, for example, the benchmarking study of [1]). However, as 'black box' models, they do not provide an easy means to understand how a decision is reached. The aims of this paper are first to demonstrate the power of a 'white box' model and second to directly compare this model to a deep neural network. Using data from the Groningen Gas Field in the Netherlands, we build on the logistic regression model of [2] by the addition of four further features. We then evaluate the performance of the augmented model relative to the CNN of [3], pre-trained on the Groningen data, on progressively increasing noise-to-signal ratios. We discover that, for each ratio, our logistic regression model correctly detects every earthquake, while the deep model fails to detect nearly 20 % of seismic events. It may thus be concluded that deep models should be treated with a degree of caution, since, in addition to their lack of transparency, they may not always perform ideally well.

2 Data and Methods

2.1 Data acquisition, pre-processing, and partitioning

The data used in this study, as in [2] and [3], were derived from the G-network, a network of 70 seismic stations set up by the Dutch government after public protest over increased induced seismicity due to gas extraction from the

Groningen Gas Field. Each G-network station has four geophones, between 50 m and 200 m, and we gather data from all four geophone levels as the comparison model of [3] requires this. We downloaded seismograms for event examples from web services hosted by the Royal Netherlands Meteorological Institute (KNMI), applying the same date and magnitude (≥ 0.2 , resulting mean 1.05) selection criteria as in [2]. This resulted in 2300 seismograms recording 47 events. We were unable to obtain 4000 noise samples, as used in [2], from the 2017-18 period of the studies of [2] and [3] and so used the same time of year in 2020-21. Additional negative examples for the increasing noise-to-signal phase of this study were drawn from this same period, with 38100 new samples collected. All waveform data were detrended, demeaned, and bandpass-filtered to frequencies of 5-25 Hz as in [3]. In addition, in order to be compatible with the CNN of [3], the resulting seismograms were then downsampled by a factor of two. We chose to segregate our train (60 %), validation (20 %), and test (20 %) datasets by event rather than seismogram (seismogram-based segregation having been used in the studies of [2] and [3]) in order to prevent data leakage from test to train datasets. Hence, for example, our test set may not contain 20 % of the seismograms, as events are detected by varying numbers of G-network stations. This difference, and the differing time period from which noise data were derived, are not problematic for internal comparisons within this study. However, they should be borne in mind when comparing our results directly with those of [2] and [3].

2.2 Machine Learning Models

The primary model used in this study is logistic regression (LR). The major appeal of this linear model is that its fitted weights can be used to understand the role and importance of different features. An elastic net penalty

$$\left[\alpha \sum_{i=1}^P |\beta_i| + (1 - \alpha) \sum_{i=1}^P \beta_i^2 \right]$$

where the β_i are the (external inputs + 1) weights of the model, may be added to the LR loss function; for a suitable setting of the hyperparameter α this can encourage not only smaller weights but variable selection. In our work the elastic net penalty was used only for this latter purpose, as a means to filter new candidate features according to their importance, our final model being a simple, interpretable LR model as in [2].

We benchmarked against a convolutional neural network (CNN), a type of deep learning model originally designed for image data, and used in [3]. The CNN was not retrained as it had already been trained on data from the Groningen Gas Field. The CNN of [3] was designed to take advantage of the multiple geophone levels used in the G-network, leveraging the potential of the moveout pattern of energy to distinguish between disturbances originating underground (more likely to be seismic in origin) and ones originating at the surface (more likely to be noise).

2.3 Performance measures

In order to explore the behaviour of the models for increasing amounts of noise it was important to use a performance measure that is robust for imbalanced data. Accuracy is not a good choice in this respect, though we quote it as a secondary metric in our results as it is widely used elsewhere even for imbalanced datasets. Our preferred measure will be the more robust Matthews Correlation Coefficient, defined by

$$MCC = [TP * TN - FP * FN] / [(TP + FP) * [TP + FN] * [TN + FP] * [TN + FN]]^{1/2}$$

where TP denotes true positives, TN true negatives, FP false positives, and FN false negatives. An MCC of 1.0 corresponds to a perfect prediction, a value of 0.0 either random prediction or an assignment of all examples to a single class. It is this last behaviour that makes the MCC valuable for imbalanced datasets, as the tendency of most machine learning models is to over-assign examples to the majority class.

2.4 Use of the HCSTA and catch22 software packages

Most of the work that has been done in earthquake detection with LR models has used seismologically derived features (see, for example, [4], and references therein). But it is also of interest to explore statistically-derived features, as was done by [2], who used features from the HCTSA (Highly Comparative Time-Series Analysis) package of [5]. HCTSA offers around 7,700 candidate features, with the four used in [2] being selected by a mixture of mechanisms integral to the package and qualitative judgement. The catch22 MATLAB package [6] contains a subset of features the HCTSA authors found to be best-performing for time series classification. Because there are only 22 of these features it is easy to add them to the inputs of an LR model with an elastic net penalty and use this model's inherent feature selection ability to highlight potentially useful new features. This means of feature selection differs from that used in [2] to choose the original four features, and it seemed possible it could discover some potentially valuable features the selection method of the earlier paper had missed. In fact this was so, as will be shown in the Results section to follow.

3 Results

3.1 Feature discovery using elastic net




Figure 1: Distribution of feature values for the 72 best-performing elastic net models.

200 runs of the elastic net model (which had 26 inputs, the original four HCTSA features from [2] plus 22 new features from the catch22 package) were carried out. It was discovered that there was no single best set of weights; there were in fact many elastic net models with identical best performance in terms of the validation MCC. The distribution of weights for each feature for these 72 equally well-performing models is shown in Figure 1, above. This figure confirms the value of the original four HCTSA features (W1 to W4), but also suggests the potential value of four catch22 features (C10, C11, C14, and C15), with C11 especially promising. These four new features were added to the original four. The 18 less influential catch22 features and the use of the elastic net penalty were then discarded in order to have an interpretable LR model as in [2].

3.2 Comparison of the logistic regression and CNN models at low noise-to-signal ratio

Both the LR model of [2] and our augmented LR model were trained at the noise-to-signal ratio of 1.73:1 used in [2]. The CNN of [3], which was not retrained by us, had been pre-trained at a ratio of 1.47:1. Both models were tested in this subsection at the ratio of 1.73:1 used for testing in [2]. From Table 1 it can be seen the original LR model is able to achieve an MCC of 0.9691, but that our augmented model is even so able to statistically significantly improve on it. However, the CNN does substantially less well. From the confusion matrices in Table 2 it can be seen that the CNN is very effective (having no false positives) in classifying noise, but makes a substantial number of errors (false negatives) on earthquake data, failing to detect around 20 % of the earthquake examples. As will be seen in the next subsection, the CNN's focus on correctly detecting noise becomes only more problematic as the proportion of noise to signal increases.

Model	Model description	Model type	Test Acc.	Test MCC	Significant at 5 %
1	Model of [2]	LR	98.56	0.9691	-
2	As above + selected catch22	LR	99.18	0.9824	Yes (2 > 1)
3	Model of [3]	CNN	82.94	0.8648	Yes (3 < 2, 1)

Table 1: Test performance at noise ratio 1.73:1 of the three models considered, in terms of both accuracy and MCC, where > (<) denotes that one model is statistically significantly better (worse) than another.

	Predicted negative	Predicted positive
True negative	1313	5
True positive	0	763

Confusion Matrix: LR Model of [2] + Selected catch22 Features

	Predicted negative	Predicted positive
True negative	1318	0
True positive	127	627

Confusion Matrix: CNN Model of [3]

Table 2: Confusion matrices for the better-performing of the two logistic regression models (model of [2] + selected catch22) and for the benchmark CNN model of [3].

3.3 Performance of the models on less balanced datasets

Table 3 shows that as the noise ratio is increased all the models, which were trained on low ratios as described above, find classification more difficult. It can be seen, however, that the four catch22 features are of increasing value as the noise increases: at the highest noise-to-signal ratio considered in this study, 50:1, the augmented model’s MCC is around 36 % better than that of the original model.

No.	Model description	1.73:1	5:1	10:1	25:1	50:1
1	Model of [2]	0.9691	0.7663	0.4062	0.2388	0.1642
2	As above + selected catch22	0.9824	0.8203	0.515	0.3164	0.2227
3	Model of [3]	0.8648	0.8921	0.8998	0.9045	0.9061

Table 3: Test MCCs for the two LR models of [2] and the benchmark CNN of [3] as the noise-to-signal ratio is increased from the baseline of 1.73:1 used in [2].

It is not a surprise that the LR models find prediction on less balanced data harder. What is a surprise is that the CNN appears to find it easier. However, inspection of the confusion matrices for the higher noise ratios explains this: at these higher ratios the CNN makes the same errors on the earthquake examples as it did at the starting ratio of 1.73:1, but continues to classify all noise examples perfectly; as the proportion of noise increases, the CNN therefore necessarily appears to do better.

4 Conclusions

This study benchmarked an eight-input logistic regression (LR) model (four new input features, discovered via the use of the catch22 software package combined with an elastic net feature selector, having been added to the model of [2] against a convolutional neural network (CNN) pre-trained on data from the same source, the Groningen Gas Field in the Netherlands [3]. Tests were carried out both at the low noise-to-signal ratio of [2] and at higher noise ratios. While the CNN scored higher as this ratio increased, as measured by the Matthews Correlation Coefficient (MCC), the mistakes made by the LR model were notably all false positives, which can be reduced by seismological agencies using pre-existing or easily adoptable [4] processes, while the mistakes made by the CNN were the more

serious and problematic false negatives. We do not claim deep learning models have no place in seismology. They can be remarkably powerful; the EQTransformer model of [7] has been shown particularly effective in recent benchmarking trials [1]. However, deep models need to be trained with care and have yet to be proven on substantially imbalanced datasets. A degree of caution may thus be warranted and the time may not yet be right for such models to be deployed in practice.

5 Acknowledgments

The authors would like to thank Umair bin Waheed for helpful discussions and advice, and Ahmed Shaheen for the provision of the pre-trained CNN model used in [3].

References

- [1] Jannes Münchmeyer, Jack Woollam, Andreas Rietbrock, Frederik Tilmann, Dietrich Lange, Thomas Bornstein, Tobias Diehl, Carlo Giunchi, Florian Haslinger, Dario Jozinović, Alberto Michelini, Joachim Saul, and Hugo Soto. Which picker fits my data? A quantitative evaluation of deep learning based seismic pickers. *Journal of Geophysical Research: Solid Earth*, 127(1), 2022.
- [2] Umair bin Waheed, Ahmed Shaheen, Mike Fehler, and Ben Fulcher. Winning with simple learning models: Detecting earthquakes in Groningen, the Netherlands. In *82nd EAGE Annual Conference & Exhibition*, pages 1–5, 2020.
- [3] Ahmed Shaheen, Umair bin Waheed, Michael Fehler, Lubos Sokol, and Sherif Hanafy. GroningenNet: Deep learning for low-magnitude earthquake detection on a multi-level sensor network. *Sensors*, 21:8080, 2021.
- [4] Julian D. Miranda, Ciro A. Gamboa, Angelica Florez, and Miguel Altuve. Voting-based seismic data classification system using logistic regression models. In *2019 XXII Symposium on Image, Signal Processing and Artificial Vision (STSIVA)*, pages 1–5, 2019.
- [5] Ben Fulcher and Nick Jones. Highly comparative feature-based time-series classification. *IEEE Transactions on Knowledge and Data Engineering*, 26:3026–3037, 2014.
- [6] Carl Henning Lubba, Sarab Sethi, Philip Knaute, Simon Schultz, Ben Fulcher, and Nick Jones. catch22: CAnonical Time-series CHaracteristics selected through highly comparative time-series analysis. *Data Mining and Knowledge Discovery*, 33:1821–1852, 2019.
- [7] S. Mostafa Mousavi, William L. Ellsworth, Weiqiang Zhu, Lindsay Y. Chuang, and Gregory C. Beroza. Earthquake transformer: An attentive deep-learning model for simultaneous earthquake detection and phase picking. *Nature Communications*, 11(1):3952, 2020.