

Teleseismic Discrimination Using Deep Learning

Rayna Arora¹, Ronan Le Bras²

¹Comprehensive Nuclear-Test-Ban Treaty Organization (CTBTO) Youth Group,

<https://orcid.org/0000-0002-9563-9995>

²Comprehensive Nuclear-Test-Ban Treaty Organization (CTBTO), Vienna, Austria,

<https://orcid.org/0000-0003-2439-6938>

Key Points:

- Deep Learning can be used on seismic waveforms to discriminate between earthquakes and explosions at teleseismic distances.
- A model built on waveform inputs rather than spectrograms can achieve better results with fewer parameters.
- This work can be used to monitor treaty compliance and build global seismic hazard maps using a sparse seismic network.

Abstract

We study the problem of discrimination between earthquakes and explosions on the basis of seismic signals detected at teleseismic distances (over 2000 km). Most work in the field of discrimination has been limited to signals detected within a few hundred kilometers which limits their utility from the perspective of sparse global seismic networks for either treaty monitoring or seismic hazard analysis. We show that existing Deep Learning architectures that have been proposed for discrimination or related tasks such as phase classification or signal detection can be repurposed for teleseismic discrimination. Using hyperparameter tuning methods we have been able to improve the performance relative to the original architectures while reducing the model complexity. We present empirical analysis of a number of different methods, and demonstrate that our proposed Deep Learning architecture performs the best at teleseismic discrimination and is able to reliably identify rockburst events.

Plain Language Summary

Seismic events are caused mostly by naturally occurring earthquakes or manmade explosions. The capability to discriminate between event types is very important. It can be used to build maps of regions that are prone to earthquake damage and to monitor whether signatories of a treaty banning nuclear explosions are following through with their commitment. Previously, research in discrimination has focused mainly on waveforms detected by seismic stations within a few hundred kilometers of the event. This distance limitation implies that the methods can only be used in regions of the earth that have dense seismic networks which tend to be concentrated in developed countries. Thus, our work seeks to address the discrimination problem using waveforms detected at more than 2000 km away (known as teleseismic distances), guaranteeing that all locations on the earth's surface are within a few thousand kilometers of some station in global seismic networks. We explore various Deep Learning methods that have been developed for discrimination and systematically enhance them for teleseismic discrimination. We demonstrate our methods work well not only on the standard discrimination task for which they were trained, but also on completely unseen seismic events that have an energy release mechanism similar to explosions.

1 Introduction

An accurate catalog of earthquakes is required for probabilistic seismic hazard analysis (McGuire, 2008). Such hazard analysis is the basis for building codes that help determine the safety of human occupants as well as critical infrastructure. A common problem with such seismic catalogs is the proliferation of mining activity that gets inadvertently captured and contaminates the catalog as noted by (Mackey et al., 2003), for example. Consequently, an important post-processing step after the creation of a catalog is to classify the explosion events, and to remove them from the bulletin. This task of classifying an event as an explosion versus an earthquake, is known as the *discrimination* problem in seismology.

Another important application of this discrimination task is for monitoring compliance with treaties banning nuclear testing. In the last century, the major nuclear powers have signed various treaties limiting the size and number of nuclear tests that they could conduct. These limits were introduced to curb the development of increasingly lethal weapons which have endangered our existence as a species as well as to protect earth's ecosystems from the effects of the radiation released by these tests. These efforts to rein in nuclear testing culminated in the Comprehensive Nuclear-Test-Ban Treaty (CTBT)(UN, 1996), which bans all nuclear testing anywhere on earth, and has been signed by 185 countries. The CTBTO is an international organization charged with the verification of the CTBT.

While earthquakes and explosions can both generate a vast amount of energy, the forces involved are very different. Explosions release energy in a small volume around the source, and this causes mostly compressional waves, or P phases, to radiate outwards isotropically. Earthquakes, on the other hand, release energy over multiple kilometers along a fault line and mostly in the form of shear waves, or S phases, that could have an anisotropic radiation pattern. These differences show up in the waveforms that are detected at seismic stations, and form the basis of work on discrimination.

At distances less than 2000 km, a number of seismic *phases* (Bormann et al., 2013), each of which corresponds to a distinct wave path through the Earth, are normally detected. Many methods rely on contrasting the detections of multiple phases from the same event. For example, the spectral characteristics of the P versus the L_g phase (a type of surface wave that can be detected up to a few hundred kilometers) are an easy marker for discrimination purposes as shown in (Dysart & Pulli, 1990). At teleseismic distances, however, time differences between various phases are large and they are well separated.

There are some unfortunate ramifications of having good discrimination methods that only work well at short distances. The main issue is that these methods can only be applied to events that occur in regions of the earth that are covered with very dense seismic networks. In other words, if a country wishes to use the latest techniques for seismic hazard analysis it must have the resources to deploy and maintain many stations. This creates an inequity in preparedness which causes poorer countries to be affected significantly more by natural disasters than richer countries. (Nairobi, 2005).

For the purpose of monitoring compliance with a global treaty such as the CTBT, these regional discrimination methods have limited utility. The International Monitoring System (IMS), installed and maintained by the CTBTO, includes seismic networks which are too sparse to provide an assurance that a nuclear explosion test will be detected by at least one of the stations within 200 km. For example, (Stump et al., 2002) has done an analysis of the IMS network when it is fully operational. This work shows that 90% of the earth's land mass will be covered by at least one station within 2000 km, but only 10% is within 200 km of an IMS station. Of course, member states that have signed the CTBT are free to use data obtained from other seismic networks to make their own determination as to the source of an event. However, there is no guarantee that such additional data will be available for all regions of the world. This brings into question the ability of the treaty to be effectively monitored. Given that the US Senate has yet to ratify the treaty citing concerns of monitoring capability (Pifer, 2016), it is important to address these concerns for the ultimate safety of the earth's population from further development of atomic weapons.

For the above reasons, our work aims to extend the capabilities of existing discrimination methods. Although it is much harder to accurately classify seismic waveforms at distances over 2000 km, we note that it is not entirely intractable. For example, we show waveforms in Figures 1 from an earthquake and an explosion detected at teleseismic distances. In these two examples, it is easy to determine the source of the event from the waveform. For example, the impulsive nature of the explosion event shows up in the sharp onset spike in Figure 1a while the earthquake in Figure 1b has a slower climb to its peak amplitude. Of course, for smaller magnitude events at greater distances such clear features are not immediately obvious. For example in Figure 2 the difference between the waveforms are not immediately apparent to a visual inspection. In fact, it is not even apparent in these last two waveforms whether there is a detection at all. In all the examples in the paper we show a 90 second snippet of the waveform with the actual onset of the event at the 10 second mark. All the waveforms are passed through a 1 Hz high-pass filter and normalized.




Figure 1. Examples of waveforms which are visually easy to discriminate.

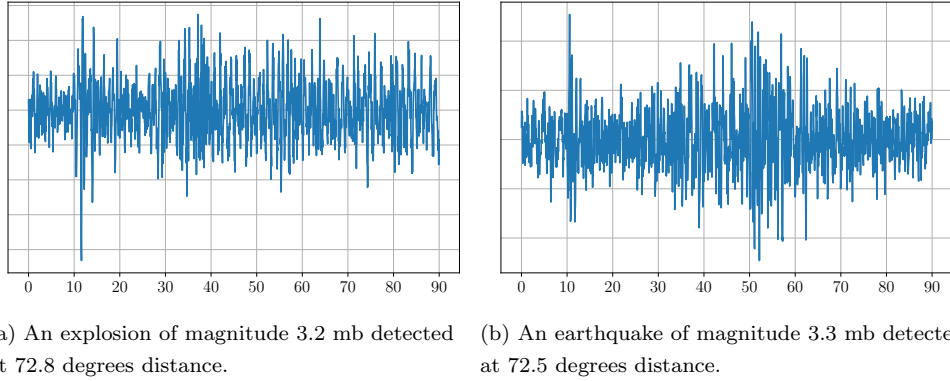


Figure 2. Examples of visually ambiguous waveforms.

In Section 2 we describe some of the previous work related to analyzing seismic waveforms. Next we describe our data sources in Section 3 and our discrimination algorithm in Section 4. Finally, we describe our experiments in Section 5.

2 Related Work

The importance of discrimination between earthquakes and explosions for the purpose of seismic hazard analysis as well as nuclear explosion monitoring is well documented in the literature. We refer the reader to (Rabin et al., 2016) for a good overview of this topic. As mentioned in that paper, many of the previously developed methods rely on computing parameters computed from input waveforms such as event magnitude. Further, these parametric methods often require the detection of different types of seismic phases to be used for discrimination. For example, the $M_s:mb$ ratio method (Blandford, 1982) requires the detection of a surface wave and a body wave in order to compute two different magnitude estimates. The limitation of such methods for a global seismic network are apparent in the statistics. (Rabin et al., 2016) goes on to report that such parametric methods are only applicable to 60% of all events reported by the CTBTO in years 2011-2013.

In recent years, the focus has shifted to applying Deep Learning directly on spectrograms for the purpose of discrimination, for example (Magana-Zook & Ruppert, 2017). Also, in (Linville et al., 2019) the authors show results on architectures based on Con-

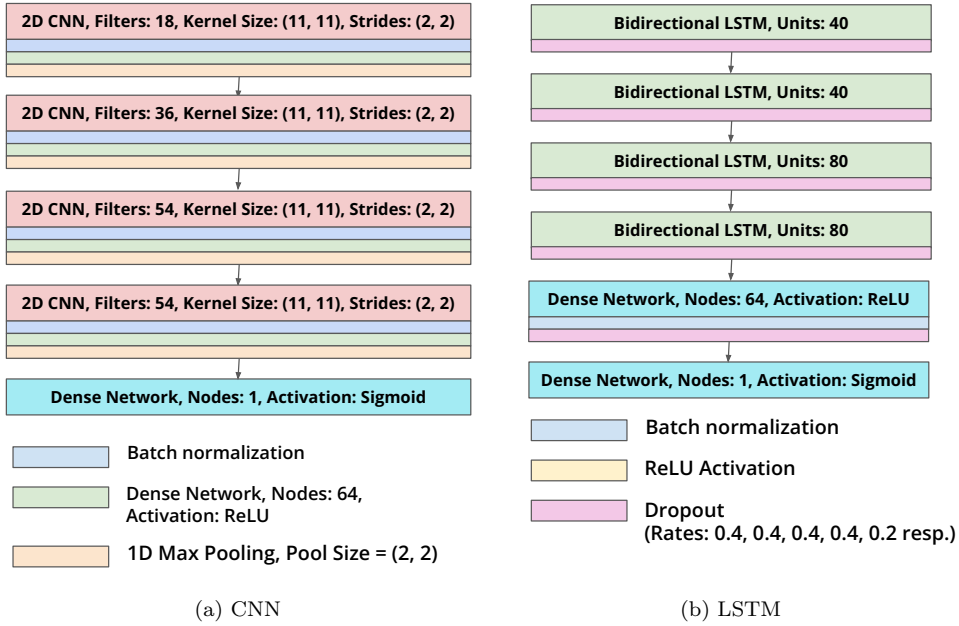


Figure 3. Neural Network architectures in Linville et al., 2019.

volutional Neural Nets (CNNs) as well as Long Short-term Memory (LSTM). We show these architectures in Figures 3a and 3b respectively. Although these architectures have only been tested for discrimination at regional distances, we believe that these can be adapted to teleseismic distances since they require nothing more than a spectrogram of the waveform as an input. However, other approaches such as (Ranasinghe et al., 2019), which are also based on CNNs, but require both a *P* and an *S* phase to be detected are not as easy to use since *S* phases are rarely detected at larger distances. Similarly, approaches based on Support Vector Machines such as (Kim et al., 2020) that require features from both *P* and *S* phases are inapplicable for teleseismic discrimination.

We note that there are a number of methods based on waveforms that we could extend to teleseismic distances such as (Pezzo et al., 2003), (Li et al., 2018), (Ross et al., 2018), (Miao et al., 2020), (Wei et al., 2020), or (Ray et al., 2017). All of them are designed to solve various classification tasks in seismology without any assumptions on explicit feature extraction. In this paper we focus on Cnn-Rnn Earthquake Detector (CRED) (Mousavi et al., 2019) since its architecture is the most versatile. As shown in Figure 4a, it uses a mix of convolutional and recurrent layers in a residual structure. CRED was proposed to discriminate seismic phases from noise, but there is no reason it can't be used for other classification tasks. Similar to most other models it uses the spectrograms of the waveforms as input, and has only been tested on regional detections.

3 Data

There is no existing dataset for teleseismic discrimination, so as part of our work we had to build one. We collected seismic event data from the online bulletins published by the International Seismic Centre (ISC)(Storchak et al., 2017, 2020) and corresponding waveforms from the waveform repository published by the Incorporated Research Institutions for Seismology (IRIS).



Figure 4. Side by side comparison of CRED and our architecture.

The ISC event bulletins have an optional field describing the source of each event. For a small subset of these events this field is filled in with either earthquake, explosion, or rockburst (Lu et al., 2013). We selected all the events with these three possible sources and then used the IRIS waveform repository to obtain the waveforms from stations which had detected the event according to the ISC bulletin. However, some waveforms didn't show any clear signal and had to be discarded. We used the STA/LTA (short-term average of 5 seconds divided by the long-term average of 20 seconds) method (Allen, 1978) to determine if there was a suitable detection in the waveform. For each valid waveform, we kept 90 seconds of data; 10 seconds before the onset and 80 seconds after the onset. This waveform snippet was high-pass filtered at 1 Hz and normalized. We only kept the waveform corresponding to the vertical channel and downsampled it to 20 samples per second.

The following steps describe the overall data gathering procedure.

- For each ISC event bulletin from 1970 to 2018:
 - For each event in bulletin, if event source is either earthquake, explosion, or rockburst:
 - * For each station at a distance of 20 degrees or more that detects the event:
 1. Download a waveform snippet around the onset time.
 2. High-pass filter at 1 Hz.
 3. If the STA/LTA reaches or exceeds 2 within 5 seconds of arrival time:
 - Add this waveform to the dataset. Record the event parameters and distance of station to event.

In order to train a balanced Deep Learning model, we used stratified sampling of the earthquake waveforms to ensure that our data included the same number of earthquakes and explosions in each distance bucket from 20 to 180 degrees in steps of 10 degrees.

In total, we have 7608 data points compiled from data published by ISC and IRIS. All of the source code and dataset used in our project is available at Zenodo via DOI 10.5281/zenodo.5167966 with MIT license (RaynaArora, 2021). This includes the code to download and compile the data set, build models on the data, and to analyze the results.

4 Methods

In our work, we use the CRED architecture as a basis. However, instead of using a spectrogram of the waveform as input we found the performance to be better by using the waveforms directly. We also made a number of other simplifications to the model by systematically identifying the best architecture. We used grid search to tune the following hyperparameters on the validation data:

1. The kernel size in each convolutional layer, varied from 10 to 40 in steps of 2.
2. The number of filters in each convolutional layer, varied from 16 to 320 in steps of 16.
3. The units (output dimension) in each LSTM or Bidirectional LSTM layer, varied from 32 to 128 in steps of 32.
4. The number of convolutional, LSTM, and Bidirectional LSTM layers, varied from 1 to 5 in steps of 1.
5. Whether or not to use a residual layer (only when considering more than one convolutional layer).
6. The dropout rates in the LSTM and dense layers, varied from 0.1 to 0.4 in steps of 0.1.
7. The number of nodes in the dense layer — 64 or 128.
8. Whether to use a spectrogram or waveform as input.

The resulting optimal architecture is described in Figure 4b.

5 Experiments

We used the Keras software package with the Adam optimizer for training all the Deep Learning architectures that we evaluated. In our experiments we divided the data into 80% training, 10% validation, and 10% test. All the hyperparameters were tuned on the validation set. The models were trained with a batch size of 32 and training was stopped when the validation accuracy plateaued for 60 epochs.

Table 1a shows the results of our architecture as well as 3 other architectures in the literature. It also includes a baseline logistic regression model computed from the onset time and dominant frequency of each waveform for reference. Our model achieves the highest accuracy and Area Under the Curve (AUC) and has the least number of parameters by a factor of 2.

Figure 5a shows our model’s accuracy versus distance detected. This graph demonstrates that our model generalizes well across all distances. We also plot the accuracy of our model by event magnitude in Figure 5b. Our accuracy is quite low in the 3 to 3.5 range, which comprises less than 1% of our data, and hence we have insufficient samples to train or evaluate. On the other hand, our accuracy of 90% in the 3.5 to 4.0 mb range is very promising because this range is critical for treaty monitoring purposes.

Model	Accuracy	AUC	Parameters
Proposed Method	89.0 %	0.95	96,641
Mousavi et al. (CRED)	87.2 %	0.94	208,273
Linville et al. CNN	84.7 %	0.91	456,237
Linville et al. RNN	82.7 %	0.91	375,425
Baseline	69.8 %	0.72	3

(a) Accuracy on test data comprising only earthquakes and explosions.

Model	Rockburst Accuracy
Proposed Method	95.8 %
Mousavi et al. (CRED)	98.3 %
Linville et al. CNN	95.4 %
Linville et al. RNN	91.3 %

(b) Accuracy on test data comprising of only rockbursts.

Table 1. Accuracy of various approaches.




Figure 5. Accuracy of our architecture along distance and magnitude dimensions.

5.1 Rockburst

In the previous experiments, we trained and tested our models on data from explosions and earthquakes only. However, in order to test the generalization capability of our models, we took these models and tested them on rockburst data. As explained in (Ma et al., 2015)

Rockburst is the sudden release of elastic strain energy in rock masses under high local stresses, as a result of rock fragmentation, ejection, projection and even earthquakes.

Given the energy-release mechanism of rockbursts is similar to explosions we expect our models to classify them as such. As shown in Table 1b, our model performs reasonably well on this unseen dataset.

6 Conclusion

We have demonstrated for the first time that it is possible to discriminate between earthquakes and explosions detected at teleseismic distances. Our approach of using seismic waveforms rather than spectrograms and tuning the model hyperparameters and architecture has led to a novel architecture with far fewer parameters than existing work.

Our work achieves the best results on the discrimination task and competitive performance on unseen rockburst data. Even with limited data we were able to achieve nearly 90% accuracy on teleseismic discrimination. The dataset synthesized from existing sources can be used to expand upon this work.

We have taken an important step towards making treaty monitoring feasible with a sparse global seismic network. Furthermore, we expect that our work can be applied to improve seismic hazard analysis in developing nations with a sparse seismic network, thereby mitigating the disproportionate casualties counts these countries suffer due to earthquakes.

Acknowledgments

The views expressed herein are those of the author(s) and do not necessarily reflect the views of the CTBTO Preparatory Commission.

Dr. Nimar S. Arora provided valuable guidance on the general direction of this project.

The facilities of IRIS Data Services, and specifically the IRIS Data Management Center, were used for access to waveforms, related metadata, and/or derived products used in this study. IRIS Data Services are funded through the Seismological Facilities for the Advancement of Geoscience (SAGE) Award of the National Science Foundation under Cooperative Support Agreement EAR-1851048.

All seismic data were downloaded through the IRIS Wilber 3 system <https://ds.iris.edu/wilber3/> or IRIS Web Services <https://service.iris.edu/>, including the following seismic networks: (1) the AZ (ANZA; UC San Diego, 1982); (2) the TA (Transportable Array; IRIS, 2003); (3) the US (USNSN, Albuquerque, 1990); (4) the IU (GSN; Albuquerque, 1988).

This research doesn't create a new dataset. It is based entirely on existing publicly available datasets. We have provided a Jupyter notebook to download the public data and train all the models in this paper as well as all the analysis. The downloaded dataset has also been provided as an alternative. The notebooks and dataset are available at (RaynaArora, 2021).

References

- Allen, R. V. (1978). Automatic earthquake recognition and timing from single traces. *Bulletin of the Seismological Society of America*, 68(5), 1521–1532.
- Blandford, R. R. (1982). Seismic event discrimination. *Bulletin of the Seismological Society of America*, 72(6B), S69–S87.
- Bormann, P., Storchak, D. A., & Schweitzer, J. (2013). The iaspei standard nomenclature of seismic phases. In *New manual of seismological observatory practice 2 (nmsop-2)* (pp. 1–20). Deutsches GeoForschungsZentrum GFZ.
- Dysart, P. S., & Pulli, J. J. (1990). Regional seismic event classification at the noress array: seismological measurements and the use of trained neural networks. *Bulletin of the Seismological Society of America*, 80(6B), 1910–1933.
- Kim, S., Lee, K., & You, K. (2020, Mar). Seismic discrimination between earthquakes and explosions using support vector machine. *Sensors*, 20(7), 1879. Retrieved from <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7180981/> doi: 10.3390/s20071879
- Li, Z., Meier, M.-A., Hauksson, E., Zhan, Z., & Andrews, J. (2018, May). Machine learning seismic wave discrimination: Application to earthquake early warning. *Geophysical Research Letters*, 45(10), 4773–4779. Retrieved from <https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2018GL077870> doi: 10.1029/2018gl077870
- Linville, L., Pankow, K., & Draelos, T. (2019, Feb). Deep learning models augment analyst decisions for event discrimination. *Geophysical Research Letters*, 46(7), 3643–3651. Retrieved from <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2018GL081119> doi: 10.1029/2018gl081119
- Lu, C.-P., Dou, L.-M., Zhang, N., Xue, J.-H., Wang, X.-N., Liu, H., & Zhang, J.-W. (2013). Microseismic frequency-spectrum evolutionary rule of rockburst triggered by roof fall. *International Journal of Rock Mechanics and Mining Sciences*, 64, 6–16.
- Ma, T., Tang, C., Tang, L., Zhang, W., & Wang, L. (2015). Rockburst characteristics and microseismic monitoring of deep-buried tunnels for jinping ii hydropower station. *Tunnelling and Underground Space Technology*, 49, 345–368. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0886779815000814> doi: <https://doi.org/10.1016/j.tust.2015.04.016>
- Mackey, K. G., Fujita, K., Gounbina, L. V., Koz'min, B. M., Imaev, V. S., Imaeva, L. P., & Sedov, B. M. (2003, 04). Explosion contamination of the northeast siberian seismicity catalog: Implications for natural earthquake distributions and the location of the tanlu fault in Russia. *Bulletin of the Seismological Society of America*, 93(2), 737–746. Retrieved from <https://doi.org/10.1785/0120010196> doi: 10.1785/0120010196
- Magana-Zook, S. A., & Ruppert, S. D. (2017, December). Explosion Monitoring with Machine Learning: A LSTM Approach to Seismic Event Discrimination. In *Agu fall meeting abstracts* (Vol. 2017, p. S43A-0834).
- McGuire, R. K. (2008). Probabilistic seismic hazard analysis: Early history. *Earthquake Engineering & Structural Dynamics*, 37(3), 329–338.
- Miao, F., Carpenter, N. S., Wang, Z., Holcomb, A. S., & Woolery, E. W. (2020, Mar). High-accuracy discrimination of blasts and earthquakes using neural networks with multiwindow spectral data. *Seismological Research Letters*, 91(3), 1646–1659. Retrieved from https://www.researchgate.net/publication/339856492_High-Accuracy_Discrimination_of_Blasts_and_Earthquakes_Using_Neural_Networks_With_Multiwindow_Spectral_Data doi: 10.1785/0220190084
- Mousavi, S. M., Zhu, W., Sheng, Y., & Beroza, G. C. (2019, Jul). Cred: A deep residual network of convolutional and recurrent units for earthquake signal detection. *Scientific Reports*, 9(1). Retrieved from <https://www.nature.com/articles/s41598-019-45748-1> doi: 10.1038/s41598-019-45748-1

- Nairobi. (2005, May). Natural disasters —a heavy price to pay. *The New Humanitarian*. Retrieved from <https://www.thenewhumanitarian.org/fr/node/222212>
- Pezzo, E. D., Esposito, A., Marinaro, M., Martini, M., & Scarpetta, S. (2003, Feb). Discrimination of earthquakes and underwater explosions using neural networks. *Bulletin of the Seismological Society of America*, 93(1), 215–223. doi: 10.1785/0120020005
- Pifer, S. (2016, Sep). What’s the deal with senate republicans and the test ban treaty? *Brookings*. Retrieved from <https://www.brookings.edu/blog/order-from-chaos/2016/09/26/whats-the-deal-with-senate-republicans-and-the-test-ban-treaty/>
- Rabin, N., Bregman, Y., Lindenbaum, O., Ben-Horin, Y., & Averbuch, A. (2016, Jun). Earthquake-explosion discrimination using diffusion maps. *Geophysical Journal International*, 207(3), 1484–1492. Retrieved from https://www.researchgate.net/publication/304176091_Earthquake-Explosion_Discrimination_Using_Diffusion_Maps doi: 10.1093/gji/ggw348
- Ranasinghe, N. R., Huang, L., Clee, T., & Kemp, J. A. (2019, December). A deep learning approach to discriminate between explosions and earthquakes. In *Agu fall meeting abstracts* (Vol. 2019, p. S43D-0674).
- Ray, J., Hansen, C., Forrest, R., & Young, C. J. (2017, Apr). *Using discrete wavelet transforms to discriminate between noise and phases in seismic waveforms*. Retrieved from <https://www.osti.gov/servlets/purl/1456582>
- RaynaArora. (2021, Aug). Dataset for teleseismic discrimination [dataset and jupyter notebooks]. *Zenodo*. Retrieved from <https://zenodo.org/record/5167966> doi: 10.5281/zenodo.5167966
- Ross, Z. E., Meier, M., Hauksson, E., & Heaton, T. H. (2018, Aug). Generalized seismic phase detection with deep learning. *Bulletin of the Seismological Society of America*, 108(5A), 2894–2901. Retrieved from <https://pubs.geoscienceworld.org/ssa/bssa/article-abstract/108/5A/2894/546740/Generalized-Seismic-Phase-Detection-with-Deep> doi: 10.1785/0120180080
- Storchak, D. A., Harris, J., Brown, L., Lieser, K., Shumba, B., & Giacomo, D. D. (2020, Nov). Rebuild of the bulletin of the international seismological centre (ISC)—part 2: 1980–2010. *Geoscience Letters*, 7. doi: 10.1186/s40562-020-00164-6
- Storchak, D. A., Harris, J., Brown, L., Lieser, K., Shumba, B., Verney, R., . . . Korger, E. I. M. (2017, Dec). Rebuild of the bulletin of the international seismological centre (ISC), part 1: 1964–1979. *Geoscience Letters*, 4. doi: 10.1186/s40562-017-0098-z
- Stump, B. W., Hedlin, M. A., Pearson, D. C., & Hsu, V. (2002). Characterization of mining explosions at regional distances: Implications with the international monitoring system. *Reviews of Geophysics*, 40(4), 2–1.
- UN. (1996). *Comprehensive Nuclear-Test-Ban Treaty (CTBT)*. <https://www.ctbto.org/the-treaty/treaty-text/>. (Accessed: 2021-5-10)
- Wei, H., Shu, W., Dong, L., Huang, Z., & Sun, D. (2020, Aug). A waveform image method for discriminating micro-seismic events and blasts in underground mines. *Sensors*, 20(15), 4322. Retrieved from <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7436190/> doi: 10.3390/s20154322