

DeepQuake – An application of CNN for seismo-acoustic event classification in The Netherlands

Luca Trani¹, Giuliano Andrea Pagani¹, João Paulo Pereira Zanetti¹, Camille Chapeland^{1,2} and Láslo Evers^{1,2}

¹Royal Netherlands Meteorological Institute (KNMI)
²Delft University of Technology (TU Delft)

Key Points:

- Assessing feasibility of DL-based detection of low magnitude induced seismo-acoustic event
- Evaluating performance in comparison with an operational seismic detection system
- Introducing a visualisation mechanism for inspection of CNN intermediate results

Abstract

Recent developments of infrastructures and methods are major driving forces in the advances of solid Earth sciences. The deployment of large and dense sensor networks enables data centres to acquire data of increased volume and quality. The analysis of such data provides scientists with a better understanding about natural phenomena in the subsurface. Nevertheless new challenges arise to exploit the growing information potential. Innovative methods based on Artificial Intelligence offer concrete opportunities to tackle those challenges. In this paper we present an investigation of Convolutional Neural Networks (CNN) for seismo-acoustic event classification in the Netherlands. We designed, trained and evaluated two CNN models. Our results suggest that as CNN inputs spectrograms are more suitable than continuous waveforms. We discuss our findings' potential and requirements for their operational adoption. We focus on explainability aspects and offer an approach to pave the way for a broader uptake of Artificial Intelligence based methods.

Plain Language Summary

Seismic monitoring services have the critical mission to detect, locate and analyse ground motion recorded by networks of seismometers deployed around the globe. Such motion can be derived from waveforms triggered by natural phenomena such as earthquakes as well as by human-related activities (e.g., explosions, drilling works). Automated systems offer a fundamental aid to detect seismic events, however, they are often focused on one specific type of event and require intensive human validation activities. Here, we present an approach for building an operational intelligent system that helps operators by automatically discriminating seismic waveforms into three categories: noise, earthquake or other event type.

1 Introduction

Seismo-acoustic waves are regularly observed from earthquakes (Shani-Kadmiel et al., 2018), mining blasts (Evers et al., 2012), demolition of old ordnance (Ruigrok et al., 2019), nuclear tests (Assink et al., 2016) and (underwater) volcanoes (Green et al., 2013). Such natural and human-made sources, often have a signature in more than one medium, *i.e.*, seismic-waves in the solid Earth can couple to the oceans and/or atmosphere and generate acoustic waves and vice versa. The analysis of the seismo-acoustic wavefield, as it is captured by the various sensors, provides unprecedented insight into the source characteristics and the medium through which the waves have propagated.

Source detection, identification and characterisation in terms of the type of event, its location, origin time and size is the ultimate challenge. At present the characterisation of seismo-acoustic sources and the discrimination of earthquakes from other type of events is predominantly done by human analysts trained to cross-link diverse types of information. With the evolution and growing density of geophysical monitoring networks, the number of events that need to be analysed rapidly increases, making it impractical to fully rely on human analysts. Thus, an automated system that can correctly classify events as a natural or an induced earthquake or as other, is desirable. Such a system, could for instance, improve the critical decision-making process of an earthquake Early Warning service. This is crucial in order to timely and correctly inform society about an unraveling geophysical hazard.

In recent years Artificial Intelligence (AI) and methods such as Machine Learning (ML) and Deep Learning (DL) have become increasingly popular. Their successful application in several geophysical contexts has demonstrated great potential. For instance, in the solid Earth domain those methods have been applied to classify volcano deformations from InSAR images (Anantrasirichai et al., 2018), to detect, locate and characterise earthquakes (Perol et al., 2018; Lomax et al., 2019; Kriegerowski et al., 2019; Zhang et al., 2020), to pick and associate phase arrivals (Ross et al., 2018; Dokht et al., 2019; Ross et al., 2019; Wang et al., 2019) and to help analysts discriminate between different types of seismic events (Linville et al., 2019).

In this paper we present an investigation of Convolutional Neural Networks (CNN) for seismo-acoustic event classification in the Netherlands. We report the results obtained in the DeepQuake project that focused on the classification of continuous seismic waveforms recorded in the Groningen region into three categories: *earthquake*, *noise* and *other seismo-acoustic event-type* (e.g. ordnance explosion and sonic booms).

The Royal Netherlands Meteorological Institute (KNMI) operates an extensive seismo-acoustic network in order to monitor and assess anthropogenic hazards induced by the extraction of natural gas from the Groningen reservoir (KNMI, 1993). The seismicity of that region is characterised by shallow events of low to moderate magnitude (typically, depth around 3km and magnitude ≤ 3.6) (Dost et al., 2017). The automated detection of such microseismic events is an inherently challenging task, this is especially true in the Groningen gas field due to the peculiar site characteristics and noise conditions (Spica et al., 2018). Combining CNN and validation techniques, we devised a novel approach to tackle those challenges, thus laying the foundations for an explainable and trustworthy automated seismo-acoustic classification system.

2 Data collection and preprocessing

The effectiveness of CNN and the reliability of their predictions strongly depend on the quality and amount of available training data. We build our dataset using 1172 events (including earthquakes and other types of events, e.g. explosions and sonic booms) catalogued in the Netherlands between 2014 and 2018 (Aardbevingscatalogus, 2016). These events have been recorded by the 599 stations of the Dutch network (KNMI, 1993) presented in Figure 1.

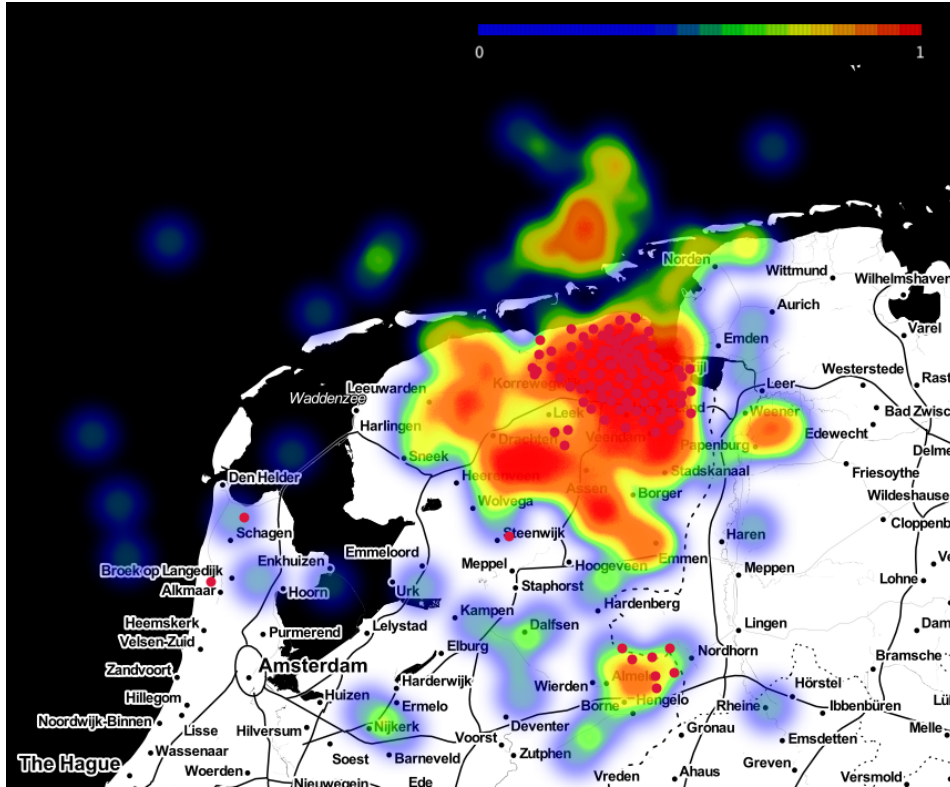


Figure 1: Overview map showing the distribution of stations (red dots) and normalised number of events of our dataset.

For each event in the catalogue, a 20-second waveform stream S is extracted from all stations within a certain radius depending on the event magnitude. For events with magnitude $M \leq 1.5$, the maximum distance considered is 10 km; when $1.5 < M \leq 2.5$, events up to 25 km are considered; and larger events have a maximum distance of 50 km.

To inhibit the neural network from “expecting” a signal at a specific time in the stream, we first estimate the direct P-wave arrival time t_P , using a local velocity model (see supp. mat.). We then select our 20-second window between t_P^- , randomly chosen between 1.5 and 8.5 seconds before t_P , and $t_P^+ = t_P^- + 20$ seconds. For each selected S , the signal-to-noise ratio (SNR) is calculated as

$$SNR(S) = \frac{\max(|S[t_P - 0.5 : t_P + 6.5]|)}{\max(|S[t_P^- : t_P - 0.5]|)}.$$

Only streams with $SNR \geq 4$ are used in the training set.

Noise samples are selected by extracting streams between events, with at least a 15 minutes difference from any event. They are not tested for SNR. This ensures that all types of noise, including high-amplitude irregular noise, are included for training within the noise dataset.

After the streams are collected, the following pre-processing steps are applied:

1. rotation into vertical, North, and East components (Z-N-E)
2. resampling to 100Hz (using Fourier method)
3. linear detrending
4. a bandpass filter between 0.5Hz and 22Hz (to avoid 50Hz electronic distortion)
5. normalisation by the absolute maximum value in the stream

Additionally, the short-time Fourier transform (STFT) of each stream is computed using a basic Hanning window (with length 120 and overlap 60). Both the pre-processing and the frequency computation processes are equally applied to either event and noise streams. The process from raw waveforms to the input samples for the CNN model is summarised in Figure 2.

At the end of this process, each sample has three one-dimensional channels of length 2000, corresponding to amplitude, and three two-dimensional channels of size 61x35, corresponding to frequency data. The final dataset contains 83863 samples: 41931 for events and 41932 for noise. They are split into *training*, *validation* and *test* sets at a ratio of approximately 8:1:1. To minimise biases each set is created with streams from distinct days. This enables a balanced time spreading and avoids to use the same event recorded by different stations for training and validation.

3 Methods

The approach adopted and described in this paper builds on recent advances in seismology (Perol et al., 2018; Lomax et al., 2019; Kriegerowski et al., 2019). We leverage machine learning and CNN to perform supervised classification of continuous seismic waveforms. CNN are composed by several convolutional layers and their corresponding filters that learn different representations of the input data by activating on features of increasing complexity. The presence of several layers characterises CNN as Deep Architectures (Bengio, 2009). Therefore, they are also associated with the concept of Deep Learning (LeCun et al., 2015).

Deep Architectures, such as LeCun’s CNN constituted by convolutional layers and sub-sampling layers, are particularly effective for recognition and detection tasks (LeCun & Bengio, 1995; LeCun et al., 1999). They exploit inherent *compositional hierarchies* present in many signals whereby high-level features can be derived by composing low-level ones (LeCun et al., 2015). Compared to fully connected neural networks they enable a more efficient use of resources (e.g. computation, memory), thus making them suitable for operational, real-time systems and Early Warning applications (Li et al., 2018).

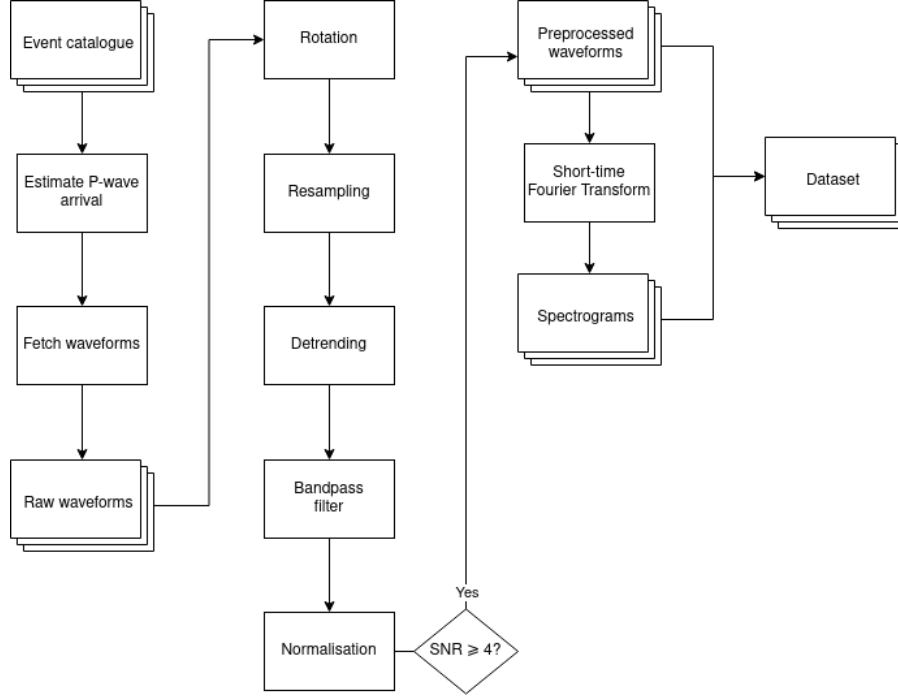



Figure 2: Pre-processing steps for the preparation of the training, validation and test datasets.

Our architecture (depicted in Figure 3) builds on those features – inspired by CNN’s ability to recognise objects in multidimensional arrays (*e.g.* 2D images) (Krizhevsky et al., 2012; Hemanth & Estrela, 2017), we apply them to time-series data (*i.e.* continuous seismic waveforms) in order to recognise and learn seismo-acoustic features.

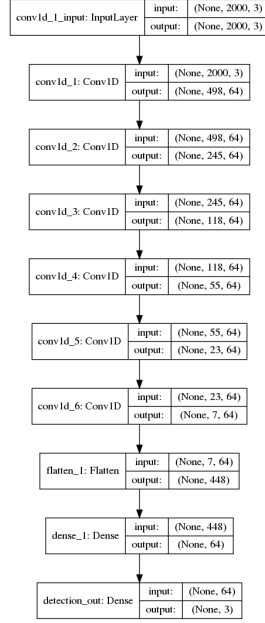
Figure 3a illustrates the steps and operations involved in the classification of a seismic waveform into three categories: *earthquake*, *noise* and *other event-type*. In a first phase, *Feature Learning*, relevant elements, known as feature maps, are extracted from an input data sample (*i.e.* a seismic waveform) by applying a convolution operation combined with a Rectified Linear Unit (ReLU) activation function (Goodfellow et al., 2016). The process is iterated on the resulting feature maps – they are inputs to a next layer of convolution where new filters are applied to generate new feature maps. At each step the map dimensionality is reduced by exploiting suitable stride values, a similar result can be obtained by using pooling operations (*e.g.* average, max).

In a second phase, *Classification*, the feature maps produced by the last convolution are flattened into a vector which is provided as input to a fully connected layer of neurons. That layer is followed by a final one containing the three output categories; scores for each category are computed by a SoftMax activation function (Goodfellow et al., 2016).

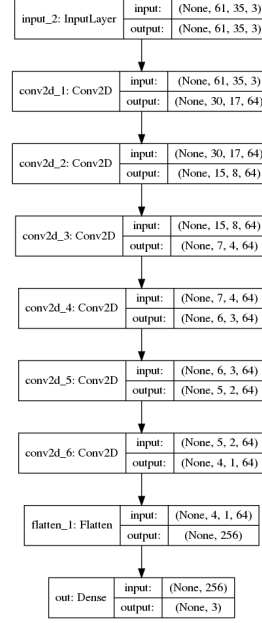
The design of an optimal CNN architecture is usually a long and complex task. It involves finding and tuning suitable configurations, an operation also known as hyper-parameter search (Donoghue & Roantree, 2015). Such a process leverages variables like: type and amount of layers, amount of neurons in each layer, regularisation parameters, optimisation of learning rate, stride values *etc.* Methodologies exist to perform random searches in a hyper-parameter space (Bergstra & Bengio, 2012). However, heuristic approaches are often preferred for pragmatic reasons. In our case we performed a manual search to get insight of the effects of different configurations on the detection accuracy. This yielded the two architectures as shown in Figure 3: *arch-time* with time-series as input (Figure 3b); *arch-spect* with spectrograms as input (Figure 3c).



(a) Schematic view of a CNN architecture for seismo-acoustic event classification.



(b) Architecture *arch-time*. The input layer contains 2000 values (20 seconds of signal sampled at 100Hz) and 3 channels corresponding to the seismometer orientations (Vertical, North-South, East-West). The input layer is followed by 6 convolutional layers with 64 filters. Data reduction is achieved by using a combination of filter size and its displacement (*i.e.* stride). After the last convolution layer a fully connected layer and a SoftMax function are applied to obtain an output score for each of the 3 detection categories.



(c) Architecture *arch-spect*. The input contains 3 channels represented with a 61x35 matrix followed by 6 layers of convolutions with 64 filters. The data reduction is performed in a similar way as for *arch-time* by using stride. Once the convolutions are performed on the inputs a fully connected layer is applied.

Figure 3: CNN architectures.

After selecting a specific architecture, a CNN needs to be trained in order to be usable. The process of training a CNN consists in finding the appropriate set of weights (*i.e.* values for the filters) that minimise a chosen loss function – it is an iterative process that requires several *epochs*. Typically the minimisation is realised by applying a form of gradient descent to the loss function (Ruder, 2016). The training process stops either after a number of predetermined epochs or when a certain tolerance value has been obtained in the reduction of the loss function between consecutive epochs.

Our solution is implemented in Python (van Rossum, 1995), using ObsPy (Beyreuther et al., 2010) for seismic processing, numerical facilities from NumPy and SciPy (Virtanen et al., 2020), and the TensorFlow (Abadi et al., 2015) platform for machine learning.

4 Results and Validation

We evaluated the performance of our approach by combining automated statistics, validation and benchmarking techniques.

Figure 4 illustrates the classification results on the test set by adopting a confusion matrix representation (Sammut & Webb, 2011). We can notice that the properly classified categories emerge on the other diagonal whereas the miss-classifications are scattered around it.

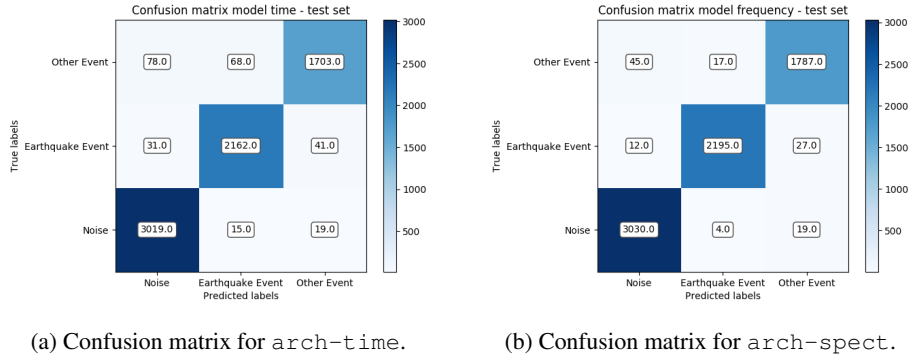


Figure 4: Confusion matrix representation of classification results on the test set.

Overall *arch-time* (Figure 4a) achieves an accuracy of 0.965 while *arch-spect* (Figure 4b) reaches an accuracy of 0.983. Noteworthy are the improvements in the classification of the “Other Event” class and the “Earthquake Event” class. To evaluate the effectiveness of our CNN we run predictions on continuous waveform data recorded in 2019. Figure 5 shows the classification results for a few earthquakes of diverse intensity. It also includes a comparison with the current operational seismic detection system adopted at KNMI *i.e.* Seis-Comp3 (SC3) (Weber et al., 2007).

We can notice that with lower magnitudes the results are less prominent, conceivably such effects are due to the noise conditions. However, our system seems to be less sensitive to daily noise variations. We repeated such an experiment with data recorded in the lockdown due to the COVID-19 pandemic – a unique period characterised by a substantial ambient noise reduction (Lecocq et al., 2020). Example results are illustrated in Figure 5g – despite a slight reduction of false positives the detection of low magnitude events remains a challenging task.

5 Discussion

The results achieved suggest that an operational system based on our CNN could outperform current automated detection systems (*e.g.* SC3) by improving the detection accuracy (*e.g.* less false positive) and by extending the range of application to other types of events beyond earthquakes. We notice that the accuracy performance improves when using *arch-spect*. We attribute this behaviour to (i) an intrinsic difference of the frequency content of the three classification categories; and (ii) the time-frequency representation that, for similarity with image data, is more suitable for CNN.

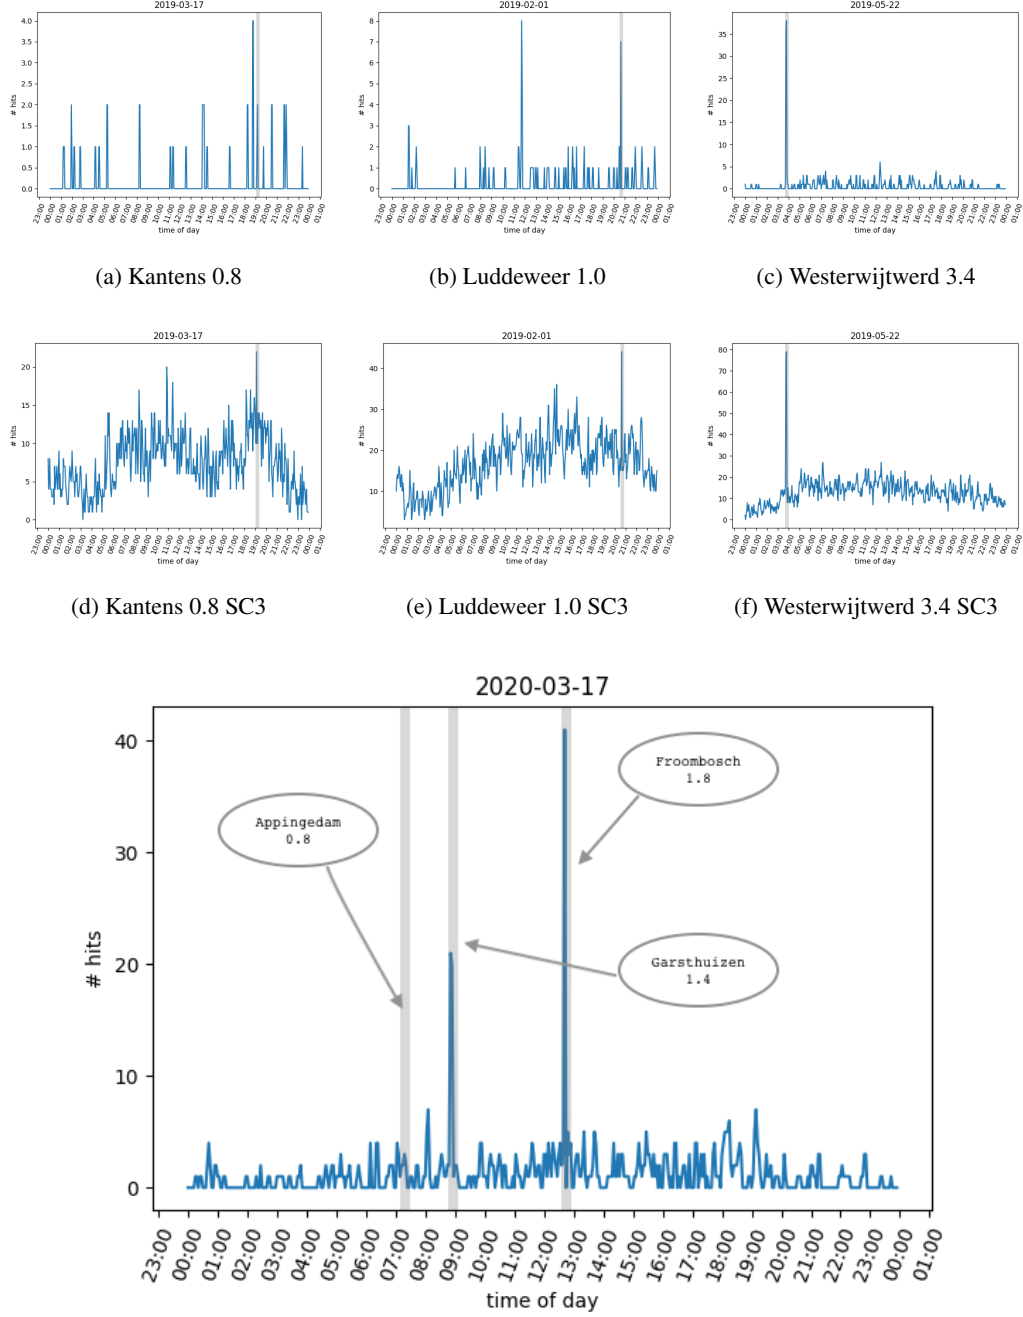


Figure 5: Results of predictions obtained with *arch-spect*. Recordings of 90 stations during 3 days containing catalogued events (highlighted in grey) are analysed in (a, b, c). A comparison with SC3 is done in (d, e, f) by counting the number of picks detected in the same days. An analysis of a lockdown day is illustrated in (g)

The higher accuracy of *arch-spect* comes with additional pre-processing steps for the computation of the STFT. Therefore, the final choice of the trade-off between accuracy and response time depends on the application requirements.

Moving a DL-based tool to production requires building understanding, trust and confidence in such methods and their results. In other words they need to be made “explainable” (Adadi & Berrada, 2018) to seismologists, operators and analysts. DL is often perceived as a black box for the lack of transparency of its processes and decision logic. As a first attempt to better understand the features learned by the trained CNN model, we devised a visual tool that presents the activations of each filter in the convolutional layers integrated with the original input signal – an example of such a representation is illustrated in Figure 6.

When using the trained model, we can notice how the heatmap for a noise signal appears randomly scattered (Figure 6a) whereas in the case of an earthquake (Figure 6b) or an acoustic event (Figure 6c) it appears more focused on specific features *e.g.* amplitude variations. Such a focus is even more evident in the time-frequency domain. In the case of a noise sample (Figure 6d) the activation focuses on a low range of frequencies and it is spread across the whole time window. In the cases of seismic and acoustic events (Figure 6e and Figure 6f) there are clear frequency patterns that the network is able to recognise and activate upon. Such a representation confirms that the CNN model is extracting relevant seismological features. When further developed, this tool could be used for dynamic inspection and to support fine tuning of the CNN. Also, it could be adopted to relate predictions with expert knowledge.

6 Conclusions and future work

We have presented a DL approach for the classification of seismo-acoustic events and showed potential advantages with respect to current detection systems. We have reported the results obtained with two CNN architectures, `arch-time` working with seismic waveforms and `arch-spect` working with spectrograms; with the latter achieving a higher accuracy. Low magnitude events remain a challenge, however, the COVID-19 lockdown period showed us that low ambient noise conditions result in fewer false alarms.

We validated the results and provided a visual inspection mechanism based on the activation of CNN layers to highlight the data features captured by the network. Such a tool would enable domain experts to gain trust in DL-based methods.

Our future work will focus on improving seismo-acoustic source discrimination and characterisation *e.g.* magnitude, location. For that goal we are investigating multi-model solutions, for instance by combining our CNN with Graph (GNN) and Recurrent Neural Networks (RNN). Also, we are considering to enhance our dataset with synthetic samples *e.g.* by using Generative Adversarial Networks (GANs). Furthermore, we will continue working on explainability aspects towards the establishment of a trusted, reliable and reproducible operational tool.

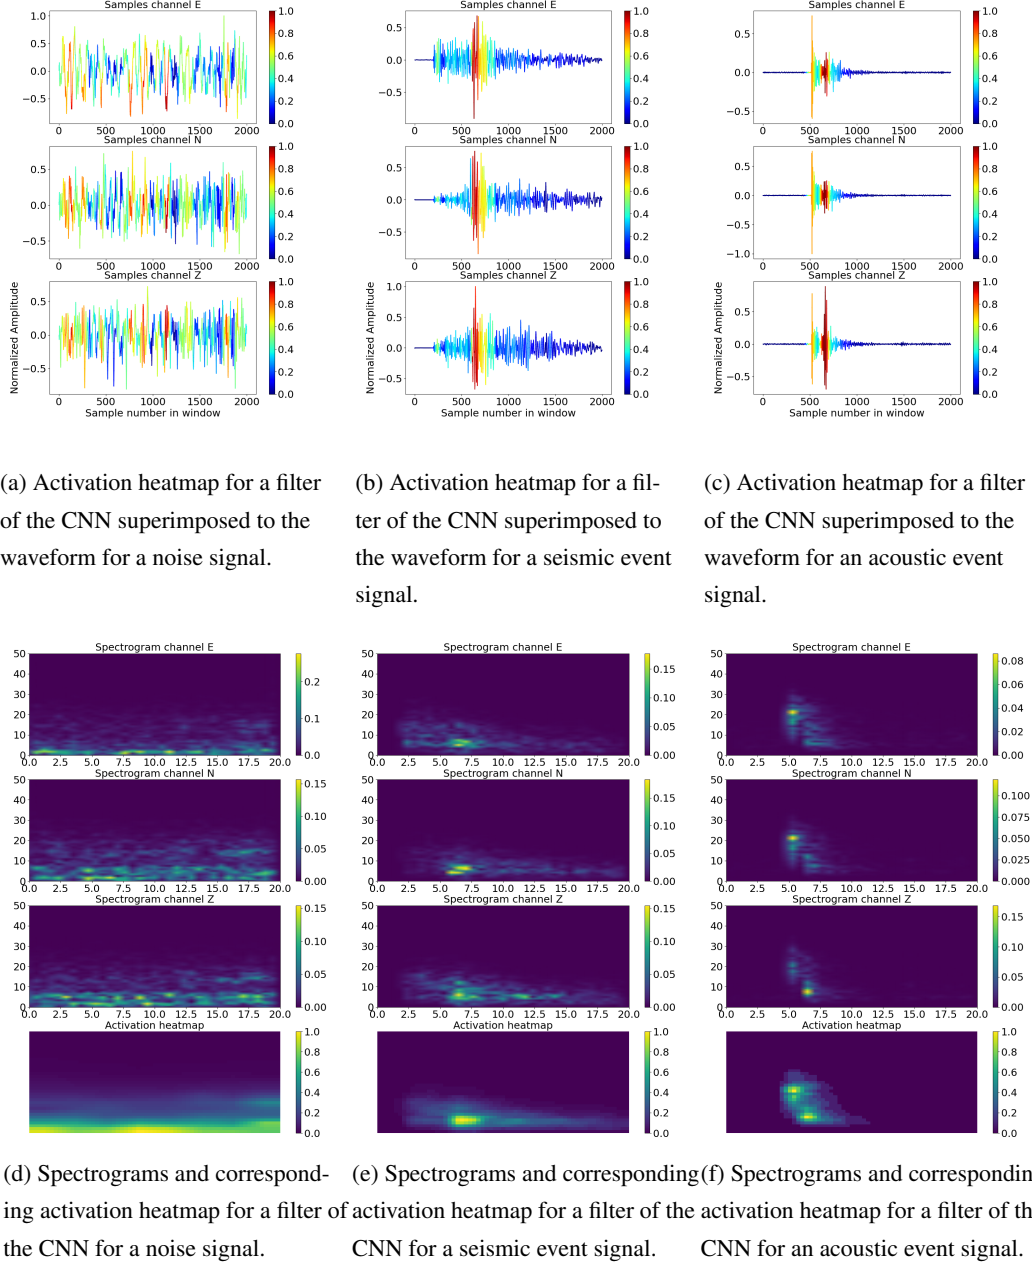


Figure 6: Example of a filter's activation relative heatmap superimposed to the three categories of waveform in time domain (a,b,c) and spectrogram with filter heatmap (d,e,f). Other filters might focus on different features.

Acknowledgments

Data and software utilised in this manuscript will be made publicly available in a FAIR compliant repository (e.g., Zenodo). The authors would like to thank Shahar Shani-Kadmiel who provided valuable comments and suggestions to improve this manuscript. This work is co-funded by the EOSC-hub project (Horizon 2020) under Grant number 777536. Láslo Evers' contribution is funded through a VIDI project from the Dutch Research Council (NWO), project number 864.14.005. This work was carried out on the Dutch national e-infrastructure with the support of SURF Cooperative.

References

- Aardbevingscatalogus. (2016). KNMI. Retrieved from <https://www.knmi.nl/kennis-en-datacentrum/dataset/aardbevingscatalogus>
- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., ... Zheng, X. (2015). *TensorFlow: Large-scale machine learning on heterogeneous systems*. Retrieved from <https://www.tensorflow.org/> (Software available from tensorflow.org)
- Adadi, A., & Berrada, M. (2018). Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI). *IEEE Access*, 6, 52138–52160. doi: 10.1109/ACCESS.2018.2870052
- Anantrasirichai, N., Biggs, J., Albino, F., Hill, P., & Bull, D. (2018). Application of Machine Learning to Classification of Volcanic Deformation in Routinely Generated InSAR Data. *Journal of Geophysical Research: Solid Earth*, 123(8), 6592–6606. doi: 10.1029/2018JB015911
- Assink, J. D., Averbuch, G., Smets, P. S. M., & Evers, L. G. (2016). On the infrasound detected from the 2013 and 2016 dprk's underground nuclear tests. *Geophysical Research Letters*, 43(7), 3526–3533. Retrieved from <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2016GL068497> doi: 10.1002/2016GL068497
- Bengio, Y. (2009). Learning Deep Architectures for AI. *Foundations and Trends in Machine Learning*, 2(1), 1–127. doi: 10.1561/22000000006
- Bergstra, J., & Bengio, Y. (2012). Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13(10), 281–305. Retrieved from <http://jmlr.org/papers/v13/bergstra12a.html>
- Beyreuther, M., Barsch, R., Krischer, L., Megies, T., Behr, Y., & Wassermann, J. (2010, 05). ObsPy: A Python Toolbox for Seismology. *Seismological Research Letters*, 81(3), 530–533. Retrieved from <https://doi.org/10.1785/gssrl.81.3.530> doi: 10.1785/gssrl.81.3.530
- Dokht, R. M., Kao, H., Visser, R., & Smith, B. (2019). Seismic event and phase detection using time-frequency representation and convolutional neural networks. *Seismological Research Letters*, 90(2 A), 481–490. doi: 10.1785/0220180308
- Donoghue, J. O., & Roantree, M. (2015). A framework for selecting deep learning hyper-parameters. In S. Maneth (Ed.), *Data science* (pp. 120–132). Cham: Springer International Publishing.
- Dost, B., Ruigrok, E., & Spetzler, J. (2017). Development of seismicity and probabilistic hazard assessment for the groningen gas field. *Netherlands Journal of Geosciences*, 96(5), s235–s245. doi: 10.1017/njg.2017.20
- Evers, L. G., van Geyt, A. R. J., Smets, P., & Fricke, J. T. (2012). Anomalous infrasound propagation in a hot stratosphere and the existence of extremely small shadow zones. *Journal of Geophysical Research: Atmospheres*, 117(D6). Retrieved from <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2011JD017014> doi: 10.1029/2011JD017014
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- Green, D. N., Evers, L. G., Fee, D., Matoza, R. S., Snellen, M., Smets, P., & Simons, D. (2013). Hydroacoustic, infrasonic and seismic monitoring of the submarine eruptive activity and sub-aerial plume generation at south sarigan, may 2010. *Journal of Volcanology and Geothermal Research*, 257, 31 - 43. Retrieved

- from <http://www.sciencedirect.com/science/article/pii/S0377027313000735> doi: <https://doi.org/10.1016/j.jvolgeores.2013.03.006>
- Hemanth, D. J., & Estrela, V. V. (2017). *Deep learning for image processing applications* (Vol. 31). IOS Press.
- KNMI. (1993). *Netherlands seismic and acoustic network*. Royal Netherlands Meteorological Institute (KNMI). Retrieved from <https://doi.org/10.21944/e970fd34-23b9-3411-b366-e4f72877d2c5>
- Kriegerowski, M., Petersen, G. M., Vasyura-Bathke, H., & Ohrnberger, M. (2019). A deep convolutional neural network for localization of clustered earthquakes based on multistation full waveforms. *Seismological Research Letters*, 90(2 A), 510–516. doi: 10.1785/0220180320
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet Classification with Deep Convolutional Neural Networks. *Advances In Neural Information Processing Systems*, 1–9. doi: <http://dx.doi.org/10.1016/j.protcy.2014.09.007>
- Lecocq, T., Hicks, S. P., Van Noten, K., van Wijk, K., Koelemeijer, P., De Plaen, R. S. M., ... Xiao, H. (2020). Global quieting of high-frequency seismic noise due to covid-19 pandemic lockdown measures. *Science*. Retrieved from <https://science.sciencemag.org/content/early/2020/07/22/science.abd2438> doi: 10.1126/science.abd2438
- LeCun, Y., & Bengio, Y. (1995). Convolutional networks for images, speech, and time-series. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks*. MIT Press.
- LeCun, Y., Bengio, Y., Hinton, G., Y., L., Y., B., & G., H. (2015). Deep learning. *Nature*, 521(7553), 436–444. doi: 10.1038/nature14539
- LeCun, Y., Haffner, P., Bottou, L., & Bengio, Y. (1999). Object recognition with gradient-based learning. In D. Forsyth (Ed.), *Feature grouping*. Springer.
- Li, Z., Meier, M. A., Hauksson, E., Zhan, Z., & Andrews, J. (2018). Machine Learning Seismic Wave Discrimination: Application to Earthquake Early Warning. *Geophysical Research Letters*, 45(10), 4773–4779. doi: 10.1029/2018GL077870
- Linville, L., Pankow, K., & Draelos, T. (2019). Deep Learning Models Augment Analyst Decisions for Event Discrimination. *Geophysical Research Letters*, 46(7), 3643–3651. doi: 10.1029/2018GL081119
- Lomax, A., Michelini, A., & Jozinović, D. (2019). An Investigation of Rapid Earthquake Characterization Using Single Station Waveforms and a Convolutional Neural Network. *Seismological Research Letters*, 90(2A), 517–529. doi: 10.1785/0220180311
- Perol, T., Gharbi, M., & Denolle, M. (2018). Convolutional neural network for earthquake detection and location. *Science Advances*, 4(2). doi: 10.1126/sciadv.1700578
- Ross, Z. E., Meier, M. A., & Hauksson, E. (2018). P Wave Arrival Picking and First-Motion Polarity Determination With Deep Learning. *Journal of Geophysical Research: Solid Earth*, 123(6), 5120–5129. doi: 10.1029/2017JB015251
- Ross, Z. E., Yue, Y., Meier, M. A., Hauksson, E., & Heaton, T. H. (2019). PhaseLink: A Deep Learning Approach to Seismic Phase Association. *Journal of Geophysical Research: Solid Earth*, 124(1), 856–869. Retrieved from <https://doi.org/10.1029/2018JB016674> doi: 10.1029/2018JB016674
- Ruder, S. (2016). An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*.
- Ruigrok, E., Domingo-Ballesta, J., van den Hazel, G.-J., Dost, B., & Evers, L. (2019). Groningen explosion database [Journal Article]. *First Break*, 37(8), 37–41. Retrieved from <https://www.earthdoc.org/content/journals/10.3997/1365-2397.n0053> doi: <https://doi.org/10.3997/1365-2397.n0053>
- Sammut, C., & Webb, G. I. (2011). *Encyclopedia of machine learning*. Springer Science & Business Media.
- Shani-Kadmiel, S., Assink, J., Smets, P., & Evers, L. (2018, January 16). Seismoacoustic coupled signals from earthquakes in central Italy: Epicentral and secondary sources of infrasound. *Geophysical Research Letters*, 45(1), 427–435. doi: 10.1002/2017GL076125

- 333 Spica, Z. J., Nakata, N., Liu, X., Campman, X., Tang, Z., & Beroza, G. C. (2018, 05). The
 334 Ambient Seismic Field at Groningen Gas Field: An Overview from the Surface to
 335 Reservoir Depth. *Seismological Research Letters*, 89(4), 1450-1466. Retrieved from
 336 <https://doi.org/10.1785/0220170256> doi: 10.1785/0220170256
 337 van Rossum, G. (1995, May). *Python tutorial* (Tech. Rep. No. CS-R9526). Amsterdam:
 338 Centrum voor Wiskunde en Informatica (CWI).
 339 Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., ...
 340 Contributors, S. . . (2020). Scipy 1.0: fundamental algorithms for scientific computing
 341 in python. *Nature Methods*, 17(3), 261–272. Retrieved from [https://doi.org/](https://doi.org/10.1038/s41592-019-0686-2)
 342 [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2) doi: 10.1038/s41592-019-0686-2
 343 Wang, J., Xiao, Z., Liu, C., Zhao, D., & Yao, Z. (2019). Deep Learning for Picking Seis-
 344 mic Arrival Times. *Journal of Geophysical Research: Solid Earth*, 124(7), 6612–6624.
 345 doi: 10.1029/2019JB017536
 346 Weber, B., Becker, J., Hanka, W., Heinloo, A., Hoffmann, M., Kraft, T., ... Thoms, H.
 347 (2007). Seiscomp3—automatic and interactive real time data processing. In *Geophys-
 348 ical research abstracts* (Vol. 9).
 349 Zhang, X., Zhang, J., Yuan, C., Liu, S., Chen, Z., & Li, W. (2020). Locating induced earth-
 350 quakes with a network of seismic stations in Oklahoma via a deep learning method.
 351 *Scientific Reports*, 10(1), 1–12. Retrieved from [http://dx.doi.org/10.1038/](http://dx.doi.org/10.1038/s41598-020-58908-5)
 352 [s41598-020-58908-5](http://dx.doi.org/10.1038/s41598-020-58908-5) doi: 10.1038/s41598-020-58908-5