

# PolarCAP – A Deep Learning approach for First Motion Polarity Classification of Earthquake Waveforms

Megha Chakraborty<sup>1,2</sup>, Claudia Quinteros Cartaya<sup>1</sup>, Wei Li<sup>1</sup>, Johannes Faber<sup>1,3</sup>, Horst Stoecker<sup>1,3,4,5</sup>, Georg Rumpker<sup>1,2</sup>, Nishtha Srivastava<sup>1,2</sup>

<sup>1</sup>Frankfurt Institute for Advanced Studies, 60438 Frankfurt am Main, Germany

<sup>2</sup>Institute of Geosciences, Goethe-University Frankfurt, 60438 Frankfurt am Main, Germany

<sup>3</sup>Institute for Theoretical Physics, Goethe Universität, 60438 Frankfurt am Main, Germany

<sup>4</sup>Xidian-FIAS international Joint Research Center, Giersch Science Center, 60438 Frankfurt am Main, Germany

<sup>5</sup>GSI Helmholtzzentrum für Schwerionenforschung GmbH, 64291 Darmstadt, Germany

## Key Points:

- We present PolarCAP, a deep learning model that can classify the polarity of a waveform with known P-arrival time with a 98% accuracy.
- The first-motion polarity of seismograms is a useful information, but its manual determination can be laborious and imprecise.
- We demonstrate that in several cases the model can assign trace polarity more accurately than a human analyst.

## Abstract

The polarity of first P-wave arrivals plays a significant role in the effective determination of focal mechanisms specially for smaller earthquakes. Manual estimation of polarities is not only time-consuming but also prone to human errors. This warrants a need for an automated algorithm for first motion polarity determination. We present a deep learning model - PolarCAP that uses an autoencoder architecture to identify first-motion polarities of earthquake waveforms. PolarCAP is trained in a supervised fashion using more than 130,000 labelled traces from the Italian seismic dataset (INSTANCE) and is cross-validated on  $\sim 22,000$  traces to choose the most optimal set of hyperparameters. We obtain an accuracy of 0.98 on a completely unseen test dataset of almost 33,000 traces. Furthermore, We check the model generalizability by testing it on the datasets provided by previous works and show that our model achieves a higher recall on both positive and negative polarities.

## Plain Language Summary

The polarity of an earthquake waveform describes the vertical direction of ground motion when the seismic waves first arrive at a station. This information is used to determine the mechanism of the faulting responsible for the earthquake. Traditionally polarity is determined manually by looking at the waveforms. However this is not just time consuming but often suffers from human errors. Hence there is a need for an algorithm that can perform this task automatically and reliably. In this paper we present a deep learning model PolarCAP (Polarity determining Convolutional Auto-encoder for first motion P-waves) which leverages the manually assigned polarity information of over 130,000 earthquake traces provided by the Italian seismic dataset (INSTANCE) for machine learning to learn the polarity estimation of any seismogram with an accuracy above 98%. We also demonstrate that the predictions from the model are more reliable than human-assigned labels.

## 1 Introduction

The first motion polarity of earthquake waveforms is an important parameter in determining focal mechanisms, particularly for smaller earthquakes. Traditionally the first-motion polarity is assigned manually by expert analysts. However, based on their observation on a dataset from the Northridge region, Hardebeck and Shearer (2002) reported that the picked polarities are inconsistent with the true polarities about 10% (for impulsive onset where polarities are more easily determined) to 20% (for emergent onset where polarities are more ambiguous) of the times. This, coupled with the growing volumes of seismological data, warrants the need for a faster, more precise and efficient method for the picking of polarities.

An automated polarity picking algorithm proposed by C. Chen and Holland (2016) is based on comparing the signal amplitude with the background noise and checking whether it crosses a user-defined threshold. Pugh et al. (2016) presented a Bayesian inference approach to polarity determination. Such numerical approaches, however, require (i) intensive human involvement, (ii) are heavily dependent on a limited number of parameters, and (iii) fail to account for the complex nature of seismograms; and hence cannot compete with manual picks (Ross et al., 2018).

Data-driven computer vision techniques, such as convolutional neural networks have been shown to be capable of analysing spatially independent information by mimicking the perception of images by the human brain (Voulodimos et al., 2018; Lundervold & Lundervold, 2019; Brachmann et al., 2017). Like in most research fields, deep learning has been successfully applied to seismology for tasks such as event detection & location (Perol et al., 2018), seismic phase identification & picking (Y. Chen, 2018; Zhu & Beroza, 2019; Li et al., 2021, 2022), magnitude characterization (Mousavi & Beroza, 2020; Chakraborty et al., 2022, 2021, 2004). The applicability of simple Convolutional Neural Networks (CNNs) in

67 the picking of first-motion polarities has been demonstrated by Ross et al. (2018), Hara et  
68 al. (2019) and Uchide (2020).

69 In this study we use an autoencoder model for determining the first-motion polarities.  
70 Unlike Mousavi et al. (2019) who also use autoencoders for polarity classification in an  
71 unsupervised fashion, we adopt a supervised approach and leverage the polarity information  
72 provided in the metadata of the INSTANCE dataset (Michelini et al., 2021). Since there are  
73 several algorithms that are capable of picking P-arrival times with an accuracy of the order  
74 of 0.01s (Mousavi et al., 2020; Li et al., 2021; Zhu & Beroza, 2019; Liao et al., 2021) we  
75 focus solely on classification of polarities and not the picking of P-arrival times. Therefore,  
76 we assume the P-arrival time to be an a-priori knowledge and use data windows of fixed  
77 length centred around the known P-arrival sample. We perform extensive analysis on our  
78 results to investigate its potential at outperforming human analysts and to investigate the  
79 possible scenarios that can lead to an error in assignment of polarity by the model.

## 80 2 Methodology

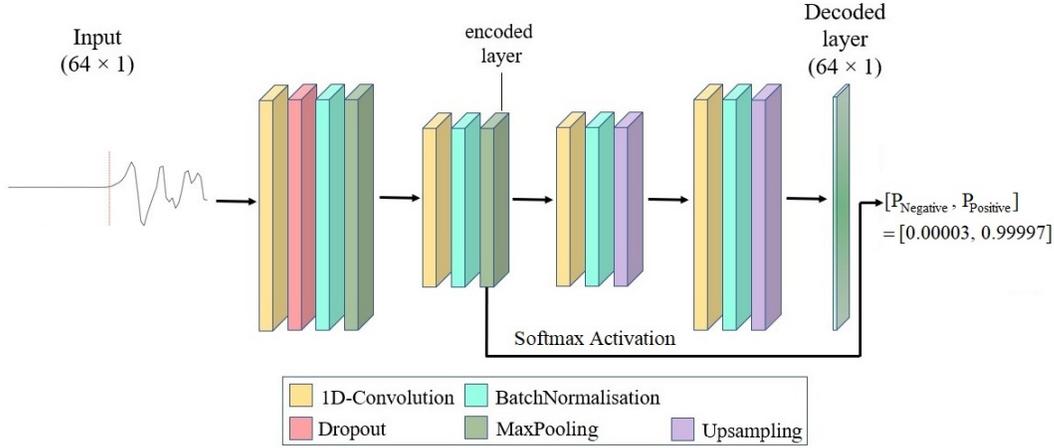
### 81 2.1 Data

82 Nearly 1.2 million waveforms recorded primarily by the Italian National Seismic Net-  
83 work between January 2005 and January 2020 and obtained from the INSTANCE dataset  
84 (Michelini et al., 2021) are used to train and validate our model. For simplicity we leave out  
85 those traces for which the number of detected P-wave arrivals, as reported in the metadata,  
86 is not equal to 1. It has been observed by Ross et al. (2018) that the accuracy in the predic-  
87 tion of polarity falls exponentially with signal-to-noise ratio (SNR), when the SNR is below  
88 10dB. Thus, we ignore traces with SNR less than 10dB to ensure minimum discrepancy in  
89 the training data. This leaves us with 443,002 traces out of which 109,748 have polarity in-  
90 formation identified by expert analysts and made available through catalogs. These 109,748  
91 traces are divided in the ratio 60:10:30 for training, validation and testing purposes.

92 Each trace in the training and validation sets is included twice in the respective set –  
93 once in its original form and once by flipping it (i.e. multiplying it by -1). This augmen-  
94 tation technique, previously used by Hara et al. (2019), not only doubles the volume of training  
95 data but also helps in balancing out the two classes (‘positive’ and ‘negative’ polarity), which  
96 is known to benefit the learning of a classifier (Batista et al., 2004). No such augmentation  
97 is applied to the test dataset to retain the original distribution of classes. We do not apply  
98 any pre-processing steps, other than normalising each waveform by dividing it with its  
99 maximum absolute value. As mentioned earlier, we work with the assumption that the first  
100 P-arrival time is already known to us. We use a fixed data window centred around the  
101 P-arrival sample. The length of the window was chosen experimentally to be 64 samples  
102 after exploring several values between 32 to 1024 samples. Note that only powers of 2 were  
103 used since the auto-encoder reduces the data dimensionality by a factor of 2 at each step.

### 104 2.2 Model architecture and training

105 We use an autoencoder model (Rumelhart et al., 1986) whose architecture is shown in  
106 Figure 1. It uses two sets of 1D Convolution (Kiranyaz et al., 2015) and Maxpooling (Nagi  
107 et al., 2011) layers to map the data into a 16-dimensional latent space (encoded layer). The  
108 decoder for reconstructing the data consists of two sets of Convolutional and Upsampling  
109 layers. Further details on the hyperparameters used can be found on the caption for Figure  
110 1. A softmax function is applied to the encoded layer to perform the classification. The  
111 model is implemented using Keras (Charles, 2013) and trained and tested on an NVIDIA  
112 A100 GPU. An Adam Optimiser (Kingma & Ba, 2014) is used for backpropagation. The  
113 loss function is a weighted sum of the reconstruction and classification losses with weights  
114 of 1 and 200 respectively. Since, we are more interested in the classification performance  
115 than the reconstruction performance, and the reconstruction is only used to facilitate the



**Figure 1.** The autoencoder architecture used for our study. The 1D convolutional layers in the encoder use ‘relu’ activation and have 32 and 8 filters respectively and kernel sizes of 32 and 16, respectively. Each Maxpooling layer reduces the data dimension by 2. The drop out rate used is 0.3. The convolutional layers in the decoder use ‘tanh’ and ‘relu’ activations respectively and 8 and 32 filters and kernel sizes of 16 and 32, respectively. The final decoder layer has a ‘tanh’ activation. A softmax function is applied on the encoded layer to get the probability of the waveform having a ‘negative’ and a ‘positive’ polarity respectively.

116 learning, higher weight is assigned to classification loss. To calculate the reconstruction and  
 117 classification losses we use the Keras (Charles, 2013) inbuilt loss functions for mean squared  
 118 error and huber loss with  $\delta = 0.5$ , respectively. We use early stopping (Prechelt, 2012)  
 119 to prevent overfitting, whereby the training stops automatically if the validation loss does  
 120 not decrease for 15 consecutive epochs, and the best set of learned model weights (one with  
 121 lowest validation loss) is saved iteratively. We also use *ReduceLROnPlateau* function to  
 122 reduce the learning rate by a factor of 10 if the validation loss does not decrease for 10  
 123 epochs, starting with a learning rate of  $10^{-3}$  and letting it have a minimum value of  $10^{-6}$ .  
 124 Using these conditions, the model trains for 75 epochs with an average training time of 1s  
 125 per epoch.

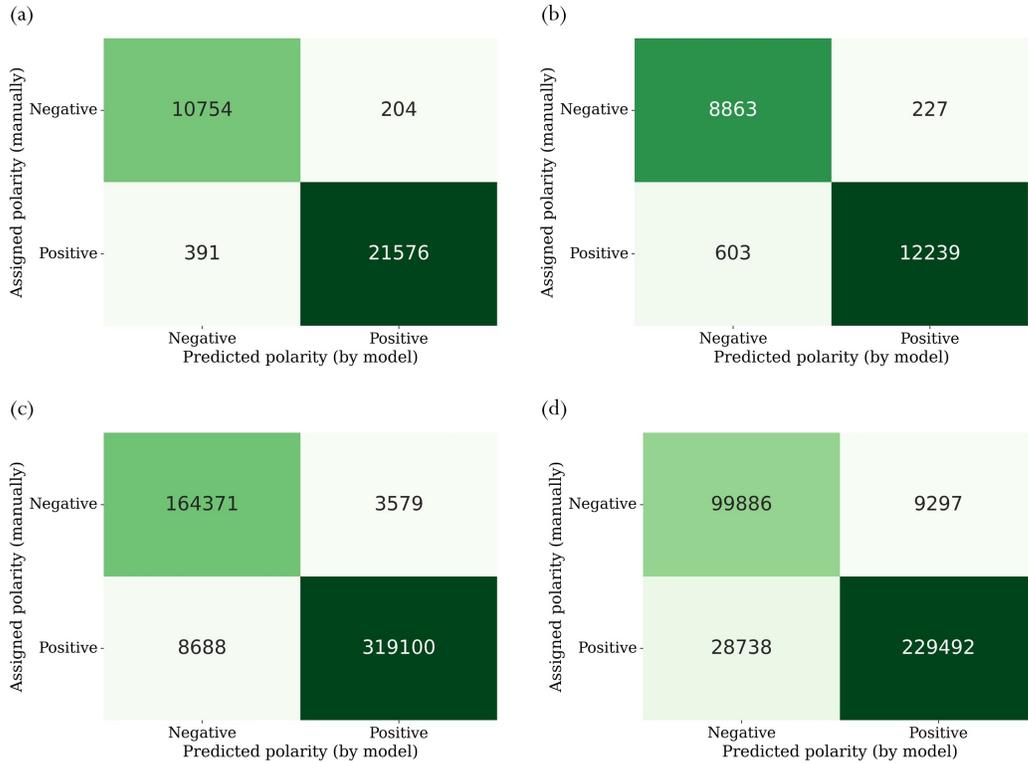
### 126 3 Results and Discussion

127 As stated previously, PolarCAP was trained only on data with signal-to-noise ratio  
 128 (SNR) above 10dB to ensure a good quality of training data, so we test it separately on data  
 129 with SNR above and below 10dB. This testing was performed on two independent datasets:  
 130 (i) the portion of the INSTANCE dataset (Michelini et al., 2021) not used for training or  
 131 validation and (ii) test dataset from Southern California Seismic Network (SCSN) used in  
 132 Ross et al. (2018). For the latter dataset, however, we ignored the waveforms for which the  
 133 polarity was ‘undetermined’. The results have been summarised in the first half of Table 1  
 134 and the corresponding confusion matrices (Ting, 2017) can be found in Figure 2. As one  
 135 can see, the accuracy for traces with SNR above 10dB is around 98% for both the datasets  
 136 which means for about 98% of the traces the polarity labels determined by our model agrees  
 137 with the polarity labels assigned through manual analysis. A few examples of such traces  
 138 are shown in Figure 3a. As expected, the accuracy is lower for smaller SNR since a higher  
 139 noise level makes the polarity information ambiguous and hence difficult to determine either  
 140 manually or using a deep learning model.

**Table 1.** Summary of Model Performance when trained with and without data augmentation on the two datasets for traces with SNR above and below 10dB

		Test Dataset	Accuracy (%)	Precision (%)		Recall (%)	
				Positive	Negative	Positive	Negative
Without Data Augmentation	INSTANCE <sup>1</sup>	SNR $\geq$ 10dB	98.19	99.06	96.49	98.22	98.12
		SNR $<$ 10dB	96.22	98.18	93.63	95.30	97.50
	SCSN <sup>2</sup>	SNR $\geq$ 10dB	97.53	98.89	94.98	97.35	97.87
		SNR $<$ 10dB	89.65	96.11	77.67	88.87	91.48
With Data Augmentation	INSTANCE <sup>1</sup>	SNR $\geq$ 10dB	97.65	98.61	95.77	97.86	97.23
		SNR $<$ 10dB	94.24	97.68	89.98	92.36	96.90
	SCSN <sup>2</sup>	SNR $\geq$ 10dB	97.78	98.99	95.50	97.64	98.05
		SNR $<$ 10dB	92.46	96.86	83.55	92.26	92.93

<sup>1</sup>Michelini et al. (2021); <sup>2</sup>Ross et al. (2018)



**Figure 2.** Confusion matrices for testing model on INSTANCE data (Michelini et al., 2021) with (a) SNR  $\geq$  10dB (b) SNR  $<$  10dB and SCSN data used by Ross et al. (2018) with (c) SNR  $\geq$  10dB (d) SNR  $<$  10dB.

### 3.1 Manual inspection of traces with mismatched assigned and predicted polarity

We further investigate the 595 traces from INSTANCE dataset with SNR above 10dB for which the polarity predicted by PolarCAP and the manually assigned polarities were in disagreement (refer to Figure 2). We could classify these cases into one of the following three categories:

- The polarity assigned by the analyst was correct.
- The polarity predicted by the model was correct.
- The polarity information was ambiguous due to high noise levels, or incorrect P-picking.

Based on our analysis we observed that for 40.8% of the cases the polarity predicted by the model was correct whereas in 27.6% of the cases the model predicted incorrect polarities. In the remaining 31.6% cases the polarity information was hard to determine through manual inspection, mostly due to high noise levels. Some examples of the former case can be found in Figure 3b. We further found that incorrect classification of traces by the model, was either due to incorrect picking of the first P-arrival (although one can see from Figure 3b that some degree of time shift is accounted for by the model based on what is encountered in the training data) or the P-arrivals being emergent in nature. We went on to inspect the traces where we identified the P-arrival sample to be incorrectly picked, and picked the P-arrivals ourselves using the EPick model (Li et al., 2021). Figure 4a shows five traces where the P-arrival samples determined by EPick seemed more accurate than those provided in the metadata. We then fed 64 sample windows centred around the picked P-phases and in each of these cases the polarities predicted by the model now matched with the assigned polarities. Figure 4b shows some examples of emergent arrivals. It is also observed that the probability of prediction for emergent onsets are usually lower than that for impulsive onsets.

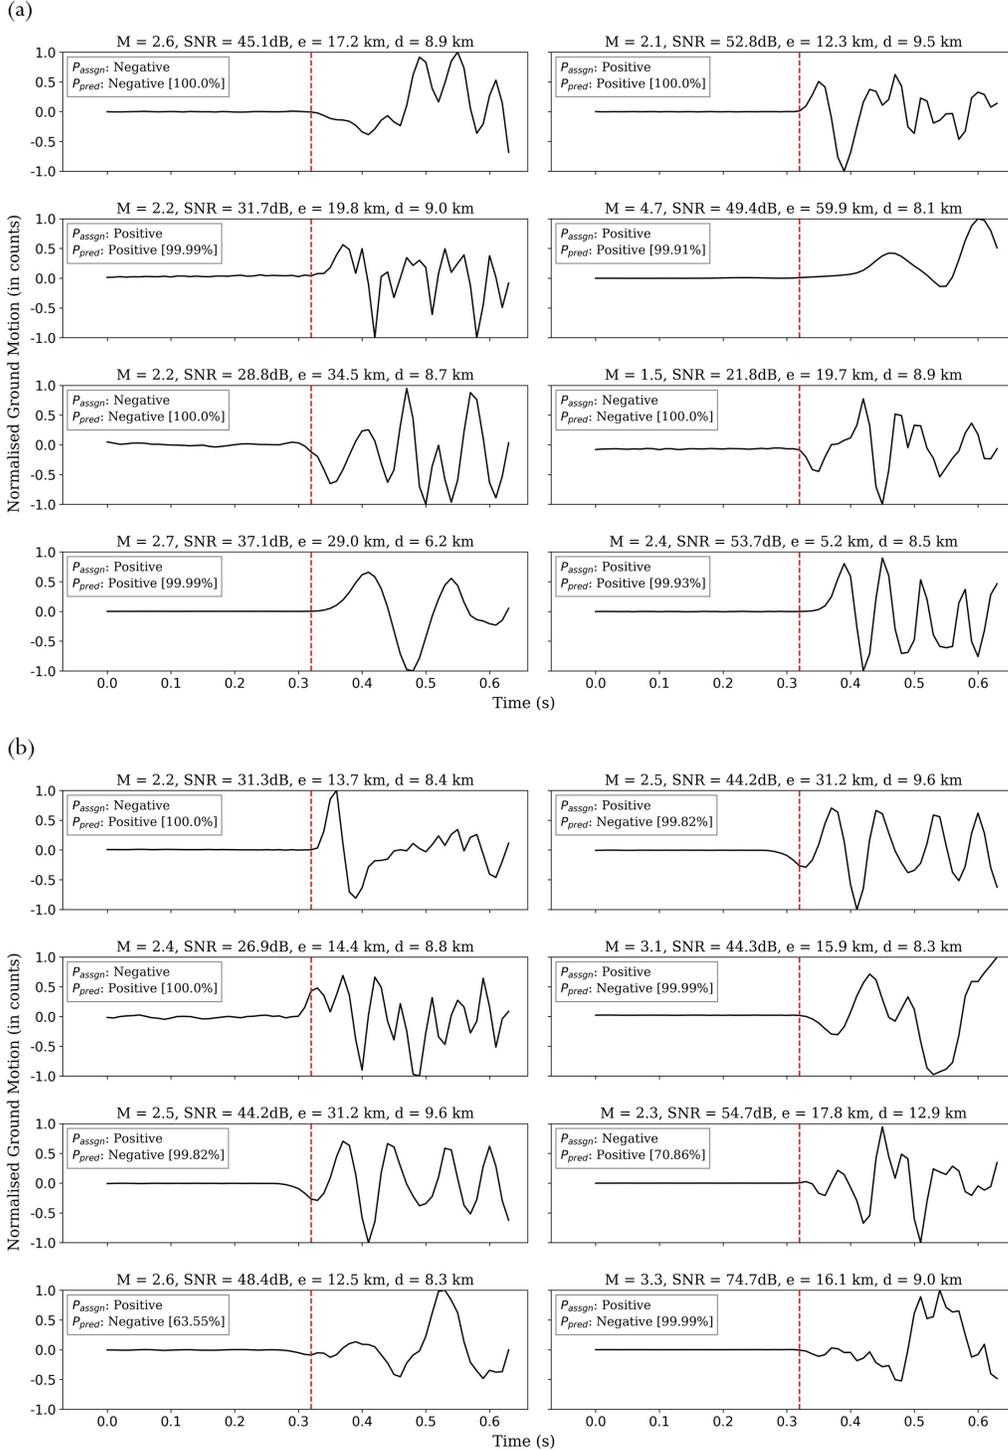
### 3.2 Factors affecting model accuracy

We looked at the distribution of the incorrectly classified traces in terms of signal-to-noise ratio, magnitude, focal depth, and epicentral distances. Signal-to-noise ratio (SNR) can influence the ease with which the first P-arrival is picked and hence the first motion polarity can be determined, and magnitude can be correlated with the SNR. However, we did not find any observable correlation between signal-to-noise ratio, magnitude and model accuracy, as shown in Figure 5a, which means the model is capable of performing polarity determination across a wide range of SNR and magnitudes. One can also see from Figure 5b that the incorrect classifications are restricted to shallower events ( $<80$  km) even though these are most represented in the training data. This is likely to be because deeper earthquakes tend to have a more impulsive nature as compared to shallow earthquakes of similar magnitude (Bormann et al., 2014).

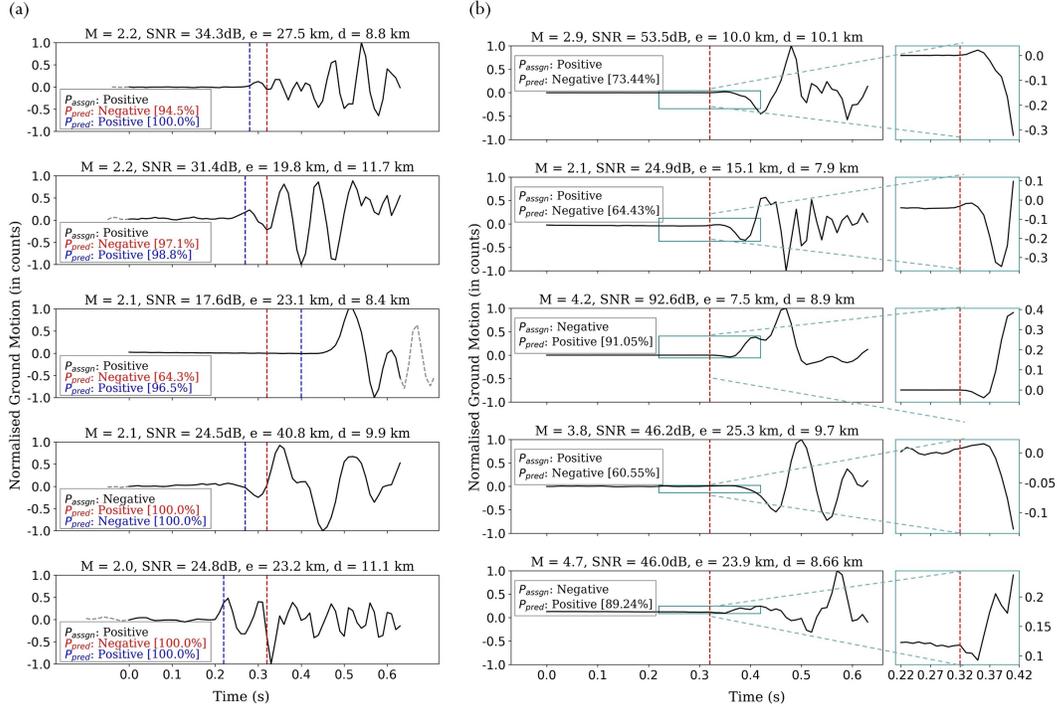
We also looked at the fraction of incorrect classification for different site conditions as characterised by the average shear wave velocity of the top 30 meters of subsurface ( $V_{S30}$ ). We find that the tendency for incorrect polarity classification is slightly higher for  $V_{S30} < 420ms^{-1}$ , although for the 11 events recorded at stations with  $V_{S30} < 240ms^{-1}$  there is no observed misclassification (Figure 5c).

### 3.3 Further improvement through augmentation

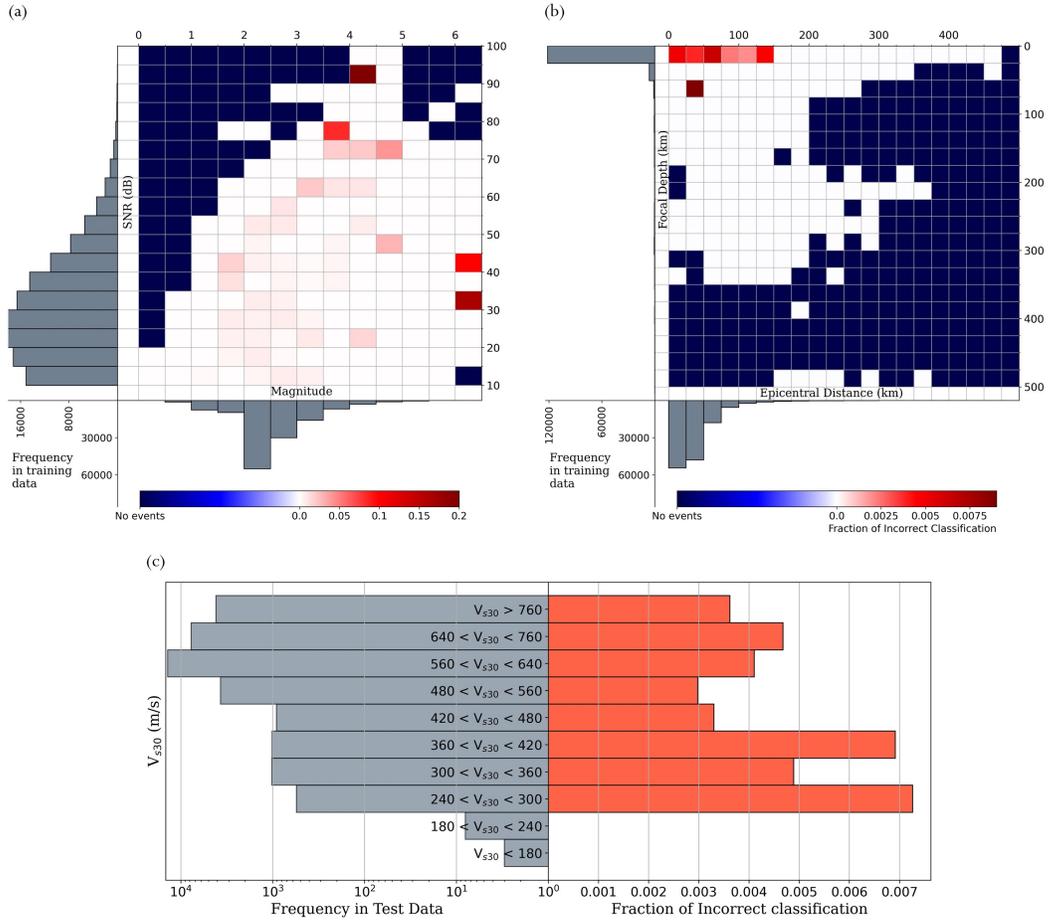
As outlined above, the incorrect determination of polarity was caused by incorrect picking of P-wave arrival time or emergent nature of P-onsets. In order to tackle the first issue, we used data augmentation to add a time shift to some of the traces in the training data. This augmentation technique was also explored in Uchide (2020). Since we did not



**Figure 3.** Examples of traces where (a) the polarity predicted by the model matches with the polarity assigned by human analysts, (b) the model predicts the correct polarity as opposed to manually assigned polarity. The red dashed line shows the P-arrival sample provided in the metadata. In the figure titles, M stands for event magnitude, SNR for signal-to-noise ratio, e for epicentral distance and d for focal depth.  $P_{assign}$  and  $P_{pred}$  represent the assigned and predicted polarities respectively and the percentages in square brackets represent probabilities corresponding to the predicted polarity.



**Figure 4.** (a) Examples of traces where the polarity prediction was incorrect due to an error in the P-arrival time provided in the metadata (showed with red-dashed line). The P-arrival sample picked by the model from Li et al. (2021) is shown with a dotted blue line. The dotted gray line shows the new time window. The predicted polarities in red and blue correspond to the arrival times provided in the metadata and those picked using Li et al. (2021) respectively (b) Examples of traces where the polarity was predicted incorrectly by the model due to the emergent nature of the P-arrival. The panel on the right shows a magnified plot of just 10 samples around the P-arrival sample (provided in the metadata). In the figure titles,  $M$  stands for event magnitude, SNR for signal-to-noise ratio,  $e$  for epicentral distance and  $d$  for focal depth.  $P_{assign}$  and  $P_{pred}$  represent the assigned and predicted polarities respectively and the percentages in square brackets represent probabilities corresponding to the predicted polarity.



**Figure 5.** Distribution of incorrectly classified traces as a fraction of total number of traces, with respect to (a) Signal-to-noise ratio and magnitude (b) Focal depth and epicentral distance (c) Average shear wave velocity of the top 30 meters of subsurface ( $V_{s30}$ ).

189 have any information on the quality of P-wave onsets in the metadata, it was not possible  
 190 to apply augmentation to increase the amount of emergent traces in the training data. The  
 191 time shift was applied to 1/5th of the traces (which were then added back to the dataset);  
 192 the amount of time shift was chosen from a normal distribution with mean 0 and standard  
 193 deviation of 10 samples. A maximum shift of 30 samples was allowed to ensure that the  
 194 first P-arrival is included in the window.

195 After re-training the model on the augmented data, we tested it again on the same test  
 196 sets. The corresponding evaluation of the model performance is shown in the second half  
 197 of Table 1. As one can see, the use of data augmentation resulted in lower accuracy on the  
 198 INSTANCE dataset (for both low and high SNR) This could be because in the dataset, more  
 199 often than not the traces had correctly picked P-arrival times and hence the augmentation  
 200 was not needed. In case of the test data from Ross et al. (2018), on the other hand, the  
 201 accuracy increased to 0.25% for high SNR data and 2.81% for low SNR data upon using  
 202 data augmentation.

## 203 4 Conclusion

204 In this study, we explored the potential of a deep learning model - PolarCAP to de-  
 205 termine first-motion polarity of earthquake waveforms when the P-arrival information is  
 206 available, faster and more accurately as compared to human analysts. We show that when  
 207 tested on unseen traces, the polarity predicted by the model, matches the ones assigned by  
 208 human analysts over 98% of the times. We observed that the major reasons behind incorrect  
 209 assignment of polarity by the model were incorrect P-arrival picks and emergent arrivals;  
 210 to that end we also found the event depth to indirectly affect the fraction of incorrect pre-  
 211 dictions by affecting the quality of P-wave arrival. However in spite of these hurdles we  
 212 find that when the polarities predicted by the model differed from those presented in the  
 213 metadata, it was usually the model that was correct (almost 41% of the times, while the  
 214 model was wrong only about 28% of the times), thus demonstrating its ability to overcome  
 215 human errors.

## 216 5 Open Research

217 The seismic waveforms used in our research are a part of two datasets – INSTANCE  
 218 (Michellini et al., 2021) which was downloaded from <https://doi.org/10.13127/instance>  
 219 (last accessed January 2022) and the dataset described in Ross et al. (2018) which can be  
 220 downloaded from <https://scedc.caltech.edu/data/deeplearning.html#picking.polarity>  
 221 (last accessed April 2022).

## 222 Acknowledgments

223 This research is supported by the “KI-Nachwuchswissenschaftlerinnen” - grant SAI 01IS20059  
 224 by the Bundesministerium für Bildung und Forschung - BMBF. Calculations were performed  
 225 at the Frankfurt Institute for Advanced Studies’ new GPU cluster, funded by BMBF for  
 226 the project Seismologie und Artifizielle Intelligenz (SAI). The authors are also thankful to  
 227 Darius Fenner and Jonas Köhler for their kind suggestions. H.St. gratefully acknowledges  
 228 the Judah M. Eisenberg Laureatus - Professur at Fachbereich Physik, Goethe Universität  
 229 Frankfurt, funded by the Walter Greiner Gesellschaft zur Förderung der physikalischen  
 230 Grundlagenforschung e.V.

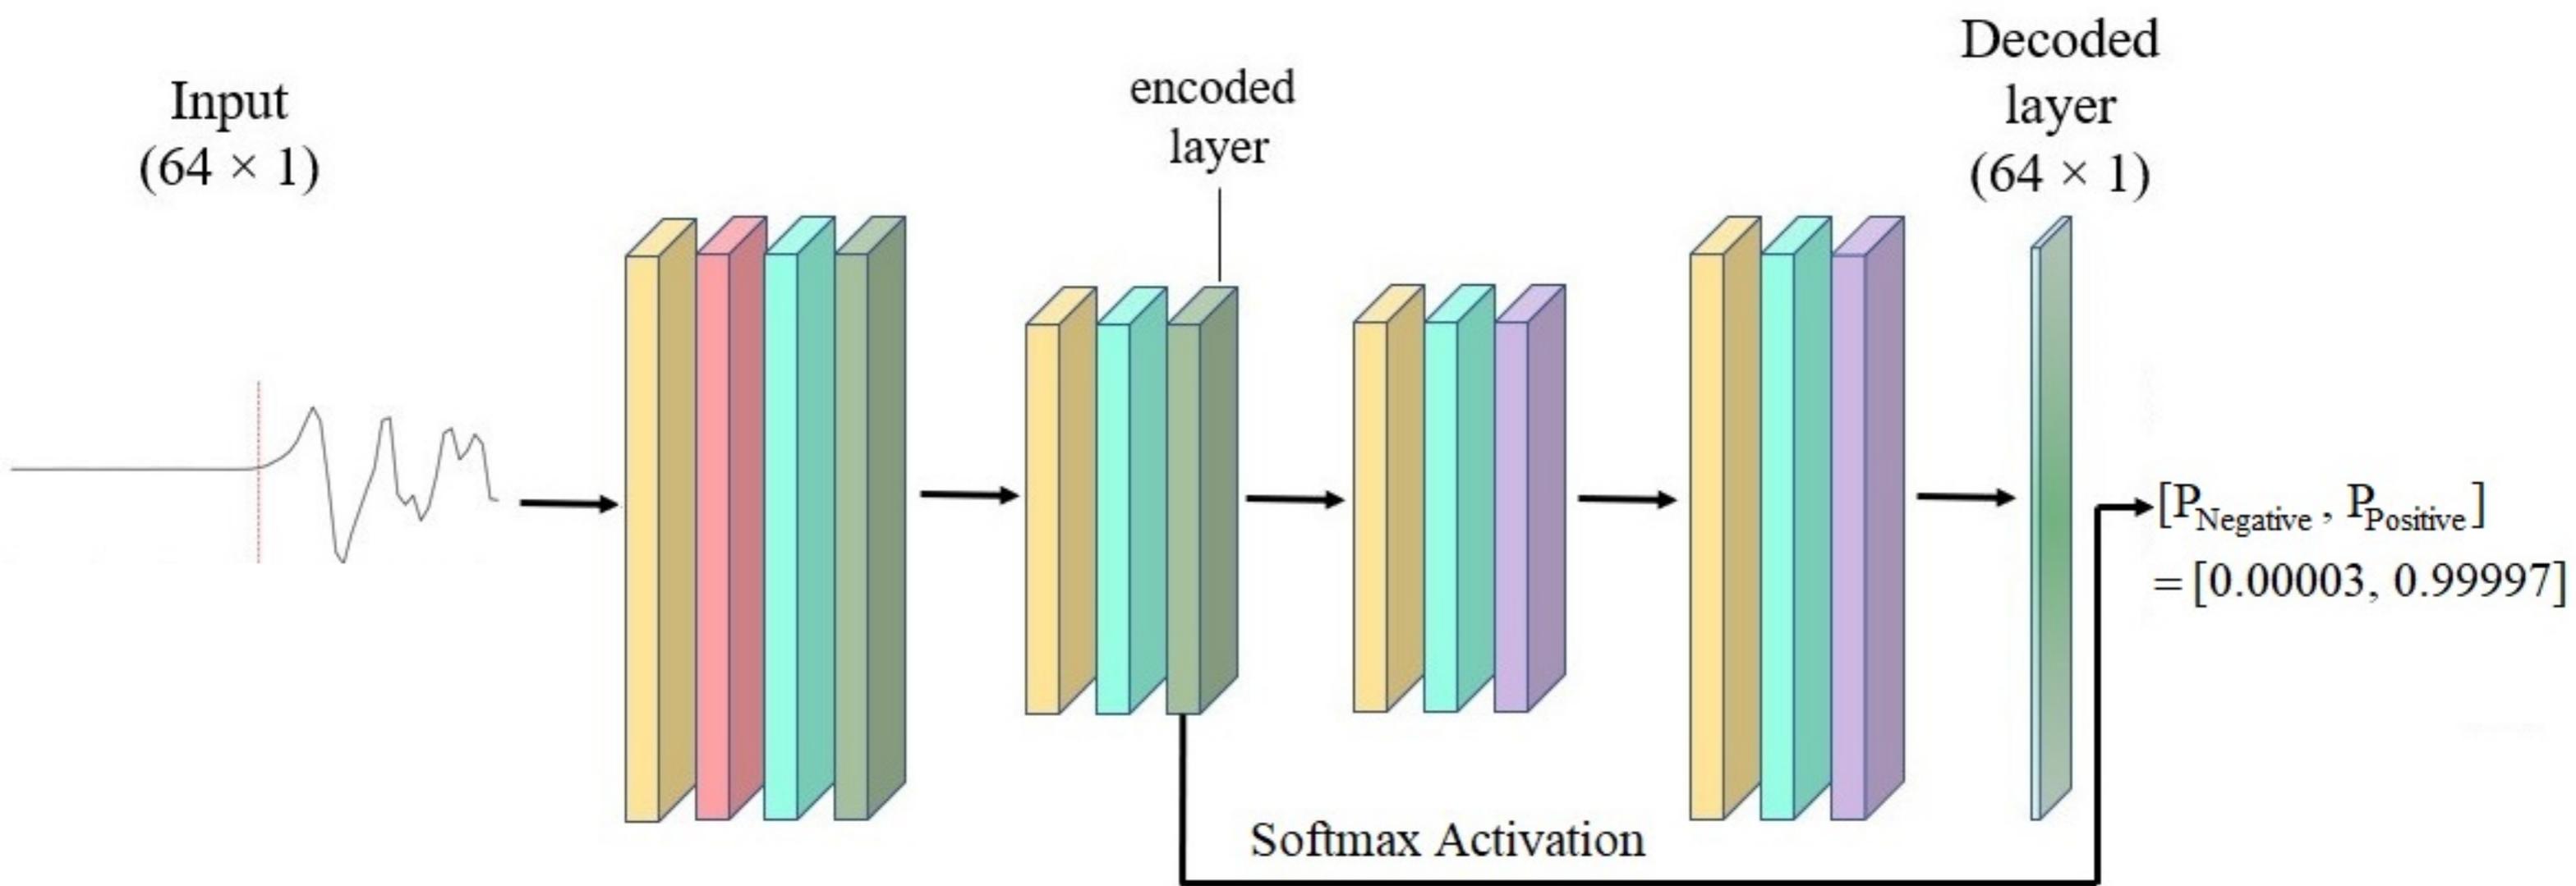
## 231 References

232 Batista, G. E. A. P. A., Prati, R. C., & Monard, M. C. (2004, June). A study of the  
 233 behavior of several methods for balancing machine learning training data. *SIGKDD*  
 234 *Explorations Newsletter*, 6(1), 20–29. Retrieved from <https://doi.org/10.1145/>

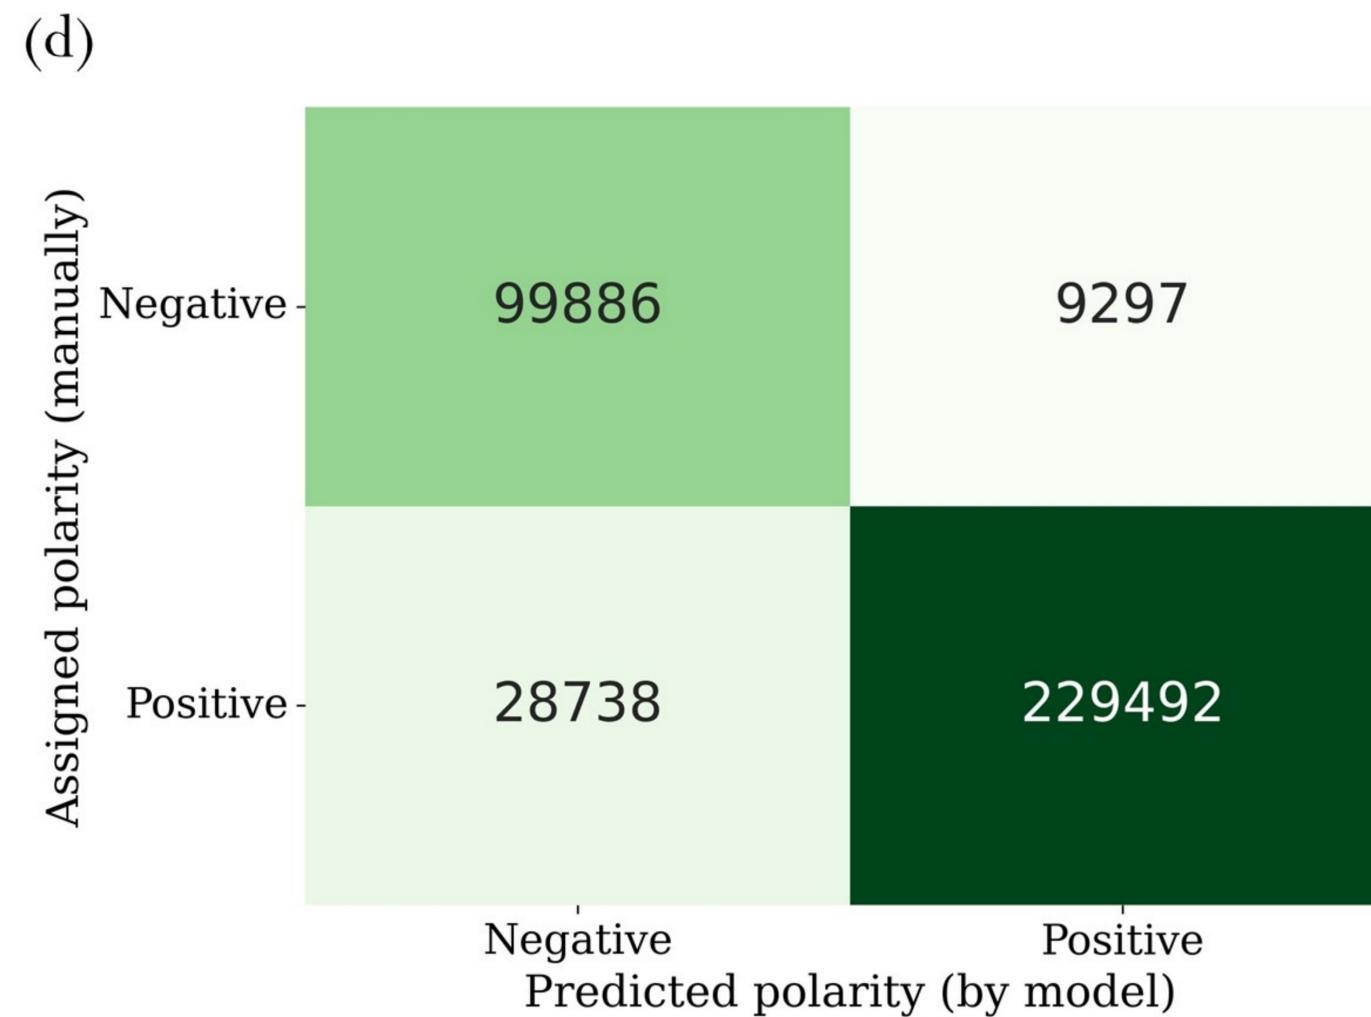
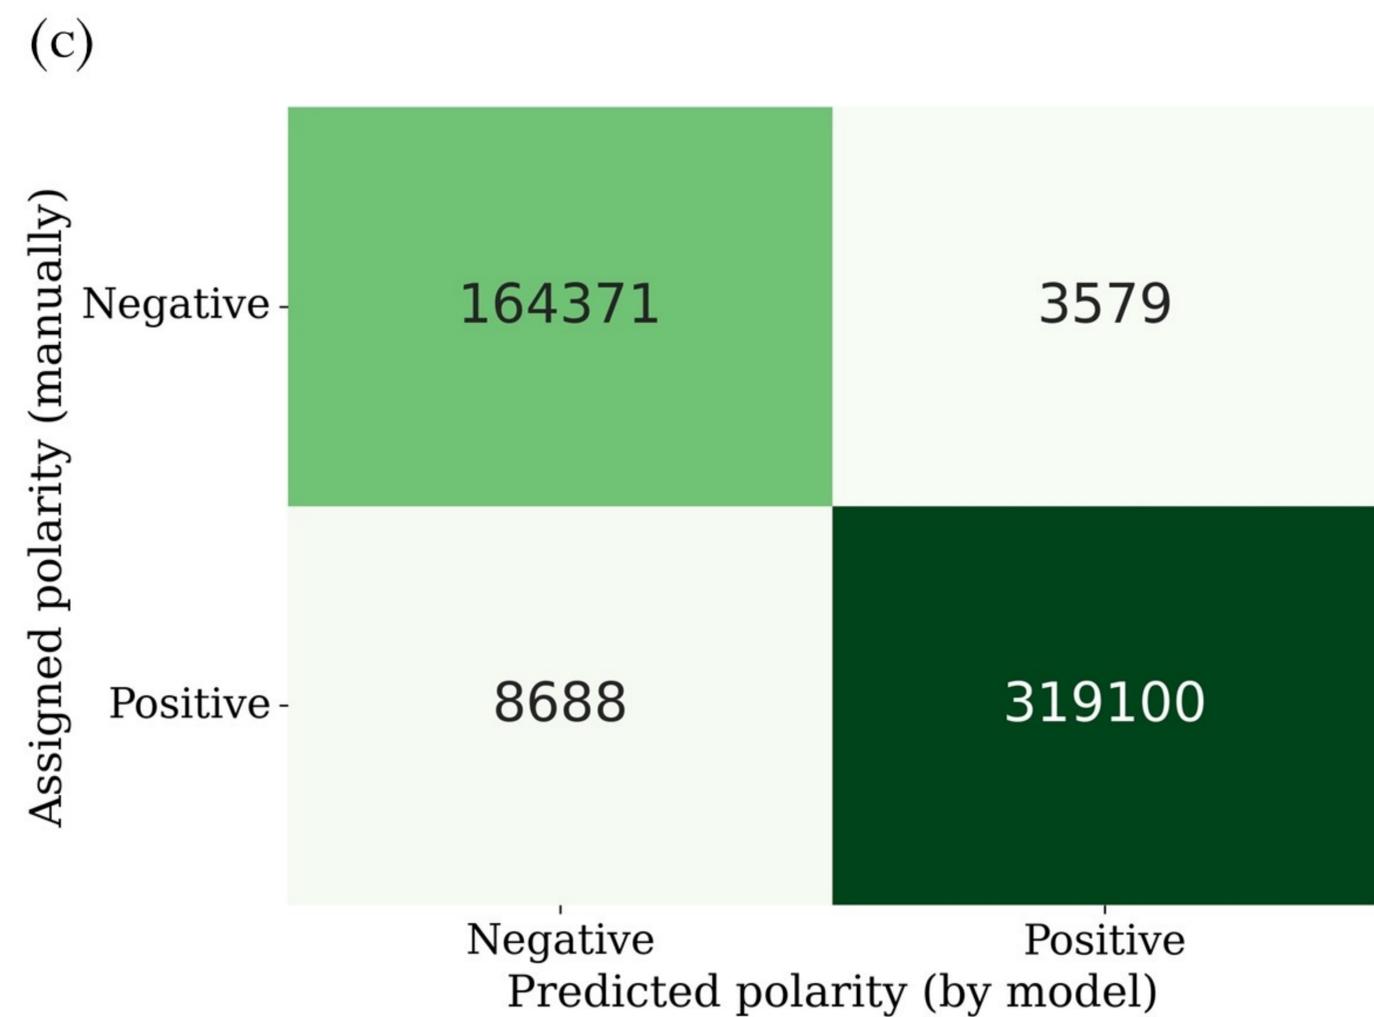
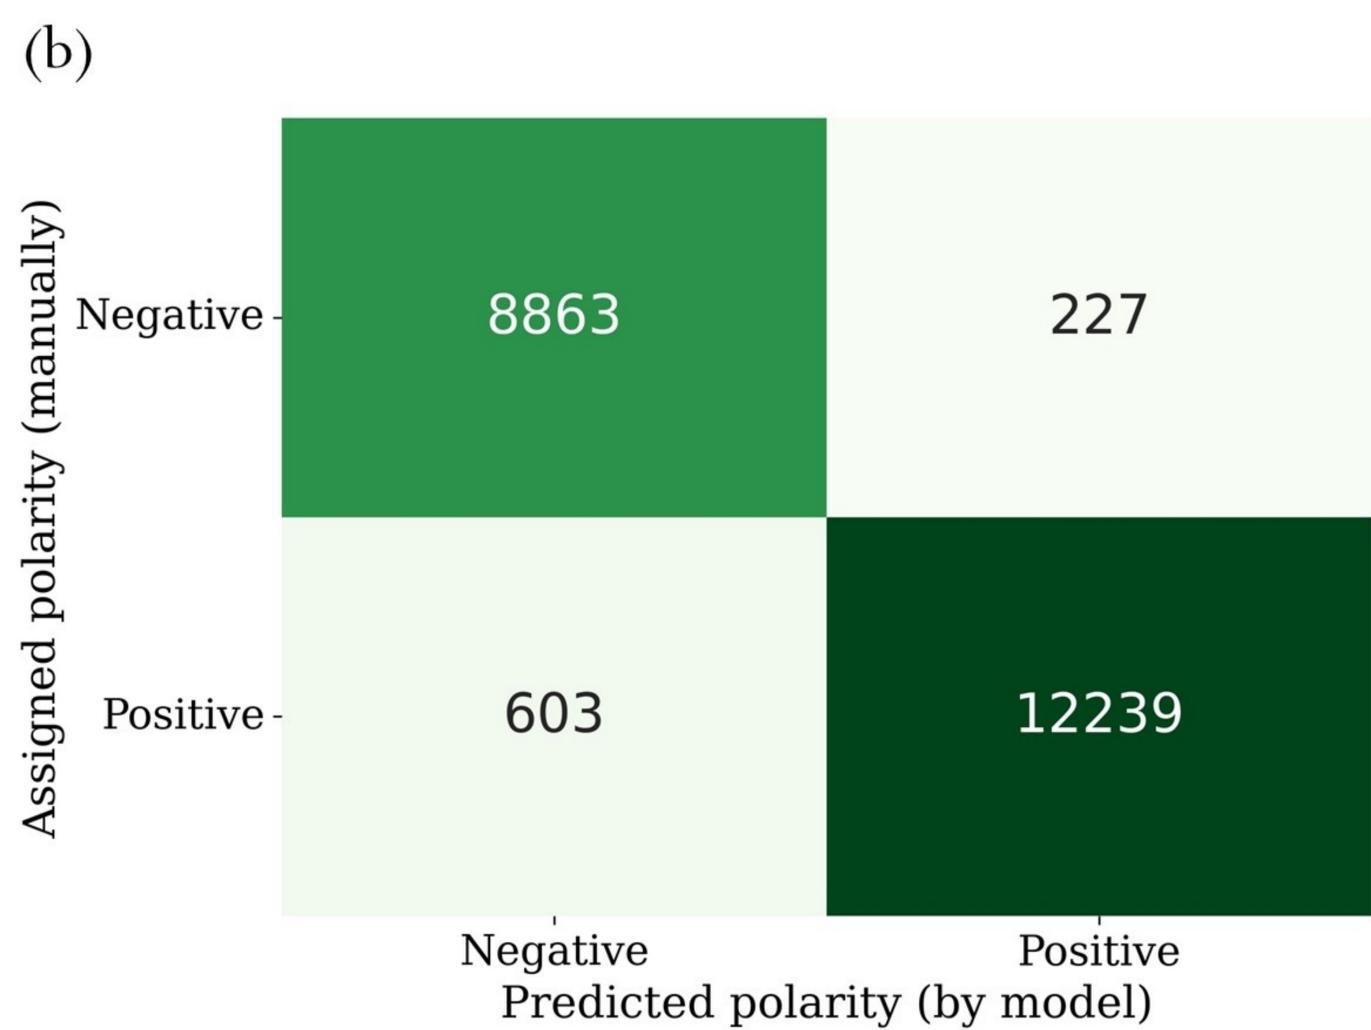
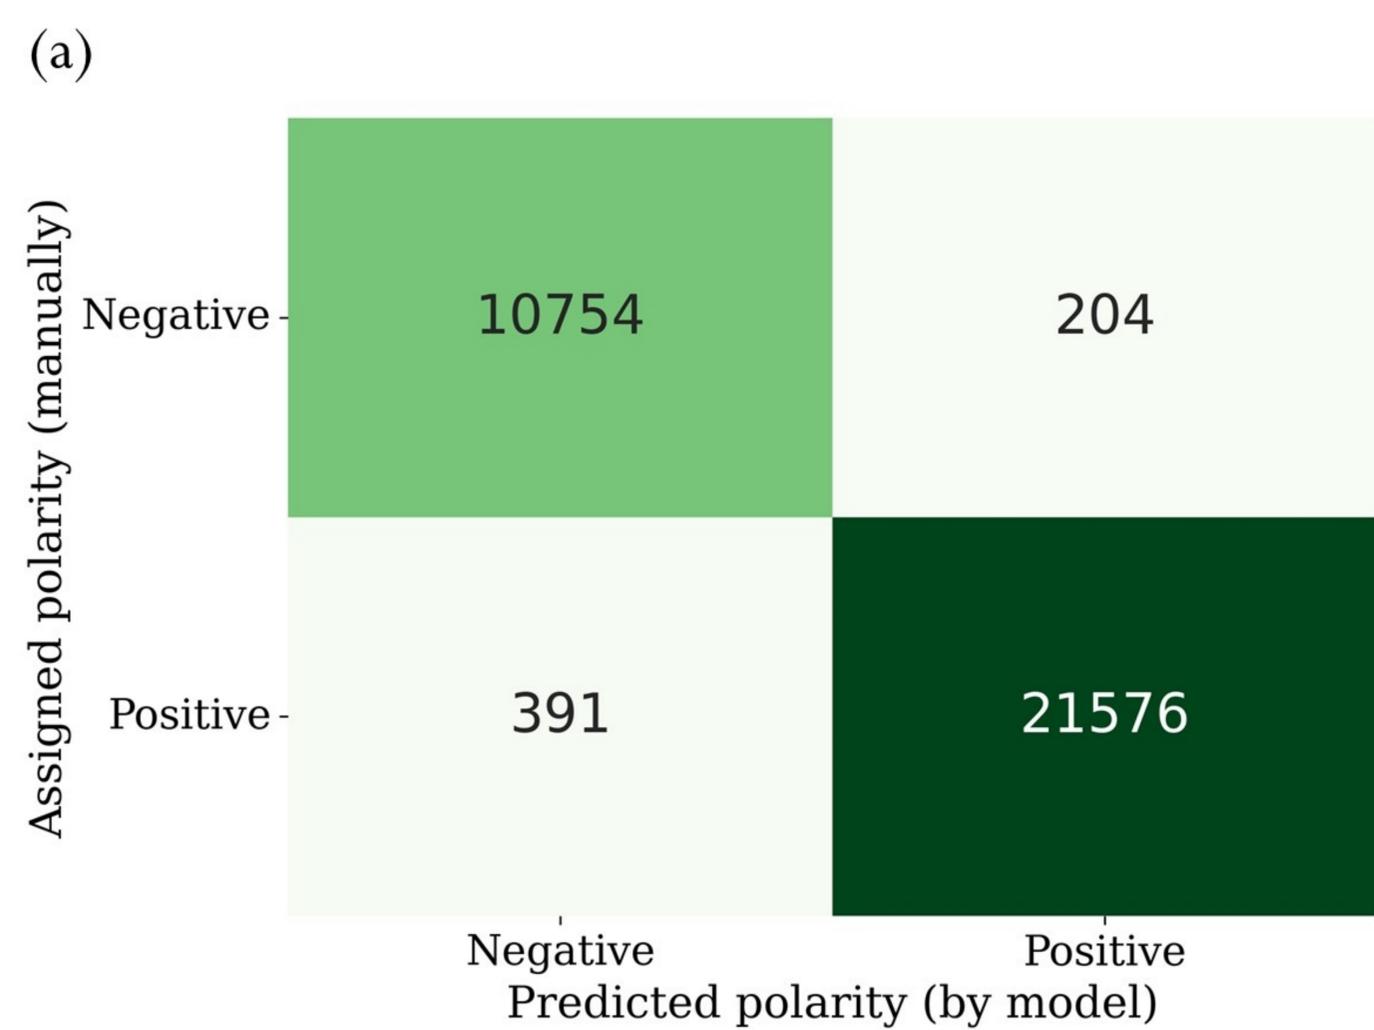
- 1007730.1007735 doi: 10.1145/1007730.1007735
- 235 Bormann, P., Klinge, K., & Wendt, S. (2014). Data analysis and seismogram interpretation.  
236 In *Bormann, p. (ed.), new manual of seismological observatory practice 2 (nmsop-2)*,  
237 *potdam : Deutsches geoforschungszentrum gfz* (p. 1-126).  
238
- 239 Brachmann, A., Barth, E., & Redies, C. (2017). Using cnn features to better under-  
240 stand what makes visual artworks special. *Frontiers in Psychology*, 8. Retrieved  
241 from <https://www.frontiersin.org/article/10.3389/fpsyg.2017.00830> doi:  
242 10.3389/fpsyg.2017.00830
- 243 Chakraborty, M., Fenner, D., Li, W., Faber, J., Zhou, K., Rumpker, G., ... Srivastava, N.  
244 (2022). Creime – a convolutional recurrent model for earthquake identification and  
245 magnitude estimation. *Earth and Space Science Open Archive*, 49. Retrieved from  
246 <https://doi.org/10.1002/essoar.10511140.1> doi: 10.1002/essoar.10511140.1
- 247 Chakraborty, M., Li, W., Faber, J., Ruempker, G., Stoecker, H., & Srivastava, N. (2021).  
248 *A study on the effect of input data length on deep learning based magnitude classifier*.  
249 arXiv. Retrieved from <https://arxiv.org/abs/2112.07551> doi: 10.48550/ARXIV  
250 .2112.07551
- 251 Chakraborty, M., Rumpker, G., Stöcker, H., Li, W., Faber, J., Fenner, D., ... Srivastava, N.  
252 (2004, June). Real time magnitude classification of earthquake waveforms using deep  
253 learning. *SIGKDD Explor. Newsl.*, 6(1), 20–29. Retrieved from [https://doi.org/  
254 10.1145/1007730.1007735](https://doi.org/10.1145/1007730.1007735) doi: 10.1145/1007730.1007735
- 255 Charles, P. (2013). *Project title*. GitHub. Retrieved from [https://github.com/  
256 charlespwd/project-title](https://github.com/charlespwd/project-title)
- 257 Chen, C., & Holland, A. A. (2016). Phasepapy: A robust pure python package for automatic  
258 identification of seismic phases. *Seismological Research Letters*, 87(6), 1384–1396.
- 259 Chen, Y. (2018). Automatic microseismic event picking via unsupervised machine learning.  
260 *Geophys. J. Int.*, 212(1), 88–102.
- 261 Hara, S., Fukahata, Y., & Iio, Y. (2019). P-wave first-motion polarity determination of  
262 waveform data in western Japan using deep learning. *Earth Planets Space*, 71(127).  
263 doi: <https://doi.org/10.1186/s40623-019-1111-x>
- 264 Hardebeck, J. L., & Shearer, P. M. (2002). A new method for determining first-motion focal  
265 mechanisms. *Bulletin of the Seismological Society of America*, 92(6), 2264–2276.
- 266 Kingma, D. P., & Ba, J. (2014, 12). *Adam: A method for stochastic optimization*.
- 267 Kiranyaz, S., Ince, T., Hamila, R., & Gabbouj, M. (2015). Convolutional neural networks  
268 for patient-specific ecg classification. , 2608-2611. doi: 10.1109/EMBC.2015.7318926
- 269 Li, W., Chakraborty, M., Fenner, D., Faber, J., Zhou, K., Rumpker, G., ... Srivastava,  
270 N. (2021). *Epick: Multi-class attention-based u-shaped neural network for earthquake  
271 detection and seismic phase picking*. Retrieved from [https://arxiv.org/abs/2109  
272 .02567](https://arxiv.org/abs/2109.02567)
- 273 Li, W., Sha, Y., Zhou, K., Faber, J., Ruempker, G., Stoecker, H., & Srivastava, N.  
274 (2022). *Deep learning-based small magnitude earthquake detection and seismic phase  
275 classification*. arXiv. Retrieved from <https://arxiv.org/abs/2204.02870> doi:  
276 10.48550/ARXIV.2204.02870
- 277 Liao, W., Lee, E., Mu, D., Chen, P., & Rau, R. (2021, 03). ARRU Phase Picker: Attention  
278 Recurrent-Residual U-Net for Picking Seismic P- and S-Phase Arrivals. *Seismol. Res.  
279 Lett.*, 92(4), 2410-2428. Retrieved from <https://doi.org/10.1785/0220200382> doi:  
280 10.1785/0220200382
- 281 Lundervold, A. S., & Lundervold, A. (2019). An overview of deep learning in  
282 medical imaging focusing on mri. *Zeitschrift für Medizinische Physik*, 29(2),  
283 102-127. Retrieved from [https://www.sciencedirect.com/science/article/pii/  
284 S0939388918301181](https://www.sciencedirect.com/science/article/pii/S0939388918301181) (Special Issue: Deep Learning in Medical Physics) doi:  
285 <https://doi.org/10.1016/j.zemedi.2018.11.002>
- 286 Michellini, A., Cianetti, S., Gaviano, S., Giunchi, C., Jozinović, D., & Lauciani, V. (2021).  
287 Instance the italian seismic dataset for machine learning, seismic waveforms and as-  
288 sociated metadata. *Istituto Nazionale di Geofisica e Vulcanologia (INGV)*. Retrieved  
289 from <https://doi.org/10.13127/instance>

- 290 Mousavi, S. M., & Beroza, G. C. (2020). A machine-learning approach for earthquake  
291 magnitude estimation. *Geophys. Res. Lett.*, *47*, e2019GL085976. Retrieved from  
292 <https://doi.org/10.1029/2019GL085976>
- 293 Mousavi, S. M., Ellsworth, W. L., Zhu, W., Chuang, L., & Beroza, G. (2020). Earthquake  
294 transformer—an attentive deep-learning model for simultaneous earthquake detection  
295 and phase picking. *Nat. Commun.*, *11*(3952). doi: <https://doi.org/10.1038/s41467-020-17591-w>
- 296
- 297 Mousavi, S. M., Zhu, W., Ellsworth, W., & Beroza, G. (2019). Unsupervised clustering of  
298 seismic signals using deep convolutional autoencoders. *IEEE Geoscience and Remote  
299 Sensing Letters*, *16*(11), 1693-1697. doi: 10.1109/LGRS.2019.2909218
- 300 Nagi, J., Ducatelle, F., Di Caro, G. A., Cireşan, D., Meier, U., Giusti, A., ... Gambardella,  
301 L. M. (2011). Max-pooling convolutional neural networks for vision-based hand gesture  
302 recognition. In *2011 IEEE International Conference on Signal and Image Processing  
303 Applications (ICSIPA)* (p. 342-347). doi: 10.1109/ICSIPA.2011.6144164
- 304 Perol, T., Gharbi, M., & Denolle, M. (2018). Convolutional neural network for earthquake  
305 detection and location. *Sci. Adv.*, *4*(2), e1700578. doi: 10.1126/sciadv.1700578
- 306 Prechelt, L. (2012). Early stopping — but when? In *Neural networks: Tricks of the trade:  
307 Second edition* (pp. 53–67). Springer Berlin Heidelberg. Retrieved from [https://doi.org/10.1007/978-3-642-35289-8\\_5](https://doi.org/10.1007/978-3-642-35289-8_5) doi: 10.1007/978-3-642-35289-8\_5
- 308
- 309 Pugh, D., White, R., & Christie, P. (2016). Automatic bayesian polarity determination.  
310 *Geophysical Journal International*, *206*(1), 275–291.
- 311 Ross, Z. E., Meier, M.-A., & Hauksson, E. (2018). P wave arrival picking and first-motion po-  
312 larity determination with deep learning. *Journal of Geophysical Research: Solid Earth*,  
313 *123*(6), 5120-5129. Retrieved from <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2017JB015251> doi: <https://doi.org/10.1029/2017JB015251>
- 314
- 315 Rumelhart, D., Hinton, G., & Williams, R. (1986). Learning internal representations by error  
316 propagation. In *Explorations in the microstructure of cognition* (Vol. 1, p. 318–362).  
317 MIT Press, Cambridge, MA, USA. Retrieved from [http://dl.acm.org/citation](http://dl.acm.org/citation.cfm?id=104279.104293)  
318 [.cfm?id=104279.104293](http://dl.acm.org/citation.cfm?id=104279.104293)
- 319 Ting, K. M. (2017). Confusion matrix. In *Encyclopedia of machine learning and data  
320 mining* (pp. 260–260). Boston, MA: Springer US. Retrieved from [https://doi.org/10.1007/978-1-4899-7687-1\\_50](https://doi.org/10.1007/978-1-4899-7687-1_50) doi: 10.1007/978-1-4899-7687-1\_50
- 321
- 322 Uchide, T. (2020, 08). Focal mechanisms of small earthquakes beneath the Japanese is-  
323 lands based on first-motion polarities picked using deep learning. *Geophysical Journal  
324 International*, *223*(3), 1658-1671. doi: 10.1093/gji/ggaa401
- 325 Voulodimos, A., Doulamis, N., Doulamis, A., & Protopapadakis, E. (2018). Deep learning  
326 for computer vision: A brief review. *Computational Intelligence and Neuroscience*,  
327 *2018*(7068349). doi: <https://doi.org/10.1155/2018/7068349>
- 328 Zhu, W., & Beroza, G. C. (2019). Phasenet: a deep-neural-network-based seismic arrival-  
329 time picking method. *Geophys. J. Int.*, *216*(1), 261-273. Retrieved from <https://doi.org/10.1093/gji/ggy423>
- 330

Figure\_1.jpg.

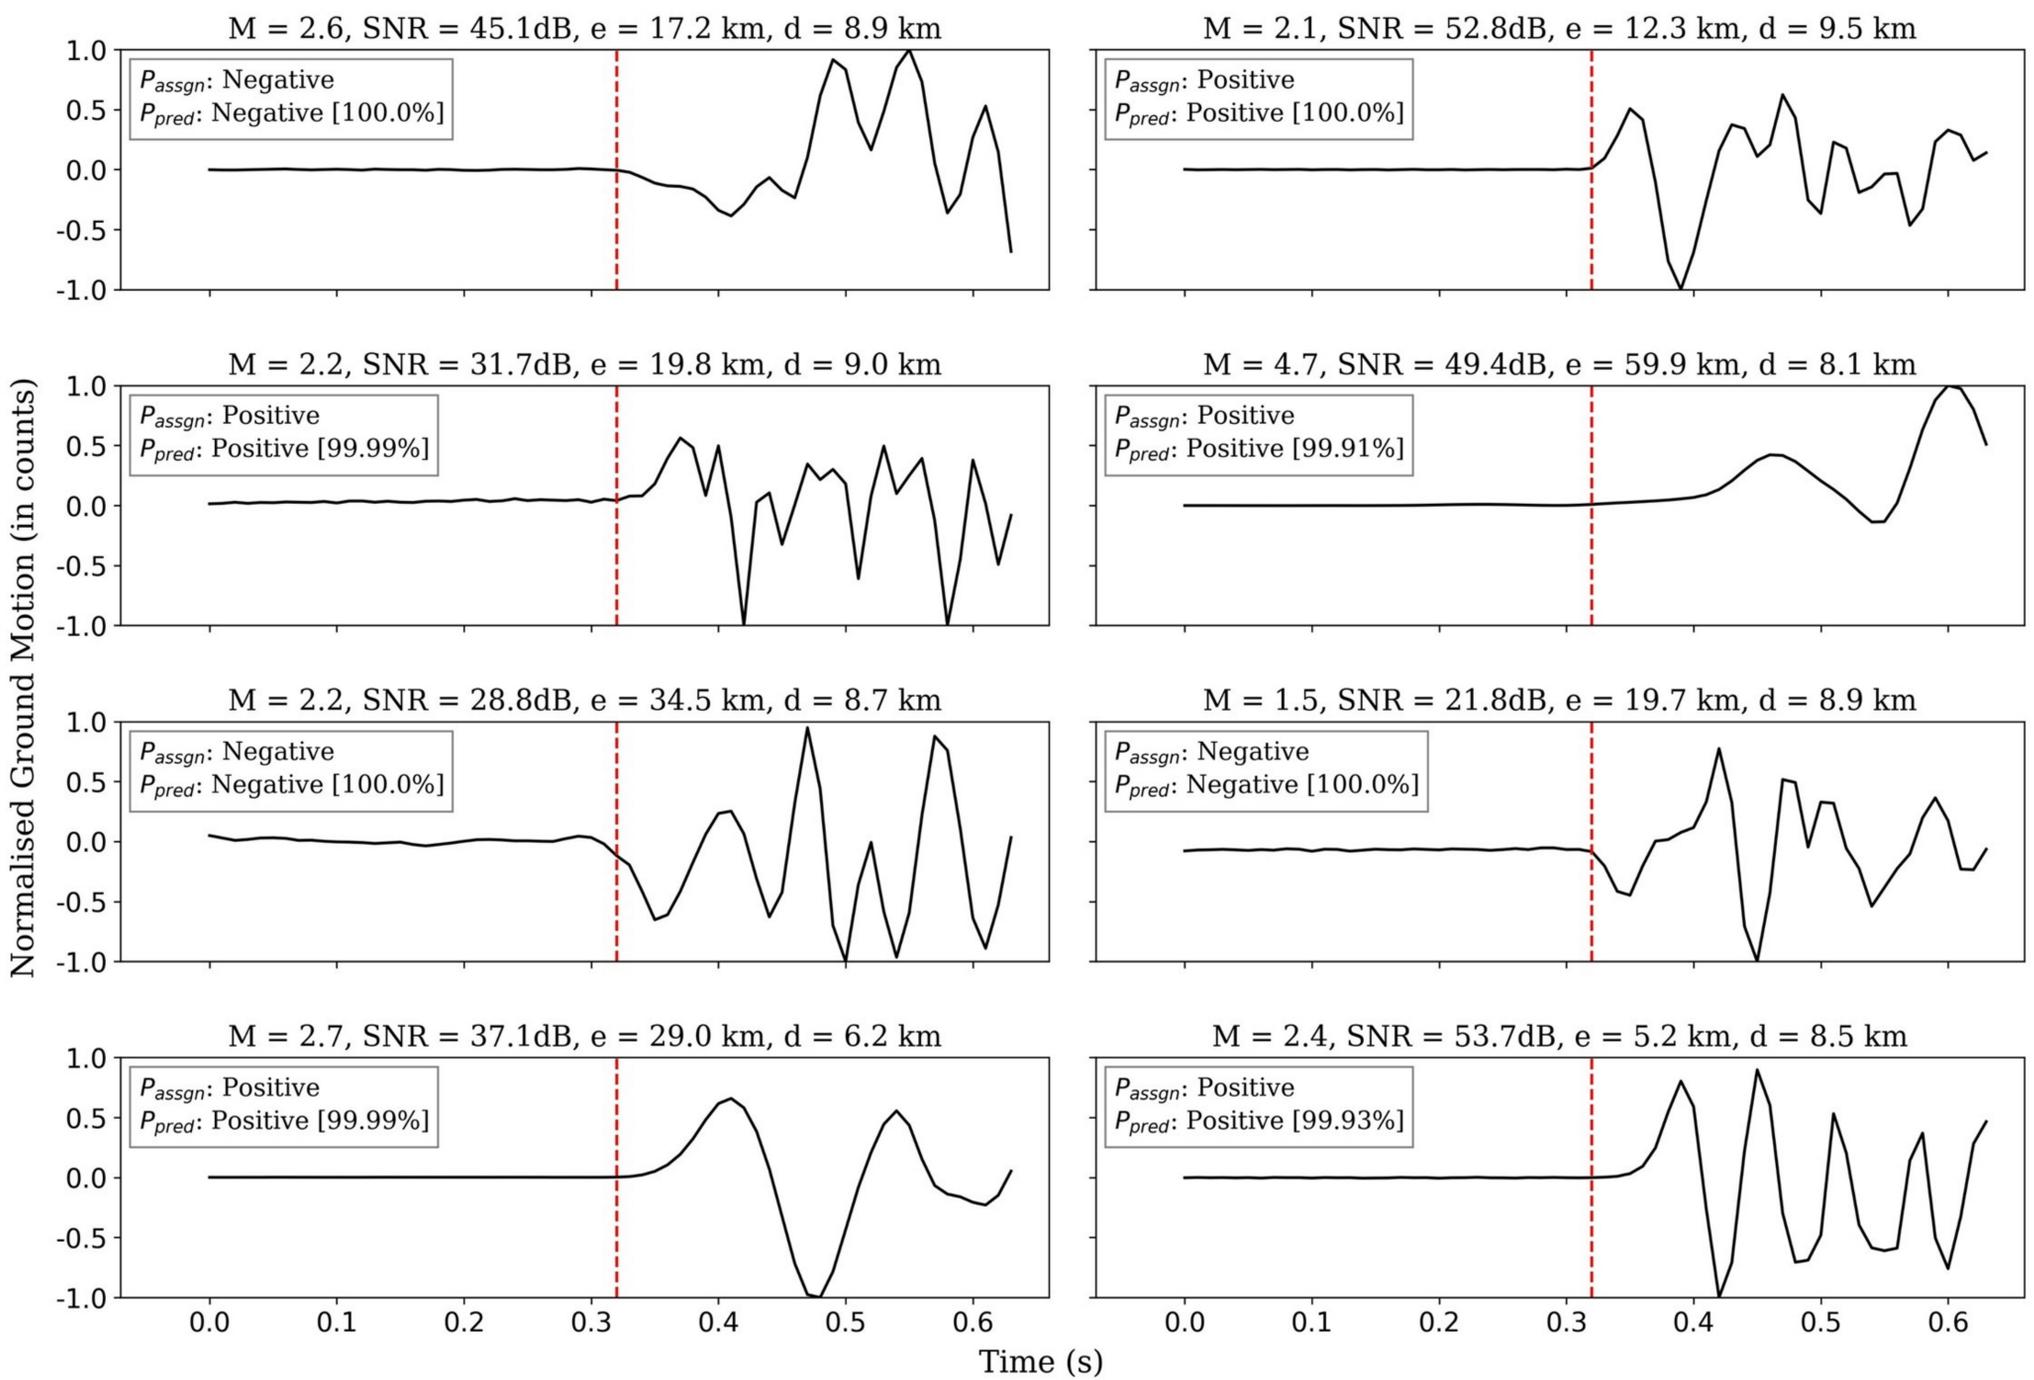


Figure\_2.jpg.

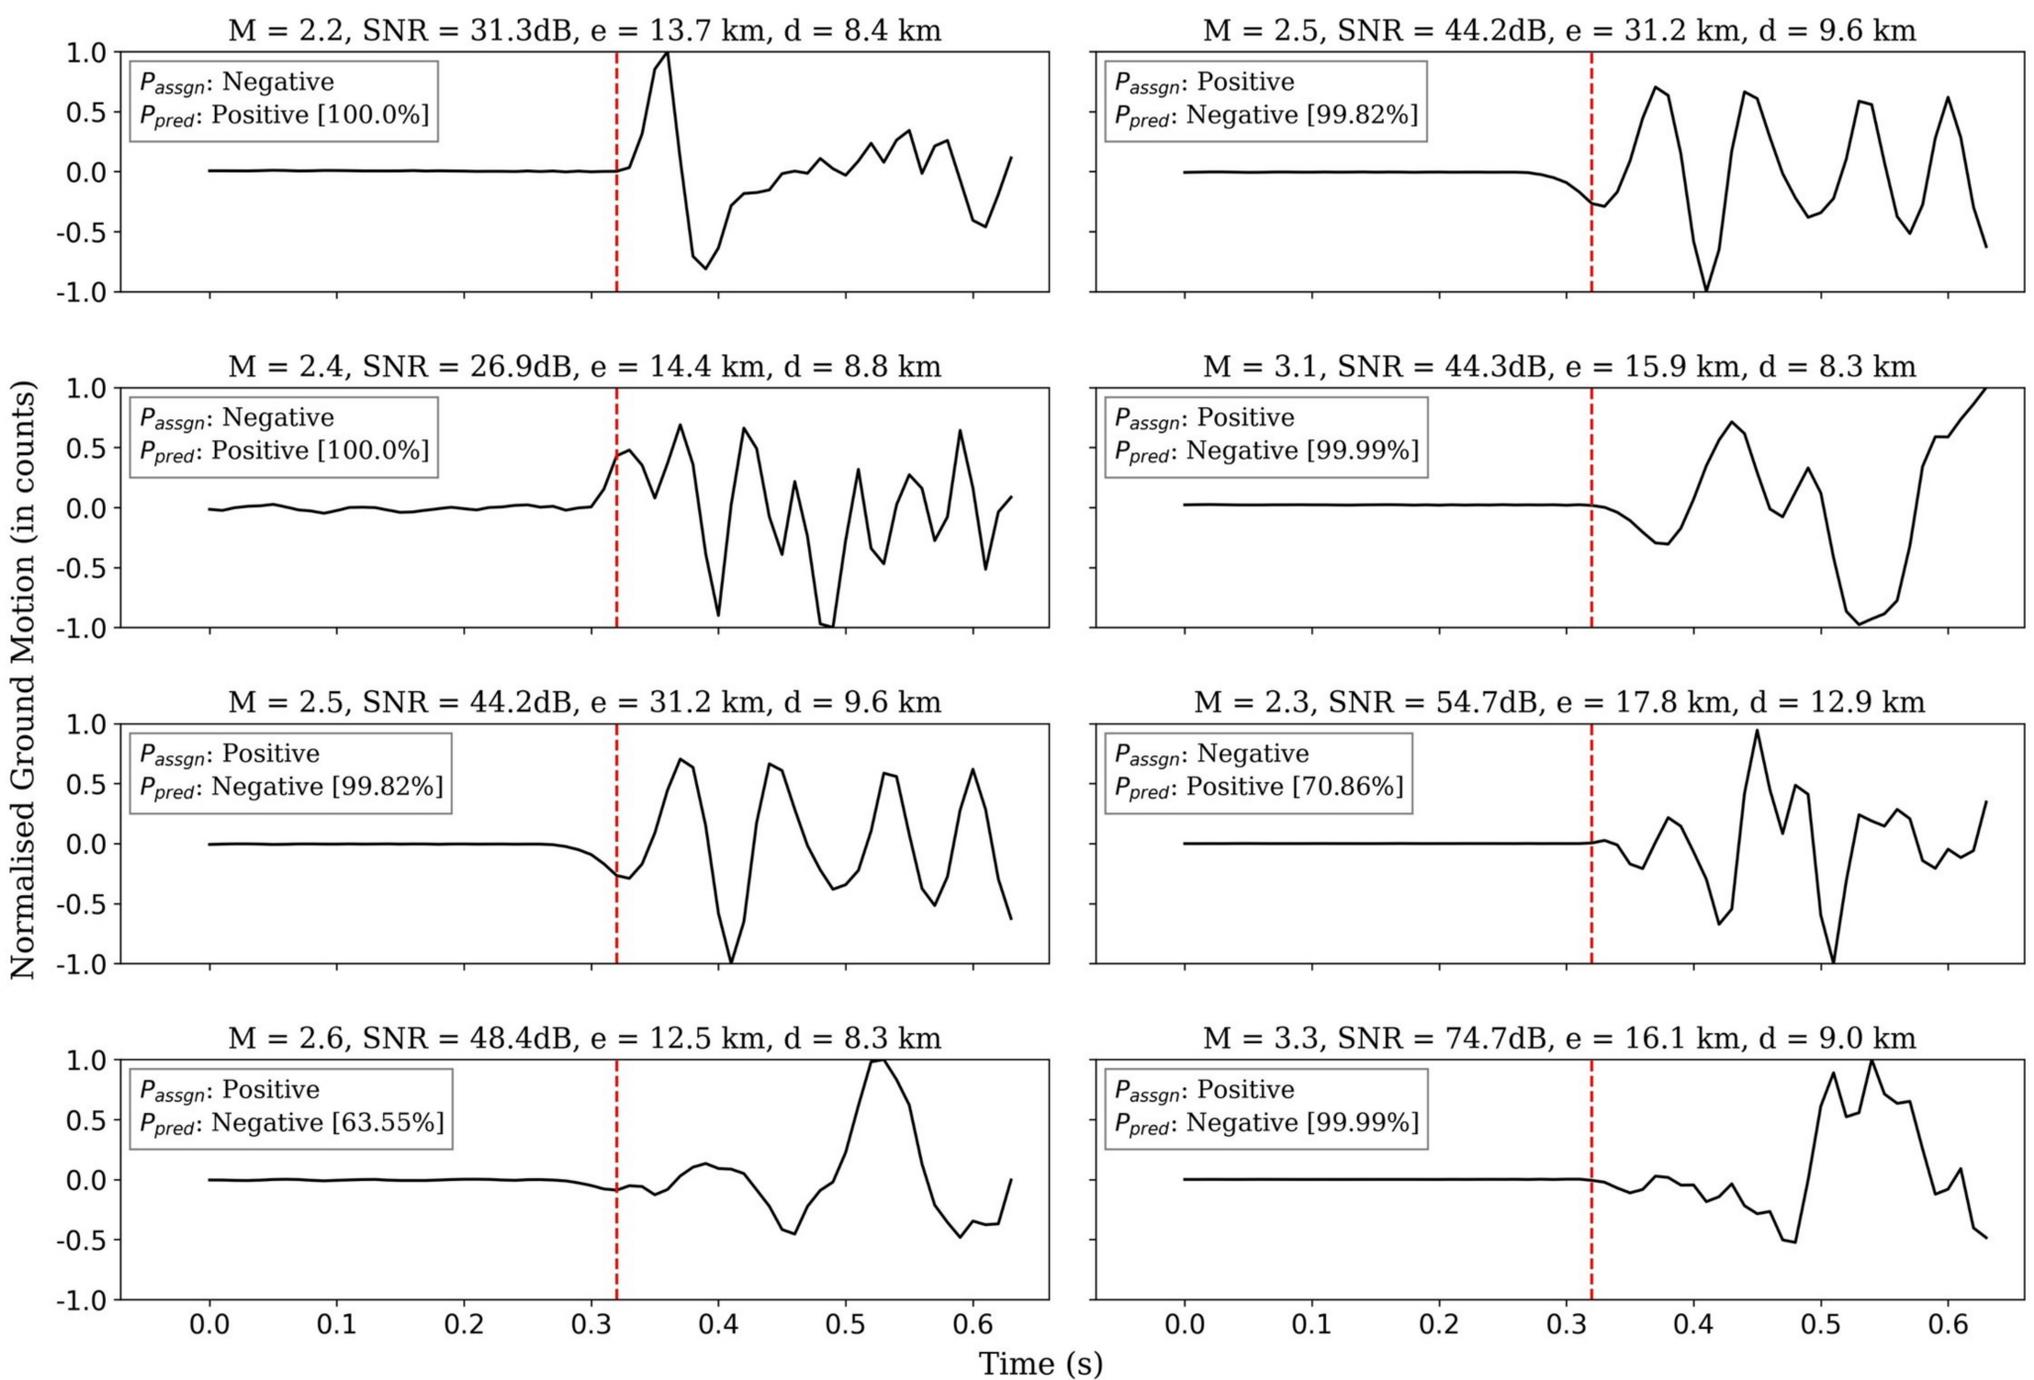


Figure\_3.jpg.

(a)

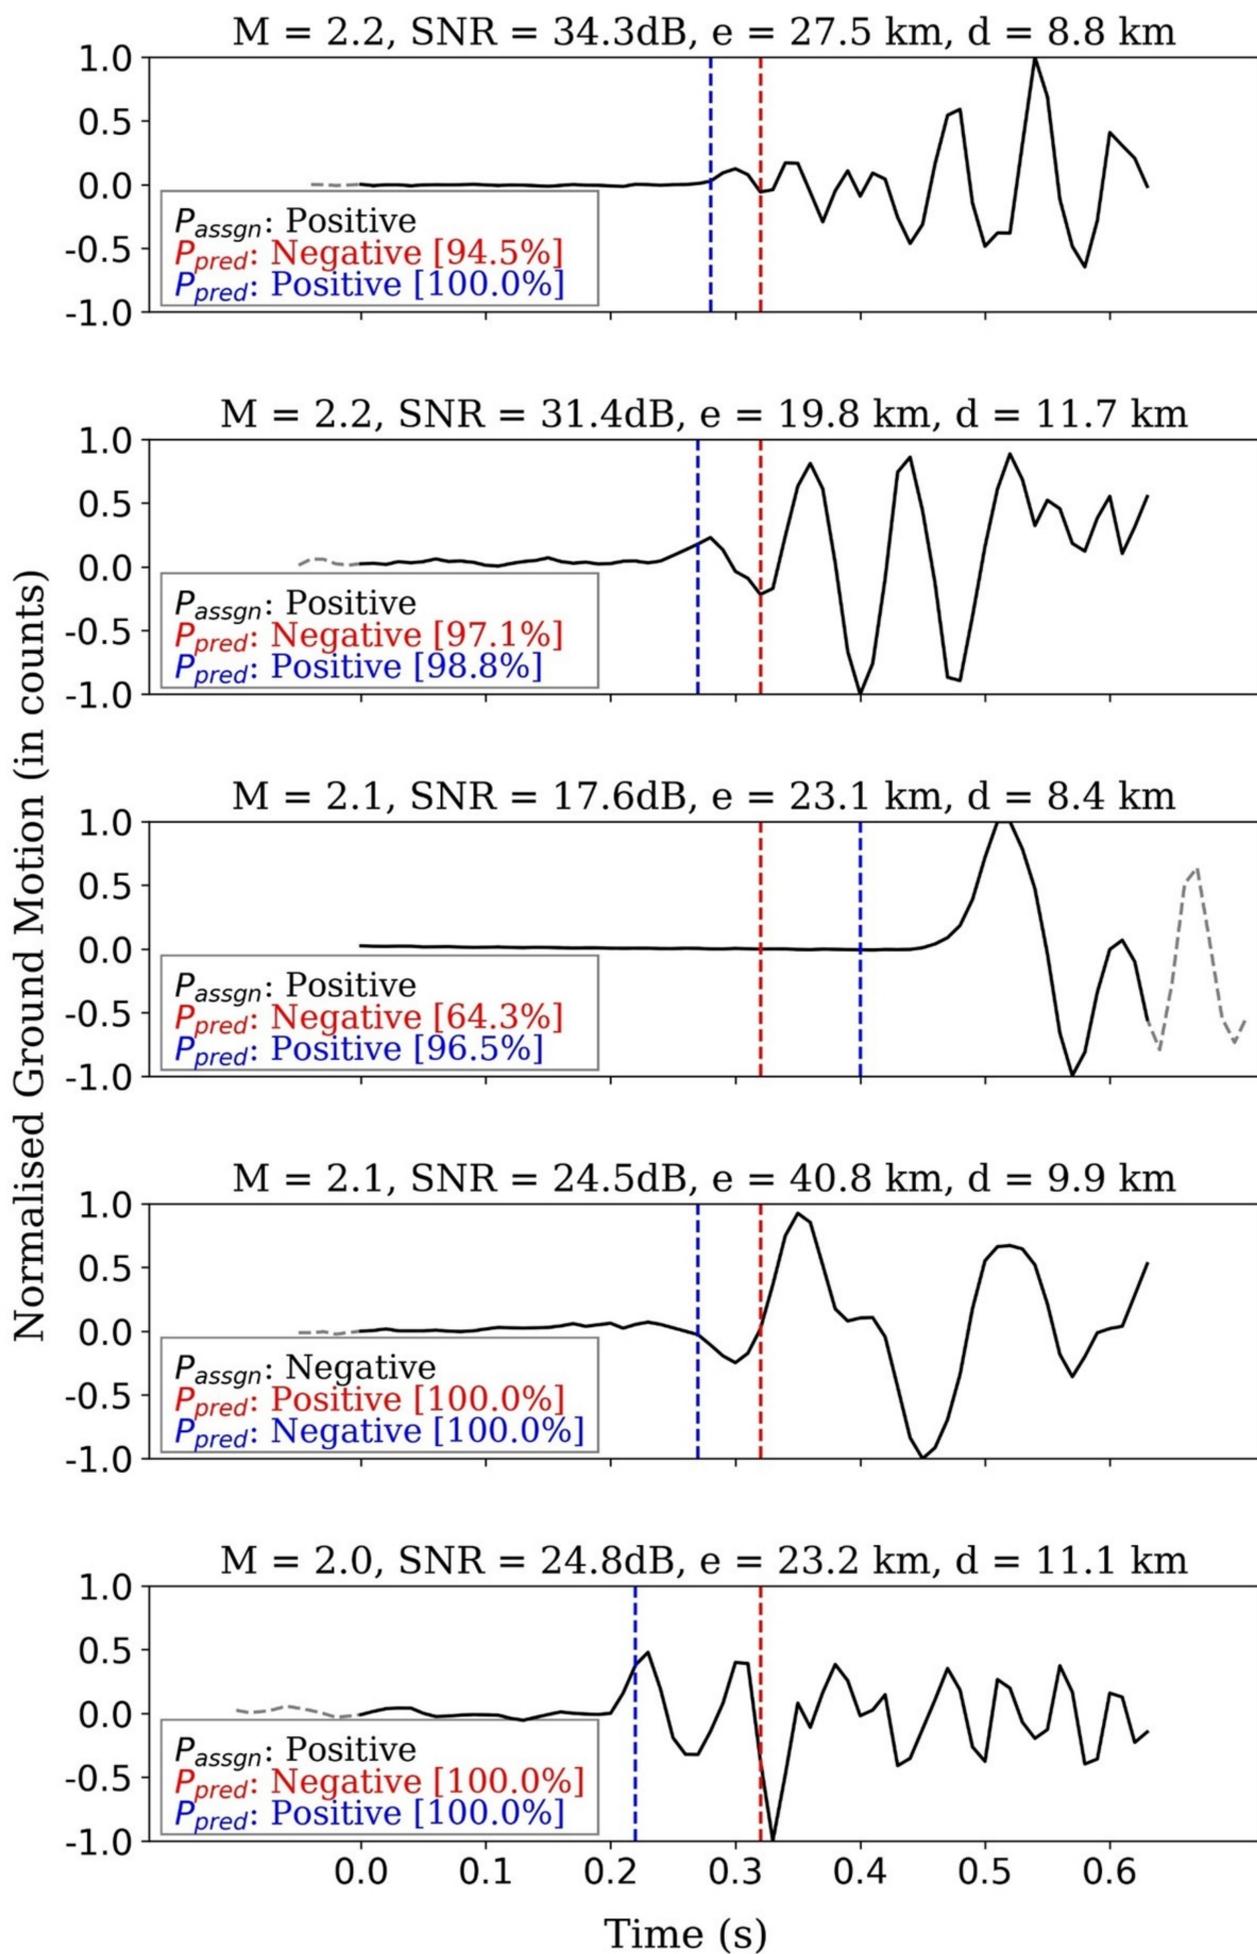


(b)

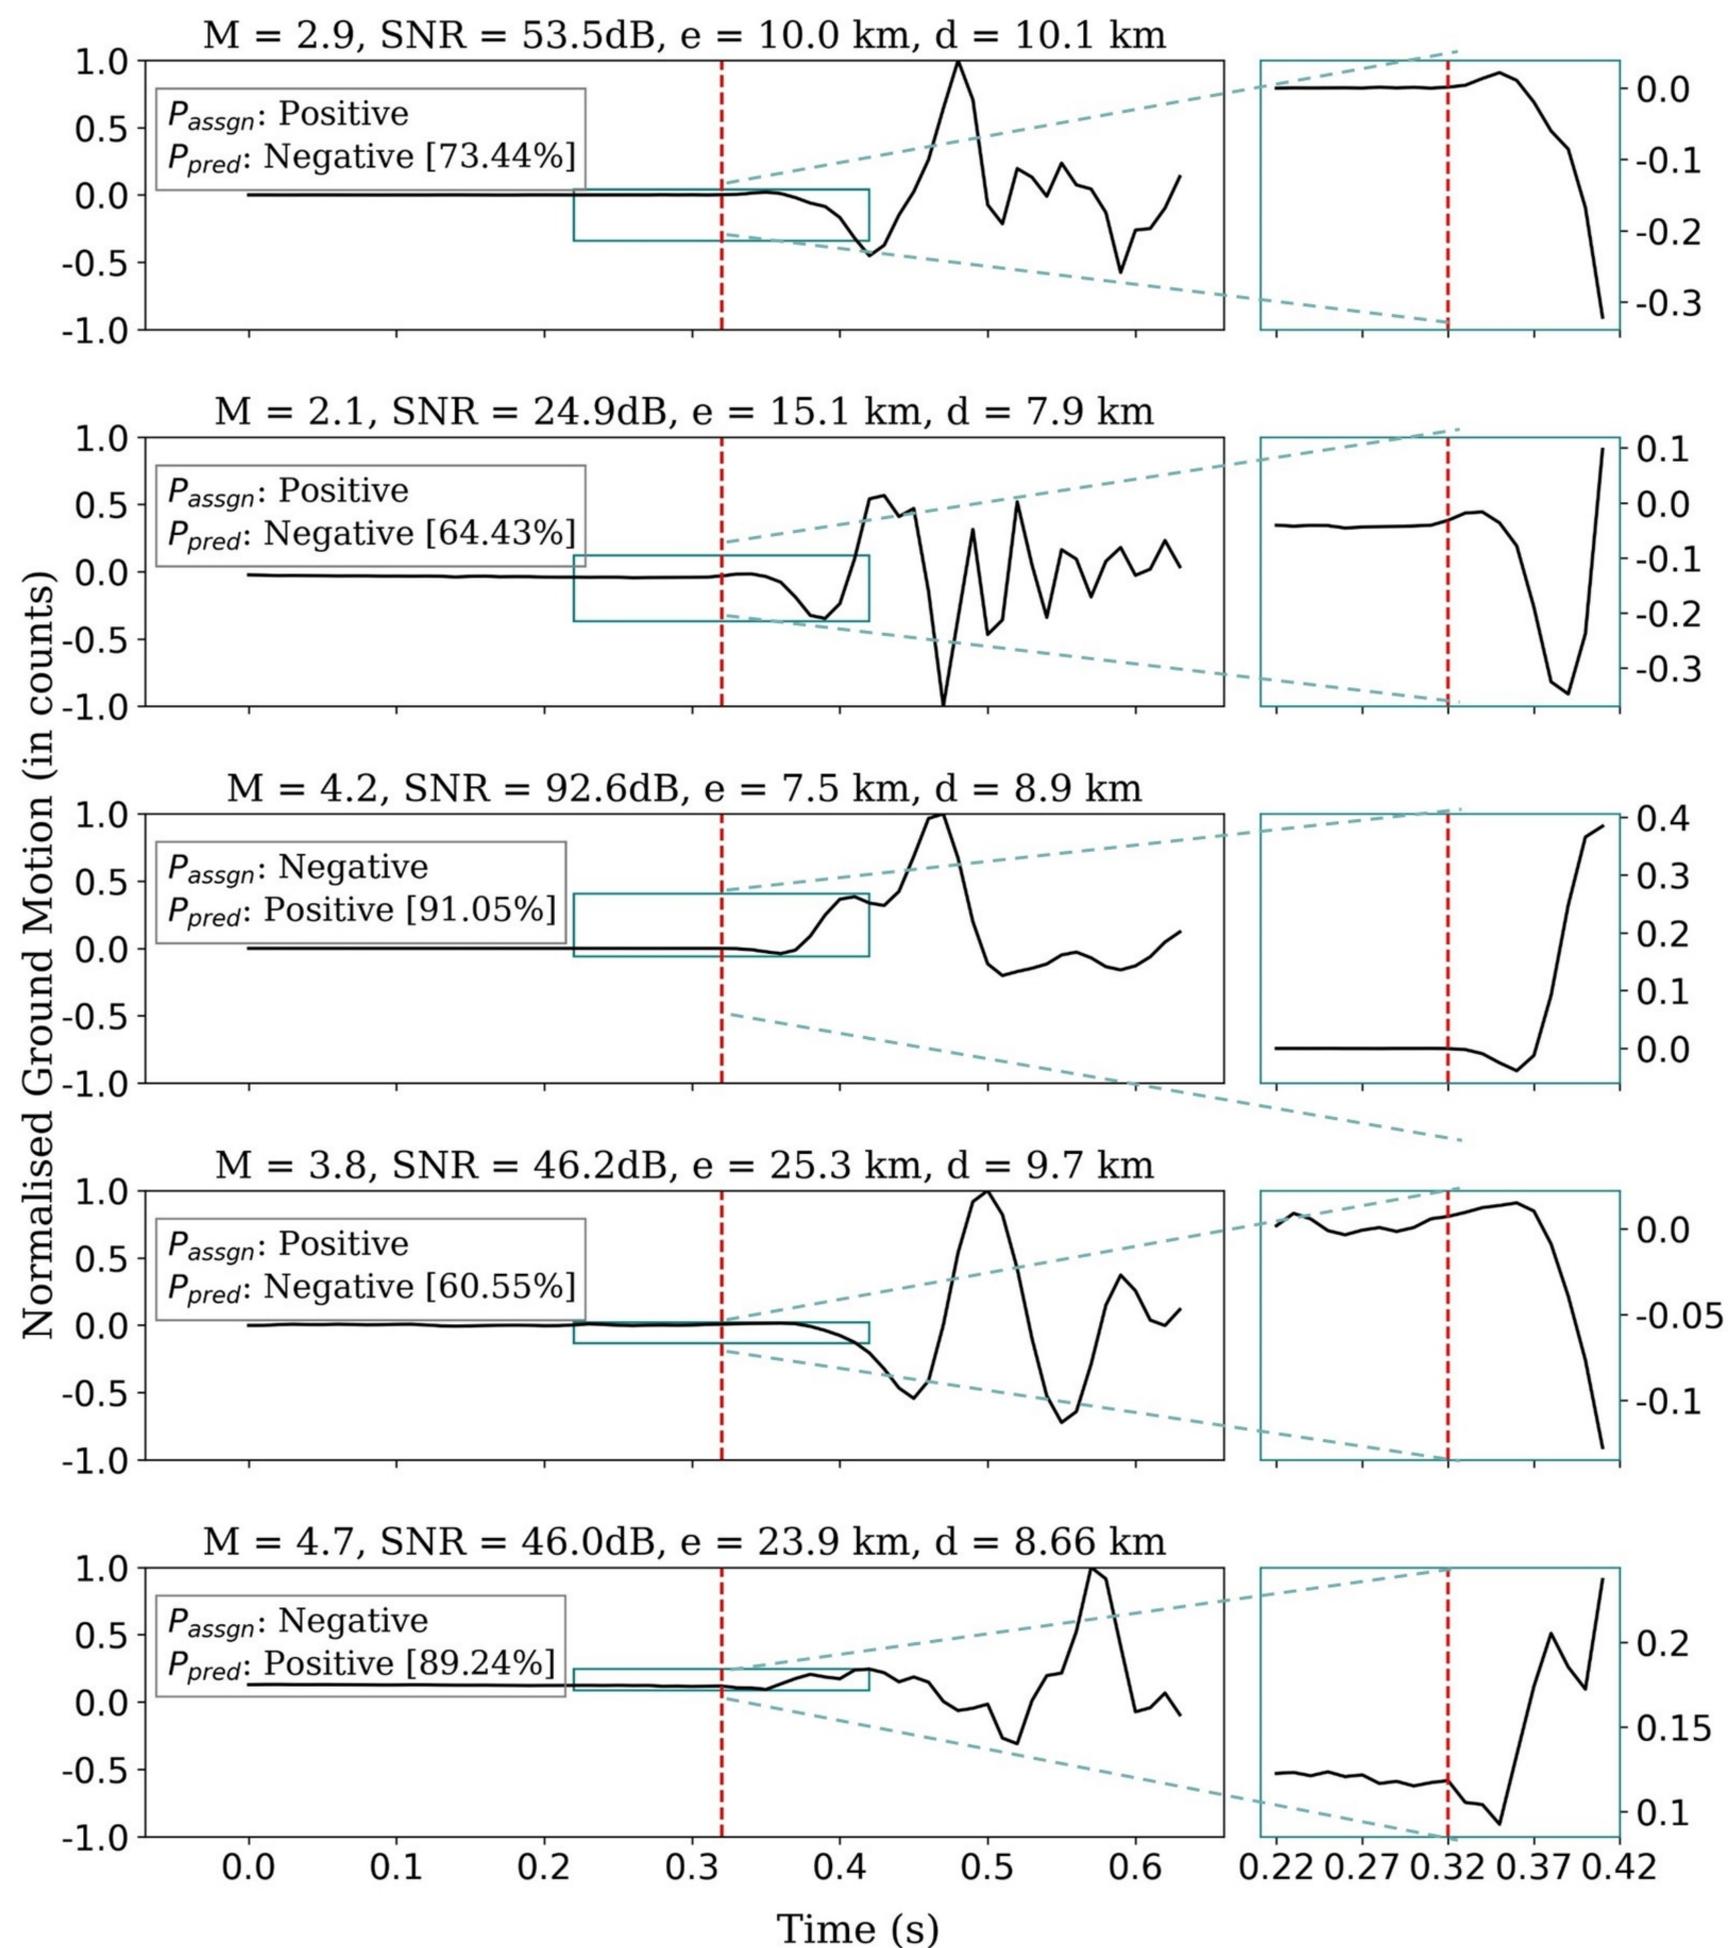


Figure\_4.jpg.

(a)

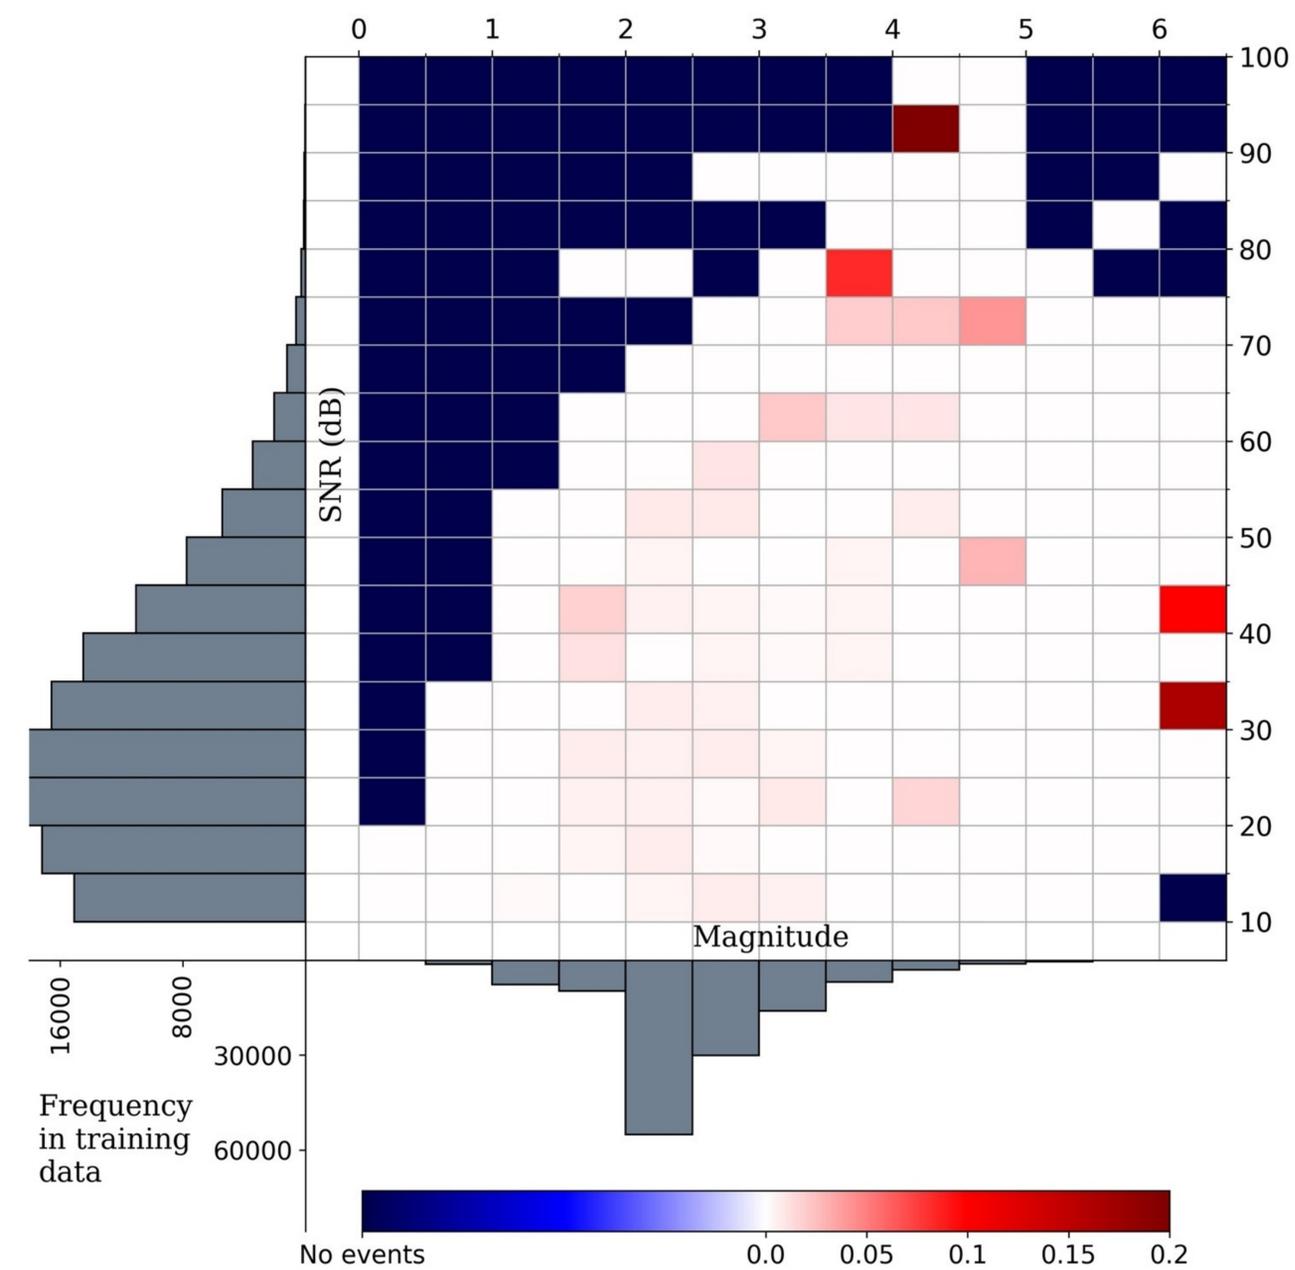


(b)

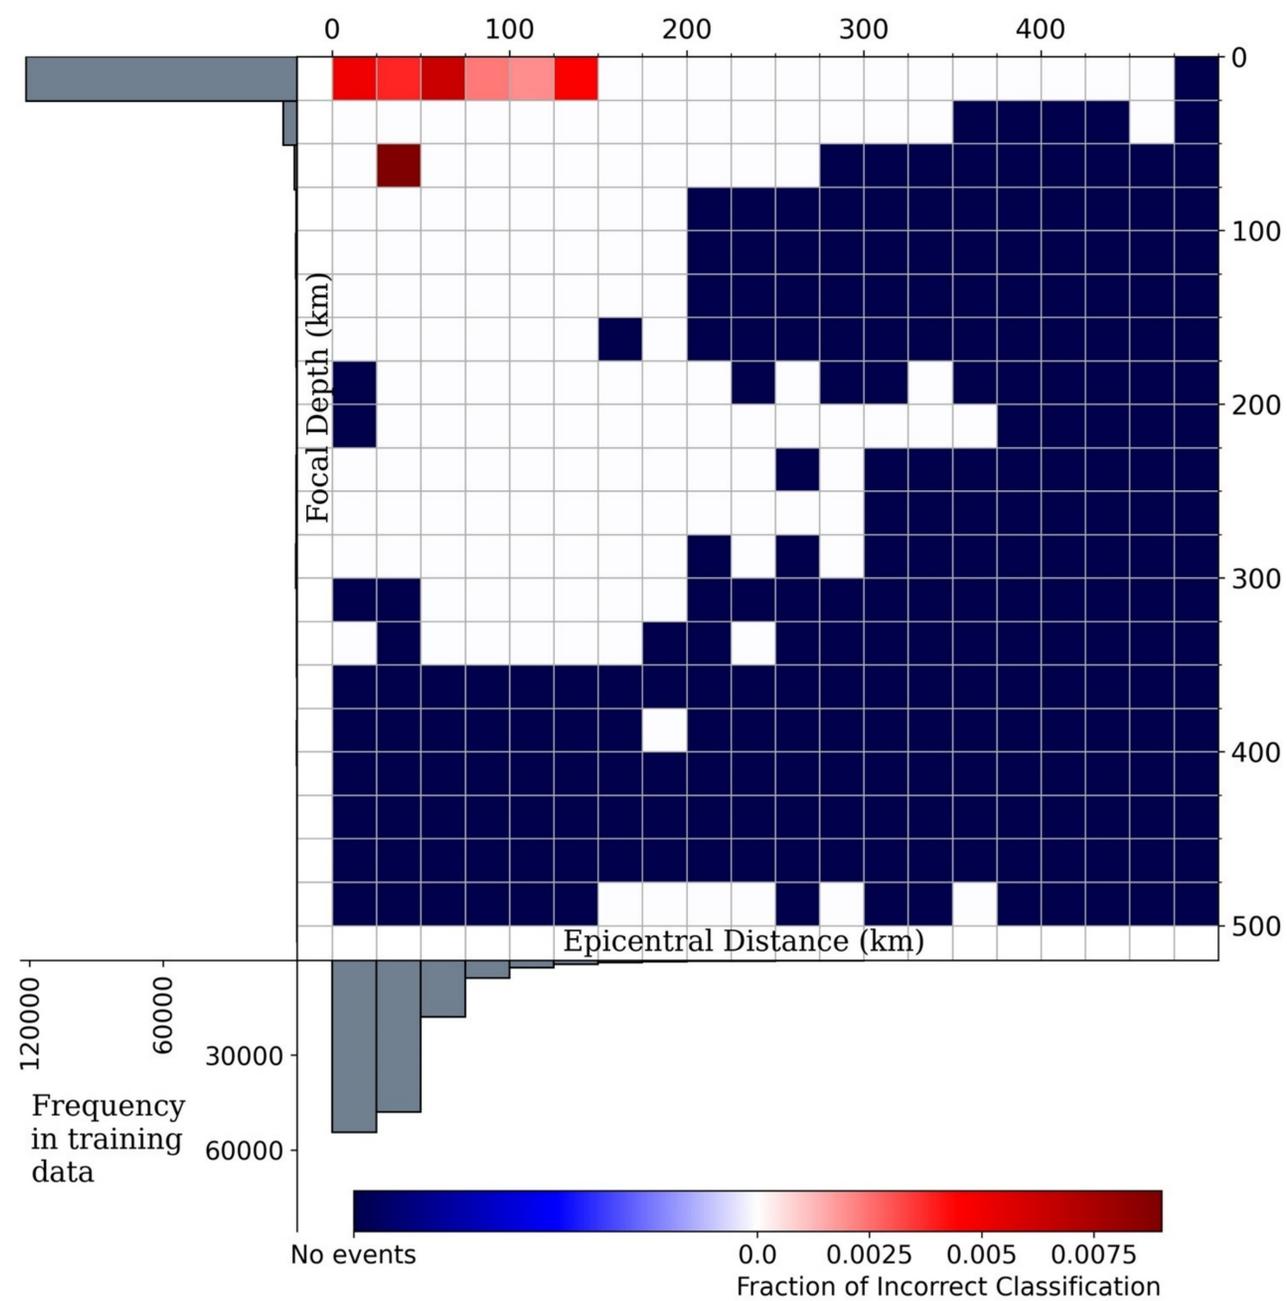


Figure\_5.jpg.

(a)



(b)



(c)

