Geophys. J. Int. (2025) **1**, 1–15

Leveraging synthetic data for deep learning denoising and prediction of measured earthquake waveforms

Alexander Bauer[1] and Conny Hammer[1]
[1]*Institute of Geophysics, University of Hamburg, Bundesstraße 55, D-20146 Hamburg (Germany). E-mail: alex.bauer@uni-hamburg.de*

Received 30 October 2025; in original form 30 October 2025

SUMMARY
Single-station recordings of teleseismic earthquakes are inherently complex due to the superposition of numerous seismic phases and their contamination with noise, which can be particularly problematic in urban environments. A detailed knowledge of the wavefield generated by teleseismic earthquakes is critical for high-precision research facilities like those involved in photon science or gravity wave detection. However, seismological stations are often sparsely distributed, especially in regions with low seismic hazard, such as Northern Germany. To address this, we introduce a deep learning approach to predict low-frequency earthquake waveforms at arbitrary locations within and around the Hamburg metropolitan area. Our aim is to train a convolutional neural network (CNN) to predict measured earthquake waveforms from synthetic ones. For this, we use measured earthquake waveforms from a seismic station close to Hamburg that has been in operation for almost 30 years. However, using the measured earthquake waveforms as labels requires denoising them first. Hence, we propose a two-step strategy: first, we generate noise-free synthetic waveforms, add artificial noise that emulates the station's actual noise characteristics, and train a first CNN to denoise them. Second, we apply the first CNN to the noisy measured waveforms to obtain noise-free labels. With these, we train a second CNN to translate synthetic waveforms into the noise-free measured ones. Applications of the second CNN to earthquakes not included in the training dataset show that this method effectively predicts measured waveforms not only for the training station but also for stations in and around Hamburg. This approach represents a significant step towards accurately modeling the seismic wavefield in three dimensions without the need for densely distributed seismological stations.

Key words: earthquake ground motions; earthquake interaction, forecasting, and prediction; seismic noise; machine learning; neural networks, fuzzy logic; site effects

1 INTRODUCTION

The wavefield generated by teleseismic earthquakes is a highly complex superposition of direct and reflected or scattered arrivals of body and surface waves. While a detailed knowledge of this wavefield is naturally important for hazard assessment in many regions of the Earth, in regions with low seismic hazard it is still crucial for the calibration of high-precision measurements, for example in the fields of photon science and gravity wave detection, where the seismic wavefield is a noise contribution that needs to be removed. However, in regions with low seismic hazard such as Northern Germany, seismological stations are often sparsely distributed, which impedes the direct measurement of the earthquake wavefield in three dimensions. In addition, seismological stations located in urban areas measure a lot of anthropogenic noise, which complicates the extraction of the earthquake wavefield from the data. Therefore, we propose a scheme in which we leverage deep learning to predict the regional earthquake wavefield at arbitrary locations around a long-standing seismological station with a good signal-to-noise ratio.

In the recent years, numerous studies in the field of seismology have tried to leverage deep learning for earthquake analysis. Zhu & Beroza (2019) established PHASENET, a neural network to aid the picking of arrival times. Chakraborty et al. (2022b) developed CREIME, a convolutional recurrent model that jointly performs earthquake identification and magnitude estimation while Chakraborty et al. (2022a) presented POLARCAP, a deep learning approach for first motion polarity classification of earthquake waveforms. Yang et al. (2022) introduced the URBANDENOISER algorithm to remove urban anthropogenic noise from seismological data. Saad et al. (2022) have used unsupervised deep learning to denoise single-channel earthquake data. Deep learning models have also demonstrated promising results in detecting, characterizing, and locating earthquakes: Perol et al. (2018) introduced CONVNETQUAKE, a convolutional neural network that detects and locates earthquakes and Mousavi et al. (2020) have introduced the Earthquake Transformer model, which simultaneously detects earthquake signals and picks seismic phases. Kubo et al. (2024) summarize various recent machine learning applications in earthquake seismology, highlighting improvements in detection, local-

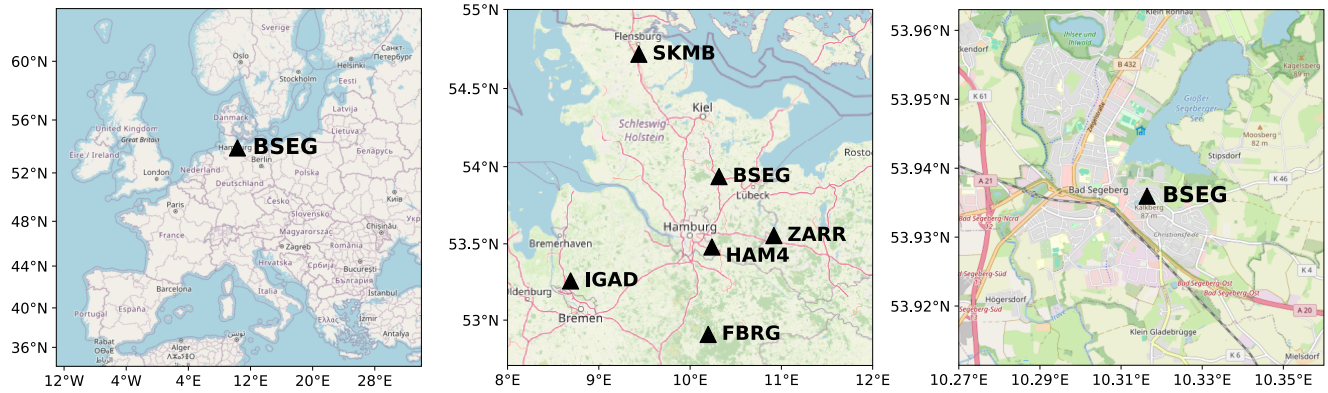

Figure 1. (left) Location of the training station BSEG in Northern Europe; (center) Existing seismological stations of the German Regional Seismic Network (GRSN) in and around Hamburg (Germany); (right) close-up on the location of the station BSEG located in Bad Segeberg.

ization, and characterization, and Mousavi & Beroza (2022) provide a comprehensive overview of the transformative potential of deep learning in seismology. However, there are only few studies, in which deep learning has been exploited to predict measured earthquake waveforms from synthetic data. Monterrubio-Velasco et al. (2024) trained a machine learning estimator on synthetic data from physics-based simulations to predict real-time earthquake ground-shaking in Southern California. Deep learning models have also been employed to predict laboratory earthquakes and forecast fault zone stress using synthetic data generated from laboratory experiments, suggesting their potential application to natural earthquakes (Laurenti et al. 2022; Wang et al. 2025).

The aim of our study is to predict the wavefield generated by teleseismic earthquakes in and around Hamburg (Germany), which hosts Deutsches Elektronen-Synchrotron (DESY) and the European X-Ray Free-Electron Laser Facility (European XFEL), two research facilities that are interested in an accurate estimation of the regional seismic wavefield to be able to remove it from their measurements. In this context, within the WAVE Initiative (wave-hamburg.eu), a few seismometers and 12 km long DAS fiber have been installed within these facilities (Bauer et al. 2025b). As a first step towards the desired 3D prediction of the earthquake wavefield, in this study we introduce a deep-learning scheme for the prediction of noise-free measured earthquake waveforms from synthetic data. Capitalizing on the versatility and portability of neural networks, we use a deep-learning framework that we have originally developed for the denoising of multi-channel seismic data (Knispel et al. 2022; Zehn et al. 2023) and wavefield decomposition for diffraction separation (Bauer et al. 2025a) and which has recently been applied to estimate the transfer function between two seismic stations close to each other (Klinge et al. 2025).

The workflow we propose is two-fold: first, we train a neural network to denoise the earthquake waveforms measured at the selected training station by means of noise-free synthetic waveforms combined with artificially generated noise that emulates the actual noise at the station. Second, we train another neural network to predict the noise-free measured earthquake waveforms from synthetic data. Although we train this second neural network on data of a specific station, it can be applied to synthetic waveforms generated for arbitrary locations in the regional vicinity of that station, thereby paving the way towards predicting the regional-scale earthquake wavefield without the need for a denser distribution of seismological stations.

2 METHODS AND WORKFLOW

The workflow we propose first requires the selection of a seismological station that fulfills two main criteria: first, it should be in operation for as many years as possible in order to be able to gather data from as many earthquakes as possible; second, the station data should ideally not be contaminated with too much noise. Fig. 1 (center) shows a map of existing seismological stations in Northern Germany, which are all part of the German Regional Seismic Network (GRSN) operated by the Federal Institute of Geosciences and Natural Resources of Germany (BGR). The station that best fulfills the mentioned criteria is the station BSEG, which has been measuring data since 1996 January 1 and is one of the longest-standing stations in Germany. Except for the station IGAD, all other stations on the map were installed after 2010, and all of them are more noisy then BSEG. Although the station is located in the small city of Bad Segeberg, it is characterized by a relatively low level of anthropogenic noise compared to other stations in Northern Germany and in addition, its distance from Hamburg is not too large. A close-up on the location of BSEG is shown on the right map in Fig. 1.

Fig. 2 sketches the proposed workflow that can be divided into two main steps (denoted as 1 and 2) and five substeps (marked with letters A–E). The goal of Step 1 is to setup a deep learning model to denoise data from the training station and the goal of Step 2 is to setup a deep learning model that can predict noise-free measured earthquake waveforms from synthetic ones. While Step 1 is entirely based on synthetic waveforms, in Step 2 we utilize both synthetic and measured waveforms. At the beginning of the workflow, we generate synthetic waveforms for a large number of past large earthquakes using INSTASEIS (van Driel et al. 2015) and the IRIS Synthetics Engine, SYNGINE (Krischer et al. 2017), at the coordinates of the training station BSEG. In substep A, we then try to characterize the noise that is usually present at the training station BSEG by analyzing data of time periods where no earthquakes occurred and trying to generate random noise with a similar spectrum as the measured noise. Using this noise generator, we then add random noise that emulates the station noise to the noise-free synthetic earthquake waveforms, thereby gathering a dataset with pairs of noisy synthetic earthquake waveforms and corresponding noise free labels. In substep B, we use this dataset to train a first neural network, which we will refer to as CNN1 in the following, to denoise the synthetic waveforms. With the trained CNN1, we proceed to Step 2 of the workflow.

For Step 2 of the workflow, we first download data from the training station BSEG for all past large earthquakes, for which data

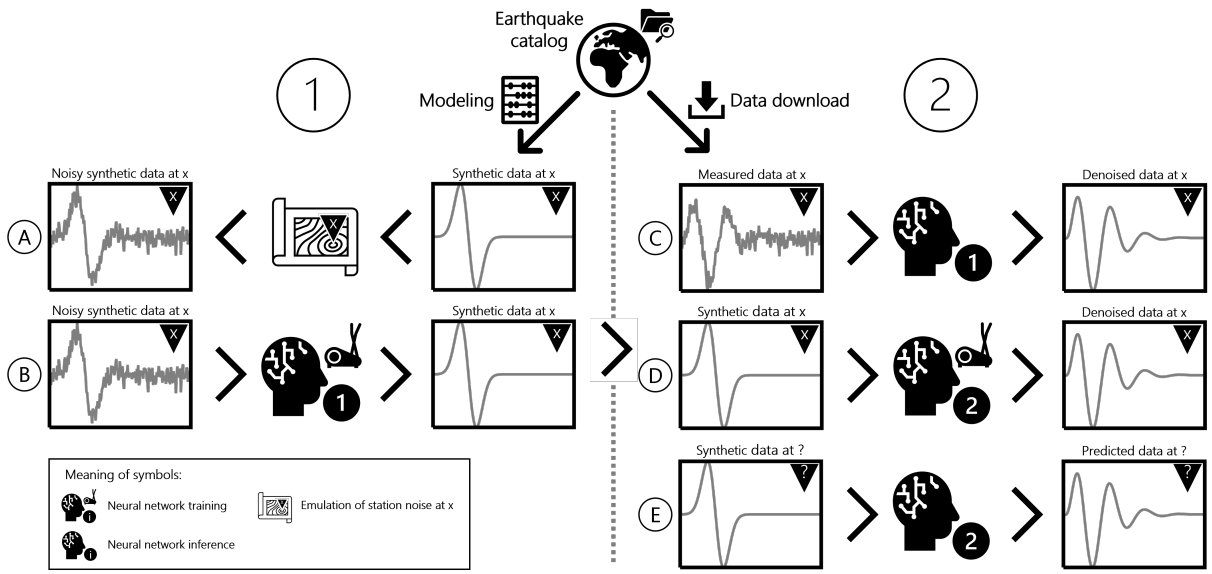

Figure 2. Sketch of the proposed two-step strategy. In the first step (left side), we contaminate synthetic earthquake waveforms with artificial noise emulating the station's actual noise characteristics (A) and train a first neural network, CNN1, to denoise them (B). In the second step of the workflow (right side), we apply CNN1 to measured earthquake waveforms to denoise them (C). We then use the denoised data as labels to train a second neural network, CNN2, to predict noise-free measured earthquake waveforms from synthetic ones (D). CNN2 can then be used to predict measured earthquake waveforms at other locations in the station's region (E).

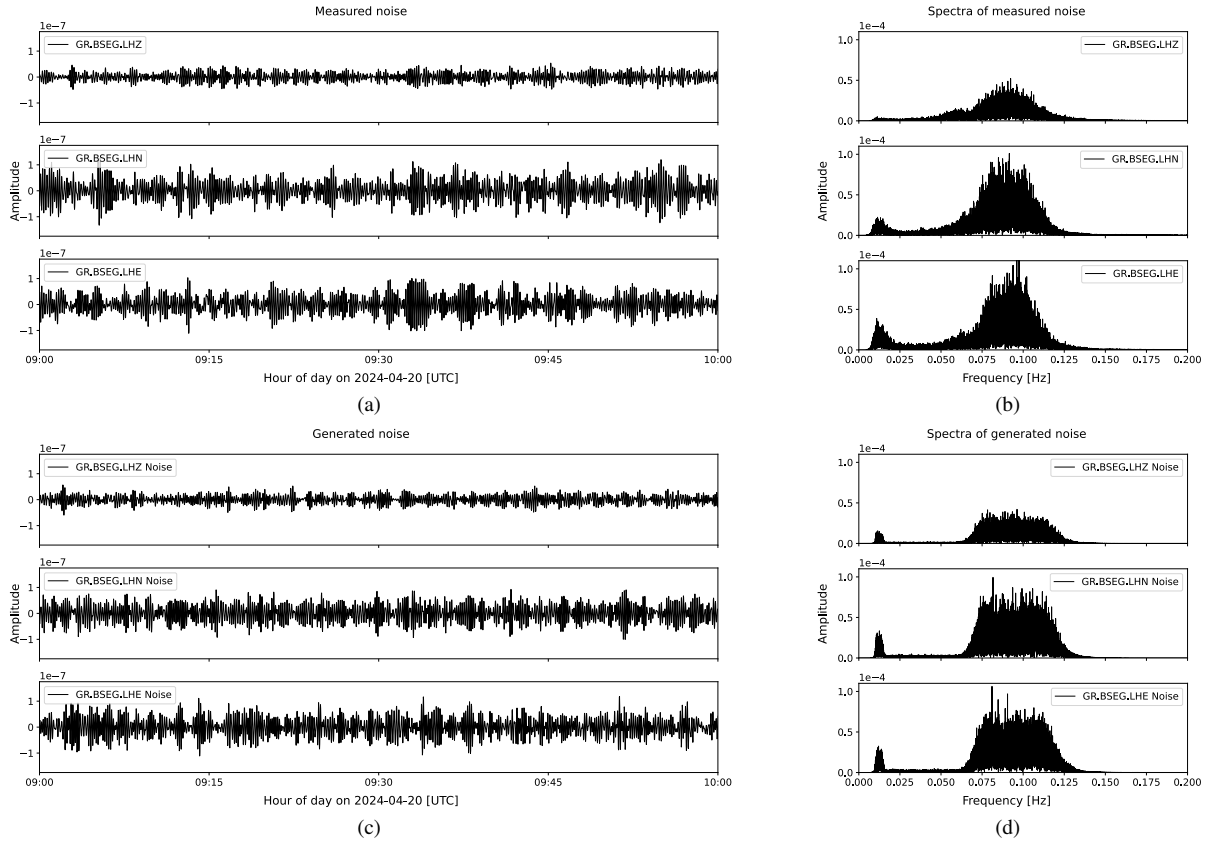

Figure 3. Noise analysis and modeling at the station BSEG. Panel (a) shows one hour of measured noise and (c) shows the generated synthetic noise based on Equation (1), here with $\alpha = 1$; (b) and (d) are the corresponding spectra of measured and generated noise calculated from 24 hours of noise.

Index i	$f_i^{\min}$ [Hz]	$f_i^{\max}$ [Hz]	Scaling factor γ_i
1	0.07	0.12	10
2	0.01	0.015	4
3	0.01	0.1	0.5

Table 1. Filter bands for synthetic noise generation.

is available and generate corresponding synthetic waveforms, thus compiling a dataset with pairs of synthetic and measured waveforms for all available earthquakes. However, it is not advisable to use the raw data from the station as labels, since it contains random noise, which would severely compromise the training process, as the model would try to learn predicting the unwanted noise. Therefore, in substep C, we apply CNN1 to the downloaded measured earthquake waveforms to denoise them. In substep D, we use these denoised measured earthquake waveforms as labels for training CNN2 to predict noise-free measured earthquake waveforms from their synthetic counterparts. The trained CNN2 can then be utilized to predict noise-free measured waveforms from synthetic waveforms generated at arbitrary locations in the region of the training station BSEG, for example in Hamburg or at any other location where measured data is either too noisy or no seismological station is installed.

In the following sections, we will describe the training process of CNN1 and CNN2 in detail and present results of the application of both models.

3 STEP 1: DENOISING

In this section, we explain the training process of the denoising neural network, CNN1, and present applications to both unseen noisy synthetic waveforms and measured data from the station BSEG.

3.1 Emulation of station noise at BSEG

For substep A, we first tried to characterize the noise that is usually present at the chosen training station BSEG by searching and analyzing a number of time periods between 1996 and 2024 where no earthquakes occurred and calculating the standard deviations of the three components of the data in each period. We used a total of 14 time periods of 24 hours, taken from different years and different seasons. Using the mean value of the calculated standard deviations, we then attempted to generate random noise with a similar spectrum as the measured noise. To achieve this, we generated three vectors of random numbers for each of the three components, scaled and bandpass-filtered them in three distinct frequency bands before summing them up:

$$\mathbf{n}_c = \sum_{i=1}^{3} \mathrm{bp}_i(\mathbf{r}_i \, \sigma_c \, \beta \, \gamma_i) \; , \tag{1}$$

where $\mathbf{n}_c$ is the noise vector corresponding to the component $c \in (\mathrm{Z}, \mathrm{N}, \mathrm{E})$, $\mathbf{r}_i$ is a vector of random numbers sampled from a normal distribution, σ_c is the standard deviation estimated from the measured noise, γ_i is a predefined factor that scales the noise in the three frequency bands, $\beta \in [0.9, 1.1]$ is a random factor that varies the noise level by ± 10 per cent, and bp_i is a bandpass filter. Table 1 summarizes the ranges of the three frequency bands i and the corresponding scaling factors γ_i. The standard deviations of the noise estimated from the station data are $\sigma_{\mathrm{Z}} = 1.6 \cdot 10^{-8}$ m/s for the vertical component and $\sigma_{\mathrm{N,E}} = 3.4 \cdot 10^{-8}$ m/s for the horizontal components.

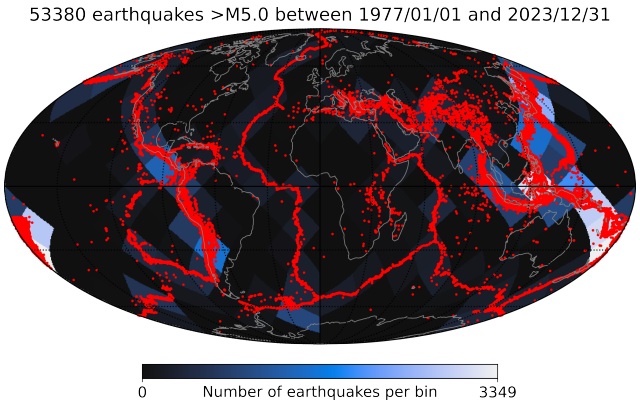

Figure 4. Geographic distribution of the 53,380 earthquakes $\geq$ M5.0 in the training data for the denoising neural network (CNN1) in Step 1 of the workflow. The 192 HEALPix bins are color-coded by the number of earthquakes that occured in them.

Fig. 3 shows an example of the noise present at BSEG and the synthetic noise generated with our approach. Fig. 3a is one hour of noise measured on 2024 April 20, a day on which no major Earthquakes occurred. The data is taken from the low-frequent L channel (1 Hz sampling frequency) and bandpass-filtered between 0.01 and 0.1 Hz. Fig. 3b is the spectrum of 24 hours of noise on the same date. Based on this spectrum and the standard deviations σ_c we adapted the factors γ_i and the filter frequencies $f_i^{\min}$ and $f_i^{\max}$ in Equation (1) to generate synthetic noise with a similar spectrum. Fig. 3d shows the resulting spectrum of the generated noise and Fig. 3c the corresponding time series. Although the generated noise is entirely artificial, it resembles very well the real noise measured at BSEG.

3.2 CNN1: Data preparation and training

Since Step 1 of the workflow is entirely based on synthetic data, we are not restricted to the measuring period of the station BSEG. To generate a dataset as large and realistic as possible, we therefore downloaded all available moment tensors for earthquakes worldwide with a minimum magnitude of 5.0 that occurred between 1977 and 2023. Fig. 4 shows the geographic distribution of the epicenters of these 53,380 earthquakes. The bins in the background are HEALPix pixels (Gorski et al. 2005) that cover an area of equal size on the Earth's surface. The bins are color-coded by the number of earthquakes that occurred in them. For all these events, we generated synthetic waveforms with INSTASEIS (van Driel et al. 2015) and the IRIS Synthetics Engine, SYNGINE, (Krischer et al. 2017) at the coordinates of the station BSEG. We used the 2 s database based on the AK135 velocity model with density and Q by Montagner & Kennett (1996) and bandpass-filtered the computed seismograms into a range of $[0.01, 0.1]$ Hz to suppress unwanted high-amplitude surface wave artifacts that we observed in certain events. The noise-free synthetic waveforms are the labels for the training of CNN1. To obtain the corresponding noisy input data for the neural network, we added simulated noise of the station BSEG to the synthetic waveforms with the random noise generator we implemented based on Equation (1). We repeated this five times to augment the training data to a total of 266,900 events, such that each event is

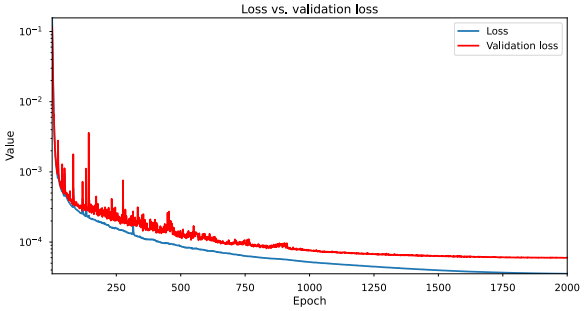

Figure 5. Loss and validation loss curves of the training of the denoising neural network, CNN1.

present five times in the labels, but the corresponding noisy input differs.

Each earthquake waveform in the training dataset contains the three components (Z, N, E) with 3584 time samples at 1 Hz sampling frequency, resulting in a data size of $(266900, 3584, 3)$, which amounts to 11 GB of data each for input and labels. For the neural network, we have normalized each input-label pair into the range $[-1, 1]$ by applying a logarithmic transformation,

$$\mathbf{s}(t)_{\text{norm}} = \frac{\text{sign}(\mathbf{s}(t)) \log(1 + |\mathbf{s}(t)|)}{\max |\mathbf{s}(t)|} , \tag{2}$$

where $\mathbf{s}(t)$ is a three-component seismogram. In contrast to conventional min-max scaling or zero-mean normalization (standardization), the logarithmic transformation emphasizes smaller values in the data. We have chosen this type of normalization to aid the preservation of smaller-amplitude phases in the seismograms. Due to the large dynamic range of amplitudes among all earthquake waveforms, a global normalization of the entire dataset is not advisable, because this would result in the majority of the values concentrating around zero. We therefore normalized each input-label pair individually, but maintained the amplitude ratio between them by scaling both with their common maximum absolute value.

The neural network architecture we used for the denoising neural network, CNN1, is based on DIFFRACTIONET (Bauer et al. 2025a), since we use the same deep-learning framework for this study. However, in contrast to DIFFRACTIONET, here we have not used the classical U-Net architecture (Ronneberger et al. 2015), but a version enhanced with ResNeXt blocks (He et al. 2016; Xie et al. 2017), which performed better in the test runs we carried out. The neural network has a depth of four and each depth level contains two convolutional layers. In addition, we implemented a dense layer with a dimension of 256 at the bottleneck followed by a dropout layer with a dropout rate of 0.2. The shape of both input and output is $(3584, 3)$, where 3584 is the number of time samples per earthquake waveform and 3 is the number of components. We use the scaled exponential linear unit (SELU) activation function (Klambauer et al. 2017) given by

$$\text{SELU}(x) = \lambda \begin{cases} x & \text{if } x > 0, \\ \alpha(e^x - 1) & \text{if } x \leq 0, \end{cases} \tag{3}$$

with $\alpha \approx 1.67$ and $\lambda \approx 1.05$ (Klambauer et al. 2017). The main advantage of SELU over the commonly used ReLU activation function is that it does not decay to zero for $x \leq 0$ and can therefore help to avoid the dying neuron issue (Lu et al. 2019) that has been observed for the ReLU activation function. In the final layer of the neural network we use the tanh activation function. For training the network, we chose the mean squared error (MSE) loss function.

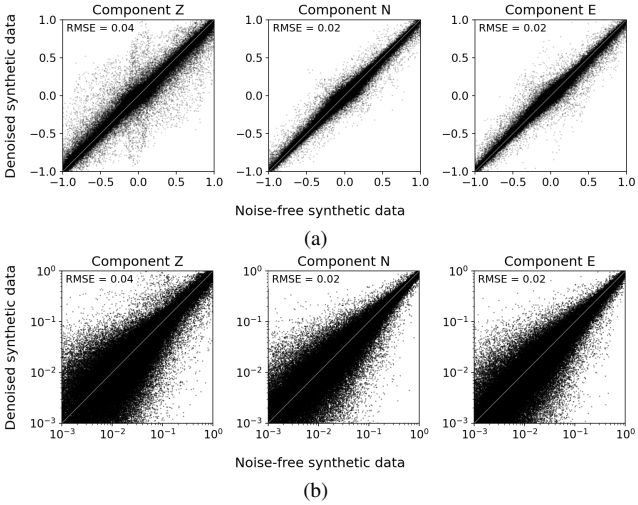

Figure 6. Application of CNN1 to unseen noisy synthetic earthquake waveforms. Comparison of amplitudes of noise-free synthetic (reference) versus denoised synthetic data for all three components, plotted with (a) normal and (b) logarithmic axes. We normalized the amplitudes of each event into the range $[-1, 1]$ before plotting. In the top left corner of each plot the root mean squared error (RMSE) between denoised synthetic data and the noise-free reference is shown.

We trained CNN1 for 2000 epochs with an initial learning rate of $\eta_0 = 10^{-4}$ and a decay schedule given by

$$\eta(i) = \eta_0 \cdot d^{\lfloor i/(10n) \rfloor} , \tag{4}$$

where $\eta(i)$ is the learning rate at epoch i, n is the number of batches in the training data, $d = 0.96$ is a decay factor and $\lfloor x \rfloor$ is the floor function. Using this schedule, the learning rate is multiplied by the decay factor d every 10 epochs. We used 80 per cent of the data for training and 20 per cent for validation, resulting in 835 batches of 256 earthquake waveforms each. Training CNN1 on a single NVIDIA Quadro GV100 GPU with 32 GB of memory took 842 hours (approximately 35 days). Fig. 5 shows the loss and validation loss curves that are both characterized by a consistent decay, indicating a successful training. In the following two subsections, we present applications to unseen noisy synthetic earthquake waveforms and to the measured data at BSEG, respectively.

3.3 Applications of CNN1

Since the training dataset only contains earthquakes until the end of 2023, we could use the earthquakes that occurred in 2024 for testing the performance of the trained CNN1. For this, we generated synthetic waveforms for 217 earthquakes with magnitudes larger than 5.0 in 2024, added random noise following Equation (1) and applied the trained CNN1 to them. Fig. 6 shows plots of labels (the noise-free synthetic waveforms) versus predictions (the waveforms denoised by CNN1) for all three components of all events combined, plotted with normal (Fig. 6a) and logarithmic axes (Fig. 6b). The plots with normal axes show a clear alignment of the amplitudes along the line from the bottom left to the top right corner, indicating an overall good performance of the neural network. The logarithmic plot reveals that the agreement between the amplitudes is better for larger values.

Fig. 7 shows the results for two selected earthquakes, one with slightly larger amplitudes than the station noise and one with lower amplitudes than the station noise. The first event (Fig. 7a) simu-

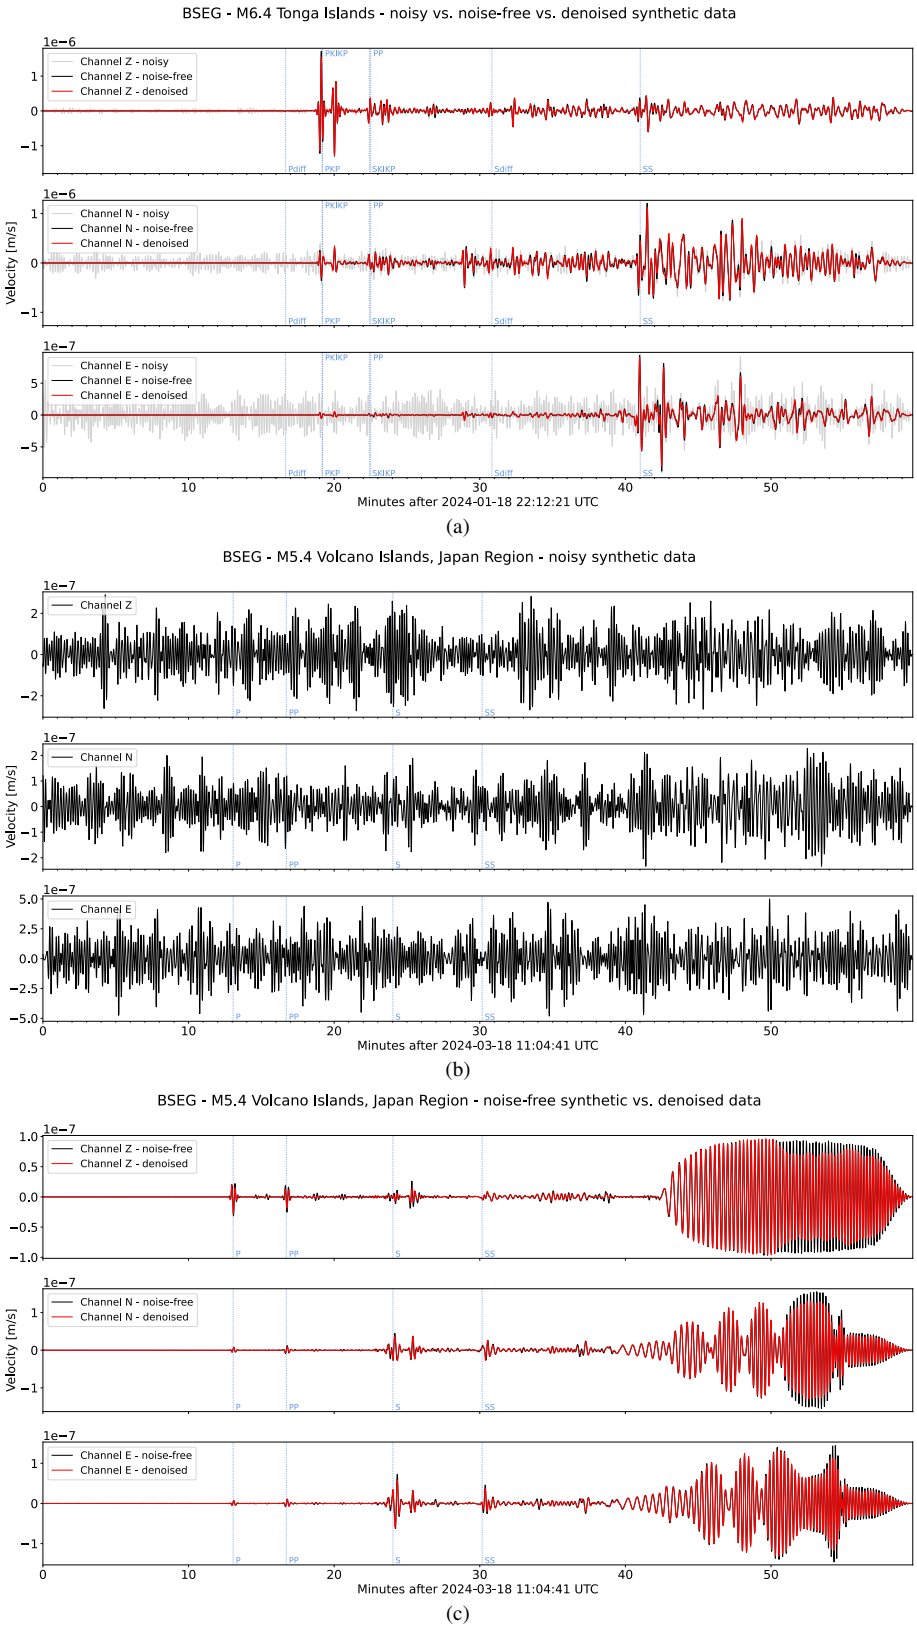

Figure 7. Application of CNN1 to unseen noisy synthetic data for two earthquakes. Panel (a) shows the simulated waveforms of a M6.4 earthquake that occurred in Tonga. The noisy synthetic input waveforms are plotted in light gray, the waveforms denoised by CNN1 in red and the noise-free labels in black; (b) shows the noisy synthetic input data for a M5.4 earthquake that occurred near the Volcano Islands in the Japan region and (c) the data denoised by CNN1 (red) along with the noise-free reference (black). Note the unrealistic surface wave arrivals modeled by INSTASEIS. Each plot contains the three components (Z, N, E). The light blue dotted lines mark the expected arrival times of a few prominent phases.

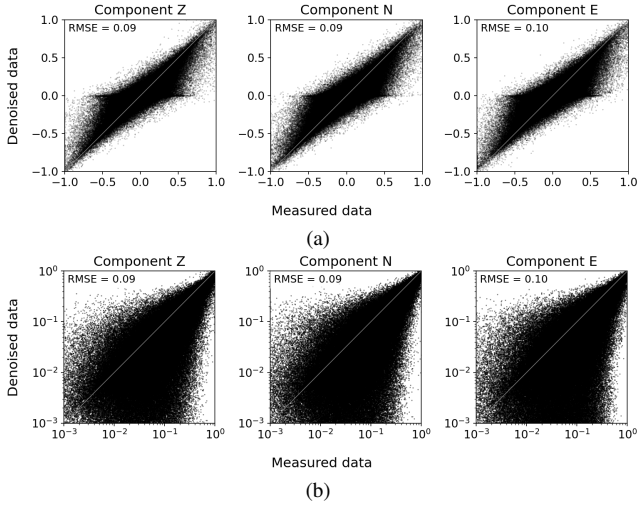

Figure 8. Application of CNN1 to noisy measured earthquake waveforms at BSEG. Comparison of amplitudes of measured versus denoised data for all three components, plotted with (a) normal and (b) logarithmic axes. We normalized the amplitudes of each event into the range $[-1, 1]$ before plotting. In the top left corner of each plot the root mean squared error (RMSE) between measured and denoised data is shown.

lates a magnitude 6.4 earthquake that occurred in Tonga on 18 January 2024. The noisy synthetic input waveforms are plotted in light gray, the noise-free reference in black and the waveforms denoised by CNN1 in red. The light blue dotted lines mark the expected arrival times of a few prominent body wave phases, which we calculated with the TAUPY tool that is part of the OBSPY toolbox. The results show a very good agreement between the noise-free reference and the denoised data except for some minimal differences in amplitude. The second and more challenging example simulates a magnitude 5.4 earthquake that occurred near the Volcano Islands in the Japan region on 18 March 2024. Fig. 7b shows the noisy input data, in which the earthquake signatures are entirely masked by noise. Fig. 7c shows the denoised data by CNN1 (red) plotted above the noise-free reference waveforms (black). Despite the high level of noise, the results reveal a good agreement between reference and denoised waveforms. In particular, CNN1 was able to correctly denoise the P, PP and S phases, although they have considerably smaller amplitudes than the noise. Also note the unrealistic surface wave arrivals modeled by INSTASEIS.

As a next step, we applied CNN1 to measured data from the station BSEG for the same 217 earthquakes larger than magnitude 5.0 that occurred in 2024. However, note that in this case, there are no labels and consequently, it is not possible to entirely assess the quality of the neural network's predictions. Fig. 8 shows plots of the amplitudes of measured data versus denoised data for all three components of all events combined. The plots with normal axes (Fig. 8a) reveal a trend towards smaller values in the denoised data, which is a signature of the removed noise. While the logarithmic plots (Fig. 8b) confirm this by an accumulation of values towards the lower right, both plots also show a good agreement for large values.

Fig. 9 shows the results of an application of CNN1 to the measured data from BSEG for the same two earthquakes as in Fig. 7. As before, the dotted light blue lines mark the expected arrival times of a few prominent body wave phases. For the M6.4 Tonga earthquake (Fig. 9a), the station data is plotted in black and the denoised data in red. In this case, the amplitudes are considerably larger than the station noise, but a good denoising performance can be observed particularly within the first 20 minutes before the first arrivals. Where the amplitudes allow an assessment of the results, a very good agreement between the phases is evident. For the M5.4 Volcano Islands earthquake, an assessment of the denoising performance is more complicated, because there is no reference. Fig. 9b shows the data measured at BSEG. While the surface waves are visible, the smaller-amplitude body wave phases are hidden behind the noise. Fig. 9c is the data denoised by CNN1. The phases of the surface waves show a good agreement between measured and denoised data. A closer look at the waveforms reveals that CNN1 also succeeded to extract the P and S phases from the data despite their low amplitudes. Considering the differences between the noise-free synthetic waveforms and the actual data of the same event (compare the unrealistic surface wave arrivals in Fig. 7c), it is indeed remarkable that the neural network manages to denoise the measured data that successfully. Although CNN1 naturally struggles to extract lower-amplitude phases if the noise level is very high, in general the results reveal a very promising denoising performance even on measured data, despite the reduced complexity of the synthetic waveforms and the relative simplicity of our noise emulation approach compared to the complex noise patterns observed in actual measurements.

4 STEP 2: EARTHQUAKE WAVEFIELD PREDICTION

The goal of the second step of our workflow is to train a second neural network, CNN2, to predict noise-free measured earthquake waveforms from synthetic data at the station BSEG (substep D). This step requires CNN1 for the generation of noise-free labels (substep C). In the following subsections, we explain the training process of CNN2 and present applications to unseen earthquakes at the training station BSEG and to a set of stations at other locations across Northern Germany (substep E).

4.1 CNN2: Data preparation and training

For Step 2 of the workflow, we initially considered all 41,934 earthquakes larger than magnitude 5.0 with moment tensors available on *The Searchable Product Repository* (SPUD) offered by IRIS-DMC in the period 1996–2024. After checking for possibly overlapping earthquakes, that is, event pairs with a difference in source time $\Delta t_0 < 4000$ s (for this, we also considered earthquakes with magnitudes ≥ 4.0 and ≤ 5.0) and skipping events where no data is available at the station BSEG or it has gaps, we downloaded a total of 16,931 earthquake waveforms from the low-frequent L channel (1 Hz sampling frequency) and computed the corresponding synthetic waveforms with INSTASEIS (van Driel et al. 2015) using the 2 s SYNGINE (Krischer et al. 2017) database based on the AK135 velocity model with density and Q (Montagner & Kennett 1996). We bandpass-filtered both measured and synthetic data between 0.01 and 0.1 Hz and subsequently applied the neural network trained in the first step, CNN1, to the waveforms to denoise them. Afterwards, we performed a number of automated checks on the denoised waveforms. We sorted out events with too low amplitudes, a too large amplitude difference between synthetic and measured data and spurious signatures before the first arrival by rejecting all

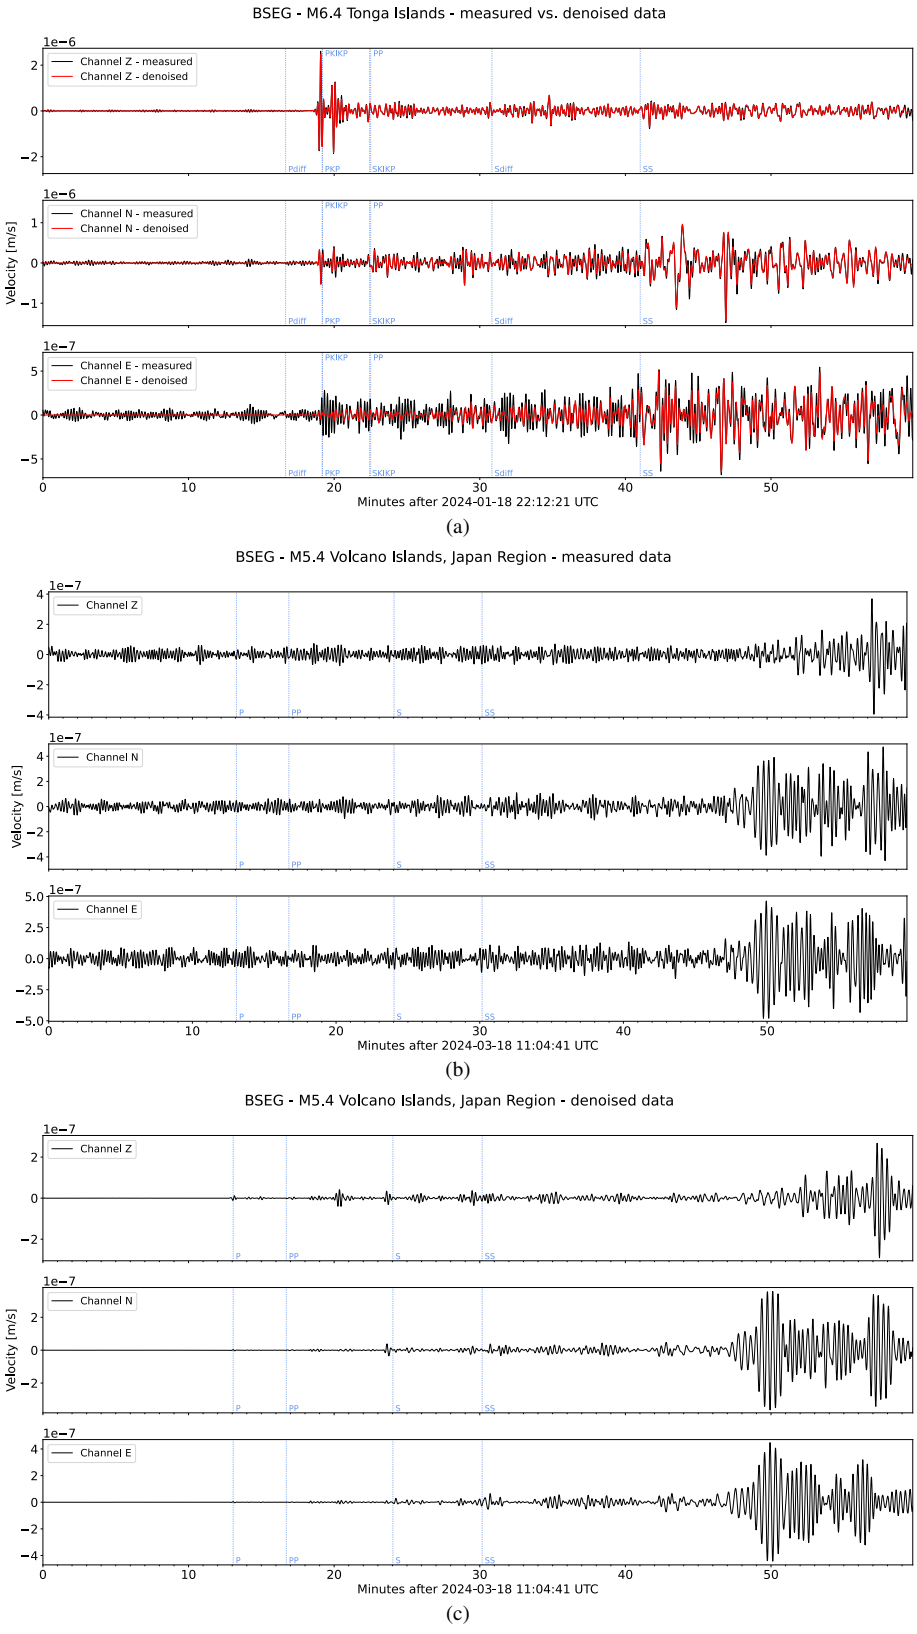

Figure 9. Application of CNN1 to data of the same two earthquakes as in Fig. 7 measured at BSEG. Panel (a) shows the waveforms of a M6.4 earthquake that occurred in Tonga. The measured waveforms are plotted in black and the waveforms denoised by CNN1 in red; (b) shows the measured data at BSEG of a M5.4 earthquake that occurred near the Volcano Islands in the Japan region and (c) shows the data denoised by CNN1. Note the large differences in the surface wave arrivals compared to the synthetic data in Fig. 7c. Each plot contains the three components $(\mathrm{Z}, \mathrm{N}, \mathrm{E})$. The measured waveforms are plotted in light gray and the data denoised by CNN1 in red. The light blue dotted lines mark the expected arrival times of a few prominent phases.

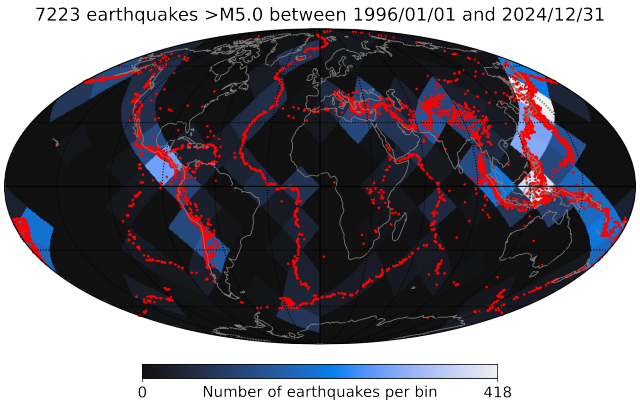

Figure 10. Geographic distribution of the 7223 earthquakes $\geq$ M5.0 in the training data for the predicting neural network (CNN2) in Step 2 of the workflow. The 192 HEALPix bins are color-coded by the number of earthquakes that occured in them.

events that fulfilled at least one of four predefined criteria:

$$\max|\tilde{\mathbf{s}}_c(t)| \quad < \quad 1\cdot 10^{-7}\,\mathrm{ms}^{-1}\,, \tag{5}$$

$$\frac{\max|\mathbf{s}_c(t)|}{\mathrm{mean}|\mathbf{s}_c(t)|} \quad < \quad 5\,, \tag{6}$$

$$\frac{\max|\mathbf{s}_c(t)|}{\max|\tilde{\mathbf{s}}_c(t)|} \quad > \quad 10\,, \tag{7}$$

$$\frac{\max|\mathbf{s}_c(t>\tilde{t}_a)|}{\max|\mathbf{s}_c(t<\tilde{t}_a)|} \quad < \quad 7.5\,, \tag{8}$$

where $\mathbf{s}_c(t)$ is a (denoised) measured seismogram in the component $c \in (\mathrm{Z,N,E})$, $\tilde{\mathbf{s}}_c(t)$ its synthetic counterpart and $\tilde{t}_a$ is the theoretical time of the first arrival calculated with the TAUPY tool.

The resulting checked dataset contains a total of 7,223 earthquakes with minimum magnitudes of 5.0. Fig. 10 shows the distribution of these earthquakes on the Earth color-coded by the number of occurrences per HEALPix bin. These denoised waveforms are the labels for the training of CNN2. Accordingly, the size of the training dataset of CNN2 is $(7223, 3584, 3)$, which amounts to 297 MB of data for each input and labels. As in Step 1, we normalized the data into the range $[-1,1]$ by applying a logarithmic transformation (Equation (2)) to each pair of waveforms. To take into account the inhomogeneous distribution of earthquakes on the Earth, we computed weights based on the HEALPix bins shown in Fig. 10. For this, we counted the earthquakes per bin and divided the maximum count by them,

$$\mathbf{w_{bins}} = \frac{\max(\mathbf{counts})}{\mathbf{counts}}\,, \tag{9}$$

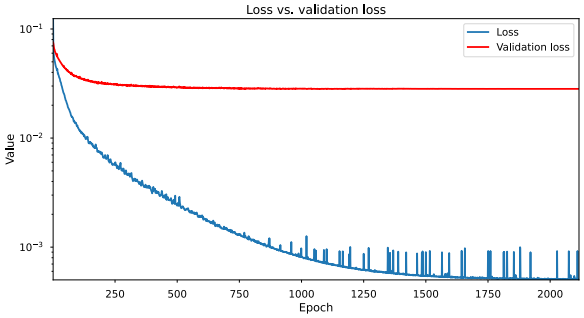

Figure 11. Loss and validation loss curves of the training of the prediction neural network, CNN2.

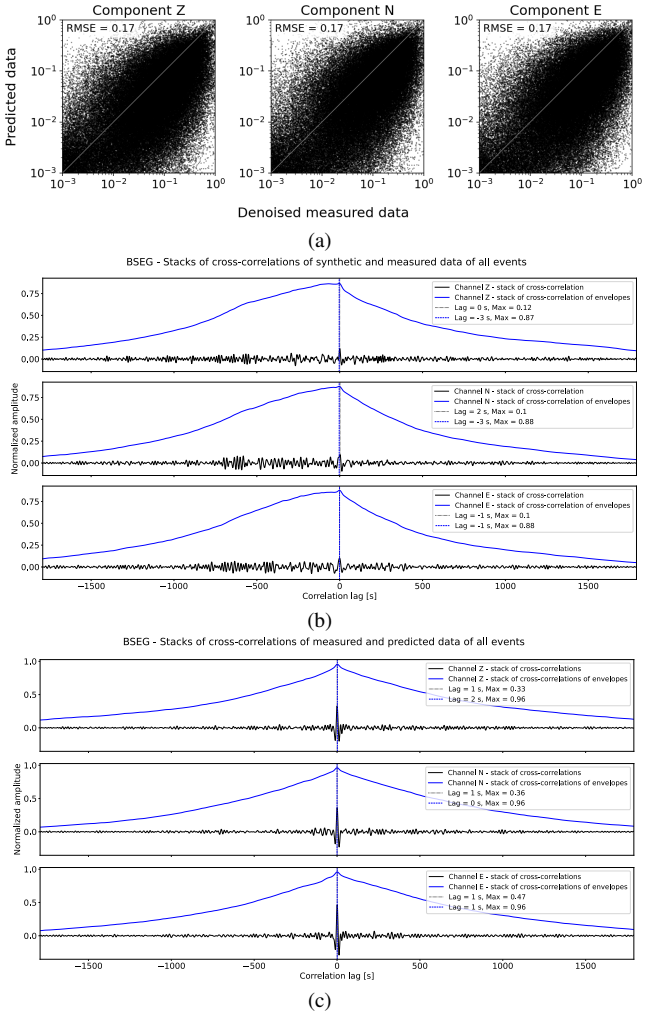

Figure 12. Prediction of noise-free measured earthquake waveforms at BSEG from synthetic data with CNN2; (a) comparison of amplitudes of denoised measured data versus predicted data, both plotted with logarithmic axes. We normalized the amplitudes of each event into the range $[-1,1]$ before plotting. The root mean square error (RMSE) between both is displayed in the upper left corners. Panel (b) shows stacks of cross-correlations (black) and stacks of cross-correlations of envelopes (blue) of synthetic input data and the denoised measured data of all 217 events for each component; (c) shows the same, but for denoised measured data and predicted data. The clearly defined maxima in (c) indicate the overall good agreement between the predicted and actual earthquake waveforms and a comparison between (b) and (c) indicates what CNN2 has learned during the training.

where **counts** is a vector with the number of earthquakes per bin and $\mathbf{w_{bin}}$ is the resulting vector with the weights for all n_{bins} HEALPix pixels (192 in our case) that each cover an area of equal size on the Earth's surface. Then, we assigned the corresponding weight to each event in the dataset, resulting in a vector $\mathbf{w_{events}}$ with a weight for each event, and finally ensured that the sum of all weights is equal to the number of events,

$$\mathbf{w_{final}} = \frac{n_{events}}{\sum \mathbf{w_{events}}}\,, \tag{10}$$

where $\mathbf{w_{final}}$ is the final vector of weights and n_{events} is the number of events in the dataset (7223 in our case). We used these weights during the neural network training such that events from regions with less frequent occurrence of earthquakes had a higher impact on the value of the loss function.

In contrast to CNN1, for CNN2 we did not use ResNeXt blocks, but a classical U-Net architecture (Ronneberger et al. 2015) with the same stack of convolutional layers as DIFFRACTIONET (Bauer et al. 2025a): the neural network has a depth of 3 with one convolutional layer in the network head and a stack of convolutional layers of $(2, 4, 1)$ in the deeper layers. As for CNN1, we also implemented a dense layer with 256 neurons followed by a dropout layer with a dropout rate of 0.2 at the network's bottleneck. We used the LeakyReLU activation function, a linear activation in the output layer and the mean squared error (MSE) loss function. We used an initial learning rate of 10^{-4} and an adaptive learning rate schedule that decreases the learning rate by a factor of 0.96 if the validation loss does not decrease for 10 epochs. Furthermore, we implemented an early stopping criterion that stops the training once the training loss does not decrease for 100 consecutive epochs. The training ran for 2116 epochs on a single NVIDIA Quadro GV100 GPU with 32 GB of memory and took almost 10 hours. Fig. 11 shows the loss and validation loss curves of the training process. Compared to the training of CNN1 (Fig. 5), the divergence of the loss curves is more pronounced, which we attribute to three main reasons. First, the higher complexity of the task of CNN2; predicting complex measured earthquake waveforms from their simpler synthetic counterparts is a considerably more complex task than removing random noise. Second, the smaller number of earthquakes in the training data; the number of unique earthquakes in the training dataset of CNN2 is only 13.5 per cent of the number of events in the training dataset of CNN1. Third, possible imperfections in the labels of the training data of CNN2, for example events where the denoising did not work as expected, because the station data contained a different type of noise.

4.2 Applications of CNN2 to unseen earthquakes

To assess the performance of the trained CNN2, we applied it to synthetic earthquake waveforms generated for 148 earthquakes with magnitudes ≥ 5.0 that occurred between January and end of July 2025. However, since there are no reference waveforms, an assessment of the quality of the predicted waveforms is challenging. While for higher-amplitude events, we can compare the predictions directly to the station data, for lower-amplitude events the station data contains too much noise for that. We therefore denoised the station data of the 148 earthquakes with CNN1 to obtain a reference. Note, however, that the denoising itself may introduce artifacts such as spurious arrivals. Fig. 12 shows a logarithmic plot of the amplitudes of the denoised measured data against the predicted data (Fig. 12a). The overall spread of the amplitudes is clearly more pronounced than in the applications of CNN1, which we mainly attribute to the higher complexity of the task of CNN2 and the smaller number of events in the training data. Nevertheless, amplitudes still tend to be centered around the line from the lower left to the top right corner. Since the simple spread of amplitudes is not too meaningful in this case, a different way to assess the accuracy of the predictions can be achieved by cross-correlating the predicted waveforms with the denoised measured data and stacking them. Fig. 12c shows this stack of the cross-correlations of predicted and denoised measured data (black) along with stacks of the cross-correlations of the envelopes of both (blue). All curves show distinct maxima at the zero-lag (marked by the vertical lines) indicating an overall good agreement of the predicted waveforms and the denoised measured data across the 148 events. As a comparison, Fig. 12b shows the same stacks of cross-correlations between the synthetic input data and the denoised measured data, where the lack of clear maxima in the stacks of cross-correlations and the asymmetric shape of the stacks of the cross-correlated envelopes reveal the difference between the simpler synthetic waveforms and the more complex measured ones, and in turn, a comparison between Figs. 12c and 12b indicates what CNN2 has learned during the training process.

Fig. 13 shows results obtained with CNN2 for a higher-amplitude earthquake, a magnitude 6.2 earthquake that occurred at the Southern Mid-Atlantic Ridge on 25 June 2025. Fig. 13a is the synthetic input data computed at the coordinates of BSEG (note again the unrealistic surface wave arrivals), Fig. 13b is the data measured at BSEG and Fig. 13c is the data predicted by CNN2 based on the synthetic input. While a comparison between the synthetic input with its unrealistically modeled surface waves (Fig. 13a) and the measured data (Fig. 13b) exhibits the higher complexity of the measured data, a comparison of the measured data with the predicted waveforms (Fig. 13c) reveals a very good agreement between both along the majority of the observed phases of both body and surface waves. In addition, the waveforms predicted by CNN2 are noise-free, which can be observed in the part prior to the first arrival.

As a more challenging example, Fig. 14 shows the results for a lower-amplitude event, a magnitude 5.4 earthquake that occurred in Sulawesi, Indonesia on 14 March 2025. In this case, Fig. 14a is a combined plot of the synthetic input data (light gray), the measured data at BSEG (black) and the data predicted by CNN2 (red). In this example, the station data has a noise level similar to the amplitudes of the surface waves. Moreover, a large deviation between measured and synthetic waveforms is again evident in the surface wave arrivals. Still, a good agreement between measured and predicted data can already be observed on the arrivals after approximately 23 minutes in Fig. 14a. Fig. 14b shows a comparison of the station data denoised with CNN1 (black) and the predicted data (red), which reveals an overall good agreement between reference and prediction also for the low-amplitude PP arrival. Fig. 14c is a close-up on the earlier arrivals, comparing synthetic input (light gray), denoised measured data (black) and predicted data (red), which confirms the previous findings.

4.3 Wavefield prediction at other locations

To be able to predict the wavefield generated by teleseismic earthquakes in three dimensions on a regional scale, the neural network has to be able to predict measured earthquake waveforms not only at the training station, but at arbitrary coordinates in the region of the training station. To test this, we have generated synthetic waveforms for a magnitude 6.4 earthquake that occurred in the Drake Passage between South America and Antarctica at the coordinates of the six stations in Northern Germany that are shown on the map in the center of Fig. 1 and applied CNN2 to these waveforms. Fig. 15 shows plots that compare synthetic input data (light gray), measured data (black) and predicted data (red) for all six stations, zoomed into the time range that covers the body wave arrivals. For the station HAM4 located in the East of Hamburg (Fig. 15b), the results show both a good agreement of the predicted phases and the strong anthropogenic station noise. Also the results for the stations SKMB to the North of Hamburg (Fig. 15c) and FBRG to the South of Hamburg (Fig. 15e) show a similarly good agreement of measured and predicted data, while the stations ZARR to the East of Hamburg (Fig. 15d) and IGAD to the West of Hamburg (Fig. 15f) exhibit a worse signal-to-noise ratio, but most of the visible phases still align with the predicted data. Considering that

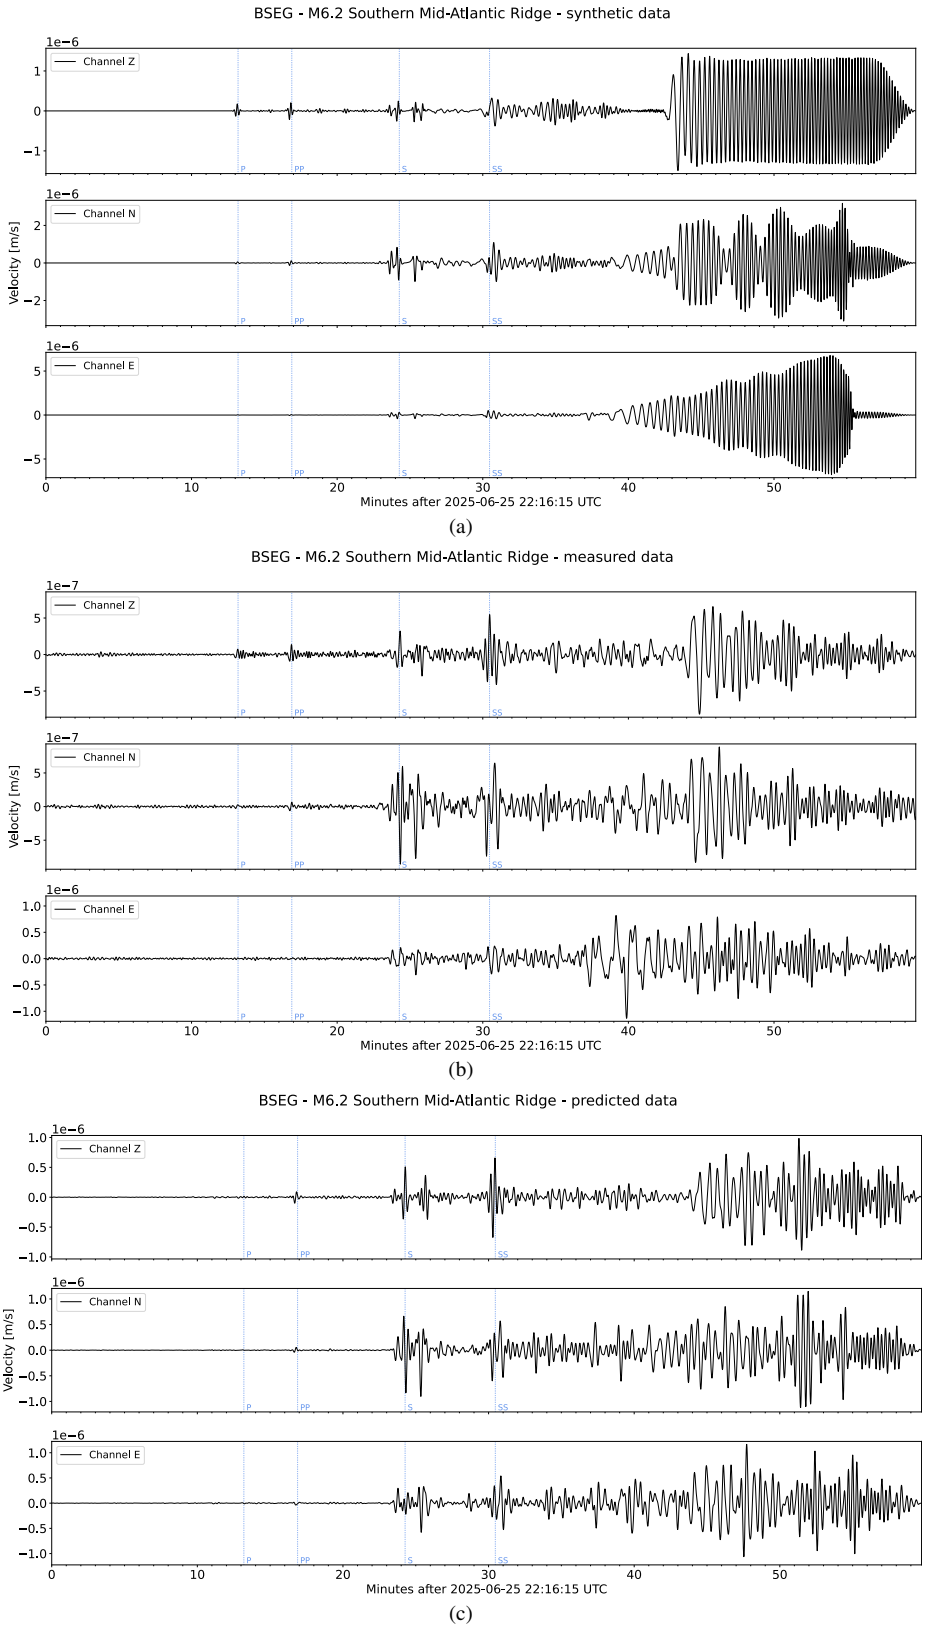

Figure 13. Prediction of noise-free measured earthquake waveforms at BSEG from synthetic data with CNN2. Example for a higher-amplitude event, a M6.2 earthquake that occurred near the Southern Mid-Atlantic Ridge on 23 June 2025. Panel (a) shows the synthetic input data (note again the unrealistic surface wave arrivals), (b) shows the measured data at BSEG and (c) shows the earthquake waveforms predicted by CNN2. Each plot contains the three components (Z, N, E). The light blue dotted lines mark the expected arrival times of a few prominent phases.

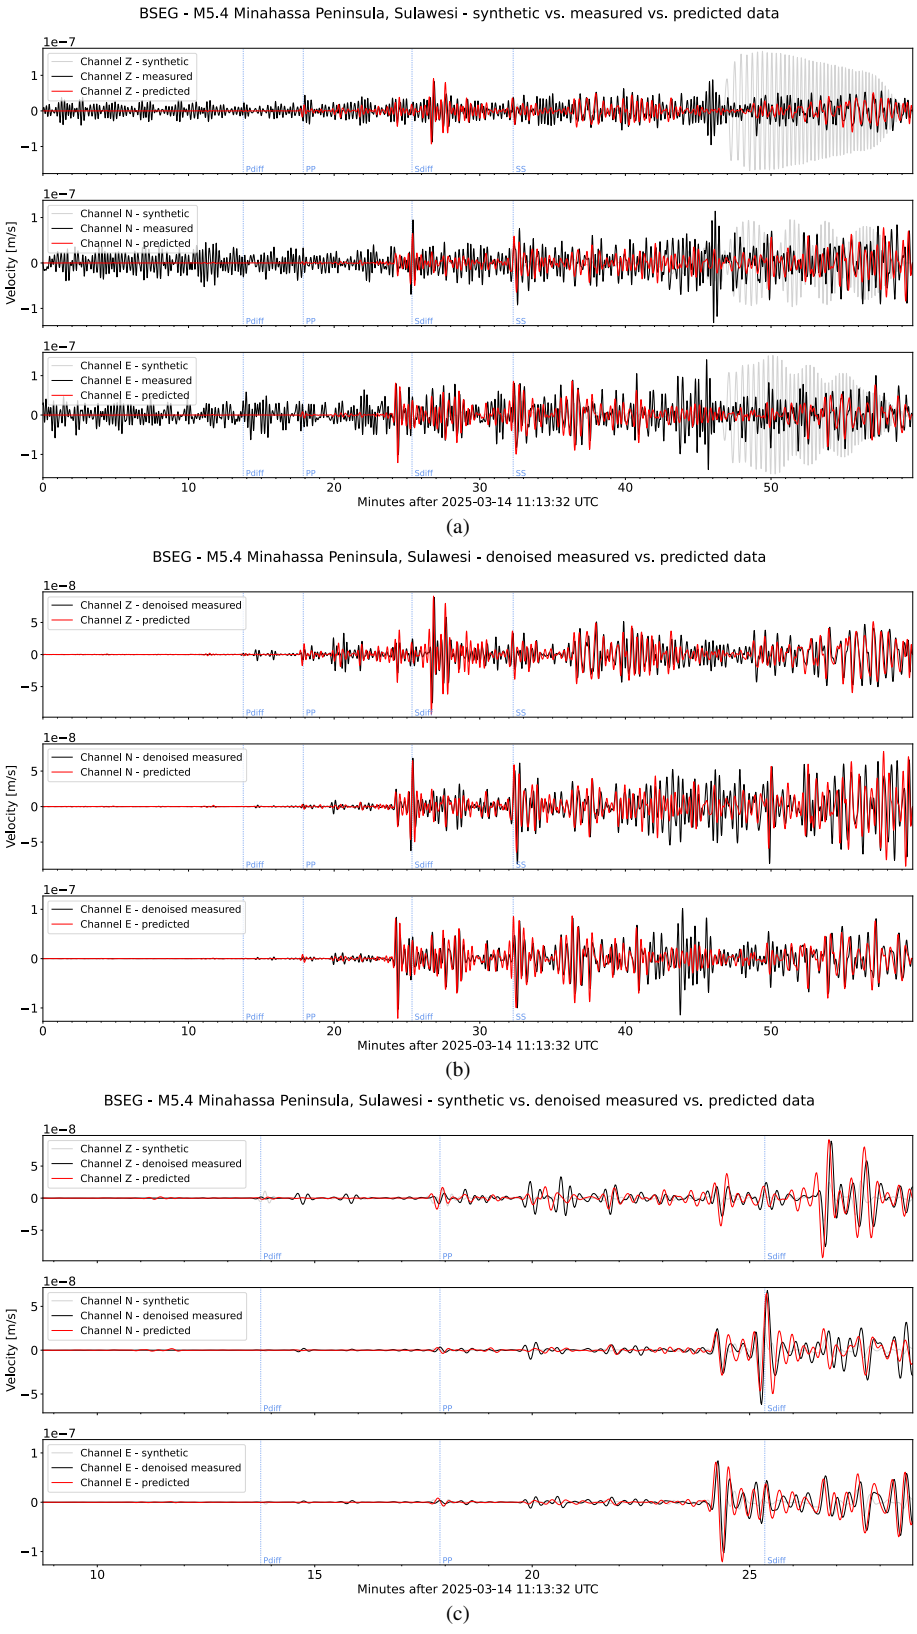

Figure 14. Prediction of noise-free measured earthquake waveforms at BSEG from synthetic data with CNN2. Example for a lower-amplitude event, a M5.4 earthquake that occurred in Sulawesi, Indonesia on 14 March 2025. Panel (a) shows a comparison of synthetic input data (light gray), measured data (black) and predicted data (red); (b) compares measured data denoised by CNN1 (black) with the predictions (red) and (c) shows a close-up on the earlier phases of the earthquake with synthetic input (light gray), denoised measured data (black) and predicted data (red). Each plot contains the three components (Z,N,E). The light blue dotted lines mark the expected arrival times of a few prominent phases.

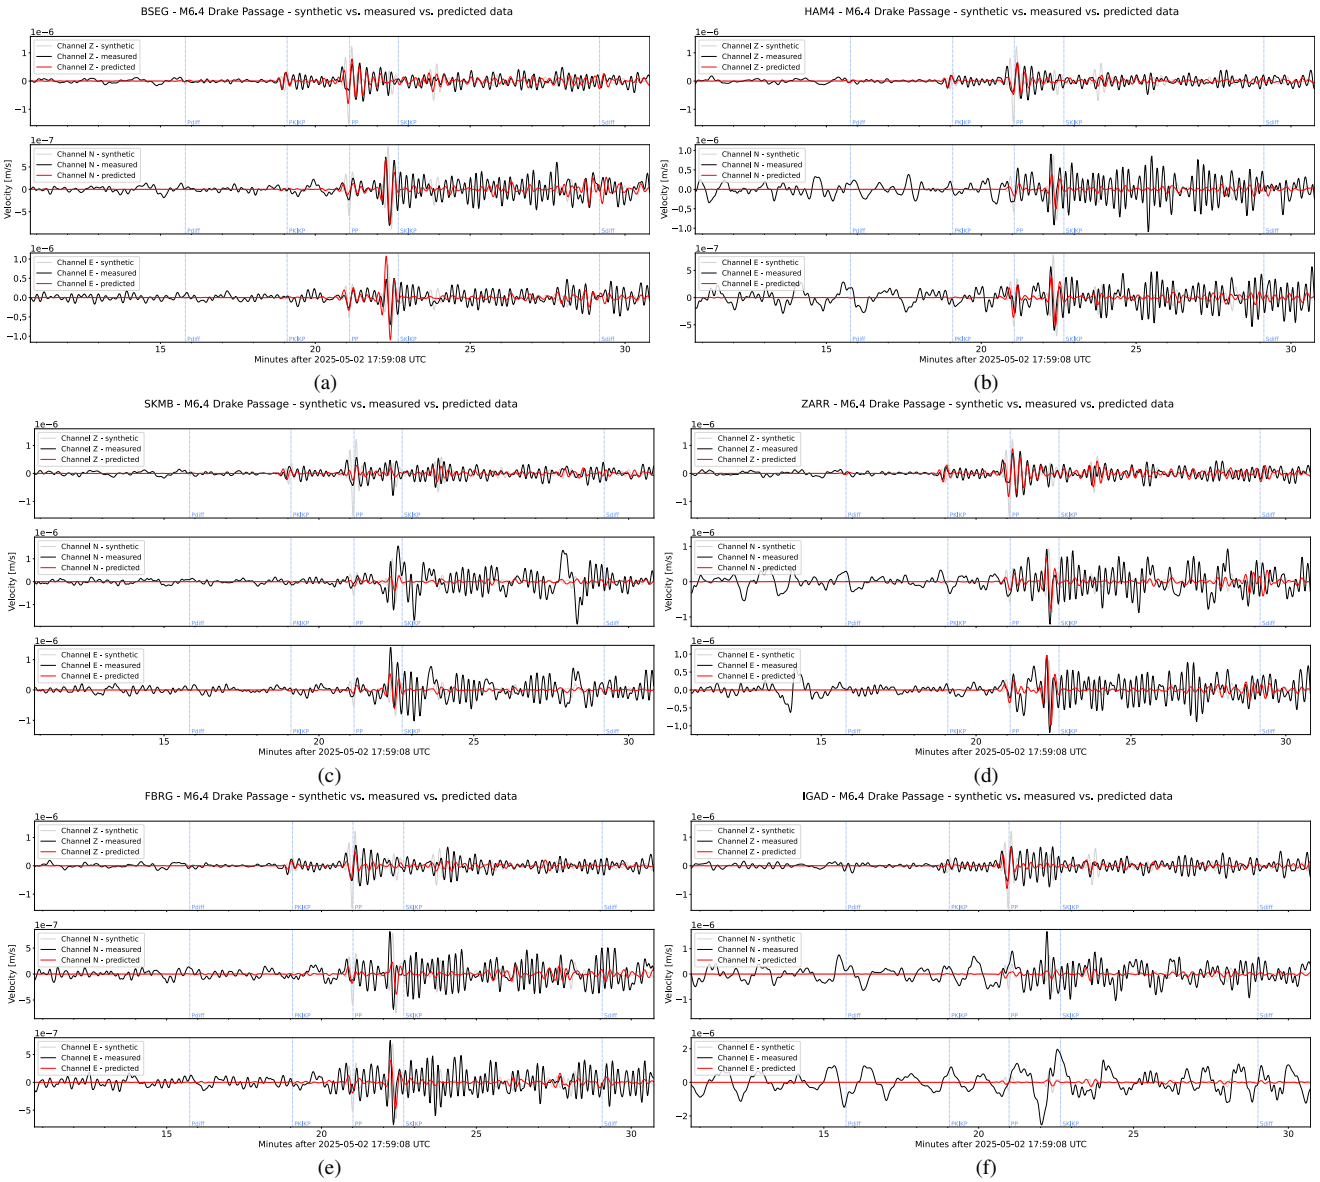

Figure 15. Prediction of noise-free measured earthquake waveforms of a M6.4 earthquake that occurred in the Drake Passage between South America and Antarctica on 2 May 2025 at six stations in Northern Germany with CNN2. Panel (a) shows the results for the training station BSEG, (b) for the station HAM4 in Hamburg, (c) for the station SKMB near Flensburg, (d) for the station ZARR in Zarrentin, (e) for the station FBRG in Faßberg and (f) for the station IGAD near Garlstedt. Each plot contains the synthetic input data (light gray), the measured data (black) and the predicted data (red) for the three components (Z, N, E). The light blue dotted lines mark the expected arrival times of a few prominent phases.

CNN2 was trained on a single station and no site effects of the other stations have been taken into account, these results are a promising step towards the prediction of the regional wavefield generated by teleseismic earthquakes in three dimensions, even in regions with a sparse distribution of stations.

5 DISCUSSION

Deep neural networks have demonstrated their enormous potential to solve complex, highly nonlinear problems in many fields of earth and data science. The task of this study – to predict measured earthquake waveforms from synthetic ones – is a very challenging nonlinear problem that may be considered unsolvable with deterministic methods available today, as it would require forward modeling with an ultra-realistic, highly-resolved Earth model and detailed knowledge of the geological conditions in the station's region and the impact of additional factors such as ground coupling on the measurements. In contrast, a deep neural network trained to predict the more complex measured waveforms from the simpler synthetic ones implicitly learns both the complex earth model and the station's site effects, given a sufficiently large and diverse training dataset. Naturally, a drawback of this approach is that it is unknown, what exactly the neural network learns. In addition, the neural network always outputs results, but their uncertainties are unknown, similar to hallucination phenomena observed in large language models. However, we argue that a result based on deterministic forward modeling also contains unquantifiable uncertainties, as the Earth model is only known to a very limited extent and

in addition, the inputs, that is, the moment tensors, are product of a non-linear inversion that in turn also requires the assumption of a simplified Earth model, introducing even more uncertainties into the modeled waveforms. In addition, numerical instabilities may cause also cause artifacts in the modeled waveforms (compare the aforementioned unrealistic surface wave arrivals, for example in Fig. 13a).

The two largest challenges of the proposed deep-learning-based prediction of measured earthquake waveforms are the limited number of earthquakes and the lack of noise-free reference waveforms. To obtain the latter, we added the first step, the deep-learning-based denoising, to the workflow. In contrast to the prediction step, the denoising offered controlled conditions, in which we could generate a very large number of synthetic earthquake waveforms and add emulated station noise to them to assemble a sufficiently large and diverse training dataset with perfect labels. These near-optimum training conditions naturally reflected in the denoising results, where the neural network achieved almost perfect predictions when applied to unseen noisy synthetic data. When applied to measured waveforms, a perfect denoising performance could not be expected due to the systematic differences between synthetic and measured waveforms and the presence of other types of station noise that we could not emulate, such as coherent anthropogenic noise or ocean microseism. Nevertheless, the denoising neural network, CNN1, achieved an impressive denoising performance for many earthquakes, even though the lack of labels impedes a quantification. To sort out events with a bad denoising performance, we included a quality control step, in which we checked for a combination of amplitude ratios that indicate bad data quality.

With the measured earthquake waveforms denoised by CNN1, we could assemble the training dataset for the second step of the workflow, the prediction of measured waveforms from synthetics. The previously mentioned factors including the limited number of earthquakes and possible imperfections in the denoising results rendered the prediction step a task of very high complexity. Consequently, CNN2 was harder to train and we could not expect a similar training convergence and generalization performance as in the denoising case. Although this reflected in both training metrics and applications to unseen data, CNN2 still achieved to generate quite convincing results for most of the tested earthquakes at the training station BSEG. In addition, the model was also able to predict quite accurate waveforms for the coordinates of other seismological stations across Northern Germany, indicating a good generalization potential. This feature paves the way for the next step, a prediction of the 3D regional earthquake wavefield across Northern Germany, even in areas where no seismological stations are available.

In general, we expect the results of the entire workflow to improve, if the quality of the synthetic data can be improved. Especially the surface arrivals of many synthetic earthquake waveforms appear unrealistic (compare Figs. 7c and 13a). This requires a Green's functions database compatible with INSTASEIS based on a more detailed global Earth model. In the long term, this could be combined with a more detailed regional model for Hamburg and its surroundings, in which the teleseismic wavefield could be injected (Monteiller et al. 2021). Furthermore, recent and upcoming developments in the fast-evolving field of deep learning may help to improve the generalization performance, for example variational autoencoders (e.g. Kingma & Welling 2013), attention mechanisms (Vaswani et al. 2017) or few-shot learning (e.g. Wang et al. 2020).

6 CONCLUSIONS

We have presented a deep-learning scheme for the prediction of the wavefield generated by large teleseismic earthquakes from synthetic data. The workflow consists of two steps: first, we trained a neural network, CNN1, to denoise earthquake waveforms at the long-standing training station BSEG, located close to Hamburg in Northern Germany, by emulating the station noise and adding it to noise-free synthetic waveforms. With the trained CNN1, we generated noise-free labels for the second step, in which we trained a second neural network, CNN2, to predict noise-free measured earthquake waveforms from synthetic data. With applications of CNN2 to unseen earthquakes at the training station BSEG, we demonstrated a good agreement between measured and predicted waveforms. Promising prediction results of CNN2 at the locations of other stations across Northern Germany constitute a first step towards the prediction of the large-scale regional wavefield generated by teleseismic earthquakes in three dimensions.

ACKNOWLEDGMENTS

We used INSTASEIS (van Driel et al. 2015) and SYNGINE (Krischer et al. 2017) for the generation of synthetic earthquake waveforms and The Searchable Product Repository (SPUD) offered by the Incorporated Research Institutions for Seismology Data Management Center (IRIS-DMC) to obtain the moment tensors for past earthquakes. We used OBSPY for seismological data processing and TAUPY for the calculation of arrival times. We are grateful to J. Walda and S. Knispel for their important contributions in the implementation of the deep learning framework, for which we used TENSORFLOW 2. The seismological station BSEG is part of the German Regional Seismic Network (GRSN) operated by the Federal Institute of Geosciences and Natural Resources (BGR). This work was supported by the Federal Ministry of Education and Research of Germany (BMBF, 05D23GU5). We further thank the Machine Learning Methods in Geophysics Group at the University of Hamburg and the members of the ErUM-WAVE project for continuous discussion.

DATA AVAILABILITY

The presented datasets and the trained neural networks are available and will be provided upon reasonable request.

REFERENCES

Bauer, A., Schwarz, B., Walda, J., & Gajewski, D., 2025a. DiffractioNet: Deep learning seismic and GPR diffraction separation, *Geophysics*, **90**(3), V255–V264.

Bauer, A., Schwarz, B., Walda, J., Hammer, C., Schlarb, H., Isleif, K.-S., Gerberding, O., & Hadziioannou, C., 2025b. Urban DAS data enhancement and coherent noise removal with deep learning, in *86th EAGE Annual Conference & Exhibition*, European Association of Geoscientists & Engineers.

Chakraborty, M., Cartaya, C. Q., Li, W., Faber, J., Rümpker, G., Stoecker, H., & Srivastava, N., 2022a. PolarCAP–A deep learning approach for first motion polarity classification of earthquake waveforms, *Artificial Intelligence in Geosciences*, **3**, 46–52.

Chakraborty, M., Fenner, D., Li, W., Faber, J., Zhou, K., Rümpker, G., Stoecker, H., & Srivastava, N., 2022b. CREIME–A convolutional recurrent model for earthquake identification and magnitude estimation, *Journal of Geophysical Research: Solid Earth*, **127**(7), e2022JB024595.

Gorski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., & Bartelmann, M., 2005. HEALPix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere, *The Astrophysical Journal*, **622**(2), 759.

He, K., Zhang, X., Ren, S., & Sun, J., 2016. Deep residual learning for image recognition, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778.

Kingma, D. P. & Welling, M., 2013. Auto-encoding variational bayes, *arXiv preprint arXiv:1312.6114*.

Klambauer, G., Unterthiner, T., Mayr, A., & Hochreiter, S., 2017. Self-normalizing neural networks, *arXiv preprint*, **abs/1706.02515**.

Klinge, J., Schippkus, S., Walda, J., Hadziioannou, C., & Gajewski, D., 2025. Predictive modelling of seismic wave fields: learning the transfer function using encoder–decoder networks, *Geophysical Journal International*, **240**(3), 1611–1621.

Knispel, S., Walda, J., Zehn, R., Bauer, A., & Gajewski, D., 2022. A self-attention enhanced encoder-decoder network for seismic data denoising, in *Second International Meeting for Applied Geoscience & Energy*, pp. 2922–2926, Society of Exploration Geophysicists.

Krischer, L., Hutko, A. R., van Driel, M., Stähler, S., Bahavar, M., Trabant, C., & Nissen-Meyer, T., 2017. On-demand custom broadband synthetic seismograms, *Seismological Research Letters*, **88**(4), 1127–1140.

Kubo, H., Naoi, M., & Kano, M., 2024. Recent advances in earthquake seismology using machine learning, *Earth, Planets and Space*, **76**(1), 36.

Laurenti, L., Tinti, E., Galasso, F., Franco, L., & Marone, C., 2022. Deep learning for laboratory earthquake prediction and autoregressive forecasting of fault zone stress, *Earth and Planetary Science Letters*, **598**, 117825.

Lu, L., Shin, Y., Su, Y., & Karniadakis, G. E., 2019. Dying relu and initialization: Theory and numerical examples, *arXiv preprint arXiv:1903.06733*.

Montagner, J.-P. & Kennett, B., 1996. How to reconcile body-wave and normal-mode reference earth models, *Geophysical Journal International*, **125**(1), 229–248.

Monteiller, V., Beller, S., Plazolles, B., & Chevrot, S., 2021. On the validity of the planar wave approximation to compute synthetic seismograms of teleseismic body waves in a 3-D regional model, *Geophysical Journal International*, **224**(3), 2060–2076.

Monterrubio-Velasco, M., Callaghan, S., Modesto, D., Carrasco, J. C., Badia, R. M., Pallares, P., Vázquez-Novoa, F., Quintana-Ortí, E. S., Pienkowska, M., & de la Puente, J., 2024. A machine learning estimator trained on synthetic data for real-time earthquake ground-shaking predictions in Southern California, *Communications Earth & Environment*, **5**(1), 258.

Mousavi, S. M. & Beroza, G. C., 2022. Deep-learning seismology, *Science*, **377**(6607), eabm4470.

Mousavi, S. M., Ellsworth, W. L., Zhu, W., Chuang, L. Y., & Beroza, G. C., 2020. Earthquake transformer—an attentive deep-learning model for simultaneous earthquake detection and phase picking, *Nature communications*, **11**(1), 3952.

Perol, T., Gharbi, M., & Denolle, M., 2018. Convolutional neural network for earthquake detection and location, *Science Advances*, **4**(2), e1700578.

Ronneberger, O., Fischer, P., & Brox, T., 2015. U-Net: Convolutional networks for biomedical image segmentation, in *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241.

Saad, O. M., Chen, Y., Savvaidis, A., Chen, W., Zhang, F., & Chen, Y., 2022. Unsupervised deep learning for single-channel earthquake data denoising and its applications in event detection and fully automatic location, *IEEE Transactions on Geoscience and Remote Sensing*, **60**, 1–10.

van Driel, M., Krischer, L., Stähler, S. C., Hosseini, K., & Nissen-Meyer, T., 2015. Instaseis: instant global seismograms based on a broadband waveform database, *Solid Earth*, **6**(2), 701–717.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I., 2017. Attention is all you need, *Advances in neural information processing systems*, **30**.

Wang, C., Xia, K., Yao, W., & Marone, C., 2025. Generalizable deep learning models for predicting laboratory earthquakes, *Communications Earth & Environment*, **6**(1), 219.

Wang, Y., Yao, Q., Kwok, J. T., & Ni, L. M., 2020. Generalizing from a few examples: A survey on few-shot learning, *ACM computing surveys (csur)*, **53**(3), 1–34.

Xie, S., Girshick, R., Dollár, P., Tu, Z., & He, K., 2017. Aggregated residual transformations for deep neural networks, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1492–1500.

Yang, L., Liu, X., Zhu, W., Zhao, L., & Beroza, G. C., 2022. Toward improved urban earthquake monitoring through deep-learning-based noise suppression, *Science advances*, **8**(15), eabl3564.

Zehn, R., Zehn, K., Schuenemann, E., & Walda, J., 2023. Signal preservation while performing seismic denoising with neural networks, in *84th EAGE Annual Conference & Exhibition*, vol. 2023, pp. 1–5, European Association of Geoscientists & Engineers.

Zhu, W. & Beroza, G. C., 2019. PhaseNet: a deep-neural-network-based seismic arrival-time picking method, *Geophysical Journal International*, **216**(1), 261–273.