

Detection of Hidden Low-Frequency Earthquakes in Southern Vancouver Island with Deep Learning

Jiun-Ting Lin ^{*1}, Amanda M. Thomas ², Loïc Bachelot ², Douglas R. Toomey ², Jake Searcy ³, Diego Melgar ²

¹Lawrence Livermore National Laboratory, Livermore, CA, USA, ²Department of Earth Sciences, University of Oregon, Eugene, OR, USA, ³School of Computer and Data Sciences, University of Oregon, Eugene, OR, USA

Author contributions: *Conceptualization:* J.-T. L., A. M. T., *Data Curation:* A. M. T., *Formal Analysis:* J.-T. L., A. M. T., *Funding Acquisition:* J.-T. L., A. M. T., *Investigation:* All authors., *Methodology:* J.-T. L., A. M. T., L. B., D. R. T., J. S., *Project Administration:* A. M. T., *Resources:* J.-T. L., A. M. T., *Software:* J.-T. L., A. M. T., J. S., *Supervision:* A. M. T., D. M., *Validation:* All authors., *Visualization:* J.-T. L., A. M. T., *Writing – original draft:* J.-T. L., A. M. T., *Writing – review & editing:* All authors.

Abstract Low-frequency earthquakes (LFEs) are small-magnitude earthquakes that are depleted in high-frequency content relative to traditional earthquakes of the same magnitude. These events occur in conjunction with slow slip events (SSEs) and can be used to infer the space and time evolution of SSEs. However, because LFEs have weak signals, and the methods used to identify them are computationally expensive, LFEs are not routinely cataloged in most places. Here, we develop a deep-learning model that learns from an existing LFE catalog to detect LFEs in 14 years of continuous waveform data in southern Vancouver Island. The result shows significant increases in detection rates at individual stations. We associate the detections and locate them using a grid search approach in a 3D regional velocity model, resulting in over 1 million LFEs during the performing period. Our resulting catalog is consistent with a widely used tremor catalog during periods of large-magnitude SSEs. However, there are time periods where it registers far more LFEs than the tremor catalog. We highlight a 16-day period in May 2010, when our model detects nearly 3,000 LFEs, whereas the tremor catalog contains only one tremor detection in the same region. This suggests the possibility of hidden small-magnitude SSEs that are undetected by current approaches. Our approach improves the temporal and spatial resolution of the LFE activities and provides new opportunities to understand deep subduction zone processes in this region.

Non-technical summary Relative to regular earthquakes, low-frequency earthquakes (LFEs) release their energy in a “slower” way and can help us to understand seismic activity on deep (30 km+) faults and implications for earthquake hazards. However, because of their weak signals, detecting LFEs efficiently is challenging. In this paper, we develop a deep learning model that detects more than 1 million LFEs in southern Vancouver Island in 14 years. Our resulting LFE catalog is generally consistent with the tremor and slow-slip event (SSE) catalogs. This is expected because they share related, if not similar, processes. What is unexpected is that we find LFEs that are not in the tremor and slow-slip event catalogs. This suggests that our method can find hidden small-magnitude SSEs that are undetected by existing approaches. Our method can help advance our understanding of seismic activity in this region.

1 Introduction

Slow slip events (SSEs) are a type of transient fault slip during which the slip rate accelerates to speeds that are 1-2 orders of magnitude faster than the background tectonic loading rate (e.g., Bürgmann, 2018; Behr and Bürgmann, 2021). SSEs occur frequently in subduction zones around the globe (Saffer and Wallace, 2015). In the past two decades, much effort has been dedicated to documenting their spatial and temporal characteristics in different tectonic environments (Obara, 2002; Rogers and Dragert, 2003; Beroza and Ide, 2011; Obara and Kato, 2016; Bürgmann, 2018; Behr and Bürgmann, 2021). Because slow slip events occur over significantly longer timescales than typical earthquakes, they generate very weak seismic waves that are both lower in am-

plitude and depleted in high-frequency (i.e., > 1 Hz) content relative to regular earthquakes (e.g. Thomas et al., 2016).

Obara (2002) first recognized what he dubbed non-volcanic tremor (NVT) beneath the Shikoku and Kii peninsulas in Japan. It has a low-amplitude signal with a predominant frequency content of 1-10 Hz lasting a few hours to a few days. Obara (2002) also recognized that NVT signals propagated with a velocity most consistent with that of S-waves and located deep on the plate interface. Shortly thereafter nonvolcanic tremor was recognized as the seismic manifestation of deep slow slip (Rogers and Dragert, 2003). NVT can be rapidly detected and is a useful tool for identifying and tracking SSE evolution.

Many researchers have developed algorithms to detect NVT. Kao et al. (2005) used the source scanning al-

Production Editor:
Gareth Funning
Handling Editor:
Brendan Crowell
Copy & Layout Editor:
Hannah F. Mark

Received:
October 18, 2023
Accepted:
August 28, 2024
Published:
October 18, 2024

*Corresponding author: lin51@llnl.gov

gorithm, which estimates a source brightness by stacking normalized seismograms across a network that are shifted in time according to predicted traveltimes from each potential source location. Multiple studies have built on this foundation leveraging waveform classification, similarity of waveform envelopes, and the source scanning algorithm to identify tremor (Kao et al., 2009; Bombardier et al., 2023). One of the most widely used catalogs is that of Wech and Creager (2008), which identifies tremors in real-time by cross-correlating waveform envelopes and grid searching the location, which shifts the S-wave time until the summed cross-correlation functions for all the station pairs reach the maximum value (Wech, 2021). Because there is no P-wave information alongside the S-wave delays, the depth estimates may be unreliable (Wech, 2021). Furthermore, detections employ 5-minute time windows, which do not allow for analysis of shorter timescale phenomena or resolution of energy coming from multiple locations. Another approach to identify tremor is known as the cross-station correlation approach that relies on finding coherent shear waves across stations by searching over a range of polarization angles and time lags to maximize the cross-correlation (Rubin and Armbruster, 2013; Armbruster et al., 2014; Savard and Bostock, 2015).

NVT is made up, in whole or in part, of low-frequency earthquakes (LFEs, Shelly et al., 2007). LFEs are more traditional seismic sources that have identifiable P- and S-waves but are deficient in high-frequency content (above a few Hz) relative to shallow earthquakes of similar magnitude (e.g. Thomas et al., 2016). Traditionally LFEs are detected by template matching approaches (e.g. Bostock et al., 2012; Chamberlain et al., 2014; Royer and Bostock, 2014; Bostock et al., 2015; Shelly et al., 2007) or cross-station approaches (e.g. Savard and Bostock, 2015). These methods utilize waveform coherency either cross-time or cross-station to construct a summed cross-correlation function. When it exceeds a threshold (e.g. eight times the median absolute deviation, Shelly et al., 2007), the window is considered a detection. This process can be refined by stacking all the detected waveforms to generate new LFE templates with an increased signal-to-noise ratio. Within this framework, groups of LFEs that occur at different times but have similar waveform characteristics are grouped into families that reflect slip at the same or nearly the same location.

Although the physical process responsible for their generation is still a matter of debate (Obara, 2002; Obara and Hirose, 2006; Seno and Yamasaki, 2003), LFEs are generally thought to reflect surrounding, largely aseismic fault slip during SSEs (e.g. Thomas et al., 2018) and permit analysis of space and time evolution of slip on short timescales and in high spatial resolution. They can be used to study slip evolution in individual SSEs (e.g. Frank et al., 2014; Inbal et al., 2021), resolve inferred smaller magnitude SSEs that are not easily identifiable in high-rate Global Navigation Satellite System data (e.g. Rousset et al., 2019), and to constrain the velocity structure of the forearc crust (e.g. Savard et al., 2018; Calvert et al., 2020; Delph et al., 2021). Despite all the

potential uses of LFEs, they are not routinely cataloged in Cascadia because of their low signal-to-noise ratio.

Thomas et al. (2021) proposed a machine-learning (ML) approach that can identify LFE waveforms in noisy timeseries data from a single station. They have successfully applied this model in Parkfield, CA, and shown that it identified new events that are not in the original catalog, suggesting the potential of utilizing such an approach. More recently, Münchmeyer et al. (2024) applied a similar approach and demonstrated its capability of detecting LFEs in various regions. They showed the transferability of the model, suggesting that the general characteristic of LFEs can be learned. Here we train a Convolutional Neural Network (CNN) to detect LFEs in a long-term scale using the catalog of Bostock et al. (2015) which was originally assembled via template matching using continuous seismic data from southern Vancouver Island. We find that the model can reliably detect LFEs with a false positive rate of <1% when applied on multiple stations. We apply the model to 14 years of continuous seismic data recorded in southern Vancouver Island (Figure 1) to detect LFEs on individual stations. We associate detections and locate them using a 3D regional velocity model. For the period with large SSEs (Michel et al., 2018), the resulting catalog is generally consistent with the tremor catalog. Furthermore, the new catalog also identifies many LFEs that do not have corresponding tremors. Overall this technique may be useful for efficient, operational detection of LFEs and further understanding of the seismic radiation that occurs during SSEs.

2 Methods

2.1 Training data for phase picks

The first goal of this work is to develop a phase picker that can distinguish LFEs from noise and make arrival time picks on records deemed to contain signal. Accomplishing this task requires obtaining training data that includes many representative examples of noise, P-, and S-waves with associated picks. We obtained phase picks from known LFEs that were originally identified by a combination of autocorrelation and template matching (e.g. Bostock et al., 2012; Royer and Bostock, 2014). The catalog we use is that of Bostock et al. (2015), downloaded from the slow earthquake database (Kano et al., 2018). For each arrival time pick in the catalog, we download a 30 s window of data centered on the pick time using the Obspy package (Krischer et al., 2015). To distinguish earthquakes from noise, representative noise samples are included in the training data for the CNN. As such we download a similar number of noise windows (defined as the time period prior to the P arrival time pick). Noise data are randomly selected when there are no known LFEs before and after the time with a minimum separation of 180 s. This process results in more than 500,000 waveforms for P-wave and S-wave picks, and approximately the same amount of noise data. We save the data at 100 Hz to achieve a temporal resolution that can ideally detect differences as small as 0.01 s. For the target, we use a Gaussian

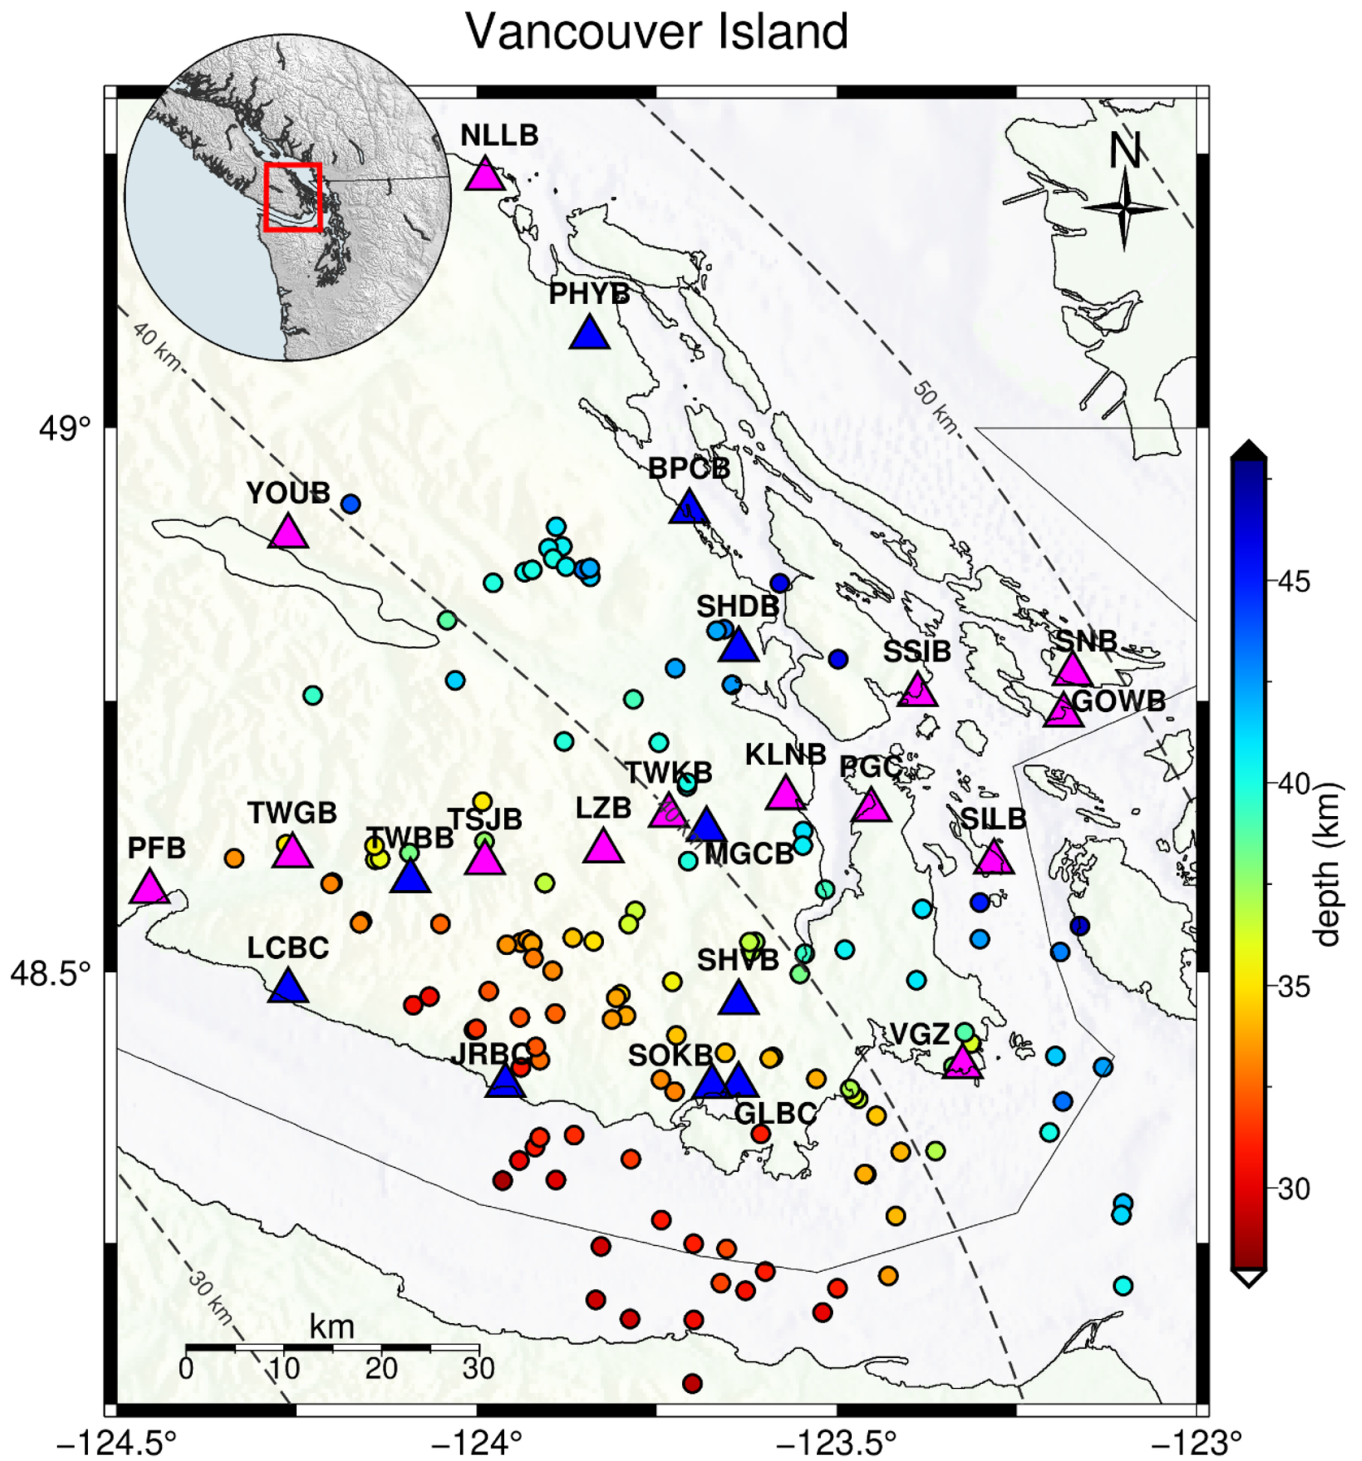


Figure 1 Map view of the study area. Magenta triangles show the stations used for model training and testing; blue triangles represent the unseen stations, which are not involved during the training process, for model testing. Circles denote the LFEs locations from the [Bostock et al. \(2015\)](#) catalog, color-coded by their depth.

function with a standard deviation of 0.4 s with the same length and sampling rate as the input waveform centered at the P or S wave arrival time. This Gaussian function allows some errors in the arrival time pick in the catalog, but it is not sufficiently large that it smears the detection resolution. For noise waveforms, we set the target to zero (Figure 2).

2.2 Convolutional Neural Network architecture and training

The input to the network is three component waveform data. Since data windows are 30 s long and we employ a sample rate of 100 Hz, the input data has a length of 3000 samples. Leaving the training data in the original form, with the pick in the middle, would result in the CNN learning to pick the middle sample each time. As such, similar to [Thomas et al. \(2021\)](#) we use a data generator during training that randomly selects subsamples

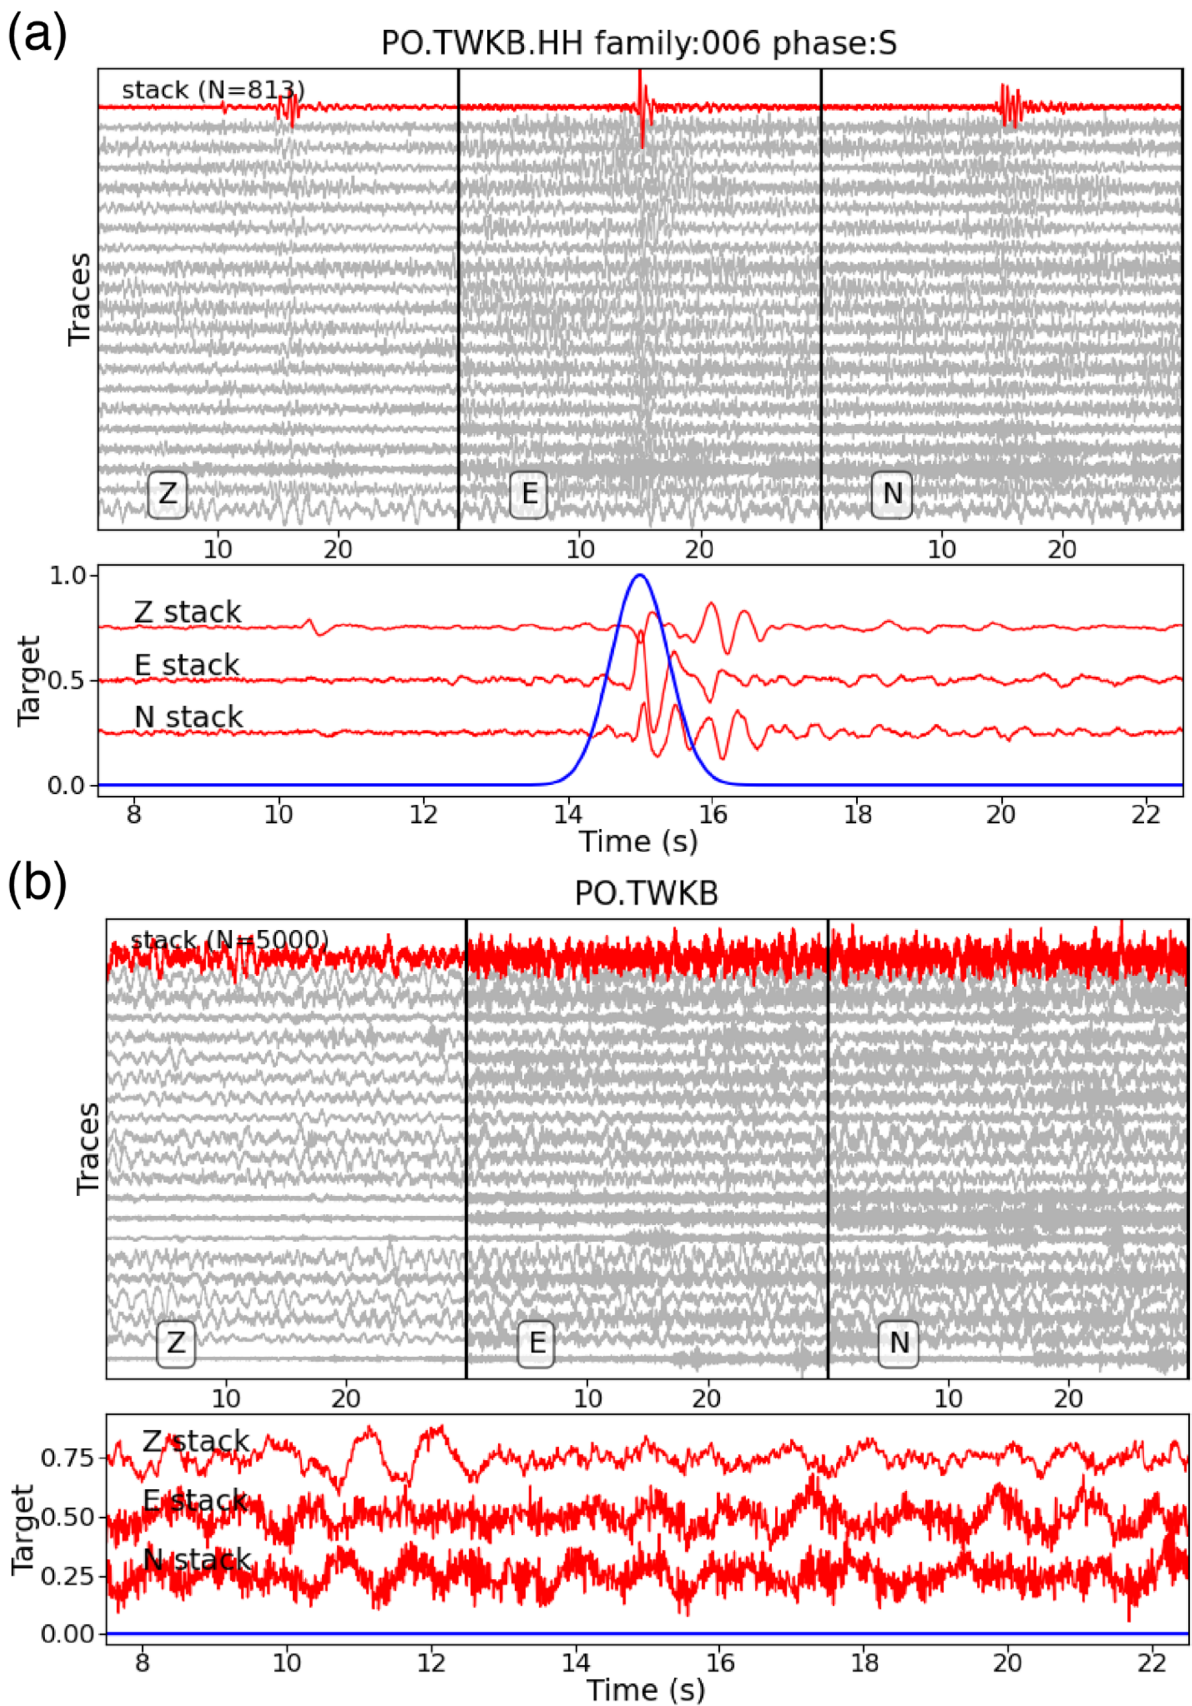


Figure 2 Examples of LFE and noise data at station TWKB. (a) 3-component waveforms from LFE catalog (family: 006) from Bostock et al. (2015). Gray lines show the raw data normalized by their amplitude (only a few examples are shown here); red lines show the stacking of the family (i.e. gray lines) with a total of 813 waveforms. Blue line shows the Gaussian function as the possibility of the S-wave arrival, which is the target for model training. Note that the model applies to individual waveforms (i.e. gray lines), not the stacked data which is only for demonstration purposes. (b) Similar to (a) but for noise data.

of traces from the training data, called batches, and applies the following modifications to the data prior to input. First, we randomly select a start time in the first half of the trace and include only 15 seconds of data beginning at that time. This has the effect of randomly shifting the pick in time such that it can occur at any time during the window. Second, to account for variable amplitudes in the training data, we normalize the three component data with the maximum amplitude of all three components and apply a logarithmic transformation to the input data, as in Thomas et al. (2021). This transformation maps each value, x , in the original traces to two numbers: the first is $\text{sgn}(x)$ while the second is the $\ln(\|x\| + \epsilon)$ where $\epsilon = 1 \times 10^{-6}$. This has the effect of scaling the features such that input amplitudes do not vary over orders of magnitude, and preserving information on the sign. The data generator supplies six channels (3 components with a normalized amplitude and sign for each) in batches to the CNN during training and augments the training data by shifting the pick times.

For the ML model, we employ the U-Net architecture from Thomas et al. (2021). U-Nets are composed of several convolutional layers (Lecun et al., 1998) and links, which allow the raw and early information to be accessible to the later decision layers. This architecture has been shown to be successful in biomedical image processing (Ronneberger et al., 2015) and in seismic phase identification (e.g. Zhu and Beroza, 2018). The model contains a size factor to control the number of convolution filters per layer (double the number, original and half the number of filters). Here, we only test three network sizes, called size 0.5, size 1, and size 2 model in Thomas et al. (2021), and fix the standard deviation of arrival time label of 0.4 s because our goal is to build and test the feasibility of applying such a method in a noisy environment. We find that the size 2 model works the best for P-wave and S-wave detection in our case (Supplementary Figure S1). We do not fine-tune the hyperparameters as they have been shown to have a minor influence on the performance (Thomas et al., 2021).

Data partitioning is important to prevent potential data leakage, a serious issue in ML models. One modification that we make is instead of mixing all the waveforms from all the events (Thomas et al., 2021), we split the data by the event ID so that traces from the same event will not participate in both the training and testing datasets, potentially minimizing the model memorization. A total of 269,422 events are used in the study. We hold 25% of the events for model testing. We set our model batch size to 32, meaning that the model updates its weights based on 32 samples in each training step, and with a total of 30 training epochs (Supplementary Figure S2). Once the training is completed, we evaluate it with both the testing dataset and the continuous data.

2.3 LFE association and location

To associate the detections from our model, we require a minimum of three detections occurring within the same 15-second time window. This criterion results in 1,058,114 candidate LFEs that can be located. We use

a direct grid search approach to locate the LFEs. Intuitively, the problem requires 4 detections to avoid ambiguity in location and origin time. However, by using the grid search approach, we limit the possible source locations to a specific region, which serve as an additional constraint to alleviate this ambiguity for events with fewer detections. Furthermore, it enables us to locate the global minimum without the need to handle derivatives at sharp velocity boundaries (Lomax et al., 2009). We first calculate travel times to each station from each potential source in a 3D grid centered at -123.75 and 48.7 for longitude and latitude, respectively. The spacing of the grid is 1 km in each direction with a total of 120 and 140 grid points in longitude and latitude directions, respectively, and up to 60 km depth. Velocities are defined on this grid by interpolating the velocity model from Savard et al. (2018). We calculate the travel times based on the method described in Toomey et al. (1994).

For each set of associated detections, we search over all possible source locations seeking to minimize the difference between the observed and simulated travel times. Specifically, we calculate

$$\delta_k^{i=1, N; j=1, M} = (OT^i + T_k^i) - \hat{T}_k^j \quad (1)$$

Here $\delta_k^{i=1, N; j=1, M}$ is the travel time difference between the i -th set of observed travel times and a potential source located at the j -th grid node. OT^i is the origin time of the LFE source responsible for the associated detections, T_k^i is the observed travel time from this source to station k , and \hat{T}_k^j is the modeled travel time of a source located at the j -th grid node to the station. Although the origin time OT is unknown, it is a constant applied to each associated set of detections. This constant shift for all stations can be removed by subtracting the mean value. Equation (1) can be modified to

$$\hat{\delta}_k^{i=1, N; j=1, M} = \delta_k^{i, j} - \frac{\sum_{k=1}^K \delta_k^{i, j}}{K} \quad (2)$$

Where K is the number of available stations for each associated events. We find the preferred location j by searching the grid node with minimum misfit.

$$j = \underset{j}{\operatorname{argmin}} \left[\frac{\sum_{k=1}^K \left| \hat{\delta}_k^{i=1, N; j=1, M} \right|}{K} \right] \quad (3)$$

In total we search over $N=1,058,114$ associated sets of detections and $M=1,024,800$ possible source locations.

3 Results

3.1 Assessing model performance

We test the model with the unexposed 25% testing data, as introduced above. We first treat it as a simple binary classification problem (i.e. LFE or noise) and calculate the model accuracy, precision, and recall (Figure 3). We will analyze the performance of arrival time in a later section. The metrics are defined below:

$$\text{Accuracy} (\%) = \frac{TP + TN}{TP + TN + FP + FN} \times 100\% \quad (4)$$

where true positive (TP) is defined by the number of positive detections that are actual LFEs; true negative (TN) is the number of negative detections that are noise; false positive (FP) and false negative (FN) are the numbers of incorrect LFE and noise predictions, respectively. We calculate the accuracy as a function of the decision threshold for the P- and S-wave models and find that the S-wave model has slightly higher accuracy (~92%) than the P-wave (~90%) at threshold=0.1. Next, we calculate precision as

$$Precision (\%) = \frac{TP}{TP + FP} \times 100\% \quad (5)$$

Unlike accuracy, precision ignores the number from negative predictions, and the value simply represents the rate of positive predictions and that are actually positive. Both our P and S-model have a precision of ~95% at threshold=0.1 which means the predicted LFEs are generally true, and only 5% of the detections are false detection i.e. noise. Furthermore, to understand the rate of misclassification of actual LFEs we calculate recall, or the true positive rate (TPR)

$$Recall (\%) = \frac{TP}{TP + FN} \times 100\% \quad (6)$$

Recall evaluates the rate of actual LFEs and that are successfully detected. For example both our P and S-model have a recall of ~90% at threshold=0.1, this means 90% of the LFEs can be identified, and 10% of the LFEs are misclassified as noise.

A receiver operating characteristic (ROC) curve is another metric to evaluate the overall model performance (Figure 3b). The ROC curve varies the decision thresholds of a binary classifier and examines the TPR against the false positive rate (FPR). The Area Under the ROC Curve (AUC) is a more common representation of the ROC curve. AUC spans a value from 0.5 to 1, where 0.5 represents randomly guessing, and 1 indicates a perfect model. To further validate our model, we perform three different tests: testing the model with the full testing dataset (v1); testing with only large (>M2.2) events (v2); and recording at close (<30 km) epicentral distances (v3). We randomly select data from the above criteria and pass them into the generator to generate 1,000 LFEs and 1,000 noise samples and repeat the procedure 20 times to assess the distribution of ROC curves and AUC values (Figure 3). In comparison to the v1 test, which had AUCs of 0.92 and 0.97 for the P- and S-wave models, respectively, we find that the model performs better when testing it using only large events with an AUC of 0.96 and 0.98, for P- and S-wave models, respectively. This suggests that the ML model performs better with the higher signal-to-noise ratio data, representative of larger LFEs. The v3 test shows that the model does not perform significantly better than the v1 test. This is because most of the LFEs are located beneath the stations with depth of ~40 km (Figure 1) and thus the difference in horizontal distances is insignificant.

3.2 Application to continuous seismic data

After evaluation of the aforementioned metrics, we set a decision threshold at 0.1 for both the P- and S-wave

model and run our ML model on 14 years of continuous waveform data from 2005 to 2018. We evaluate the model and find that it can reliably identify known LFEs. Figure 4 shows an example of a known LFE with S-wave detections at multiple stations. The model clearly picks the arrival at stations TWKB, LZB, PGC, and SSIB. For station SILB and VGZ, the model detects the event but with a few seconds of arrival time difference. Overall the model is adept at identifying existing LFEs. Furthermore, we find that our ML model routinely detects events that are not in the original catalog (Figure 5). Assuming there is only one LFE in the 15 s time window and all the detections are made independently, the chance that such detections are false detections is smaller than 1% given the high precision of the model (Figure 3).

Beyond individual detections, Figure 6 shows time series of daily detection counts for the 14 stations that were used to train the network. High LFE rates manifest across the network during times of known SSEs, while detection rates are low during inter-SSE time periods. Furthermore, despite being trained on data from 2003-2014, the model shows promising results when applying it to data outside of this period, suggesting its temporal extrapolation capabilities. We also apply the trained model to 10 stations that were not used during model training (Supplementary Figure S3). The detection counts on these stations have the same low daily detection counts during inter-SSE periods that increase abruptly during times of known SSEs for time periods when data is available. The CNN has not seen any of the LFE data from these stations, yet it can still robustly detect LFEs, and the patterns are consistent with the original Bostock catalog (Supplementary Figure S3). This result demonstrates the CNNs ability to extrapolate learned information to new settings and that path, site effects, and noise character for the unseen stations are likely to be similar to the training stations.

3.3 P and S-wave arrival time estimates

As shown in Figure 4, arrival time prediction can be challenging, especially for low signal-to-noise ratio data. We find that by setting a decision threshold of 0.1, the model has an averaged S-wave travel time misfit of -0.2 s, with a standard deviation of 3.6 s (Figure 7). This slightly decreases to -0.17 s and a standard deviation of 2.6 s when setting a higher threshold of 0.5. The negative mean value is mainly because the model identifies some of the earlier P-wave arrivals. This is shown in Figure 7a in the 0-40 km distance groups, where the arrival time misfits show a secondary peak at -6 s, near the expected S-P wave arrival time difference for the depth of ~40 km. Similarly, this can be also seen for the larger magnitude events shown in Figure 7b where the P-wave amplitude is expected to be more obvious. For long-distance groups (40-80 km), this effect becomes insignificant because of the attenuation of the P-wave at such distances. We find that the misfits do not decrease when events are less than 40 km deep, this is likely because all the sources are deep and thus the difference in the horizontal distance is insignificant, similar to the result of

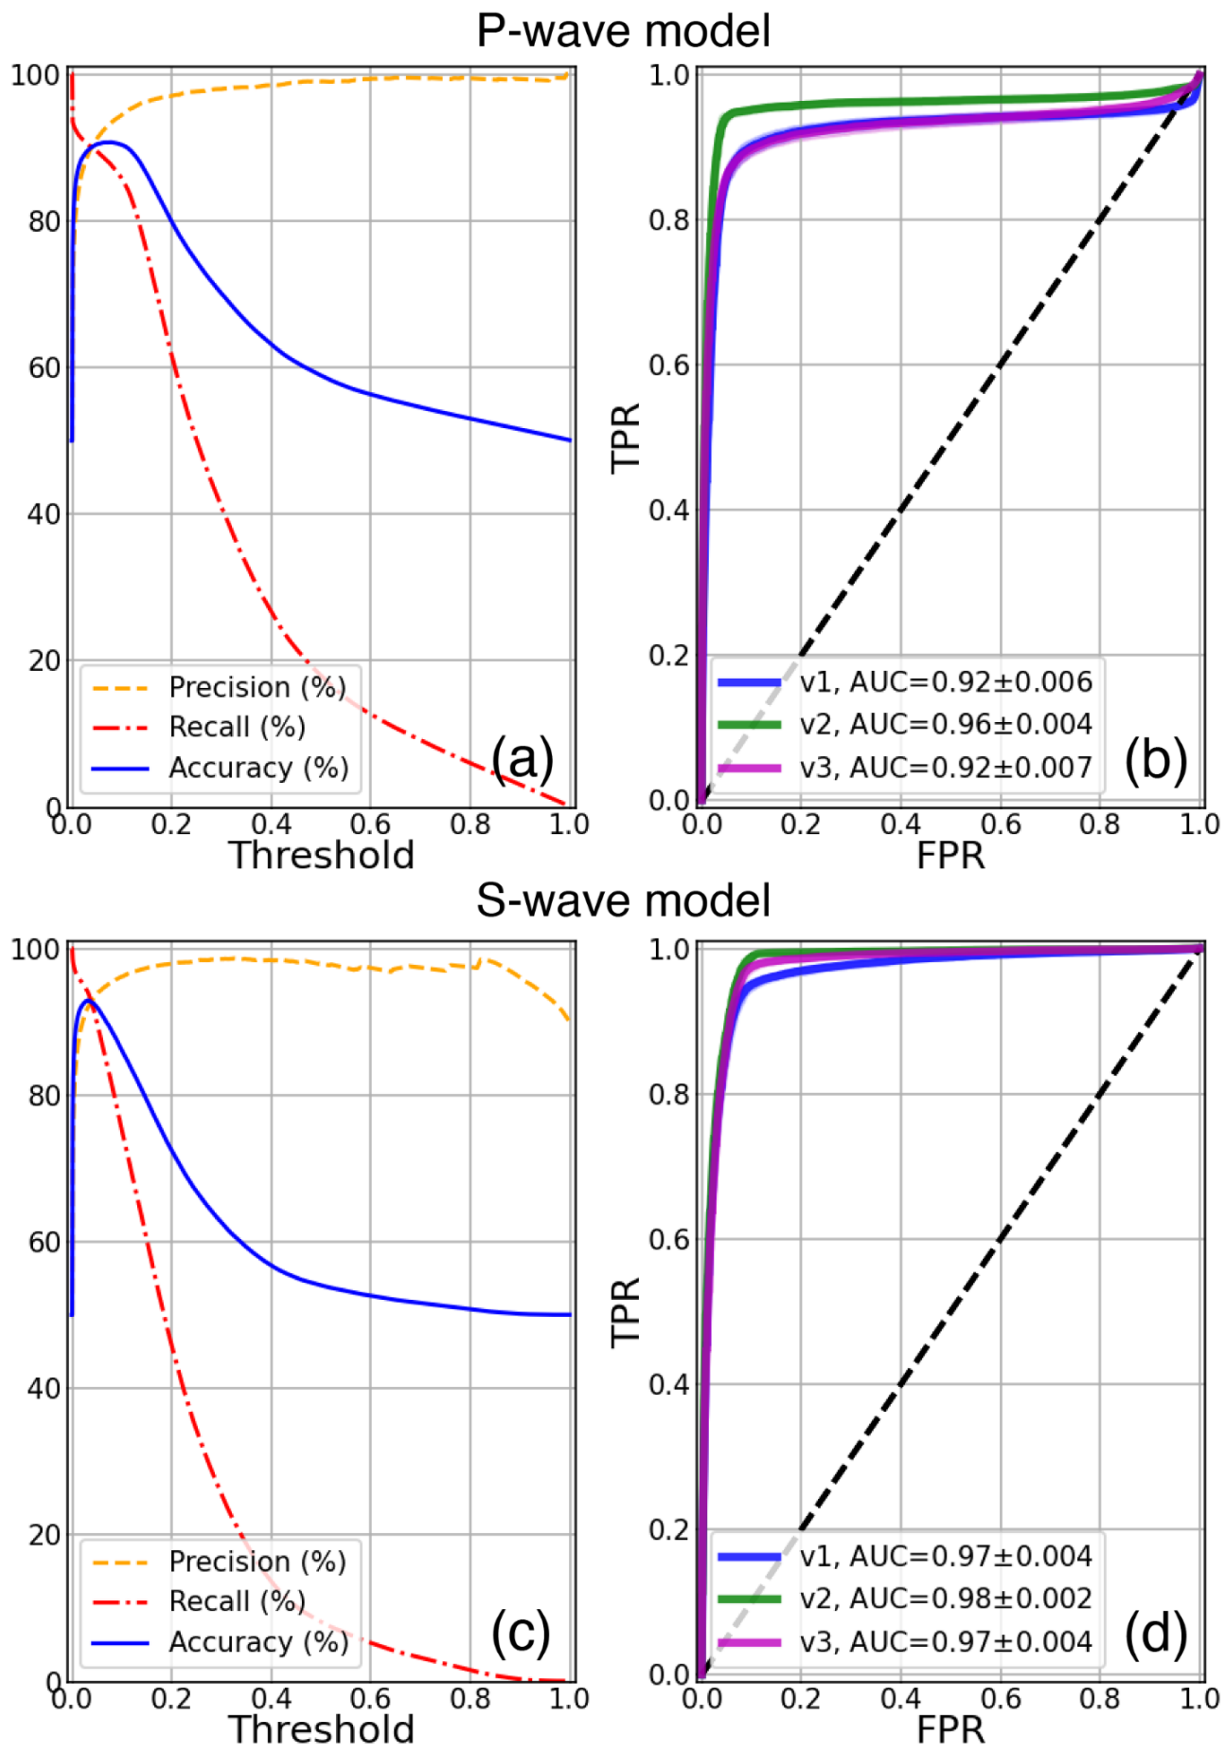


Figure 3 Performance analysis for our selected model (size=2). (a) Precision, recall, and accuracy curve as a function of decision threshold. (b) ROC curve for testing with v1: full testing data, v2: large events ($M>2.2$) only, and v3: close epicentral distance ($<30\text{km}$) events only. The AUC values and their standard deviations are calculated from 20 groups of 2,000 random samples, mixing with half (i.e. 1,000) of noise data, from the testing dataset. (c), (d), same as (a), (b) but for S-wave model.

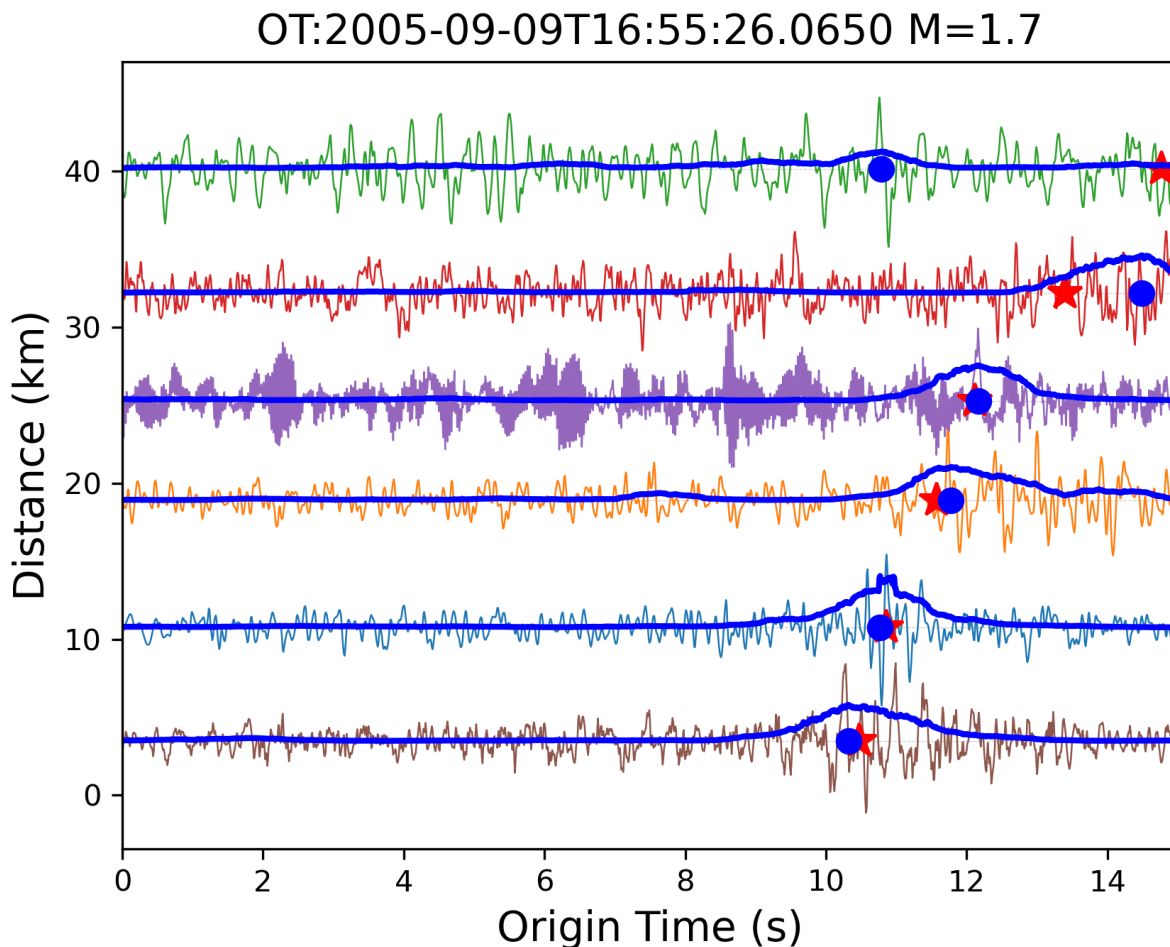


Figure 4 Example of S wave detections of a known LFE (family: 022, origin time: 2005-09-09T16:55:26.065) from testing dataset (only showing East-component). All the waveforms are normalized by their amplitude and plotted with their epicentral distance along the y-axis. Bolded lines show the model prediction. Blue dots mark the detected arrivals from the model, red stars show the actual arrivals.

the $v1$ and $v3$ tests in Figure 3. For the P-wave model, we do not find the predicted arrival time useful because the predictions are frequently mixed with the S-wave arrivals, yielding large misfits with a standard deviation of 4.2 s (Supplementary Figure S4). This is expected, as shown in Figure 2, P-wave arrivals usually have such low signal-to-noise ratio that they are rarely detected. Thus in our daily seismicity analysis presented in Figure 6 and the later location analysis, we do not include the detections from the P-wave model.

3.4 Location uncertainties

To assess the location uncertainty on our detected LFEs we perform the following sensitivity test. We define a line of 100,000 locations extending from the SE to NW at 30 km depth (Figure 8a). We then randomly select four stations, determine the S-wave arrival times from each location using the travel time grid, and add a travel time perturbation by randomly selecting a travel time shift based on the distributions of arrival time misfits shown in Figure 7c. We then grid search the location of the perturbed arrival times to find the best fit solution for each synthetic event and remove all events with $\delta > 0.5$ s. The

results of this analysis are shown in Figure 8b. We find significant scatter in individual locations with an average difference in actual and estimated location of ~ 22 km. Unfortunately these locations provide little resolution in depth since we utilize only S-waves and significant changes in source depth have similar distributions of arrival times. Considering that an LFE cluster is made up of a group of events with a similar spatial distribution (e.g. Shelly et al., 2007) averaging the locations of LFEs (i.e. centroid location) can significantly reduce location uncertainties to 10 km for $N=10$ sources (Figure 8c) and 8 km for $N=30$ sources (Figure 8d, Figure S5-S7).

3.5 LFE catalog

After associating all the detections from Figure 6 and applying the grid search approach using Equations (1)-(3), with a requirement of at least three stations for each location, we compile a catalog with 1,058,114 LFEs recorded between Jan 1, 2005 and Feb. 21, 2017. This catalog can be downloaded from Lin (2023a). The differences in detection criteria, timespans covered, and stations utilized makes a direct comparison of this catalog with either the template-matched catalog or the




Figure 5 Example of S wave detections of a new event, which is not in the original catalog of [Bostock et al. \(2015\)](#). Waveforms are normalized by their amplitude (only showing East-component). Bolded lines and dots show the model prediction and the detected arrivals from the model, respectively.

tremor catalog challenging. However, we do believe that the CNN-derived catalog contains true LFEs that were missed by the other two detection methods. For example, if we compare the template-matched and CNN-derived catalogs during September 3-26, 2005 period, the LFE catalog contains nearly double the total number of events ($N=119,064$) from the [Bostock et al. \(2015\)](#) catalog ($N=57,054$). As a proxy for the events represented in both catalogs, we determine which LFEs in the new catalog have a corresponding detection within 15 s of an LFE in the [Bostock et al. \(2015\)](#) catalog. By this metric, only 62.5% of events in the new catalog have a corresponding detection in the template-matched catalog. As mentioned above the false detection rate is $<1\%$ for events associated across three or more stations hence we believe there are many more LFEs to be discovered utilizing the CNN.

While the total number of detections varies between catalogs, time periods with (relatively) large LFE rates in the CNN-derived catalog are consistent with those in the tremor catalog ([Wech, 2021](#), Figure 9) and the original [Bostock et al. \(2015\)](#) catalog. Figure 9 shows that the CNN-derived catalog extends further back in time and has high event rates during times of known

SSEs identified by [Bostock et al. \(2015\)](#). It also has good detection rate agreement with the tremor catalog – meaning time periods when there are hundreds of daily tremor detections are in agreement with those that have thousands of LFEs daily – beginning in late 2009 until early 2014 when the LFE detection rates decrease significantly. This decrease is due to a lack of stations in our data set, with only three stations (i.e. NLLB, PFB, PGC) available. Because application of the trained CNN is not computationally intensive, the CNN can be easily applied to continuous seismic records hence the CNN-derived LFE catalog contains many LFEs that occur during inferred smaller magnitude SSEs that aren't readily apparent in surface geodetic records. There are multiple time periods over which the tremor catalog has few or no detected tremors whereas the CNN-derived catalog contains hundreds or thousands of events over 1-2 day time periods. Also, ambient LFE activity, i.e. 1 or more per day, is common.

For example, Figure 10 shows all high-quality (i.e. $\delta < 0.5s$) detections in the study area between May 4 and May 20 of 2010. In this time period the CNN detects 2,882 LFEs. In this same time period and spatial extent the tremor catalog contains only one tremor which oc-

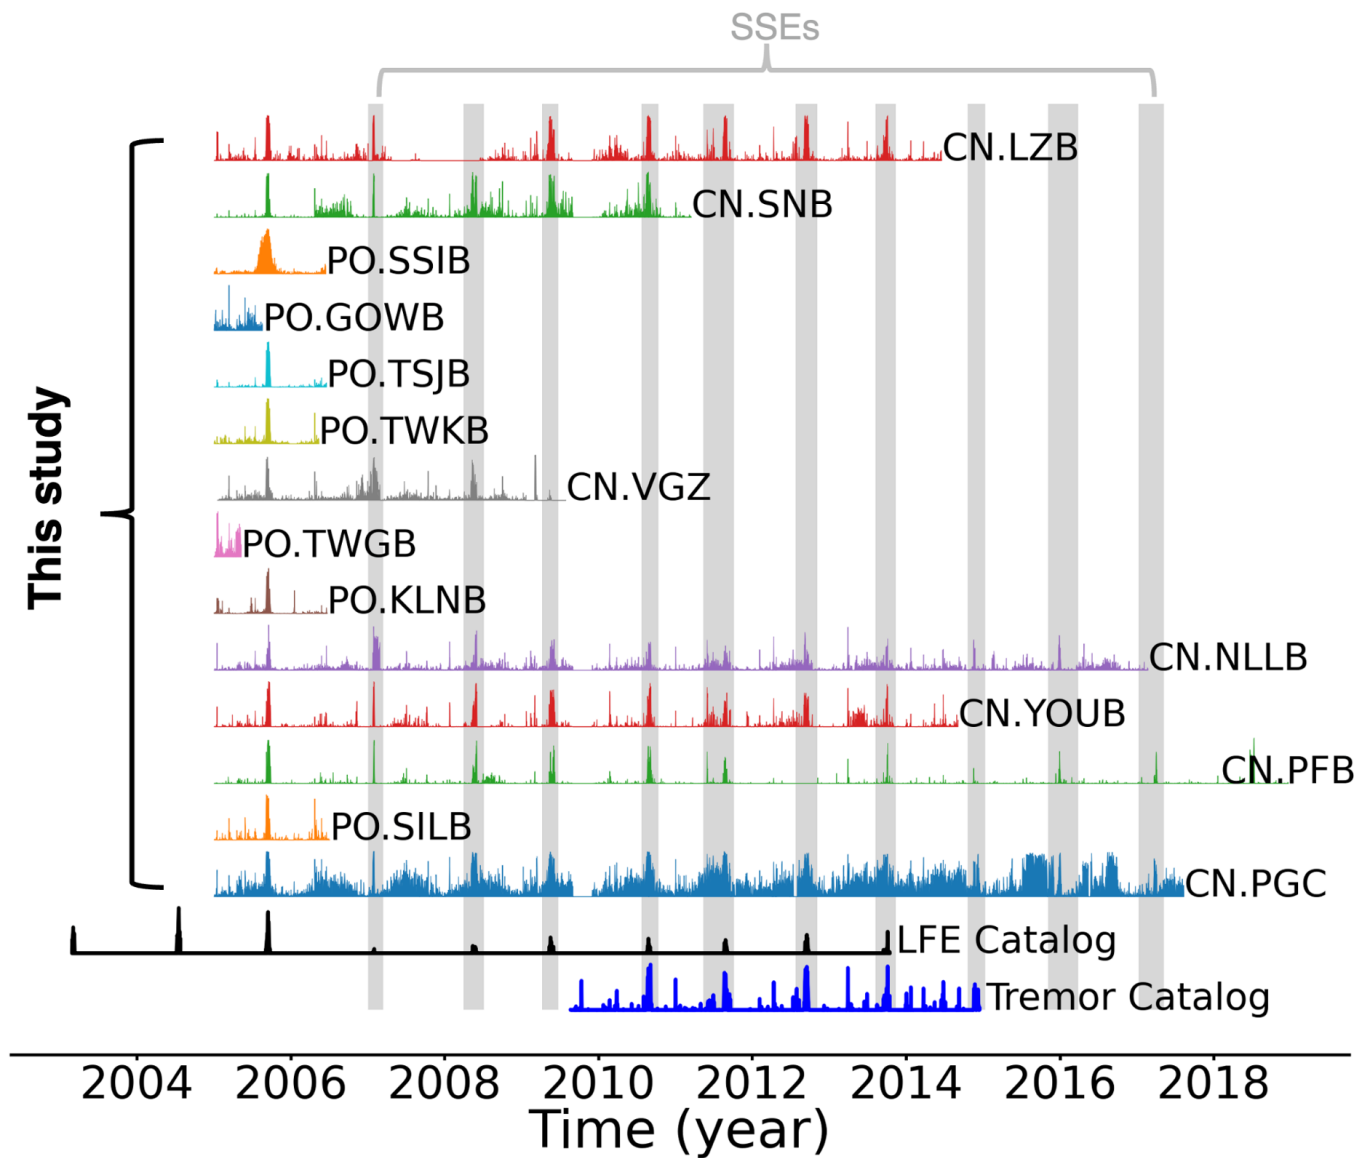


Figure 6 Model performance on 14 years of continuous data at the stations shown in Figure 1. The time series show the daily detection number for all the stations, normalized by their maximum value. The bottom rows show the original catalog from [Bostock et al. \(2015\)](#) and the tremor catalog from [Wech \(2021\)](#). Light gray shaded areas mark the SSEs in this region from [Michel et al. \(2018\)](#)

curred on May 14th (Figure 10F). The cumulative number of LFEs vs time graph shown in Figure 10D reveals a rich character with highly variable detection rates. To explore this time period further, we apply a density based clustering algorithm, DBSCAN ([Ester et al., 1996](#); [Schubert et al., 2017](#)), as implemented in the scikit-learn package ([Pedregosa et al., 2011](#)) to the detections in this time period. We set a distance threshold of 10 km and convert time to distance by scaling time by a velocity of 10 km/day. We also require a minimum of 15 LFEs in each cluster. Clusters comprise both core samples that have a minimum number of LFEs in their neighborhood and edge samples which are events within the neighborhood of core samples but do not have the minimum number of samples within their own neighborhood.

The clustering approach identifies groups of LFEs localized in space and time which are shown in Figure 10 A-C and E-F. We infer each of these clusters is generated by SSEs that produced tremors below the detection

threshold. In particular, the final event, cluster 10 (Figures 10D and G), appears to be an intermediate magnitude SSE with both a higher LFE rate (~1500 in a two-day period) and a larger spatial footprint that extends over most of the study area (Figure 10G). We confirmed that the dearth of tremor in this time period was not because of any abnormality in the tremor detection algorithm run by the Pacific Northwest Seismic Network nor is there any reason to believe the catalog is incomplete during this time period (A. G. Wech, personal communication). In fact, tremors appear to be seen in the time series data from 5 stations, spanning a maximum distance of 50 km, on May 18th (Supplementary Figure S8). In summary, we believe the CNN based detection method may have identified multiple small and one intermediate magnitude slow slip event that only produced small tremors below the detection threshold.

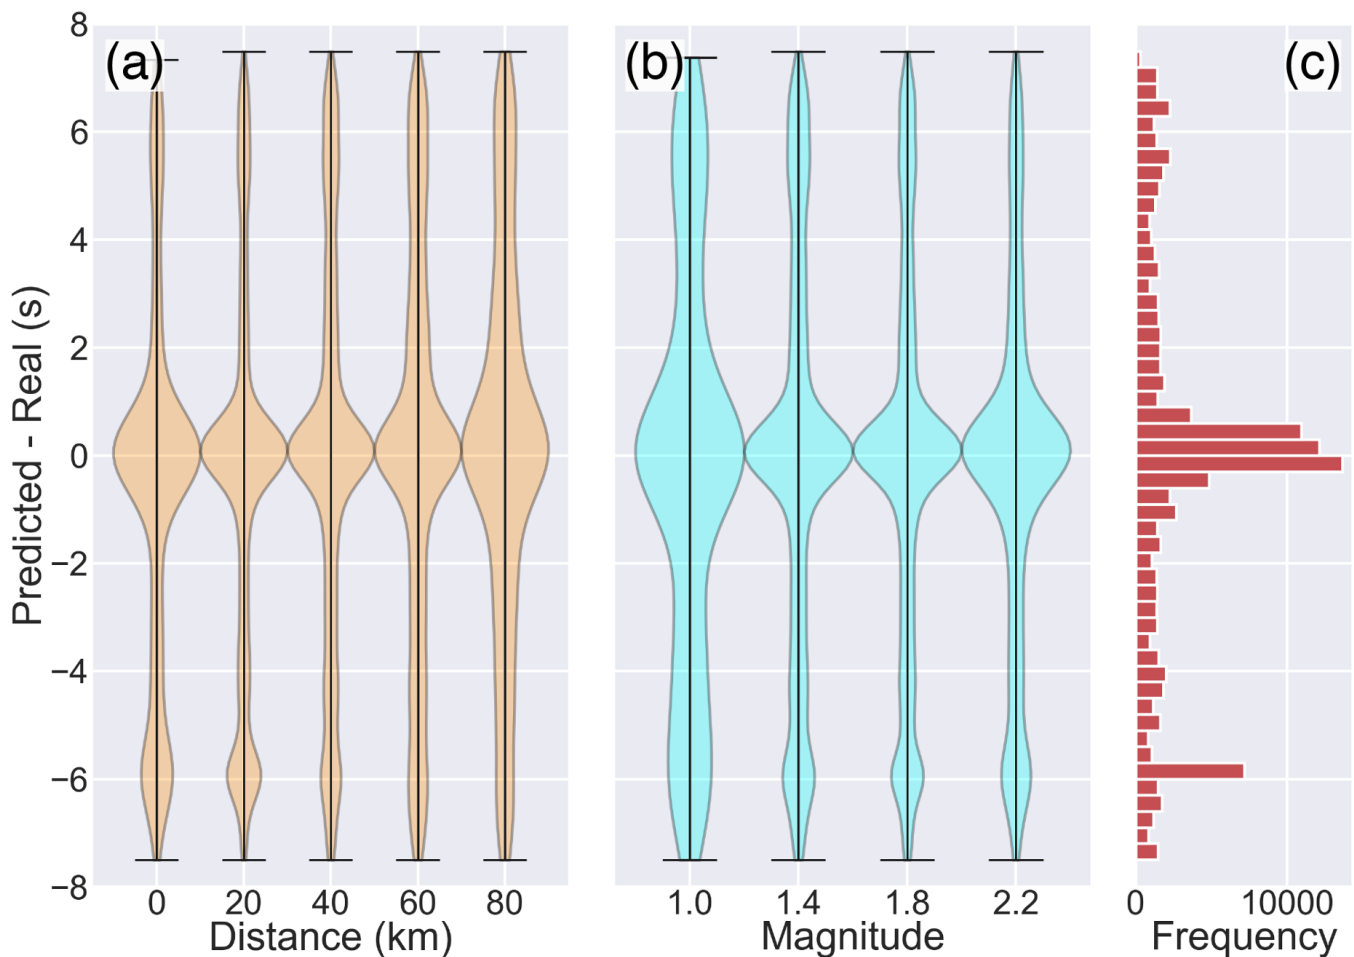


Figure 7 Distribution of S arrival time misfits in different distance and magnitude groups, evaluated by ~150,000 testing data. The model has better predictions for those close and large magnitude events.

4 Discussion

4.1 Comparison to traditional approaches

Past studies have shown that traditional template matching methods are an effective tool for identifying repeating LFEs in continuous seismic data (e.g. [Thomas and Bostock, 2015](#); [Shelly et al., 2007](#)). Despite its success, template matching has several limitations: it requires templates to be selected *a priori*, it finds only known signals and cannot extrapolate to waveforms of similar character, it requires similar station distributions through time and is computationally intensive. The CNN we develop here has several advantages over template matching. First, it is capable of identifying new and known LFEs as described above. Second, it can be applied to new stations in the same geographic region to detect existing and new LFEs. It remains to be seen how far beyond the study region the CNN can reliably identify LFEs; this likely depends on the spatial variability of the LFE source and high-frequency noise. Finally, it is computationally efficient. After training, the time complexity of the model is linear, directly proportional to the amount of data. In contrast, the computational time of the template matching method scales with both the volume of data and the number of templates.

4.2 Implications of new detections

Comprehensive analysis of the LFE catalog we generate is beyond the scope of the current work. However, the catalog appears largely consistent with the tremor catalog in that time period; relatively high detection rates in the tremor catalog correspond to times of relatively high detection rates in the LFE catalog during the time period between 2010 and 2014 (during which the two catalogs can be compared). The LFE catalog contains many examples of high LFE occurrence rates (e.g. 100s per day) over short time periods (e.g. 1-2 days). Given the high precision of 95% and the requirement that detections occur on at least 3 stations, the false detection rate is less than 1%. This suggests that the vast majority of detections are robust even though their arrivals are difficult to accurately determine. We infer that these are small magnitude SSEs that generate LFEs but did not exceed the detection threshold of the tremor's detector (i.e. Supplementary Figure S8). Additionally, cluster 10 in Figure 10D and G appears to be an intermediate magnitude SSE that was entirely unrepresented in the tremor catalog. Previous studies have suggested that there is a slip rate threshold for tremor genesis ([Wech and Bartlow, 2014](#)) so perhaps this event simply never reached sufficiently large slip speeds. Similarly, [Hulbert et al. \(2022\)](#) applied a deep learning approach to extract tremor waveforms in this region. They were able to lo-

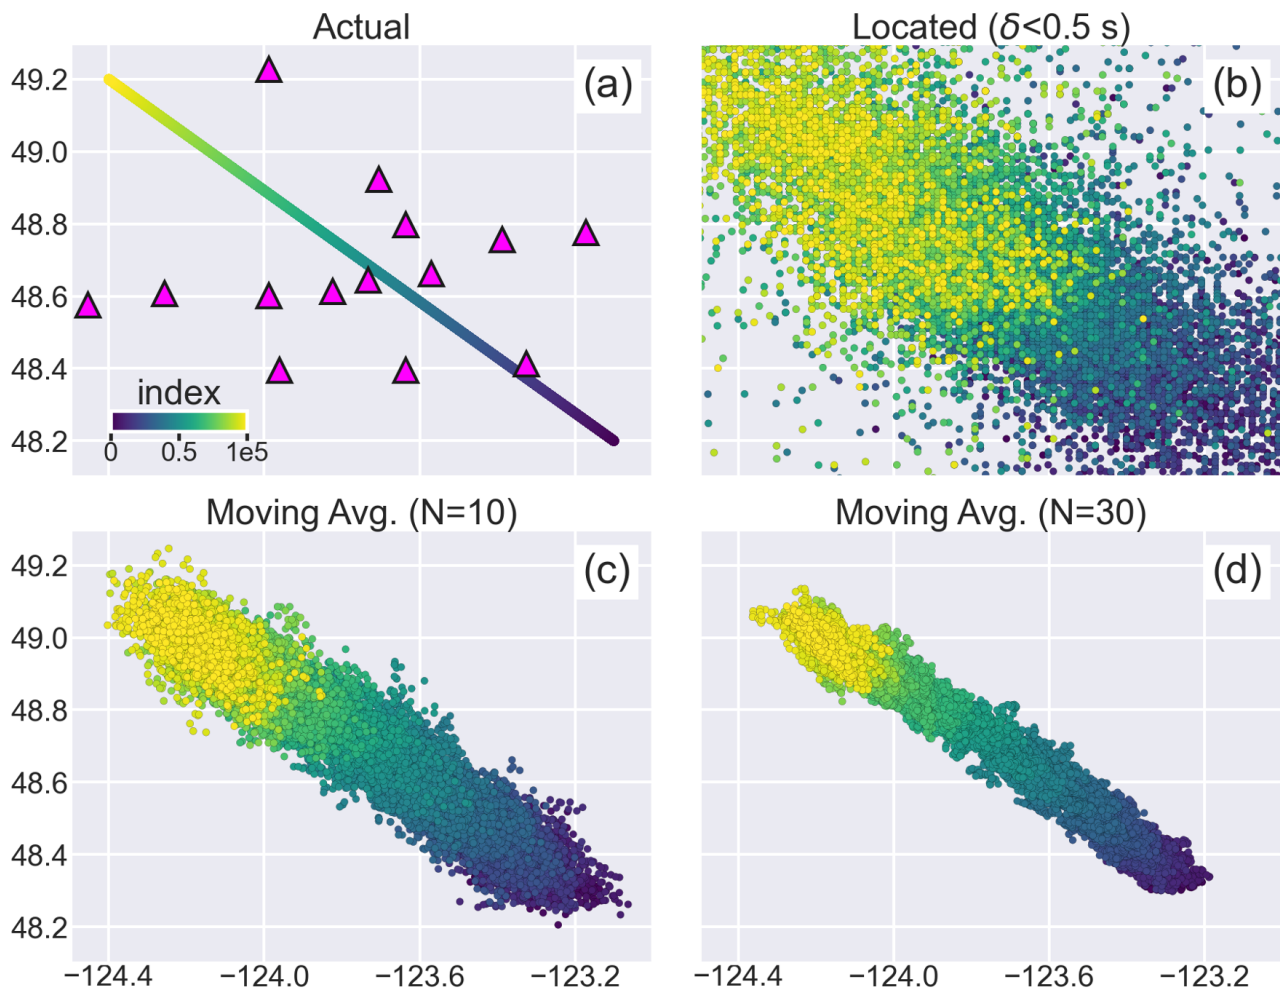


Figure 8 Location sensitivity test of 100,000 simulated events moving from SE to NW. (a) Location of the 100,000 events color-coded by their index number. All the events are set to 30 km depth, recorded by a random set of stations (triangles) ranging from a minimum of 4 stations to a maximum of 10 stations. (b) Locating result with averaged travel time residual < 0.5 s. (c) Moving average of the located result with $N=10$ sources. (d) Same as (c), but for $N=30$ sources.

cate more tremors that were not detected in the original catalog. These missing events are important for understanding SSE nucleation processes and for extending the SSE catalog to smaller magnitudes. Finally, in the time period between January 1 2010 and January 1 2014, when there are several stations available to detect LFEs, we find that only 7% of days contained no LFE detections whatsoever. This suggests that ambient LFE activity may be widespread, as has also been suggested for tremor activity (Rouet-Leduc et al., 2018).

4.3 Limitations and future opportunities

We test the model performance in a broader region using the Plate Boundary Observatory stations (PB) stations, spanning from central to southern Cascadia. The comparison with the SSEs (Michel et al., 2018) reveals that although the peak LFE detections align with the timing of large SSE occurrences, specifically for the northern stations close to southern Vancouver Island (i.e. B001-B013), the overall results are noisy and may include many false detections (Supplementary Figure S9). We note that even though the model's statistics indicate a low false detection rate, this assessment is based

on testing data recorded in southern Vancouver Island similar to the training data (i.e. the same set of stations and high signal-to-noise ratio events from the LFE catalog of Bostock et al., 2015). The accuracy in different regions will need additional assessment. In fact, noisier performance is expected due to different noise characteristics and source-station paths compared to the training data in southern Vancouver Island. Transfer learning (e.g. Chai et al., 2020) with additional training data in this region may improve the detections in this region.

The CNN is successful at identifying LFEs in continuous seismic data, however precise arrival time picks are a challenge for the detector as it routinely makes picks that are seconds different than the known LFE arrival time in the testing data. This is undoubtedly due to the low signal-to-noise character of LFEs and may also be complicated by the tendency of LFEs to occur in rapid succession. It will require additional work to accurately locate LFEs, but we anticipate that the predictions can be added as an additional constraint for a more robust detection i.e. only consider events when both the P- and S-wave are high confidence, have reasonable S-P times, and moveout consistent with a physical source. Recent

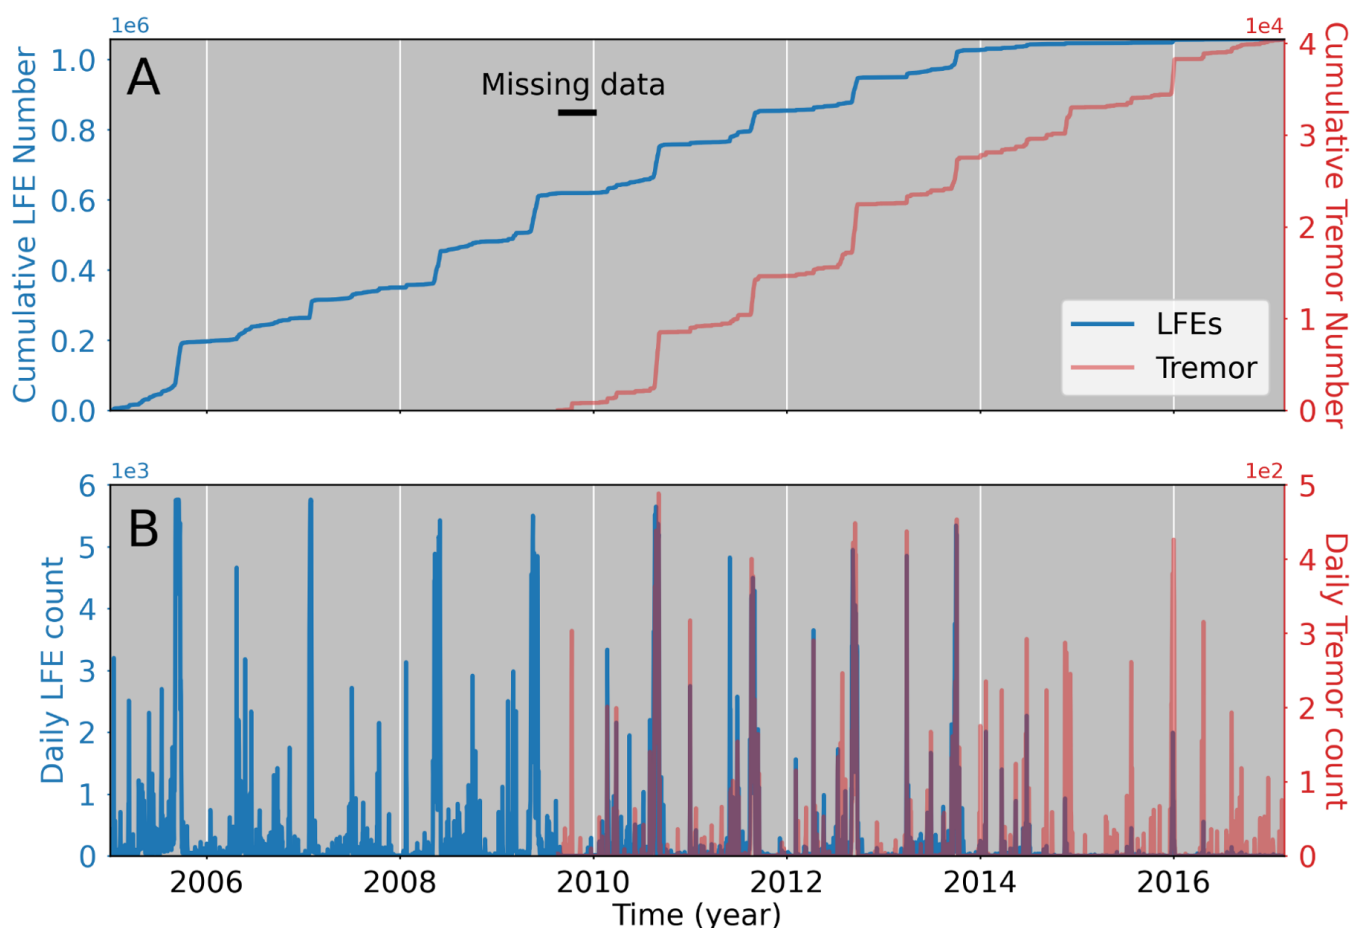


Figure 9 Panel A shows the cumulative number of LFEs in the CNN-derived catalog and cumulative number of tremors located on southern Vancouver Island from (Wech, 2021). LFEs require detections on a minimum of three stations in the same 15 s time window. Panel B shows the daily LFE and Tremor counts.

work such as grid-based Bayesian inversion may provide one solution to better locate the LFEs (Bombardier et al., 2023). Another possible solution that would simultaneously validate detected LFEs and permit precise locations is to combine the CNN with template matching by utilizing the times of associated detections as initial templates and cross correlating them with other time periods in which the CNN detector registers detections on multiple stations.

5 Conclusions

LFE activity provides a tool to track fault slip evolution during SSEs. Traditional methods for detecting LFEs are computationally expensive and they are usually limited by the assumption that sources repeat. Here we train a CNN to detect LFEs and identify their P- and S-wave arrivals in southern Vancouver Island. When applied to the testing dataset, our model has a high accuracy of 92% and 90% for discriminating S-waves and P-waves from noise at a decision threshold of 0.1, respectively. This is remarkable considering the low signal-to-noise ratio of the data. We applied the CNN to 14 years of continuous data and find that the model detects more LFEs during times of known slow slip events present in the tremor catalog. We then located the LFEs with a grid search approach in a 3D regional velocity model. The

resulting new catalog found LFEs that are not present in the tremor catalog. Notably, on May 17th, 2010, a cluster contains nearly 1500 LFEs with the locations of these events localize to a region nearly half the size of the study area. In contrast, the tremor catalog contains no detection at the same period in this area. This suggests the possibility of small magnitude SSEs that fall below the tremor detection threshold. In summary, the CNN approach to LFE detection is promising in both its efficiency and its ability to detect small amounts of seismic radiation from SSEs that does not satisfy the tremor detection criteria, providing new opportunities to understand deep subduction zone processes in this region.

6 Acknowledgment

We thank the editor, Brendan Crowell, Michael Bostock, and an anonymous reviewer for their valuable comments and suggestions, which have helped us greatly improve our manuscript. The work was funded by National Science Foundation Award 1848302 and United States Geological Survey National Earthquake Hazards Reduction Program Award G22AP00299. The work was performed in part under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. This is LLNL Contribution Number LLNL-JRNL-855845.

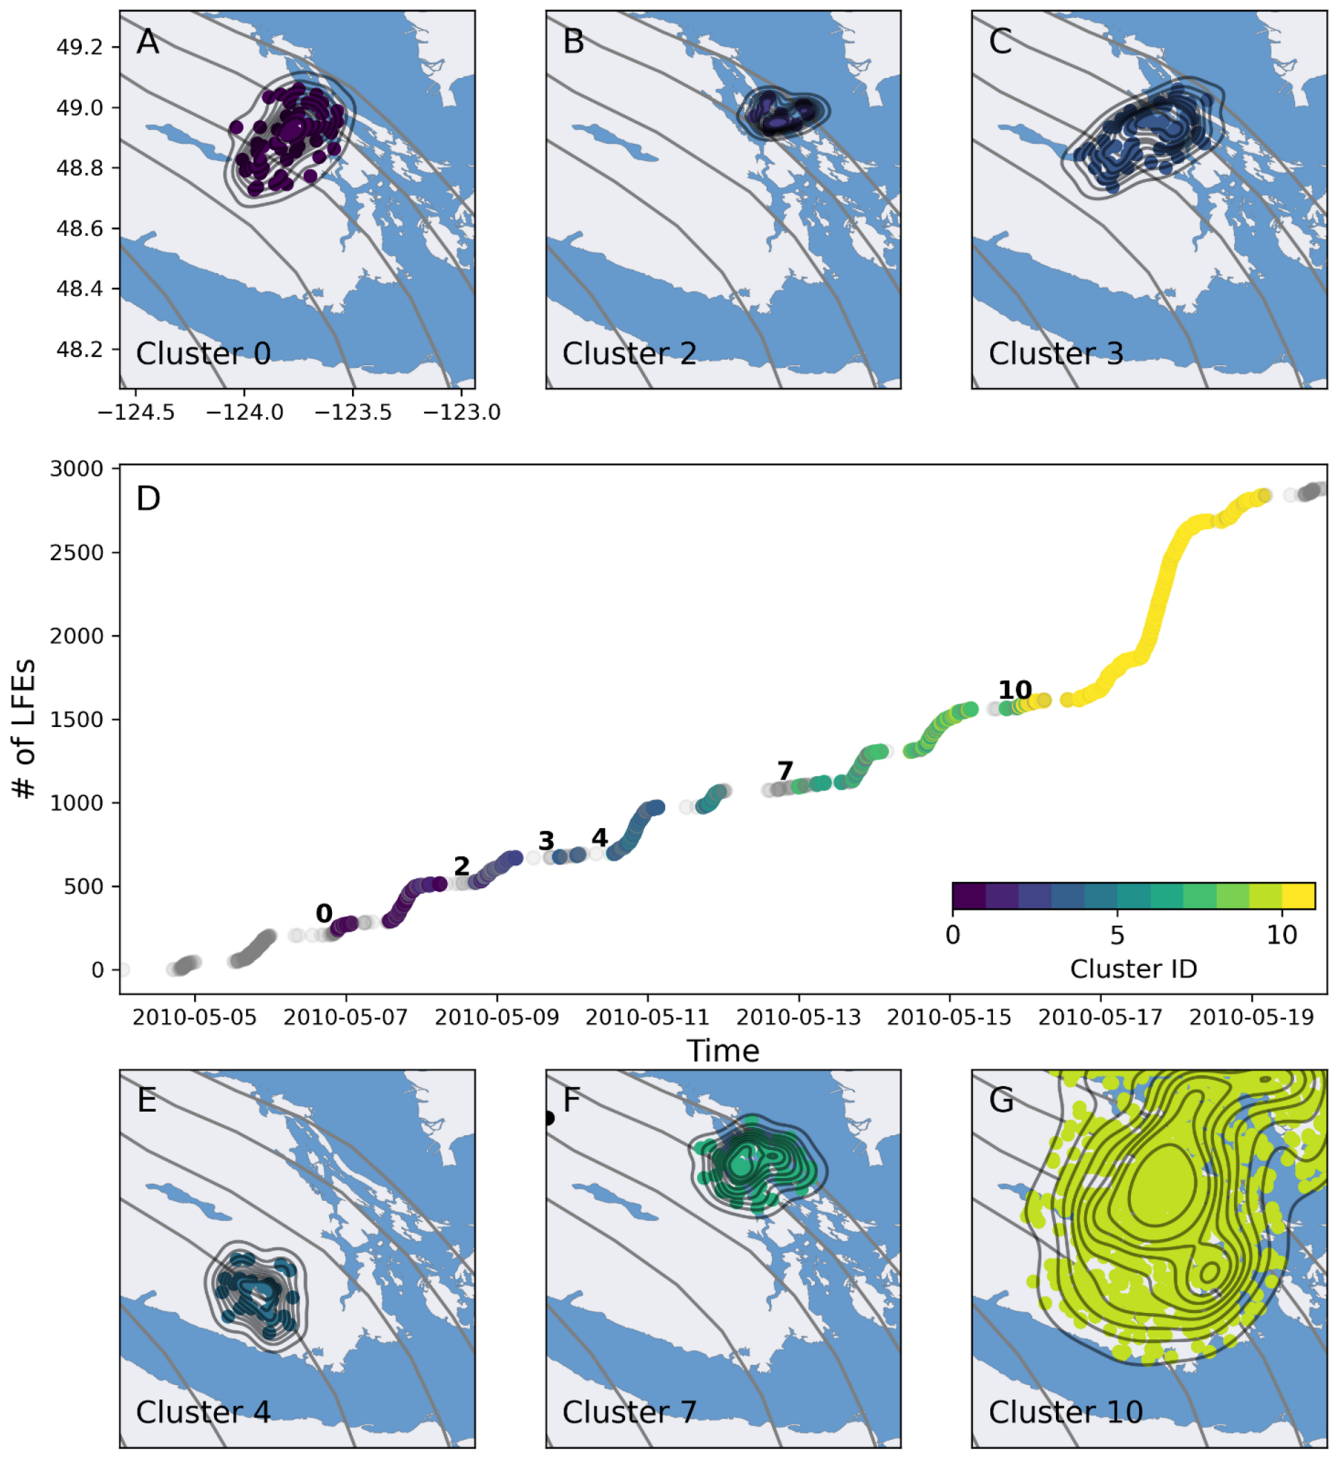


Figure 10 Panels A-C and E-G show spatial extents of LFE clusters. Black lines are density contours. Grey lines are slab isodepth contours. Panel D shows the cumulative number of LFEs as a function of time in the study area. Events are color coded by their cluster ID. This same time period contains only one tremor (shown as a black dot in Panel F) which occurred on May 14th. Geographic area is the same as in Figure 1.

7 Data and code availability

Most of the waveform data used for this study were accessed through the EarthScope Data Management Center. The CN and C8 data can be accessed from the Canadian National Data Centre. The original LFE catalog can be downloaded from the slow earthquake database (Kano et al., 2018). The codes for LFE detection were taken from Thomas et al. (2021) and are avail-

able at <https://zenodo.org/records/10076720> (Lin, 2023b). LFE catalog can be downloaded from <https://doi.org/10.5281/zenodo.10016020> (Lin, 2023a).

References

Armbruster, J. G., Kim, W., and Rubin, A. M. Accurate tremor locations from coherent S and P waves. *Journal of Geophysi-*

- cal Research: Solid Earth*, 119(6):5000–5013, June 2014. doi: 10.1002/2014jb011133.
- Behr, W. M. and Bürgmann, R. What's down there? The structures, materials and environment of deep-seated slow slip and tremor. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 379(2193): 20200218, Feb. 2021. doi: 10.1098/rsta.2020.0218.
- Beroza, G. C. and Ide, S. Slow Earthquakes and Nonvolcanic Tremor. *Annual Review of Earth and Planetary Sciences*, 39 (1):271–296, May 2011. doi: 10.1146/annurev-earth-040809-152531.
- Bombardier, M., Dosso, S. E., Cassidy, J. F., and Kao, H. Tackling the challenges of tectonic tremor localization using differential traveltimes and Bayesian inversion. *Geophysical Journal International*, 234(1):479–493, Feb. 2023. doi: 10.1093/gji/ggad086.
- Bostock, M. G., Royer, A. A., Hearn, E. H., and Peacock, S. M. Low frequency earthquakes below southern Vancouver Island. *Geochemistry, Geophysics, Geosystems*, 13(11), Nov. 2012. doi: 10.1029/2012gc004391.
- Bostock, M. G., Thomas, A. M., Savard, G., Chuang, L., and Rubin, A. M. Magnitudes and moment-duration scaling of low-frequency earthquakes beneath southern Vancouver Island. *Journal of Geophysical Research: Solid Earth*, 120(9):6329–6350, Sept. 2015. doi: 10.1002/2015jb012195.
- Bürgmann, R. The geophysics, geology and mechanics of slow fault slip. *Earth and Planetary Science Letters*, 495:112–134, Aug. 2018. doi: 10.1016/j.epsl.2018.04.062.
- Calvert, A. J., Bostock, M. G., Savard, G., and Unsworth, M. J. Cascadia low frequency earthquakes at the base of an overpressured subduction shear zone. *Nature Communications*, 11(1), Aug. 2020. doi: 10.1038/s41467-020-17609-3.
- Chai, C., Maceira, M., Santos-Villalobos, H. J., Venkatakrishnan, S. V., Schoenball, M., Zhu, W., Beroza, G. C., and Thurber, C. Using a Deep Neural Network and Transfer Learning to Bridge Scales for Seismic Phase Picking. *Geophysical Research Letters*, 47(16), Aug. 2020. doi: 10.1029/2020gl088651.
- Chamberlain, C. J., Shelly, D. R., Townend, J., and Stern, T. A. Low-frequency earthquakes reveal punctuated slow slip on the deep extent of the Alpine Fault, New Zealand. *Geochemistry, Geophysics, Geosystems*, 15(7):2984–2999, July 2014. doi: 10.1002/2014gc005436.
- Delph, J. R., Thomas, A. M., and Levander, A. Subcretionary tectonics: Linking variability in the expression of subduction along the Cascadia forearc. *Earth and Planetary Science Letters*, 556: 116724, Feb. 2021. doi: 10.1016/j.epsl.2020.116724.
- Ester, M., Kriegel, H., Sander, J., and Xu, X. A density-based algorithm for discovering clusters in large spatial databases with noise. *Proceedings of 2nd International Conference on Knowledge Discovery and Data Mining (KDD-96)*, page 226–231, 1996.
- Frank, W. B., Shapiro, N. M., Husker, A. L., Kostoglodov, V., Romanenko, A., and Campillo, M. Using systematically characterized low-frequency earthquakes as a fault probe in Guerrero, Mexico. *Journal of Geophysical Research: Solid Earth*, 119(10): 7686–7700, Oct. 2014. doi: 10.1002/2014jb011457.
- Hulbert, C., Jolivet, R., Gardonio, B., Johnson, P. A., Ren, C. X., and Rouet-Leduc, B. Tremor Waveform Extraction and Automatic Location With Neural Network Interpretation. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–9, 2022. doi: 10.1109/tgrs.2022.3156125.
- Inbal, A., Thomas, A. M., Newton, T., and Bürgmann, R. Complex Migration of Tremor Near Cholame, CA, Resolved by Seismic Array Analysis. *Journal of Geophysical Research: Solid Earth*, 126 (9), Sept. 2021. doi: 10.1029/2021jb022174.
- Kano, M., Aso, N., Matsuzawa, T., Ide, S., Annoura, S., Arai, R., Baba, S., Bostock, M., Chao, K., Heki, K., Itaba, S., Ito, Y., Kamaya, N., Maeda, T., Maury, J., Nakamura, M., Nishimura, T., Obana, K., Ohta, K., Poiata, N., Rousset, B., Sugioka, H., Takagi, R., Takahashi, T., Takeo, A., Tu, Y., Uchida, N., Yamashita, Y., and Obara, K. Development of a Slow Earthquake Database. *Seismological Research Letters*, 89(4):1566–1575, June 2018. doi: 10.1785/0220180021.
- Kao, H., Shan, S., and Dragert, H. A wide depth distribution of seismic tremors along the northern Cascadia margin. *Nature*, 436: 841–844, 2005. doi: 10.1038/nature03903.
- Kao, H., Shan, S., Dragert, H., and Rogers, G. Northern Cascadia episodic tremor and slip: A decade of tremor observations from 1997 to 2007. *Journal of Geophysical Research: Solid Earth*, 114 (B11), Nov. 2009. doi: 10.1029/2008jb006046.
- Krischer, L., Megies, T., Barsch, R., Beyreuther, M., Lecocq, T., Caudron, C., and Wassermann, J. ObsPy: a bridge for seismology into the scientific Python ecosystem. *Computational Science & Discovery*, 8(1):014003, May 2015. doi: 10.1088/1749-4699/8/1/014003.
- Lecun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791.
- Lin, J.-T. A Deep Learning-Based Low-Frequency Earthquake Catalog in Southern Vancouver Island, 2023a. doi: 10.5281/ZENODO.10016020.
- Lin, J.-T. *jiunting/Cas_LFE: First release of Cascadia LFE detection scripts*, 2023b. doi: 10.5281/ZENODO.10076720.
- Lomax, A., Michelini, A., Curtis, A., and Meyers, R. *Earthquake Location, Direct, Global-Search Methods*, page 2449–2473. Springer New York, 2009. doi: 10.1007/978-0-387-30440-3_150.
- Michel, S., Gualandi, A., and Avouac, J.-P. Interseismic Coupling and Slow Slip Events on the Cascadia Megathrust. *Pure and Applied Geophysics*, 176(9):3867–3891, Sept. 2018. doi: 10.1007/s00024-018-1991-x.
- Münchmeyer, J., Giffard-Roisin, S., Malfante, M., Frank, W., Poli, P., Marsan, D., and Socquet, A. Deep learning detects uncataloged low-frequency earthquakes across regions. *Seismica*, 3(1), May 2024. doi: 10.26443/seismica.v3i1.1185.
- Obara, K. Nonvolcanic Deep Tremor Associated with Subduction in Southwest Japan. *Science*, 296(5573):1679–1681, May 2002. doi: 10.1126/science.1070378.
- Obara, K. and Hirose, H. Non-volcanic deep low-frequency tremors accompanying slow slips in the southwest Japan subduction zone. *Tectonophysics*, 417(1–2):33–51, Apr. 2006. doi: 10.1016/j.tecto.2005.04.013.
- Obara, K. and Kato, A. Connecting slow earthquakes to huge earthquakes. *Science*, 353(6296):253–257, July 2016. doi: 10.1126/science.aaf1512.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., and Duchesnay, E. Scikit-learn: Machine learning in Python. *the Journal of machine Learning research*, 12: 2825–2830, 2011.
- Rogers, G. and Dragert, H. Episodic Tremor and Slip on the Cascadia Subduction Zone: The Chatter of Silent Slip. *Science*, 300 (5627):1942–1943, June 2003. doi: 10.1126/science.1084783.
- Ronneberger, O., Fischer, P., and Brox, T. *U-Net: Convolutional Networks for Biomedical Image Segmentation*, page 234–241. Springer International Publishing, 2015. doi: 10.1007/978-3-319-24574-4_28.
- Rouet-Leduc, B., Hulbert, C., and Johnson, P. A. Continuous chatter of the Cascadia subduction zone revealed by machine learning. *Nature Geoscience*, 12(1):75–79, Dec. 2018. doi:

- 10.1038/s41561-018-0274-6.
- Rousset, B., Bürgmann, R., and Campillo, M. Slow slip events in the roots of the San Andreas fault. *Science Advances*, 5(2), Feb. 2019. doi: 10.1126/sciadv.aav3274.
- Royer, A. and Bostock, M. A comparative study of low frequency earthquake templates in northern Cascadia. *Earth and Planetary Science Letters*, 402:247–256, Sept. 2014. doi: 10.1016/j.epsl.2013.08.040.
- Rubin, A. M. and Armbruster, J. G. Imaging slow slip fronts in Cascadia with high precision cross-station tremor locations. *Geochemistry, Geophysics, Geosystems*, 14(12):5371–5392, Dec. 2013. doi: 10.1002/2013gc005031.
- Saffer, D. M. and Wallace, L. M. The frictional, hydrologic, metamorphic and thermal habitat of shallow slow earthquakes. *Nature Geoscience*, 8(8):594–600, July 2015. doi: 10.1038/ngeo2490.
- Savard, G. and Bostock, M. G. Detection and Location of Low-Frequency Earthquakes Using Cross-Station Correlation. *Bulletin of the Seismological Society of America*, 105(4):2128–2142, June 2015. doi: 10.1785/0120140301.
- Savard, G., Bostock, M. G., and Christensen, N. I. Seismicity, Metamorphism, and Fluid Evolution Across the Northern Cascadia Fore Arc. *Geochemistry, Geophysics, Geosystems*, 19(6): 1881–1897, June 2018. doi: 10.1029/2017gc007417.
- Schubert, E., Sander, J., Ester, M., Kriegel, H. P., and Xu, X. DBSCAN Revisited: Why and How You Should (Still) Use DBSCAN. *ACM Transactions on Database Systems*, 42(3):1–21, July 2017. doi: 10.1145/3068335.
- Seno, T. and Yamasaki, T. Low-frequency tremors, intraslab and interplate earthquakes in Southwest Japan—from a viewpoint of slab dehydration. *Geophysical Research Letters*, 30(22), Nov. 2003. doi: 10.1029/2003gl018349.
- Shelly, D. R., Beroza, G. C., and Ide, S. Non-volcanic tremor and low-frequency earthquake swarms. *Nature*, 446(7133):305–307, Mar. 2007. doi: 10.1038/nature05666.
- Thomas, A. and Bostock, M. Identifying low-frequency earthquakes in central Cascadia using cross-station correlation. *Tectonophysics*, 658:111–116, 2015. doi: 10.1016/j.tecto.2015.07.013.
- Thomas, A. M., Beroza, G. C., and Shelly, D. R. Constraints on the source parameters of low-frequency earthquakes on the San Andreas Fault. *Geophysical Research Letters*, 43(4):1464–1471, Feb. 2016. doi: 10.1002/2015gl067173.
- Thomas, A. M., Beeler, N. M., Bletery, Q., Burgmann, R., and Shelly, D. R. Using Low-Frequency Earthquake Families on the San Andreas Fault as Deep Creepmeters. *Journal of Geophysical Research: Solid Earth*, 123(1):457–475, Jan. 2018. doi: 10.1002/2017jb014404.
- Thomas, A. M., Inbal, A., Searcy, J., Shelly, D. R., and Bürgmann, R. Identification of Low-Frequency Earthquakes on the San Andreas Fault With Deep Learning. *Geophysical Research Letters*, 48(13), July 2021. doi: 10.1029/2021gl093157.
- Toomey, D. R., Solomon, S. C., and Purdy, G. M. Tomographic imaging of the shallow crustal structure of the East Pacific Rise at 9°30'N. *Journal of Geophysical Research: Solid Earth*, 99(B12): 24135–24157, Dec. 1994. doi: 10.1029/94jb01942.
- Wech, A. G. Cataloging Tectonic Tremor Energy Radiation in the Cascadia Subduction Zone. *Journal of Geophysical Research: Solid Earth*, 126(10), Oct. 2021. doi: 10.1029/2021jb022523.
- Wech, A. G. and Bartlow, N. M. Slip rate and tremor genesis in Cascadia. *Geophysical Research Letters*, 41(2):392–398, Jan. 2014. doi: 10.1002/2013gl058607.
- Wech, A. G. and Creager, K. C. Automated detection and location of Cascadia tremor. *Geophysical Research Letters*, 35(20), Oct. 2008. doi: 10.1029/2008gl035458.
- Zhu, W. and Beroza, G. C. PhaseNet: A Deep-Neural-Network-Based Seismic Arrival Time Picking Method. *Geophysical Journal International*, 2018. doi: 10.1093/gji/ggy423.

The article *Detection of Hidden Low-Frequency Earthquakes in Southern Vancouver Island with Deep Learning* © 2024 by Jiun-Ting Lin is licensed under CC BY 4.0.