

JGR Solid Earth

RESEARCH ARTICLE

10.1029/2021JB023830

Deep Neural Networks for Creating Reliable PmP Database With a Case Study in Southern California

Wen Ding¹ , Tianjue Li^{2,3} , Xu Yang⁴ , Kui Ren¹ , and Ping Tong^{2,3,5} 

Special Section:

Machine learning for Solid Earth observation, modeling and understanding

Wen Ding and Tianjue Li have made equal contributions to this work.

Key Points:

- A deep-neural-network-based algorithm, PmPNet, is developed to automatically pick the Moho-reflected seismic waves PmP
- PmPNet efficiently achieves high precision and recall simultaneously to automatically identify PmP waves from a massive seismic database
- PmPNet creates a large PmP database with a total of 28,093 high-quality PmP picks in southern California

Supporting Information:

Supporting Information may be found in the online version of this article.

Correspondence to:

X. Yang,
xuyang@math.ucsb.edu

Citation:

Ding, W., Li, T., Yang, X., Ren, K., & Tong, P. (2022). Deep neural networks for creating reliable PmP database with a case study in southern California. *Journal of Geophysical Research: Solid Earth*, 127, e2021JB023830. <https://doi.org/10.1029/2021JB023830>

Received 14 DEC 2021

Accepted 7 APR 2022

Author Contributions:

Conceptualization: Xu Yang
Data curation: Tianjue Li, Ping Tong
Methodology: Wen Ding, Xu Yang, Kui Ren
Resources: Xu Yang, Ping Tong
Supervision: Xu Yang, Kui Ren, Ping Tong
Validation: Tianjue Li
Visualization: Wen Ding, Tianjue Li
Writing – original draft: Wen Ding, Tianjue Li

¹Department of Applied Physics and Applied Mathematics, Columbia University, New York, NY, USA, ²Division of Mathematical Sciences, School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore, Singapore, ³Earth Observatory of Singapore, Nanyang Technological University, Singapore, Singapore, ⁴Department of Mathematics, University of California, Santa Barbara, CA, USA, ⁵Asian School of the Environment, Nanyang Technological University, Singapore, Singapore

Abstract Recent progresses in artificial intelligence and machine learning make it possible to automatically identify seismic phases from exponentially growing seismic data. Despite some exciting successes in automatic picking of the first P- and S-wave arrivals, auto-identification of later seismic phases such as the Moho-reflected PmP waves remains a significant challenge in matching the performance of experienced analysts. The main difficulty of machine-identifying PmP waves is that the identifiable PmP waves are rare, making the problem of identifying the PmP waves from a massive seismic database inherently unbalanced. In this work, by utilizing a high-quality PmP data set (10,192 manual picks) in southern California, we develop PmPNet, a deep-neural-network-based algorithm to automatically identify PmP waves efficiently; by doing so, we accelerate the process of identifying the PmP waves. PmPNet applies similar techniques in the machine learning community to address the unbalancement of PmP datasets. The architecture of PmPNet is a residual neural network (ResNet)-autoencoder with additional predictor block, where encoder, decoder, and predictor are equipped with ResNet connection. We conduct systematic research with field data, concluding that PmPNet can efficiently achieve high precision and high recall simultaneously to automatically identify PmP waves from a massive seismic database. Applying the pre-trained PmPNet to the seismic database from January 1990 to December 1999 in southern California, we obtain nearly twice more PmP picks than the original PmP data set, providing valuable data for other studies such as mapping the topography of the Moho discontinuity and imaging the lower crust structures of southern California.

Plain Language Summary The success of machine learning in computer sciences, medical sciences, and many other fields has accelerated the implementation and development of machine learning techniques in seismology, making it possible to automatically identify seismic phases from the exponentially growing seismic data. At present, the auto-identification of later seismic phases, such as the Moho-reflected PmP waves, remains a significant challenge in matching the performance of experienced analysts. The main difficulty lies in the rare identifiable PmP waves, which makes the identification problem inherently unbalanced. In this work, by utilizing a high-quality PmP data set in southern California, we develop a deep-neural-network-based algorithm, PmPNet, to accelerate the process of identifying the PmP waves. We conduct systematic research with field data, and conclude that the PmPNet can efficiently achieve high precision and high recall simultaneously for automatically identifying the PmP waves from a massive seismic database. Applying the pre-trained PmPNet to the seismic database from January 1990 to December 1999, we have tripled the PmP data set in southern California.

1. Introduction

Continuously improving our knowledge of the crust is of fundamental importance, since it preserves a more than 3.4 Gy-old record of the planet's evolution, provides us with natural resources, and presents social challenges in the form of various natural hazards (Christensen & Mooney, 1995; Mooney, 2010). Seismic direct arrivals and those resulting from single or multiple reverberations within the crust contain important information on crustal structures and dynamics. Among them, the Moho-reflected PmP waves (Figures 1a and 1b) have been frequently used in active-source studies to constrain the mid-lower crustal structure (Mooney, 2010). However, the scarcity of identifiable high-quality PmP waves in earthquake seismograms, that is, ~1% (Table S1 in the Supporting Information S1; Li et al., 2022; Sun et al., 2008; Wang et al., 2018; Xia et al., 2007), hampers its wide-ranged

Writing – review & editing: Wen Ding, Tianjue Li, Xu Yang, Kui Ren, Ping Tong

utilities in passive-source seismic studies, especially when one needs to prepare the data through the labor-intensive manual picking work. To date, this labor-intensive picking work is inevitable.

The success of machine learning in computer sciences, medical sciences, and many other fields has accelerated the implementation and development of machine learning techniques in seismology. During the past few years, we have seen that tasks of earthquake monitoring, including detection, hypocenter location, phase identification, and arrival time picking, can be more efficiently performed by adopting machine learning techniques than traditional approaches (Beroza et al., 2021). In particular, scientists have a long interest in automatic arrival time picking due to its labor-intensive and repeating characteristics. Thanks to decades of accumulation of high-quality seismic data and dedicated data labeling work by skilled analysts, there are quite a few successful machine learning approaches recently developed for picking seismic arrivals automatically. Zhu and Beroza (2019) proposed a deep neural network-based arrival-picking method called “PhaseNet,” which can pick the P- and S-wave arrival times of local earthquakes. Wang et al. (2019) proposed a rich side-output residual network-based method called “PickNet” to pick the P- and S-wave arrivals of both local and regional earthquakes. Ross et al. (2018) developed two separate convolutional neural networks (CNNs) with each CNN for one task, dealing with the P-wave arrival time picking and the first-motion polarity identification. Later on, they presented a deep recurrent neural network-based framework named “PhaseLink,” which aims at assigning both the P- and S-wave picks to the earthquakes that generate them (Ross et al., 2019). Moreover, Garcia et al. (2021) tried to use a CNN-based method to identify the surface-reflected mantle seismic phase SS and its precursors.

Despite all the aforementioned successes of machine learning in automatic P and S phases picking and the automatic identification of global reflected/scattered seismic phases (with the frequency of 15–50 s), it is still a challenge for machines to match the performance of experienced analysts in the identification of local and regional seismic reflected waves such as PmP (with the frequency of 1–7 Hz). The main difficulty lies in the fact that the identifiable high-quality PmP waves are extremely rare in the amount of 100–10,000 (Li et al., 2022; Sun et al., 2008; Wang et al., 2018; Xia et al., 2007) compared to the P- and S-wave data in the amount of more than 100,000 (Ross et al., 2018, 2019; Wang et al., 2019; Zhu & Beroza, 2019), and the local PmP waves are much more difficult to identify in such a high-frequency band compared to the long-period global seismic phase data. Because of this, the data set for PmP phase identification is inherently unbalanced. When one handles the unbalanced data set, it contains very few samples belonging to the interested category, and it becomes tricky to decide how to weigh the loss function in machine learning algorithms. Therefore, the resulting lack of reliable hand-labeled data makes it hard to train standard machine learning models to pick the PmP phase automatically. This is unlike the P and S phase picking where one in general has enough well-labeled data to train regular neural networks.

In this study, we propose PmPNet, one deep learning algorithm enhanced from the autoencoder framework (Rumelhart & McClelland, 1987) with convolutional residual connections (He et al., 2016), to accelerate the process of identifying the PmP phase. Moreover, our algorithm can overcome the challenge caused by unbalanced training datasets in imaging and other applications (Johnson & Khoshgoftaar, 2019). Based on a high-quality PmP data set (10,192 manual picks) in southern California, We demonstrate that the proposed algorithm can achieve high precision and recall for the automatic identification of PmP waves from a large number of seismograms. Applying the developed PmPNet to the more than 30-year accumulated seismic data (Figures 1c and 1d), we update the PmP database for southern California with more than 28,000 picks.

2. Method

2.1. PmPNet Architecture

In this study, we implemented PmPNet as a convolutional residual autoencoder (RAE) with an additional prediction module.

2.1.1. Standard Autoencoders

Standard AEs are trained to regenerate their inputs. Training an AE aims to make its output as close to the corresponding input as possible. To be more specific, we treat an AE as a black box, and denote the input signal as x where x could be a seismic signal in our application. We denote the loss function as $l(\cdot, \cdot)$. The loss function could be the usual L_2 loss function $l(x_1, x_2) := \|x_1 - x_2\|_2^2$ or any alternative. With this notation, the loss of AE at

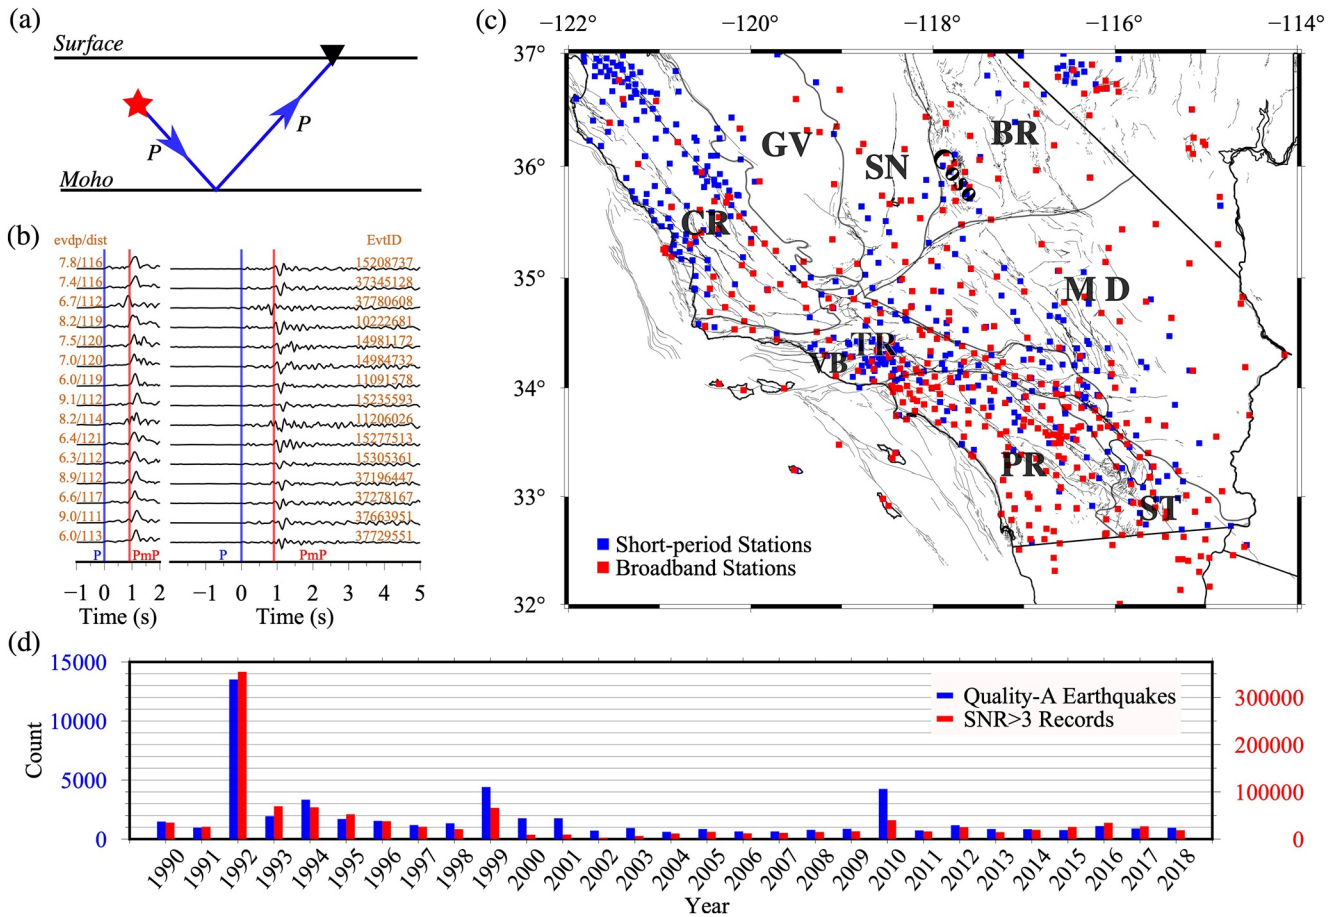


Figure 1. (a) Schematic diagram showing the raypath of PmP wave. The PmP wave (blue line) leaves the source (red star) downward, then impinges onto the Moho interface and being reflected upward, last being recorded by a surface seismic station (black triangle). (b) Typical PmP waves and their associated envelopes. These vertical-component waveforms are recorded by the broadband seismic station CI.GMR. They have been bandpass filtered (1–7 Hz). The blue line denotes the onset of the first P wave, and the red line shows the onset of PmP wave. For each waveform, its related event ID (EvtID), event focal depth (evdp) and the source-station distance (dist) are also shown at the top in orange. (c) The map shows the distribution of short-period seismic stations (blue squares) and broadband seismic stations (red squares). Main subterranean in southern California include Coast Ranges (CR), Great Valley (GV), Sierra Nevada (SN), Coso volcanic field (Coso), Basin and Range (BR), Transverse Ranges (TR), Ventura Basin (VB), Mojave Desert (MD), Peninsular Ranges (PR) and Salton Trough (ST). (d) Available seismic events each year with high-quality location, that is, <1 km in the horizontal direction and <2 km in the vertical direction (in blue) and available seismic records each year with SNR >3 (in red).

input x is $l(AE(x), x)$. In other words, AE is trained to resemble the identity operator. A more detailed structure of AE consists of two substructures: an encoder and a decoder. The encoder takes in an input x and maps the input to a latent variable z , while the decoder takes in the latent variable and maps it to the output $AE(x)$. We write the process abstractly as $z = \text{Encoder}(x)$ and $AE(x) = \text{Decoder}(z)$.

The benefit of AEs lies in their ability to perform nonlinear feature embedding from the high dimensional input space to the low dimensional latent space. Suppose that an AE is well trained (in the sense that x is indeed very close to $AE(x)$, in appropriate metric). Then $z = \text{Encoder}(x)$ and $x \approx \text{Decoder}(z)$. Therefore knowing either x or z is the same as knowing both. In most cases, the dimension of z is far lower than x . Therefore z is a low dimensional representation of original input x . Due to this simple fact that x and z contain the same amount of information (since we could reproduce one from the other), replacing the high-dimensional x with the low-dimensional z is possible and is computationally beneficial (again, since z has a lower effective dimension).

2.1.2. Overview of PmPNet Structure

In our construction, the input of PmPNet is a one-dimensional vector of length 297. It combines three parts: *envelope*, *dist*, *evdp*. The length 281 *envelope* is the normalized envelope of the vertical component of seismic signals,

which has been resampled at 40 Hz and covers the time window from 2 s before to 5 s after the observed P-wave arrival. *Dist* refers to the epicentral distance, and *evdp* refers to the focal depth, both being repeated 8 times and concatenated to the end of the signal. The number of duplication eight is chosen according to our experiments, and the final results presented in this paper are not sensitive to this selection. To summarize, the input of PmPNet is of the form:

$$\text{input } x = (\underbrace{\text{envelope}_1, \dots, \text{envelope}_{281}}_{297}, \underbrace{\text{dist}, \dots, \text{dist}}_8, \underbrace{\text{evdp}, \dots, \text{evdp}}_8) \quad (1)$$

PmPNet outputs three quantities: (a) the recovered input, (b) the PmP travel time t (a positive real number), and (c) the PmP probability p , a real number in $[0, 1]$ representing the probability that the input seismic signal contains a PmP phase.

PmPNet includes three major substructures: an encoder, a decoder, and a predictor. Similar to a standard AE, the encoder and the decoder are trained to regenerate the input. We train the predictor to read the latent variable (generated by the encoder) to predict the PmP probability p and travel time t . To be more precise, let \tilde{x} be the output of the decoder, that is, the recovered input. Then $z = \text{Encoder}(x)$, $\tilde{x} = \text{Decoder}(z)$, and $(p, t) = \text{Predictor}(z)$. Graphical data flows inside the PmPNet are described in Figure 2.

2.1.3. Substructures of PmPNet

Each substructure of PmPNet is a convolutional neural network with residual connections. Here are more details on each of the substructures, as illustrated in Figure 2:

1. The encoder includes several convolutional layers (Lecun et al., 1999), batch normalization layers (Ioffe & Szegedy, 2015), ReLU layers (Nair & Hinton, 2010), pooling layers (LeCun et al., 1990), and n_1 encoder ResNet blocks (He et al., 2016).
2. The decoder consists of several de-convolutional layers, batch normalization layers, ReLU layers, pooling layers, and n_2 decoder ResNet blocks.
3. The predictor involves several convolutional layers, batch normalization layers, ReLU layers, pooling layers, and n_3 predictor ResNet blocks as well as fully connected layer to make PmP probability and travel time prediction.

The exact number of blocks n_1 , n_2 , and n_3 are determined through benchmarking the performance of PmPNet when training with real data. We refer interested readers to (Schmidhuber, 2015) for a quick overview of the various components of the network we used in our construction.

2.2. Data Set Preparation

We have developed a two-stage workflow of identifying and picking PmP waves in a semiautomatic way (Li et al., 2022). Briefly speaking, the two-stage workflow includes two parts: At the first stage, high-quality PmP waves are automatically picked on selected seismograms. At the same time, a visual check on three-component waveforms is conducted to confirm that the chosen signals indeed come from Moho reflection. At the second stage, the volume of the PmP data set is expanded by involving other waves traveling along similar paths as those picked at the first stage. By utilizing the newly developed two-stage workflow, we have built the first PmP database with 10,192 PmP waves from the broadband vertical-component seismic data retrieved from southern California Earthquake Data Center (SCEDC, 2013). Here, we extract 5,000 waveforms that contain PmP waves from the PmP database and assign a PmP label to each waveform. We also find additional 100,000 waveforms that do not contain obvious PmP waves and assign a non-PmP label to each waveform. Thus our training data contain those labeled with either PmP or non-PmP with the data amount ratio of 1:20 to mimic the rare high-quality PmP waves in the real case. The used seismic data spanning from January 2000 to July 2010 are triggered by 6,636 local earthquakes with a magnitude between 2.0 and 5.0, epicentral distance (*dist*) between 50 and 200 km, and focal depth (*evdp*) shallower than 20 km. We have resampled the raw seismic waveforms to a uniform rate of 40 Hz and bandpass filtered (1–7 Hz) the waveforms to prepare the network's input. We then change the velocity

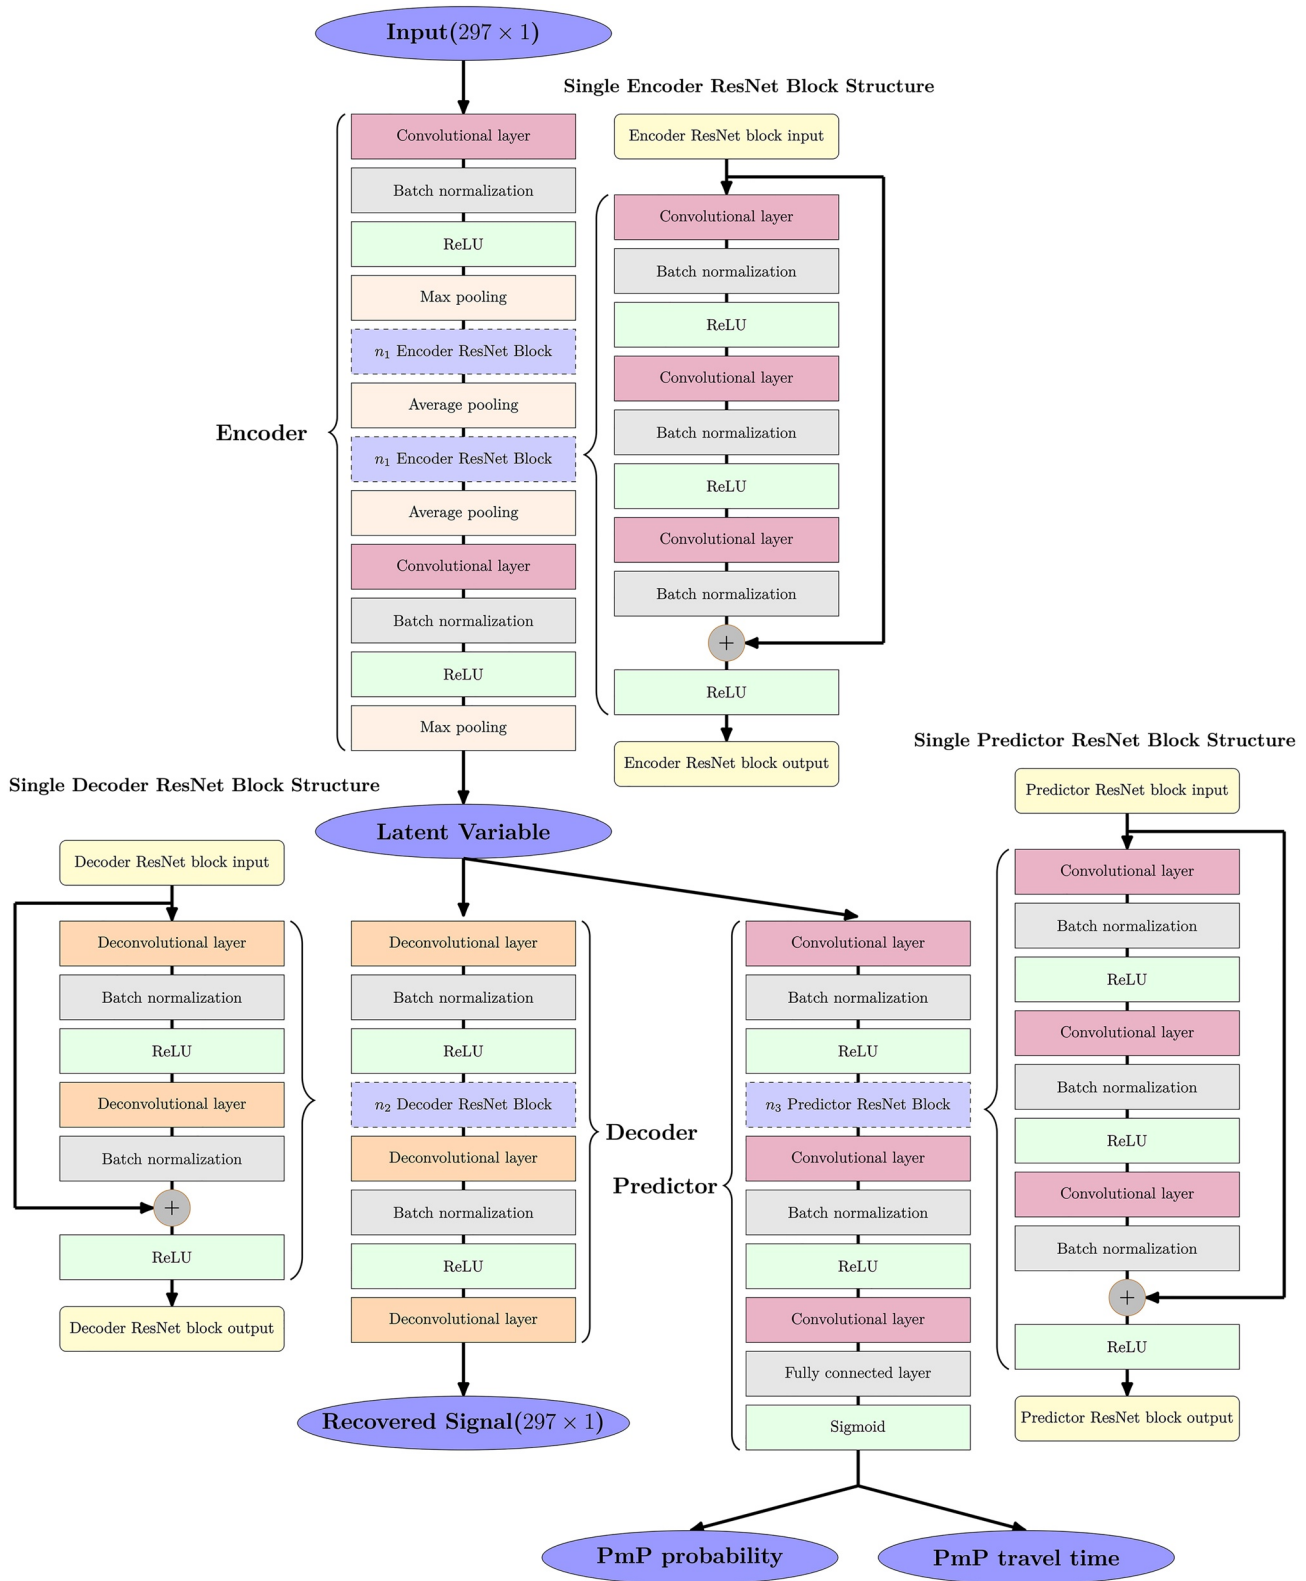


Figure 2. The main architecture of PmPNet and the data flow inside the network: The input of PmPNet is a one-dimensional vector combining three parts: signal envelope, epicentral distance, and focal depth; PmPNet outputs three quantities: the recovered signal (including signal envelope, epicentral distance, and focal depth), the PmP probability p and the PmP travel time t ; PmPNet includes three major substructures: an encoder, a decoder, and a predictor.

seismograms into envelopes, and only the portion of an envelope in a time window from 2 s before to 5 s after the first P is used. The PmP travel times for the waveforms labeled with PmP are manually picked, whereas for those labeled with non-PmP, the associated PmP travel times are calculated by using a uniform 1D P-wave velocity model (the HK model, Hadley & Kanamori, 1977) with a fixed Moho at 30 km depth.

Before the training data set is fed to PmPNet, we first normalize each sample of the input envelope to have a maximum one, and then independently standardize each time step of the input envelope to mean zero and variance one, and then linearly normalize the dist, evdp features into the range of [0, 1].

2.3. Dealing With the Issue of Unbalanced Data

With the help of an expert hand-picked PmP phase data set, we can formulate the PmP phase-detection problem as a supervised classification problem where the objective is to train a binary classifier that would allow us to divide a given set of seismic signals into a category of signals with a PmP phase and another category without it. The main challenge in the training of the classifier is that the training data set contains a significantly larger amount of signals with no PmP phase, the majority class, than signals with a PmP phase, the minority class. This unbalance of the data set is intrinsic since PmP waves are less abundant than the P and S waves or maybe equally abundant but less frequently identified by domain experts. Naive classifiers trained with such an unbalanced data set will typically over-classify the majority class due to their excessive exposure to the majority samples. This means that the commonly used predictive accuracy is not a sensible performance measure for a classifier trained with unbalanced data. To give an example, consider a PmP data set with 0.1% positive PmP phase signals. Any naive PmP classifier can achieve 99.9% accuracy on this data set by simply labeling all signals as non PmP.

Developing methods to handle the challenge of the unbalanced data set is an important research topic in the machine learning literature; see for instance (Guo et al., 2017; Johnson & Khoshgoftaar, 2019) for recent reviews on the subject. In this work, we adopt an algorithmic level method to handle the issue of unbalanced data. The method is based on precision and recall values, instead of predictive accuracy, as the evaluative metric for our PmPNet to be trained.

2.3.1. Performance Measures for Classification With Unbalanced Data

To introduce the definition of precision and recall, we define the true positive (TP), false positive (FP), true negative (TN), and false negative (FN) of the PmPNet prediction from a seismic signal as follows, respectively. TP denotes the event that the PmPNet predicts the signal has a PmP phase and the prediction is correct (i.e., the signal indeed has a PmP phase). FP denotes the event that PmPNet predicts the signal has a PmP phase, but the prediction is incorrect (i.e., the signal actually has no PmP phase). TN denotes the event that PmPNet predicts the signal has no PmP phase and the prediction is correct. FN denotes the event that PmPNet predicts the signal has no PmP phase, but the prediction is incorrect (meaning that the signal actually has a PmP phase). We recall the following three metrics based on TP, FP, TN, FN:

- Precision: $= \frac{TP}{TP+FP}$ measures the proportion of the PmP labeled samples by PmPNet that actually contain PmP phase.
- Recall: $= \frac{TP}{TP+FN}$ measures the proportion of actual PmP phase signals that are correctly labeled to be PmP by PmPNet.
- F1 score: $= 2 \frac{\text{Precision} * \text{Recall}}{\text{Precision} + \text{Recall}}$ as a balanced measurement of both precision and recall.

A perfect PmP classifier would have a 100% precision and 100% recall simultaneously. However, in reality, there is a trade-off between precision and recall. If we aim at a high recall PmP phase classifier, we would want to lose as few actual PmP waves as possible, but the risk of mistaking non-PmP wave as PmP is higher, hence lowering the precision. Similarly, if the goal is to train a high precision PmP phase classifier, then we would want our labeled PmP to have a very high probability of being an actual PmP wave, that is, the “concentration” of PmP wave from the original data set, but then we could be at risk of losing a larger proportion of actual PmP wave, hence lowering the recall.

2.3.2. Loss Function for Classification With Unbalanced Data

In the training process of PmPNet, we appropriately adopt weighted binary cross-entropy loss function for classification to reduce ultimate prediction bias toward the majority class, that is, the class of signals with no PmP phase. The true label p_{true} of datum x is either 0 or 1. We force the value of $p(x)$ to be between 0 and 1, which represents the probability that x has a PmP phase. The loss for classification purpose for that particular prediction under the weighted binary cross-entropy loss is set to be:

$$\text{loss}_c(p(x); p_{\text{true}}) := -\omega p_{\text{true}} \log(p(x)) - (1 - p_{\text{true}}) \log(1 - p(x)) \quad (2)$$

Here $\omega > 0$ is a parameter for emphasizing precision or recall for the training process, where large ω increases the recall for PmPNet, whereas small ω increases the precision for PmPNet. When $\omega \gg 1$, the loss function will emphasize $-p_{\text{true}} \log(p(x))$ term, the resulting effect is that $p(x)$ is pushed toward 1 when $p_{\text{true}} = 1$, that is, thus increase the recall. On the contrast, when $\omega \ll 1$, the loss function will emphasize $-(1 - p_{\text{true}}) \log(1 - p(x))$ term, therefore $p(x)$ is pushed down to 0 when p_{true} is 0, hence increase precision. In our training practices in the rest of the paper, we set $\omega = 20$ to be the unbalance coefficient of the data set to balance the importance of precision and recall. The training and validation results do not change much when we slightly change ω .

2.4. Loss Function for PmPNet

We are now ready to formalize the PmPNet training process. Given a training input set $\{x_i\}_{i=1}^N$ with N data points, the corresponding PmP label $\{p_{\text{true},i}\}_{i=1}^N$ and PmP travel time $\{t_{\text{true},i}\}_{i=1}^N$, we can optimize a PmPNet with trainable parameters set θ . Three different loss functions (l_1, l_2, l_3) are utilized for encoder-decoder, classification and travel time training respectively. The total training loss of PmPNet is the sum of three individual losses

$$\text{Loss}(\theta) := \frac{1}{N} \sum_{i=1}^N [l_1(\tilde{x}_i(\theta, x_i), x_i) + l_2(p_i(\theta, x_i), p_{\text{true},i}) + l_3(t_i(\theta, x_i), t_{\text{true},i})] \quad (3)$$

The variables involved are summarized as follows:

1. x_i is the input datum and $\tilde{x}_i(\theta, x_i) := \text{Decoder}_\theta(\text{Encoder}_\theta(x_i))$ is the recovered datum by the encoder-decoder pair.
2. $p_{\text{true},i}$ is the true PmP label picked by experts, which is either 0 or 1. Here $p_{\text{true},i} = 1$ means that x_i has a PmP phase, while $p_{\text{true},i} = 0$ means that x_i does not have a PmP phase.
3. $t_{\text{true},i}$ is the true PmP travel time, which is either manually picked for those labeled with PmP or theoretically computed by using the HK model for those labeled with non-PmP, as described in Section 2.2.
4. $(p_i, t_i) := \text{Predictor}_\theta(\text{Encoder}_\theta(x_i))$ with $p_i(\theta, x_i)$ and $t_i(\theta, x_i)$ being the PmP probability and the PmP travel time respectively.
5. The first component of the loss function $l_1(x, \tilde{x}) := \|x - \tilde{x}\|_1$ is the L_1 loss between input datum and recovered datum.
6. The second component of the loss function $l_2(p, p_{\text{true}}) := \text{loss}_c(p; p_{\text{true}})$ is the weighted cross-entropy loss discussed above in Equation 2 with the weight ω chosen to be 20.
7. The third component of the loss function $l_3(t, t_{\text{true}}) := |t - t_{\text{true}}|$ is absolute difference between the true travel time and predicted travel time.

We emphasize that \tilde{x}_i, p_i , and t_i are outputs of PmPNet, and therefore depend on the input datum x_i as well as the network parameters θ .

2.5. Training and Validation

Besides data preparation, the most critical step in the learning approach we proposed is the training of PmPNet. For the sake of reproducibility, we provide here the hyperparameters we used in the training process.

1. Learning rate: We initialize it as 0.001 with exponential decay of factor 2 for every 10 epochs.
2. Number of ResNet blocks for encoder, decoder and predictor are set respectively as $n_1 = 2, n_2 = 4$, and $n_3 = 1$.

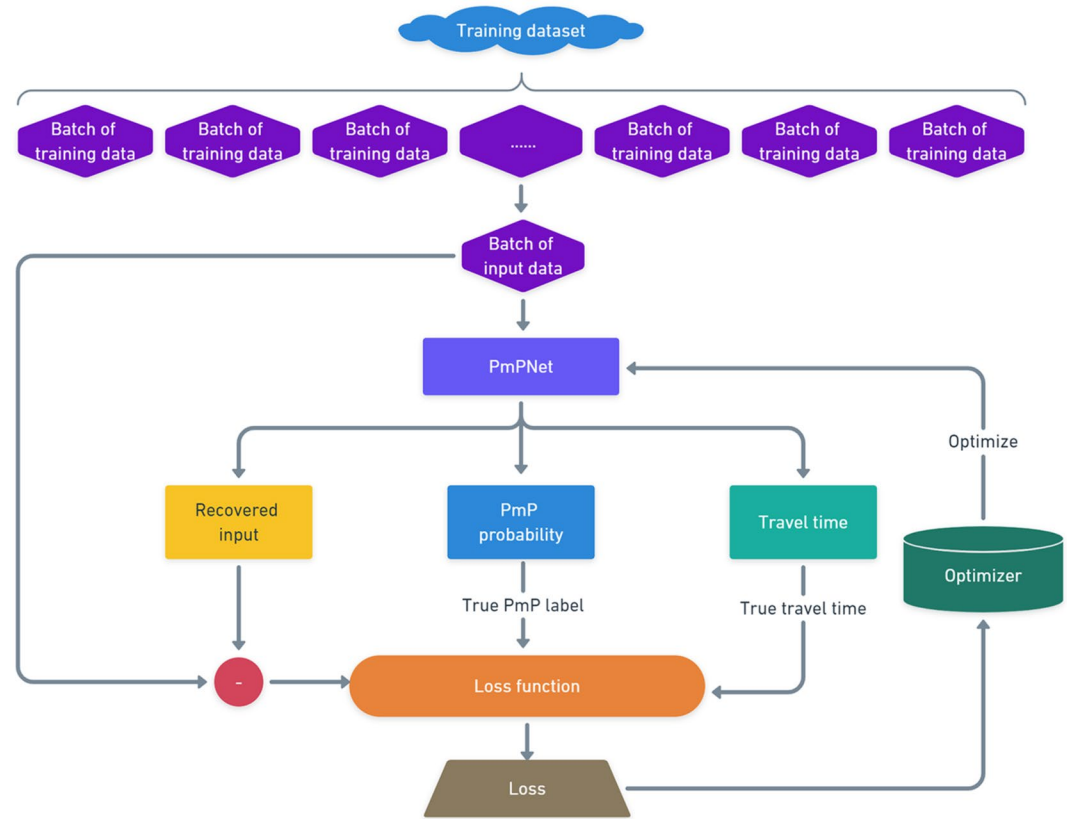


Figure 3. PmPNet training flow: (a) one batch of data points is fed into PmPNet, and the loss between the PmPNet output and the true labels is computed; and (b) the optimizer reads in the loss and update the trainable parameters of PmPNet. One epoch of training consists of continuing this iteration until the whole data set has been traversed. The training phase for PmPNet is complete when the pre-selected maximum number of epochs is reached.

3. Formation of the training set, that is, number of training data points and PmP/non-PmP ratio: we use 5, 000 PmP-labeled waveforms and 100, 000 non-PmP labeled waveforms.
4. Batch size (i.e., the number of training data points utilized in one iteration) is set as 200 in our experiments.
5. Number of epochs (i.e., the number of rounds of passing entire training data set into PmPNet) is set to be 80 in our experiments.
6. Training optimization algorithm: we use the Adam stochastic optimization algorithm (Kingma & Ba, 2017).

Once all these hyperparameters are chosen, PmPNet is fully determined by trainable parameters θ . Trainable parameters in batch normalization layers are initialized with mean 0 and variance 1. Trainable parameters in all other layers are initialized with Gaussian distribution with means 0 and variances depending on the width of the layers. After we initialize all layers, training PmPNet is essentially searching θ that minimize the loss function (Equation 3) for PmPNet.

To start the training process, we randomly choose a batch of data points without replacement, and feed the data into PmPNet. PmPNet will output the corresponding reproduced signal as well as the (p, t) pair with the current trainable parameter θ . After computing the loss, we apply one Adam step to update the trainable parameters θ . We then randomly choose another batch of data points without replacement to move another training step forward. This process continues until the whole training data set is traversed. The training process described is often called one “epoch”. Next, we adjust the learning rate to repeat the process for another epoch. The training process is complete when the predetermined number of maximum epochs are reached. The training process for PmPNet is illustrated in Figure 3.

Following common practices in the machine learning community, we perform a training-validation cycle on our PmPNet. Before the training process starts, we randomly split the input data set (described above in Section 2.2) into a training set (which is 80% of the original data set) and a validation set (which is 20% of the original data set). The training set and the validation set have no intersection, meaning that no data points in the validation set are present in the entire training process. During the validation process, we input every data point in the validation set into the post-trained PmPNet, and it outputs the corresponding t , p and recovered input. We visually compare the recovered input and the input and compute the absolute difference between predicted and true travel times. In order to make classification, we set a probability threshold $p_{\text{threshold}}$, that is, any signal with $p \geq p_{\text{threshold}}$ is predicted as PmP wave by PmPNet. Otherwise, it is predicted as a non-PmP wave. Given one probability threshold, we can compute one pair of precision and recall.

2.6. PmPNet Performance Validation

To systematically characterize the performance of PmPNet, we construct the precision-recall curve for the validation data set. To do that, we evaluate the precision-recall value corresponding to each classification threshold probability $p_{\text{threshold}}$. We then vary the probability threshold to get the precision-recall value as a function of the threshold $p_{\text{threshold}}$. The curve of precision versus recall is the precision-recall curve. This curve can help us choose the optimal value of $p_{\text{threshold}}$ for future use.

The validation performances of PmPNet are as follows. The proposed PmPNet can reach high precision (96.6%) and recall (85.3%) simultaneously, hence a high F1 score on validation set, see Table 1 for detailed precision-recall pairs and F1 score with different probability threshold and Figure 4b for the corresponding precision-recall curve. The average travel time absolute difference is around 0.33s, while maximum difference constantly stays within 5s (Figure 4c). The recovered input can capture most of the patterns from the input signal (Figures 4d and 4e), which indicates the latent variable is indeed a good representation of the input. Details of the validation performance of PmPNet are listed in Figure 4.

3. Real Applications

For the PmP database spanning from January 2000 to December 2018 built by the two-stage PmP-picking workflow (Li et al., 2022), we observe that travel times of the PmP waves increase linearly with respect to epicentral distance, amplitudes of the PmP waves are mostly larger than the P waves and particle motions of most PmP waves polarize in a similar way as the P waves. Here, we apply the trained PmPNet to the same 19-year long vertical-component seismic data to automatically identify the waveforms which could contain high-quality PmP waves. To achieve the goal, we first select the waveforms with the PmP label with a probability of larger than 0.8. Then, we directly remove the overlapped waveforms from the selected waveforms with those picked already through the two-stage PmP-picking workflow. After that, we visually check and only keep the retained waveforms when main crustal phases like PmP, S, and SmS can be observed simultaneously on the associated three-component waveforms. As shown in Figure 4f, the trained PmPNet has successfully recalled the most PmP waves (larger than 96%) before 2011, and even for the seismic data after 2010 which are not involved in training PmPNet, there is also a high recall value of larger than 85%. In addition, we have newly picked 570 high-quality PmP waves based on those picks recognized by PmPNet (e.g., Figure 5h), which have a similar data coverage (Figure 5b) and a linear relationship between the PmP travel time and epicentral distance (Figure 5e). All of these show the high performance of our trained PmPNet in identifying the PmP waves from broadband seismic data.

Table 1
Precision-Recall Pairs From PmPNet With Different Probability Threshold for the Validation Data Set

$p_{\text{threshold}}$	0.0001	0.01	0.1	0.2	0.5	0.8	0.9	0.99	0.999
Precision	0.4111	0.8387	0.9190	0.9259	0.9464	0.9659	0.9783	0.9845	0.9865
Recall	0.9322	0.8908	0.8757	0.8701	0.8644	0.8531	0.8493	0.8399	0.8267
F1 score	0.5706	0.8640	0.8968	0.8971	0.9035	0.9060	0.9092	0.9065	0.8996

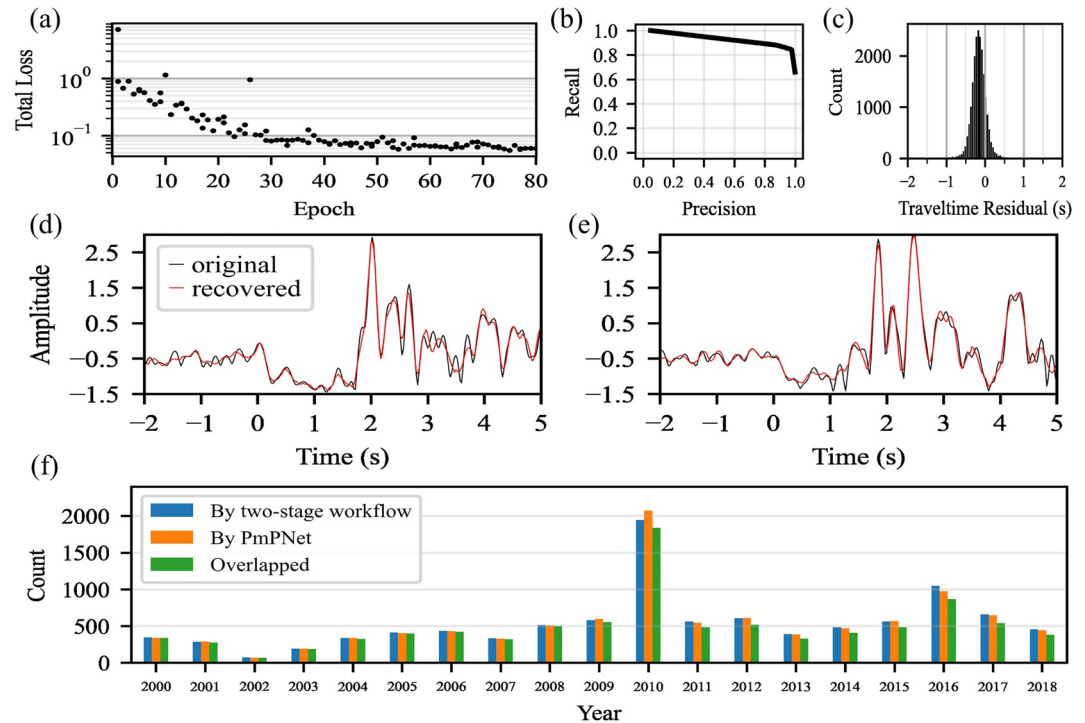


Figure 4. Training and validation performance of PmPNet when applied to real data. (a) The total training loss decreases as the epoch increases and it becomes stable after about 40-epoch training. (b) The precision-recall curve on validation set. (c) The PmP travel time residual between the predicted and manually picked ones on validation set. (d)–(e) The PmPNet recovered input and the input component on validation set. (f) PmP picks when applying the trained PmPNet to real data. Blue bars show the picked PmP waves each year by the two-stage workflow, orange bars show the identified PmP waves each year by PmPNet with the probability of greater than 0.8, and green bars show the overlapped PmP waves each year between the two identifiers.

To further release the power of our trained PmPNet in identifying the PmP waves from the massive seismic data, we also apply the trained network to the short-period vertical-component seismic data spanning from January 1990 to December 1999 retrieved also from SCEDC. The retrieved seismic waveforms are triggered by earthquakes with a magnitude between 2.0 and 5.0, and they are first cleared with the vertical-component SNR larger than 3.0, focal depth shallower than 20 km and epicentral distance between 50 and 200 km. Then, we align waveforms to the first P arrivals using the STA/LTA method (Trnkoczy, 2009) and cut them into a time window from 2 s before to 5 s after the first P arrivals. The associated envelopes are fed to the trained PmPNet, and through the network's automatic judgment, we obtain the probability of the PmP label and the PmP travel time in each seismic waveform. Waveforms that have the PmP probability of larger than 0.8 are selected out. And then we visually check the waveforms and pick the point with an abrupt amplitude jump near the predicted PmP arrival time as the PmP onset. After finishing the above process, we obtain a total of more than 27,000 PmP picks. For the obtained short-period PmP picks, their travel times increase linearly with respect to epicentral distance (Figure 5f), and they share similar waveform characters (Figure 5i) to those PmP waves picked from broadband data. In addition, the newly obtained short-period PmP picks improve the data coverage significantly in the central-eastern Transverse Ranges and the Mojave Desert.

4. PmP Database for Southern California

To update the high-quality PmP database in southern California, we take additional quality control on the obtained PmP waves. We first group all PmP waveforms obtained so far according to their Moho reflection points into common reflection gathers (Figures 6a and 6b), where the Moho reflection points of PmP waves are calculated in the HK model and to form a gather the distance between any two reflection points is no more than 30 km. Then

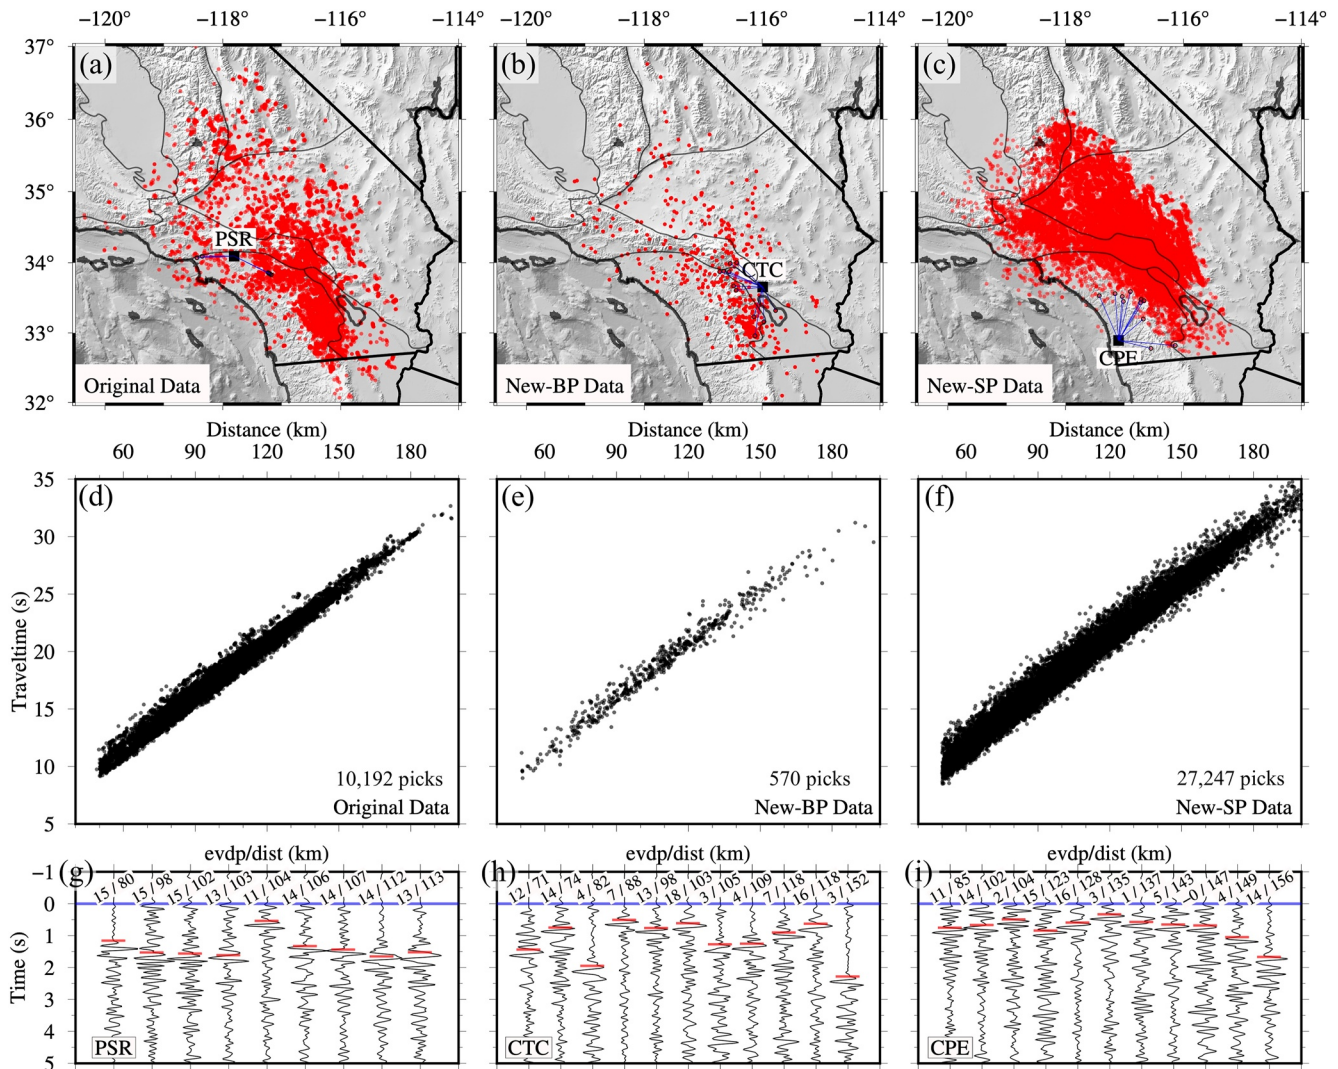


Figure 5. Comparison of the newly picked PmP waves with those built using the two-stage PmP-picking workflow. (a) The spatial distribution of PmP waves of the original database by (Li et al., 2022). Red points refer to the PmP reflection points which are calculated in the HK model. The black square and black circles denote the seismic station and reflection points of the PmP waveforms shown in (g). (b) Same as (a) but for the newly picked PmP waves by PmPNet on broadband (BP) data. (c) Same as (a) but for the newly picked PmP waves by PmPNet on short-period (SP) data. (d)–(f) The travel times of PmP waves with respect to epicentral distance. (g)–(i) Waveforms are aligned to the first P-wave arrivals and sorted according to earthquake depth (evdp) and epicentral distance (dist). Red bars denote the onsets of PmP waves.

within each gather, we calculate the Moho depth assuming a constant $V_p = 6.3$ km/s (Zhu & Kanamori, 2000) for every PmP wave and discard the outliers which deviate largely in Moho depth prediction (>3.0 km) from the other co-grouped picks. After that, we retain a total of more than 28,000 high-quality PmP picks. We also group all retained PmP waveforms according to their propagation paths (Figures 6c and 6d), and within the same group, the distance between any two sources of PmP waves is no more than 30 km in the horizontal direction and 3 km in the vertical direction, and the distance between any two receivers is no more than 30 km. Then within each group, we calculate the similarity between any two waveforms in a time window from 2 s before to 5 s after the first P arrival using a time-domain cross-correlation technique (Schaff & Waldhauser, 2005), and we use the averaged value as the representative waveform similarity in that group. The narrow Moho depth prediction range for each common reflection gather (Figure 7c) and high representative waveform similarity for each group (Figure 7d) show that the newly updated PmP database are of high quality and self-consistent. More field waveforms of the updated PmP database in the form of receiver gather (recorded at the same seismic station) are presented in Figures S1–S4 in the Supporting Information S1. As shown in Figure 7a, the retained PmP waves have dense data

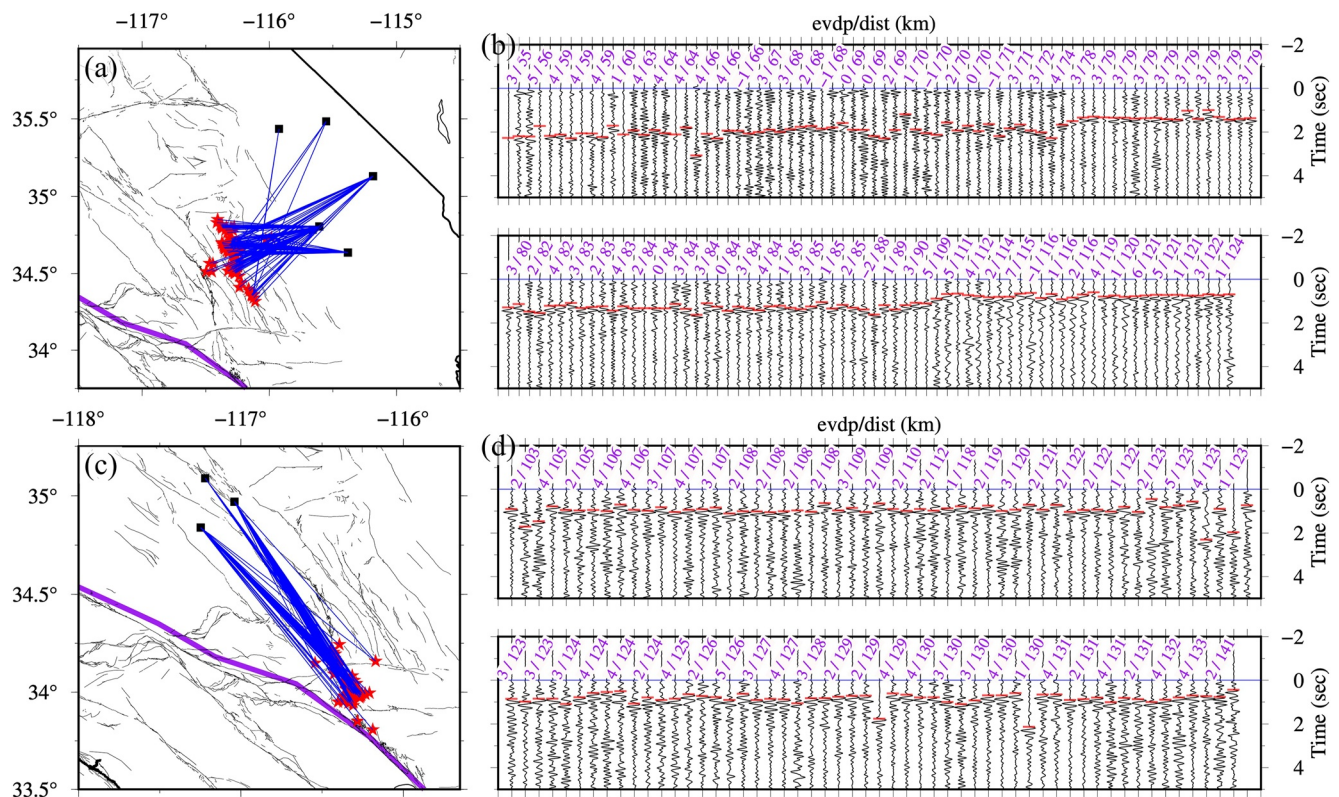


Figure 6. Using common reflection gather and grouped PmP waveforms traveling along similar paths to build the consistent PmP database. (a) An example of common PmP reflection gather in map view. Seismic stations are denoted as black-filled squares, and local earthquakes are represented as red stars. The bold purple line shows the plate boundary between the Pacific Plate and the North American Plate. (b) The PmP waveforms within a common reflection gather. Red bars show the PmP arrival times. Waveforms are aligned to the first P-wave arrivals and sorted according to earthquake depth (evdp) and epicentral distance (dist). (c) Same as (a) but for an example of the PmP waveforms traveling along similar paths in map view. (d) Same as (b) but for the PmP waveforms traveling along similar paths.

coverage in the eastern Peninsular Ranges, the central-eastern Transverse Ranges, the Mojave Desert, and the southern part of Basin and Range.

5. Discussion

Acquiring the optimal parameters for a PmPNet with good performance requires enormous computational power, since various experiments are needed to conduct. In our experiments (Figure S5 in the Supporting Information S1), we only vary one hyperparameter with others fixed when assessing the parameter's effect on the performance of PmPNet. We found that though the number of blocks $n_1 = 2$, $n_2 = 4$, $n_3 = 1$ is simple, it results in a PmPNet performing well similar to the more complicated blocks. Using more than 2,000 waveforms labeled with PmP can generate a PmPNet performing well, whereas involving more waveforms labeled with PmP will result in a PmPNet with better performance. When using 5,000 PmP-labeled waveforms to train the PmPNet, the amount ratio between the waveforms labeled with PmP and those labeled with non-PmP has little influence on the performance of PmPNet. Small initialized learning rates (<0.001) play a vitally important role in training a PmPNet with good performance. The various batch size can be chosen in training a good-performance PmPNet. Using the above chosen optimal hyperparameter configuration, it takes ~ 1.2 hr to finish the 80-epochs training on one NVIDIA GeForce RTX 2080 graphics processing units (GPUs), and in most cases, the training loss stabilizes at around 40th epoch (Figure 4a).

It is worth mentioning that the trained PmPNet performs less well on the short-period seismic data than on the broadband seismic data. For example, PmPNet identifies a higher percentage of waveforms on the short-period seismic data that most likely contain high-quality PmP waves, for example, about 13.4%, while we find that only about 24.3% are good PmP candidates through our manual check. Whereas on the broadband seismic data,

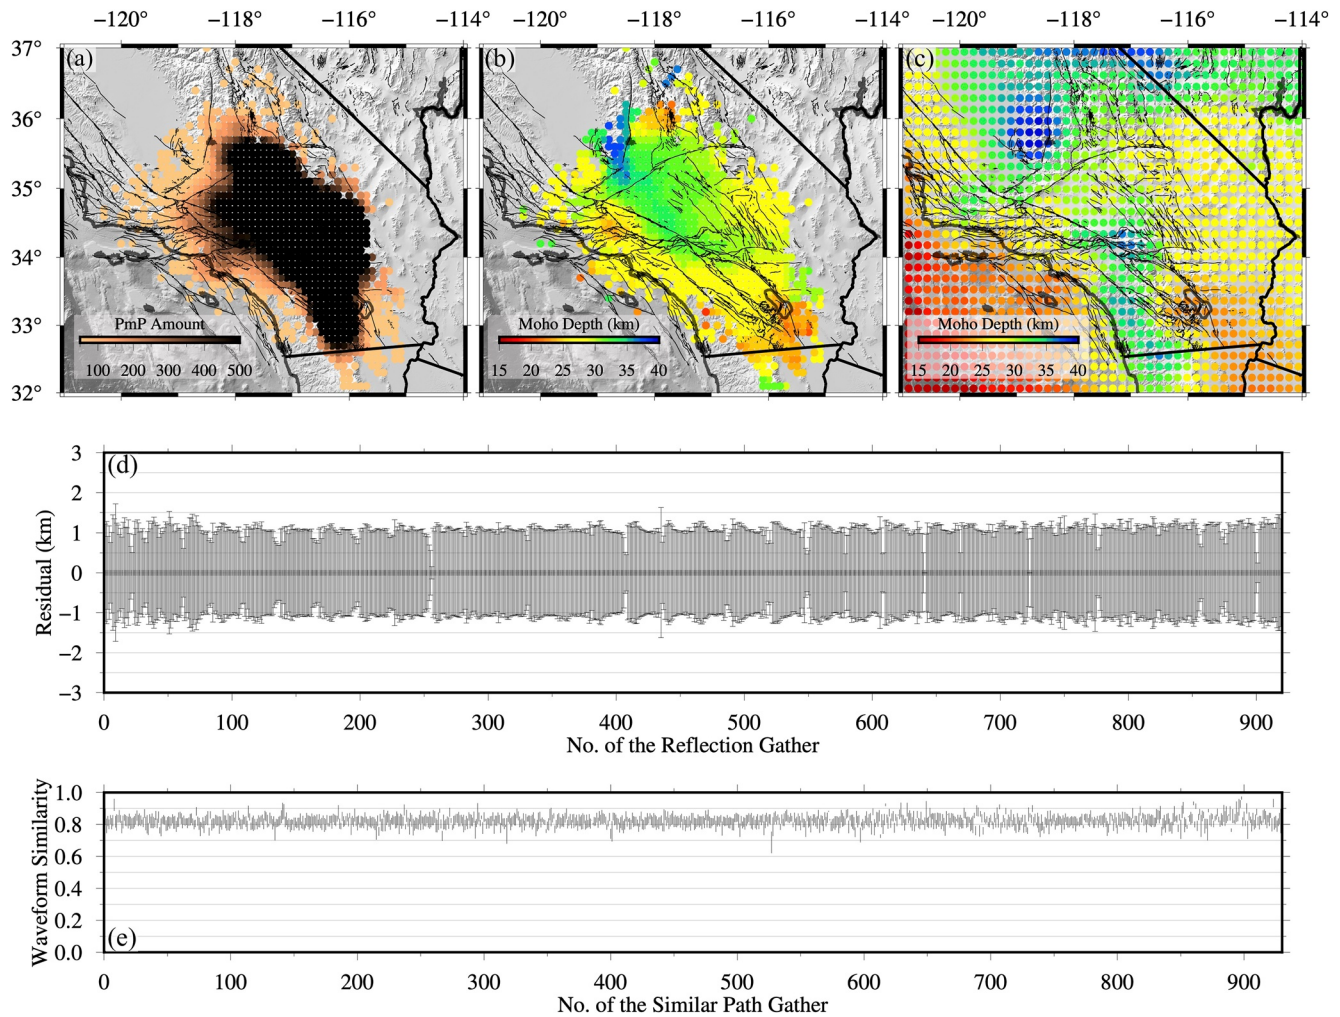


Figure 7. The newly updated PmP database for southern California. (a) The number of PmP reflection points counted at each node of a 30 km by 30 km grid. (b) The Moho geometry estimated from the new PmP database at every node of the same grid in (a). The Moho depth is the average over a reflection gather, assuming a constant $V_p = 6.3$ km/s. (c) The Moho geometry compiled in CMM-1.0, which is built mainly based on teleseismic Ps data and active-source surveys (Tape et al., 2012). (d) The standard deviation of Moho depth for each reflection gather. (e) The waveform similarity (cross-correlation coefficient) for each similar raypath gather.

PmPNet identifies about 3.0% of the total available waveforms that most likely contain high-quality PmP waves (Figure 8a), and through our manual check, we find that almost 95.4% are good PmP-wave candidates. The relatively poor performance of PmPNet on the short-period seismic data is likely caused by the higher noise level compared to that on the broadband seismic data both in original waveforms (Figures S6a and S6b in the Supporting Information S1) and in the bandpass filtered waveforms (1–7 Hz, Figures 8b and 8c). Our training data (broadband seismic data) may less characterize the PmP wave on the short-period data. Involving some waveforms, either labeled PmP or non-PmP, from the short-period data in the future study could improve the performance of PmPNet in identifying the PmP waves on the short-period seismic data.

The main task for PmPNet is to identify tens of thousands of PmP waves from millions of raw seismograms. In the meantime, PmPNet also gives the absolute traveltime prediction for PmP waves. The traveltime prediction performance of PmPNet is highly comparable with the HK model (Figures 4c, 9a and 9d), except for slight bias when predicting the PmP traveltimes at short distances. While we do observe that the traveltime residual between the predicted PmP traveltime by PmPNet and that picked manually is somewhat large (Figures 9b and 9e): there are only 47% of picks with traveltime residuals no larger than 0.5 s. To solve this problem, we further train a separate model called “PmP-traveltime-Net” to predict the differential traveltime between PmP and first P waves. The structure of PmP-traveltime-Net is very similar to PmPNet, detailed introduction can be

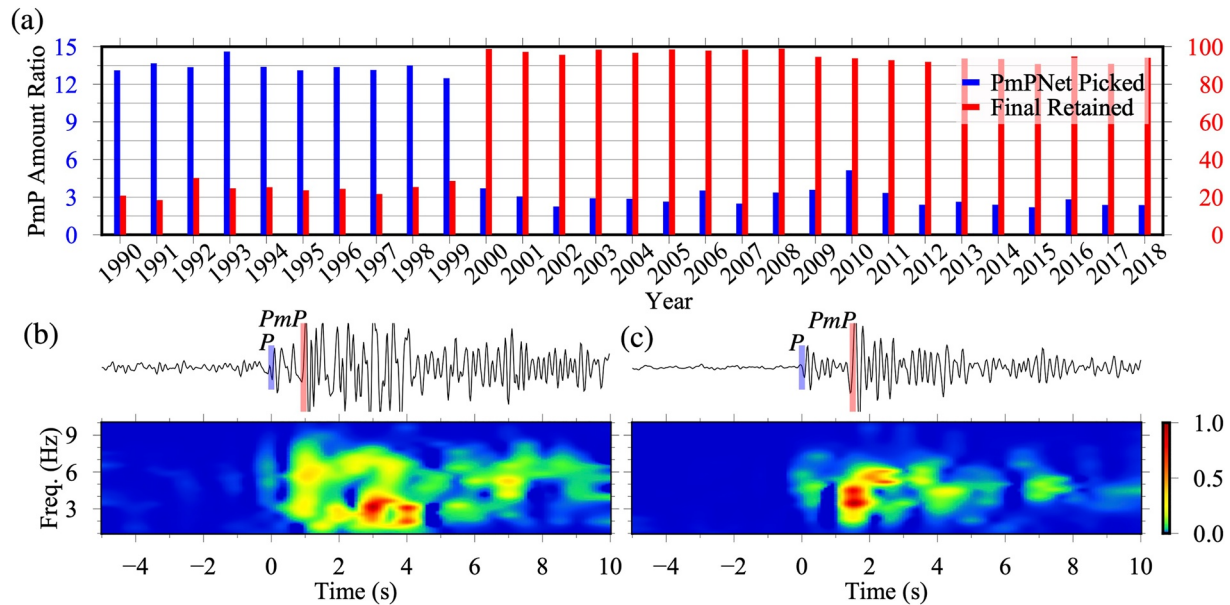


Figure 8. (a) The data amount ratio between the waveforms containing high-quality PmP with the probability of larger than 0.8 and the total available seismic data each year (in blue). And the data amount ratio between the finally retained PmP waves and the waveforms containing high-quality PmP with the probability of larger than 0.8 each year (in red). (b) An example of short-period waveform and its spectrum. Here, the shown waveform has been bandpass filtered (1–7 Hz). The waveform is recorded at station CI.LJB and it is triggered by an earthquake with the magnitude of 2.0, focal depth of 3 km, and epicentral distance of 112 km. The blue bar indicates the onset time of the first P wave, and the red bar indicates the onset time of PmP wave. (c) An example of a broadband waveform and its spectrum. The waveform is recorded at station CI.HOL and it is triggered by an earthquake with the magnitude of 2.1, focal depth of 5 km, and epicentral distance of 146 km.

found in Text S1 in the Supporting Information S1. We train PmP-traveltime-Net on 10,000 waveform data with clear PmP waves prepared by the two-stage workflow (Li et al., 2022) and validate it on all the manually picked data (during the period from 1990 to 2018). Test performance of PmP-traveltime-Net shows that there are 95% of picks with traveltime residuals no larger than 0.5 s and even 88% of picks with traveltime residuals no larger than 0.1 s (Figures 9c and 9f and Figure S7 in the Supporting Information S1). We recommend conducting the necessary manual verification for further research purposes, though the traveltime prediction can reach a high-level accuracy.

For the Moho geometry constrained by the newly built PmP database (Figure 7b), we observe a shallow Moho beneath the Salton Trough, and a deep Moho beneath the eastern Transverse Ranges and the southern Sierra Nevada, which are consistent with the California Moho Model version 1.0 (CMM-1.0, Tape et al., 2012). However, we find the Moho beneath the Mojave Desert is slightly deeper compared to the CMM-1.0, which reflects a lower-velocity anomaly in the lower crust beneath this region (Lee et al., 2014; Shaw et al., 2015). In other regions like the Ventura Basin and the Coso volcanic field, we observe some small-scale Moho undulations, which is related to the complicated local crust structure (Wilson et al., 2003; Yan & Clayton, 2007). In the western Peninsular Ranges, the Moho depth estimated by sparse PmP data is much shallower compared to the CMM-1.0, which may be related to the oversimplified one-layer crust model assumed here or/and a gradual transition from the crust to the upper mantle there (Li et al., 2022).

6. Conclusions

In this work, we have proposed a deep-neural-network-based algorithm, PmPNet, to help automatically identify PmP waves. Taking advantage of the manually prepared PmP database (10,192 picks during the period 2000.01 ~ 2018.12) by Li et al. (2022), we use 5,000 signal envelopes with clear PmP waves (during the period 2000.01 ~ 2010.07) and 100,000 signal envelopes without clear PmP waves to train (80% of envelopes) and validate (20% of envelopes) PmPNet. We find the optimal model parameters of PmP-Net through checking three metrics: precision, recall, and F1 score. Our trained optimal PmPNet can reach high precision (96.6%) and recall (85.3%) simultaneously, hence also a high F1 score (0.906). Further test of the trained PmPNet on the 2000.01 ~ 2018.12

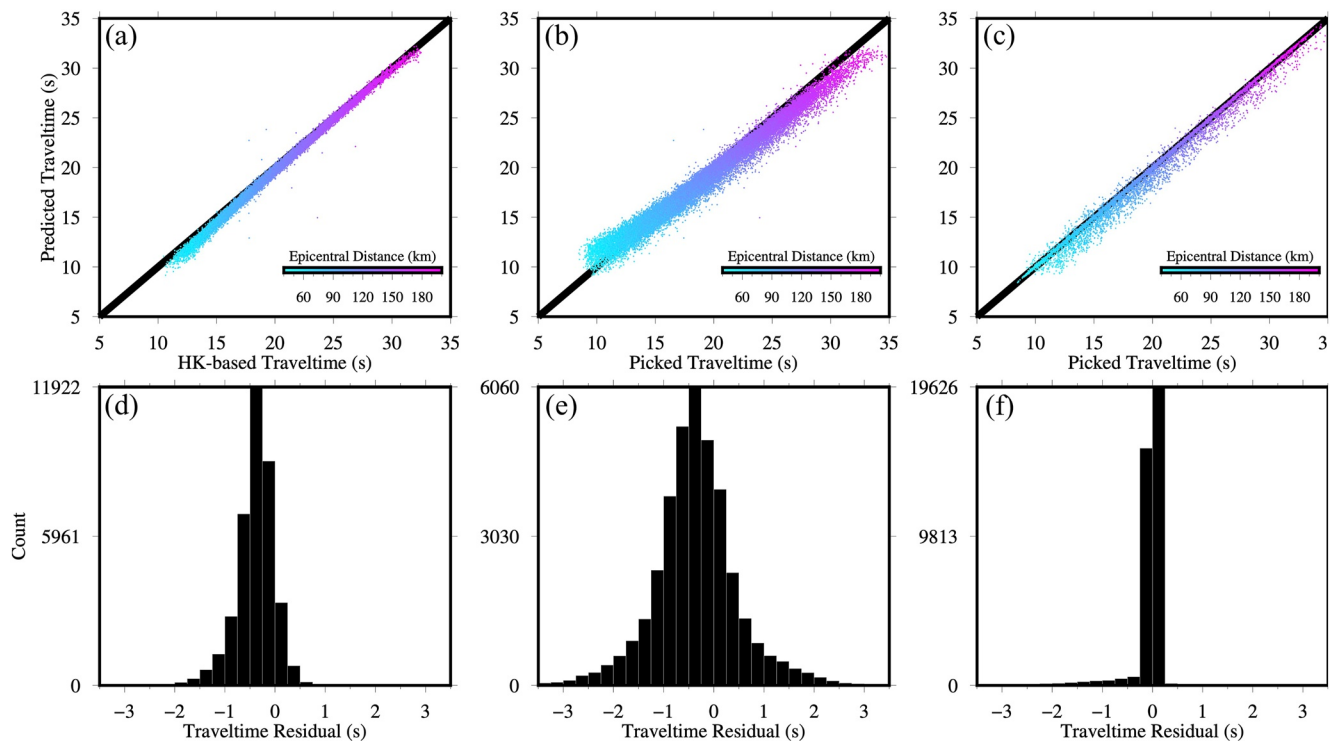


Figure 9. Accuracy of predicted PmP traveltimes by PmPNet. (a) Comparison between the predicted PmP traveltimes by PmPNet (y-axis) and those calculated using the HK model with a fixed Moho at 30 km depth (x-axis). All manually picked data are used to carry out the test and points are color-coded with epicentral distance. (b) Same as (a) but for comparison between the predicted PmP traveltimes by PmPNet and those picked manually. (c) Same as (a) but for comparison between the predicted PmP traveltime by PmP-traveltime-Net and those picked manually. (d)–(f) Distribution of PmP traveltime residuals between the two datasets shown in (a)–(c), respectively. Traveltime residual is calculated by subtracting the *x*-axis data from the *y*-axis data.

seismic data shows that the algorithm can successfully identify almost all PmP waves (larger than 96%) before 2011, and still more than 85% of the PmP waves after 2010 where the test data are not involved in training PmPNet. Applying the trained PmPNet to the seismic data during the period 1990.01 ~ 1999.12, we have nearly tripled the volume of the PmP database in southern California (28,093 picks).

The updated PmP database in this study provides valuable seismic observations that can help us image the whole crustal P-wave structure beneath southern California when complemented with the local first P-wave travel time data. Furthermore, the resulting high-resolution P-wave structure, together with the S-wave structure constrained by ambient noise data, will lay a solid foundation for our understanding of the various deformation process within the lower crust (Bürgmann et al., 2008; Shinevar et al., 2018). Besides, the high-quality PmP database in this study can serve as the template to help train PmPNet and make it suitable for identifying the PmP wave from both the broadband and short-period seismic data. Moreover, we could build the PmP database in various tectonic settings following this way. Our developed PmPNet can be easily modified to recognize other later seismic phases, especially when they are relatively rarely used in routine seismic studies, because of the excellent performance of our designed algorithm on the unbalanced data.

Data Availability Statement

The seismic data and local earthquake catalog are requested from Southern California Earthquake Data Center (<https://scedc.caltech.edu/>). For this study, the picked PmP waves and PmPNet codes can be accessed at <https://doi.org/10.21979/N9/00380V>. More detailed information on PmP database and developed workflow/code to identify PmP phase can be found at <https://github.com/Seismic-Data-imaging-the-Earth/PmPWorld>.

Acknowledgments

The authors thank the Editor, Michael Bostock, and the Associate Editor for handling this paper. The authors are grateful to the Associate Editor, the reviewer Albert Leonardo and another anomalous reviewer for their constructive comments that have greatly improved the paper. The work of W. Ding and K. Ren was partially supported by the NSF grants DMS-1937254 and EAR-2000850. T. Li and P. Tong were partly supported by Singapore MOE AcRF Tier-2 Grant (MOE2019-T2-2-112) and the National Research Foundation Singapore and the Singapore Ministry of Education under the Research Centers of Excellence Initiative (Project Code Number: 04MNS001953A620). X. Yang was partially supported by the NSF grants DMS-1818592 and DMS-2109116. Part of the work was initiated during X. Yang's visit to the Department of Applied Physics and Applied Mathematics, Columbia University, and X. Yang is grateful to the department for their hospitality.

References

- Beroza, G., Segou, M., & Mostafa Mousavi, S. (2021). Machine learning and earthquake forecasting—Next steps. *Nature Communications*, 12(1), 4761. <https://doi.org/10.1038/s41467-021-24952-6>
- Bürgmann, R., Dresen, G., & Georg (2008). Rheology of the lower crust and upper mantle: Evidence from rock mechanics, geodesy, and field observations. *Annual Review of Earth and Planetary Sciences*, 36(1), 531–567. <https://doi.org/10.1146/annurev.earth.36.031207.124326>
- Christensen, N. I., & Mooney, W. D. (1995). Seismic velocity structure and composition of the continental crust: A global view. *Journal of Geophysical Research*, 100(B6), 9761–9788. <https://doi.org/10.1029/95jb00259>
- Garcia, J. A., Waszek, L., Tauzin, B., & Scherr, N. (2021). Automatic identification of mantle seismic phases using a convolutional neural network. *Geophysical Research Letters*, 48(18), e2020GL091658. <https://doi.org/10.1029/2020gl091658>
- Guo, H., Li, Y., Shang, J., Gu, M., Huang, Y., & Gong, B. (2017). Learning from class-imbalanced data: Review of methods and applications. *Expert systems with applications*. <https://doi.org/10.1016/j.eswa.2016.12.035>
- Hadley, D., & Kanamori, H. (1977). Seismic structure of the transverse ranges, California. *Geological Society of America Bulletin*, 88(10), 1469–1478. [https://doi.org/10.1130/0016-7606\(1977\)88<1469:ssottr>2.0.co;2](https://doi.org/10.1130/0016-7606(1977)88<1469:ssottr>2.0.co;2)
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *IEEE conference on computer vision and pattern recognition*. <https://doi.org/10.1109/cvpr.2016.90>
- Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd international conference on machine learning*.
- Johnson, J. M., & Khoshgoftaar, T. M. (2019). Survey on deep learning with class imbalance. *Journal of Big Data*, 6(27). <https://doi.org/10.1186/s40537-019-0192-5>
- Kingma, D. P., & Ba, J. (2017). *Adam: A method for stochastic optimization*. arXiv:1412.6980.
- LeCun, Y., Boser, B., Denker, J., Henderson, D., Howard, R., Hubbard, W., & Jackel, L. (1990). Handwritten digit recognition with a back-propagation network. In D. Touretzky (Ed.), *Advances in neural information processing systems* (Vol. 2). Morgan-Kaufmann.
- Lecun, Y., Haffner, P., Bottou, L., & Bengio, Y. (1999). Object recognition with gradient-based learning. In D. Forsyth (Ed.), *Feature grouping*. Springer. https://doi.org/10.1007/3-540-46805-6_19
- Lee, E.-J., Chen, P., Jordan, T. H., Maechling, P. B., Denolle, M. A., & Beroza, G. C. (2014). Full-3-d tomography for crustal structure in southern California based on the scattering-integral and the adjoint-wavefield methods. *Journal of Geophysical Research: Solid Earth*, 119(8), 6421–6451. <https://doi.org/10.1002/2014jb011346>
- Li, T., Yao, J., Wu, S., Xu, M., & Tong, P. (2022). Moho complexity in southern California revealed by local pmp and teleseismic ps waves. *Journal of Geophysical Research: Solid Earth*, 127(2), e2021JB023033. <https://doi.org/10.1029/2021jb023033>
- Mooney, W. (2010). 1.11 crust and lithospheric structure—global crustal structure. *Seismology and Structure of the Earth: Treatise on Geophysics*, 1, 361.
- Nair, V., & Hinton, G. E. (2010). Rectified linear units improve restricted Boltzmann machines. In *Proceedings of the 27th international conference on machine learning* (pp. 807–814).
- Ross, Z. E., Meier, M.-A., & Hauksson, E. (2018). P wave arrival picking and first-motion polarity determination with deep learning. *Journal of Geophysical Research: Solid Earth*, 123(6), 5120–5129. <https://doi.org/10.1029/2017jb015251>
- Ross, Z. E., Yue, Y., Meier, M., Hauksson, E., & Heaton, T. H. (2019). PhaseLink: A deep learning approach to seismic phase association. *Journal of Geophysical Research: Solid Earth*, 124, 856–869. <https://doi.org/10.1029/2018JB016674>
- Rumelhart, D. E., & McClelland, J. L. (1987). Learning internal representations by error propagation. In *Parallel distributed processing: Explorations in the microstructure of cognition: Foundations* (pp. 318–362).
- SCEDC (2013). *SCEDC (2013): Southern California earthquake center*. Caltech Dataset. <https://doi.org/10.7909/C3WD3xH1>
- Schaff, D. P., & Waldhauser, F. (2005). Waveform cross-correlation-based differential travel-time measurements at the northern California seismic network. *Bulletin of the Seismological Society of America*, 95(6), 2446–2461. <https://doi.org/10.1785/0120040221>
- Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks*, 61, 85–117. <https://doi.org/10.1016/j.neunet.2014.09.003>
- Shaw, J. H., Plesch, A., Tape, C., Suess, M. P., Jordan, T. H., Ely, G., et al. (2015). Unified structural representation of the southern California crust and upper mantle. *Earth and Planetary Science Letters*, 415, 1–15. <https://doi.org/10.1016/j.epsl.2015.01.016>
- Shinevar, W. J., Behn, M. D., Hirth, G., & Jagoutz, O. (2018). Inferring crustal viscosity from seismic velocity: Application to the lower crust of southern California. *Earth and Planetary Science Letters*, 494, 83–91. <https://doi.org/10.1016/j.epsl.2018.04.055>
- Sun, A., Zhao, D., Ikeda, M., Chen, Y., & Chen, Q. (2008). Seismic imaging of southwest Japan using p and pmp data: Implications for arc magmatism and seismotectonics. *Gondwana Research*, 14(3), 535–542. <https://doi.org/10.1016/j.gr.2008.04.004>
- Tape, C., Plesch, A., Shaw, J. H., & Gilbert, H. (2012). Estimating a continuous Moho surface for the California unified velocity model. *Seismological Research Letters*, 83(4), 728–735. <https://doi.org/10.1785/0220110118>
- Trnkoczy, A. (2009). Understanding and parameter setting of sta/Ita trigger algorithm. In *New manual of seismological observatory practice (nmsop)* (pp. 1–20). Deutsches GeoForschungsZentrum GFZ.
- Wang, H., Zhao, D., Huang, Z., Xu, M., Wang, L., Nishizono, Y., & Inakura, H. (2018). Crustal tomography of the 2016 kumamoto earthquake area in west Japan using p and pmp data. *Geophysical Journal International*, 214(2), 1151–1163. <https://doi.org/10.1093/gji/ggy177>
- Wang, J., Xiao, Z., Liu, C., Zhao, D., & Yao, Z. (2019). Deep learning for picking seismic arrival times. *Journal of Geophysical Research: Solid Earth*, 124(7), 6612–6624. <https://doi.org/10.1029/2019JB017536>
- Wilson, C. K., Jones, C. H., & Gilbert, H. J. (2003). Single-chamber silicic magma system inferred from shear wave discontinuities of the crust and uppermost mantle, Coso geothermal area, California. *Journal of Geophysical Research*, 108(B5). <https://doi.org/10.1029/2002jb001798>
- Xia, S., Zhao, D., Qiu, X., Nakajima, J., Matsuzawa, T., & Hasegawa, A. (2007). Mapping the crustal structure under active volcanoes in central tohoku, Japan using p and pmp data. *Geophysical Research Letters*, 34(10). <https://doi.org/10.1029/2007gl030026>
- Yan, Z., & Clayton, R. (2007). Regional mapping of the crustal structure in southern California from receiver functions. *Journal of Geophysical Research*, 112(B5). <https://doi.org/10.1029/2006jb004622>
- Zhu, L., & Kanamori, H. (2000). Moho depth variation in southern California from teleseismic receiver functions. *Journal of Geophysical Research*, 105(B2), 2969–2980. <https://doi.org/10.1029/1999jb900322>
- Zhu, W., & Beroza, G. C. (2019). Phasenet: A deep-neural-network-based seismic arrival-time picking method. *Geophysical Journal International*, 216(1), 261–273. <https://doi.org/10.1093/gji/ggy423>