An automated, deep-learning-based method for investigating spatial-temporal evolution of seismicity

Enze Zhang[1] and Hongfeng Yang[2]

[1]Institute for Geophysics, The University of Texas at Austin
[2]The Chinese University of Hong Kong

November 22, 2022

Abstract

Earthquake migration patterns are important to reveal various triggering mechanisms, including the tectonic process and those caused by anthropogenic activities. Mapping out the spatial-temporal seismicity pattern is traditionally conducted using reference marks either in spatial or time. However, such mapping is particularly challenging for induced earthquakes because most industrial records that provide reference marks are unavailable to the public. Moreover, advances in earthquake detection techniques proliferate earthquake catalogs and thus require labor-intensive investigation. Therefore, a new methodology is demanded to automatically investigate spatial-temporal patterns of seismicity without reference marks. Here, we present a deep learning-based method to automatically identify the timings and locations of anomalous seismicity, defined by the sudden change of earthquakes in a region. We first rasterize multi-dimensional earthquake catalogs into 2-D distribution maps. Then, we identify the maps with anomalous seismicities and extract their timings and locations to generate condensed catalogs to reduce the manual effort in further investigation. We choose Changning and Weiyuan in Sichuan Basin as our study areas due to their high seismicity rates in recent years. We use the Changning catalog to train the method and the Weiyuan catalog to test the method's spatial transferability. Our approach successfully condenses both the Changning and Weiyuan catalogs with the accuracy of 0.87 based on the F1 score. The anomalous seismicities identified by our network include both earthquakes associated with hydraulic fracturing and aftershocks following strong quakes. As such, our method could be applied to broader areas with more complex migration patterns, including natural earthquake sequences.

An automated, deep-learning-based method for investigating spatial-temporal evolution of seismicity

Enze Zhang[1,2†] and Hongfeng Yang[1,3]

[1]Earth System Science Programme, The Chinese University of Hong Kong, Sha Tin, NT, Hong Kong, China.

[2]Institute for Geophysics, The University of Texas at Austin, Austin, TX, USA.

[3]Shenzhen Research Institute, The Chinese University of Hong Kong, Shenzhen, China.

Corresponding author: Hongfeng Yang (hyang@cuhk.edu.hk)

†Current address: Institute for Geophysics, The University of Texas at Austin, Austin, TX, USA.

Key Points:

- The method can detect anomalous events from induced earthquake catalogs.
- The detected anomalous events can be used to further investigate the triggering mechanism of earthquakes.
- The method can be applied to regions beyond the training data.

Abstract

Earthquake migration patterns are important to reveal various triggering mechanisms, including the tectonic process and those caused by anthropogenic activities. Mapping out the spatial-temporal seismicity pattern is traditionally conducted using reference marks either in spatial or time. However, such mapping is particularly challenging for induced earthquakes because most industrial records that provide reference marks are unavailable to the public. Moreover, advances in earthquake detection techniques proliferate earthquake catalogs and thus require labor-intensive investigation. Therefore, a new methodology is demanded to automatically investigate spatial-temporal patterns of seismicity without reference marks. Here, we present a deep learning-based method to automatically identify the timings and locations of anomalous seismicity, defined by the sudden change of earthquakes in a region. We first rasterize multi-dimensional earthquake catalogs into 2-D distribution maps. Then, we identify the maps with anomalous seismicities and extract their timings and locations to generate condensed catalogs to reduce the manual effort in further investigation. We choose Changning and Weiyuan in Sichuan Basin as our study areas due to their high seismicity rates in recent years. We use the Changning catalog to train the method and the Weiyuan catalog to test the method's spatial transferability. Our approach successfully condenses both the Changning and Weiyuan catalogs with the accuracy of 0.87 based on the F1 score. The anomalous seismicities identified by our network include both earthquakes associated with hydraulic fracturing and aftershocks following strong quakes. As such, our method could be applied to broader areas with more complex migration patterns, including natural earthquake sequences.

Plain Language Summary

Earthquakes migrate in space and time, sometimes forming clusters due to various mechanisms. Detecting anomalies in seismicity patterns helps understand why they occur and can play critical roles in seismic hazard mitigation. Although finding anomalies in a single dimension is straightforward, it is challenging to detect anomalies in earthquake patterns when spatial and temporal information is coupled together. A prior information either in space or time is normally required to analyze seismicity, but such information is not always available for induced earthquakes. Therefore, a new method is required to decouple the spatial and temporal information, providing references in at least one domain without prior knowledge. Here, we develop a deep-learning-based method to identify timings of abnormal seismicity. With these timings as references, the anomalies in the spatial domain will be apparent and can be easily extracted. With our method, the investigation of induced earthquakes will no longer depend on prior knowledge from industrial records.

1 Introduction

Earthquakes evolve spatially over time, and some patterns of the evolution provide important insights into the mechanisms driving earthquakes and their interactions (Freed, 2005). Different mechanisms such as static triggering (King et al., 1994), triggering due to afterslip (Barbot et al., 2009; Peng and Zhao, 2009), and dynamic triggering (Anderson et al., 1994; Hill et al., 1993; Kilb et al., 2000; Yun et al., 2019) can cause various earthquake migration patterns. For instance, King et al. (1994) discovered that the aftershocks of the 1992 Landers earthquake were distributed at sites where Coulomb stresses have risen. In a longer timescale, afterslip may have a more critical role in triggering aftershocks than static triggering. Barbot et al. (2009) observed greater moment release from afterslip than coseismic slip of the 2004 Parkfield earthquake. The aftershocks of this earthquake migrated along the fault with logarithmic time since the mainshock, showing the distinct migration pattern of afterslip-triggered aftershocks (Peng and Zhao, 2009). In addition to aftershocks triggered by local mainshocks, dynamic triggering could also cause aftershocks by long-period waves. For instance, the 1992 M_W 7.3 Landers earthquake triggered three magnitude 3.4+ events and numerous small events at distances of several hundred kilometers (Anderson et al., 1994). Prejean et al. (2004) observed that the remotely triggered seismicity initiated with the arrival of the surface wave of the 2002 M_W 7.8 Denali Fault earthquake.

In addition to the tectonic process, anthropogenic activities such as fluid injection may also cause clear migration of earthquakes, which serve as one of the vital information to infer their inducing mechanisms. In the past decade, the rate of induced earthquakes has increased abruptly in the US and Canada (Atkinson et al., 2016; Bao and Eaton, 2016; Friberg et al., 2014; Holland, 2013) due to activities like wastewater disposal that are associated with hydrocarbon production. Hydraulic fracturing, a technique commonly used in stimulating fracture growth, can also trigger moderate (M_L 3–5) and strong ($M_L > 5$) earthquakes, causing substantial damage. Since 2010, China has been conducting shale gas exploration, drilling over 500 production wells in the Sichuan Basin (Tan et al., 2020). Meanwhile, the Sichuan Basin has experienced frequent earthquakes (Lei et al., 2019a; Yang et al., 2020; Wong et al., 2021; Zhou et al., 2021). The fluid injection may activate the local faults and lead to large earthquakes. For instance, the M_W 5.7 earthquake that occurred on June 17th, 2019, Changning, killed 13 people, injured more than 200 people, and damaged numerous buildings, was considered as an interaction among hydraulic fracking, salt mining, and smaller magnitude earthquakes in the region (Lei et al., 2019b; Jia et al., 2020; Liu and Zahradník, 2020). In September and December 2019, two moderate earthquakes with magnitudes >5 occurred in Weiyuan, killing four people and injuring 75 (Lei et al., 2019a; Wang et al., 2020; Sheng et al., 2020). These sequences illustrated the urgent need to closely monitor the seismicity induced by hydraulic fracturing and better understand the underlying triggering mechanisms.

Mapping out seismicity evolution in high resolution has been commonly used to infer the mechanisms driving earthquakes and their migrations for both natural earthquakes (Anderson et al., 1994; Peng and Zhao, 2009; Zhang et al., 2022; Zhu et al., 2022) and induced seismicity (Lei et al., 2017; Haffener et al., 2018; Grigoli et al., 2018). The typical approach to analyzing spatio-temporal migration relies on a spatial or temporal mark, such as a mainshock for natural earthquakes or industrial records for induced ones. Compared with natural earthquakes, induced seismicity is sometimes more challenging to investigate partly due to its small spatio-temporal scales and the necessity of including low magnitude earthquakes with limited accuracy of hypocenters. Despite these difficulties, many studies have conducted the temporal correlation between earthquake occurrence timings and injection records to link the induced seismicity with hydraulic fracturing (Haffener et al., 2018; Lei et al., 2017; Lei et al., 2019b; Meng et al., 2019; Tan et al., 2020). Some other studies have used earthquake hypocenters as a critical indicator to identify the induced earthquakes with large magnitude ($M_L > 4$) (Grigoli et al., 2018; Sheng et al., 2020). Incorporating both spatial and temporal information, Johann and Shapiro (2020) applied a multidimensional cross-correlation technique to investigate the spatio-temporal relationship between induced seismicity and injection volumes. However, all the studies mentioned above are based on correlation with industrial activities, which are not entirely accessible to the public (Schultz et al., 2020). Therefore, a new methodology is demanded to map out seismicity migration without prior information.

Here, focusing on induced earthquakes that often exhibit spatial clustering yet abrupt changes of low-magnitude events in a short period, we develop a deep-learning-based, automated method to extract anomalous spatial-temporal information from earthquake catalogs. Powered by deep learning, our method does not require prior knowledge (e.g., background seismicity rates from historical catalogs, industrial operation records) and thus is applicable for induced earthquakes. We first train our neural network from a well-identified induced earthquake sequence in the Changning shale gas block within the Sichuan Basin. Then we apply the trained network to the dataset in the Weiyuan shale gas block to identify anomalous changes in the pattern of seismicity. Moreover, due to the transferability of deep learning, our method has the potential to be applied to a broader area and detect abnormal changes in seismicity during natural earthquake sequences, including those associated with foreshock sequences or aftershock triggering.

2. Catalog data and characteristics

Our study region is located in the southern Sichuan Basin (Fig. 1), where several shale gas blocks have been rapidly developed since 2011. Among the shale gas blocks in Sichuan, Changning and Weiyuan blocks are the two major gas production sources (Zou et al., 2018). Since 2014, the pace of shale gas production has been accelerated, and frequent earthquakes, including events with magnitudes larger than five, have been reported (Meng et al., 2019; Yang et al., 2020; Zhou et al., 2021).

We first adopt an earthquake catalog in Weiyuan, which contains 24,719 earthquakes from September 2018 to August 2020 (Wong et al., 2021; Fig. 1a), bounded by longitudes 104.21° and 105° and latitudes 29.2° and 29.8°. From 2018 to February 2019, seismic waveforms were recorded by nine short-period seismometers (Yang et al., 2020). From 2019 to 2020, 14 additional seismometers were deployed in the region, bringing the total number of stations to 23. The phase data are picked by a machine-learning phase picker (Zhu and Beroza, 2019), and earthquakes are relocated through the double-difference (HypoDD) algorithm (Waldhauser and Ellsworth, 2000).

The catalog in the Changning shale gas field is from Meng et al. (2019), who derived high-resolution earthquake locations from local temporary seismic stations. The catalog contains 18,507 earthquakes from July 2015 to January 2020 (Fig. 1c) with magnitudes up to M_W 4.7, bounded by longitudes 104.2° and 105.4° and latitudes 27.8° and 28.6°. Between February 2015 to April 2017, 6 temporary seismometers were deployed, and additional 15 seismometers were added afterward. The stations were distributed evenly within and surrounding the study region, yielding a high-resolution catalog with a completeness magnitude of M_L 1.1 that was derived from a double-difference tomographic method (tomoDD, Zhang and Thurber 2003). Furthermore, the seismicity in the catalog shows a close relationship with hydraulic fracturing (Meng et al., 2019), making it a reliable training dataset to extract the features of injection-induced earthquakes.

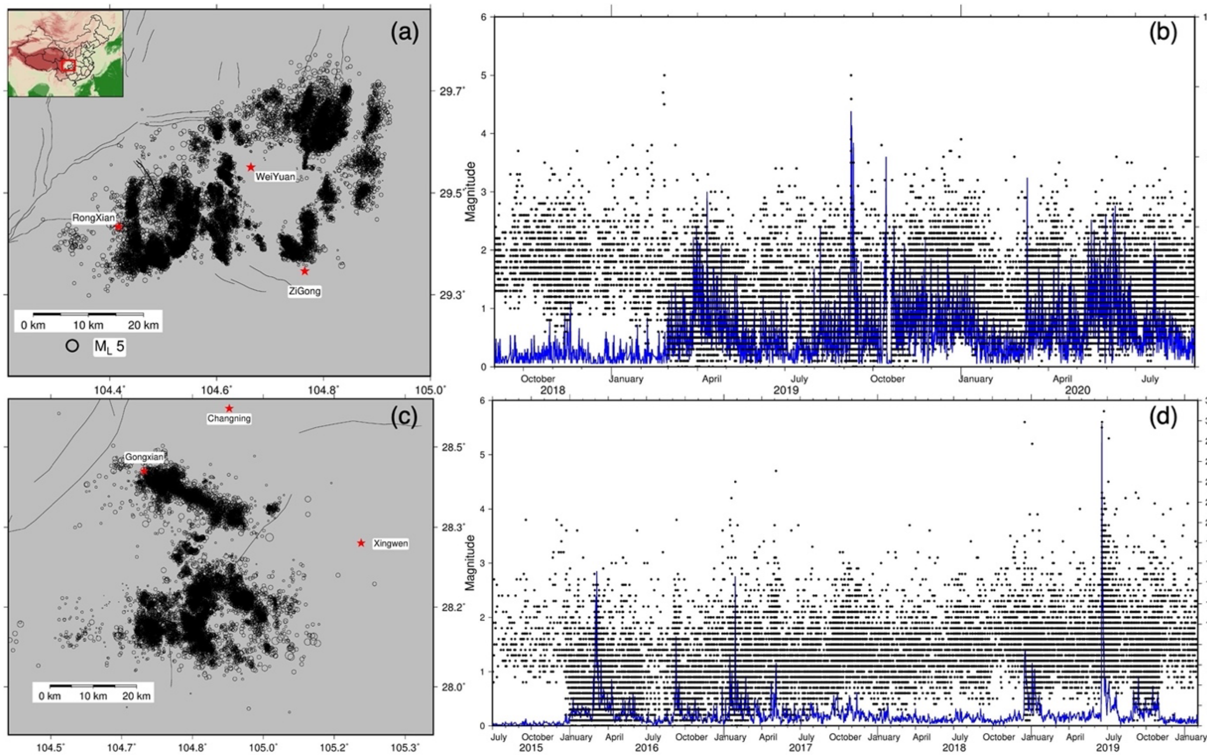

Figure 1. Earthquake distribution in our study areas Weiyuan (a) and Changning (c). (b) and (d) shows the time series of earthquake numbers for Weiyuan and Changning over the entire study period.

Both catalogs contain seismic features distinct from background seismicity. Here, we define three types of behaviors as abnormal seismicity: (1) earthquake migration, (2) sudden, and (3) gradual increase of earthquake number in small subregions (Figure 2). In the first scenario, earthquake migration, the previous earthquake swarms vanish, and new swarms appear (Figure 2a), but the total number of earthquakes in the whole region does not change much (Figure 2e). In the second scenario, a group of earthquakes may emerge quickly in a small region and then vanish in a short term (Figure 2b), exhibiting a clear signature of temporal clustering (Figure 2f). In contrast, the earthquake number may change gradually over time but overall maintain at a high level for days (Figure 2g). We classify such phenomenon as type 3.

The coupling of spatial and temporal information in the catalogs complicates the detection of individual clusters (Figure 1). For instance, the three types of abnormal features could occur simultaneously in various subregions (Figure 2d & h). Therefore, we choose the deep learning algorithm to solve this complexity. Additionally, the spatial transferability of deep learning could enable the method to be applied to places beyond the training region, meaning that we could apply the network to extract similar abnormal features in various regions. To demonstrate, we use the Changning catalog to train the deep learning network and then use the Weiyuan catalog to test the spatial transferability of the network.

In the two catalogs, we focus on the spatio-temporal changes in earthquake number and epicenter to extract abnormal features. The magnitudes of the earthquakes are not used because (1) most induced earthquakes have small magnitudes, and (2) large-magnitude earthquakes will naturally become good references providing timings and locations for detailed investigations. We do not use depth information since induced earthquakes caused by hydraulic fracturing are usually concentrated in a specific depth range. Also, the usage of earthquake hypocenters requires higher accuracy in event depths, which are however often not resolved as good as did for epicenters.

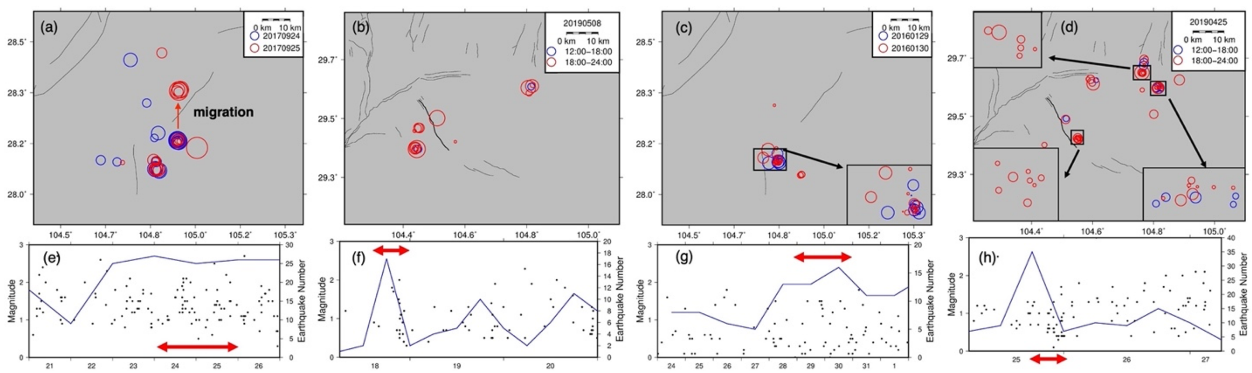

Figure 2. Example of anomalous seismicity within a 2-day time window in Changning (a & c) and 12-hour window in Weiyuan (b & d). Circles represent the locations and magnitudes of earthquakes. (a) shows the example of earthquake migration, (b) shows the sudden change in seismicity, (c) shows the gradual changes in seismicity in the region, and (d) shows the combination of the second and the third scenarios. (e-h) show the earthquake magnitude and the time series of earthquake numbers each day (a & c) or every 6 hours (b & d) in a longer time window. The red arrows indicate the period of the figures above.

3 Method

Temporal and spatial changes in earthquake numbers have been widely used to illustrate the causal relationship between anthropogenic activities and induced seismicity (Lei et al., 2017; Lei et al., 2019a; Yang et al., 2020). Without a large magnitude earthquake or industrial records as a reference, the abnormal increase in earthquake rates becomes essential to distinguish the induced seismicity from the background. In Changning, the average seismic rate before hydraulic fracturing was less than four events every three years from 1970 to 2014. From 2015 to 2017, 15,057 earthquakes with M_L>0 were identified after shale gas production began (Meng et al., 2019). A similar phenomenon occurred in Weiyuan, with infrequent seismicity before mid-2015 and a dramatic increase in the number of earthquakes afterward (Yang et al., 2020).

Here, based on the variation of earthquake numbers, we automatically extract spatio-temporal anomalous information based on deep learning technology (Fig. 3). The deep learning network is designed to identify and extract the timing of abnormal events, and a post-process procedure extracts the locations of abnormal events. More specifically, our method has three main steps:

- Rasterizing an earthquake catalog into distribution maps that are discretized in space and time.
- Applying the network to a sequence of consecutive maps to identify the abnormal ones and their timings.

- Extracting the locations where the earthquakes are concentrated from each abnormal map.

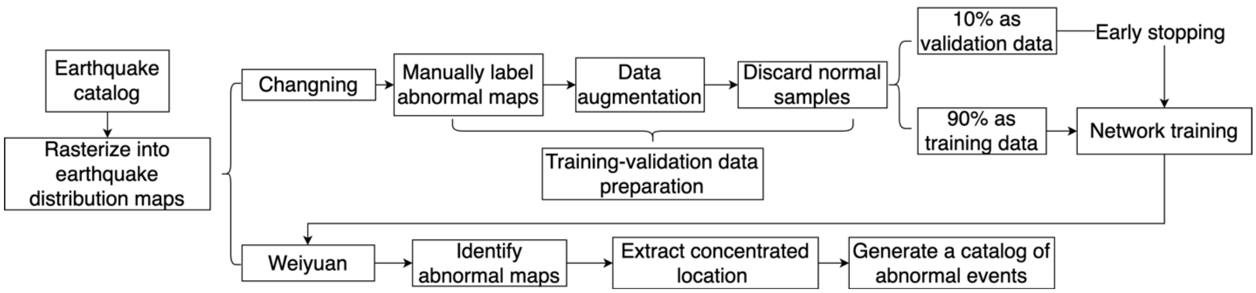

Figure 3. Diagram of the deep-learning-based framework described in the Method section.

3.1 Generating distribution maps

To generate distribution maps, we count the number of earthquakes in each grid with a given temporal and spatial resolution. Here we set the spatial resolution as 2 kilometers, the commonly recognized maximum diffusion range of an injection well (Atkinson et al., 2020). Thus, the image sizes are 59×44 in Changning and 38×33 in Weiyuan. After rasterization, we pad images with zeros to keep an identical size, 75×75, for each image. Apparently, the seismic density determines the temporal resolution. A too-wide temporal resolution could cause many timings to be abnormal, which violates our objective of saving manual practice. However, a too-narrow time window might lead to missing some abnormal timings. In Changning, we rasterize the catalog into daily distribution maps, while the temporal resolution for the Weiyuan catalog is six hours from a trial-error process, and each map has its timing.

3.2 Principles of labeling anomalous seismicity and preparing the training set

The labeling of anomalous seismicity is based on the definition of three types of abnormal behaviors (Fig. 2). The maps with migration and sudden increases are self-evident. When the earthquake number increases gradually, we only consider the map with the local peak as abnormal to reduce the number of detected events and the manual effort required for further investigation. Although our objective is to detect the abnormal induced seismicity, we do not exclude the anomalous events caused by large natural earthquakes because it will not hurt our objective and could also provide unique abnormal seismic features to enrich the training set. For labeling all the types of abnormal maps, we use a threshold of six, meaning all the abnormal maps should have at least six earthquakes in a small cluster. The cluster size varies, and different types of behavior might occur at the same time. For instance, one cluster is vanishing but still has a large earthquake number, while another cluster suddenly appears in a different place. Due to these complexities, deep learning technology is more appropriate for detecting abnormal events. The label preparation is conducted iteratively. We first prepare the initial labels and train the network. Then, we manually check the differences between the network's output and refine the labels. For instance, the network can detect some missing abnormal events by manual labeling. Further, for ambiguous events that we consider normal, we will reconsider them if the network thinks differently. Such iterative progress would mitigate the manual error and subjectiveness in the label preparation.

After label preparation, we integrate a series of consecutive maps and their labels to compose a training example. We move the temporal window by one image each step to make multiple training examples. The number of labels equals the number of consecutive maps minus one since the abnormality of the first map will not be determined without a prior map as a reference. The training set contains all the abnormal distribution maps in Changning from July 2015 to January 2020, with 1624 distribution maps with 212 abnormal ones. Due to the imbalance between normal examples and abnormal ones, we randomly discard 1/3 of normal samples to balance the training data. From the training samples, we randomly choose 10% as the validation dataset. The validation dataset mitigates overfitting and adjusts hyperparameters such as the learning rate.

We use the data in Weiyuan from March to October 2019 as our test dataset. The test region is beyond the training region to test the spatial transferability of the network. Further, the test set can provide a more representative accuracy of the network than the training set since a high accuracy of the training set is expected. The preparation for the test set is the same as the training set. The test set includes 937 distribution maps with 154 abnormal ones.

We apply the data augmentation to increase the diversity of our training dataset and the generalization of the network while keeping the manual labeling of rasterized maps to a manageable level. We flip and rotate the distribution image by degrees ranging from 45° to 315° with an interval of 45°. We also shift images vertically by 5 and 10 pixels and horizontally by steps ranging from 5 to 25 with an interval of 5 pixels. In the original training data, the earthquakes are located in the maps' central part. The earthquakes will sample more areas in the image domain by shifting images. All the augmentations are conducted independently and share the same list of labels.

To further enrich the training dataset, we generate three types of artificial distribution maps. The first type has scatter distributions with an earthquake number less than two in a single pixel (Fig. S1a). The other two types are both superimposed on the first one. The second type has pixels with earthquake numbers larger than six (Fig. S1b). The third type has pixels where the earthquake number increases gradually in the same pixel, and the peak has more than six earthquakes (Movie S1). The maps from the second and the peak from the third types will be classified as anomalous maps, while the maps from the first type are normal ones.

3.3 Deep learning network for identifying anomalous maps

Deep learning has been widely applied in seismology, such as automating phase picking (Zhu and Beroza, 2019; Johnson et al., 2021), locating earthquakes (Zhang et al., 2020), and determining focal mechanisms in real-time (Kuang et al., 2021). Here, we adopt the idea of image classification to identify abnormal earthquake distribution maps. The conventional way of classifying images is to apply the deep learning network to a single image. The output would be a list of binary numbers indicating which class the image belongs to (He et al., 2016). Here, we apply the network to a sequence of distribution maps and output binary numbers indicating their abnormality, i.e., seismicity changes as defined above. To examine the effect of input map numbers, we use five, seven, and ten consecutive maps as the input.

Here, we use the ResNet deep learning architecture, which has achieved outstanding performance in image classification (He et al., 2016). ResNet includes a building block of residual learning (Fig. S2), which can avoid the typical problem that the accuracy becomes saturated and degrades as the network depth increases. With such a design, ResNet allows to greatly increase the network depth and the learning capacity (He et al., 2016). Therefore, it could be applied to broader areas with more complex and diverse migration patterns, including natural earthquakes. ResNet has different branches with various numbers of blocks and layers. The more layers a network has, the more learning capability it contains. However, a too large network might yield overfitting issues, depending on the complexity level of the task. We examine the performance of ResNet-18, ResNet-34, and ResNet-101 (Fig. 4) and choose the best one. ResNet-18 has eight blocks, containing 17 convolutional layers; ResNet-34 has 16 blocks, containing 33 convolutional layers; ResNet-101 has 33 blocks, containing 100 convolutional layers. All the networks have a fully connected layer at the end of the architecture to generate the output labels. We use binary cross-entropy as the loss function and the stochastic gradient descent method as the optimizer. We also use the L2 regularization factor to mitigate the overfitting issue. We train the networks using different learning rates and L2 regularization factors and choose the most proper values for each network based on the final validation loss. The details of selecting the learning rate and L2 regularization factor are described in the supplementary material (Text S1).

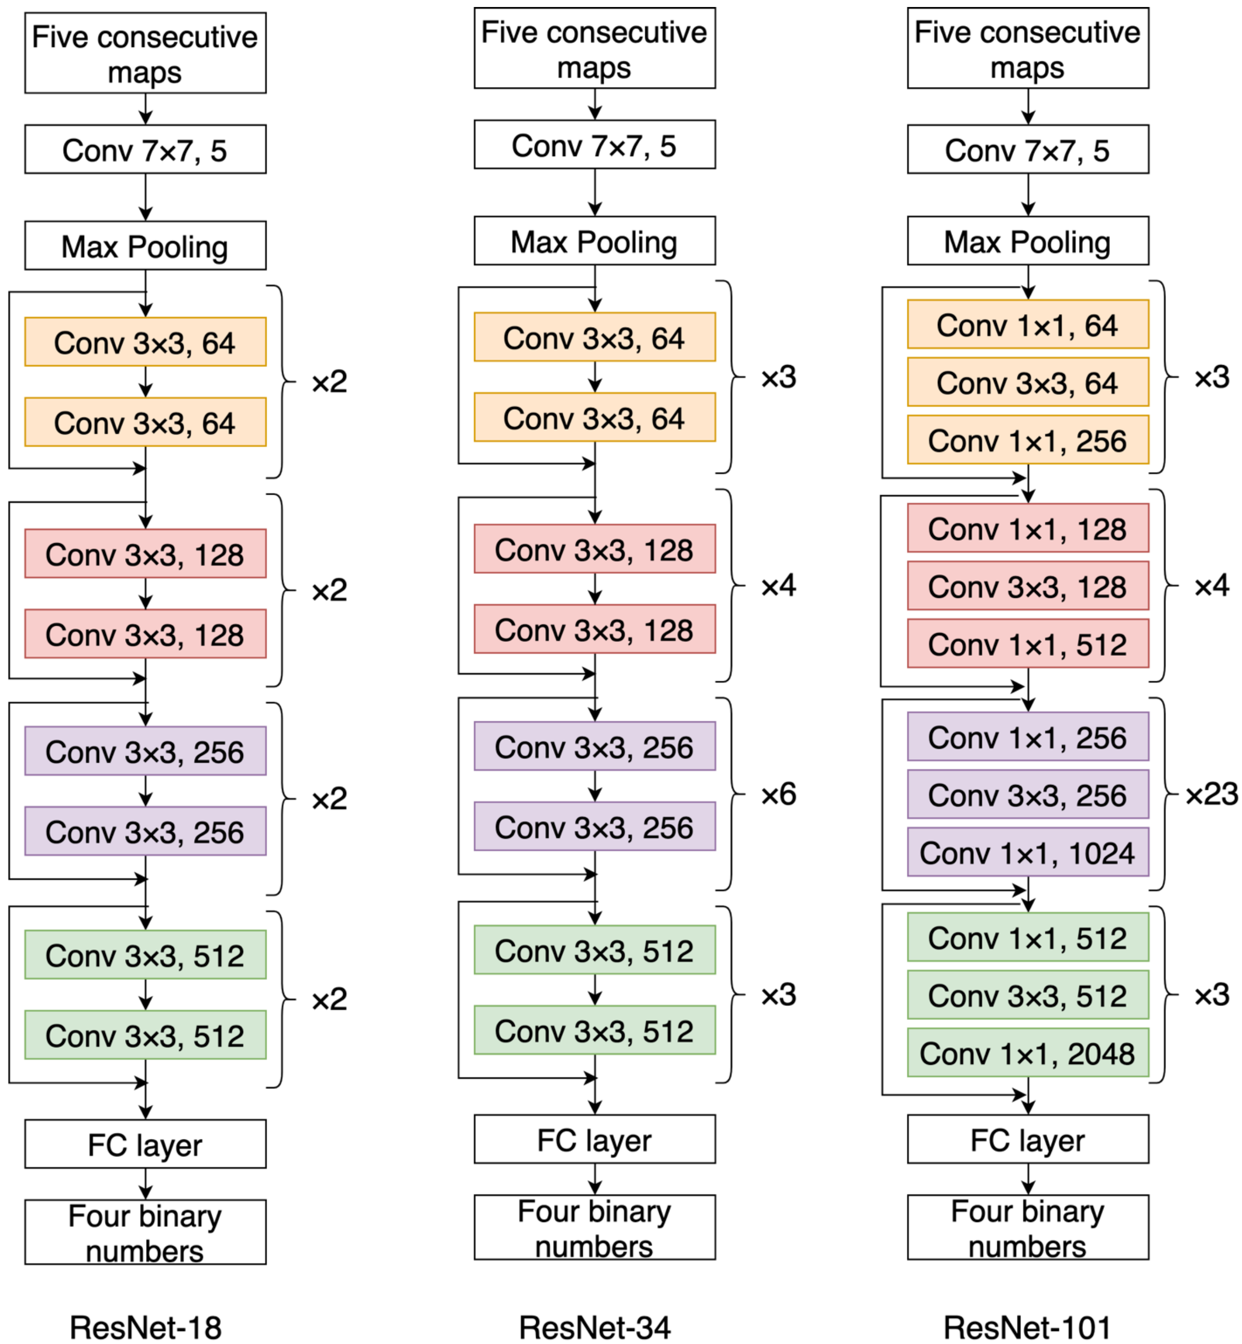

Figure 4. Architecture of the ResNet-18, ResNet-34, and ResNet-101. Different colors represent different building blocks. Conv 3×3, 64 means a convolutional layer with a 3×3 kernel and 64 channels. FC means fully connected layers. The inputs are a sequence of distribution maps, and outputs are binary numbers indicating the abnormality of the last four maps.

We start the training by initializing the network parameter using the He initialization method (He et al., 2015). We adopt an early-stopping strategy to mitigate overfitting. After each training epoch, we calculate the validation error and stop the training if the validation error starts to increase for ten consecutive epochs. The inferencing after the training will sample a single image multiple times since we move the temporal window by one image each time. We calculate the averaged output for each image and consider it abnormal if the value is larger than 0.4, an empirical threshold. The inferencing will be conducted on test data—the Weiyuan catalog from March to October 2019—to quantify the accuracy of the networks. Specifically, we calculate the F1 score on the test dataset to measure the consistency between manual labels and the network's predication and use it as the accuracy:

$$F1 = \frac{TP}{TP + 0.5 \times (FN + FP)} \quad (1)$$

TP is true positive representing the maps that both the manual label and the network consider as abnormal. FN is false negative, and FP is false positive. Both are network misidentifications: FN is the network's negative predication and TP is the positive predication. We choose to use the F1 score since it focuses on the network's ability to detect abnormal events and considers both types of misidentifications. We conduct nine experiments that combine three networks (ResNet-18, ResNet34, and ResNet-101) and three numbers (five, seven, and ten) of input maps. We calculate the F1 score of each experiment and choose the one with the largest F1 score.

3.4 Extracting anomalous locations from the identified maps

The post-processing procedure further extracts the anomalous locations and generates a catalog of abnormal events. For the identified abnormal maps, we apply a thresholding method to extract the location with concentrated earthquake distribution. We first filter the image with a 3×3 matrix of ones to sum up all the values in the surrounding pixels. Second, we extract two sets of locations: (1) the locations with a value larger than the threshold (six) in the filtered maps and (2) the locations in the original maps where the pixel value is larger than half of the threshold. Third, we take the intersection of the two location sets as the abnormal locations. Instead of using a single threshold, these processes could extract locations of clusters of various sizes. Moreover, the post-processing could filter out some misidentified maps with no concentrated earthquakes, further increasing our method's accuracy. Finally, we generate a catalog of abnormal events by taking each abnormal location at a specific timing as an event.

4. Results

The key results are presented in the following order: (1) F1-score-based comparison among the nine experiments as in Table 1; (2) the application on the Weiyuan catalog; (3) the application on the Changning catalog.

4.1 The best network based on the F1 score of the test set

Table 1. F1 scores of the nine experiments. TP means true positive, and FP means false positive.

	Five maps	Seven maps	Ten maps
ResNet-18	TP: 141	TP: 128	TP: 115
	FP: 58	FP: 28	FP: 16
	F1: 0.79	F1: 0.81	F1: 0.79

ResNet-34	***<u>TP: 137</u>***	TP: 136	TP: 122
	<u>FP: 19</u>	FP: 30	FP: 24
	<u>F1: 0.87</u>	F1: 0.84	F1: 0.80
ResNet-101	TP: 127	TP: 112	TP: 132
	FP: 13	FP: 12	FP: 32
	F1: 0.85	F1: 0.79	F1: 0.82

We estimate the accuracy of the nine experiments by comparing manual labels and network predications. The uncertainties estimation is conducted in Weiyuan, a different shale gas field than Changning where the training was conducted, which can therefore provide a more representative estimation of the model's accuracy. Table 1 summarizes the F1 score of the nine experiments with higher F1 scores showing the better results. All the experiments have similar and promising performances, while ResNet-34 is better than the other two branches for our task. More convolutional layers lead to greater learning capability. However, the complexity level is limited due to the small input image size (75×75). Therefore, a too deep network could easily overfit the training data. In other words, ResNet-101 has a larger learning capability than what this task needs and the overfitting issues make the performance of ResNet-101 worse than ResNet-34. However, ResNet-101 might be more appropriate when applying to datasets that have longer durations and larger special coverage (e.g., natural earthquake catalogs). The best experiment here is the combination between ResNet-34 and five consecutive maps, which is used to derive all the following results. As an example, Fig. 5a shows the visual comparison between manual labels and the predications of the best network. The two datasets are consistent with each other in the number of days when we found anomalous seismicity in September 2019. The visual comparison in other periods of the test data is shown in Fig. S3.

Using these identified timings, we can then find the location where the anomalous seismicity occurred (Fig. 5b). Some earthquakes, e.g., those in the east of the study region associated with the 2019 M_S 5.4 earthquake, are spatially and temporally clustered (Fig. 5b). However, there are spatially separated locations where seismicity nearly emerged at the same time (Fig. 5b), making it difficult to automatically identify by traditional methods.

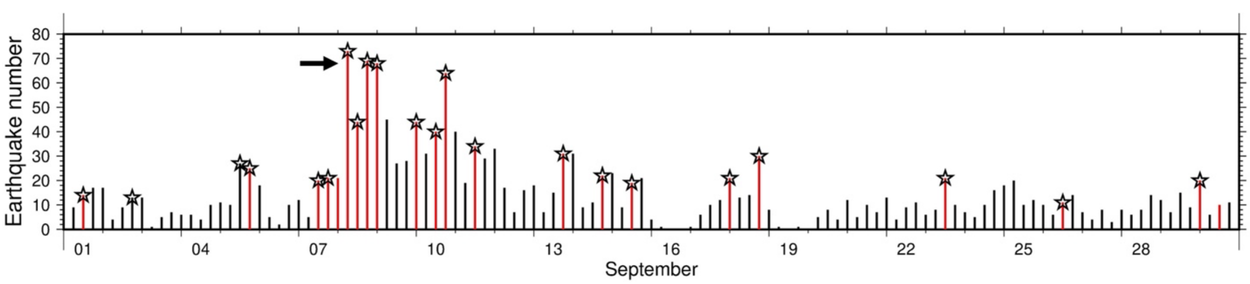

Figure 5. Comparison between manually picked anomalous timings and network identifications in Weiyuan, September 2019. The red bars represent the manually identified anomalous timings.

The stars show the identified timings by the well-trained network. The arrow points to the anomalous seismicity caused by the M_S 5.4 earthquake on September 8th, 2019.

4.2 Application on the Weiyuan catalog

In Weiyuan, we condense the catalog from 24,719 events to 831 abnormal ones after running the detection by our trained network. The locations and timings in the condensed anomalous catalog provide critical information for further investigation of the triggering mechanisms of each earthquake cluster. For example, we identify an anomalous cluster starting from April 30th, 2019, in northeastern Weiyuan (Fig. 6a). Using a 60-hour temporal window and a 2-km spatial grid, we find that earthquakes near the anomalous event show a distinct spatio-temporal pattern (Fig. 6a, c, and e). Most earthquakes were within the pressure diffusion front with a hydraulic diffusivity of 0.8 m2/s, which is consistent with the value estimated in the region (Wong et al., 2021; Sheng et al., 2022), indicating that this earthquake swarm is likely driven by pore pressure diffusion. The other example is a cluster starting from August 08th, 2019, in Weiyuan (Fig. 6b). Before August 08th, 2019, there were no earthquakes around the anomalous location (2 km spatial coverage), but the earthquake number increased drastically later (Fig. 6b). This cluster contains no events with magnitude larger than 3 (Fig. 6f), and the magnitude-time pattern does not suggest an aftershock sequence.

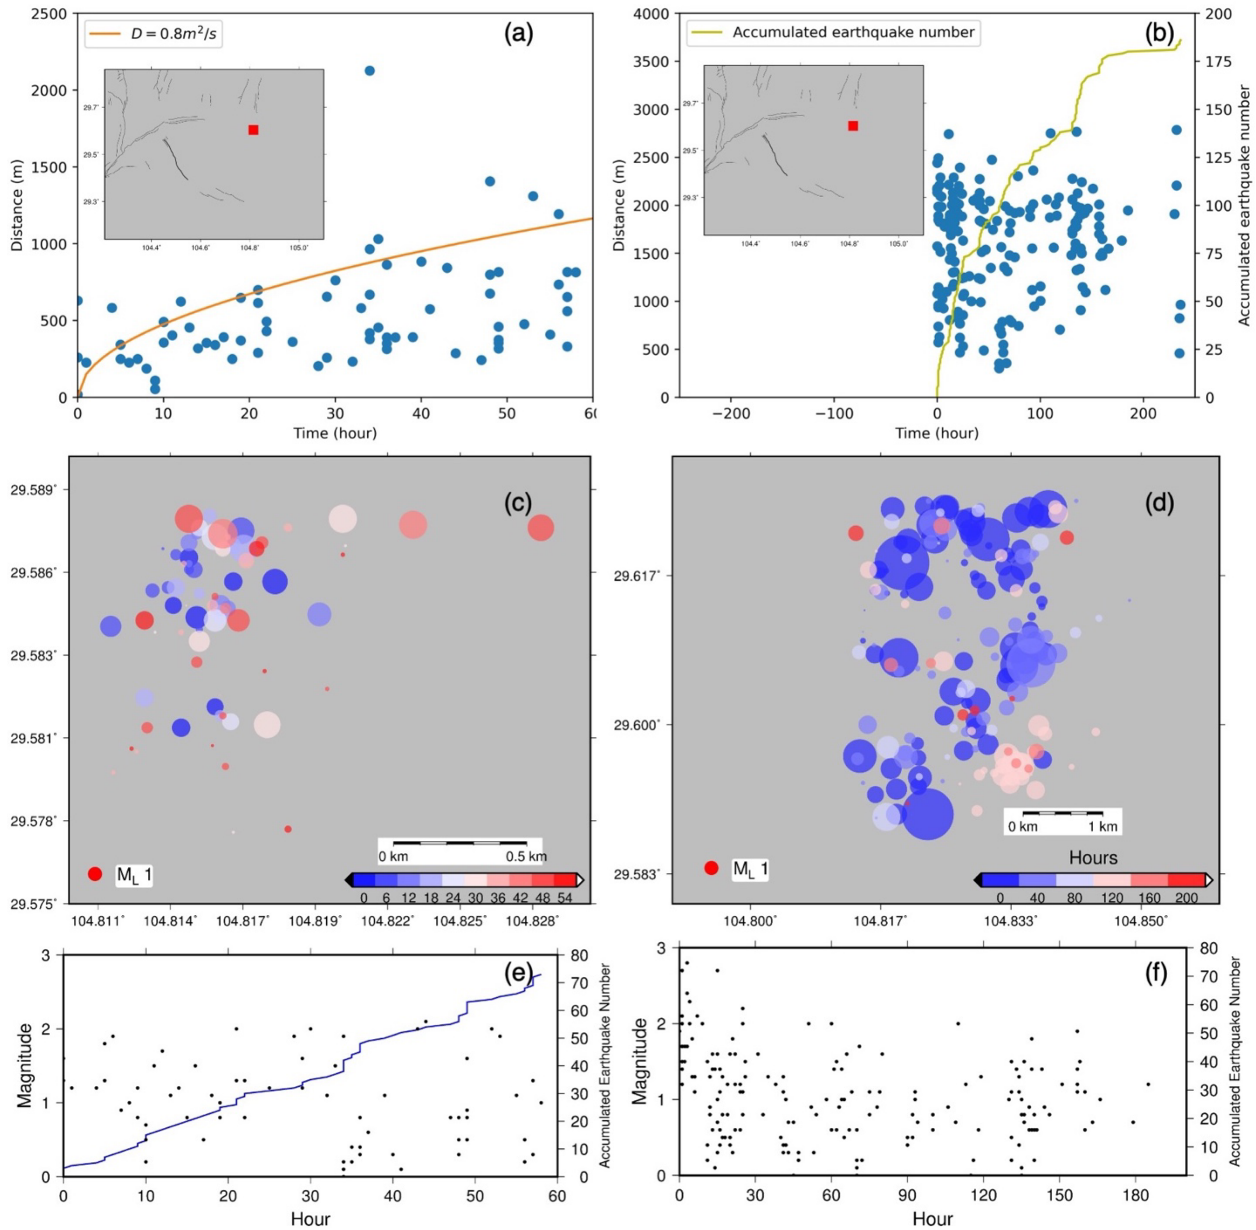

Figure 6. Two examples of identified anomalous earthquake clusters in Weiyuan. (a) shows a cluster occurred on April 30th, 2019, and was potentially driven by pore pressure diffusion. The red block in the map indicates the location of the cluster. (b) shows a sequence occurred on August 08th, 2019, and the earthquake number increased dramatically. The red blocks in the maps of (a) and (b) indicate the locations of the clusters. (c) and (e) show the seismicity of the cluster in (a), while (d) and (f) indicate the one in (b). Both clusters have no large-magnitude earthquakes.

In addition to detecting locations and timings for individual clusters, the condensed catalog could provide insights into the overall earthquake migration during the entire study period (Movie S2). To illustrate, we count the total number of "anomalous" events in each grid to generate a hot map of anomalous seismicity (Fig. 7b), in which a few distinct subregions can be identified in Weiyuan (Fig. 7b). In comparison, identifying such "anomalous" subregions was not straightforward in a traditional map of seismicity, despite that the earthquakes were colored by their occurrence times (Fig. 7a). For instance, we identify one subregion (subregion 1) on the hot map (Fig. 7b), where numerous seismicities occurred in the subregion within two years (Fig. 7a). As identified by our network, the anomalous seismicities were concentrated from March to May and September 2019, respectively (Fig. 8). From March to May 2019, seismicity in subregion 1 clearly exhibited a few clusters, emerging at different times (Fig. 8a). But there were no earthquakes with magnitudes larger than 3 (Fig. 8c). In September 2019, an M_S 5.4 earthquake occurred in the subregion, leading to a group of aftershocks that were identified by our neural network (Fig. 8b & d). In the west of the study region, we also identified one subregion 2 (Fig. 7), which exhibited two anomalous behaviors from April to July 2020 (Fig. 9). First, an earthquake cluster emerged in the northeastern part of subregion 2 (blue dots in Fig. 9a) and lasted until the end of April (Fig. 9b). After a few days of a few earthquakes (Fig. 9b), numerous earthquakes started to occur in the central part of the subregion and then migrated towards its northeast, northwest, and southwest directions.

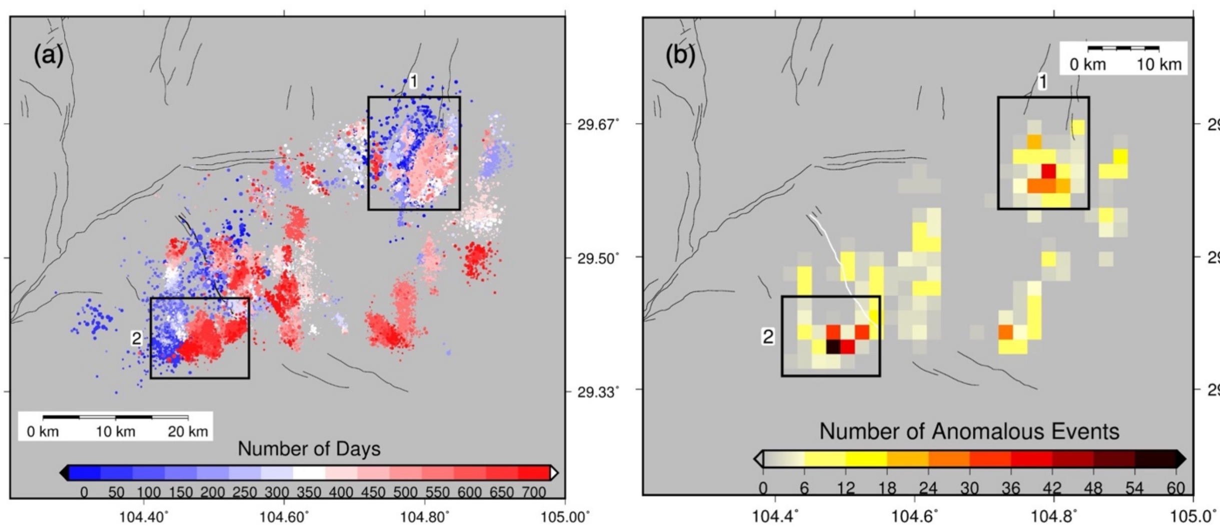

Figure 7. The density of abnormal seismicity in Weiyuan. (a) shows the seismicity of the entire Weiyuan catalog. (b) a hot map of our identified "abnormal" regions, showing the total number of anomalous events in each grid. The subregions are shown as black boxes and are identified based on the density of abnormal events.

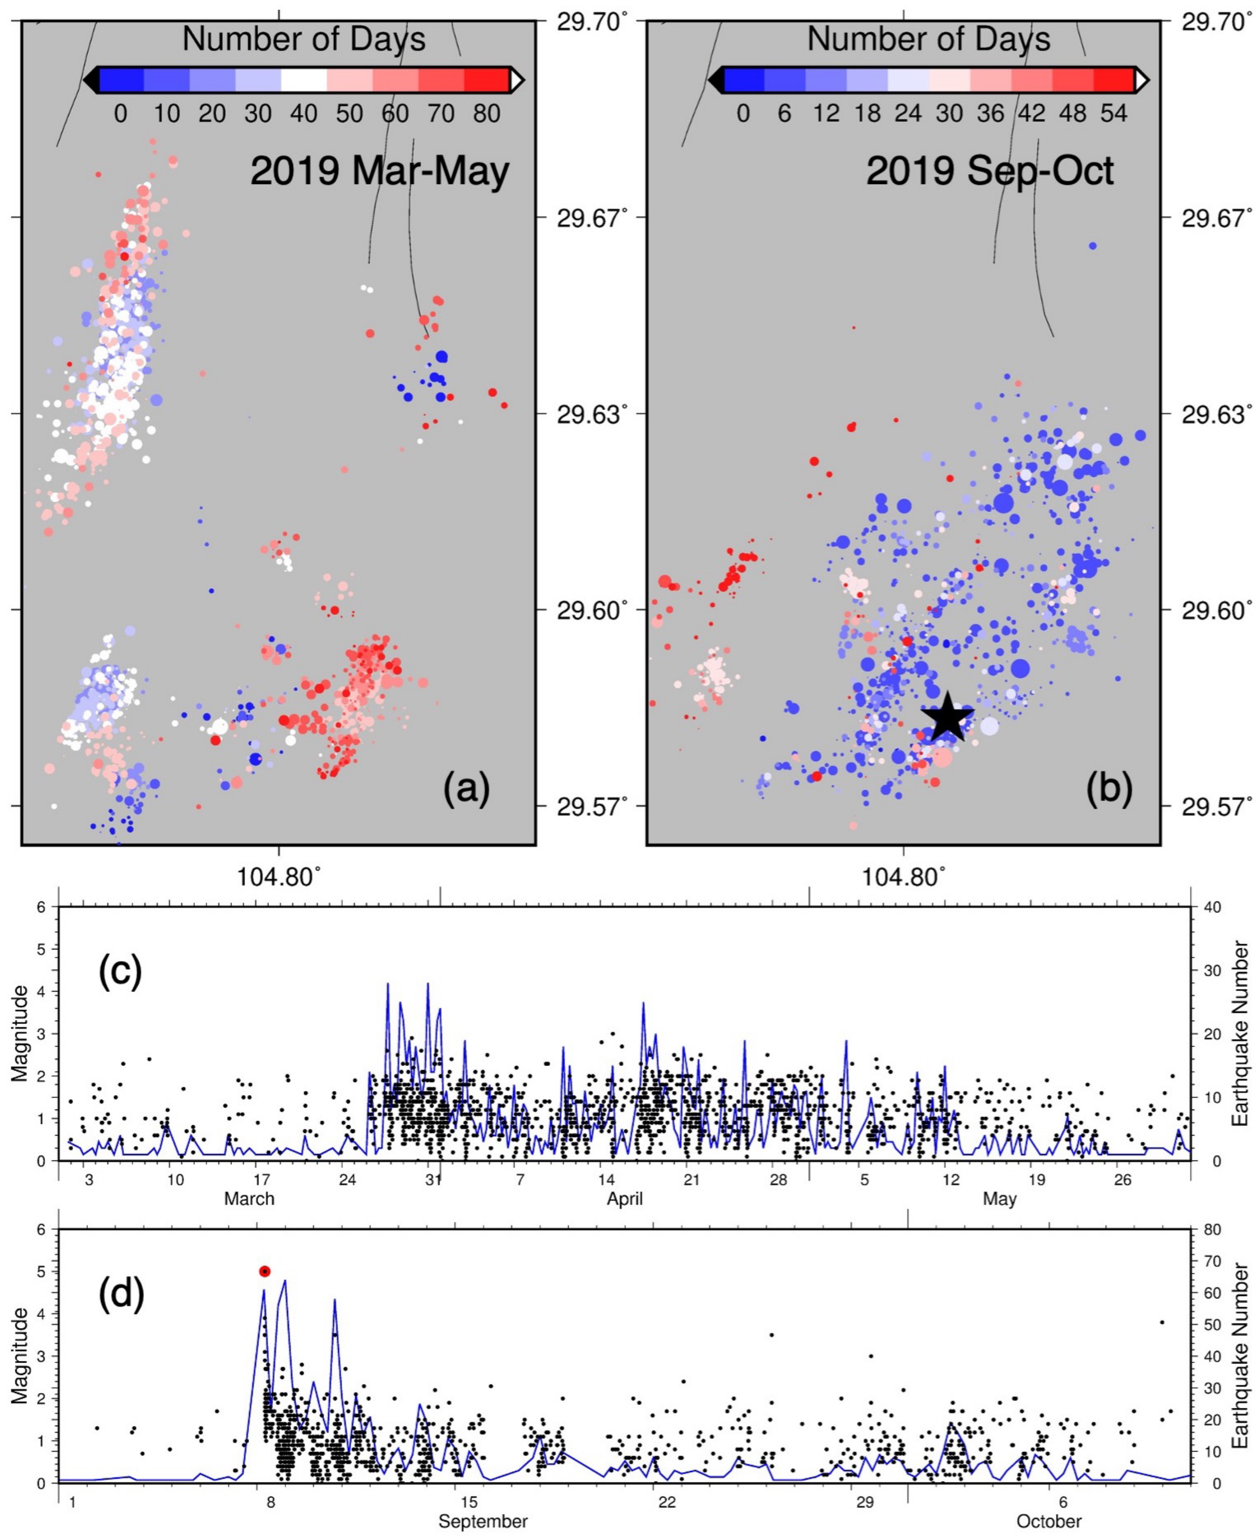

Figure 8. Seismicity in subregion 1 of Weiyuan during abnormal periods. (a) and (b) show the seismicity in subregion 1 from Mar to May 2019 and from Sep to Oct 2019, respectively. (c) and (d) are the corresponding time series of (a) and (b), showing the earthquake magnitudes and

earthquake numbers every six hours in subregion 1. The M_S 5.4 earthquake is shown as the black star in (b) and the red dot in (c).

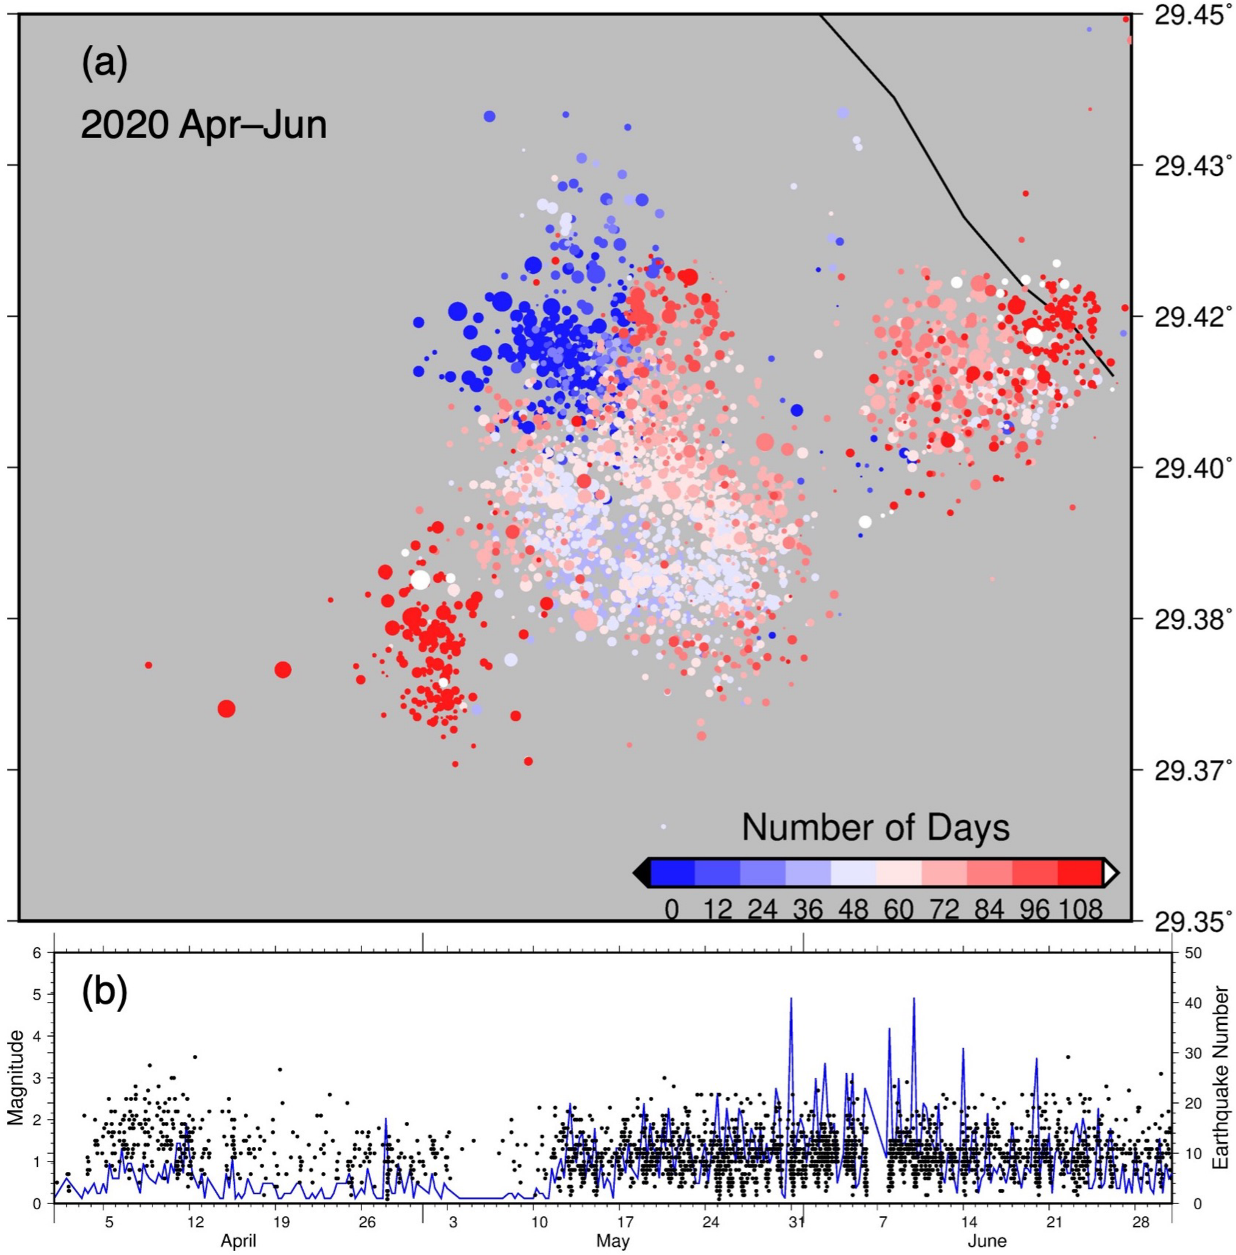

Figure 9. Seismicity in subregion 2 of Weiyuan from April to June 2020 (a) and the corresponding time series of earthquake magnitudes and earthquake numbers every six hours (b). The northeastern cluster (blue) occurred in April, and the other cluster started to emerge in May and migrated in three directions afterward.

4.2 Application on the Changning catalog

After training our network from a subset of seismicity in Changing, we then conduct the network detection and condense the catalog from 18,507 events to 498 based on the manually identified labels. In Changning, the event locations were usually concentrated in small subregions for several months (Movie S3). The hot map of anomalous seismicity in Changning shows four distinct subregions (Fig. 10). For instance, from February to April 2016, anomalous locations were always in subregion 1, forming two obvious earthquake clusters in the northern and southern parts (Fig. S4). From April to May 2017, subregion 1 became active again, exhibiting a cluster in the central zone. In subregion 2, numerous clusters emerged in September 2016 and from January to February 2017. These clusters have no earthquakes with magnitudes larger than five and no distinct migration features (Fig. S5). It was suggested that the anomalous seismicity in subregion 2 was likely caused by hydraulic fracturing operations, which happened from 3rd September to 8th October 2016 and from 12th to 19th January 2017, respectively (Meng et al., 2019).

In addition to anomalous low-magnitude-induced seismicity, our method could also detect anomalies caused by large earthquakes and their aftershocks. For instance, our method detects the anomalies in seismicity caused by an M_W 5.2 earthquake on December 16th, 2018, in subregion 3 (Lei et al., 2019b) and an M_W 5.8 earthquake that occurred on June 17th, 2019, in subregion 4 (Fig. 11). Furthermore, the aftershocks of the large earthquakes could also cause the proliferation of earthquake numbers. Therefore, the anomalous event continuously occurred in subregion 3 from December 2018 to January 2019 (Fig. 12a & c) and in subregion 4 from June to July 2019 (Fig. 12b & d). Although it is not our primary motivation to detect aftershocks and practically it is not necessary to use such an algorithm because the large earthquakes already serve as landmarks, the ability to detect anomalies caused by large-magnitude earthquakes shows the feasibility of applying our algorithm in natural earthquake sequences. It also demonstrates the effectiveness of using earthquake numbers without magnitudes and depth.

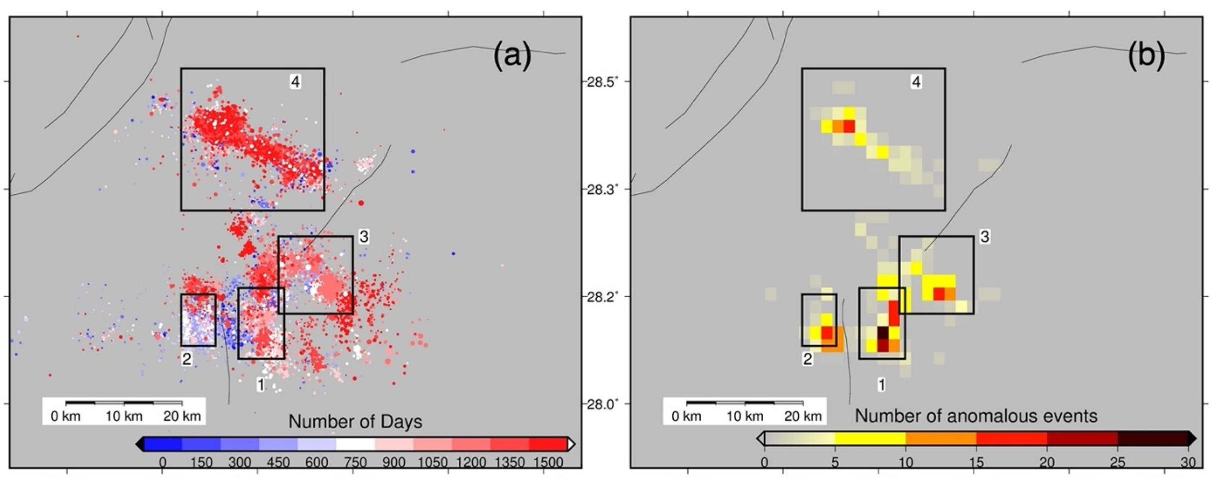

Figure 10. The density of abnormal seismicity in Changning. (a) shows the seismicity of the entire Changning catalog. (b) hot map of abnormal seismicity, showing the total number of anomalous events in each grid. The subregions are shown as black boxes and are identified based on the density of abnormal events.

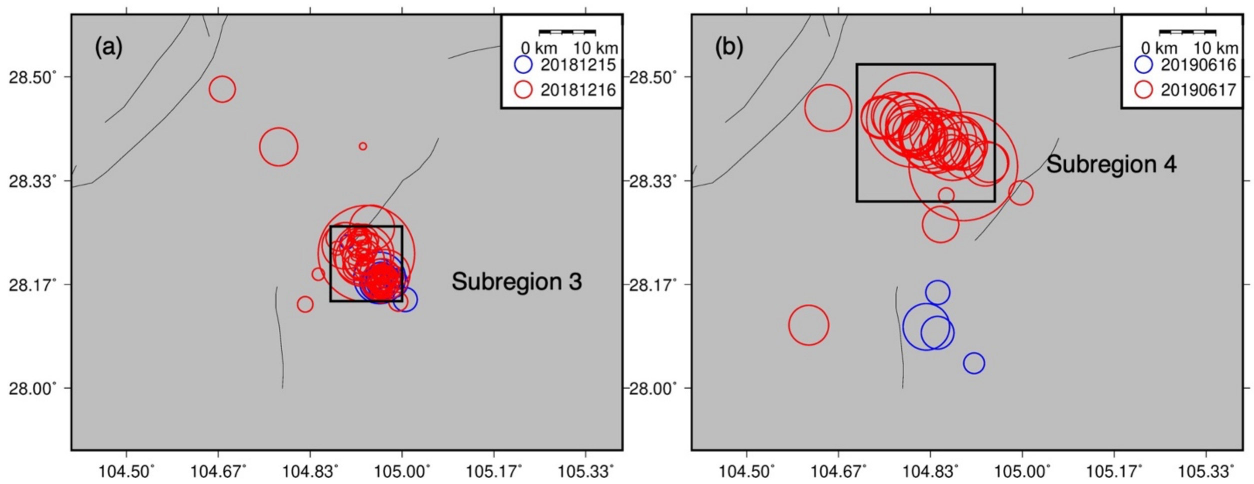

Figure 11. Anomalies caused by large-magnitude earthquakes in Changning. (a) shows the anomaly on December 16th, 2018 in subregion 3. (b) shows the anomaly on June 17th, 2019 in subregion 4. The circles in the legend box indicate the M_L 1 earthquakes.

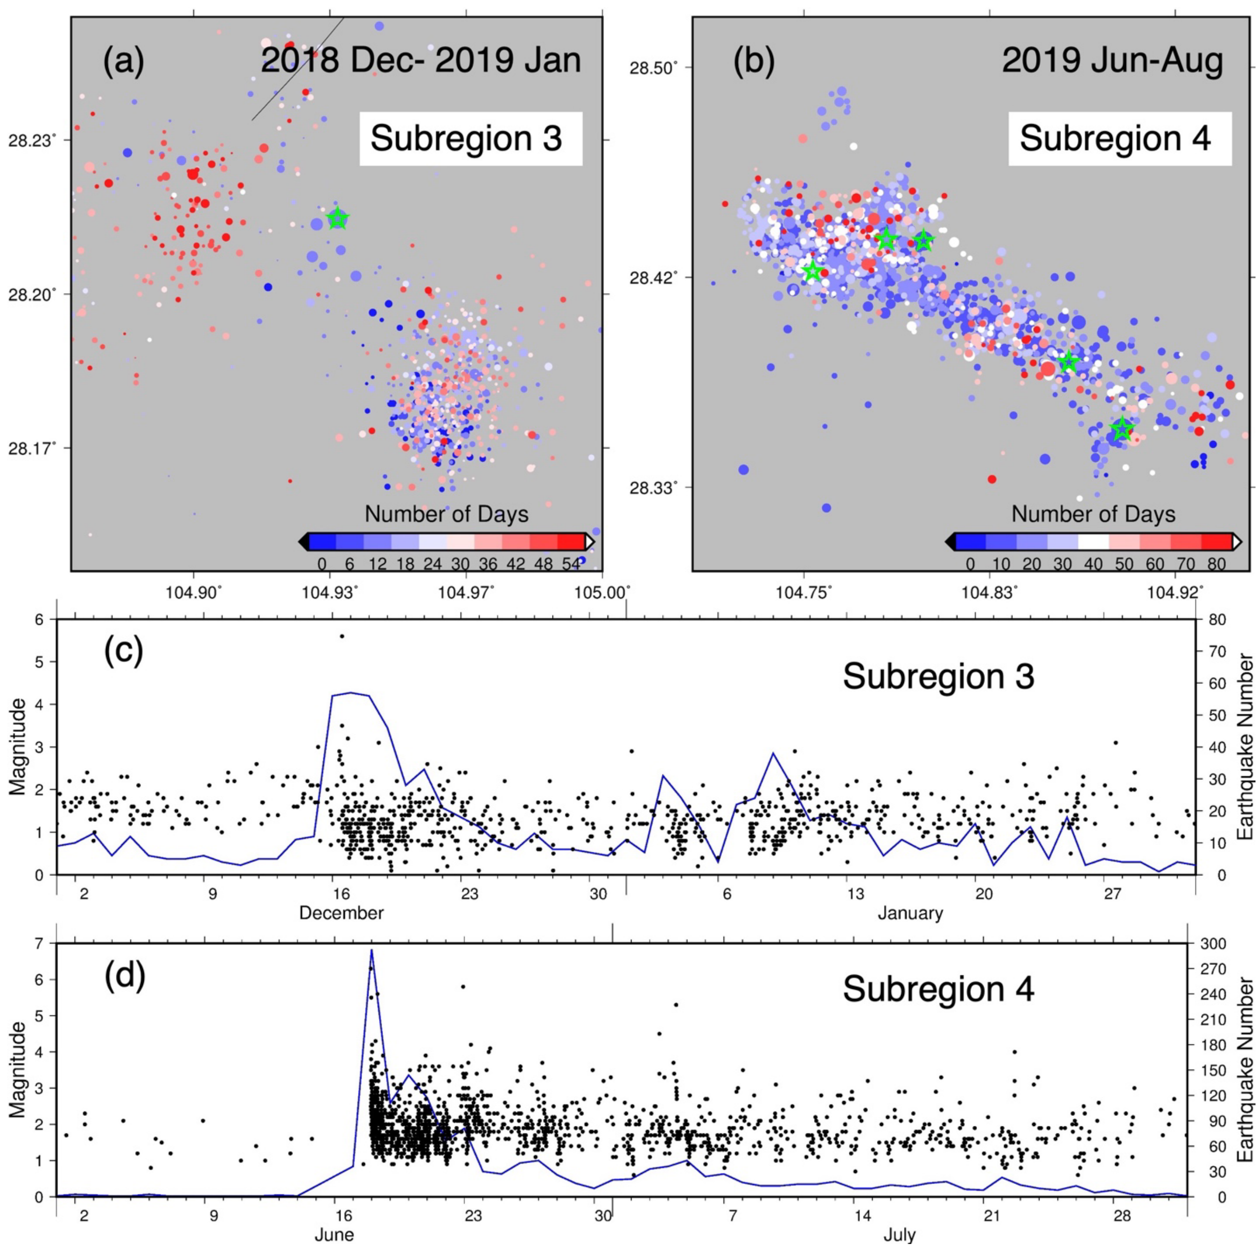

Figure 12. Seismicity in subregions 3 and 4 of Changning. (a) and (b) show the seismicity in subregions 3 and 4 during their active period, respectively. (c) and (d) are the corresponding time series of (a) and (b), showing the earthquake magnitudes and daily earthquake numbers in subregions 3 and 4. The abnormal seismicities in both subregions 3 and 4 are triggered by large earthquakes, shown as green stars in both (a) and (b).

5. Discussion

5.1 Choosing the appropriate time interval

The choice of temporal resolution directly impacts the results, as the temporal density of earthquakes determines the time interval of distribution maps. A large temporal resolution could cause a large proportion of abnormal timings, while a too-small interval could miss some abnormal timings. Here, we test the effect of the time interval for the Weiyuan catalog by setting it as 24 and 4 hours, respectively, and compare it with our optimal 6-hour interval. For the 24-hour interval, we identify 315 abnormal timing from 724 distribution maps. While most of the reported timings were correct, the large proportion of anomalous maps violates our objective of saving manual practice. In addition, the reported timings are less precise than using a six-hour interval. In comparison, we condense the original catalog to 572 events when using the 4-hour interval, less than the number (831) using the 6-hour interval. The 4-hour interval misses some events because a shorter temporal interval dilutes the earthquake distribution for each map, and some changes are not intense enough to be detected (Fig. 13b).

The results of the above test show that the choice of time interval to generate the distribution maps can lead to a trade-off between the accuracy of reported timing and the amount of manual practice. To overcome this, we may choose the strategy of adaptive time interval, i.e., find the anomalous timing in a relatively large time window first and then use a finer temporal resolution. Furthermore, the absolute value of the time interval is subject to the total number of earthquakes in a selected region, which depends on the background seismicity rate in the region and the total duration of the catalog. For induced earthquakes, the study region and catalog duration are usually selected according to the area of interest and time scales of industrial activities. Therefore, the temporal resolution in our algorithm can be set and adjusted subject to research focus.

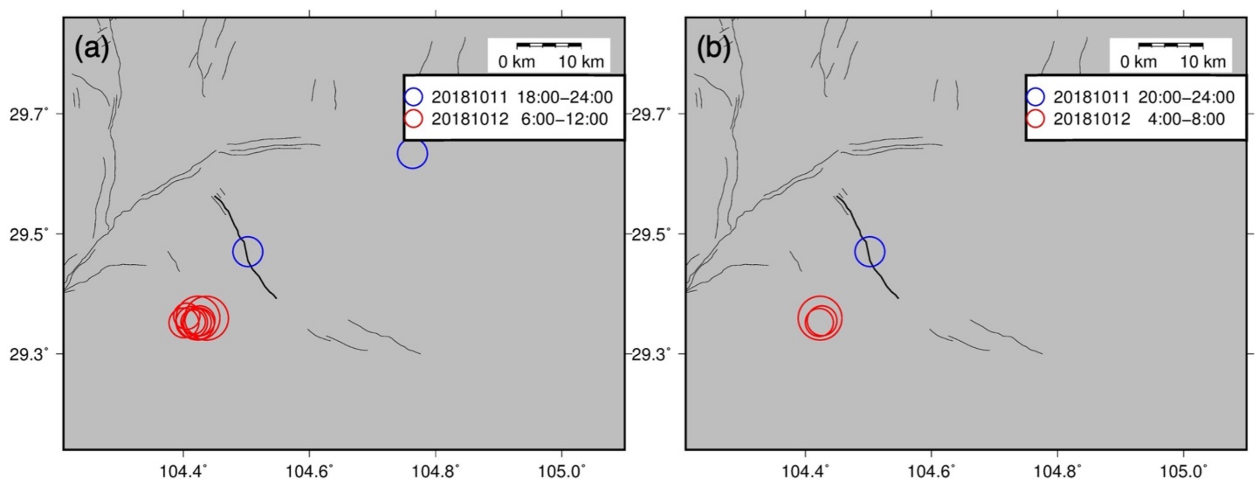

Figure 13. Example of missing events due to the finer temporal resolution. (a) is identified as abnormal but (b) is not because the changes in (b) are not intense enough. The circles in the legend box indicate the M_L 1 earthquakes.

5.2 Potential applications on natural earthquakes

Compared with induced earthquakes in the same term, catalogs of natural earthquakes may have fewer earthquake numbers and thus a lower frequency of anomalies. However, our algorithm can be directly applied on detecting “anomalous” changes in natural earthquakes. As demonstrated by the aftershock sequences of the 2019 M_W 5.8 Changning and the 2018 M_W 5.2 Xingwen earthquakes, our algorithm is effective in detecting such changes in the amount of seismicity in a relatively small region. Although it is not necessary to identify the emergence of aftershocks with such an advanced technique, it may be applicable to investigate detailed aftershock evolution provided that some well-identified training datasets are available.

As we do not need to include magnitude information of earthquakes, it is obvious that our algorithm can be effective to identify earthquake swarms, an earthquake sequence in which no clear large-magnitude events (mainshock) exist. Normally earthquake swarms have been considered associated with fluid migration (Shelly et al., 2013); thus, mapping out swarms may advance our understanding of subsurface fluid transportation.

In addition, earthquakes may exhibit in foreshock-mainshock sequences, i.e., a series of small magnitudes preceding a large event in a time window from days to weeks (Kato and Nakagawa, 2014; Yao et al., 2020; Zhang et al., 2022; Zhu et al., 2022). Indeed, nearly 50% of large earthquakes had foreshocks, particularly for interplate events (Jones and Molnar, 1976; Bouchon et al., 2013). Despite the mechanisms driving foreshocks remain controversial (Zhu et al., 2022), foreshocks have been considered as the most reliable precursors that are hopeful for earthquake prediction. Because foreshocks often have low magnitudes, similar to the earthquakes used in this study, it is anticipated that our algorithm is capable of detecting the foreshock migration.

However, it is extremely challenging to recognize an ongoing foreshock sequence (Brodsky and Lay, 2014). Although our algorithm is able to identify an emerging earthquake sequence, it is a well-known difficult problem to distinguish them from the background seismicity. If there are well-recorded catalogs of background seismicity and large events with profound foreshock sequences, we may train our network to learn their features, respectively. Should there be distinct features between events leading to large earthquakes and these background ones, our network is hopeful to gain the capability of labeling potential foreshock sequences. The network's generalization ability of course needs to be tested in various regions and should be done in future studies.

5.3 Limitations and possible solutions

Although our method achieves promising results in the Sichuan Basin, it has several limitations. First, the network might fail to identify small-distance migration, depending on the spatial grid in our model and the location resolution of the training catalog. Second, when anomalous seismicity occurs frequently, the network might not automatically pick the precise onset of the anomaly, depending on the temporal resolution. These two cases are caused by the leak of corresponding training examples. The Changning catalog does not contain enough cases where anomalous seismicity migrated in a small distance or occurred frequently. Including more catalogs into the training data could increase the method's generalization and accordingly help to solve these misidentifications. Third, the current study does not consider the hypocenter depth of earthquakes due to the difficulties of visualizing 3-D matrixes and preparing the training labels. With more manual practice and higher catalog accuracy in the future, we could consider depth information in the network.

6. Conclusion

This study designs a novel method based on deep learning to automate the detection of anomalous seismicity. The detected locations and timings provide important information for investigating the triggering mechanisms of each earthquake cluster. Our approach could condense a large earthquake catalog to a focused catalog containing only anomalous events, saving intensive manual practice. Based on the condensed catalogs, our further analysis reveals seismically abnormal subregions in the Changning and Weiyuan shale gas field and their corresponding active periods. In addition to the anomalous seismicity caused by anthropic activities, the method could also detect anomalies caused by large natural earthquakes. Owing to the large learning capability of deep learning, we could apply the method to broader areas with more complex and diverse earthquake migration patterns.

Acknowledgment:

This study is supported by Natural Science Foundation of China (U2139203), Hong Kong Research Grants Council (14303721), CRIMS and Direct Grants from Faculty of Science at the Chinese University of Hong Kong. We also thank Dr. Lingyuan Meng for providing the Changning catalog, Jeremy Wing Ching Wong and Jinping Zi for providing the Weiyuan catalog.

Open Research:

The catalogs and codes are available on https://github.com/enzezhang/DLSeismicAnomaly.

Reference

Anderson, J. G. (1994). Seismicity in the western Great Basin apparently triggered by the Landers, California, earthquake, 28 June 1992. *Bulletin - Seismological Society of America*, *84*(3), 863–891.

Atkinson, G. M., Eaton, D. W., Ghofrani, H., Walker, D., Cheadle, B., Schultz, R., et al. (2016). Hydraulic fracturing and seismicity in the western Canada sedimentary basin. *Seismological Research Letters*, *87*(3), 631–647. https://doi.org/10.1785/0220150263

Atkinson, G. M., Eaton, D. W., & Igonin, N. (2020). Developments in understanding seismicity triggered by hydraulic fracturing. *Nature Reviews Earth and Environment*, *1*(5), 264–277. https://doi.org/10.1038/s43017-020-0049-7

Bao, X., & Eaton, D. W. (2016). Fault activation by hydraulic fracturing in western Canada. *Science*, *354*(6318), 1406–1409. https://doi.org/10.1126/science.aag2583

Barbot, S., Fialko, Y., & Bock, Y. (2009). Postseismic deformation due to the M_W 6.0 2004 Parkfield earthquake: Stress-driven creep on a fault with spatially variable rate-and-state friction parameters. *Journal of Geophysical Research: Solid Earth*, *114*(7), 1–26.

https://doi.org/10.1029/2008JB005748

Bouchon, M., Durand, V., Marsan, D., Karabulut, H., & Schmittbuhl, J. (2013). The long precursory phase of most large interplate earthquakes. *Nature Geoscience*, *6*(4), 299–302. https://doi.org/10.1038/ngeo1770

Brodsky, E. E., & Lay, T. (2014). Recognizing foreshocks from the 1 April 2014 Chile earthquake. *Science*, *344*(6185), 700–702. https://doi.org/10.1126/science.1255202

Freed, A. M. (2005). Earthquake triggering by static, dynamic, and postseismic stress transfer. *Annual Review of Earth and Planetary Sciences*, *33*, 335–367. https://doi.org/10.1146/annurev.earth.33.092203.122505

Friberg, P. A., Besana-Ostman, G. M., & Dricker, I. (2014). Characterization of an earthquake sequence triggered by hydraulic fracturing in Harrison county, Ohio. *Seismological Research Letters*, *85*(6), 1295–1307. https://doi.org/10.1785/0220140127

Grigoli, F., Cesca, S., Rinaldi, A. P., Manconi, A., López-Comino, J. A., Clinton, J. F., et al. (2018). The November 2017 M_w 5.5 Pohang earthquake: A possible case of induced seismicity in South Korea. *Science*, *360*(6392), 1003–1006. https://doi.org/10.1126/science.aat2010

Haffener, J., Chen, X., & Murray, K. (2018). Multiscale Analysis of Spatiotemporal Relationship Between Injection and Seismicity in Oklahoma. *Journal of Geophysical Research: Solid Earth*, *123*(10), 8711–8731. https://doi.org/10.1029/2018JB015512

He, K., Zhang, X., Ren, S., & Sun, J. (2015). Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *Proceedings of the IEEE International Conference on Computer Vision*, *2015 Inter*, 1026–1034. https://doi.org/10.1109/ICCV.2015.123

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (Vol. 2016-Decem, pp. 770–778). IEEE Computer Society. https://doi.org/10.1109/CVPR.2016.90

Hill, D. P., Reasenberg, P. a, Michael, a, Arabaz, W. J., Beroza, G., Brumbaugh, D., et al. (1993). Remotely Triggered. *Science*, *260*(June), 1617–1623.

Holland, A. A. (2013). Earthquakes triggered by hydraulic fracturing in south-central Oklahoma. *Bulletin of the Seismological Society of America*, *103*(3), 1784–1792. https://doi.org/10.1785/0120120109

Jia, K., Zhou, S., Zhuang, J., Jiang, C., Guo, Y., Gao, Z., et al. (2020). Nonstationary background seismicity rate and evolution of stress changes in the changning salt mining and shale-gas hydraulic fracturing region, Sichuan basin, China. *Seismological Research Letters*, *91*(4), 2170–2181. https://doi.org/10.1785/0220200092

Johann, L., & Shapiro, S. A. (2020). Understanding vectorial migration patterns of wastewater-induced earthquakes in the united states. *Bulletin of the Seismological Society of America*, *110*(5), 2295–2307. https://doi.org/10.1785/0120200064

Johnson, S. W., Chambers, D. J. A., Boltz, M. S., & Koper, K. D. (2021). Application of a convolutional neural network for seismic phase picking of mining-induced seismicity. *Geophysical Journal International*, *224*(1), 230–240. https://doi.org/10.1093/gji/ggaa449

Jones, L., & Molnar, P. (1976). Frequency of foreshocks. *Nature*, *262*(5570), 677–679. https://doi.org/10.1038/262677a0

Kato, A., & Nakagawa, S. (2014). Multiple slow-slip events during a foreshock sequence of the 2014 Iquique, Chile M_W 8.1 earthquake Aitaro. *Geophysical Research Letters*, *41*(15), 5420–5427. https://doi.org/10.1002/2014GL061138

Kilb, D., Gomberg, J., & Bodin, P. (2000). Triggering of earthquake aftershocks by dynamic stresses. *Nature*, *408*(6812), 570–574. https://doi.org/10.1038/35046046

King, G. C. P., Stein, R. S., & Lin, J. (1994). *Static Stress Changes and the Triggering of Earthquakes*. *Bulletin of the Seismological Society of America* (Vol. 84). Retrieved from http://pubs.geoscienceworld.org/ssa/bssa/article-pdf/84/3/935/5341861/bssa0840030935.pdf

Kuang, W., Yuan, C., & Zhang, J. (2021). Real-time determination of earthquake focal mechanism via deep learning. *Nature Communications*, *12*(1). https://doi.org/10.1038/s41467-021-21670-x

Lei, X., Huang, D., Su, J., Jiang, G., Wang, X., Wang, H., et al. (2017). Fault reactivation and earthquakes with magnitudes of up to M_W 4.7 induced by shale-gas hydraulic fracturing in Sichuan Basin, China. *Scientific Reports*, *7*(1), 1–12. https://doi.org/10.1038/s41598-017-08557-y

Lei, X., Wang, Z., & Su, J. (2019a). Possible link between long-term and short-term water injections and earthquakes in salt mine and shale gas site in Changning, south Sichuan Basin, China. *Earth and Planetary Physics*, *3*(6), 510–525. https://doi.org/10.26464/epp2019052

Lei, X., Wang, Z., & Su, J. (2019b). The December 2018 M_L 5.7 and January 2019 M_L 5.3 earthquakes in South Sichuan basin induced by shale gas hydraulic fracturing. *Seismological Research Letters*, *90*(3), 1099–1110. https://doi.org/10.1785/0220190029

Liu, J., & Zahradník, J. (2020). The 2019 M_W 5.7 Changning Earthquake, Sichuan Basin, China: A Shallow Doublet With Different Faulting Styles. *Geophysical Research Letters*, *47*(4), 1–9. https://doi.org/10.1029/2019GL085408

Meng, L., McGarr, A., Zhou, L., & Zang, Y. (2019). An investigation of seismicity induced by hydraulic fracturing in the Sichuan Basin of China based on data from a temporary seismic network. *Bulletin of the Seismological Society of America*, *109*(1), 348–357. https://doi.org/10.1785/0120180310

Peng, Z., & Zhao, P. (2009). Migration of early aftershocks following the 2004 Parkfield earthquake. *Nature Geoscience*, *2*(12), 877–881. https://doi.org/10.1038/ngeo697

Prejean, S. G., Hill, D. P., Brodsky, E. E., Hough, S. E., Johnston, M. J. S., Malone, S. D., et al. (2004). Remotely triggered seismicity on the United States west coast following the Mw 7.9 Denali fault earthquake. *Bulletin of the Seismological Society of America*, *94*(6 SUPPL. B), 348–359. https://doi.org/10.1785/0120040610

Schultz, R., Skoumal, R. J., Brudzinski, M. R., Eaton, D., Baptie, B., & Ellsworth, W. (2020). Hydraulic fracturing-induced seismicity. *Reviews of Geophysics*, *58*(3), 1–43. https://doi.org/10.1029/2019RG000695

Shelly, D. R., Hill, D. P., Massin, F., Farrell, J., Smith, R. B., & Taira, T. (2013). A fluid-driven earthquake swarm on the margin of the Yellowstone caldera. *Journal of Geophysical Research E: Planets*, *118*(9), 4872–4886. https://doi.org/10.1002/jgrb.50362

Sheng, M., Chu, R., Ni, S., Wang, Y., Jiang, L., & Yang, H. (2020). Source Parameters of Three Moderate Size Earthquakes in Weiyuan, China, and Their Relations to Shale Gas Hydraulic Fracturing. *Journal of Geophysical Research: Solid Earth*, *125*(10), 1–14. https://doi.org/10.1029/2020JB019932

Sheng, M., Chu, R., Peng, Z., & Wei, Z. (2022). Earthquakes Triggered by Fluid Diffusion and Boosted by Fault Reactivation in Weiyuan , China Due to Hydraulic Fracturing *Journal of Geophysical Research : Solid Earth*, 1–18. https://doi.org/10.1029/2021JB022963

Tan, Y., Hu, J., Zhang, H., Chen, Y., Qian, J., Wang, Q., et al. (2020). Hydraulic Fracturing Induced Seismicity in the Southern Sichuan Basin Due to Fluid Diffusion Inferred From Seismic and Injection Data Analysis. *Geophysical Research Letters*, *47*(4), 1–10. https://doi.org/10.1029/2019GL084885

Waldhauser, F., & Ellsworth, W. L. (2000). A Double-difference Earthquake location algorithm: Method and application to the Northern Hayward Fault, California. *Bulletin of the Seismological Society of America*, *90*(6), 1353–1368. https://doi.org/10.1785/0120000006

Wang, M., Yang, H., Fang, L., Han, L., Jia, D., Jiang, D., & Yan, B. (2020). Shallow faults reactivated by hydraulic fracturing: The 2019 Weiyuan Earthquake Sequences in Sichuan, China. *Seismological Research Letters*, *91*(6), 3171–3181. https://doi.org/10.1785/0220200174

Wong, W. C. J., Zi, J., Yang, H., & JinRong Su, A. (2021). Spatial-temporal Evolution of Injection Induced Earthquakes in Weiyuan Area by Machine-Learning Phase Picker and Waveform Cross-correlation. *Earth and Planetary Physics*, *5*(0), 0–0. https://doi.org/10.26464/epp2021055

Yang, H., Zhou, P., Fang, N., Zhu, G., Xu, W., Su, J., et al. (2020). A shallow shock: The 25 February 2019 M_L 4.9 Earthquake in the Weiyuan Shale Gas Field in Sichuan, China. *Seismological Research Letters*, *91*(6), 3182–3194. https://doi.org/10.1785/0220200202

Yao, D., Huang, Y., Peng, Z., & Castro, R. R. (2020). Detailed Investigation of the Foreshock Sequence of the 2010 M_W 7.2 El Mayor-Cucapah Earthquake. *Journal of Geophysical Research: Solid Earth*, *125*(6). https://doi.org/10.1029/2019JB019076

Yun, N., Zhou, S., Yang, H., Yue, H., & Zhao, L. (2019). Automated Detection of Dynamic Earthquake Triggering by the High-Frequency Power Integral Ratio. *Geophysical Research*

Letters, *46*(22), 12977–12985. https://doi.org/10.1029/2019GL083913

Zhang, H., & Thurber, C. H. (2003). Double-difference tomography: The method and its application to the Hayward Fault, California. *Bulletin of the Seismological Society of America*, *93*(5), 1875–1889. https://doi.org/10.1785/0120020190

Zhang, X., Zhang, J., Yuan, C., Liu, S., Chen, Z., & Li, W. (2020). Locating induced earthquakes with a network of seismic stations in Oklahoma via a deep learning method. *Scientific Reports*, *10*(1), 1–12. https://doi.org/10.1038/s41598-020-58908-5

Zhang, Y., Y. An, F. Long, G. Zhu, M Qin, Y. Zhong, Q. Xu, and H. Yang (2022), Short-term foreshock and aftershock patterns of the 2021 *Ms* 6.4 Yangbi earthquake sequence, *Seismological Research Letters*, doi:10.1785/0220210154.

Zhou, P. C., Ellsworth, W. L., Yang, H. F., Tan, Y. J., Beroza, G. C., Sheng, M. H., & Chu, R. S. (2021). Machine-learning-facilitated earthquake and anthropogenic source detections near the Weiyuan Shale Gas Blocks, Sichuan, China. *Earth and Planetary Physics*, *5*(6), 532–546. https://doi.org/10.26464/epp2021053

Zhu, G., H. Yang, Y. Tan, M. Jin, X. Li, and W. Yang (2022) The cascading foreshock sequence of the Ms 6.4 Yangbi earthquake in Yunnan, China, *Earth and Planetary Sci. Lett.,* https://doi.org/10.1016/j.epsl.2021.117594

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

Zou, C., Ni, Y., Li, J., Kondash, A., Coyte, R., Lauer, N., et al. (2018). The water footprint of hydraulic fracturing in Sichuan Basin, China. *Science of the Total Environment*, *630*, 349–356. https://doi.org/10.1016/j.scitotenv.2018.02.219

AGU PUBLICATIONS

[Journal of Geophysical Research: Solid Earth]

Supporting Information for

[An automated, deep-learning-based method for investigating spatial-temporal evolution of seismicity]

[Enze Zhang[1,2†] and Hongfeng Yang[1,3]]

[1Earth System Science Programme, The Chinese University of Hong Kong, Sha Tin, NT, Hong Kong, China.

2Institute for Geophysics, The University of Texas at Austin, Austin, TX, USA.

3Shenzhen Research Institute, The Chinese University of Hong Kong, Shenzhen, China.

†Current address: Institute for Geophysics, The University of Texas at Austin, Austin, TX, USA.]

Contents of this file

Text S1
Figures S1 to S5
Tables S1 to S6

Introduction

The supporting information provides additional information about the method development, model details, model validation, and further investigation on derived results.

Text S1. Choosing a proper learning rate and L2 regularization factor

We use the final validation loss to choose the learning rate and L2 regularization factor. If the learning rate is too large, the training loss will decrease greatly at the beginning but show little change later on, and the final validation loss is high. If the learning rate is too small, the training loss curve will show a nearly linear trend. The final validation loss will be high since it needs more training epochs to decrease than the patience threshold we set. L2 regularization factor penalizes large weights, and therefore it could suppress any irrelevant components of the weight vector by choosing the smallest vector that decreases the training loss. A too small weight decay can hardly have enough regularization effect while a too large weight decay could hurt the network training progress. From Table S1 to Table S6 shows the final validation loss by using different values of weight decay and learning rate for the three networks. For each network, we choose the learning rate and weight decay with the smallest final validation loss.

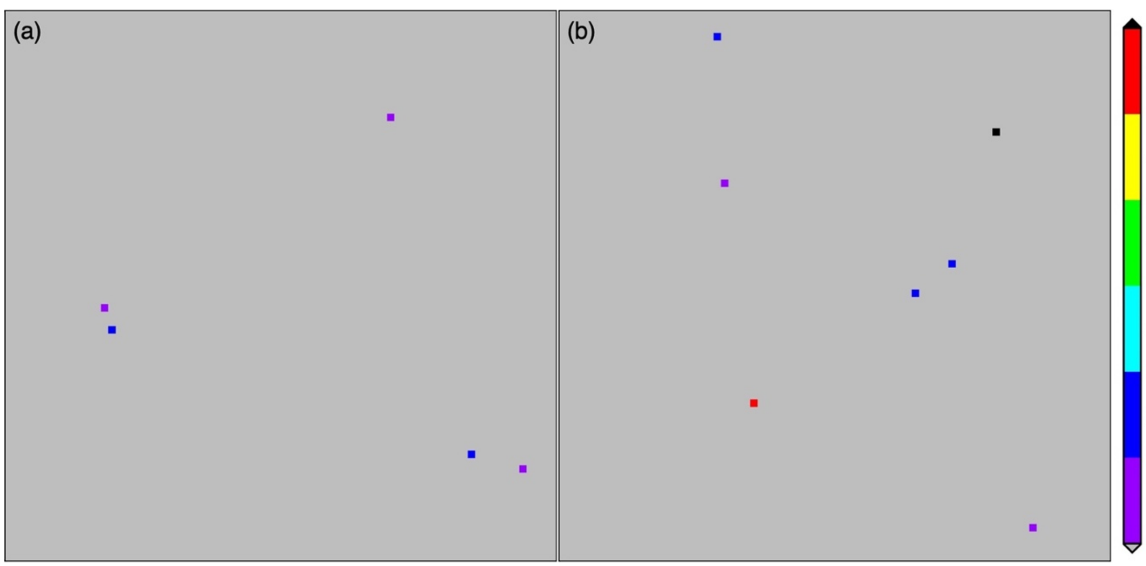

Figure S1. Example of earthquake artificial distributions maps. (a) shows the scatter distribution and (b) show the concentrated distribution with pixels having more than five earthquakes.

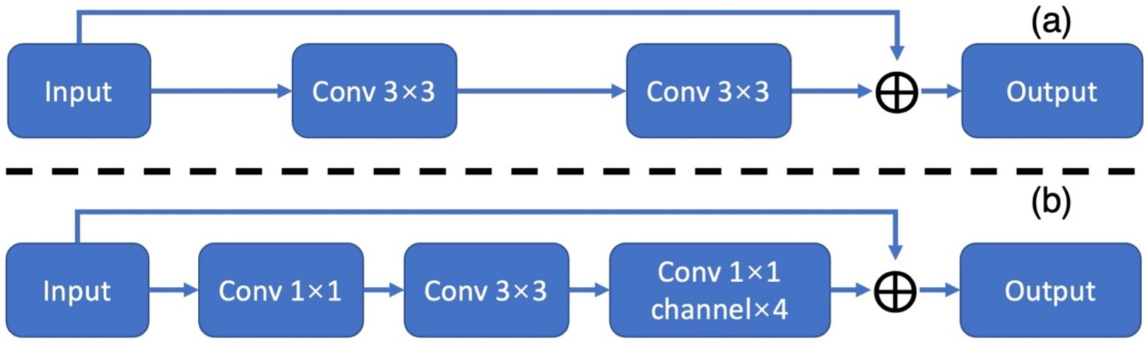

Figure S2. (a) The building block of residual learning used in ResNet-18 and ResNet-34. (b) The building block of residual learning used in ResNet-101. For both types of building blocks, the block adds its input with the feature map after going through convolutional layers to obtain the output.

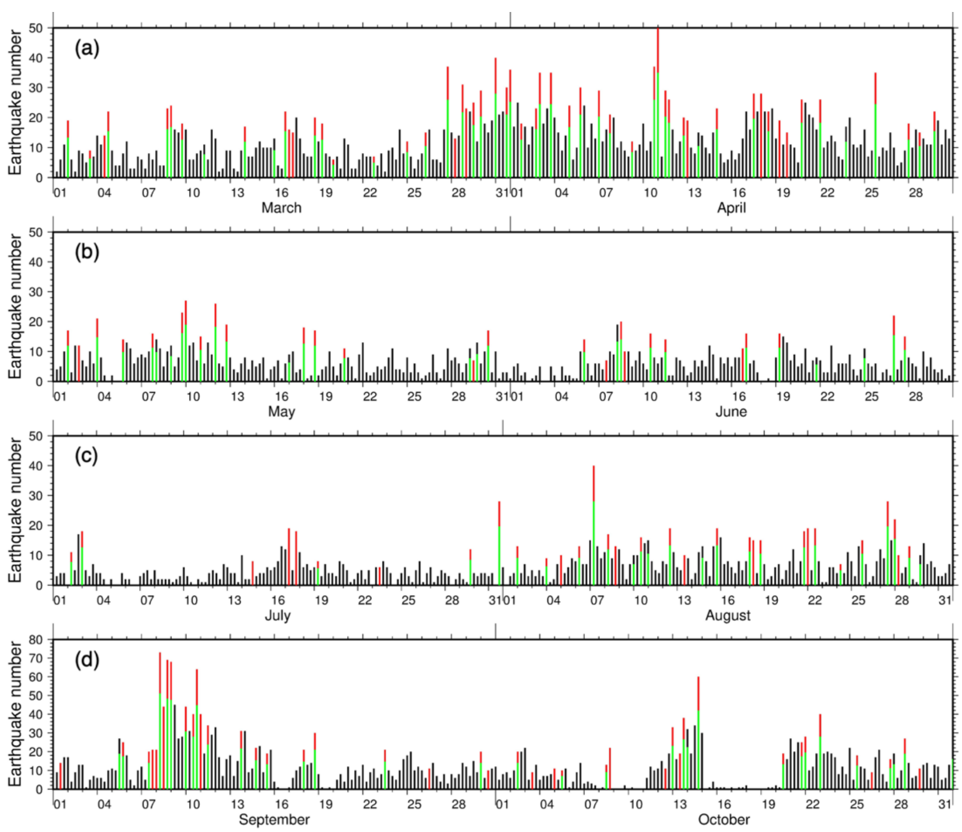

Figure S3. Comparison between manually-picked anomalous timings and network identifications in Weiyuan, 2019. The red bars represent the manually-identified anomalous timings. The green ones are identified by the well-trained network.

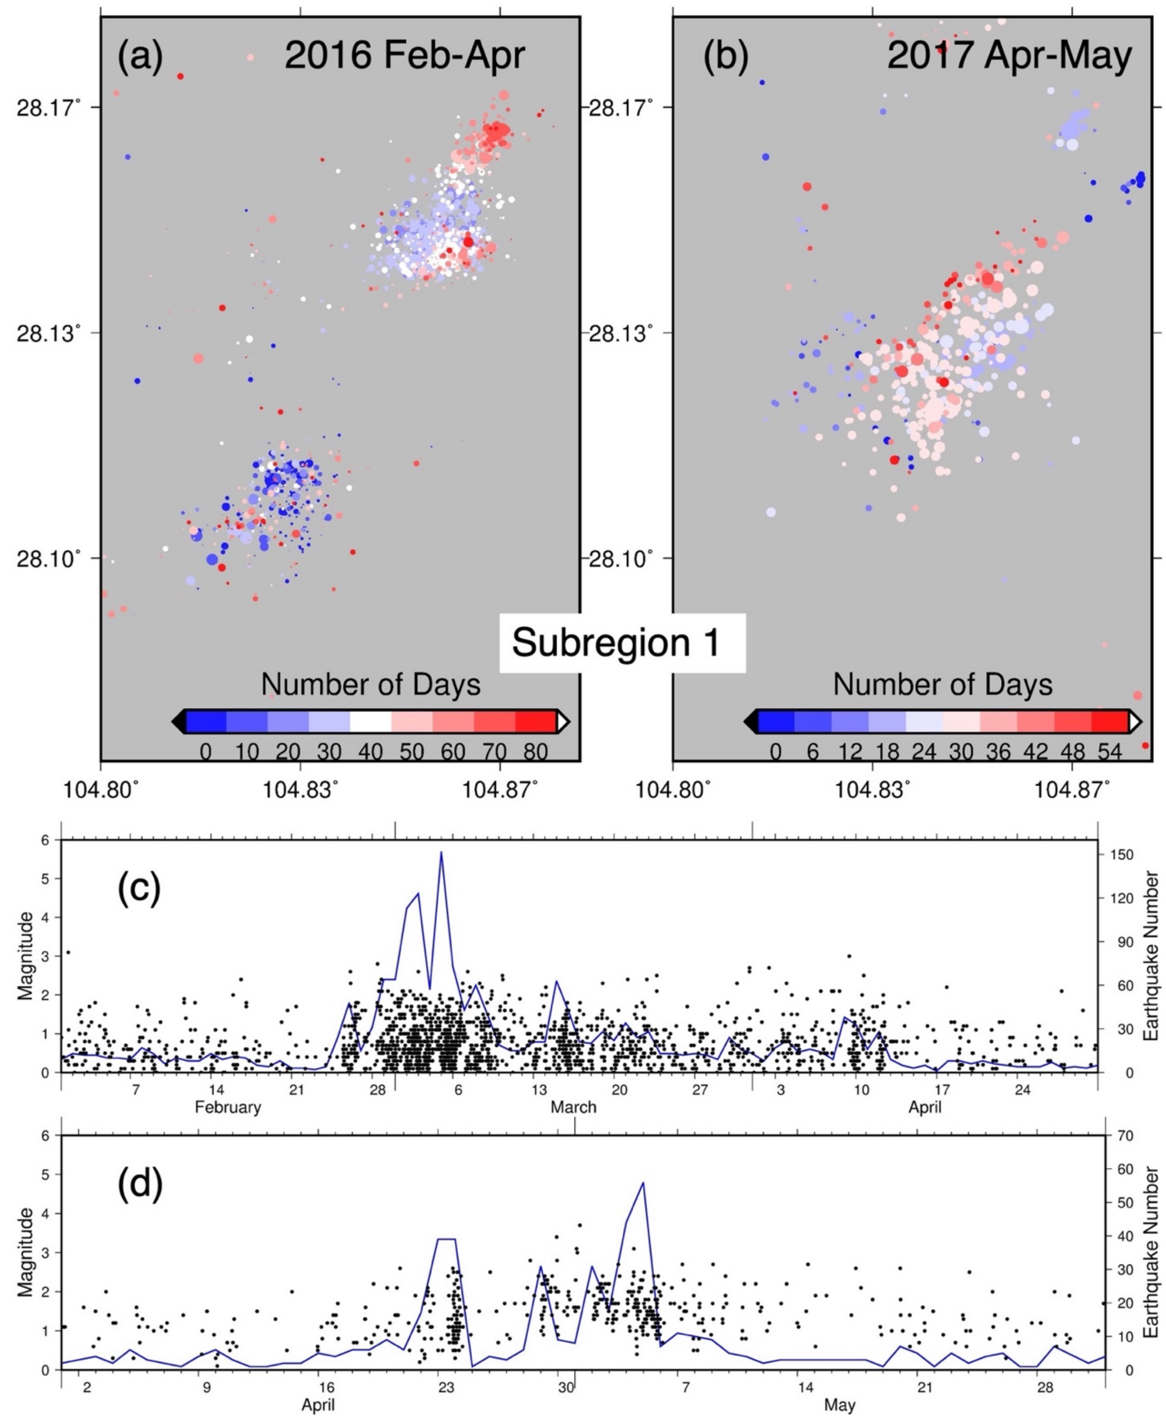

Figure S4. Seismicity in subregion 1 of Changning region during abnormal periods. (a) and (b) show the seismicity in subregion 1 from Feb to Apr 2016 and from Apr to May

2017, respectively. (c) and (d) are the corresponding time series of (a) and (b), showing the earthquake magnitudes and daily earthquake numbers in subregion 1.

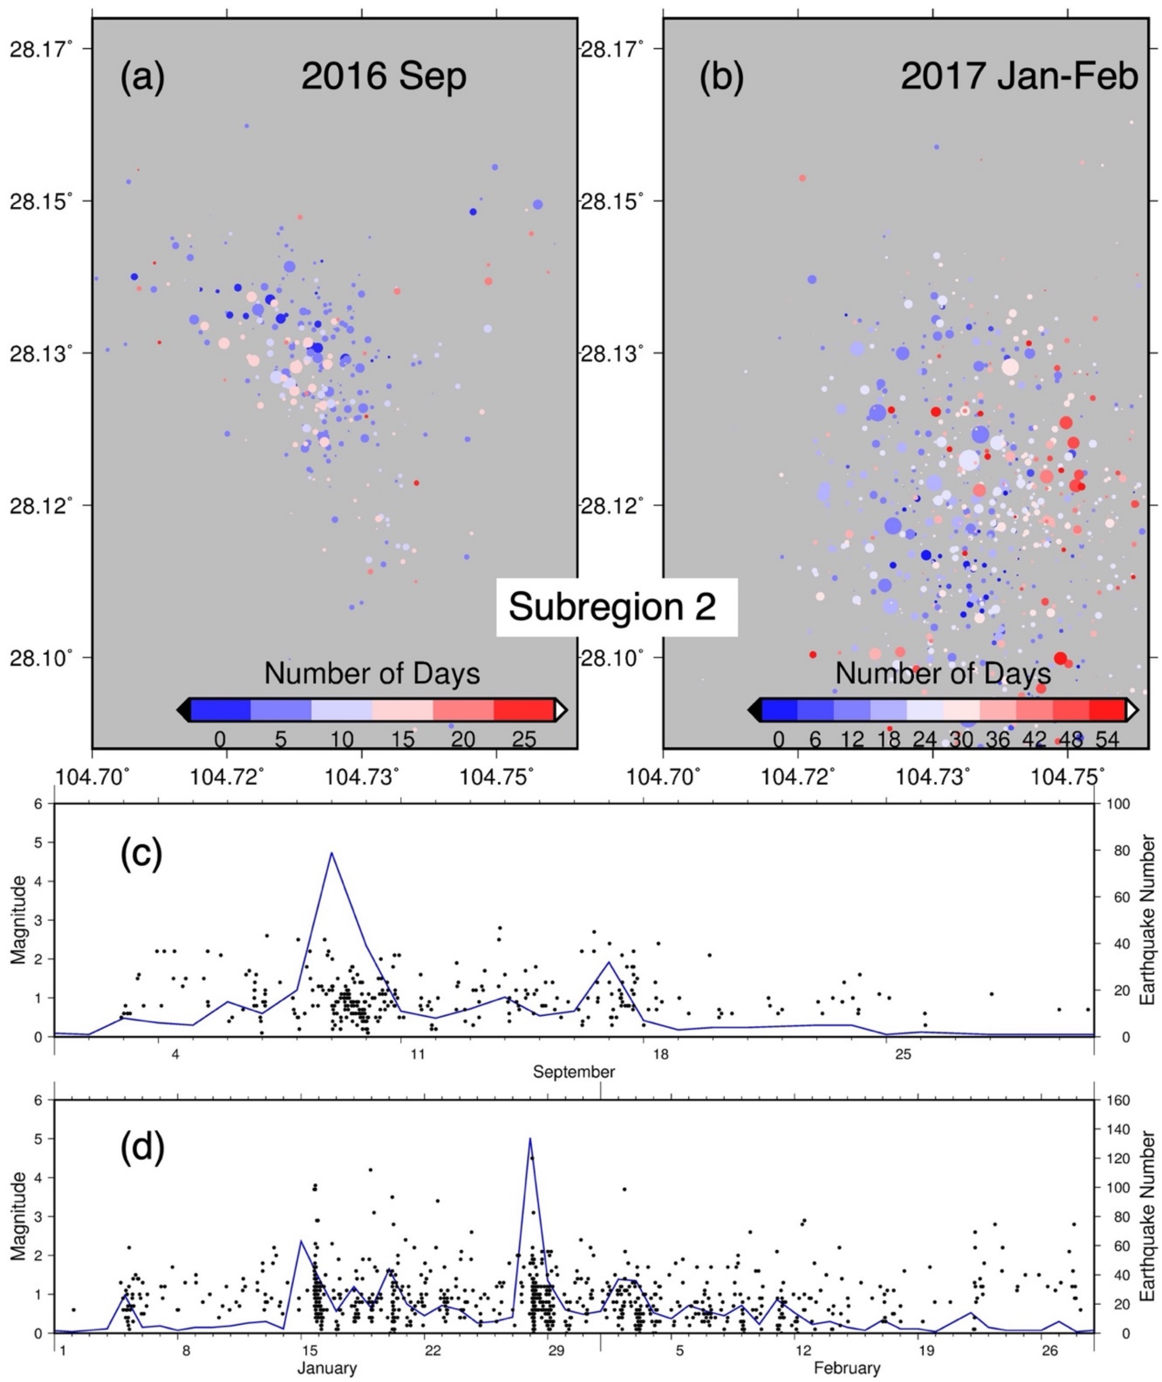

Figure S5. Seismicity in subregion 2 of Changning region during abnormal periods. (a) and (b) show the seismicity in subregion 1 in Sep 2016 and from Jan to Feb 2017, respectively. (c) and (d) are the corresponding time series of (a) and (b), showing the earthquake magnitudes and daily earthquake numbers in subregion 2.

Weight decay	0	0.001	0.003	***0.005***	0.01	0.03	0.05
Final validation loss	0.22	0.215	0.212	***0.161***	0.166	0.199	0.26

Table S1. Final validation losses using different L2 regularization factors (weight decay) for ResNet-18.

Learning rate	0.001	0.003	***0.005***	0.007	0.01
Final validation loss	0.161	0.152	***0.128***	0.167	0.228

Table S2. Final validation losses using different learning rates for ResNet-18.

Weight decay	0	0.001	0.003	***0.005***	0.01	0.03	0.05
Final validation loss	0.195	0.187	0.137	***0.133***	0.174	0.272	0.327

Table S3. Final validation losses using different L2 regularization factors (weight decay) for ResNet-34.

Learning rate	0.001	***0.003***	0.005	0.007	0.01
Final validation loss	0.175	***0.133***	0.162	0.139	0.161

Table S4. Final validation losses using different learning rates for ResNet-34.

Weight decay	0	0.001	0.003	***0.005***	0.01	0.03	0.05
Final validation loss	0.232	0.201	0.204	***0.145***	0.154	0.277	0.372

Table S5. Final validation losses using different L2 regularization factors (weight decay) for ResNet-101.

Learning rate	0.0005	***0.0007***	0.001	0.003	0.005	0.007
Final validation loss	0.426	***0.145***	0.149	0.153	0.182	0.193

Table S6. Final validation losses using different learning rates for ResNet-101.

Movie S1. A movie shows the artificial distribution maps where earthquake number increase gradually.

https://www.youtube.com/watch?v=DZDhHBKZdtY

Movie S2. Similar to Movie S1 but in Weiyuan, from Sep 2018 to Aug 2020.

https://www.youtube.com/watch?v=-MjzMq9Om_Y

Movie S3. A movie shows the spatial-temporal migration of anomalous seismicity in Changning, from 2016 to 2019.

https://www.youtube.com/watch?v=0JcnT08F-2g&t=1s