
Generalized Seismic Phase Detection with Deep Learning

Zachary E. Ross*, Men-Andrin Meier, Egill Hauksson, Thomas H. Heaton
Seismological Laboratory
California Institute of Technology
Pasadena, CA 91125

Abstract

To optimally monitor earthquake-generating processes, seismologists have sought to lower detection sensitivities ever since instrumental seismic networks were started about a century ago. Recently, it has become possible to search continuous waveform archives for replicas of previously recorded events (template matching), which has led to at least an order of magnitude increase in the number of detected earthquakes and greatly sharpened our view of geological structures. Earthquake catalogs produced in this fashion, however, are heavily biased in that they are completely blind to events for which no templates are available, such as in previously quiet regions or for very large magnitude events. Here we show that with deep learning we can overcome such biases without sacrificing detection sensitivity. We trained a convolutional neural network (ConvNet) on the vast hand-labeled data archives of the Southern California Seismic Network to detect seismic body wave phases. We show that the ConvNet is extremely sensitive and robust in detecting phases, even when masked by high background noise, and when the ConvNet is applied to new data that is not represented in the training set (in particular, very large magnitude events). This generalized phase detection (GPD) framework will significantly improve earthquake monitoring and catalogs, which form the underlying basis for a wide range of basic and applied seismological research.

1 Introduction

Ideally, seismologists would be able to detect every single earthquake that occurs to optimally monitor seismogenic processes, such as fluid migration, elastic strain build-up and accelerating crustal deformation episodes. In practice, however, the ability to detect earthquakes is strongly limited by the fact that the vast majority of earthquakes are very small [6], and by the perpetual presence of a wide range of nuisance signals that are not caused by local earthquakes.

Over the past decades, improvements in seismometer design, installation and network density have led to dramatic improvements in our capabilities for detecting small earthquakes. Methodological advances in earthquake detection and characterization continue to improve the resolution with which we can observe the seismically active parts of the crust. The magnitude of completeness—the magnitude above which all events have been detected by a network—for southern California now is $\sim M1.8$ in most areas [9].

Among the most successful earthquake detection algorithms developed to-date are those that exploit the similarity of earthquake waveforms between similarly located events. The most widely used version of these methods is template matching [19, 24, 23, 22]. By correlating the seismic waveforms of cataloged individual events against continuous waveform data, it is often possible to detect an order of magnitude more events than with routine methods [1]. These similarity-based methods typically have additional stringent detection requirements, such as forcing the detected events to have the same

*Correspondence to: zross@gps.caltech.edu

move-out pattern as the template event across the network. While this makes the detections reliable, it also makes the similarity-based methods blind to what they have not seen before, such as events in previously inactive regions, or large magnitude earthquakes. As a consequence, catalogs created with such methods are inherently incomplete in that they contain only the subset of events that meet the rather specific detection criteria.

The task of recognizing seismic phases in a waveform time-series is very similar to that of recognizing objects in 2D images. A 3-component seismogram can be thought of as a 1D image, with the three components being equivalent to the three RGB color channels. The stunning recent advances in the field of computer vision therefore have important potential in seismological applications. Deep learning algorithms, which use hierarchically-organized non-linear mapping functions for tasks such as classification and regressions [16], have been shown to be powerful tools for object recognition because of their ability to develop robust generalized representations of extremely large datasets [8, 14]. They have become the state-of-the-art in various machine learning domains, including handwriting recognition, language processing and object categorization [15]. Rather than using explicit template objects to search for in an image (e.g. template images of cats) these algorithms are able to detect the general characteristics that individual realizations of that object class share (e.g. the furry texture). With such non-explicit search they are able to reliably detect objects without ever having seen an exact or even similar object template.

In seismology, deep learning has recently been used to detect and roughly locate earthquakes in Oklahoma [20], and to determine P-wave arrival times and first-motion polarities [21]. Here we develop a convolutional neural network that can detect and classify seismic body wave phases over a broad range of circumstances. We use the vast hand-labeled data archives of the Southern California Seismic Network. We show that this method for generalized phase detection (GPD) is capable of reliably detecting P- and S-waves of very small to very large earthquakes without the need for explicit waveform templates. The ability to generalize entire waveform and metadata archives enables the method to successfully operate on data that was not covered by the training data, including data from entirely different tectonic regimes, and of earthquakes much larger than the largest events in the training set. The developed procedure will allow us to compile fundamentally improved seismicity catalogs with detection sensitivities comparable to template matching catalogs, but without their inherent biases. Furthermore, it has the potential to make earthquake early warning systems more reliable.

2 Methods and Results

2.1 Neural Networks and Deep Learning

In the most general sense, neural networks map a set of input values (e.g. a raster image or a time series) into a set of output values (e.g. the probability that the input image contains a cat, or that a seismogram contains a P-wave). The network architecture, which parameterizes the non-linear mapping function, has a large number of parameters (typically $>1e6$ for deep networks) that are empirically optimized with large amounts of training data. The optimization is performed such that the network predictions (i.e. the mapping output) are as accurate as possible across a large data set.

For standard fully-connected neural networks (FCNN), the mapping function consists of a large number of addition and multiplication operations, as well as simple non-linear activation functions. These operations are organized into discrete layers where the output of each layer is sequentially used as input to the next. The first layer of a FCNN acts on "features", which are attributes of the raw data, such as peak amplitudes and frequency measures. The choice of the features is usually somewhat arbitrary and difficult to optimize. Convolutional neural networks (ConvNets), on the other hand, typically combine a FCNN with a learnable and hierarchical feature extraction system (Fig. 1) that is trained to extract the relevant information directly from the raw data.

The feature extraction system itself consists of a series of layers in which the input data is sequentially i) convolved with a set of digital filters, ii) decimated, and iii) "activated". The digital filters parameterize the general object characteristics that the ConvNet will learn to recognize. Whenever this characteristic is present at some point in an input image, the convolution will yield a significant output value. The decimation, or "pooling", down-samples the image and causes the ConvNet to evaluate the image at different length scales. This enables the network to recognize the object regardless of its size, i.e. whether it covers the entire image, or whether it is just a small detail in the




Figure 1: (a) Illustration of a convolutional network for generalized phase detection and (b) application example. By applying a series of filtering and decimation operations, features are automatically extracted and used to classify data windows. (b) Example usage of GPD to continuous data. Probability time series for P- (red) and S-waves (blue) are used as characteristic functions for phase detection. Detection windows with probability greater than 0.98 are colored by the respective phase type.

background. Finally, an activation function is used on the filtered and pooled input data that forces the output values to be positive-only, and that introduces a non-linearity in the mapping function. After several sequential convolution, pooling and activation steps, the resulting output of the last layer is concatenated to a vector and is used as input to a FCNN in order to perform regression or classification tasks. During the training process, the parameters of the entire network are learned directly from the data to maximize the network’s performance. This includes the coefficients of the digital filters and the coefficients of the linear functions of the FCNN’s neurons.

Due to the convolutional nature of the filtering operations and the regular decimation, a ConvNet can detect a target object (e.g. a cat or a P-wave) irrespective of where it is located in an image or seismogram, or the size within the image. ConvNets are therefore ideal for seismic phase detection since the phases may occur anywhere in a seismogram, and can vary by orders of magnitudes in terms of duration and amplitude. The major limitation is that millions of labeled data samples are necessary to train the networks, which frequently have millions of free parameters.

2.2 A ConvNet for Generalized Seismic Phase Detection

We designed a convnet to scan through continuous seismic data, and for each 4 sec data window, classify the dominant category of the window as either a P-wave, S-wave, or noise (Figure 1b). The convnet was trained using millions of hand-labeled phases and their arrival times from analysts at the Southern California Seismic Network (see methods). The model architecture that we used consists of 6 layers (Table S1), which includes 4 convolution layers and 2 fully-connected layers. Rectified linear units were used as the activation function on each layer [18], and batch normalization was applied to each layer [10]. We explored many other possible models and found that overall, excellent performance was possible with both shallower and deeper models; thus, the method is generally insensitive to the specific architecture used. However, the network as described here performed the best of all those tested.

A total of 4.5 million 3-component seismic records were used for training and validation of the generalized phase detection framework. The data first were detrended and high-pass filtered above 2 Hz to remove microseismic noise, and all data were resampled at 100 Hz. Strong-motion records were integrated to velocity. These 3-component records consist of 1.5 million P-wave seismograms,

1.5 million S-wave seismograms, and 1.5 million noise windows, with each being exactly 4 s (400 samples) in duration. The magnitude range of the data was $-0.81 < M < 5.7$, while only records with epicentral distances less than 100 km were used. P- and S-wave windows were centered on the respective analyst pick, while noise windows were defined starting 5 s before each P-wave pick. These windows form the set of features used as input to the convnet, and examples of each class can be seen in Figure S2. The even distribution of records between each class ensures that the training process is not biased towards any one class. Then, each 3-component record was normalized by the absolute maximum amplitude observed on any of the three components. No other pre-processing was performed on the datasets.

To train the model, we first randomly split the seismograms into a training set (75%) and validation set (25%). Then, the model was trained using a cross-entropy loss function with the ADAM optimization algorithm [12], in mini-batches of 480 records with 3 NVIDIA GTX 1060 graphical processing units. We programmed the training process to terminate when the validation loss had not decreased in more than 5 epochs (full iterations through the dataset), and the epoch with the lowest loss value on the validation set was selected as the best model (Figure S1). For this study, the 3rd epoch had the highest prediction accuracy on the validation set (99%), with subsequent epochs resulting in less accurate predictions due to model overfitting. This forms the final model used throughout the paper.

We found that even though a high-pass filter at 2 Hz was used to train the model, that the method still worked over a broad range of frequency values above and below 2 Hz. For our subsequent testing with seismograms of the 2016 M_w 7.0 Kumamoto, Japan earthquake, the method reliably detected P- and S-waves even when using acceleration traces (rather than velocity), and a high-pass filter at 0.1 Hz. Thus, it is not a strict requirement that the same filter used to train the data need be applied to new data for calculations.

2.3 Classification Performance on the Validation Set

First, we demonstrate the classification performance of the trained model on an independent validation set of 1.1 million seismograms split evenly between the three classes. For each 3-component seismogram, the model outputs three probabilities corresponding to the likelihood of each respective class (P, S or noise). A seismic phase detection is declared if the probability for either P or S is above a threshold level. If both predictions are below the threshold, the waveform is assigned to the noise class. Figure 2 shows the performance of the convnet, in terms of precision and recall, for a range of probability thresholds. While precision gives the fraction of correct declarations (true positives) relative to all declarations that have been made (true positives + false positives), recall gives the fraction of correct declarations relative to all declarations that should have been made (true positives + false negatives).

Almost irrespective of the detection threshold, the precision is very high (>99%) for both phases, suggesting that the network very rarely labels noise as seismic phases, or confuses P- and S-phases. Recall is somewhat lower, between 96% and 99% for most threshold choices, indicating that seismic phases are sometimes misclassified as noise. Remarkably, the performance is almost as good for S-waves as it is for P-waves.

These tradeoff curves illustrate that the probability threshold is a tunable parameter that can be exploited to serve a specific application. For earthquake catalog generation, for instance, one may wish to minimize false detections, while in an earthquake early warning (EEW) it may be more important to not have any missed detections. Furthermore, if one can further reduce false positives through other means (e.g. by requiring multiple stations to declare a detection), lower probability thresholds can be chosen in order to gain higher recall values (fewer false negatives). Figure S3 demonstrates the performance of the method as a function of SNR.

2.4 Application to the 2016 Bombay Beach, California Swarm

We next demonstrate the applicability of the GPD framework to detecting phases in a streaming mode during the 2016 Bombay Beach, California swarm [7]. The trained network is used to classify 4 s windows of 3-component data (every 10 samples; 97.5% overlap) over the first 24 h of the swarm. For each window, if the network predicts a P- or S-wave as the most likely category, and the forecast probability is greater than 0.98, the time point of the center of the window is assigned a detection




Figure 2: Precision-recall tradeoff curve for different declaration probability thresholds and confusion matrix with definitions. If P- or S-probabilities exceed the threshold probability (color) the waveform is assigned to the respective class. If neither probability exceeds the threshold the waveform is assigned to the noise class. For low probability thresholds more cases of false positives occur, while with higher thresholds more false negative cases occur. The confusion matrix shows the possible combinations of classifications by human analysts ('H') and the convnet algorithm ('A').

with the respective phase label. When multiple consecutive values are above 0.98, the time point of the maximum value is taken as the detection time.

Due to the intensity of seismic activity during the swarm, we focus first on the initial 10 minute period following the onset. Figure 3 contains the 3-component velocity data at station BOM, which is located only 7 km away from the swarm itself. Here, P- and S-waves detected are colored red and blue, respectively. The forecast probability for these classes is shown as a function of time in the same respective colors. There are many earthquakes occurring within this short time window, and to more clearly illustrate the detection capabilities, we zoom in further on a segment that is just over two minutes long. It can be seen that during this short time window alone, 13 earthquakes are detected, and both phases are identified in 12 of them. The detections show higher P-wave amplitudes on the vertical component, and higher S-wave amplitudes on the horizontal components. S-wave detections are regularly preceded by P-wave detections, and S-P times are consistent with the largest events in the swarm, all of which suggest that the detections are real. The events are, on average, about 10 seconds apart; however, 6 of the events are significantly overlapping in time. We note that while the window duration used for phase detection here is 4 s long, the close proximity of station BOM to the swarm results in S-P times of about 1 s; thus, the method is able to robustly identify individual phases from very small segments of data, even when the phases overlap. An STA/LTA algorithm [2, 1] would trigger on few, if any, of these signals because the long-term average in the denominator does not reach small values during periods of elevated activity. Even a human analyst would likely have missed many of these events. The detection results are not unique to this 10 minute window; Figure 4 demonstrates nearly 1000 events were detected over the first 12 hours of the swarm alone, which is ~10 times as many as listed in the SCSN catalog; however it should also be mentioned that most of the SCSN events were included in the training data. Notably, the onset time of the swarm is measurable with extraordinary precision. Thus, GPD shows exceptional sensitivity to detecting P- and S-waves, without requiring specific templates to be correlated. We note that if the threshold of 0.98 is lowered, more events will be detected, but this will also result in additional false positives (Fig. 2).

2.5 Application to the 2016 Mw 7.0 Kumamoto, Japan Earthquake

Detection of large earthquakes is a critical task for seismic networks and earthquake early warning (EEW) systems, but methods that exploit the waveform similarity of previous events to detect new




Figure 3: Application of GPD to the 2016 Bombay Beach swarm. P-wave (red) and S-wave (blue) detections are colored for all samples of any window in which a class probability exceeds 0.98. The probability time series shows numerous high-probability detections. Inset shows a close up of a time window ~140s long with 13 earthquakes detected. P-detections occur consistently before S-detections. While P-waves tend to have higher vertical amplitudes, S-waves are stronger on the horizontal components.

ones will be incapable of detecting these damaging events. We demonstrate the viability of GPD for detecting large earthquakes reliably by applying it to the 2016 Mw 7.0 Kumamoto earthquake [3]. Figure 5 contains horizontal seismograms for all strong-motion stations within 100 km of the hypocenter. For every station, the P-wave of the mainshock is detected correctly, while for nearly every station, the S-wave is also detected. It is important to note that the training data were exclusively recorded in southern California, but here the model is detecting phases in Japan on completely different instruments. These results are even more remarkable given that the network was only trained on earthquakes with $M < 5$, and thus the model has not seen a large magnitude earthquake during the training process. These results are possible because within the short time windows used for prediction, the earthquake has not grown large enough in size yet to appear significantly different than a smaller event [17]. Therefore, GPD identifies coherent arrivals, rather than entire waveforms, which enables detection of even large events.




Figure 4: Example of GPD applied to the first 12 hours of the 2016 Bombay Beach sequence. The onset time of the swarm is sharply resolved. The total number of events detected is ~ 10 times as many as listed in the SCSN catalog. The swarm onset is distinctly visible, which often is not the case in routine catalogs, posing a common problem for seismic monitoring operators. The high amplitude signals before the swarm are nuisance signals from nearby vehicles that do not trigger the GPD algorithm.

3 Discussion

Detection of seismic phases is the first step in producing a seismicity catalog, and subsequent measurements derived from these phases form the basis for nearly all forms of seismological analyses. We showed that by using deep learning to develop generalized representations of millions of P-waves, S-waves, and noise records, it is possible to significantly improve this initial phase detection process. The method as presented here has detection sensitivity that rivals template matching but is not associated with any of the biases inherent to the similarity-based detection methods. Although it currently is not possible to understand how the convnet operates in significant detail, it is likely that the network learns to recognize attributes of the 3D particle motion both before and after the phase arrival, as well as some frequency-based information. The network seems to be primarily detecting onsets of phases, which is likely because the data windows were centered on the phase arrival time; this is probably the reason that the method can work on large magnitude earthquakes, since the onsets look similar to those events. If multiple phases are present within a data window, it tends to assign class labels based on whatever signal type is closest to the center of the window (e.g. Fig. 3). Rather than trying to formalize this operation into a parametric set of rules, the network learns the best way to discriminate between the three signal classes based on the data directly.

While the model as trained demonstrates outstanding portability, for some applications, the model would likely have to either be retrained using larger distance recordings, or adapted to incorporate additional signal classes. For the latter, since the noise windows were taken randomly from pre-event data, only a small fraction will include high amplitude transient signals like vehicles, calibration pulses, or noise spikes. Some of these signal types could therefore lead to occasional spurious




Figure 5: Application of GPD to the 2016 Kumamoto earthquake. P-waves (red) and S-waves (blue) are detected at nearly all stations over the epicentral distance range 3-100 km. Diamonds indicate theoretical arrival times. This demonstrates both the viability of GPD to detect large magnitude earthquakes, but also to work in regions and magnitude ranges that are not included in the training data.

classifications as P- or S-waves since they are poorly represented in the training data. GPD provides a natural framework for including these signal types as additional classes once a library of them is produced. This will help to both improve the P- and S-wave detection sensitivity while also improving the classification accuracy, since the decision boundaries between each signal type will become sharper. For regional distances and beyond, the presence of non-direct phases can potentially make detections more complicated and it may be desirable to eventually include these phase types explicitly. The model presented here is specifically trained to detect direct P- and S- phases at local distances.

An important potential application of GPD is in EEW, to improve discrimination between genuine earthquake signals and other spurious signals that could lead to false detections. Many EEW systems currently trigger based on large STA/LTA signals, which serve as an anchor for the real-time monitoring process [5]. GPD can be applied in a streaming mode for each packet received in an analogous manner to STA/LTA. Every time a P-wave is detected, the system could trigger other event characterization schemes (e.g. magnitude estimation [11]), and ultimately issue an alert if necessary. By looking for signals that are reminiscent of P-waves, rather than arbitrarily polarized large amplitude signals, GPD can provide a far more stable base to EEW systems. Furthermore, the additional information of what phase-type a detection corresponds to (P vs. S) is valuable information in an EEW context, since it can be used for excluding S-wave observations from magnitude-estimation schemes that have been optimized only for P-waves [4] and to assess if an ongoing rupture is still growing [13].

The next big question is how the GPD results can be optimally harnessed to generate substantially improved earthquake catalogs. The impressive classification performance and the outstanding

sensitivity of the method as applied to individual stations suggests that once information from multiple stations is combined, event detection via phase association should be robust. Furthermore, the fact that GPD provides reliable phase labels (P vs. S) can be of great value for unscrambling the seismic signals from simultaneous and potentially overlapping events, such as during swarms.

In summary, the GPD framework may sharpen our ability to characterize how earthquakes evolve in space and time over all scales. It may help us to better characterize and understand seismogenic processes, from tiny micro-earthquakes whose signals barely exceed the ambient noise level, all the way to destructive large magnitude event sequences.

4 Data and Resources

In this study, we used seismograms for 273,882 earthquakes ($-0.81 < M < 5.7$) recorded by the Southern California Seismic Network (SCSN) at 692 broadband and short-period 3-component stations from 2000-2017 [25]. The waveform data (last accessed January 2018) were associated with 1.5 million P-wave picks and 1.5 million S-wave picks that were manually determined by SCSN analysts. We also used 24 h continuous data at station CI BOM on 2016-09-26, the first day of the 2016 Bombay Beach, California swarm. Seismicity during this sequence was taken from the SCSN regional catalog. Finally, we used KiK-net and K-net accelerograms of the 2016 Mw 7.0 Kumamoto earthquake at stations within 100 km distance of the hypocenter (last accessed January 2018). The trained model and an accompanying dataset will be available through the Southern California Earthquake Data Center (scedc.caltech.edu).

5 Acknowledgements

This research was supported by grants from the Gordon and Betty Moore Foundation, the Swiss National Science Foundation and the NSF Geoinformatics program. We have used waveforms and meta-data data from Japanese networks and the Caltech/USGS Southern California Seismic Network (SCSN); doi: 10.7914/SN/CI; stored at the Southern California Earthquake Data Center. doi:10.7909/C3WD3xH1. We used Tensorflow to train the convolutional network.

References

- [1] R. Allen. Automatic phase pickers: Their present use and future prospects. *Bulletin of the Seismological Society of America*, 72(6B):S225–S242, Dec. 1982.
- [2] R. V. Allen. Automatic earthquake recognition and timing from single traces. *Bulletin of the Seismological Society of America*, 68(5):1521–1532, Oct. 1978.
- [3] K. Asano and T. Iwata. Source rupture processes of the foreshock and mainshock in the 2016 Kumamoto earthquake sequence estimated from the kinematic waveform inversion of strong motion data. *Earth, Planets and Space*, 68(1):147, Dec. 2016.
- [4] M. Böse, E. Hauksson, K. Solanki, H. Kanamori, and T. H. Heaton. Real-time testing of the on-site warning algorithm in southern California and its performance during the July 29 2008 Mw5.4 Chino Hills earthquake. *Geophysical Research Letters*, 36(5):1–5, Mar. 2009.
- [5] E. S. Cochran, M. D. Kohler, D. D. Given, S. Guiwits, J. Andrews, M.-A. Meier, M. Ahmad, I. Henson, R. Hartog, and D. Smith. Earthquake Early Warning ShakeAlert System: Testing and Certification Platform. *Seismological Research Letters*, 89(1):108–117, Jan. 2018.
- [6] B. Gutenberg and C. F. Richter. *Seismicity of the Earth*. Hafner Publishing Company, New York, 1954.
- [7] E. Hauksson, M. Meier, Z. E. Ross, and L. M. Jones. Evolution of seismicity near the southernmost terminus of the San Andreas Fault: Implications of recent earthquake clusters for earthquake risk in southern California. *Geophysical Research Letters*, 44(3):1293–1301, 2017.
- [8] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. *arXiv:1512.03385 [cs]*, Dec. 2015. arXiv: 1512.03385.
- [9] K. Hutton, J. Woessner, and E. Hauksson. Earthquake Monitoring in Southern California for Seventy-Seven Years (1932–2008). *Bulletin of the Seismological Society of America*, 100(2):423–446, Apr. 2010.

- [10] S. Ioffe and C. Szegedy. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *PMLR*, pages 448–456, June 2015.
- [11] H. Kanamori. Real-Time Seismology and Earthquake Damage Mitigation. *Annual Review of Earth and Planetary Sciences*, 33(1):195–214, 2005.
- [12] D. P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. *arXiv:1412.6980 [cs]*, Dec. 2014. arXiv: 1412.6980.
- [13] Y. Kodera. Real-Time Detection of Rupture Development: Earthquake Early Warning Using P Waves From Growing Ruptures. *Geophysical Research Letters*, 45(1):2017GL076118, Jan. 2018.
- [14] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc., 2012.
- [15] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, May 2015.
- [16] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, Nov. 1998.
- [17] M.-A. Meier, T. Heaton, and J. Clinton. Evidence for universal earthquake rupture initiation behavior. *Geophysical Research Letters*, 43(15):7991–7996, Aug. 2016.
- [18] V. Nair and G. E. Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pages 807–814, 2010.
- [19] Z. Peng and P. Zhao. Migration of early aftershocks following the 2004 Parkfield earthquake. *Nature Geoscience*, 2(12):877–881, Dec. 2009.
- [20] T. Perol, M. Gharbi, and M. Denolle. Convolutional neural network for earthquake detection and location. *Science Advances*, 4(2):e1700578, Feb. 2018.
- [21] Z. E. Ross, M.-A. Meier, and E. Hauksson. P-wave arrival picking and first-motion polarity determination with deep learning. *J. Geophys. Res.-Solid Earth*, submitted, 2018. arXiv:1804.08804 [physics.geo-ph].
- [22] Z. E. Ross, C. Rollins, E. S. Cochran, E. Hauksson, J.-P. Avouac, and Y. Ben-Zion. Aftershocks driven by afterslip and fluid pressure sweeping through a fault-fracture mesh. *Geophysical Research Letters*, page 2017GL074634, Jan. 2017.
- [23] D. R. Shelly, W. L. Ellsworth, and D. P. Hill. Fluid-faulting evolution in high definition: Connecting fault structure and frequency-magnitude variations during the 2014 Long Valley Caldera, California, earthquake swarm. *Journal of Geophysical Research-Solid Earth*, 121(3):1776–1795, Mar. 2016.
- [24] R. J. Skoumal, M. R. Brudzinski, and B. S. Currie. Earthquakes Induced by Hydraulic Fracturing in Poland Township, Ohio. *Bulletin of the Seismological Society of America*, 105(1):189–197, Feb. 2015.
- [25] Southern California Earthquake Data Center. Scsn. *California Institute of Technology, Dataset*, 2013.