Applying Machine Learning to Crowd-sourced Data from Earthquake Detective

Omkar Ranadive[1], Vivian Tang[1], Kevin Chao[1], and Suzan van der Lee[1]

[1]Northwestern University

November 22, 2022

Abstract

Dynamically triggered earthquakes and tremor generate weak seismic signals whose detection, identification, and authentication call for a laborious analysis. Citizen science project Earthquake Detective leverages the eyes and ears of volunteers to detect and classify weak signals in seismograms from potentially dynamically triggered (PDT) events. Here, we present the Earthquake Detective data set - A crowd-sourced set of labels on PDT earthquakes and tremor. We apply Machine Learning to classify these PDT seismic events and explore the challenges faced in segregating such weak signals. The algorithm confirms that machine learning can detect signals from small earthquakes, and newly demonstrates that this specific algorithm can also detect signals from PDT tremor. The data set and code are available online.

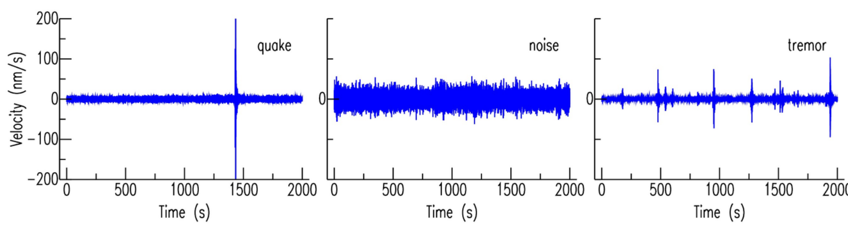

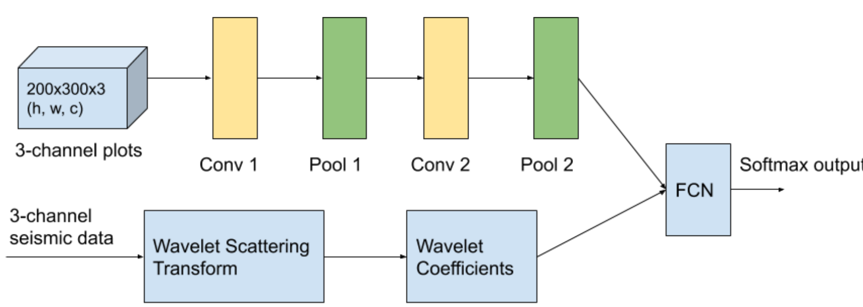

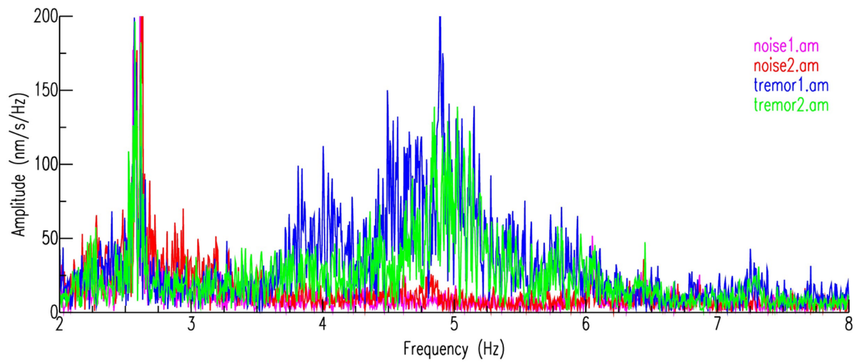
noise1.am
noise2.am
tremor1.am
tremor2.am
Amplitude (nm/s/Hz)
Frequency (Hz)

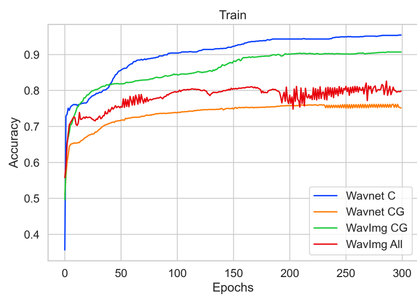
Train
Accuracy
Epochs
Wavnet C
Wavnet CG
WavImg CG
WavImg All

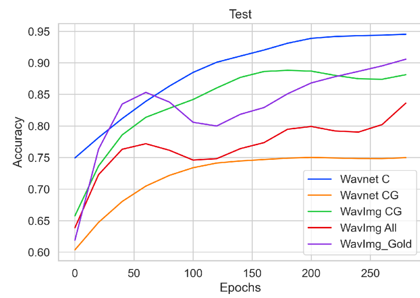
Test
Accuracy
Epochs
Wavnet C
Wavnet CG
WavImg CG
WavImg All
WavImg_Gold

Applying Machine Learning to Crowd-sourced Data from Earthquake Detective

Omkar Ranadive[1], Vivian Tang[2*], Kevin Chao[3], Suzan van der Lee[2]

[1]Department of Computer Science, Northwestern University
[2]Department of Earth and Planetary Sciences, Northwestern University
[3]Data Science Initiative/NICO, Northwestern University

Key Points:

- The data set is a crowd-sourced set of labels of weak seismic signals from potentially dynamically triggered seismic events.
- We successfully applied Machine Learning to the detection of seismic signals from dynamically triggered tremor.
- Our algorithm combines waveform-based learning with image-based learning to detect these weak seismic signals with high accuracy.

*Now is at Department of Civil and Environmental Engineering, University of Illinois Urbana-Champaign

Corresponding author: Omkar Ranadive, omkar.ranadive@u.northwestern.edu

Abstract

Dynamically triggered earthquakes and tremor generate weak seismic signals whose detection, identification, and authentication call for a laborious analysis. Citizen science project Earthquake Detective leverages the eyes and ears of volunteers to detect and classify weak signals in seismograms from potentially dynamically triggered (PDT) events. Here, we present the Earthquake Detective data set - A crowd-sourced set of labels on PDT earthquakes and tremor. We apply Machine Learning to classify these PDT seismic events and explore the challenges faced in segregating such weak signals. The algorithm confirms that machine learning can detect signals from small earthquakes, and newly demonstrates that this specific algorithm can also detect signals from PDT tremor. The data set and code are available online.

Plain Language Summary

Dynamically triggered earthquakes and tremor are seismic events that were triggered by surface waves from large magnitude and often distant earthquakes. These events generate weaker ground motion than earthquakes in agency catalogs and are hard to detect. Project Earthquake Detective has invited volunteers to help detect these events by viewing and listening to seismograms (recordings of ground motion). We applied a machine-learning algorithm to the same data. The algorithm was able to detect signals from small earthquakes and also for the first time, from triggered tremor.

1 Introduction

Over the past five years, Machine Learning (ML) has progressively grown to be a popular tool in geophysical analyses, as evidenced by three dozen papers published during that time in Geophysical Research Letters (GRL) and containing the term in their title or keywords. Much of this research demonstrates the impressive efficiencies that can be achieved by applying ML to tasks that are overwhelming for researchers from a data volume or dimensionality perspective while relatively straightforward in complexity or signal strength (Chmiel et al. (2021); Zhao et al. (2019); Lee et al. (2020); Mousavi and Beroza (2020); Z. Li et al. (2018)). Few papers demonstrate ML's success in recognizing of low-amplitude signals in seismology (Rouet-Leduc et al., 2017). Here we leverage a crowd-sourced data set of weak seismic signals classified by citizen scientists (Tang, Rösler, et al., 2020) in combination with a data set analyzed and labeled by experts (the authors) to establish a baseline for detecting different types of weak seismic signals. The ultimate goal of this work is to incorporate the ML algorithm in winnowing the data stream presented to citizen scientists and experts.

One class of weak seismic signals are seismograms from earthquakes with magnitudes below the magnitude of completion for typical earthquake catalogs. Yet these low-magnitude earthquakes belong to the same Gutenberg-Richter distribution (Gutenberg & Richter, 1954) as widely recorded earthquakes, including those that cause injuries, damage, and worse. Gutenberg and Richter's law (1954) states that for every magnitude (M) 8 earthquake that occurs, about 1 million M2 earthquakes occur. Therefore, if we can detect the abundant low-magnitude earthquakes of the kind whose signals are often buried in background seismic noise, their analyses could provide insights into the occurrence, distribution, and physics of these *and* the much sparser damaging earthquakes.

Likewise, more recently investigated low-frequency earthquakes exist that represent slower slip between two blocks of rock than that during classical earthquakes, but nevertheless generate weak seismic signals that are often labeled as "tremor" (Obara, 2002), especially when many of such events occur quasi-simultaneously.

Weak signals from tremor and from low-magnitude local earthquakes are both abundant (Rouet-Leduc et al. (2018); Hill and Prejean (2015)) and somewhat under-reported because they have been hard to detect. Past barriers to detection have included a sparse spatial distribution of seismic stations (instrumented with buried seismometers), and current barriers include the weakness and limited bandwidth of these weak seismic signals. These detection challenges have inspired the application of machine-learning algorithms to large sets of seismic waveform data. Several of these ML algorithms have successfully been trained to detect signals from local earthquakes (e.g. Ross et al. (2018); Tang, Seetharaman, et al. (2020)) while others (Liu et al., 2019) used ML to detect tremor signals. Using ML in the detection of signals from seismic activity is a rapidly growing field (e.g. Bergen et al. (2019); Meier et al. (2019); Huang, 2019; W. Li et al. (2018); Riggelsen and Ohrnberger (2014); Ruano et al. (2014); Reynen and Audet (2017); Wiszniowski et al. (2021)).

Weak waveforms from abundant minor seismic events do not only add constraints to estimates of seismic risk, which are based on regional variations in the rate at and mode in which earthquakes occur, but also provide important information on where, when, and how they strike, allowing us to learn about the conditions under which earthquakes nucleate, occur, and interact. Therefore, the more we detect weak signals from minor seismic events, the more we learn about the physics and potential hazards of seismic slip. For example, we can learn about the dynamics of earthquake triggering by first detecting seismic events that occurred simultaneously with transient strain events, then determining the likelihood that these seismic events were triggered by the strain events, followed by examining the conditions under which such triggering does and does not occur. (Tang et al., 2021)

Here, our interests lie in detecting a special sub-class of the multitude of minor seismic events, namely local earthquakes and tremor that could have been triggered by slowly-oscillating large-amplitude seismic surface waves from large-magnitude teleseismic earthquakes. Reporting and learning more about such Potentially Dynamically Triggered (PDT) events extends the spectrum of seismic slip data available for study and adds information about how fast and slow-slip earthquakes might nucleate. Traditional ways for detecting signals form PDT events are 1) seismologists interactively examining seismograms and labeling detections after a range of signal inspections (Gomberg et al., 2008), and 2) seismologists developing and applying an automated detection algorithm to seismic waveform data while controlling the quality of the detections by tweaking the algorithm's parameters and handling outliers separately (e.g. Velasco et al. (2008); Yun et al. (2021)).

In our quest to detect PDT seismic events we face a number of additional challenges:

1. The magnitudes (M) of PDT earthquakes are typically below the M of completeness of earthquake catalogs, hence their signals are weak and often buried in ambient seismic noise signals. Therefore, the database of template waveforms available for such low-M events is small at best. Signals from low-M events are not only lower in amplitude than those from higher-M events but also have narrower bandwidths, diminishing the efficacy of template matching methods.
2. Unlike signals from dynamically triggered earthquakes, signals from dynamically triggered tremor have different waveforms than those from typical tremor on account of the former signals being modulated by the teleseismic surface waves that triggered them (Chao et al., 2012)). Therefore, a database of template signals is not available for training or other purposes, although a catalog has been started (Kano et al., 2018).
3. A signal from a PDT event can arrive at any time during the time window of surface wave passage, which is much longer in duration than the PDT event signal. We have been considering up to 33 minutes of surface wave duration in labeling whether or not at least one local earthquake or tremor signal was recorded.

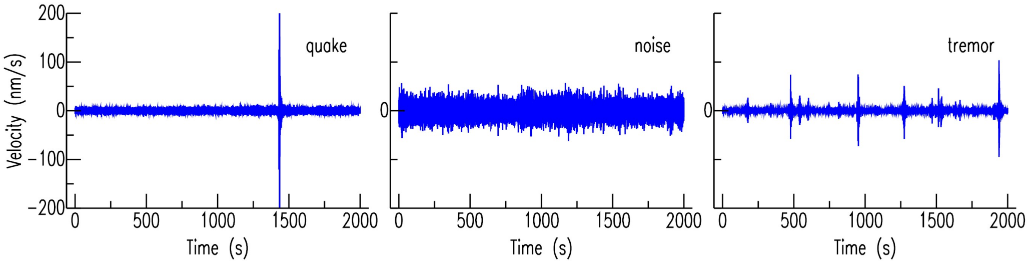

Figure 1. Example data plots shown to users on Earthquake Detective platform

4. Non-stationary noise signals often exceed or are comparable in amplitude and duration to the relatively weak signals of the PDT events we are interested in.
5. Optimal and accurate detection of signals from PDT events requires a multi-scale, multi-band, multi-component interactive analysis that is labor-intensive. The formation of a large, labeled data set for training purposes is hence not straightforward.

2 Earthquake Detective

Earthquake Detective is a crowdsourcing platform where volunteers are shown a seismogram of vertical ground velocity vs. time (Figure 1), along with a sonification of the signal, and are asked to classify the data as Earthquake/Tremor/Noise/None of the above. The platform currently has over 6000 volunteer scientists and over 130k classifications. Tang, Rösler, et al. (2020) present an analysis of how the volunteers and seismologists engage with the data.

All raw input data are time series of recorded ground motion with durations of 2000 s, which are long enough to contain the time window needed by teleseismic surface waves to pass through. The raw data is demeaned, deconvolved with the instrument response to convert digital counts to physical units of ground velocity, band-pass filtered between 2-8 Hz with a 2-pole Butterworth filter, and resampled at 20 samples per second.

3 Approach

3.1 Wavelet Scattering Transform

The wavelet scattering transform decomposes a signal using a family of wavelets. This new representation is stable against deformations and is translation invariant. This family of wavelets does not need to be learned and hence, the features can be extracted without training which can then be passed on further to Machine Learning models (Oyallon et al., 2013). The Scattering Transform has been successfully used by Seydoux et al. (2020) for clustering earthquakes in an unsupervised fashion. The scattering transform works by successively convolving wavelets with the signal and applying modulus non-linearity at each step. This can be shown as follows:

$$S_{x(t)} = ||x * \psi_{\lambda 1}| * \psi_{\lambda 2}|...| * \psi_{\lambda m}| * \phi \tag{1}$$

where $S_{x(t)}$ is the set of scattering coefficients obtained at step m, $x(t)$ is the signal, $\psi_{\lambda_m}(t)$ denotes the set of wavelets at step m and $\phi(t)$ is a low-pass filter.

In our experiments, we use the Kymatio library (Andreux et al., 2020) to perform the scattering transform and get a set of features which are passed on further to our supervised neural network model.

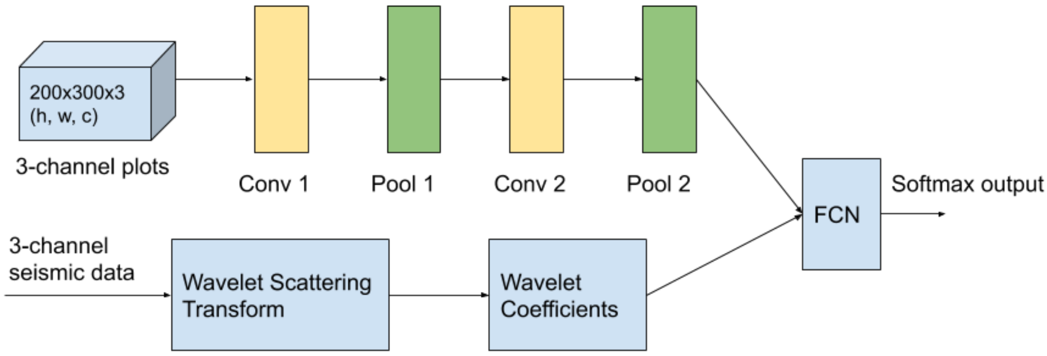

Figure 2. WavImg Model Architecture

3.2 Convolution over 3-channel plots

As the volunteers classify the data based on plots and associated audio records of vertical ground motion, we decided to try out a similar approach with our model. Along with the wavelet coefficients, we additionally provide the model with 3-channel plots (BHZ, BHE, BHN) as input. Image convolution is applied over the 3-channel plots and the resultant features are concatenated with the wavelet coefficients and then passed through a fully connected neural network (FCN).

4 Experiments

For all experiments, we use the following hyper parameters: learning rate = 1e-5, batch size = 100, epochs = 300. We perform a 3-way classification between Earthquakes/Tremors/Noise. As the number of tremor samples is considerably less, we apply a weighted cross-entropy loss where the weights are calculated as follows:

$$w_i = N_{largest}/N_i \tag{2}$$

where w_i is the weight assigned to class i, $N_{largest}$ is the number of samples of the largest class and N_i is the number of samples of class i. For all experiments, we perform a 80-20 stratified train/test split of the data. We test two models:

1. **WavNet:** In this model, we perform a wavelet scattering transform on 3-channel seismic data (BHZ, BHN, BHE) and extract relevant features from it. These features (wavelet coefficients) are then passed on to a 2-layer fully connected network.
2. **WavImg:** This model combines the Wavelet Scattering Transform with 3-channel convolution over the image plots. The combined features are then passed on through a 2-layer fully connected network. (Figure 2)

4.1 Training the Machine Learning model

Training on clean data: To test the efficacy of wavelet transform, we first run a simple experiment with WavNet. As an upper baseline, we first ran the experiment on clean data (data cleaned and filtered by our seismologists, refer Appendix A). The model converges to 95.2% training and a 94.4% testing accuracy.

Training on clean + gold users data: Compared to the clean data, the data from the Earthquake Detective is difficult to segregate due to its low amplitude signals and larger time window (~33 mins). For this experiment, we consider data from gold users (Earthquake Detective data labeled by our experts) and combine it with the clean

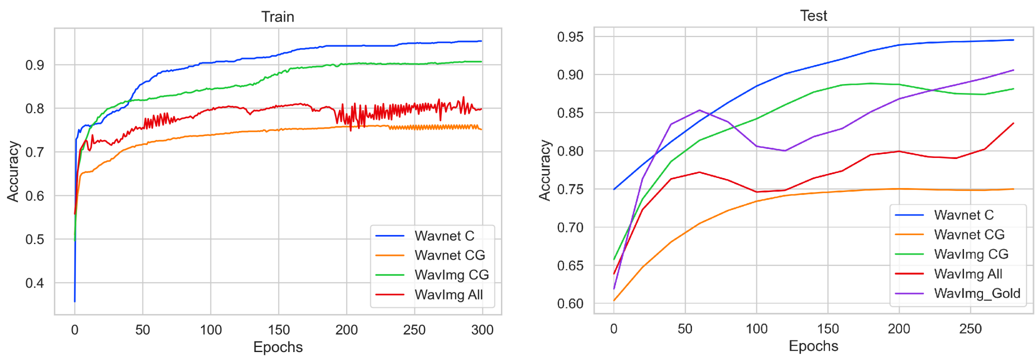

Figure 3. Model comparison chart. (C = Clean, CG = Clean+Gold, All = Clean+Gold+Volunteer)

Table 1. Comparison of model reliability with volunteers

Metric	ML Model*	Volunteers	Volunteers (top 35%)
F1-Score (Earthquake)	0.909	0.785	1.0
F1-Score (Noise)	0.914	0.738	0.975
F1-Score (Tremor)	0.888	0.415	0.924

*Here model refers to WavImg All (Clean + Gold + Volunteer)

dataset from the previous experiment. When this data was trained on the WavNet, it gave a train accuracy of 75.4% and a test accuracy of 74.9%. Next, we trained the same data using the WavImg model. This model produces a 91.4% train and 89.6% test accuracy.

Training on clean + gold + volunteer's data: Finally, we combine the previous data with data from two volunteers. For each volunteer, we calculated a reliability score which includes a precision, recall and f1 score for each class. These scores were calculated by comparing volunteer's classification with gold-set labeled by our experts. To handle the unreliability introduced in labels, we add an additional gold-test set which consists of samples labeled from our gold users that were not used for training. WavImg produces a 80.1% train accuracy, 83.6% test accuracy and **90.4%** gold-test accuracy. (Figure 3)

4.2 Results and Analysis

The WavNet model was able to perform extremely well on the clean data (95% accuracy) which proves that wavelet scattering transform extracts relevant features which can then be trained using a simple 2-layer FCN. However, due to the greater complexity of Earthquake Detective data, WavNet by itself is insufficient. The WavImg model overcomes this problem by using information from 3-channel plots. One interesting case was the last experiment in which there was variance in training due to label uncertainty (80% accuracy) but the model still performed well on the gold-test set (90.4% accuracy). (Figure 3) This shows that the model still ends up learning useful representations despite the uncertainty of the labels. For more in-depth analysis over select misclassified samples refer to Appendix B.

4.3 Model Comparison with Volunteers

To see how reliable our model is in comparison to volunteers, we calculated the f1-score (reliability score) of the model using the gold-test for each class separately. The f1-score is calculated as follows:

$$F_1 = (2 * precision * recall)/(precision + recall) \tag{3}$$

where precision is $TP/(TP{+}FP)$ and recall is $TP/(TP{+}FN)$, where TP is true positive, FP is false positive and FN is false negative.

From table 1 we can see that the model is much more reliable for all classes when compared to all volunteers. However, we found that the volunteers become more reliable than the model in all cases, when top 35% of them are selected. So this implies that while the model is better than an average volunteer, it is still not as good as the top volunteers.

5 Conclusions

The Earthquake Detective dataset is the first crowdsourced dataset for potentially dynamically triggered local earthquake and tremor signals. This is also the first time that potentially dynamically triggered tremor signals were used in and detected by ML. Our experiments provide ML baselines for the data. We have only trained our models on a small subset of the 130k+ samples available. Also, from section 4.3 we saw that the top volunteers are better than the ML model. Therefore, when all of the data is considered, better techniques to incorporate reliability scores into the model will be required.

We encourage researchers to 1) use this dataset as a catalog for potentially dynamically triggered seismic events, 2) augment the data and algorithms beyond the baseline, and 3) stream their data through Earthquake Detective for accelerating labeling of their data sets and/or for validating previously unlabeled ML results, by connecting with Earthquake Detective developers.

Appendix A Data

A1 Clean Data

The clean data used in the experiments was created as follows:

1. A set of 1000 s-long waveforms with confirmed PT tremor signals, labeled by the seismologists among us. The waveforms were resampled at 20 samples per second and band-pass filtered between 2 and 5 Hz. This data set is subdivided as follows:
   (a) Waveforms recorded by seismic stations in Taiwan or Japan, trimmed around surface waves from 6 large-M earthquakes in the eastern hemisphere (the Great Tohoku Earthquake, 1 more earthquake from Hokkaido, Japan, 1 from Qinghai, China, and 4 from Sumatra, Indonesia). These are the positive examples.
   (b) Waveforms recorded by seismic stations in Taiwan, Malaysia, Australia, and of the Global Seismographic Network (GSN), selected for having no significant signals from earthquakes or otherwise. These are the negative examples.
2. A set of 1000 s-long waveforms with confirmed PT local earthquake signals, labeled by the seismologists among us. All data was band-pass filtered between 2 and 8 Hz.
   (a) Waveforms recorded by seismic stations from USArray in the USA and the Hi-CLIMB array in Tibet, trimmed around surface waves from the 2010 M8.8 Maule Earthquake and 7 additional large-M earthquakes in the eastern hemisphere (1

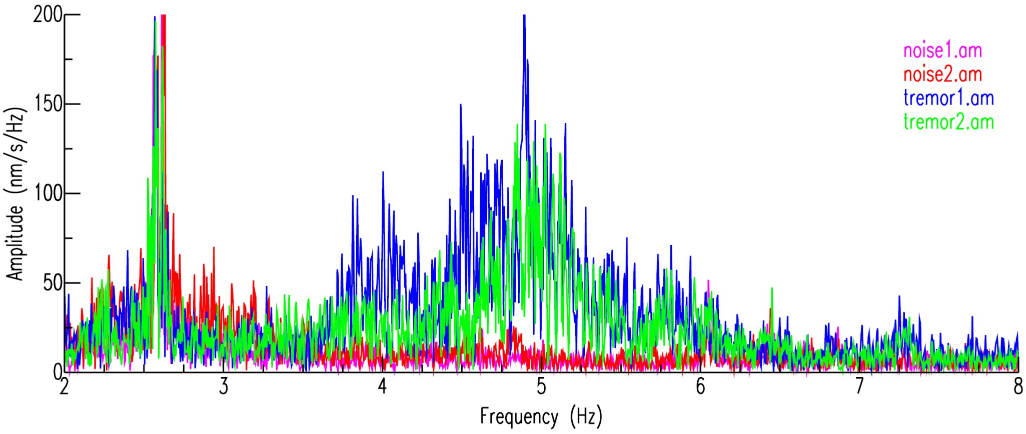

Figure A1. Amplitude spectra for four 200-s subsets of a 2000-s wave train labeled as tremor and misclassified as noise. Two subsets (red and magenta) represent noise and two (blue and green) contain a tremor-like signal.

from China, 1 from Japan, and 5 from Sumatra, Indonesia). Waveforms that showed signals from local earthquakes were labeled as positive.

(b) The waveforms from this set (2a) that were not labeled positive - they were labeled as negative examples.

(c) Additional waveforms recorded by the Hi-CLIMB array from 10 random local earthquakes with M<3.6. This auxiliary data set was used to expand the set of positive examples. An corresponding number of negative examples recorded by the same array was added to the negative examples.

A2 Data Distribution

- In the first experiment we used 551 earthquake samples, 570 noise samples and 39 tremor samples from the clean data.
- In the second experiment (clean + gold users data), we had the following data distribution: 1031 earthquake samples, 1014 noise samples and 48 tremor samples.
- In the final experiment (which additionally includes the chosen volunteer's data) had the following data distribution: 3013 earthquake samples, 2436 noise samples, 203 tremor samples.

Appendix B Model Analysis

To gain further insights into how the WavImg model classifies, we looked at select examples. One example is where the ML algorithm classified a wave train with tremor as a noise wave train. This could be the result of

- There being two separate bursts of tremor,
- Both tremor signals being relatively short in duration,
- The tremor signals being weak,
- The tremor signals sounding different from more typical tremor signals
- The presence of non-stationary noise signals.

Extracting two 200-s tremor signals and two 200-s stationary noise signals from the 2000-s wave train reveals that the the weak tremor signals have some additional power between 3 and 6 Hz, compared to the noise (Figure A1). This is not entirely characteristic but still consistent for tremor signals. Although the wave train was likely labeled correctly, the wave train is not a role model for its class and hence may have confused the ML algorithm. In one other case labeled as tremor and classified as earthquake, the tremor signal was so brief that is easy to mistake for an earthquake signal. In another case, strongly peaked signals elsewhere in the wave train might have distracted the ML from the tremor signal.

The type of waveform data used in our study contains a wild variety of noise signals, for which we did not designate a single class. However such noise signals can interfere with the ability of volunteers, and sometimes experts to correctly label wave trains. In at least 5 cases, wave trains with noise signals, labeled as noise, were misclassified as earthquakes.

Several other cases of misclassification by the ML algorithm can be traced to a mislabeling of the original data. In three cases of wave trains with noise signals mislabeled as tremor, the ML algorithm classified the wave trains as earthquakes. The ML algorithm also classified a case of mislabeled tremor correctly as noise. In at least 4 cases, the ML algorithm correctly classified wave trains as noise, while they were labeled as earthquakes. In these four cases, listening to and viewing spectral properties of the wave trains confirmed in hindsight that these signals should have indeed been labeled as noise.

Acknowledgments
This research was funded by the Integrated Data-Driven Discovery in Earth and Astrophysical Sciences (ID^{3}EAS) program under National Science Foundation grant NSF-NRT 1450006. We thank Josiah Evans and Blaine Rothrock for helpful discussions. A version of this paper was presented at the AI 4 Earth workshop at the NeurIPS 2020 conference; we thank the organizers for that opportunity to share and learn. This paper shares data generated via the zooniverse.org platform, development of which is funded by a Global Impact Award from Google and a grant from the Alfred P. Sloan Foundation.

Data Availability: Raw seismic data used in Earthquake Detective were retrieved via IRIS Data Services (DS) (http://ds.iris.edu/mda), including from the following seismic networks (1) AK (Alaska Regional Network, 1987); (2) TA (IRIS TA, 2003); (3) PB (UNAVCO, 2004); and (4) UW (PNSN, 1963). Additional seismic data were retrieved from the Broadband Array in Taiwan for Seismology (BATS) (http://bats.earth.sinica.edu.tw) by K.C., and via IRIS DS from (1) AZ (Vernon, 1982); (2) BK (NCEDC, 2014); (3) CI (CalTech/USGS, 1926); (4) II (Scripps, 1986); (5) IU (ASL/USGS, 1988); (6) TA (IRIS TA, 2003); (7) US (USNSN, ASL/USGS, 1990); (8) XF (Nepal–Himalaya–Tibet Seismic Transect, Nabelek, 2002); (9) XR (SIEDCAR, Pulliam and Grand, 2008); and (10) XV (BASE, Sheehan and Miller, 2009). The code and labeled data is available online at https://zenodo.org/record/4741553,

Cite as: Omkar Ranadive (2021), Omkar-Ranadive/Earthquake-Detective: Earthquake Detective V1, , doi:10.5281/zenodo.4741553.

and the Earthquake Detective website can be accessed at

https://www.zooniverse.org/projects/vivitang/earthquake-detective .

References

Andreux, M., Angles, T., Exarchakis, G., Leonarduzzi, R., Rochette, G., Thiry, L., ... others (2020). Kymatio: Scattering transforms in python. *Journal of Machine Learning Research*, *21*(60), 1–6.

Bergen, K. J., Chen, T., & Li, Z. (2019). Preface to the focus section on machine

learning in seismology. *Seismological Research Letters*, *90*(2A), 477–480.
Chao, K., Peng, Z., Wu, C., Tang, C., & Lin, C. (2012). Remote triggering of non-volcanic tremor around taiwan. *Geophys. J. Int.*, *188*, 301– 324. doi: doi:10.1111/j.1365246X.2011.05261.x
Chmiel, M., Walter, F., Wenner, M., Zhang, Z., McArdell, B. W., & Hibert, C. (2021). Machine learning improves debris flow warning. *Geophysical Research Letters*, *48*(3), e2020GL090874. doi: https://doi.org/10.1029/2020GL090874
Gomberg, J., Rubinstein, J. L., Peng, Z., Creager, K. C., Vidale, J. E., & Bodin, P. (2008). Widespread triggering of nonvolcanic tremor in california. *Science*, *319*(5860), 173–173.
Gutenberg, B., & Richter, C. (1954). *Seismicity of the earth 2nd ed., 310.* Princeton University Press, Princeton, NJ.
Hill, D. P., & Prejean, S. (2015). Dynamic triggering. *Treatise on Geophysics, 4, Earthquake Seismology*, 273-304. doi: doi.org/10.1016/B978-0-444-53802-4.00078-6
Kano, M., Aso, N., Matsuzawa, T., Ide, S., Annoura, S., Arai, R., … Obara, K. (2018, 06). Development of a Slow Earthquake Database. *Seismological Research Letters*, *89*(4), 1566-1575. Retrieved from https://doi.org/10.1785/0220180021 doi: 10.1785/0220180021
Lee, T. R., Phrampus, B. J., Obelcz, J., Wood, W. T., & Skarke, A. (2020). Global marine isochore estimates using machine learning. *Geophysical Research Letters*, *47*(18), e2020GL088726.
Li, W., Narvekar, N., Nakshatra, N., Raut, N., Sirkeci, B., & Gao, J. (2018). Seismic data classification using machine learning. In *2018 ieee fourth international conference on big data computing service and applications (bigdataservice)* (pp. 56–63).
Li, Z., Meier, M.-A., Hauksson, E., Zhan, Z., & Andrews, J. (2018). Machine learning seismic wave discrimination: Application to earthquake early warning. *Geophysical Research Letters*, *45*(10), 4773–4779.
Liu, Y.-H., Yeh, T.-C., Chen, K. H., Chen, Y., Yen, Y.-Y., & Yen, H.-Y. (2019). Investigation of single-station classification for short tectonic tremor in taiwan. *Journal of Geophysical Research: Solid Earth*, *124*(8), 8803–8822.
Meier, M.-A., Ross, Z. E., Ramachandran, A., Balakrishna, A., Nair, S., Kundzicz, P., … Yue, Y. (2019). Reliable real-time seismic signal/noise discrimination with machine learning. *Journal of Geophysical Research: Solid Earth*, *124*(1), 788–800.
Mousavi, S. M., & Beroza, G. C. (2020). A machine-learning approach for earthquake magnitude estimation. *Geophysical Research Letters*, *47*(1), e2019GL085976.
Obara, K. (2002). Nonvolcanic deep tremor associated with subduction in southwest japan. *Science*, *296*, 1679-1681. doi: 10.1126/science.1070378
Oyallon, E., Mallat, S., & Sifre, L. (2013). Generic deep networks with wavelet scattering. *arXiv preprint arXiv:1312.5940*.
Reynen, A., & Audet, P. (2017). Supervised machine learning on a network scale: Application to seismic event classification and detection. *Geophysical Journal International*, *210*(3), 1394–1409.
Riggelsen, C., & Ohrnberger, M. (2014). A machine learning approach for improving the detection capabilities at 3c seismic stations. *Pure and Applied Geophysics*, *171*(3), 395–411.
Ross, Z. E., Meier, M.-A., Hauksson, E., & Heaton, T. H. (2018). Generalized seismic phase detection with deep learning. *Bulletin of the Seismological Society of America*, *108*(5A), 2894–2901.
Rouet-Leduc, B., Hulbert, C., & Johnson, P. A. (2018). Breaking cascadia's silence: Machine learning reveals the constant chatter of the megathrust. *arXiv preprint arXiv:1805.06689*.

Rouet-Leduc, B., Hulbert, C., Lubbers, N., Barros, K., Humphreys, C. J., & Johnson, P. A. (2017). Machine learning predicts laboratory earthquakes. *Geophysical Research Letters*, *44* (18), 9276–9282.

Ruano, A. E., Madureira, G., Barros, O., Khosravani, H. R., Ruano, M. G., & Ferreira, P. M. (2014). Seismic detection using support vector machines. *Neurocomputing*, *135*, 273–283.

Seydoux, L., Balestriero, R., Poli, P., De Hoop, M., Campillo, M., & Baraniuk, R. (2020). Clustering earthquake signals and background noises in continuous seismic data with unsupervised deep learning. *Nature communications*, *11* (1), 1–12.

Tang, V., Chao, K., & van der Lee, S. (2021, 04). Detections of Directional Dynamic Triggering in Intraplate Regions of the United States. *Bulletin of the Seismological Society of America*. Retrieved from `https://doi.org/10.1785/0120200352` doi: 10.1785/0120200352

Tang, V., Rösler, B., Nelson, J., Thompson, J., van der Lee, S., Chao, K., & Paulsen, M. (2020). Citizen scientists help detect and classify dynamically triggered seismic activity in alaska. *Frontiers in Earth Science*, *8*, 321.

Tang, V., Seetharaman, P., Chao, K., Pardo, B. A., & van der Lee, S. (2020). Automating the detection of dynamically triggered earthquakes via a deep metric learning algorithm. *Seismological Research Letters*, *91* (2A), 901–912.

Velasco, A. A., Hernandez, S., Parsons, T., & Pankow, K. (2008). Global ubiquity of dynamic earthquake triggering. *Nature geoscience*, *1* (6), 375–379.

Wiszniowski, J., Plesiewicz, B., & Lizurek, G. (2021). Machine learning applied to anthropogenic seismic events detection in lai chau reservoir area, vietnam. *Computers & Geosciences*, *146*, 104628.

Yun, N., Yang, H., & Zhou, S. (2021). Dyntripy: A python package for detecting dynamic earthquake triggering signals. *Seismological Society of America*, *92* (1), 543–554.

Zhao, W. L., Gentine, P., Reichstein, M., Zhang, Y., Zhou, S., Wen, Y., ... Qiu, G. Y. (2019). Physics-constrained machine learning of evapotranspiration. *Geophysical Research Letters*, *46* (24), 14496–14507.

Figure 1.

200
100
0
−100
−200
Velocity (nm/s)
quake
noise
tremor
0
500
1000
1500
2000
Time (s)

Figure 2.

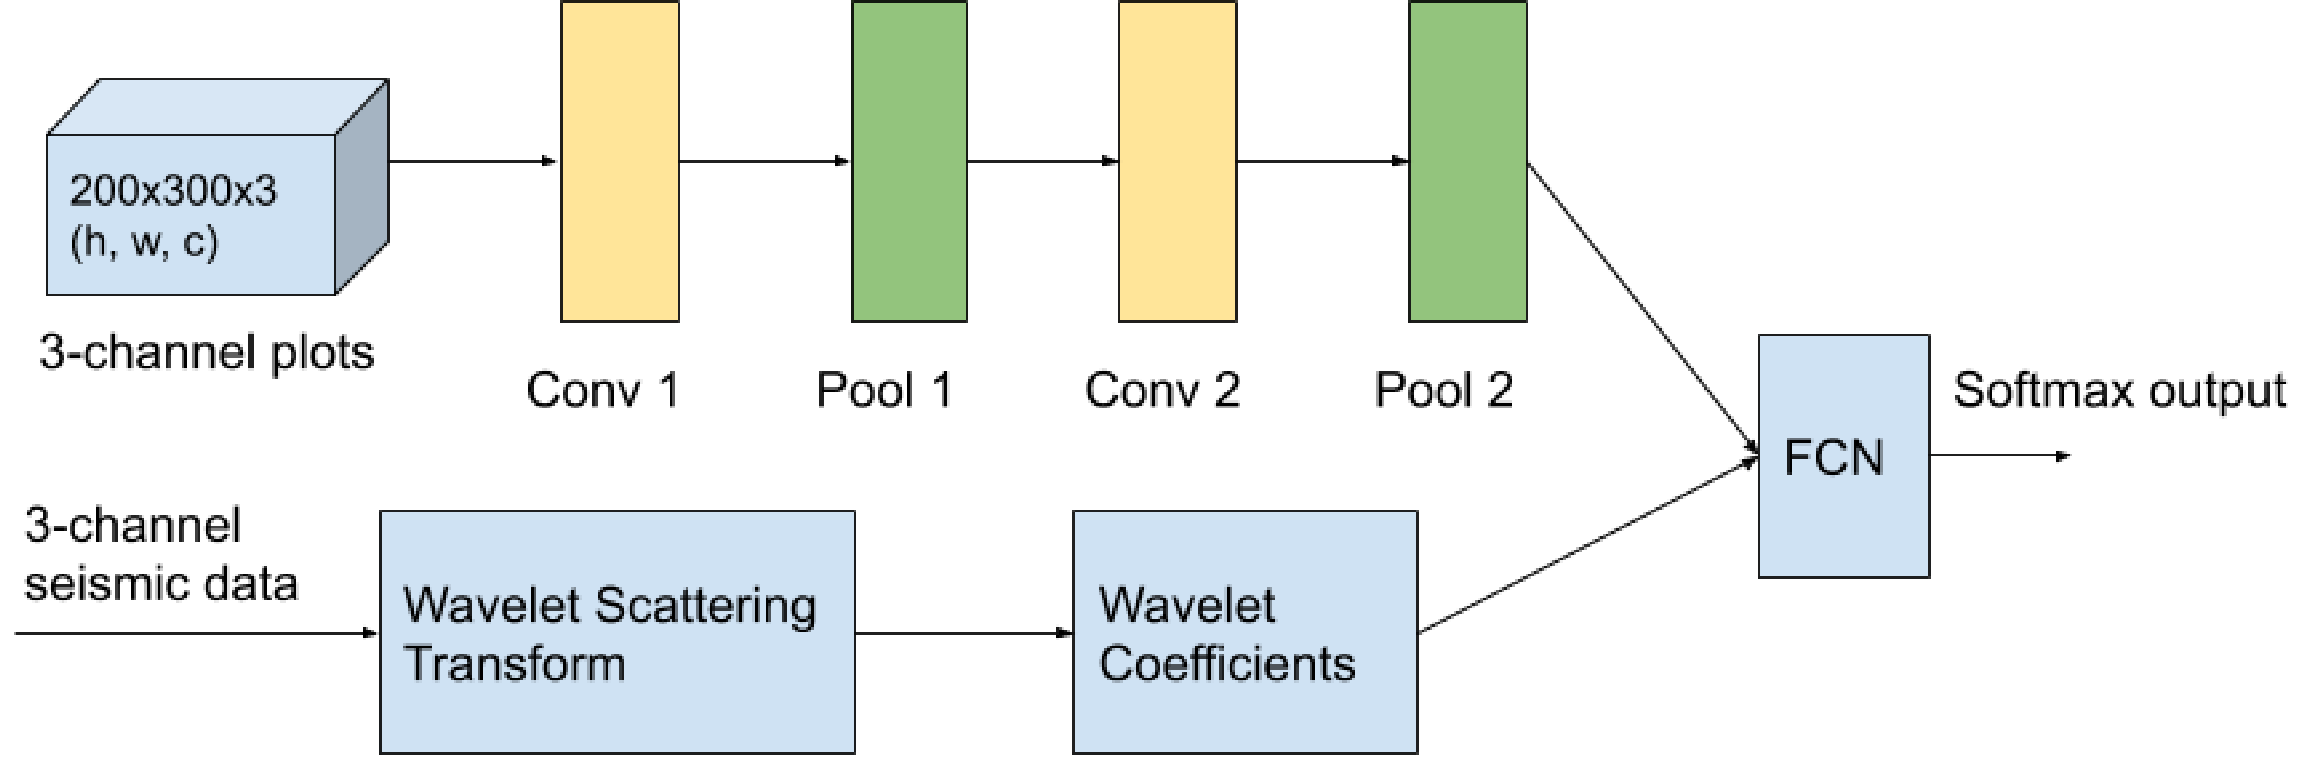
200x300x3
(h, w, c)
3-channel plots
Conv 1
Pool 1
Conv 2
Pool 2
FCN
Softmax output
3-channel
seismic data
Wavelet Scattering
Transform
Wavelet
Coefficients

Figure 3.

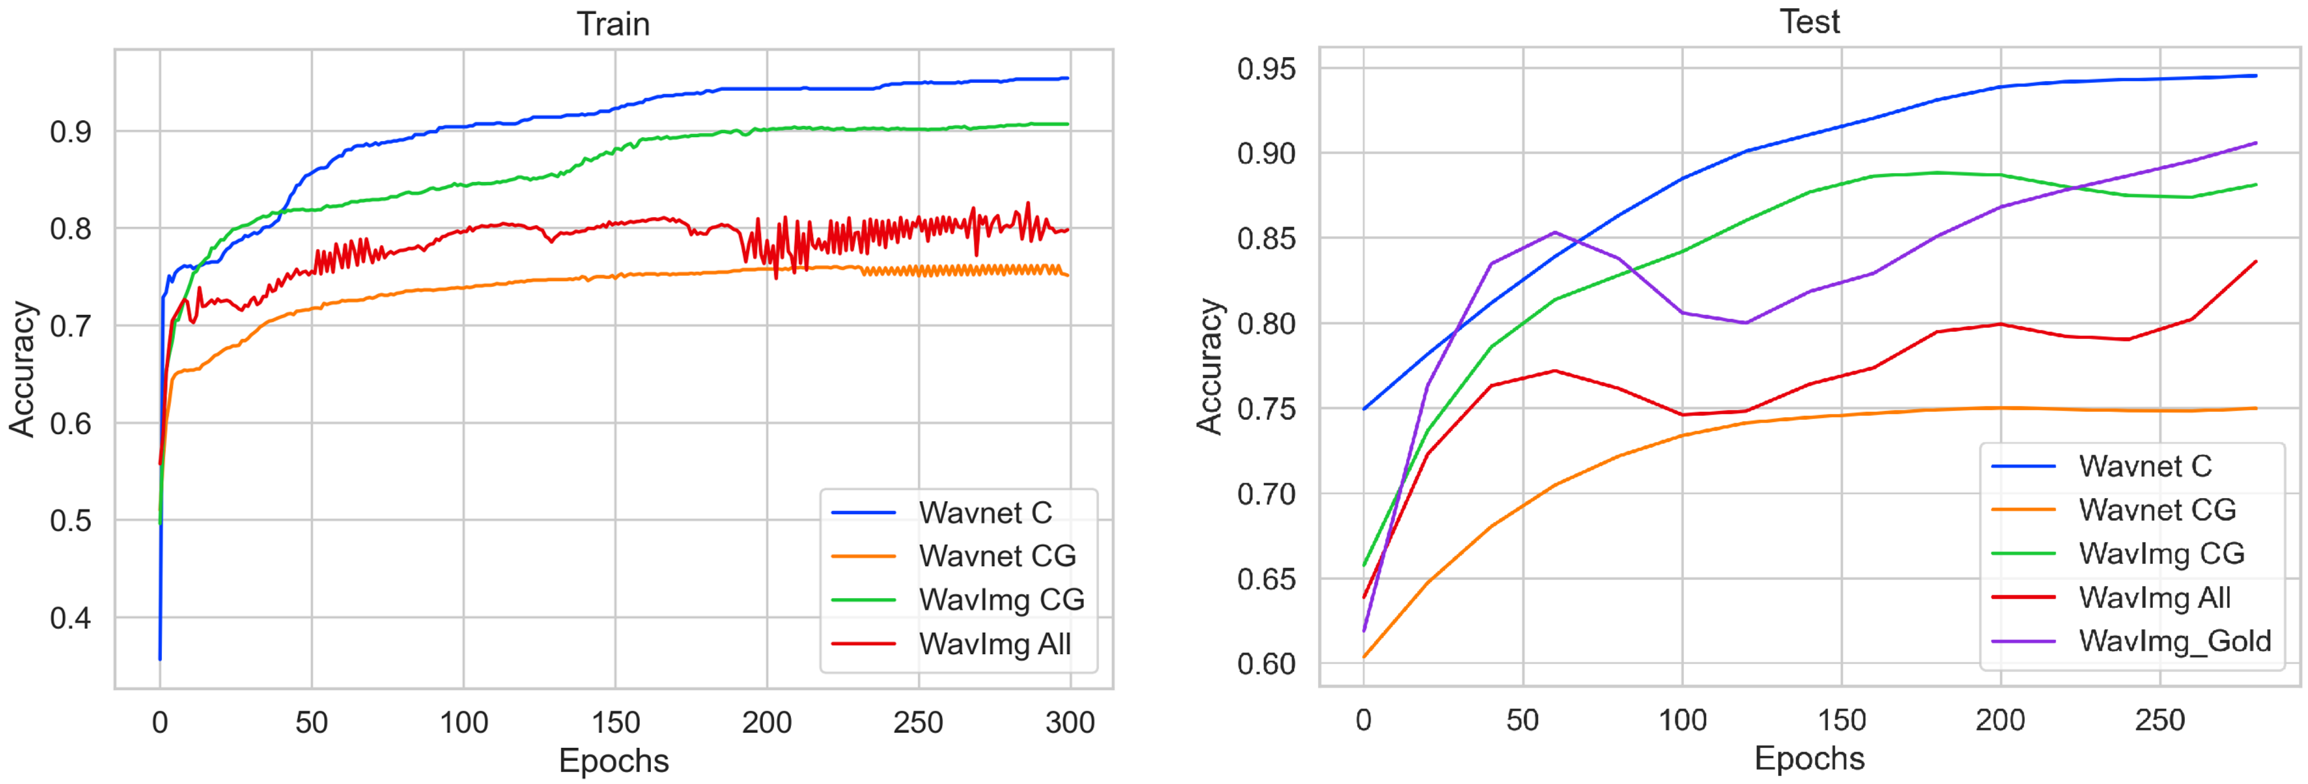
Train
Accuracy
Epochs
Wavnet C
Wavnet CG
WavImg CG
WavImg All
Test
Accuracy
Epochs
Wavnet C
Wavnet CG
WavImg CG
WavImg All
WavImg_Gold

Figure A1.


Amplitude (nm/s/Hz)
Frequency (Hz)
200
150
100
50
0
2
3
4
5
6
7
8
noise1.am
noise2.am
tremor1.am
tremor2.am