Application of a Paired Neural Network to Aftershock Identification

Andrea Conley, Brendan Donohoe, Benjamin Greene
Sandia National Laboratories

P3.6-143

POSTER

Sandia National Laboratories

NNSA National Nuclear Security Administration

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.

Application of a Paired Neural Network to Aftershock Identification

Andrea Conley, Brendan Donohoe, Benjamin Greene

ABSTRACT

Large aftershock sequences cause problems for the International Data Centre (IDC) because the seismic event rate increases dramatically during an aftershock sequence, making correct association of arrivals difficult for the automated pipeline. Aftershock sequences can continue for days or even months after a large earthquake and although aftershocks aren't events of interest for treaty monitoring purposes, they must be reviewed and eliminated by analysts, resulting in delayed release of the IDC bulletins. We turn to machine learning to automatically identify aftershock events and improve automated pipeline performance. In our research, we train a paired neural network (PNN) to automatically perform aftershock identification based on waveform similarity, even when only a few datapoints are available for training. This allows the model to be applied to classes outside of the original training dataset. We analyze the ability of our PNN to classify aftershock data constructed from signals recorded by the IMS network and several open IRIS networks added to real noises from the STanford Earthquake Dataset (STEAD) or the University of Utah network. We apply the trained model and waveform cross-correlation on the constructed test dataset and compare the performance of the two approaches.

Application of a Paired Neural Network to Aftershock Identification

Andrea Conley, Brendan Donohoe, Benjamin Greene

INTRODUCTION

- The ability to automatically identify nuisance aftershock events to reduce analyst workload when searching for events of interest is an important step in improving nuclear monitoring capabilities.
- While waveform cross-correlation methods have proven successful, they have limitations (e.g., difficulties with spike artifacts, multiple aftershocks in the same window) that machine learning may be able to overcome.
- Here we apply a Paired Neural Network (PNN) to a dataset consisting of real, high quality signals added to real seismic noises in order to work with controlled, labeled data and establish a baseline of the PNN's capability to identify aftershocks.

Application of a Paired Neural Network to Aftershock Identification

Andrea Conley, Brendan Donohoe, Benjamin Greene

METHODS

The Paired Neural Network (PNN) used in this study consists of two instances of the same underlying sub-architecture, where weights are shared between the two sub-architectures.

- One waveform from a waveform pair is input into each sub-architecture to obtain a corresponding vector.
- The Euclidean distance between the two output vectors corresponds to the estimated similarity of the waveform pair, with a far distance indicating a non-match and vice versa.
- Dropout layers are included to regularize the model and to estimate PNN score uncertainties.
- Uncertainties are estimated by performing inferences on the current data multiple times with a random set of connections deactivated with each inference, allowing for a distribution from which the mean and standard deviation can be estimated.
- Finally, a contrastive loss function is used.

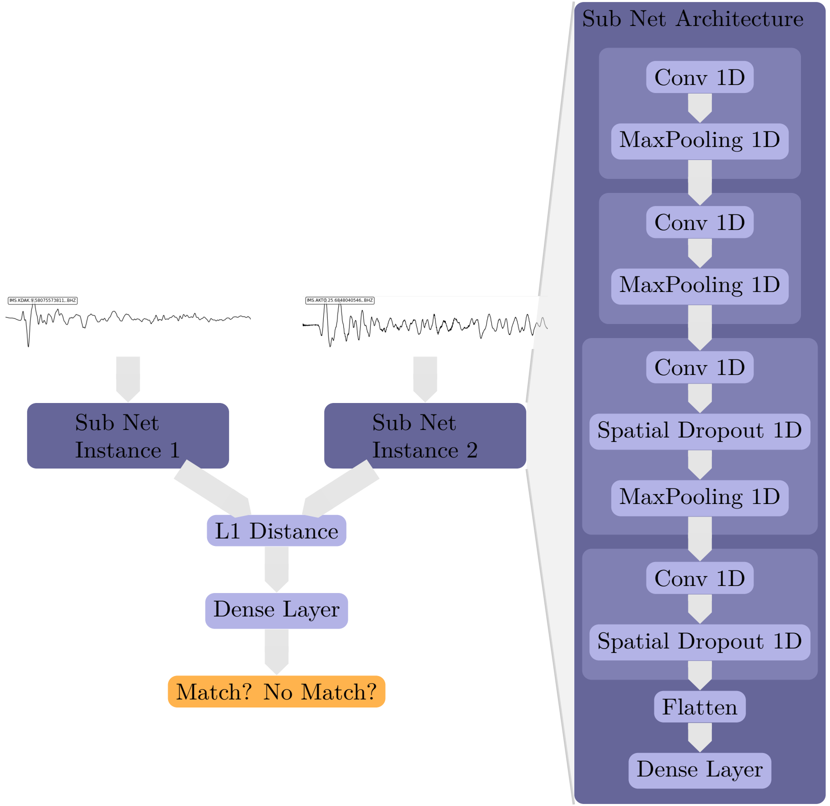

Application of a Paired Neural Network to Aftershock Identification

Andrea Conley, Brendan Donohoe, Benjamin Greene

DATASET

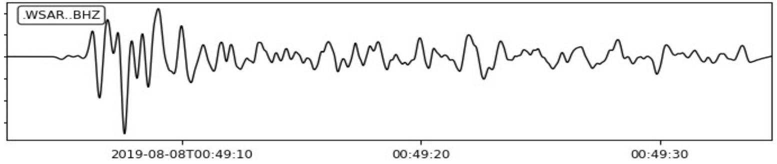

+

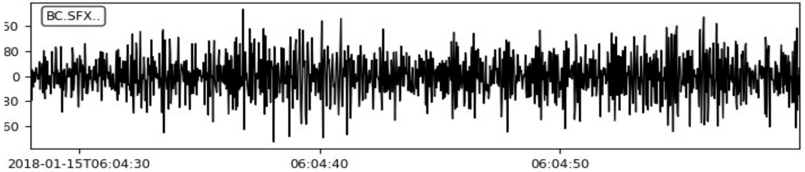

=

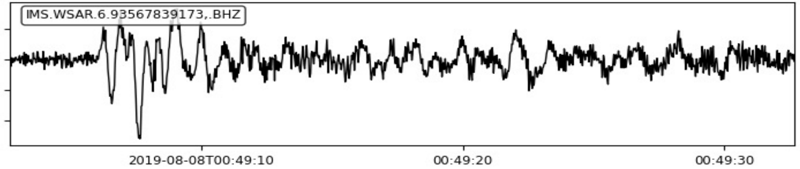

Signals and noises are added together (see above), resulting in 61,297 constructed signals from which we generate the training, validation, and test datasets.

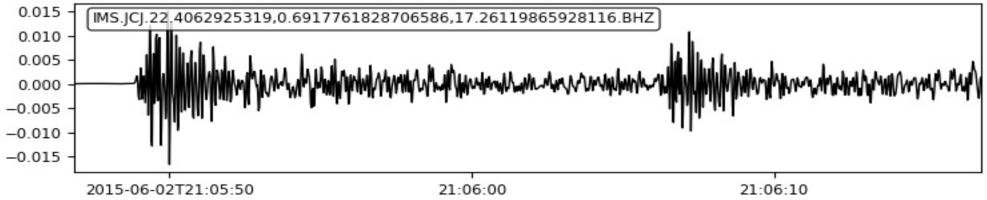

10,686 of the generated waveforms are overlapping waveforms, i.e. one window contains two copies of the same signal at different amplitudes, as shown above.

- Overlapping waveforms are a proxy for when two aftershocks occur in close proximity in time and space.
- A random signal is selected and a copy of that signal is added to the original signal after a time delay is applied.
- The copied signal is multiplied by a random factor between 0.1-10 to modify amplitude relative to the original signal.
- An overlap length between 0 and 20 is applied. An overlap length of 0 corresponds to two signals arriving simultaneously. An overlap length of 20 corresponds to the second signal arriving 20 seconds after the first signal.

SnT 2021
CTBT: SCIENCE AND TECHNOLOGY CONFERENCE

Poster No.: P3.6-143

Application of a Paired Neural Network to Aftershock Identification

Andrea Conley, Brendan Donohoe, Benjamin Greene

CTBTO
PREPARATORY COMMISSION

DATASET

Signals consist of:

- High SNR first P-arrivals picked on the BHZ component recorded by the IMS (from 2007-2020; red stations) or open IRIS networks (from 2011; black stations).
- Signals can be generated by events of any magnitude or epicentral distance.
- Signals are resampled to 40 Hz, demeaned, detrended, and high-pass filtered at a corner frequency of 0.3 Hz.
- Signal SNR compared to pre-P noise is required to be > 20 dB.
- Signals are windowed to be 30 seconds in total duration, beginning 2 seconds before the P-arrival.

Noises come from two datasets:

- The University of Utah (UU) dataset (†Tibi et al., 2021) and the STanford EArthquake Dataset (†STEAD; Mousavi et al, 2019).
- Contain no contaminating signals.
- Noises are preprocessed in the same way as signals, except no high-pass filter applied.

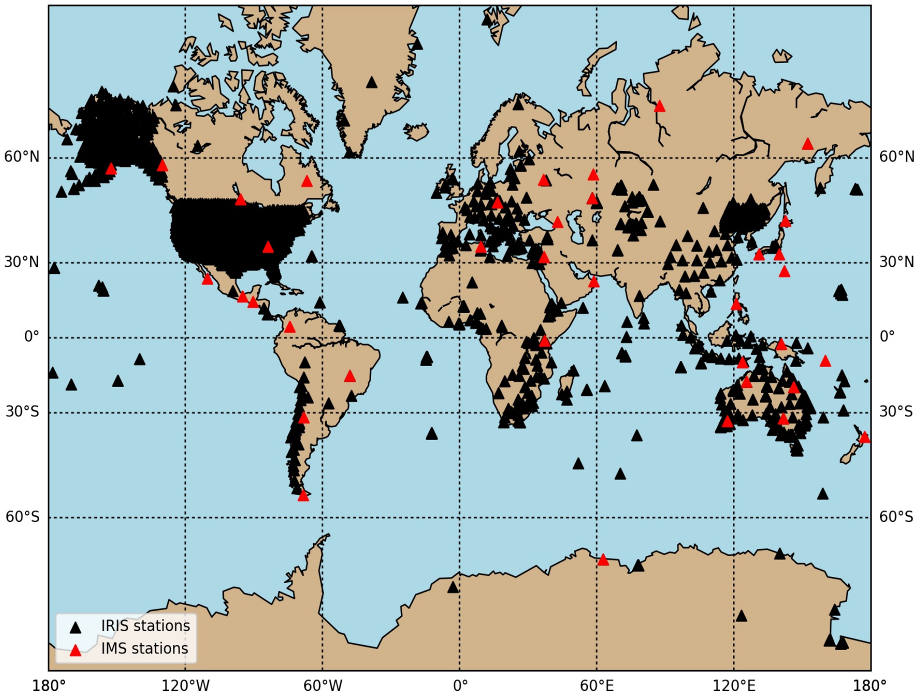

†Tibi, R., P. Hammond, R. Brogan, C. J. Young, & Koper, K. (2021), doi: 10.1785/0120200292.
†Mousavi, S. M., Sheng, Y., Zhu, W. & Beroza, G.C. (2019), doi: 10.1109/ACCESS.2019.2947848.

Application of a Paired Neural Network to Aftershock Identification

Andrea Conley, Brendan Donohoe, Benjamin Greene

DATASET

- Waveform pairs consisting of two constructed waveforms randomly selected from the 61,297 waveform dataset are input into the PNN.
  - Both waveforms in a pair may have been constructed with either the same recorded signal (a match) or different recorded signals (not a match).
- The training, validation, and test datasets consist of:
  - Training – 48000 waveform pairs (24000 matching and nonmatching)
  - Validation – 800 waveform pairs (400 matching and nonmatching)
  - Test –800 waveform pairs (400 matching and nonmatching).
- No noises or signals are shared between the training, validation, and test datasets.

Application of a Paired Neural Network to Aftershock Identification

Andrea Conley, Brendan Donohoe, Benjamin Greene

RESULTS

- We compare PNN test dataset scores to waveform Cross-Correlation (CC) scores for the same dataset.
- To get CC scores, a zero-normalized cross-correlation is performed on the test dataset.

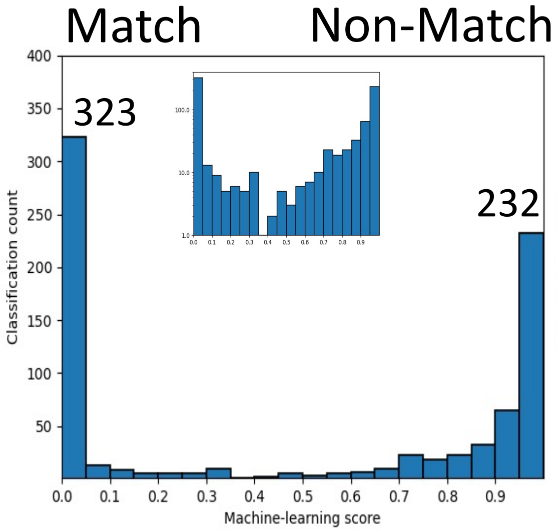

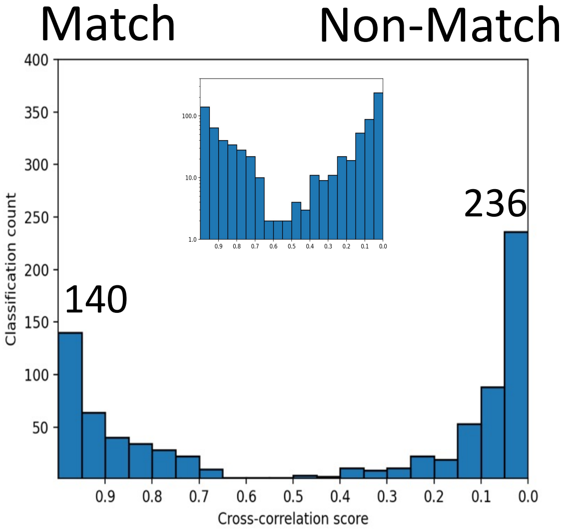

- **81% and 35% of matching pairs clearly classified by the PNN (left) and CC (right).**
- **58% and 59% of non-matching waveform pairs clearly classified by the PNN and CC, respectively.**

Application of a Paired Neural Network to Aftershock Identification

Andrea Conley, Brendan Donohoe, Benjamin Greene

RESULTS

We also compare the Receiver-Operating Characteristic (ROC) curves of the PNN (pink-to-green) and CC (blue-to-red), where the True Positive Rate (TPR; y-axis) and False Positive Rate (FPR; x-axis) are defined as:

- $True\ Positive\ Rate = \frac{TP}{TP+FN}, \quad False\ Positive\ Rate = \frac{FP}{TN+FP}$

Both the PNN and CC ROC curves have large Areas Under the Curve (AUCs), where AUC is a measure of model accuracy (AUC = 1.0: 100% correct, AUC = 0.0: 100% incorrect). The PNN AUC = 98.4% and the CC AUC = 98.1%.

However, when we zoom into the corner of the curve, we see that the PNN achieves a higher TPR than CC at all values ranging from ~87.5% to 95% TPR and ~0% to 10% FPR.

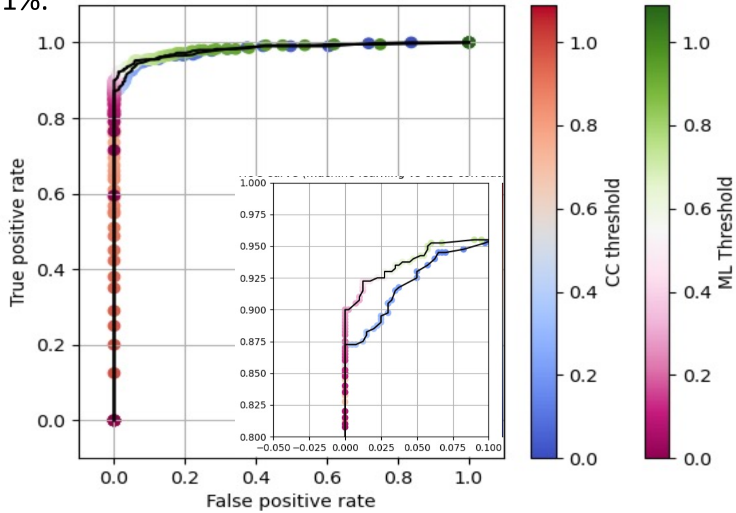

PNN thresh = 0.29, FPR = 0%, TPR = 90%	Actual Positive	Actual Negative
Predicted Positive	360	0
Predicted Negative	40	400

CC thresh = 0.29, FPR = 3%, TPR = 90%	Actual Positive	Actual Negative
Predicted Positive	362	12
Predicted Negative	38	388

At a 90% TPR, the PNN misclassifies 10% of the matching waveform pairs as non-matching and 0% of non-matching pairs as matching (see top confusion matrix).

For CC at the same TPR, 9.5% of matching waveform pairs and 3% of non-matching waveform pairs are misclassified (see bottom confusion matrix).

SnT 2021
CTBT: SCIENCE AND TECHNOLOGY CONFERENCE

Poster No.: P3.6-143

Application of a Paired Neural Network to Aftershock Identification

Andrea Conley, Brendan Donohoe, Benjamin Greene

CTBTO PREPARATORY COMMISSION

RESULTS

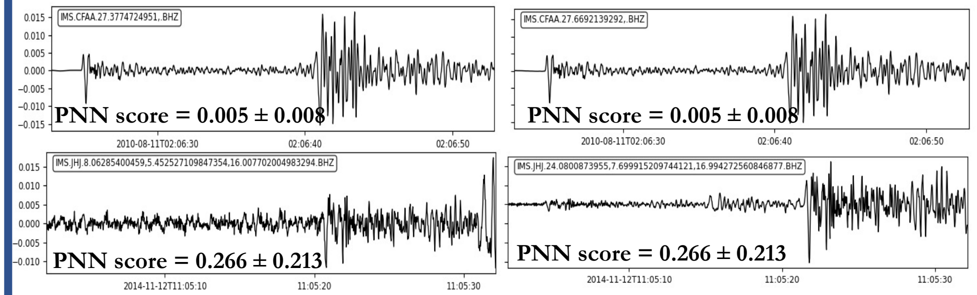

True Positive Examples

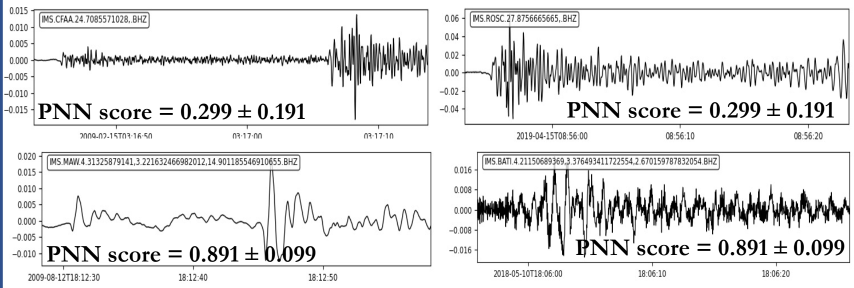

True Negative Examples

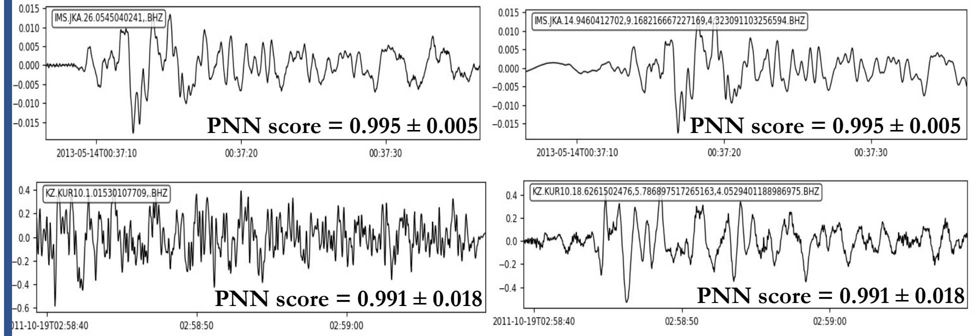

False Negative Examples

Examples of waveform pairs resulting in true positive, true negative, and false negative classifications are shown at a threshold of 0.29.

- The 0.29 threshold was selected using the previously shown ROC curve.
- This threshold is selected to optimize the FPR.
- At 0.29, TPR = 90% and FPR = 0%.

At this threshold, the PNN does well classifying waveform pairs with and without overlap.

It also does fairly well with overlapping waveforms, correctly classifying 65% and 100% of matching and non-matching waveform pairs with overlap.

All false negatives are found to involve at least one overlapping waveform.

Application of a Paired Neural Network to Aftershock Identification

Andrea Conley, Brendan Donohoe, Benjamin Greene

RESULTS

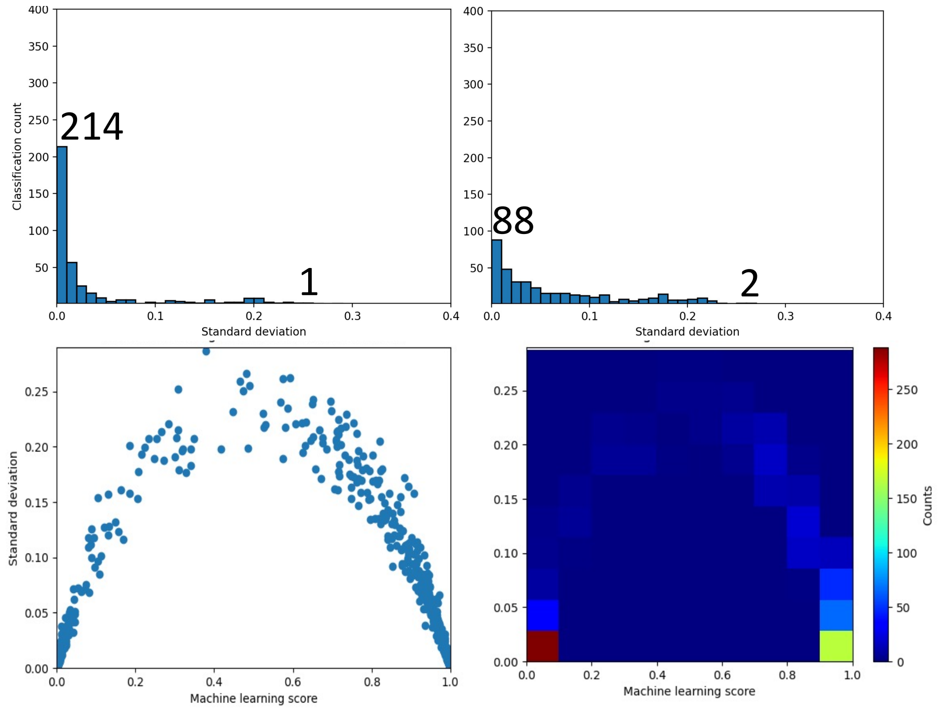

Match Non-Match Match Non-Match

Standard deviations (STD) of PNN scores are shown for known matches (top-left) and known non-matches (top-right).

Standard deviation values are found to range from ~0.0 to 0.25.

- In both cases, PNN scores for known matches and known non-matches tend to have STDs < 0.1.
- Thus, STD < 0.1 = low uncertainty and STD ≥ 0.1 = high uncertainty.

STD versus PNN score shown as a scatter plot (bottom-right) and 2D histogram (bottom-left).

- STDs < 0.1 appear to occur near PNN scores = 0.0 (clearly matching) or 1.0 (clearly non-matching).
- Thus, STDs do not provide new info over the original PNN score.

Application of a Paired Neural Network to Aftershock Identification

Andrea Conley, Brendan Donohoe, Benjamin Greene

CONCLUSIONS

A PNN can classify waveform pairs (our proxy for aftershocks) with high accuracy.

The PNN outperforms CC when classifying matching waveform pairs.

- 81% of matching pairs are clearly classified by the PNN vs only 35% for CC.
- 19.5% of matching, overlapping waveform pairs (proxies for waveforms with multiple aftershocks) are clearly classified by the PNN vs only 2% for CC.

The PNN and CC perform equally well classifying non-matching waveform pairs.

- 58% of non-matching pairs are clearly classified by the PNN vs 59% for CC.
- 23.7% of non-matching overlapping waveform pairs are clearly classified by the PNN vs 31.8% for CC.

The PNN outperforms CC in the ROC curve corner.

- The PNN has a higher TPR at all values beginning at 87.5%.
- At a 90% TPR, we find that the PNN better avoids false positives, achieving a 0% FPR versus a 3% rate for CC.

Low PNN uncertainties (< 0.1) are tied to clear matching, non-matching PNN scores. Higher uncertainties (≥ 0.1) are tied to intermediate scores.

In future work, we will investigate several remaining questions, including:

- Can our PNN model be further improved by training on more overlapping waveform data?
- How would the PNN perform when one waveform in the pair does not contain a signal?
- How would a model trained on constructed waveform data perform on real aftershock data?

We would like to acknowledge our colleagues, including but not limited to, Chris Young, Dylan Anderson, Amy Sundermier, Lisa Linville, Rigobert Tibi, Rob Porritt, Jorge Garcia, Elizabeth Berg, and Patrick Hammond, for their excellent feedback and recommendations on this work.