SAND2023-02131D

Comparative Study of the Performance of Seismic Waveform Denoising Methods

Rigobert Tibi, Christopher Young, and Robert Porritt
Sandia National Laboratories
Email: rtibi@sandia.gov

INTRODUCTION

Recorded seismic data are generally contaminated by various types of noise (cultural or natural). Despite significant progress in seismic data analysis, the separation of signal and noise remains a fundamental problem. In the seismology community, frequency filtering is the most commonly used method for noise suppression. Frequency filtering can be problematic when the signal of interest and noise occupy the same region in the frequency domain. We implemented and applied 3 classes of noise suppression methods using seismic data recorded at local to near-regional distances. The methods consist of approaches based on non-linear thresholding of continuous wavelet transforms (CWT), convolutional neural network (CNN) denoising, and frequency filtering (causal & acausal). The denoising approaches are compared by subjecting them to the same analyses and level of scrutiny using the same set of evaluation metrics.

METHODS

Nonlinear Thresholding of Continuous Wavelet Transform (CWT Denoising)

The wavelet transform of a continuous signal, *x(t)*, with respect to a wavelet function, *ψ(t)*, is defined as:

$$W(a,\tau) = \frac{1}{\sqrt{a}}\int_{-\infty}^{+\infty} x(t)\psi^*\left(\frac{t-\tau}{a}\right)dt \qquad (1)$$

a - Scale
τ - Time lag or location
ψ* - Complex conjugate of the wavelet function

As our mother wavelet, we used the Ricker wavelet

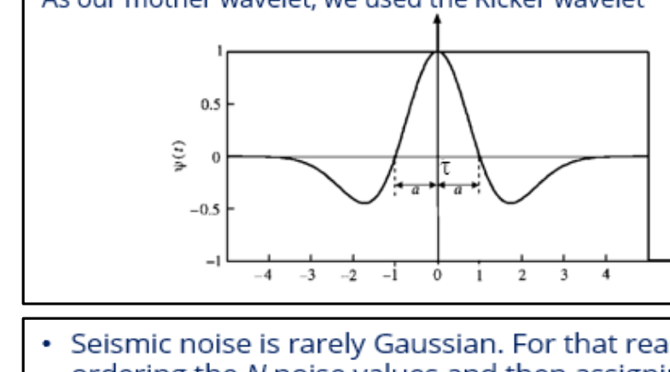

$$\psi(t) = (1-t^2)e^{-\frac{t^2}{2}} \qquad (2)$$

Key Aspects of the Denoising Approach Based on the Thresholding of CWTs:

- Noise is assumed to be stationary throughout the waveform
- Pre-event window is used to estimate the scale dependent (non-linear) threshold

Soft-thresholding (≠ hard thresholding):

The thresholded wavelet coefficients, $\widetilde{W}(a,\tau)$ are defined as:

$$\widetilde{W}(a,\tau) = \begin{cases} sign[W(a,\tau)](|W(a,\tau)| - \beta(a)) \text{ if } |W(a,\tau)| \geq \beta(a), \\ 0 \text{ otherwise} \end{cases} \qquad (3)$$

in which

$$sign[W(a,\tau)] = \frac{W(a,\tau)}{|W(a,\tau)|} \qquad (4)$$

the threshold, $\beta(a) = \text{mean}(|W(a,\tau)|) + c\ \text{stdv}(|W(a,\tau)|) \qquad (5)$

- Assuming noise coefficients follow Gaussian distribution (Donoho & Johnstone, 1994, DONO):

$$c = \sqrt{2\log_{10} N}, \qquad (6)$$

with *N* being the number of noise samples at each scale.

- Seismic noise is rarely Gaussian. For that reason, Langston & Mousavi (2019) proposed ordering the *N* noise values and then assigning a probability jump of 1/*N* when a value is attained.

$$\beta(a) = ECDF_a^{-1}(P = 0.99), \qquad (7)$$

in which $ECDF_a^{-1}$ is the inverse empirical cumulative distribution function.

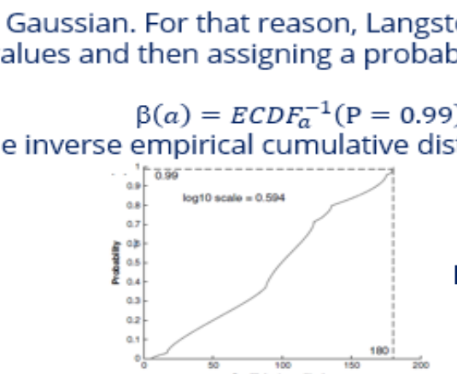

Langston & Mousavi (2019)

Denoised (thresholded) waveform, $\tilde{x}(t)$, is estimated using the inverse transform as:

$$\tilde{x}(t) = \frac{1}{c}\int_0^{+\infty}\int_{-\infty}^{+\infty}\frac{1}{\sqrt{a}}\widetilde{W}(a,\tau)\psi\left(\frac{t-\tau}{a}\right)\frac{da\,d\tau}{a^2}, \qquad (8)$$

where

$$C = \int_0^{+\infty}\frac{\hat{\Psi}^*(\omega)\hat{\Psi}(\omega)}{\omega}d\omega, \qquad (9)$$

in which $\hat{\psi}(\omega)$ is the Fourier transform of $\psi(t)$.

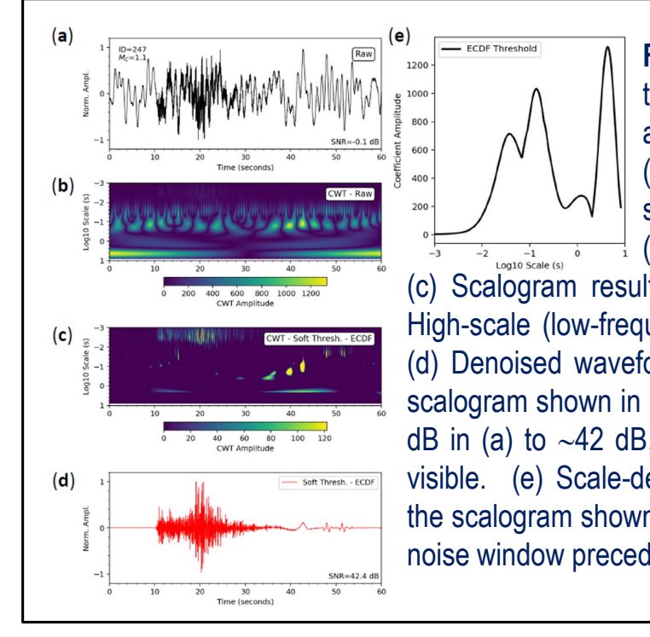

Fig. 1: (a) Example of a noisy waveform in which the signal is almost completely buried in high-amplitude noise (SNR ≈ 0). (b) Scalogram (amplitude as a function of wavelet lag time and scale) showing the continuous wavelet transform (CWT) of the waveform displayed in (a). (c) Scalogram resulting from thresholding the CWT shown in (b). High-scale (low-frequency) features have been effectively removed. (d) Denoised waveform obtained after the inverse transform of the scalogram shown in (c). The SNR has improved significantly from ~0 dB in (a) to ~42 dB, allowing for the P and S arrivals to be clearly visible. (e) Scale-dependent ECDF threshold used in thresholding the scalogram shown in (b). The threshold was estimated using a 9-s noise window preceding the P phase arrival.

Deep Learning Denoising (CNN Denoising)

The approach uses a trained deep convolutional neural network (CNN) model to decompose an input waveform into signal of interest and noise (Tibi et al., 2021).

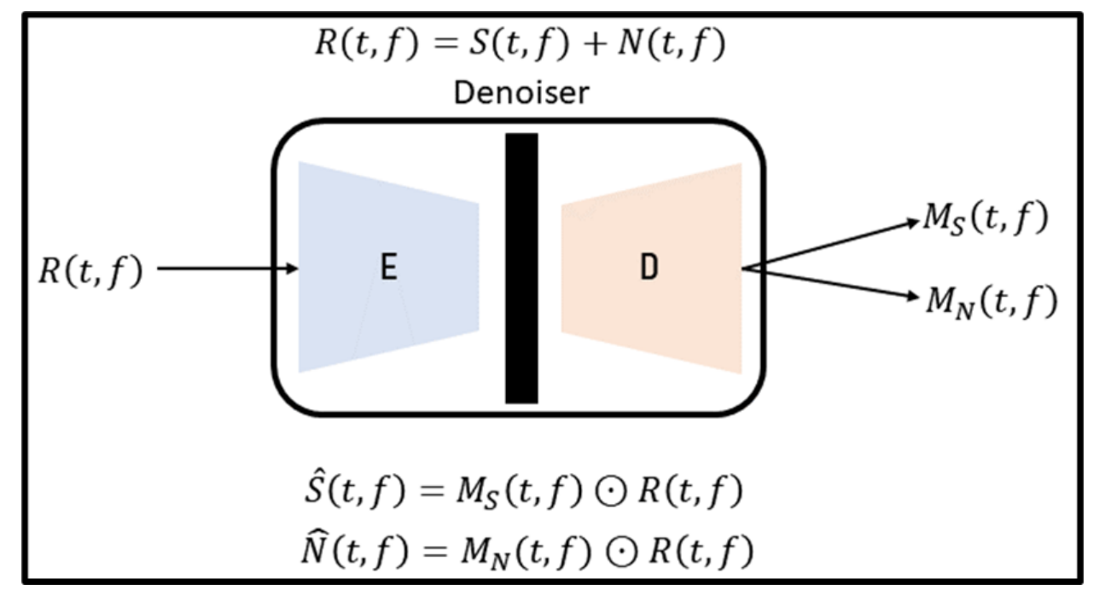

Fig. 2: For an input $R(t,f)$, the network provides a signal mask ($M_S(t,f)$) and a noise mask ($M_N(t,f)$). The estimated 'clean' signal ($\hat{S}(t,f)$) is obtained by multiplying $M_S(t,f)$ with $R(t,f)$; and the estimated noise ($\hat{N}(t,f)$) is obtained by multiplying $M_N(t,f)$ with $R(t,f)$.

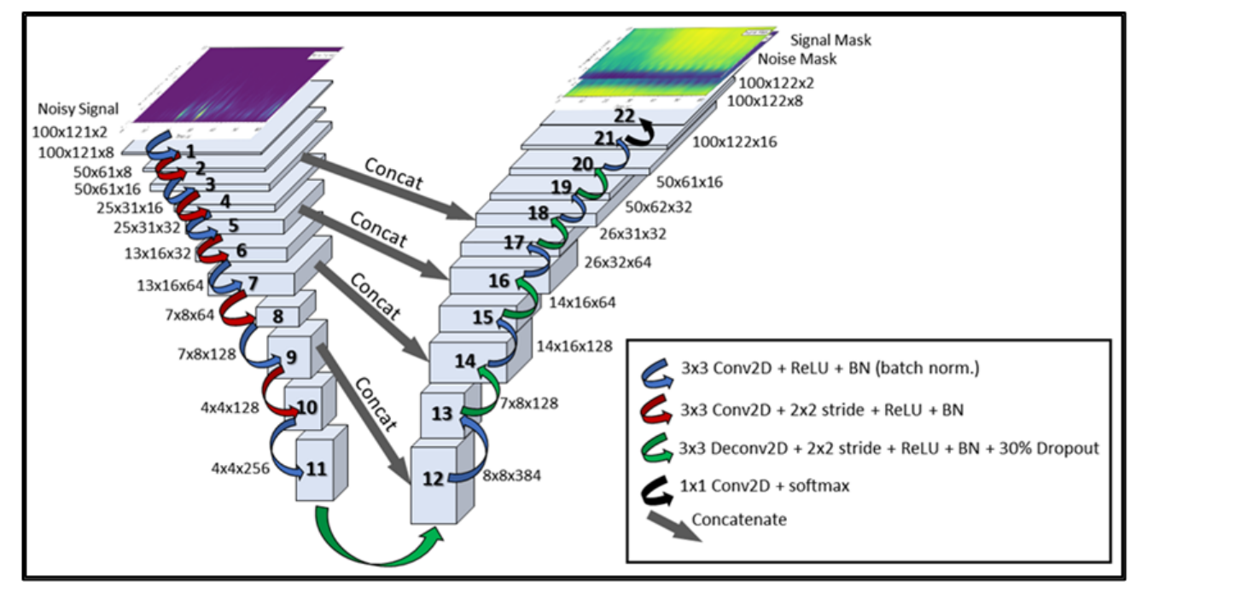

Fig. 3: The network consists of 20 hidden layers. Half of the layers make up the encoder, and the other half the decoder.

- The 'clean' signal dataset consists of 3,188 high-SNR Z-comp waveforms (60-sec long and filtered with BP 1–20 Hz) from local and near-regional events recorded at station **BRPU** of the University of Utah Seismograph Station (UUSS) network.
- The noise dataset contains 15,426 waveforms from various noise sources and various stations of the network.
- The 2 datasets are randomly divided into training, validation, and test sets using the '70-15-15' partitioning convention.
- Noisy waveforms are constructed by summing each 'clean' signal waveform and a randomly selected noise waveform. This was repeated 20 times for each set, resulting in:
  - 44,620 waveforms for the training set,
  - 9,580 waveforms for the validation set, and
  - 9,560 waveforms for the test set.
- The validation set is used in tuning the network hyperparameters, while the test set is used to assess network performance.

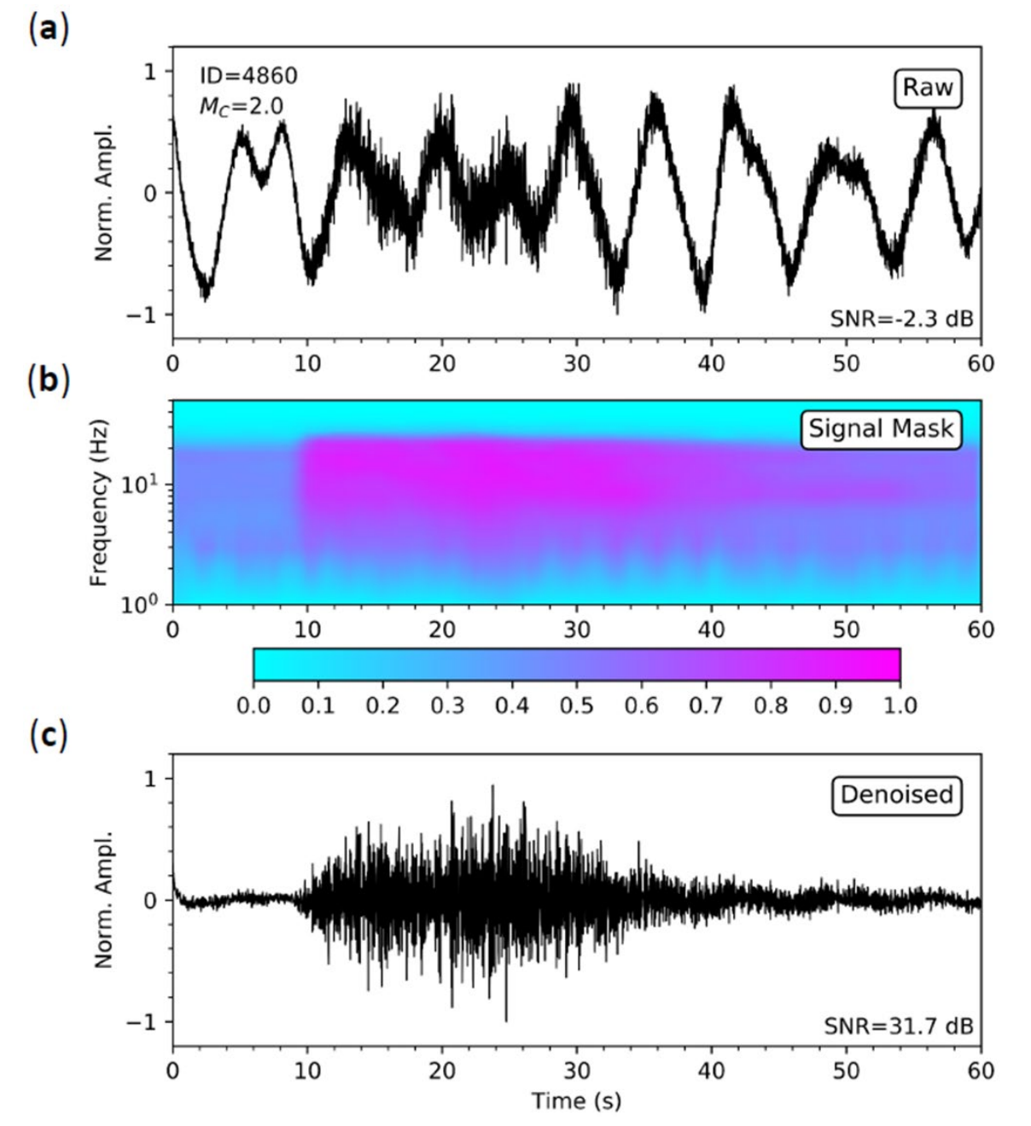

Fig. 4: (a) Example of noisy waveform. (b) Mask operator provided by the CNN network, which was used to denoised the seismogram in (a). (c) Denoised waveform. Note the improvement in SNR value from −2.3 dB for the raw seismogram in (a) to 31.7 dB for the denoised waveform in (c).

Frequency Filtering

- We used 4-pole Butterworth-bandpass filter, as implemented in Obspy (Beyreuther et al., 2010).
- Filter passband of 2–15 Hz, appropriate for the local and near-regional data used.
- To investigate the effect of filtering-related phase shift, we applied both causal and acausal (zero-phase) filter.

EVALUATION METRICS AND DATA

- Correlation Coefficient (CC) from non-zero lag cross correlation
  - Measures the similarity between the recovered waveform and the ground truth (GT)
- Signal-to-Noise Ratio (SNR in dB)
  - Using 9-sec window for both signal and noise

$$SNR = 20\log_{10}\frac{A_S}{A_N} \qquad (10)$$

- Signal-to-Distortion Ratio (SDR in dB)
  - Measures the amplitude distortion with respect to GT

$$SDR = 10\log_{10}\frac{\|W_{GT}\|^2}{\|\widehat{W}-W_{GT}\|^2} \qquad (11)$$

W_{GT} - Ground truth waveform; $\widehat{W}$ - Recovered (denoised) waveform, corrected for time shift

- Phase change (ϕ in radians)

$$\phi = 2\pi f\delta t \quad (12)$$

δt – Estimated time shift in seconds; f – Frequency set to 15 Hz (high-cut of chosen BP filter)

- We used a dataset of 4780 constructed noisy waveforms for which the underlying signal and noise waveforms are perfectly known.
- The component pure signal waveforms consists of local and near-regional recordings (distances 10.2–594.1 km) from 478 events (M_C 0.3–2) at stations of UUSS network, the same region and distance range for which the CNN denoiser was trained.
- None of these waveforms were part of the data used in the CNN model training or validation.
- By using constructed data, the characteristics of the components (signal and noise) that make up each noisy seismogram are known perfectly and can be compared against after denoising to assess the degree of fidelity at which the signal components have been recovered.

RESULTS

Effect of Input Seismogram Quality

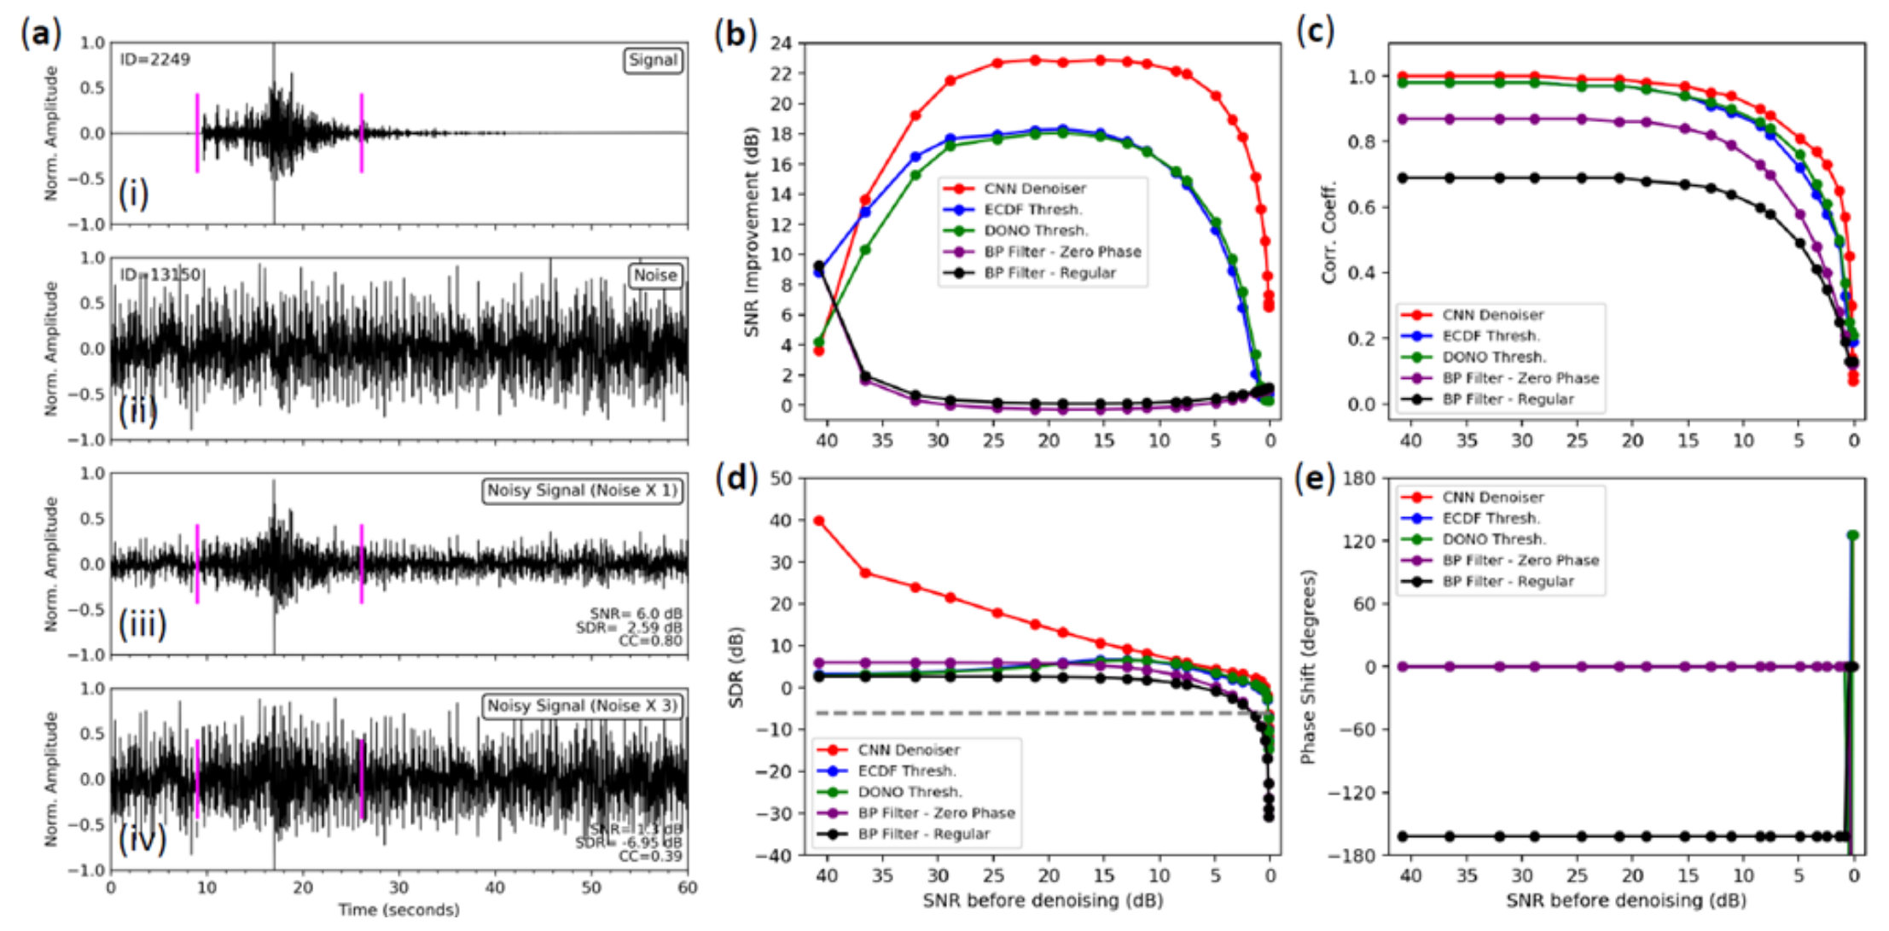

Fig. 5:

- For frequency filtering, the SNR gain after processing decreases quickly with decreasing input SNR.
- Deep Denoiser is capable of denoising a waveform with a SNR floor of approx. 0 dB.
- Causal filtering is associated with significant changes in waveform shape (CC of ~ 0.7).
- CNN denoising has unrivaled capability of conserving the amplitude information at input SNRs > 14 dB.
- In contrast to causal filtering, zero-phase filtering and the other methods do not result in phase change.

Processing Throughput

Fig. 6:

- Processing on GPU-accelerated machine using a single GPU
- For CNN denoiser, model training time not included (upfront cost)
- **Throughput:** Length of data processed in unit of time
- Average throughputs:
  - For frequency filtering: ~140 min of data per sec
  - CNN denoising: ~41 min of data per sec
  - CWT denoising: ~0.1 min of data per sec
- CWT computationally more involved than the Fourier transform and/or simple model interrogation in frequency filtering and CNN denoising

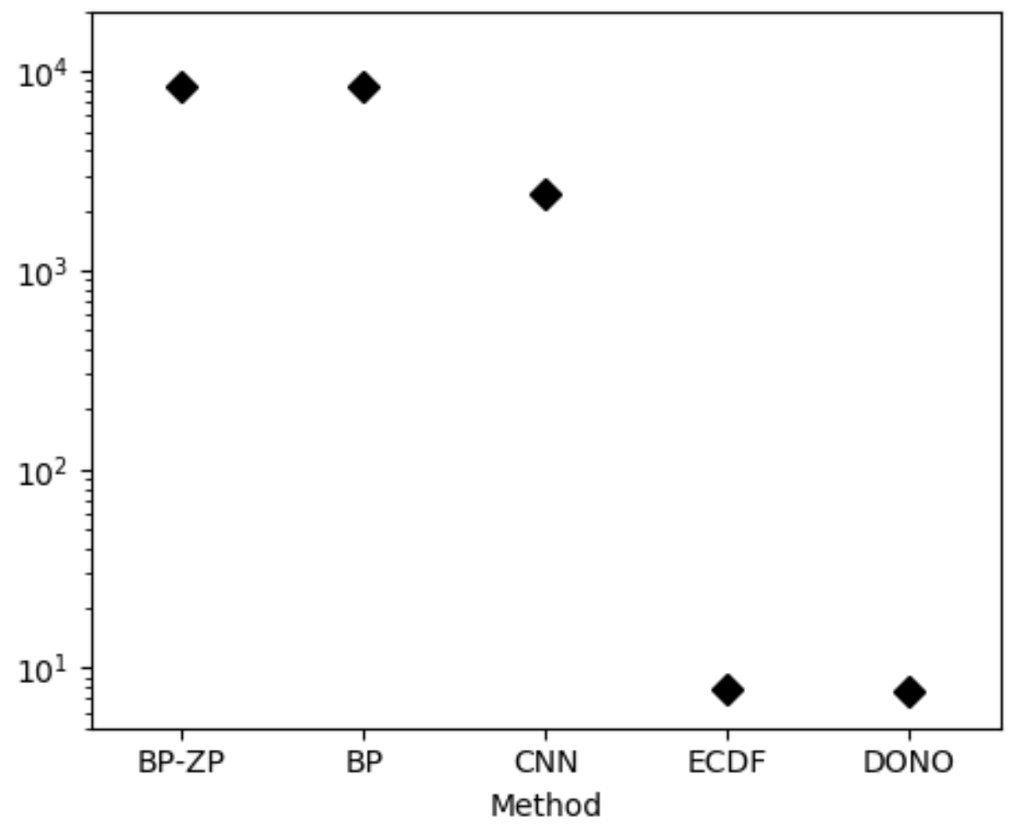

Improvement in Seismogram Quality

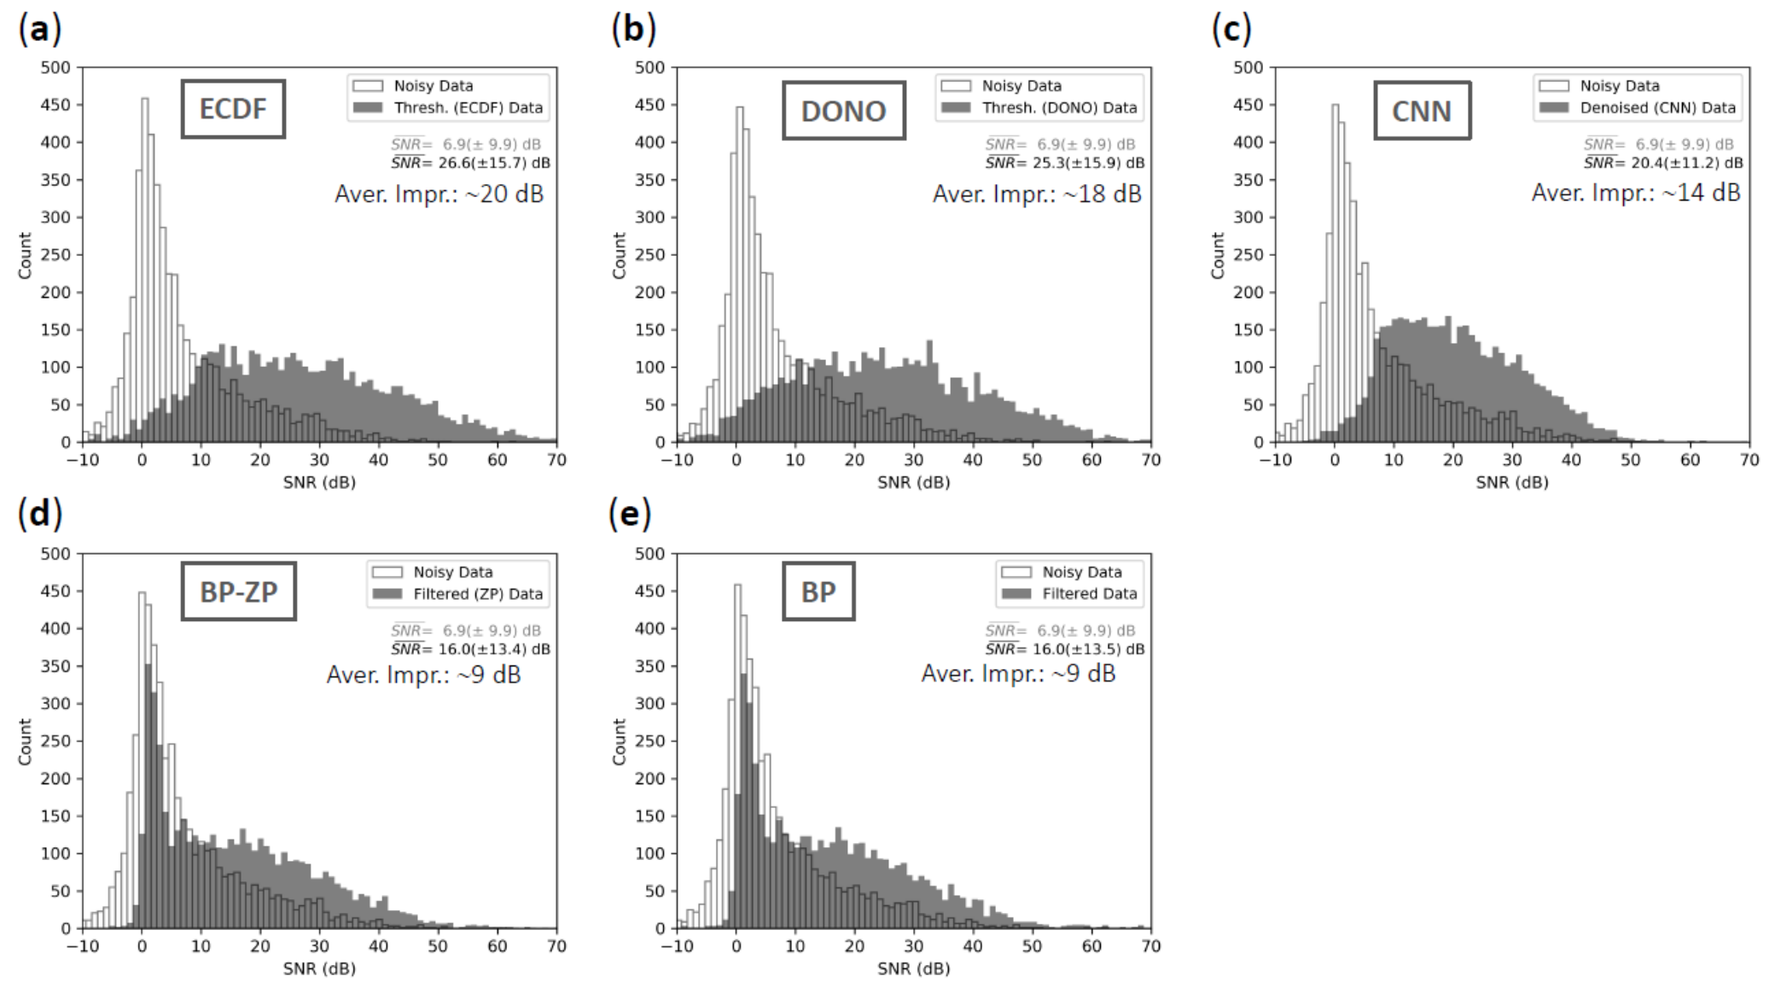

Fig. 7: The highest average improvement in SNR (~20 dB) is achieved by ECDF, followed by DONO (~18 dB) and CNN denoising (~14 dB). With only ~9 dB in average SNR improvement, frequency filtering is outperformed by the other methods. When the condition of stationarity of noise is met, CWT thresholding is a very effective method for noise suppression. However, the caveat is that the improvement in SNR comes at a high cost in terms of distortion of amplitude information.

Degrees of Fidelity to Ground Truth Waveforms

Fig. 8: For CWT thresholding, the average CC value is 0.8, suggesting that the denoised waveforms are sufficiently similar to the GTs. As implied by the low value for the average CC (~0.5), waveform shape undergoes significant changes during causal frequency filtering. For CNN denoising, we observed an average CC of ~0.9, indicating that this approach clearly outperforms CWT thresholding and frequency filtering in terms of preservation of waveform shape.

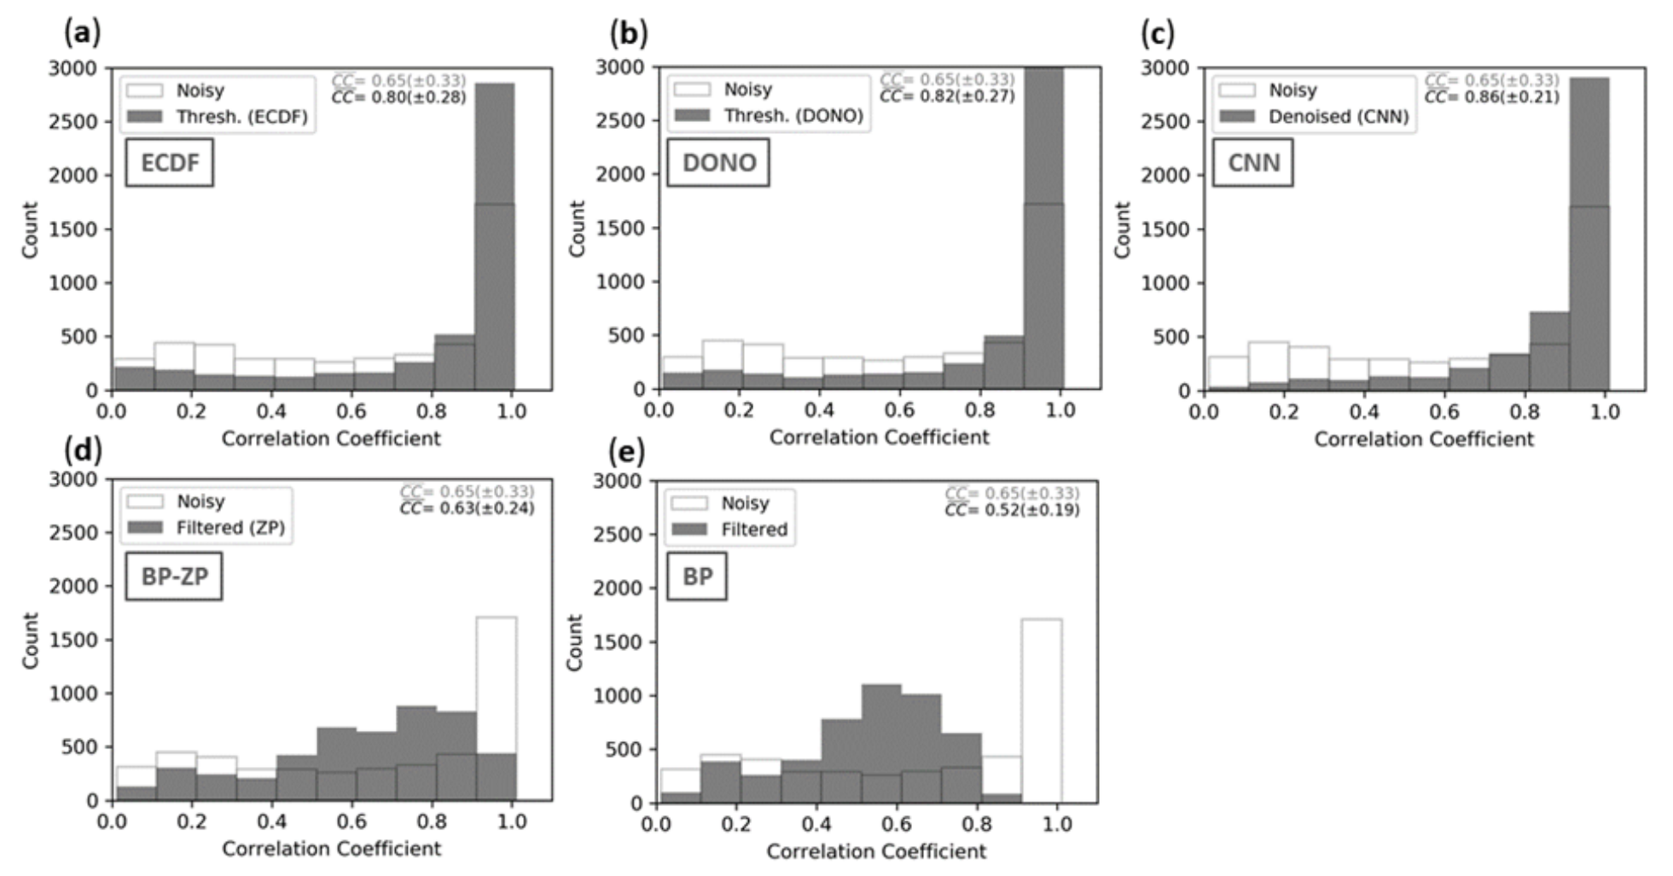

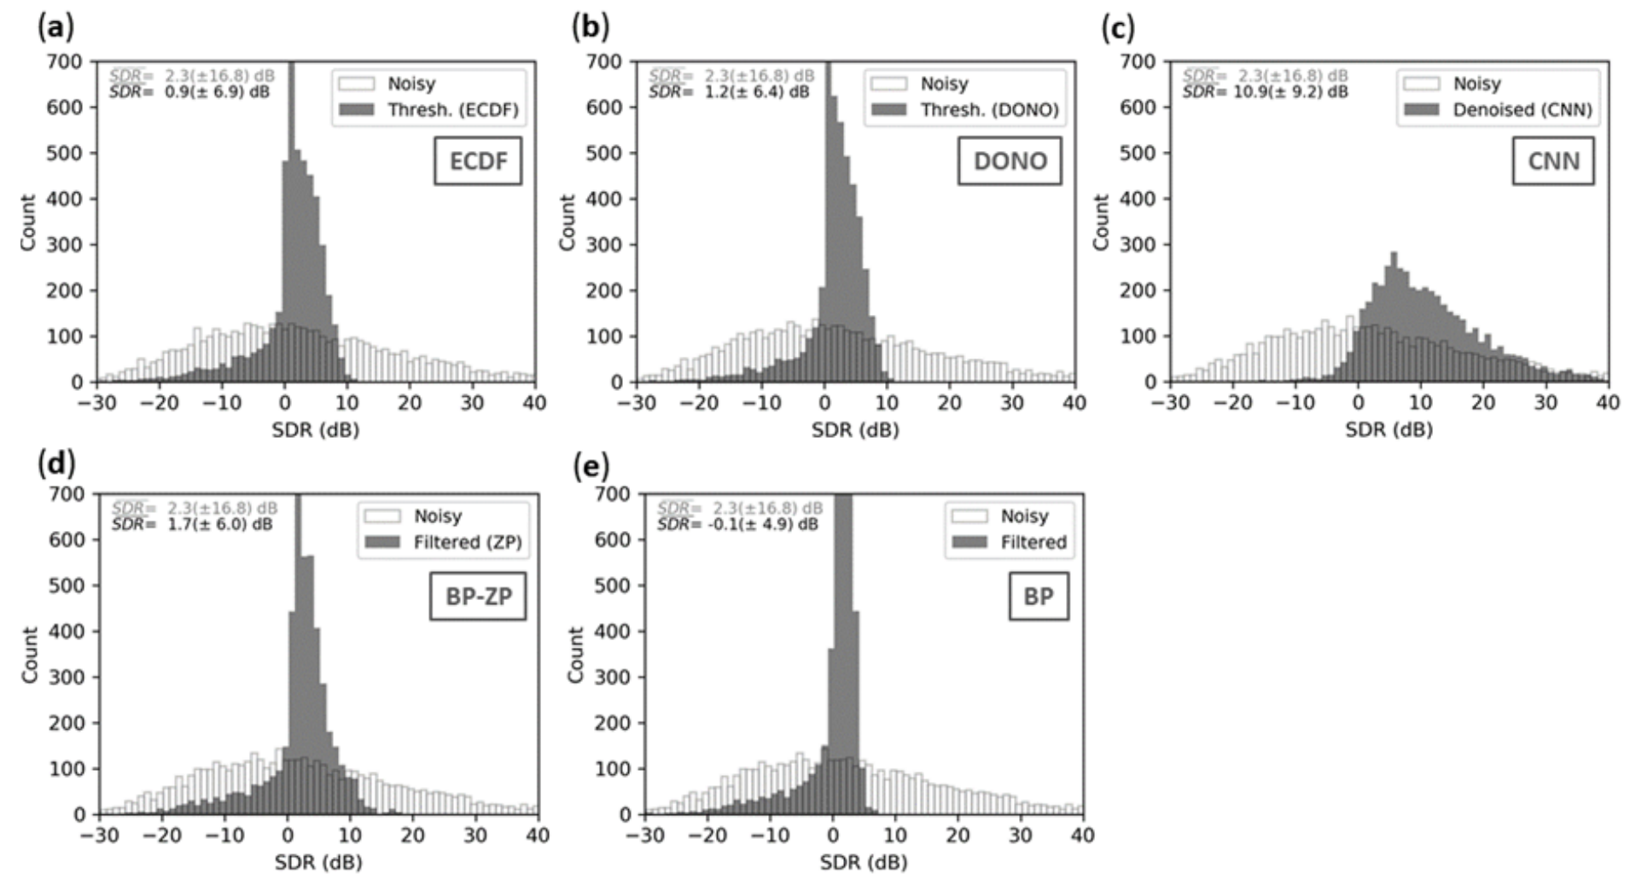

Fig. 9: The average SDR values of only ~1 dB for the CWT denoisers are a clear indication that this class of denoisers is associated with significant distortion in amplitudes. This is also the case for frequency filtering, which shows average SDR values in the range of about 0 to 2 dB. With an SDR value of ~11 dB, CNN denoiser clearly outperforms the other classes of denoisers.

Fig. 10: The estimated average phase shift of 1.1–1.4° for CWT thresholding, CNN denoising, and zero-phase filtering indicates that little to no phase change occurs for these approaches, a necessary condition for use in array analysis. This contrasts to causal frequency filtering that shows an average phase change of –166°. With an estimated standard deviation of about ±91° for that parameter, causal frequency filtering also shows the largest variability across the 4780 seismograms we processed.

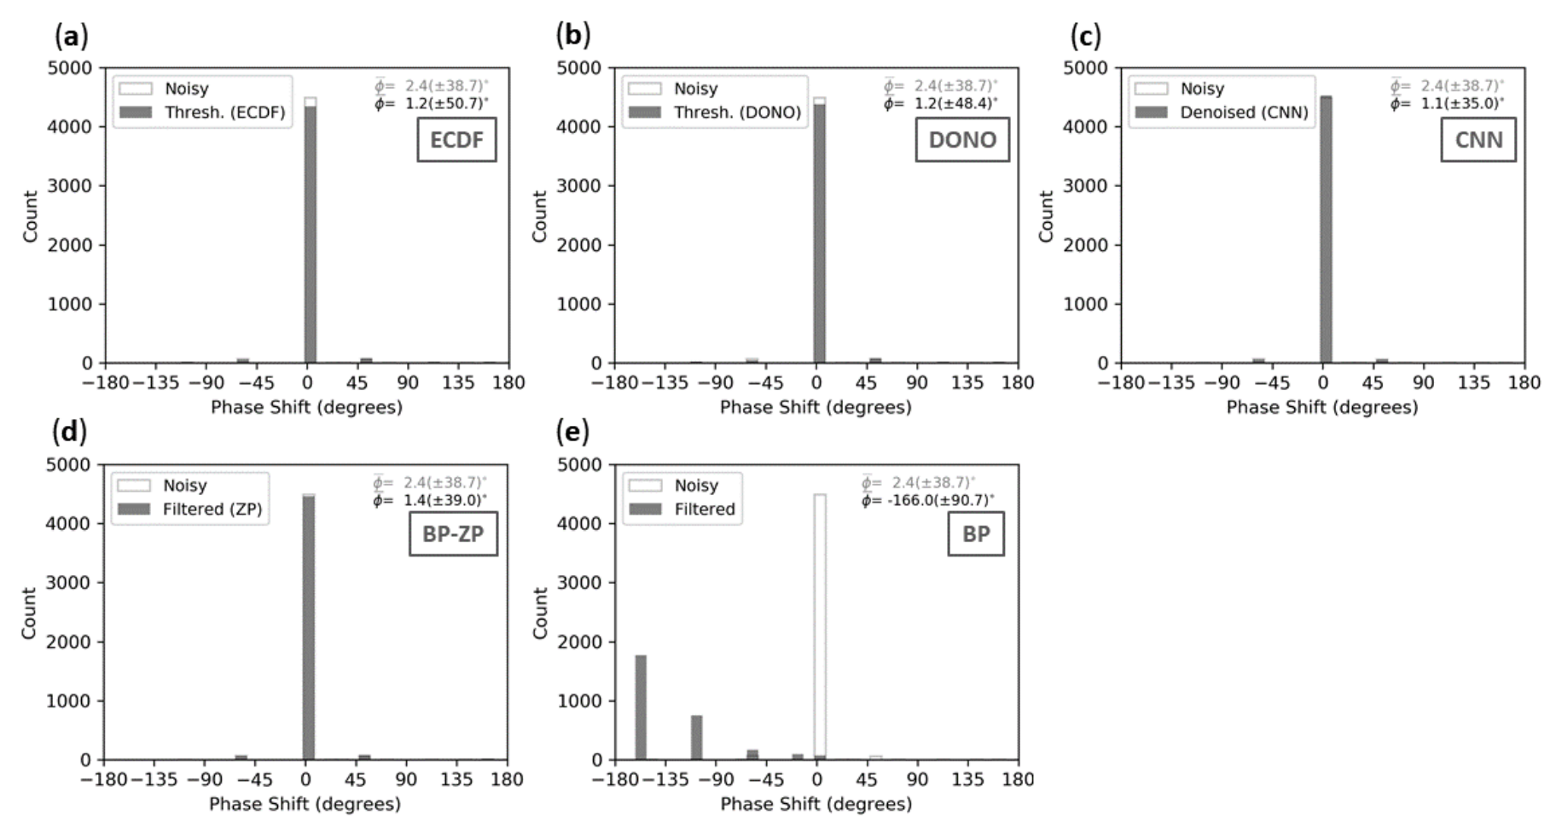

Onset-Time Determination

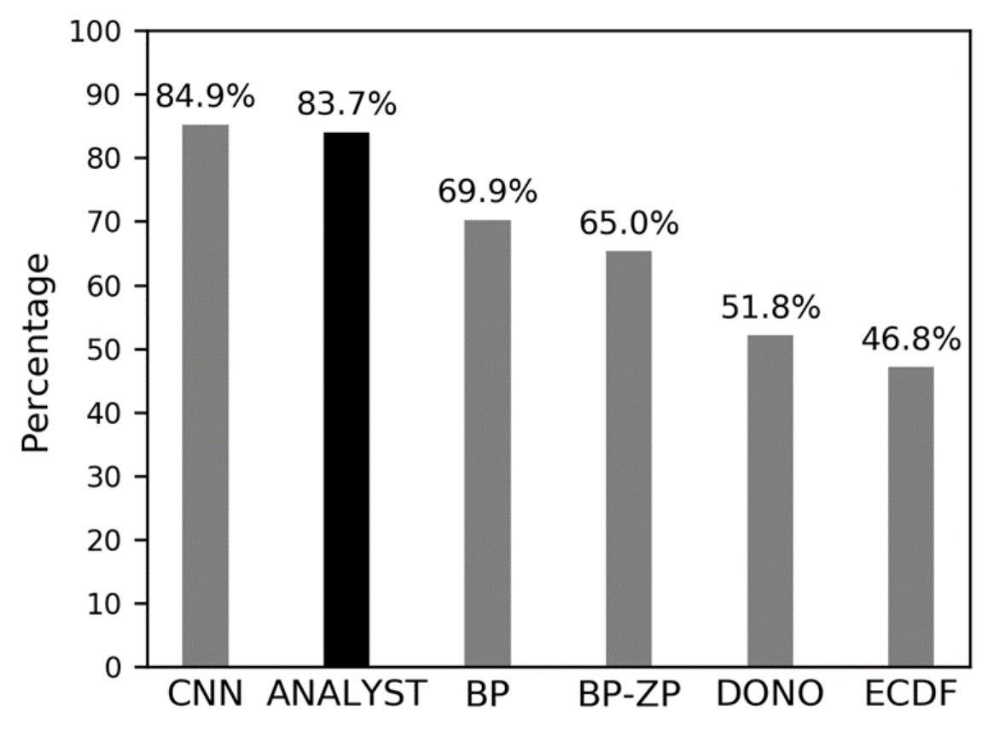

Fig. 11: The bar plot shows for each denoising approach (gray) the proportion in of the number of waveforms for which an expert-analyst was able to determine the onset time of the first arrival. The black bar labeled "ANALYST", represents that proportion after applying the analyst-determined best filters for each of the raw seismograms, using an iterative procedure, the same way data are processed in an operational setting. The CNN denoising allows more picks to be made compared with frequency filtering and CWT denoising, and is on par with the expert-analyst's best filters. The CWT techniques were more likely to introduce artifacts that made the waveforms unusable. This is reflected in the low proportion (~47–52%) of picks that our expert-analyst was able to make for these methods.

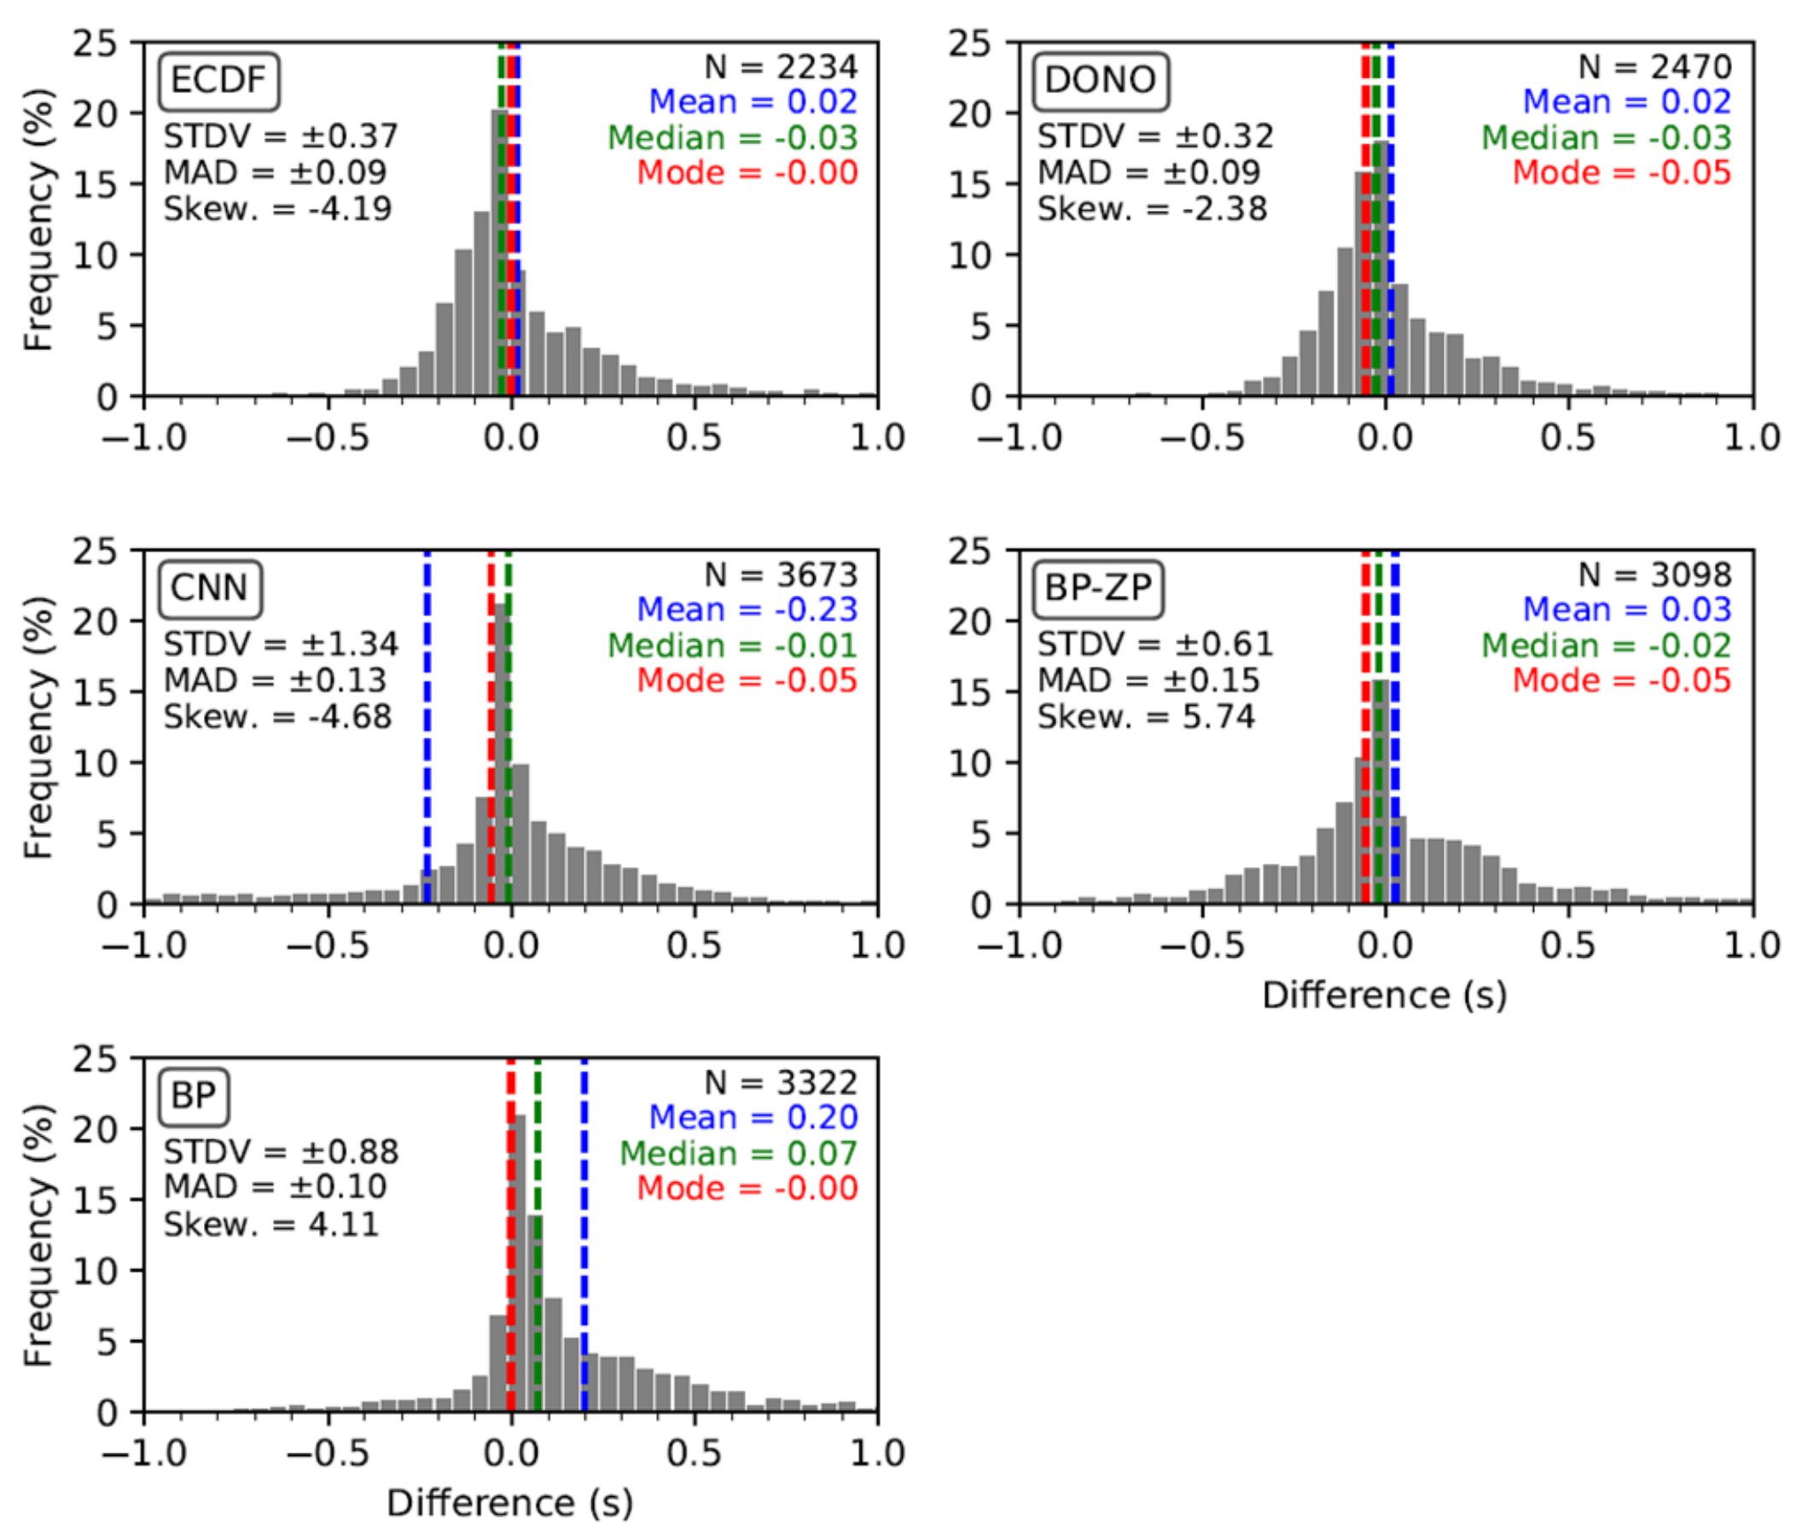

Fig. 12: Distributions of the differences between the onset times determined after applying the expert-analyst's best filters on the raw waveforms and the picks made after denoising the same waveforms using each of the methods under evaluation. As implied by the larger mean differences (0.20–0.23 seconds in absolute sense) and standard deviations (±0.88–1.34), which contrasts to the median values of ~0 s, the estimated differences for CNN denoising and causal bandpass filtering contain non-negligible numbers of outliers. Most outliers for CNN denoising result from too early picking, while for causal filtering they result from too late picking, as indicated by their negative and positive skewness values, respectively. However, for these and the other approaches the modes of the distributions are equal or close to 0, suggesting that overall, most of the picks for each of the methods are consistent with those for the expert-analyst's best filters.

CONCLUSIONS AND IMPLICATIONS

- We found that for frequency filtering, the improvement in SNR decreases quickly with decreasing input SNR and below an input SNR of ~32 dB the improvement is marginal and nearly constant.
- In contrast, for CWT and CNN denoising, the SNR gains are low at high input SNR before increasing to reach the top of the plateaus corresponding to gains of about 18 and 23 dB, respectively.
- For most of the input SNR range, the quality of the output waveform for CNN denoising in terms of output SNR and amplitudes is superior to other approaches. In addition, CNN denoising is the only approach capable of denoising seismograms with SNR floor as low as 0 dB.
- On average CWT and CNN denoising, and bandpass filtering improve the SNR by about 20, 14, and 9 dB, respectively. In terms of degree of fidelity for the denoised waveforms with respect to the GT seismograms, CNN denoising outperforms both CWT denoising and frequency filtering.
- Like zero-phase filtering, little to no phase shift occurs for CWT and CNN denoising, making these approaches suitable for use in array analysis. This contrasts with causal filtering that results in significant phase shifts.
- CNN denoising allows more picks to be made compared with frequency filtering or CWT denoising, and is on par with the expert-analyst's processing, which is consistent with the current operational procedure. The CWT techniques are more likely to introduce artifacts that made the waveforms unusable.

Examples of Analysis Purposes	CNN Denoising	CWT Thresholding	Zero-Phase Frequency Filtering 2–15 Hz	Causal Frequency Filtering 2–15 Hz
Improve SNR for signal detection	✓	✓	✓ if input is of sufficient SNR (≥ 6 dB)	✓ if input is of sufficient SNR (≥ 6 dB)
Exploit amplitude information (e.g., for magnitude, yield, or moment tensor estimation)	✓ most suitable approach	✗ significant amplitude distortion	✗ significant amplitude distortion	✗ significant amplitude distortion, changes in waveform shape & phase
Array Analysis	✓	✓	✓	✗ significant changes in phase

ACKNOWLEDGMENTS

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. The views expressed in the article do not necessarily represent the views of the U.S. Department of Energy or the United States Government. This research was funded by the National Nuclear Security Administration, Defense Nuclear Nonproliferation Research and Development (NNSA DNN R&D).

SAND2023-XXXX C