



A Self-Supervised Deep Learning Approach for Blind Denoising and Waveform Coherence Enhancement in Distributed Acoustic Sensing data

Martijn P.A. van den Ende, Itzhak Lior, Jean Paul Ampuero, Anthony Sladen, André Ferrari, Cédric Richard

► To cite this version:

Martijn P.A. van den Ende, Itzhak Lior, Jean Paul Ampuero, Anthony Sladen, André Ferrari, et al.. A Self-Supervised Deep Learning Approach for Blind Denoising and Waveform Coherence Enhancement in Distributed Acoustic Sensing data. IEEE Transactions on Neural Networks and Learning Systems, inPress, 10.1109/TNNLS.2021.3132832 . hal-03347303

HAL Id: hal-03347303

<https://hal.science/hal-03347303>

Submitted on 17 Sep 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

This manuscript has been submitted for publication in *IEEE Transactions on Neural Networks and Learning Systems*. Subsequent versions of this manuscript may have different content. If accepted, the final version of this manuscript will be available via the *Peer-reviewed Publication DOI* link printed on this webpage.

All the data, codes, and pre-trained Deep Learning models are publicly available from:

van den Ende M. P. A., Lior I., Ampuero J.-P., Sladen A., Ferrari A., Richard C. (2021):
*A Self-Supervised Deep Learning Approach for Blind Denoising and Waveform
Coherence Enhancement in Distributed Acoustic Sensing data*,
<https://doi.org/10.6084/m9.figshare.14152277>

Comments and questions are welcomed. Please contact the first author (Martijn van den Ende) via email (martijn.vandenende@oca.eu) or via Twitter (@martijnende)

A Self-Supervised Deep Learning Approach for Blind Denoising and Waveform Coherence Enhancement in Distributed Acoustic Sensing data

Martijn van den Ende, Itzhak Lior, Jean-Paul Ampuero, Anthony Sladen, André Ferrari, and Cédric Richard

Abstract—Fibre-optic Distributed Acoustic Sensing (DAS) is an emerging technology for vibration measurements with numerous applications in seismic signal analysis, including microseismicity detection, ambient noise tomography, earthquake source characterisation, and active source seismology. Using laser-pulse techniques, DAS turns (commercial) fibre-optic cables into seismic arrays with a spatial sampling density of the order of metres and a time sampling rate up to one thousand Hertz. The versatility of DAS enables dense instrumentation of traditionally inaccessible domains, such as urban, glaciated, and submarine environments. This in turn opens up novel applications such as traffic density monitoring and maritime vessel tracking. However, these new environments also introduce new challenges in handling various types of recorded noise, impeding the application of traditional data analysis workflows. In order to tackle the challenges posed by noise, new denoising techniques need to be explored that are tailored to DAS. In this work, we propose a Deep Learning approach that leverages the spatial density of DAS measurements to remove spatially incoherent noise with unknown characteristics. This approach is entirely self-supervised, so no noise-free ground truth is required, and it makes no assumptions regarding the noise characteristics other than that it is spatio-temporally incoherent. We apply our approach to both synthetic and real-world DAS data to demonstrate its excellent performance, even when the signals of interest are well below the noise level. Our proposed methods can be readily incorporated into conventional data processing workflows to facilitate subsequent seismological analyses.

Index Terms—Distributed Acoustic Sensing, self-supervised Deep Learning, blind denoising, waveform coherence

I. INTRODUCTION

Distributed Acoustic Sensing (DAS [1]) is a vibration sensing technology that turns a fibre-optic cable into an array of single-component seismometers. By connecting a DAS interrogator at one end of a (potentially) tens-of-kilometres-long optical fibre, time-series measurements of ground deformation can be made at metre-spaced intervals along the cable. This emerging technology has massively extended our range of capabilities for seismic monitoring, enabling the deployment of dense seismic arrays in areas that were previously inaccessible, such as urban [2], [3] or submarine [4], [5] locations. Fibre-optic cables can be deployed on rugged terrain on land or underwater, are temperature-robust, and they are sensed *ex-situ* (i.e. at one end of the fibre) without the need for an electrical

current to run across the cable. Moreover, DAS can be applied to standard commercial fibre-optic cables [6], [7], relieving the need for costly deployment campaigns for scientific instrumentation. With a typical spatial sampling density of 10 m and a time sampling rate of 100-1000 Hz, DAS arrays have the potential to complement or even replace existing seismic arrays [8]. This new approach of seismic data collection also provides new perspectives and challenges with regard to nuisance signals (noise) that originate from instrumental, electronic, anthropogenic, or environmental sources. Since one often has little control on the exact placement of the cable, deployments are typically not optimised for the recording of specific signals of interest, enhancing the relative contribution of noise to the recordings.

Traditional noise filtering techniques are deeply embedded in the workflow of seismological data analysis. Particularly when the signal of interest occupies a frequency band that is distinct from that of a noise source, spectral methods are highly efficient in recovering noise-suppressed signal reconstructions. Signal denoising becomes much more challenging when the signal of interest and the noise source share a common frequency band, in which case additional knowledge about the noise or the signal needs to be incorporated. One possible prior that can be exploited in seismic signal analysis is the notion that the signal of interest (e.g. an earthquake waveform) is spatio-temporally coherent, while the noise may be uncorrelated in space and/or time. Exploiting the spatial sampling density of DAS, such coherent signals can be distinguished from incoherent noise in DAS recordings.

Over the last decade, machine learning methods have been successfully applied to tackle geophysical problems and assist in laborious tasks such as earthquake detection [9], [10] and phase arrival picking [11], [12]. Given that DAS produces large volumes of data (of the order of 1 terabyte per fibre per day), it is particularly well suited for data-driven Deep Learning methods. It is therefore expected that Deep Learning can expedite various analytical workflows and accelerate the development of DAS as a low-cost seismological monitoring tool. In this contribution, we explore a Deep Learning blind denoising method that optimally leverages the spatio-temporal density of DAS recordings. Specifically, we assess the potential of this method to separate earthquake signals from the spatially incoherent background noise, which may benefit numerous seismological analysis techniques to be applied to the denoised DAS data.

M. van den Ende, I. Lior, J.-P. Ampuero, and A. Sladen are with the Université Côte d’Azur, IRD, CNRS, Observatoire de la Côte d’Azur, Géoazur, France. E-mail: martijn.vandenende@geoazur.unice.fr

M. van den Ende, A. Ferrari, and C. Richard are with the Université Côte d’Azur, OCA, UMR Lagrange, France

II. OUTLINE OF THIS PAPER

This paper is organised as follows: we begin by establishing a framework of related research in Section III, within which we define the scope of our work. Next, we detail the main concept of J -invariance underlying our proposed method (Section IV-A), and the Deep Learning model architecture and training procedures (Section IV-B). We then describe the DAS and synthetic data to be analysed, and the procedure of pre-processing in Sections IV-C and IV-D, followed by the pretraining procedure in Section IV-E. We subsequently evaluate the model performance on the synthetic data in Section V, followed by an analysis of the real-world DAS data in Section VI. We end the paper with a comparison to traditional image denoising methods, a discussion of various seismological applications, the limitations of the proposed method, and potential extensions to non-Euclidean data types (Section VII).

III. RELATED WORK

In recent years, the domain of image denoising and restoration has seen major advancements spurred by machine learning, in particular Deep Learning. In a recent review, [13] presents a summary of 200 Deep Learning studies published over the last 5 years with a focus on image denoising. In seismology, Deep Learning has likewise been utilised as a tool for denoising seismic data. But unlike for most image denoising applications, the noise recorded by seismometers often does not follow a Gaussian distribution characteristic of instrumental self-noise. Moreover, the noise variance is not stationary in time, nor is it homogeneous in space. Ambient noise sources like wind, rainfall, ocean waves, cars, and trains all contribute to recorded seismic signals of non-tectonic or volcanic origins [14], [15]. Denoising methods applied to seismic data therefore need to be robust to a wide range of nuisance signal characteristics. This renders state-of-the-art image denoising methods that assume a homogeneous, constant variance of Gaussian noise (e.g. [16], [17]) ineffective.

In seismology, various learning algorithms have been used to enhance noise removal capabilities, particularly in seismic reflection studies, e.g. [18]–[22]. The majority of these studies employ a form of compression (dictionary wavelet learning or auto-encoding) to remove uninformative (incompressible) noise from compressible or sparse signals. Inherent to lossy compression methods, a trade-off needs to be considered between the degree of denoising on the one hand, and reconstruction fidelity on the other.

For single broadband station recordings of earthquakes, DeepDenoiser [23] has demonstrated excellent performance in separating earthquake signals from a variety of noise sources, even when the noise is non-stationary or when its frequency band overlaps with that of the signal. At the training stage, this supervised method requires an a-priori, clean signal that is superimposed onto empirical noise recordings to yield a training sample. The model performance is then evaluated as the ℓ_2 -norm of the difference between the clean signal and the model output. In cases where a ground-truth (clean) signal is not available, training of the DeepDenoiser is not




Fig. 1. Intuitive explanation of the concept of J -invariance. Even if the long-range coherent contents underneath the patch on the zebra are not accessible, they can be accurately interpolated based on their surroundings. On the other hand, the details of the contents within the red-bordered patch cannot be predicted, other than the average value (dark green). The output of a function applied to the green-bordered patch therefore does not depend on the contents of that patch, which is referred to as J -invariance. Photo attribution: Yathin S. Krishnappa, CC BY-SA 3.0, via Wikimedia Commons.

possible. For on-land fibre-optic DAS, an unsupervised noise clustering [24] approach has been successfully applied to isolate repeating data patterns originating from vehicle traffic. However, this study relies on the model's ability to learn meaningful representations of the data, and requires manual inspection of the learnt representations to identify the various signal and noise sources. While the approach of [24] is useful for filtering out spatio-temporally coherent noise, it is not known whether incoherent noise, which is the focus of the present study, can be suppressed using this approach.

Building on previous work, in particular [25], we propose a blind denoising method that does not make any assumptions regarding the characteristics of the noise, other than that it is spatio-temporally incoherent. Moreover, we utilise the dense spatial sampling density of DAS to separate incoherent noise from coherent signals, even when they share a common frequency band.

IV. METHODS

A. Concept of J -invariant filtering

We begin by detailing the general concept of J -invariant filtering that underlies our proposed method, following the work of [25]. To illustrate the intuition behind this approach, we consider an image featuring a coherent signal from which a small patch in its interior has been removed (the green-bordered patch in Fig. 1). If the signal of interest exhibits sufficiently long-range coherence (with respect to the size of the patch), then the signal contents within the patch can be accurately predicted. On the other hand, uncorrelated or short-range locally-correlated features that exist within the image (e.g. the underbrush in Fig. 1) are uninformative for predicting the contents of the removed patch. A learner (human

or artificial) faced with the task to recover the hidden data will therefore only be able to use coherent signals in the input data. Consequently, it can be said that the contents of the patch are not immediately required to perform a given action on the patch (which in the present case is the unbiased estimation of the signal). This approach has recently been by [25] for blind image denoising applications, circumventing the need of clean training data.

The authors in [25] propose to train an image denoiser g using a single noisy image y that results from an unknown clean image x such that $x = \mathbb{E}[y|x]$, where $\mathbb{E}[\cdot]$ denotes the expectation operator. The denoiser derivation relies on the assumption that g is \mathcal{J} -invariant.

Definition. Let \mathcal{J} be a partition of the feature space, and let $J \in \mathcal{J}$. We write z_J for z restricted to its features in J . We say that g is \mathcal{J} -invariant if, for all J in \mathcal{J} and for all z , $g(z)_J$ does not depend on the values of z_J .

This framework implies that, under independent noise assumption, the minimiser g^* of $\mathbb{E}\|g(y) - y\|^2$ over the space of \mathcal{J} -invariant functions verifies: $g^*(y)_J = \mathbb{E}[x_J|y_{J^c}]$ for all $J \in \mathcal{J}$ where J^c denotes the complement of J . This result, when compared to the optimal denoiser $\mathbb{E}[x_J|y]$, clearly shows the couplings between the independence of the noise, the spatial coherence of the clean image and the partition \mathcal{J} .

In practice, the set of \mathcal{J} -invariant functions g is explored with a neural network f_θ with parameters θ . The neural network f_θ is made \mathcal{J} -invariant by defining it as:

$$g(\cdot) = \sum_{J \in \mathcal{J}} \Pi_J(f_\theta(\Pi_{J^c}(\cdot))) \quad (1)$$

where $\Pi_A(z)$ is the projection operator that does not modify the values of the elements of z in A but sets the elements in A^c to zero (being $\mathbb{E}[z]$ in our case). Given that (1) implies $g(\Pi_{J^c}(\cdot)) = \Pi_J(f_\theta(\Pi_{J^c}(\cdot)))$, minimisation of $\|g(y) - y\|^2$ w.r.t. θ can be performed efficiently by training f_θ with a suitable learning objective.

In [25], the authors focus primarily on single-image denoising applications, with a brief exploration of multi-image denoising using Deep Learning architectures. In the present work, we apply the concept of \mathcal{J} -invariance to batches of DAS data (which are analogous to images). As we will demonstrate, performing the training on a sufficiently diverse set of DAS data enables direct application of the trained model on new data without retraining. In the following section we describe the neural network architecture and the procedure to enforce \mathcal{J} -invariance during the training stage on batches of DAS data.

B. Model architecture

Our denoising model architecture is based on the commonly used U-Net configuration [26] featuring 4 blocks of downsampling and convolutional layers in the encoder, and 4 blocks of up-sampling, concatenation, and convolutional layers in the decoder – see Fig. 2. We begin with two convolutional layers each with 4 filters and a stride of 1. Each of the 4 downsampling blocks features an anti-aliased downsampling layer [27] with a stride of 4 along the time axis (i.e. no downsampling

is performed along the DAS channel axis), followed by two convolutional layers with a number of filters that is doubled for each block (8, 16, 32, 64). These encoder operations are reversed in the decoder by first bilinear upsampling with a factor 4, concatenating the output of the diametrically opposed block, and two convolutional layers with a decreasing number of filters (32, 16, 8, 4). The output layer is a convolutional layer with a single filter and linear activation. All convolutional layers feature a kernel of size 3×5 (DAS channels \times time samples), Swish activation functions [28] (except for the last layer), orthogonal weight initialisation [29], and no additional regularisation (dropout, batch normalisation, etc.).

We create a mini-batch sample y_k consisting of 11 neighbouring DAS channels, corresponding with 192 m of cable length with a gauge length of 19.2 m, and 2048 time samples corresponding with 41 s of recordings at 50 Hz. We define J_k as an entire single DAS waveform, chosen at random from the set of 11 channels. To enforce \mathcal{J} -invariance in the model we apply the projection operation to get $u_k := \Pi_{J_k^c}(y_k)$, as well as to the model output $f_\theta(u_k)$ defining $v_k := \Pi_{J_k}(f_\theta(u_k))$. In accordance with the theory laid out in the previous section, we define the loss \mathcal{L} computed over a mini-batch $\{y_k\}$ as:

$$\mathcal{L}(\{y_k\}) = \frac{1}{|K|} \sum_{k \in K} \|v_k - \Pi_{J_k}(y_k)\|^2 \quad (2)$$

In this study the size of the mini-batch ($|K|$) is taken to be 32 samples. This loss function is minimised using the ADAM algorithm [30].

C. DAS data acquisition and processing

The DAS data were acquired between 18 and 25 April 2019 from two submarine dark optical fibres deployed offshore Methoni, south-west Greece (see Fig. 3). The first cable connects the EMSO Hellenic Arc node, an ocean observatory deployed in the Ionian Sea, with a land station in Methoni. It is managed by the Hellenic Centre of Marine Research (HCMR), and so we will refer to this cable as the HCMR cable. The second cable was commissioned for the NESTOR project (Neutrino Extended Submarine Telescope with Oceanographic Research), and likewise connects the submarine observatory to mainland Greece at Methoni. Both fibres were sensed with a commercial Febus A1 DAS interrogator, developed by Febus Optics. The first experiment was conducted on the HCMR cable from 18 to 19 April, immediately followed by the second experiment on the NESTOR cable from 19 to 25 April. During this period several regional earthquakes were recorded (see [31]), which will be the focus of this study.

For both experiments the gauge length and DAS channel spacing were set to 19.2 m, and the data were sampled at 6 ms and 5 ms intervals for HCMR and NESTOR, respectively. For the purpose of this study, we filter the data in a 1-10 Hz frequency band and downsample in time to 50 Hz. We then selected 21 events from HCMR and 8 events from Nestor that were identified as (potential) earthquakes, 6 of which were located and catalogued (see [31]). For each of these events we extract the data within a 41 s time window (2048 time samples at 50 Hz) centred approximately around the first arrival of




Fig. 2. Deep Learning model architecture. An input sample y consists of 11 waveforms recorded at neighbouring DAS channels, each of a length of 2048 time samples. Out of these waveforms, 1 target waveform is randomly selected and defined as J . This waveform is subsequently set to zero (blanked), to enforce J -invariance. In the notation of Section IV-A, we now have $u = \Pi_{J^c}(y)$, which is passed into a U-Net auto-encoder (f_θ). Finally, we apply the projection operation to the model output such that all waveforms except for the target waveform are blanked, i.e. $v = \Pi_J(f_\theta(u))$. The loss for the mini-batch sample is computed as the ℓ_2 -norm between $\Pi_J(y)$ and v .




Fig. 3. Geographic location of the DAS cables HCMR and NESTOR. The inset shows the location of the region of interest within Greece.

the high-amplitude seismic wave trains (most likely surface waves). Lastly, the data of each channel is normalised by its standard deviation.

D. Synthetic data generation

To gain a first-order understanding of the performance of the denoiser, we generate a synthetic data set with “clean” waveforms corrupted by Gaussian white noise with a controlled signal-to-noise ratio (SNR). The clean strain rate waveforms are obtained from three-component broadband seismometer recordings of the Piñon Flats Observatory Array (PFO, [32]),

California, USA, of 82 individual earthquakes. These earthquakes are manually selected based on a visual evaluation of SNR and waveform diversity. The strain rate $\dot{\epsilon}$ recorded by DAS at a location x can be expressed as:

$$\dot{\epsilon}(x) = \frac{1}{L} \left[\dot{u} \left(x + \frac{L}{2} \right) - \dot{u} \left(x - \frac{L}{2} \right) \right] \quad (3)$$

where L is the gauge length and \dot{u} is the particle velocity [33]. To simulate DAS strain rate recordings, we take two broadband stations in the array separated by a distance of 50 m and divide the difference between their respective waveform recordings by their distance. Owing to the low noise floor of these shallow borehole seismometers, the resulting strain rate waveforms exhibit an extremely high SNR. The PFO broadband stations are sampled at a 40 Hz frequency, so in order to simulate a 1-10 Hz frequency band sampled at 50 Hz (i.e. a frequency range of 0.04-0.4 times the Nyquist frequency), we filter the synthetic waveforms in a 0.8-8 Hz frequency band and apply no resampling. All the waveforms are then scaled by their individual standard deviations. We will refer to these clean, simulated strain rate waveforms as “master” waveforms.

During training of the model, we create a synthetic sample by randomly choosing a master waveform, along with a random apparent wave speed v in the range of $\pm 0.2 - 10 \text{ km s}^{-1}$. A total of 11 copies of the selected waveform are created, and each are offset in time in accordance with the moveout, i.e. $\Delta N_i = \text{int}(iLf/v)$, ΔN_i being the time offset in number of samples of the i -th copy ($0 \leq i < 11$), L the gauge length (30 m), f the sampling frequency (50 Hz), and v the apparent wave speed. The waveforms are then cropped within a window of 2048 time samples, positioned randomly around the first arrival. Lastly, a SNR value is sampled from a log-uniform distribution over 0.01-10, and the waveforms are rescaled such that the maximum amplitude of the signal is $2\sqrt{\text{SNR}}$. These scaled waveforms are superimposed onto Gaussian white noise with unit variance, filtered in a 1-10 Hz

frequency band, and scaled by the total variance. Additionally, the data are augmented by performing random polarity flips and time reversals of the final sample.

E. Training procedure

To improve the rate of convergence on the very limited real-world DAS data set, we first train the model on the synthetic dataset. We split the dataset of master waveforms 70-30 in a training and validation set, and use these separate sets to generate training and validation synthetics as described in Section IV-D. We generate a new batch of synthetics after each training epoch, which mitigates overfitting, from which mini-batches of 32 samples are created. We continue training for 2000 epochs at which point the model performance on the validation set saturates. Training on the synthetic dataset took just over 6 hours on a single nVidia Quadro P4000. The model with the best validation set performance is saved and used in the analysis of the synthetic data.

We then continue training on the real-world DAS dataset. Out of the 21 recorded events on HCMR and 8 events on NESTOR, we manually select 4 and 2 events for validation, respectively, and keep the remaining events for training. During training, we generate a batch of samples from randomly selected events and central DAS channels. We then take 5 DAS channels on either side of a central channel and randomly select one of them as a target channel, to be blanked from the input and to be reconstructed by the model. We additionally perform polarity flips and time reversals on the set of 11 waveforms to augment the dataset. A new batch of samples is created after each epoch. Owing to the pretraining, the model performance saturates at a satisfactory validation set performance after roughly 50 epochs (15 minutes), indicating that further refinements to new DAS datasets can be made in a matter of minutes. J -invariant reconstructions of the DAS data recorded along the entire cable are generated by creating 11-channel input samples centred around a target DAS channel, and sliding that window from one DAS channel to the next until all of the channels along the cable have been reconstructed by the model.

V. RESULTS ON SYNTHETIC DATA

We begin with a qualitative assessment of the model performance, taking synthetic examples exhibiting a signal-to-noise ratio (SNR) of $\text{SNR} = \{10, 1, 0.1\}$ – see Fig. 4. These test samples are generated as described in Section IV-D taking an apparent wave speed of 1.5 km s^{-1} . In the original (clean) data, the P- and S-waves are clearly distinguishable from the background noise, and with the S-wave exhibiting a considerably higher amplitude than the P-wave (Fig. 4a). After adding a modest amount of Gaussian white noise ($\text{SNR} = 10$, filtered in a 1-10 Hz frequency band), these waveform features are still clearly visible (Fig. 4b). When this mildly corrupted sample is fed into Deep Learning model, the resulting reconstruction is near-identical to the input and the original waveform (compare Fig. 4c with panels a and b, and Fig. 4j with panels h and i).

At an intermediate $\text{SNR} = 1$, the P-wave train and portions of the S-wave train vanish within the noise, but peak strain

rates are still visible (Fig. 4d). Whereas accurate picking of P- and S-arrival times would not be possible for such SNR values, detection algorithms may still correctly identify this earthquake. After J -invariant filtering (Fig. 4e), the signals P-wave train is lifted out of the noise level. The onset of the P- and S-waves becomes much more clear, permitting a crude estimation of their arrival times (and therefore distance to the seismic source). Moreover, details of the S-wave coda are fairly well recovered.

Lastly, at an extremely poor $\text{SNR} = 0.1$, the earthquake signal is entirely obscured by the noise (Fig. 4f and k). It would be incredibly challenging to detect an event with such SNR using conventional detection algorithms (e.g. STA/LTA detection [34]). After J -invariant filtering (Fig. 4g and l), the S-wave train is recovered, albeit with a much lower amplitude than the original signal. The P-wave can no longer be distinguished in a single waveform (Fig. 4g), but from Fig. 4l it is apparent that small amounts of the P-wave energy are recovered. In spite of the imperfect reconstruction, the SNR of the reconstructed signal is sufficient to identify this event with detection algorithms.

We continue with a quantitative assessment of the model performance by computing the scaled variance of the residuals, defined as $R = \langle (y - y')^2 \rangle / \langle y^2 \rangle$, y being the clean signal and y' the reconstruction. We compute this quantity for a range of values of SNR and slowness (reciprocal of wave velocity) – see Fig. 5. As expected, the model output becomes more accurate when the SNR is high (Fig. 5a), which saturates towards the end of the SNR range. Towards the lower end of the SNR range the scaled variance approaches 1, indicating that the model essentially produces zero-centred random noise with small variance (so that $\langle (y - y')^2 \rangle \approx \langle y^2 \rangle$). This is a highly desired outcome: when provided with purely random, incoherent noise, the model should output zero. In other words, the prior learnt by the decoder of the auto-encoding network is zero, which prevents the generation of non-existing signals driven by a dominant non-zero prior.

As detailed in Section IV-D, the time-offset of the waveform between neighbouring channels is governed by the slowness (reciprocal wave velocity). For low slowness values, the offset between neighbouring waveforms is minimal, so that a reasonably accurate reconstruction can be generated from simply copying a non-blanked waveform from the model input. This is obviously undesired, and so we investigate this hypothesis by systematically varying the slowness (and correspondingly the time-offset between channels). As is apparent from Fig. 5b, this hypothesis can be safely discarded: the scaled residuals remain constant over a wide range of slowness, varying from 0.1 to 3.3 s km^{-1} (small to large time-offset, respectively). However, we do find a small but systematic drop in scaled residuals at fixed intervals. Further investigation reveals that these occur at integer multiples of $1/9.6 \text{ s km}^{-1}$. Recall from Section IV-D that the offset between neighbouring channels is given as $\Delta N = \text{int}(Lf/v)$, L being the gauge length of 19.2 m and f the sampling frequency of 50 Hz. For $v < 960 \text{ m s}^{-1}$, the offset between neighbouring channels is 0. For $960 \leq v < 1920 \text{ m s}^{-1}$ the offset is 1, etc. For channels

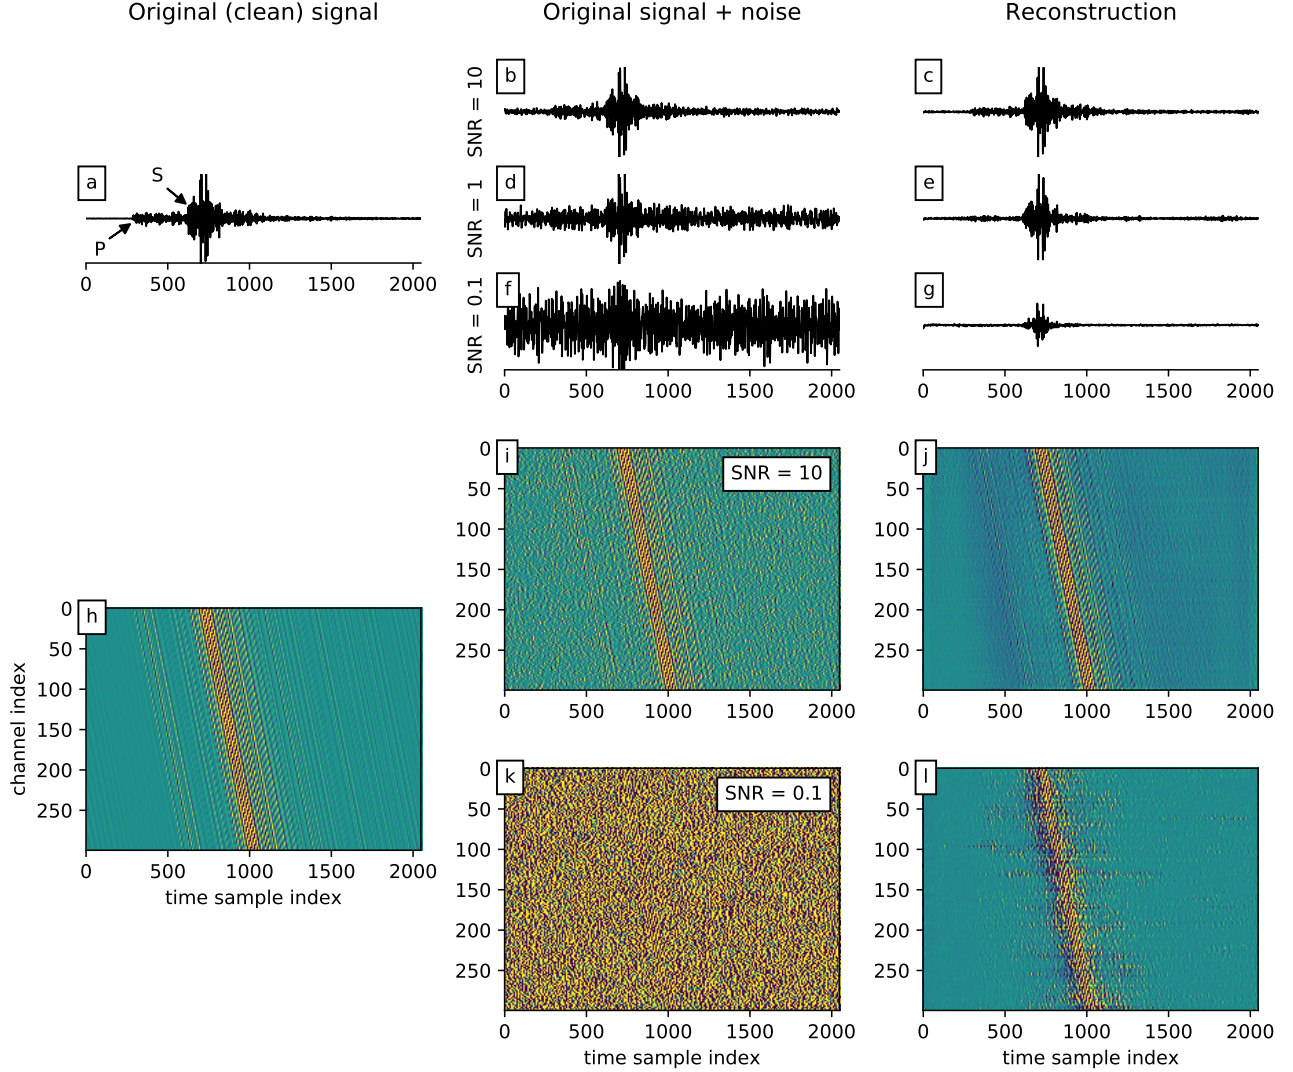


Fig. 4. Synthetic examples of model performance. a) Representative example waveform of an original (clean) earthquake signal, which is used to generate the test samples. The P- and S-wave first arrivals are indicated by “P” and “S”, respectively; b,d,f) The original signal with added noise with SNR values of 10, 1, and 0.1, respectively, which serve as model input; c,e,g) The J -invariant reconstructions corresponding with inputs b,d,f, respectively. h) The original waveform shifted and stacked according to a constant apparent wave speed of 1.5 km s^{-1} ; i,k) The shifted/stacked waveforms with added noise with SNR values of 10 and 0.1, respectively; j,l) The J -invariant reconstructions of i,k, respectively.

that are separated in distance by i gauge lengths, these jumps in time-offset occur at integer multiples of $v = iLf$. In other words, the method by which the synthetic samples are generated causes a jagged, non-exact offset between channels due to integer rounding of the time-offset. Only when the slowness is an integer multiple of $1/Lf = 1/960 \text{ s m}^{-1}$ is the offset between the channels exactly as given by the theoretical move-out. When this condition is satisfied, the offset between close and far channels is fully consistent, and correspondingly the model performance improves. This suggests that not only does the model refrain from simply copying the input data, but that it also considers both far and close channels to assess the move-out, which is then used to reproduce the correct time-offset of the reconstruction. In the real-world DAS data the wavefield is not discretised (i.e. the arrival of waves at a given channel is exact) and so this time-offset rounding does

not occur.

VI. RESULTS ON DAS DATA

After performing the pre-training on synthetic data, we retrain the model for real-world DAS data recorded by the HCMR and NESTOR cables. We first perform a qualitative assessment of the model performance by considering two events in the validation set of the HCMR data (Fig. 6) and two events in the validation set of the NESTOR data (Fig. 7). (We make a side note to the reader that these figures are prone to aliasing artefacts as a result of the PDF rendering at different magnifications.) As can also be seen in the synthetic tests (e.g. Fig. 4j), the model tends to introduce some low-frequency parasitic signals, which are particularly well visible in Fig. 6b and Fig. 7b. Fortunately these artefacts are easily removed by bandpass filtering the reconstructions in a 1-10 Hz

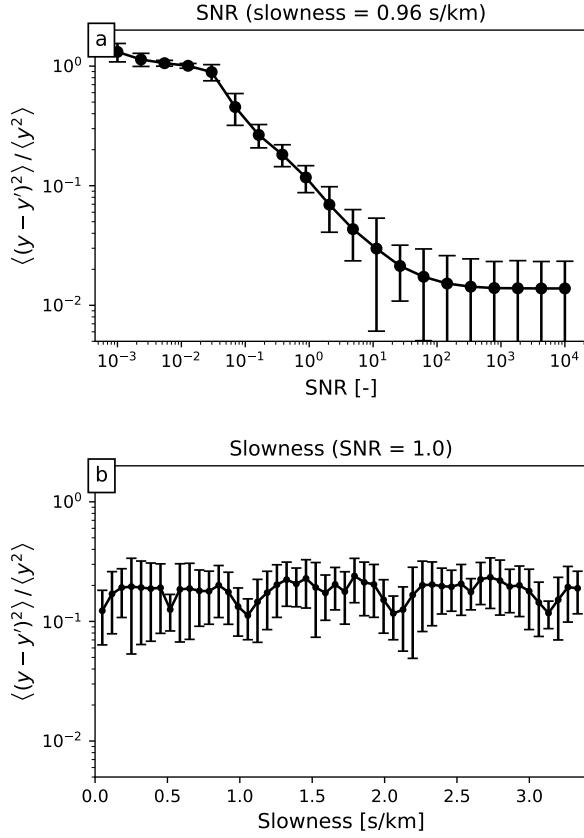


Fig. 5. Quantitative assessment of the model performance on synthetic data. a) The scaled variance of the residuals as a function of SNR, with a slowness of 0.96 s km^{-1} ; b) The scale variance of the residuals as a function of slowness, with $\text{SNR} = 1$. In both panels the range of SNR and slowness values used during training are indicated. The error bars are calculated from 300 samples generated from the same clean waveform (as shown in Fig. 4a).

frequency band. This trivial post-processing step yields high-quality reconstructions of the coherent signals present in the input data, even when the SNR is low (compare e.g. Fig. 6e with 6g).

To get a measure of the model performance without ground truth, we compute the local waveform coherence before and after J -invariant filtering and assess the gain in coherence. We define the mean local waveform coherence CC around the k -th DAS channel as:

$$CC_k = \frac{1}{4N^2} \left[\sum_{i,j=-N}^{+N} \max \left(\frac{x_{k+i} * x_{k+j}}{\sqrt{\sum_t x_{k+i}^2 \sum_t x_{k+j}^2}} \right) - 2N - 1 \right] \quad (4)$$

where x_n denotes the waveform at the n -th channel, $*$ denotes cross-correlation, and $\sum_t x^2$ denotes the sum over all time samples in x . The bin size N is set to 5. The coherence gain is then defined as the local coherence computed for the J -invariant reconstruction divided by that of the input data. As such, coherence gains above 1 indicate that the reconstruction exhibits improved waveform coherence compared to the input data, which is beneficial for coherence-based seismological analyses (template matching, beamforming). Looking at Fig. 6d,h and Fig. 7d,h, the J -invariant reconstructions

practically always exhibit (much) higher waveform coherence. Along some cable segments this quantity is inflated due to an absence of coherent signals in both the input data and the reconstruction, which is particularly apparent in Fig. 6h between 2 and 4 km distance. Nonetheless along other segments, such as between 0 and 2 km in Fig. 6h or around 15 km in Fig. 7d and h, the local coherence of recorded earthquake signals have improved substantially.

When considering the SNR of the DAS data as shown e.g. in Fig. 6a, we see that there are segments of the cable that exhibit better SNR than others (e.g. at 0.5, 3.9, and 10.8 km along the HCMR cable). This along-cable variation in SNR may be due to variations in environmental noise, cable-ground coupling degree, or orientation of the cable with respect to the wave propagation direction [33], [35]. At locations where the apparent SNR is high, we can attempt to make a wiggle-for-wiggle comparison between the input data and the reconstructions (Fig. 8). When doing so, we see that the model correctly attenuates the random noise in the first 25 s, and subsequently increases its amplitude to match the recorded signals. Overall the reconstructions exhibit a lower maximum amplitude than the input data, which is as expected (the model removes the contribution of the noise to the recorded data). The phase of the large-amplitude arrivals seems to be matched fairly well, which is important for seismological methods that rely on phase information, such as beamforming. This can be quantitatively expressed as the correlation coefficient computed for the waveforms after 25 s, which is overall satisfactory (in the range of 0.66 to 0.89 as indicated in Fig. 8).

VII. DISCUSSION

A. Comparison with traditional filtering methods

Since our proposed filtering approach for DAS data is closely related to image denoising and enhancement, we compare our method with two commonly used image processing techniques, namely non-local means (NLM) as implemented in scikit-image [36] and BM3D [37]. These two non-learning algorithms also served as a benchmark for the study of [25]. For this comparison we select the event recorded by HCMR in Fig. 6a, and process the data with NLM and BM3D using a noise variance estimated from the first 20 s of data. We then compare the denoised data with our J -invariant reconstruction after bandpass filtering – see Fig. 9. For a more detailed visual comparison, we focus on a shallow segment of the cable that recorded many fine-detailed features. While these low-amplitude details do not significantly contribute to other metrics like SNR or ℓ_2 -loss, they are critical for the detection of body wave arrivals. DAS is strongly sensitive to surface waves, both owing to their horizontal inclination [35] and slow phase velocity [38], and exhibits much less sensitivity to body waves. As a result, P- and S-waves are recorded as comparatively low amplitude features, and so the recovery and preservation of these is important.

When comparing the performance of the various denoising methods, it is immediately clear that our proposed method preserves fine-grained details with much higher fidelity than

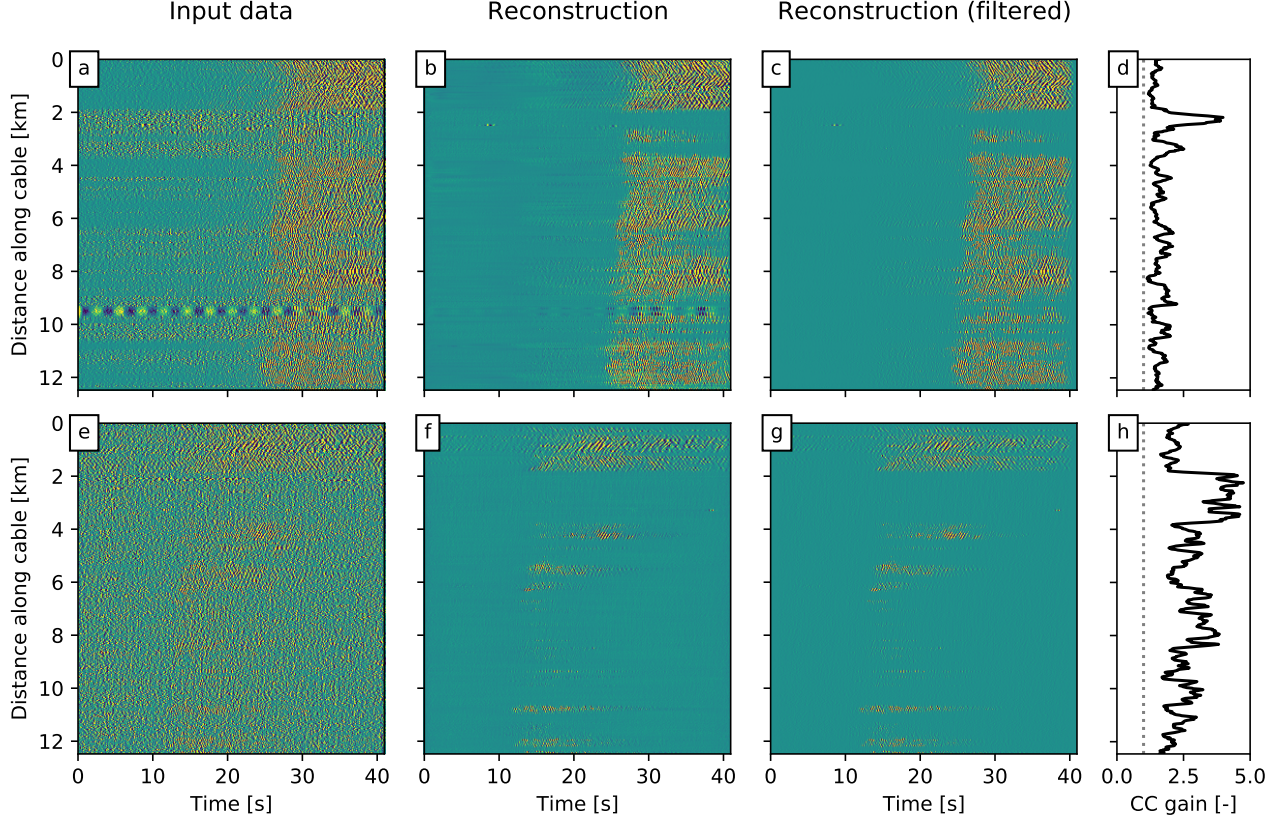


Fig. 6. J -invariant filtering results of HCMR data. a,e) The strain rate wavefield of two events in the validation set recorded by the HCMR cable; b,f) J -invariant reconstruction (model output); c,g) J -invariant reconstruction filtered in a 1-10 Hz pass band; d,h) Local increase in waveform coherence along the cable. The vertical dotted line marks a gain of 1.

the BM3D and particularly the NLM method (compare Fig. 9h with 9g and 9f, respectively). In terms of the performance on the high-amplitude wave train, NLM and BM3D essentially copy the input without altering the amplitudes, which is questionable given that these are affected by the same noise levels as prior to the arrival of these waves. Moreover, our method processes the event shown in Fig. 9 in less than 1.3 s, while NLM and BM3D require 5.0 and 28.4 s respectively (while using the “fast” method for NLM [39], which in our case is about 50 times faster than the more precise classical method). These speed gains, along with the enhanced precision of the method render our J -invariant denoiser superior to traditional image denoising methods.

B. Applications of the proposed method

DAS typically generates large volumes of data, of the order of terabytes per fibre per day. Being a data-driven approach, Deep Learning methods are ideally suited for automated processing and analysis of DAS data. However, since DAS is still an emerging technology, large *labelled* datasets are still lacking. During the experiments analysed in this study, only 6 catalogued earthquakes were recorded by the two cables [31], which prohibits a supervised approach applied to e.g. earthquake detection. In this stage of the development of DAS

unsupervised or self-supervised learning methods are more feasible.

The J -invariant filtering approach that we detailed in this study is entirely self-supervised and, after pretraining on synthetic data, can be trained using only a small dataset comprising just 29 recorded “anomalies” (only 6 of which have been formally classified and catalogued as regional earthquakes). This opens up a plethora of applications in seismology and DAS signal analysis. First and foremost, the performance of conventional earthquake detection algorithms (such as the commonly used STA/LTA method) can be dramatically improved by suppressing incoherent background noise. Taking the synthetic test shown in Fig. 4f and g as an extreme example, a conventional algorithm operating on a single waveform would be unable to detect the earthquake signal that is completely obscured by the noise. Particularly for the analysis of microseismicity (e.g. around fluid injection or extraction wells, [40], [41]), the improved SNR after J -invariant filtering would massively improve catalogue completeness, provide better earthquake location estimates, and draw a more complete picture of the evolution of seismicity in time and space. Such improved SNR could not be obtained by single-waveform frequency-based methods, since the signal of interest and the background noise share a common frequency band. By taking into account the spatial extent of the signal,

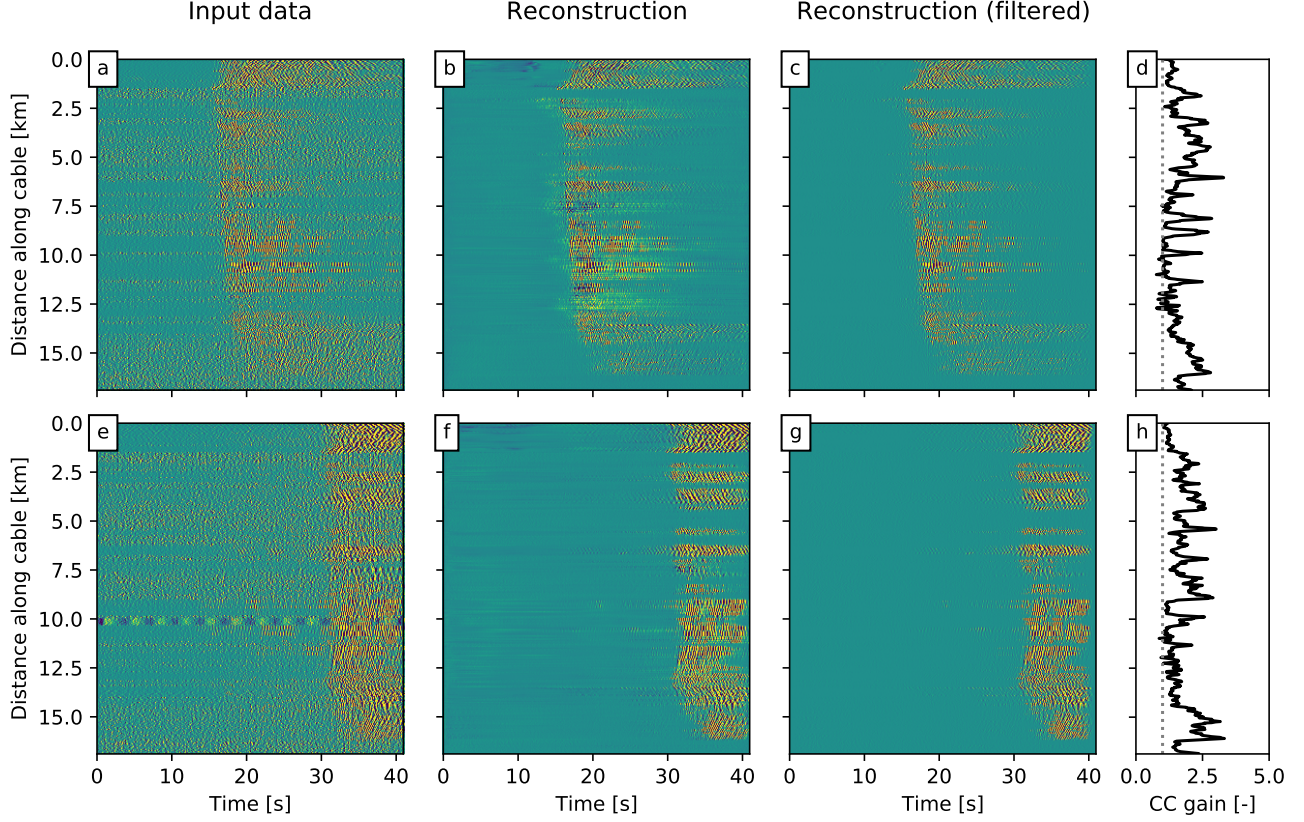


Fig. 7. J -invariant filtering results of NESTOR data. a,e) The strain rate wavefield of two events in the validation set recorded by the NESTOR cable; b,f) J -invariant reconstruction (model output); c,g) J -invariant reconstruction filtered in a 1-10 Hz pass band; d,h) Local increase in waveform coherence along the cable. The vertical dotted line marks a gain of 1.

this limitation can be overcome.

A second application of the method pertains to waveform-coherence based methods like template matching [42] and beamforming [43]: in template matching a given time-series is analysed for the occurrence of a previously recorded and identified signal (the template). By cross-correlating the time-series with the template in a sliding time window, events similar to the template can be identified through a high correlation coefficient above some predefined threshold. This approach has recently been applied successfully to DAS data [44], detecting numerous small earthquakes induced by geothermal energy extraction operations. Out of the 116 detections, 68 were identified within the level of the background noise, demonstrating the sensitivity of the template matching technique. It is possible that by preprocessing the data with our proposed J -invariant filtering approach many more events could be identified, not only due to a higher SNR of the target time-series, but also due to a higher SNR of the template waveform. The improved SNR of both cases would lower the threshold above which a detection is deemed significant. Inherent to the template matching method, a detection automatically provides a rough location estimate as well, helping to rapidly build (micro-)earthquake catalogues from DAS experiments. All of the above also applies to the detection of volcanic or tectonic tremor and very-low frequency earthquakes, two seismological

features that often exist at or below the ambient noise level [45].

Similar to template matching, seismic beamforming and backprojection relies on waveform coherence to assess the move-out of a seismic signal propagating across a seismometer or DAS array (see e.g. [8], [43], [46]). The quality of these analyses depends on the resolution with which the phase shift can be determined. This resolution may improve at higher signal frequencies, which in turn suffer from stronger attenuation and waveform decorrelation. There thus exists a trade-off between the beamforming/backprojection resolution and the coherence/SNR of the signal of interest, as a function of frequency. By employing J -invariant filtering, coherent signals at higher frequencies can be amplified, helping to shift the trade-off towards higher frequencies and to improve the resolution.

Lastly, inherent to the measurement principle of DAS, strain and strain rate recordings suffer from waveform incoherence caused by heterogeneities of the wave propagation medium that affect the phase velocity field [8], [47]. While this can be efficiently mitigated by converting the strain (rate) measurements to particle motion (displacement or velocity) [8], it requires a seismic station that records particle ground motions to be co-located with a straight section of the DAS fibre. Since this is generally not the case, other methods




Fig. 8. Wiggle-for-wiggle comparison between high SNR regions along the HCMR cable and their J -invariant reconstructions (after bandpass filtering). The waveforms correspond with the wavefield shown in Fig. 6a at the indicated DAS channel indices.

need to be explored to diminish the influence of subsurface heterogeneities in the recorded signals. Our proposed method may help, as it attempts to uncover only those signals that are coherent over some local distance. As demonstrated in Fig. 5b, the performance of the method does not depend on the offset between the waveforms recorded at neighbouring DAS channels. Therefore, one could attempt to extract signals that are coherent over longer distances by selecting DAS channels that are farther apart (instead of selecting directly adjacent channels).

C. Limitations and extensions beyond DAS

Aside from the numerous applications that may benefit from the proposed method, we acknowledge certain limitations and point out potential extensions of the method to guide future endeavours. Firstly, the amplitudes of the earthquake signals are not always fully recovered. To some extent this is expected, as the model attempts to remove the contribution of the noise to the signal. However, for signals that are at or below the noise floor, the true amplitude of the signal cannot be reliably estimated. This can be most clearly seen in Fig. 4g, exhibiting a substantially lower amplitude than Fig. 4c and e. The reason for this is simple: whether a signal's amplitude is 10 or 10 000 times smaller than the noise level, the resulting superposition of signal plus noise looks the same, i.e. the coherent signal contributes negligibly to the overall signal. In this case of a very low SNR, it cannot be inferred what the original signal amplitude was, other than that it is upper-bounded by the noise

level. Thusly, one should exercise caution interpreting the amplitudes of J -invariant reconstructions of low SNR signals. For high SNR signals this seems to be much less of a problem, as evidenced by the low scaled residuals for high SNR samples (see Fig. 5a).

Secondly, the underlying principle of our method relies on spatio-temporal signal coherence. Any parts of the signal that are incoherent will not be reconstructed and therefore be filtered from the input. This is useful for incoherent noise sources, given that the signals of interest are strongly coherent. However, it is common in (submarine) DAS to also observe strongly coherent nuisance sources like ocean gravity waves and related phenomena. For many seismological applications these are considered part of the background noise, and are therefore not desired in the output. Our method does not address the separation of multiple coherent signals. However, once incoherent noise has been removed from the input, it may be easier to separate the remaining coherent signals by other means (e.g. [24]).

Lastly, we applied our proposed method to a DAS array with constant spacing between the recording channels. Although the model makes no assumption regarding the geometry of the array (it applies to both straight and curved cables provided that the radius curvature is sufficiently large), it does implicitly require that the data are evenly distributed along the trajectory of the cable. This limits the application of the method to DAS arrays and, practically speaking, linear seismic arrays with constant inter-station distance. Fortunately, this limitation can potentially be circumvented by treating the array as a graph with the receiver station location as node attributes, or with station distances as edge attributes, and performing the learning task on the graph [48], [49]. As has been previously demonstrated for conventional seismic networks, incorporating station location information can substantially expedite geophysical learning tasks on non-Euclidean objects [50], [51]. This offers an opportunity for future work to extend this efficient filtering technique to standard seismometer arrays.

VIII. CONCLUSIONS

In this study, we present a self-supervised Deep Learning approach for blind denoising of Distributed Acoustic Sensing (DAS) data, based on the concept of J -invariance introduced by [25]. While most Deep Learning denoising methods are supervised (i.e., require a noise-free ground truth) or operate only on a single waveform, our approach leverages spatio-temporal coherence of the recorded DAS data to distinguish between incoherent signals (noise) and coherent signals (earthquakes). This permits the separation of noise and signals that share a common frequency band without the need for a noise-free ground truth. Even though the concept of J -invariance extends beyond learning algorithms, we incorporate the concept within a Deep Learning framework by training a convolutional U-Net auto-encoder with a training objective that leverages J -invariance of the signals of interest.

We first demonstrate the validity of the method on synthetic data for which a ground truth is available. This analysis shows that J -invariant filtering has the potential to faithfully

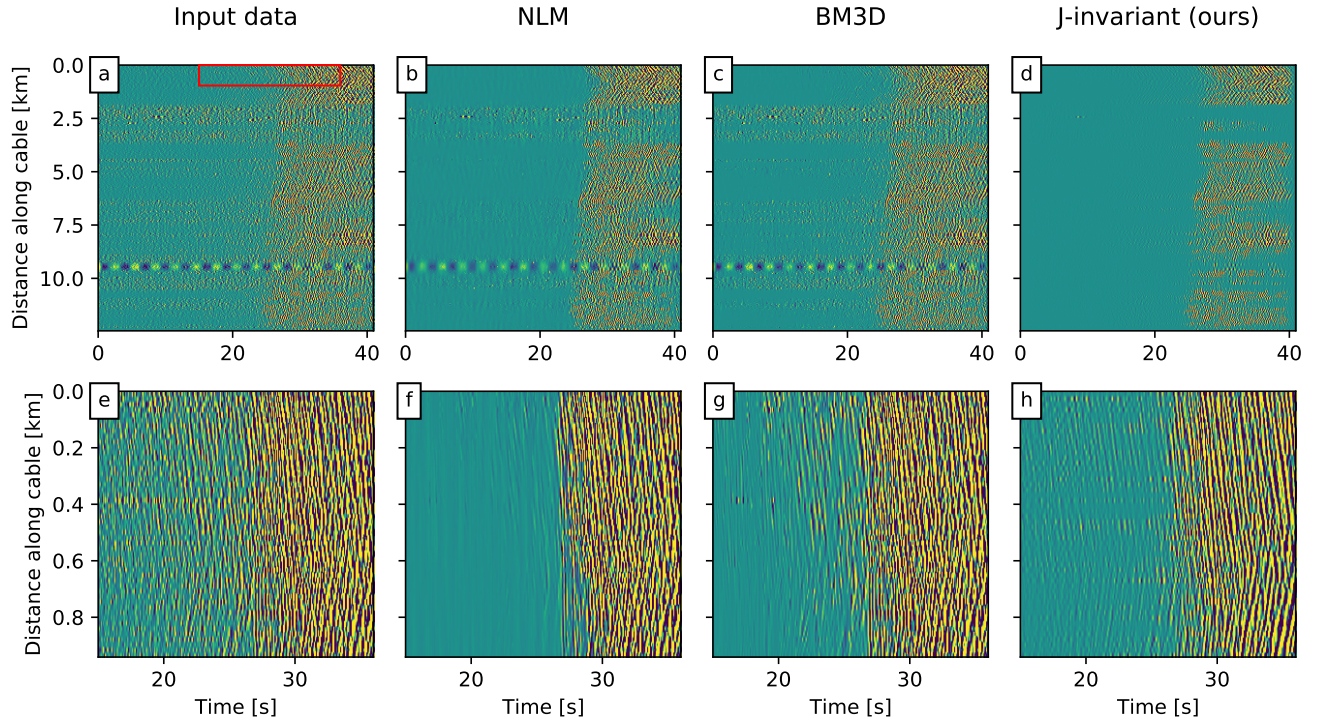


Fig. 9. Comparison of our proposed method with two conventional image denoising methods, non-local means (NLM) and BM3D. a) Input data (same as Fig. 6a); b-d) Denoised data using NLM, BM3D, and J -invariant filtering, respectively; e-h) Detailed view of the data in the region indicated by the red-bordered patch in panel a.

reconstruct coherent signals that are completely obscured by incoherent noise, with signal-to-noise ratios (SNRs) well below 1. We then continue to apply our model to real-world DAS data acquired by two submarine fibre-optic cables that are deployed off-shore Greece. Over the course of the experiments, several earthquakes were recorded with varying SNR, offering a suitable target for evaluating the model performance on real-world data. We find that in all cases the model is able to attenuate the incoherent noise that is clearly seen prior to the earthquake, isolating the earthquake signals even when the SNR is low. Moreover, the waveform coherence improves along every segment of the cables, which is beneficial for coherence-based seismological analyses.

The excellent performance of the Deep Learning denoising model expedites numerous applications in seismology, including (micro-)earthquake detection, template matching, and beamforming. Given the ease at which the proposed method can be applied to (new) DAS data, and given the numerous applications that benefit from improved SNR and waveform coherence, we suggest that J -invariant filtering could take a role in the standard workflow of DAS data processing routines.

ACKNOWLEDGEMENTS

This research was supported by the French government through the UCA^{JEDI} Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-15-IDEX-01, and through the 3IA Côte d’Azur Investments in the Future project with the reference number ANR-19-P3IA-0002. The Deep Learning was

performed within the TensorFlow framework [52] and generic data manipulations were performed with NumPy [53] and SciPy [54]. Data visualisation was done using Matplotlib [55]. Data, code, and pre-trained models are available from doi.org/10.6084/m9.figshare.14152277.

REFERENCES

- [1] A. H. Hartog, *An Introduction to Distributed Optical Fibre Sensors*. CRC Press, May 2017.
- [2] S. Dou, N. Lindsey, A. M. Wagner, T. M. Daley, B. Freifeld, M. Robertson, J. Peterson, C. Ulrich, E. R. Martin, and J. B. Ajo-Franklin, “Distributed Acoustic Sensing for Seismic Monitoring of The Near Surface: A Traffic-Noise Interferometry Case Study,” *Scientific Reports*, vol. 7, no. 1, p. 11620, Sep. 2017.
- [3] G. Fang, Y. E. Li, Y. Zhao, and E. R. Martin, “Urban Near-Surface Seismic Monitoring Using Distributed Acoustic Sensing,” *Geophysical Research Letters*, vol. 47, no. 6, p. e2019GL086115, 2020.
- [4] N. J. Lindsey, T. C. Dawe, and J. B. Ajo-Franklin, “Illuminating seafloor faults and ocean dynamics with dark fiber distributed acoustic sensing,” *Science*, vol. 366, no. 6469, pp. 1103–1107, Nov. 2019.
- [5] A. Sladen, D. Rivet, J. P. Ampuero, L. De Barros, Y. Hello, G. Calbris, and P. Lamare, “Distributed sensing of earthquakes and ocean-solid Earth interactions on seafloor telecom cables,” *Nature Communications*, vol. 10, no. 1, pp. 1–8, Dec. 2019.
- [6] G. Marra, C. Clivati, R. Luckett, A. Tampellini, J. Kronjäger, L. Wright, A. Mura, F. Levi, S. Robinson, A. Xuereb, B. Baptie, and D. Calonico, “Ultrastable laser interferometry for earthquake detection with terrestrial and submarine cables,” *Science*, vol. 361, no. 6401, pp. 486–490, Aug. 2018.
- [7] J. B. Ajo-Franklin, S. Dou, N. J. Lindsey, I. Monga, C. Tracy, M. Robertson, V. Rodriguez Tribaldos, C. Ulrich, B. Freifeld, T. Daley, and X. Li, “Distributed Acoustic Sensing Using Dark Fiber for Near-Surface Characterization and Broadband Seismic Event Detection,” *Scientific Reports*, vol. 9, no. 1, pp. 1–14, Feb. 2019.
- [8] M. P. A. van den Ende and J.-P. Ampuero, “Evaluating Seismic Beamforming Capabilities of Distributed Acoustic Sensing Arrays,” *Solid Earth Discussions*, pp. 1–24, Sep. 2020.

- [9] T. Perol, M. Gharbi, and M. Denolle, "Convolutional neural network for earthquake detection and location," *Science Advances*, vol. 4, no. 2, p. e1700578, Feb. 2018.
- [10] Z. E. Ross, M.-A. Meier, E. Hauksson, and T. H. Heaton, "Generalized Seismic Phase Detection with Deep Learning," *Bulletin of the Seismological Society of America*, vol. 108, no. 5A, pp. 2894–2901, Oct. 2018.
- [11] Z. E. Ross, M.-A. Meier, and E. Hauksson, "P Wave Arrival Picking and First-Motion Polarity Determination With Deep Learning," *Journal of Geophysical Research: Solid Earth*, vol. 123, no. 6, pp. 5120–5129, Jun. 2018.
- [12] W. Zhu and G. C. Beroza, "PhaseNet: A deep-neural-network-based seismic arrival-time picking method," *Geophysical Journal International*, vol. 216, no. 1, pp. 261–273, Jan. 2019.
- [13] C. Tian, L. Fei, W. Zheng, Y. Xu, W. Zuo, and C.-W. Lin, "Deep learning on image denoising: An overview," *Neural Networks*, vol. 131, pp. 251–275, Nov. 2020.
- [14] S. Bonnefoy-Claudet, F. Cotton, and P.-Y. Bard, "The nature of noise wavefield and its applications for site effects studies: A literature review," *Earth-Science Reviews*, vol. 79, no. 3, pp. 205–227, Dec. 2006.
- [15] A. Inbal, T. Cristea-Platon, J.-P. Ampuero, G. Hillers, D. Agnew, and S. E. Hough, "Sources of Long-Range Anthropogenic Noise in Southern California and Implications for Tectonic Tremor Detection," *Bulletin of the Seismological Society of America*, vol. 108, no. 6, pp. 3511–3527, Oct. 2018.
- [16] S. Soltanayev and S. Y. Chun, "Training deep learning based denoisers without ground truth data," in *Advances in Neural Information Processing Systems 31*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds. Curran Associates, Inc., 2018, pp. 3257–3267.
- [17] M. Zhussip, S. Soltanayev, and S. Y. Chun, "Training Deep Learning Based Image Denoisers From Undersampled Measurements Without Ground Truth and Without Image Prior," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Jun. 2019, pp. 10247–10256.
- [18] S. Beckouche and J. Ma, "Simultaneous dictionary learning and denoising for seismic data," *GEOPHYSICS*, vol. 79, no. 3, pp. A27–A31, May 2014.
- [19] Y. Chen, M. Zhang, M. Bai, and W. Chen, "Improving the Signal-to-Noise Ratio of Seismological Datasets by Unsupervised Machine Learning," *Seismological Research Letters*, vol. 90, no. 4, pp. 1552–1564, Jul. 2019.
- [20] S. Yu, J. Ma, and W. Wang, "Deep learning for denoising," *GEOPHYSICS*, vol. 84, no. 6, pp. V333–V350, Jul. 2019.
- [21] D. Liu, W. Wang, X. Wang, C. Wang, J. Pei, and W. Chen, "Poststack Seismic Data Denoising Based on 3-D Convolutional Neural Network," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 58, no. 3, pp. 1598–1629, Mar. 2020.
- [22] O. M. Saad and Y. Chen, "Deep denoising autoencoder for seismic random noise attenuation," *GEOPHYSICS*, vol. 85, no. 4, pp. V367–V376, Jun. 2020.
- [23] W. Zhu, S. M. Mousavi, and G. C. Beroza, "Seismic Signal Denoising and Decomposition Using Deep Neural Networks," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 11, pp. 9476–9488, Nov. 2019.
- [24] E. R. Martin, F. Huot, Y. Ma, R. Cieplik, S. Cole, M. Karrenbach, and B. L. Biondi, "A Seismic Shift in Scalable Acquisition Demands New Processing: Fiber-Optic Seismic Signal Retrieval in Urban Areas with Unsupervised Learning for Coherent Noise Removal," *IEEE Signal Processing Magazine*, vol. 35, no. 2, pp. 31–40, Mar. 2018.
- [25] J. Batson and L. Royer, "Noise2Self: Blind Denoising by Self-Supervision," in *Proceedings of the 36th International Conference on Machine Learning*, Long Beach, California, USA, Jun. 2019.
- [26] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," *arXiv:1505.04597 [cs]*, May 2015.
- [27] R. Zhang, "Making Convolutional Networks Shift-Invariant Again," *arXiv:1904.11486 [cs]*, Apr. 2019.
- [28] P. Ramachandran, B. Zoph, and Q. V. Le, "Searching for Activation Functions," *arXiv:1710.05941 [cs]*, Oct. 2017.
- [29] W. Hu, L. Xiao, and J. Pennington, "Provable Benefit of Orthogonal Initialization in Optimizing Deep Linear Networks," *arXiv:2001.05992 [cs, math, stat]*, Jan. 2020.
- [30] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," *arXiv:1412.6980 [cs]*, Jan. 2017.
- [31] I. Lior, A. Sladen, D. Rivet, J.-P. Ampuero, Y. Hello, C. Becerril, H. F. Martins, P. Lamare, C. Jestin, S. Tsagkli, and C. Markou, "On the Detection Capabilities of Underwater DAS," *Journal of Geophysical Research: Solid Earth*, vol. n/a, no. n/a, p. e2020JB020925, 2021.
- [32] U. S. D. Frank Vernon, "Piñon Flats Observatory (PFO) Array," 2014.
- [33] H. F. Wang, X. Zeng, D. E. Miller, D. Fratta, K. L. Feigl, C. H. Thurber, and R. J. Mellors, "Ground motion response to an ML 4.3 earthquake using co-located distributed acoustic sensing and seismometer arrays," *Geophysical Journal International*, vol. 213, no. 3, pp. 2020–2036, Jun. 2018.
- [34] R. Allen, "Automatic phase pickers: Their present use and future prospects," *Bulletin of the Seismological Society of America*, vol. 72, no. 6B, pp. S225–S242, Dec. 1982.
- [35] E. R. Martin, N. Lindsey, J. Ajo-Franklin, and B. Biondi, "Introduction to Interferometry of Fiber Optic Strain Measurements," Jun. 2018.
- [36] S. van der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, and T. Yu, "Scikit-image: Image processing in Python," *PeerJ*, vol. 2, p. e453, Jun. 2014.
- [37] Y. Mäkinen, L. Azzari, and A. Foi, "Exact Transform-Domain Noise Variance for Collaborative Filtering of Stationary Correlated Noise," in *2019 IEEE International Conference on Image Processing (ICIP)*, Sep. 2019, pp. 185–189.
- [38] T. M. Daley, D. E. Miller, K. Dodds, P. Cook, and B. M. Freifeld, "Field testing of modular borehole monitoring with simultaneous distributed acoustic sensing and geophone vertical seismic profiles at Citronelle, Alabama," *Geophysical Prospecting*, vol. 64, no. 5, pp. 1318–1334, 2016.
- [39] J. Darbon, A. Cunha, T. F. Chan, S. Osher, and G. J. Jensen, "Fast nonlocal filtering applied to electron cryomicroscopy," in *2008 5th IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, May 2008, pp. 1331–1334.
- [40] G. Kwiatek, M. Bohnhoff, G. Dresen, A. Schulze, T. Schulte, G. Zimmermann, and E. Huenges, "Microseismicity induced during fluid-injection: A case study from the geothermal site at Groß Schönebeck, North German Basin," *Acta Geophysica*, vol. 58, no. 6, pp. 995–1020, Dec. 2010.
- [41] Y. Guglielmi, F. Cappa, J.-P. Avouac, P. Henry, and D. Elsworth, "Seismicity triggered by fluid injection-induced aseismic slip," *Science*, vol. 348, no. 6240, pp. 1224–1226, Jun. 2015.
- [42] N. S. Senobari, G. J. Funning, E. Keogh, Y. Zhu, C.-C. M. Yeh, Z. Zimmermann, and A. Mueen, "Super-Efficient Cross-Correlation (SEC-C): A Fast Matched Filtering Code Suitable for Desktop Computers," *Seismological Research Letters*, vol. 90, no. 1, pp. 322–334, Nov. 2018.
- [43] P. Goldstein and R. J. Archuleta, "Array analysis of seismic signals," *Geophysical Research Letters*, vol. 14, no. 1, pp. 13–16, 1987.
- [44] Z. Li and Z. Zhan, "Pushing the limit of earthquake detection with distributed acoustic sensing and template matching: A case study at the Brady geothermal field," *Geophysical Journal International*, vol. 215, no. 3, pp. 1583–1593, Dec. 2018.
- [45] A. A. Hutchison and A. Ghosh, "Very low frequency earthquakes spatiotemporally asynchronous with strong tremor during the 2014 episodic tremor and slip event in Cascadia," *Geophysical Research Letters*, vol. 43, no. 13, pp. 6876–6882, 2016.
- [46] L. Meng, A. Inbal, and J.-P. Ampuero, "A window into the complexity of the dynamic rupture of the 2011 Mw 9 Tohoku-Oki earthquake," *Geophysical Research Letters*, 2011.
- [47] S. Singh, Y. Capdeville, and H. Igel, "Correcting wavefield gradients for the effects of local small-scale heterogeneities," *Geophysical Journal International*, vol. 220, no. 2, pp. 996–1011, Feb. 2020.
- [48] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst, "Geometric Deep Learning: Going beyond Euclidean data," *IEEE Signal Processing Magazine*, vol. 34, no. 4, pp. 18–42, Jul. 2017.
- [49] P. W. Battaglia, J. B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner, C. Gulcehre, F. Song, A. Ballard, J. Gilmer, G. Dahl, A. Vaswani, K. Allen, C. Nash, V. Langston, C. Dyer, N. Heess, D. Wierstra, P. Kohli, M. Botvinick, O. Vinyals, Y. Li, and R. Pascanu, "Relational inductive biases, deep learning, and graph networks," *arXiv:1806.01261 [cs, stat]*, Oct. 2018.
- [50] M. P. A. van den Ende and J.-P. Ampuero, "Automated Seismic Source Characterization Using Deep Graph Neural Networks," *Geophysical Research Letters*, vol. 47, no. 17, p. e2020GL088690, 2020.
- [51] J. Münchmeyer, D. Bindi, U. Leser, and F. Tilmann, "Earthquake magnitude and location estimation from real time seismic waveforms with a transformer network," *arXiv:2101.02010 [physics]*, Jan. 2021.
- [52] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah,

- M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, “TensorFlow: Large-scale machine learning on heterogeneous systems,” 2015.
- [53] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant, “Array programming with NumPy,” *Nature*, vol. 585, no. 7825, pp. 357–362, Sep. 2020.
- [54] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and S. . . Contributors, “SciPy 1.0—Fundamental Algorithms for Scientific Computing in Python,” *arXiv:1907.10121 [physics]*, Jul. 2019.
- [55] J. D. Hunter, “Matplotlib: A 2D Graphics Environment,” *Computing in Science & Engineering*, vol. 9, no. 3, pp. 90–95, 2007.