



**HAL**  
open science

## Deep Deconvolution for Traffic Analysis with Distributed Acoustic Sensing Data

Martijn P.A. van den Ende, André Ferrari, Anthony Sladen, Cédric Richard

► **To cite this version:**

Martijn P.A. van den Ende, André Ferrari, Anthony Sladen, Cédric Richard. Deep Deconvolution for Traffic Analysis with Distributed Acoustic Sensing Data. *IEEE Transactions on Intelligent Transportation Systems*, 2023, 24 (3), pp.2947-2962. 10.1109/TITS.2022.3223084 . hal-03352810v1

**HAL Id: hal-03352810**

**<https://hal.science/hal-03352810v1>**

Submitted on 23 Sep 2021 (v1), last revised 14 Oct 2023 (v2)

**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.

# Deep Deconvolution for Traffic Analysis with Distributed Acoustic Sensing Data

Martijn van den Ende, André Ferrari, Anthony Sladen, and Cédric Richard

**Abstract**—Distributed Acoustic Sensing (DAS) is a novel vibration sensing technology that can be employed to detect vehicles and to analyse traffic flows using existing telecommunication cables. DAS therefore has great potential in future “smart city” developments, such as real-time traffic incident detection. Though previous studies have considered vehicle detection under relatively light traffic conditions, in order for DAS to be a feasible technology in real-world scenarios, detection algorithms need to also perform robustly under heavy traffic conditions. In this study we investigate the potential of roadside DAS for the simultaneous detection and characterisation of the velocity of individual vehicles. To improve the temporal resolution and detection accuracy, we propose a self-supervised Deep Learning approach that deconvolves the characteristic car impulse response from the DAS data, which we refer to as a Deconvolution Auto-Encoder (DAE). We show that deconvolution of the DAS data with our DAE leads to better temporal resolution and detection performance than the original (non-deconvolved) data. We subsequently apply our DAE to a 24-hour traffic cycle, demonstrating the feasibility of our proposed method to process large volumes of DAS data, potentially in near-real time.

**Index Terms**—Distributed Acoustic Sensing, self-supervised Deep Learning, non-blind deconvolution, traffic analysis

## I. INTRODUCTION

Distributed fibre-optic sensing is an emerging technology that enables the measurement of strain and/or temperature at specific locations along fibre-optic cables. In a nutshell, the measurement principle relies on sending pulses of light into one end of an optical fibre and analysing the light that returns to the detector after having been scattered at nanometric scattering sites along the fibre (see [1] for an in-depth discussion). These scattering sites are an inevitable consequence of the fibre production process, so any commercially-available fibre-optic cable is amenable to sensing. Using interferometric techniques, the analysis of the scattered light yields insights into stretching of the fibre or temperature changes at fixed sensing points along the fibre. As such, this technology turns fibre-optic cables into arrays of sensors and has opened up a plethora of opportunities for engineering, geophysics, and environmental sciences. Examples of successful applications include structural integrity monitoring [2], wellbore seismology [3]–[5], permafrost and aquifer monitoring [6], [7], and laboratory rock mechanics experiments [8], [9].

As a subset of distributed fibre-optic sensing, Distributed Acoustic Sensing (DAS) with phase-Optical Time Domain

Reflectometry (usually referred to as  $\Phi$ -OTDR) permits distributed measurements of strain (rate) over distances of several tens of kilometres with a spatial resolution of the order of metres. Unlike counterpart technologies like Optical Frequency Domain Reflectometry, the time-sampling rate can be as high as several kHz, making it highly suitable for large-scale seismic experiments – see [10] and [11] for recent and detailed reviews of the technology and applications of DAS in seismology. While earthquakes and other seismic events like active sources or quarry blasts are of primary interest to many seismic studies using DAS, other sources of ground vibrations have likewise been considered. For example, cars have been analysed as a source of energy in roadside seismic interferometry studies [12]–[14], but also to infer the impact of COVID-19 lockdown measures [15], and to detect traffic flows [16], [17].

One strong advantage of DAS is that it can utilise existing fibre-optic cables used for telecommunication, rather than requiring dedicated cables to be deployed. This ability to tap into existing fibre-optic infrastructures greatly expedites the development of “smart” applications in urban areas, which often feature dense fibre-optic cable networks, such as real-time traffic flow monitoring. While existing technologies like traffic cameras and pneumatic tubes provide point measurements of the number and/or the speed of vehicles at a given location, DAS can provide measurements of passing vehicles every few metres along many kilometres of road. Moreover, DAS measurements are inherently anonymous (as opposed to traffic camera footage) eliminating privacy concerns and avoiding privacy regulations. Thus, DAS is a strong contender for future city-scale traffic analysis applications.

While a human analyst can often easily spot a car in DAS data, it is more challenging to automatically detect individual vehicles and extract their velocity. Some automated extraction methods have been proposed that work well when the cars are well-separated in time and space (e.g. [16], [18]), but for robust traffic analysis in densely-populated urban areas, the extraction procedures need to be able to handle heavy traffic situations (e.g., rush hour), which in turn requires high spatio-temporal precision of the employed method. In this study, we detail a self-supervised Deep Learning method that incorporates prior knowledge of the vehicles’ signals in order to achieve high-resolution detection and velocity estimation performance. In short, our model can be described as a “Deconvolution Auto-Encoder”, and we subsequently perform the detection and velocity estimation on deconvolved DAS data through conventional beamforming techniques. Our results show an improvement of the temporal resolution of the analysis with

M. van den Ende, A. Ferrari, and C. Richard are with the Université Côte d’Azur, OCA, UMR Lagrange, France. E-mail: martijn.vandenende@oca.eu

A. Sladen are with the Université Côte d’Azur, IRD, CNRS, Observatoire de la Côte d’Azur, Géoazur, France

fewer false positive and missed detections, while requiring only minimal efforts to train and deploy the Deep Learning model.

This paper is organised as follows: first, we describe in more detail in Section II the DAS measurement principle and the expression of cars in the DAS data. Also in this section we touch upon the concept of deconvolution that underlies our approach. Next, in Section III we describe our methods in full, after which we show the qualitative and quantitative results of these in Sections IV and V, respectively. Lastly, we frame the results of this study in a broader context of traffic monitoring and “smart” cities in Section VI.

## II. EXPRESSION OF VEHICLES IN DAS DATA

To ensure a full appreciation of our approach, we begin with a detailed exhibition of roadside DAS measurements of cars – see Fig. 1. Permanent deployments of fibre-optic cables are often buried (trenched) or placed within underground conduits, so we assume here that the DAS system is placed alongside a road at some depth below the surface. As a car drives past a given sensing point, the subsurface deforms due to the weight of the vehicle pressing down on the road. This deformation is transferred to the fibre-optic cable, leading to a strain of measurable amplitude. When the car is non-stationary, the changes in the strain field as the car passes by are correspondingly recorded as a strain rate. This quasi-static or geodetic deformation at a point in the subsurface is well described by the Flamant-Boussinesq approximation [14], [19]:

$$u_x(x, y, z) = \frac{F}{4\pi G} \frac{x}{r^2} \left( \frac{z}{r} + \frac{2\nu - 1}{1 + \frac{z}{r}} \right) \quad (1)$$

In this expression,  $u_x$  is the particle displacement in the  $x$ -direction, defined parallel to the road (and the fibre), at a point  $(x, y, z)$  ( $y$  being the horizontal distance perpendicular to the road and  $z$  the depth beneath the surface) relative to the location of a point load positioned at the origin. The distance from the origin is given by  $r = \sqrt{x^2 + y^2 + z^2}$ . The point load exerts a total force  $F$  onto an infinite half-space with uniform shear modulus  $G$  and Poisson’s ratio  $\nu$ . The particle velocity  $\dot{u}_x$  is obtained by differentiating (1) with respect to time, and noting that  $\dot{x}$  is the velocity of the car travelling in the  $x$ -direction. Consequently, the quasi-static signature of a car in the DAS data travels at the speed of the car, and not at seismic wave speeds (see Fig. 1b). Lastly, the DAS system does not measure particle motions, but the average longitudinal strain (rate) between two sensing points, the distance between which is called the gauge length. An expression for the equivalent DAS measurement is found simply as:

$$\varepsilon_{DAS} = \frac{1}{L} \left[ u_x \left( x + \frac{L}{2}, y, z \right) - u_x \left( x - \frac{L}{2}, y, z \right) \right] \quad (2a)$$

$$\dot{\varepsilon}_{DAS} = \frac{1}{L} \left[ \dot{u}_x \left( x + \frac{L}{2}, y, z \right) - \dot{u}_x \left( x - \frac{L}{2}, y, z \right) \right] \quad (2b)$$

for strain and strain rate, respectively ( $L$  representing the gauge length). For typical distances between the fibre and the

vehicle, the characteristic signature of a vehicle in the DAS data is relatively simple (see Fig. 2a).

Conversely, the interactions between the car tyres and the road generate surface waves that travel away from the source point at seismic speeds, and which are relatively complicated to analyse (see Fig. 1c). While the quasi-static deformation induced by a vehicle is mostly controlled by the distance of the fibre with respect to the road, the dynamic deformation patterns resulting from the surface waves are dispersive and highly sensitive to local variations in attenuation. Nonetheless, these high-frequency dispersive surface waves are useful for roadside DAS interferometry studies [12]–[14]. In a previous study of traffic flows by [18], cars were detected mostly based on the dynamic deformation they induced. However, as is clear from Fig. 1d, the DAS recordings in the quasi-static frequency band (0.1-2 Hz) are substantially simpler and more localised in time, and so for this study we focus on this frequency band.

Even though the quasi-static deformation patterns of vehicles are relatively simple and compact, they still comprise a finite duration and spatial extent (roughly 2 s at a speed of 70 km hr<sup>-1</sup> for the main lobes; Fig. 2a). Note that Eq. (1) is parametrised in terms of  $x$ , but which can be substituted for time  $t$  using  $x = vt$  ( $v$  being the speed of the car). When two cars are closely-trailing, their deformation patterns start to overlap, potentially creating more complicated superpositions. Additional complications arise when the DAS system records multiple cars travelling in opposite directions, or when trains of cars pass by. In such cases, simple thresholding and peak detection techniques may easily fail to detect and separate the vehicles in the dataset.

The approach adopted in this work is motivated by the fact that, in principle, the quasi-static deformation pattern of a single car is simple and easy to recognise. While the characteristic deformation pattern measured at a given DAS sensing point may be somewhat obscured by noise (Fig. 2b), a relatively noise-free signature of a car can be extracted by stacking the measurements of multiple sensing points after shifting the recordings according to the speed of the car (Fig. 2a). Since the waveform shown in Fig. 2a represents that of a single car situated at  $t = 0$ , one can interpret this as the impulse response of the car. And since the impulse itself is infinitely localised in time/space (i.e., the vehicle can be treated as a point load), deconvolution of the DAS data should thus yield a series of narrowly-concentrated peaks. The expectation is therefore that the detection and estimation of the velocity of vehicles can be performed in much higher resolution when the analyses are performed on the deconvolved data.

## III. METHODS

### A. DAS data acquisition and processing

The DAS data that are used in this study were acquired during a 16 day measurement campaign that took place from 12 to 28 November 2019 near the city of Mont elimar, France (Fig. 3). A commercial telecommunication fibre that was deployed along a main road connecting the villages of Albal-Romaine, Saint-Thom e, and Valvign eres, was sensed with a Febus A1-R interrogator at a temporal sampling frequency

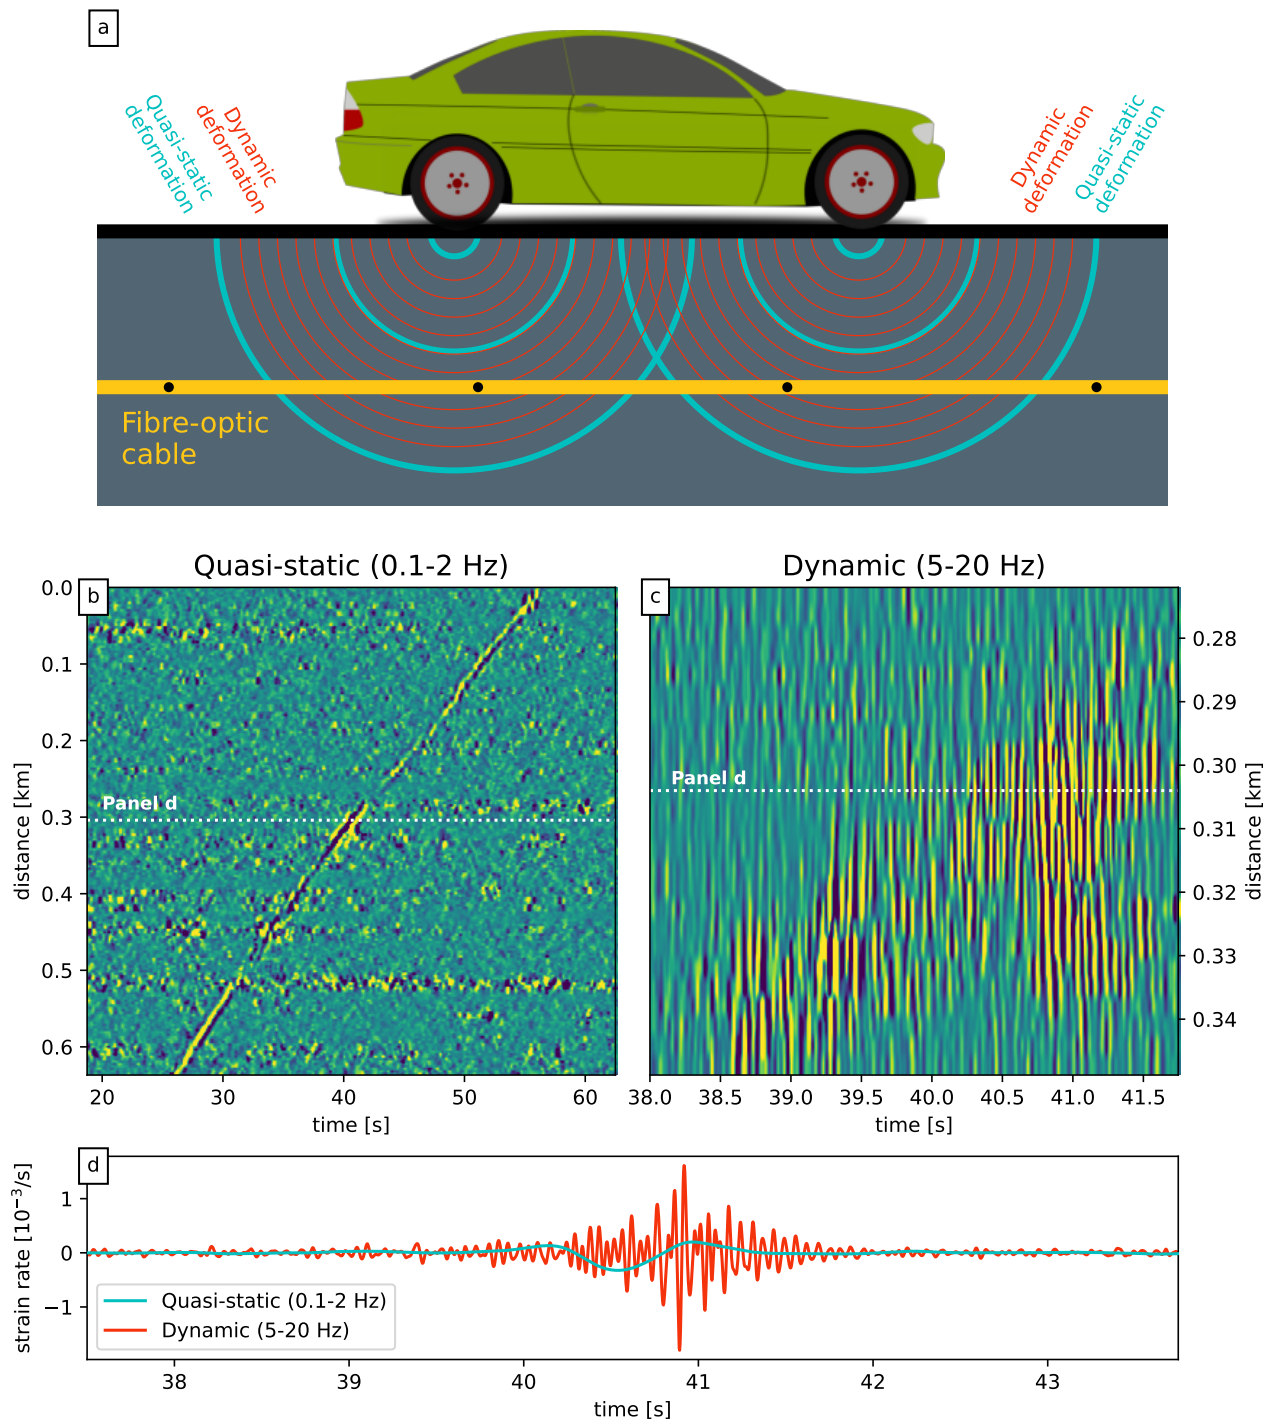


Fig. 1. Synoptic overview of roadside DAS measurements. a) The weight of the car presses down on the ground, inducing quasi-static deformation of the subsurface (cyan). While the car is driving, interactions between the tyres and the road cause dynamic deformation (red). A buried fibre-optic cable (yellow) is able to sense both. The black dots approximately indicate the distance between sensing points; b) The quasi-static deformation of a vehicle is easily recognisable and produces a relatively simple pattern that can be tracked in space and time; c) Even though its amplitude is higher than that of the quasi-static deformation, the dynamic deformation of a vehicle is more complex; d) The strain rate recordings at one specific DAS channel (indicated by the white dotted line in panels b and c).

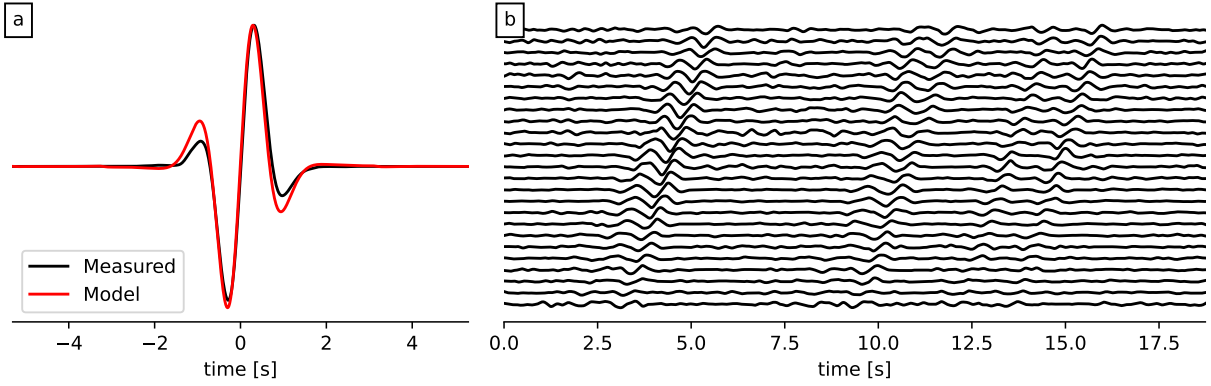


Fig. 2. a) The measured impulse response of a single car (obtained from stacking multiple DAS channels) as compared to a model based on Eq. (2b); b) Example of DAS recordings of multiple cars travelling in the same direction (towards the interrogator). At around 12.5 s two cars are seen that are closely behind one another, such that their deformation patterns start to overlap.

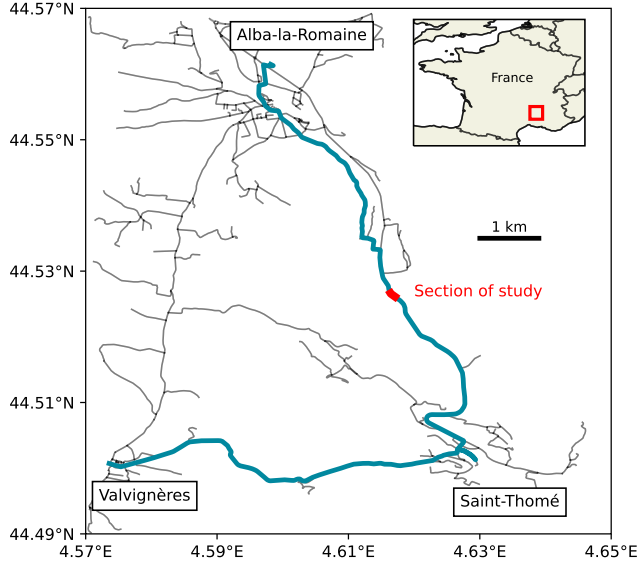


Fig. 3. Geographic location of the fibre-optic cable (blue) and the segment of interest (red). For reference, the road network is included in black. The cable is deployed alongside the main road that connects the three villages indicated.

of 400 Hz. The DAS channel spacing was set equal to the gauge length of 3.2 m, for a total fibre length of 14 km. Along this cable we identified a segment between 4.88-4.96 km (indicated in red in Fig. 3) that displayed a very good signal-to-noise ratio, most likely owing to the high-quality coupling between the trenched cable and the subsurface. This segment features 24 sensing points in total. We extracted 48-hours' worth of data for the selected segment, which was subsequently filtered in a 0.1-2 Hz pass band and downsampled to 50 Hz.

The road under consideration is a two-lane major road with a speed limit of 80 km hr<sup>-1</sup> and serves both personal vehicles and lorries of various sizes, the main difference between these two being the strain rate amplitude at which they appear in the DAS data. The DAS system clearly records vehicles coming from both directions. The selected road segment is straight and

cars typically maintain a constant speed throughout.

### B. FISTA deconvolution algorithm

Before describing the Deep Learning model and training procedure, we will detail a conventional deconvolution algorithm that was used as a reference. The task of deconvolving the DAS data recorded at the  $q$ -th sensing point,  $y_q$ , is expressed by the following objective function:

$$\hat{x}_q(t) = \arg \min_x \left\{ \frac{1}{2} \| [k * x]_t - y_q \|_2^2 + \rho \| x \|_1 \right\} \quad (3)$$

in which  $[k * x]_t$  represents the convolution between the (known) impulse response  $k$  and the underlying impulse model  $x$ , and  $\rho$  controls the strength of the  $\ell_1$ -regularisation on  $x$ . The regularisation is included to acknowledge the notion that the number cars on the road is relatively small; at most, one car every two seconds is expected to pass a given sensing point. One commonly used algorithm to solve this optimisation problem is the Iterative Shrinkage Thresholding Algorithm (ISTA; [20], [21]). For this study, we adopt an accelerated version of ISTA (Fast-ISTA or FISTA) due to [22], which exhibits faster convergence guarantees. This iterative algorithm is implemented in JAX [23] following Algorithm 1:

---

#### Algorithm 1: FISTA deconvolution algorithm

---

**Require:** Convolution kernel  $\mathbf{K}$ , DAS data  $\mathbf{y}$ , step size  $\delta$ , regularisation strength  $\rho$ ;  
**Initialise:**  $\mathbf{x}_0 \sim \mathcal{N}$ ,  $\mathbf{r}_0 \leftarrow \mathbf{x}_0$ ,  $t = 1$ ;  
**while**  $i < i_{max}$  **do**  
     $\mathbf{x}_{i+1} \leftarrow \text{soft}(\mathbf{r} - \delta \mathbf{K}^\top [\mathbf{K} \mathbf{x}_i - \mathbf{y}] | \rho)$ ;  
     $t_{i+1} \leftarrow \frac{1}{2} (1 + \sqrt{1 + 4t_i^2})$ ;  
     $\mathbf{r}_{i+1} \leftarrow \mathbf{x}_i + \left( \frac{t_i - 1}{t_{i+1}} \right) (\mathbf{x}_i - \mathbf{x}_{i-1})$ ;  
     $i \leftarrow i + 1$ ;  
**end**

---

In Algorithm 1,  $\text{soft}(\cdot|\rho)$  denotes the soft-thresholding function with threshold parameter  $\rho$ , defined as:

$$\text{soft}(x|\rho) = \begin{cases} x - \rho, & \text{if } x > \rho \\ 0, & \text{if } |x| \leq \rho \\ x + \rho, & \text{if } x < -\rho \end{cases} \quad (4)$$

To deconvolve the DAS data, we split the data into windows each 100 s in length and the optimisation is repeated for each time window, which are subsequently concatenated to yield a deconvolved dataset. The update step size of the (F)ISTA algorithm is controlled by the largest eigenvalue  $\lambda_{max}$  of  $\mathbf{K}\mathbf{K}^T$ , where  $\mathbf{K}$  is the convolution matrix constructed from the impulse response  $k$  such that  $\mathbf{K}\mathbf{x} = [k * x]_t$ . In this study we set the step size equal to  $\delta = \lambda_{max}/10$  to ensure stable iterations, and we empirically choose regularisation parameter  $\rho$  as  $1.5 \times 10^{-3}$ .

In the FISTA approach, each DAS channel is deconvolved independently of the others (i.e., the algorithm is a single-channel deconvolution procedure). This is certainly sub-optimal, since the most distinguishing feature of vehicles in the DAS data is their spatio-temporal coherence (as in Fig. 2b). We can therefore attempt to improve on this method with a Deep Learning approach.

### C. Deconvolution Auto-Encoder

1) *Architecture*: The Deep Learning architecture adopted for this study (see Fig. 4) is that of a light-weight U-Net Auto-Encoder [24], which takes as an input a set of  $N_q = 24$  consecutive waveforms of  $N_t = 1024$  time samples (20.48 s) in duration, organised in an  $N_q \times N_t$  matrix. The input is passed onto a U-Net Auto-Encoder which comprises 3 convolutional layers, followed by 3 encoder blocks, each featuring a downsampling operation and 3 convolutional layers. The learnable kernels for the convolution layers are of size  $3 \times 5$ , with the number of filters doubling after each downsampling operation, starting at 8 filters. The maxpooling operation downsamples the data by a factor 2 along the DAS sensor axis and by a factor 4 along the time axis (i.e., the maxpooling kernel and strides are of size  $2 \times 4$ ). After the encoding operations, the latent space is consequently of size  $3 \times 16 \times 64$ . The decoder reverses the encoding operations with 3 blocks of bilinear upsampling. The characteristic feature of the U-Net architecture is the presence of skip-connections, which directly connect the output of one encoder block with the corresponding (diametrically opposite) decoder block. Hence, after each upsampling operation, the output from the corresponding encoder block is concatenated to the upsampled data along the feature (channel) axis. The concatenation is followed by 3 convolutional layers. Lastly, the output layer is a single convolutional layer with 1 output channel and ReLU activation, which enforces positivity and sparsity in the model output. Each convolutional layer other than the output layer is followed by a Swish non-linearity [25], and an insertion layer of Gaussian additive noise with zero mean and unit variance for further regularisation.

The key step to turn this architecture into a self-supervised deconvolution algorithm, is to convolve the model output

$x$  with the impulse response  $k$  (Fig. 2a) along the time-axis. Hence, for a batch of inputs  $\{y_1, y_2, \dots, y_{N_b}\}$  ( $x, y \in \mathbb{R}^{N_q \times N_t}$ ,  $k \in \mathbb{R}^{1 \times N_k}$ ,  $N_b = 128$ ), the following objective function is minimised:

$$\mathcal{L} = \frac{1}{N_b} \sum_{i=1}^{N_b} (\| [k * x_i]_t - y_i \|_2^2 + \rho \| x_i \|_1) \quad (5)$$

Again,  $[k * x_i]_t$  denotes the convolution between  $k$  and  $x$  along the time-axis, and  $\rho$  is a parameter that controls the sparsity on  $x$ . Here we set  $\rho = 10$ , which seems to strike a good balance between sparsity and fidelity. Optimising this training objective will naturally lead to a model output  $\hat{x}$  which, after convolution with  $k$ , recovers  $y$ , and thus yields a deconvolution algorithm. We hence refer to this Deep Learning approach as a *Deconvolution Auto-Encoder* (or DAE for short). Also note that, like any self-supervised method, no ground truth or knowledge of the underlying impulse model  $x$  is required; the only information provided to the model is the impulse response  $k$ . Moreover, unlike the FISTA algorithm detailed in the previous section, the DAE utilises spatio-temporal correlations to construct the impulse model, and hence can incorporate spatio-temporal characteristics of vehicles (like their expected move-out and amplitude modulation) as inductive biases during training. It is expected that this gives the DAE an edge over the conventional FISTA method.

2) *Training procedure*: Empirically, we observed during exploratory experiments that the DAE converges faster and is more robust when trained on DAS strain data, rather than strain rate data. While we have not explored the underlying reason for this, we hypothesise that this might be related to symmetry: in strain, the impulse response of a vehicle is approximately symmetric, whereas in strain rate it is anti-symmetric. The superposition of multiple shifted anti-symmetric signals leads to destructive interference, which may hamper convergence. We therefore convert the input data from strain rate into strain by frequency-domain integration, i.e., the data are divided by  $-2j\pi n$  in the frequency domain, with temporal frequency  $n$ , and transformed back into the time-domain. Likewise, the impulse response is converted into a strain representation (shown in Fig. 4). Since this strain representation is surprisingly close to a Ricker wavelet (also known as the ‘‘Mexican hat’’), we replace the empirical impulse response with a Ricker wavelet, which has more favourable spectral characteristics while minimally sacrificing reconstruction fidelity.

The strain data is subsequently divided into training and validation splits. To ensure that the statistics of the training and validation datasets are similar, we split the 48-hour DAS dataset 50-50 so that both the training and validation sets include a 24-hour traffic cycle (including rush hours). These datasets are fed into a dataloader that creates batches of 128 samples  $\times$  24 DAS channels  $\times$  1024 time samples by extracting random slices from each respective dataset. In addition, random flips along the channel and time axes are performed to create more data diversity. New batches are created after each epoch, and the time slices are randomly selected each time. Because of the random sampling the concept of one epoch is not well defined, so we arbitrarily define one epoch

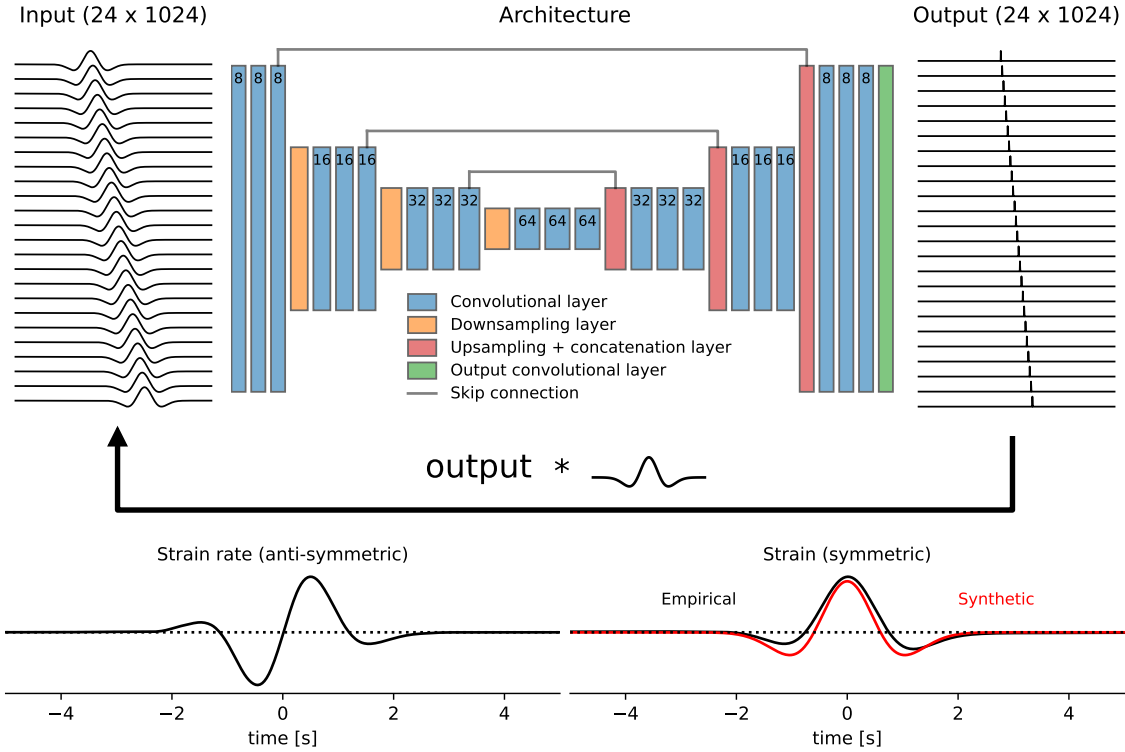


Fig. 4. Conceptual overview of the Deconvolution Auto-Encoder (DAE). The DAS data is fed into a U-Net Auto-Encoder, of which the output is convolved with a known impulse response (along the time-axis) to obtain a reconstruction of the input. The number of convolution filters in each layer are as indicated. The deconvolution performance is improved by converting the DAS data from strain rate into strain, and subsequently replacing the measured impulse response with a Ricker (“Mexican hat”) wavelet.

to comprise 10,000 samples in total. We then train the model for 1000 epochs (taking less than 12 hours on a single Nvidia Quadro P4000 GPU), at which point the performance saturates at the same level for both the training and validation set, suggesting minimal overfitting (see Supplementary Figure S1). The model is relatively small (only 305,089 trainable parameters), and so its capacity for overfitting is minimal. Increasing the size of the model did not lead to better performance.

At test time, the data in the validation set are split in regular non-overlapping 1024-sample slices, which are fed into the model. The model output  $x$  for each slice is concatenated to create the deconvolved dataset.

#### D. MUSIC beamforming and slowness correction protocol

Beamforming is a commonly-used array processing technique that estimates the direction of arrival of various signals impinging on a sensor array [26]. In seismology, seismic beamforming (and associated back-projection) is used to estimate the direction of arrival and the apparent velocity of incoming seismic waves propagating across a seismometer array [27], [28]. Since DAS constitutes a (curvi-)linear array of single-component seismic sensors, it is amenable to beamforming analysis [29], and established techniques can be adopted for the analysis of the vibrations induced by moving vehicles. In this study, we employ frequency-domain MUSIC beamforming [30], which typically exhibits better resolution

than (time-domain) delay-and-sum beamforming, particularly for multiple sources [27].

Consider a single car travelling with signed velocity  $v$ . At a given sensor  $q$  and time instant  $t$ , the signal  $y_q$  recorded by the  $q$ -th sensor can be represented in the frequency domain as:

$$Y_q(n) = z(n)e^{-j2\pi n\Delta t_q} + e_q(n) \quad (6)$$

where  $z$  is the frequency-domain representation of the signal emitted by the car (with frequency  $n$ ),  $e_q$  is the noise recorded at the  $q$ -th sensor. The exponential multiplying the signal represents the phase shift due to a time delay  $\Delta t_q$ , which in turn is related to the position of the car with respect to the sensor. For an array of equidistant sensors separated by a distance  $L$ , we can express this time delay as  $\Delta t_q = qL/v = qLs$  ( $s$  denoting the reciprocal velocity, or slowness, of the car). For the purpose of estimating a vehicle’s velocity, we discretise the range of attainable slowness values and decompose the slowness into a reference slowness  $s_0$  and a perturbation  $\Delta s_i$ , i.e.  $s_i = s_0 + \Delta s_i$ . By taking  $s_0$  to be the reciprocal of the expected mean velocity of the cars (or the road’s speed limit), this constant term can be factored out of Eq. (6), i.e.:

$$Y_q(n)e^{j2\pi nqLs_0} \equiv Y'_q(n) = z(n)a_q(n) + e'_q(n) \quad (7)$$

with  $a_q(n) \equiv e^{-j2\pi nqL\Delta s_i}$  being the  $q$ -th component of the steering vector  $a(n)$ . In the time domain,  $y'_q$  now represents a DAS measurement that is aligned according to  $s_0 = v_{\text{ref}}^{-1}$ ,

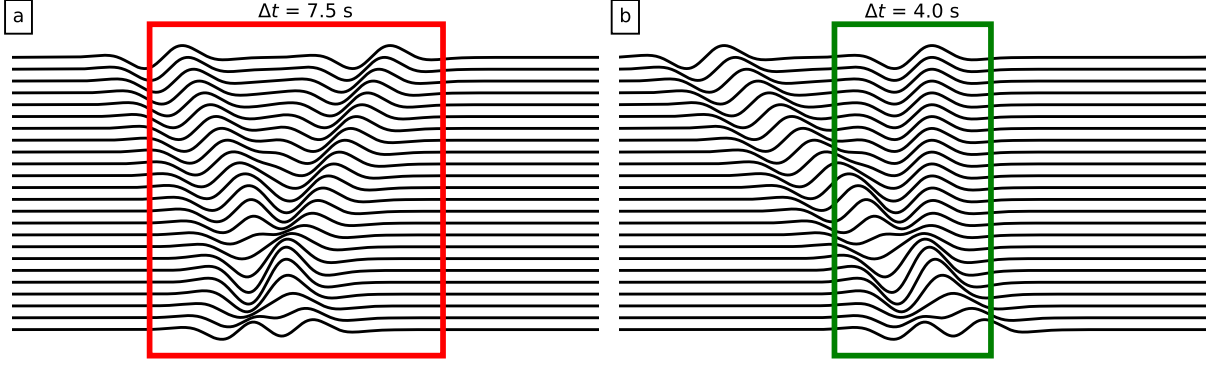


Fig. 5. Schematic illustration of the concept of factoring out the reference slowness. a) The time-window needed to comprise a single vehicle, including its move-out, is 7.5 s, and also includes most of the deformation induced by a second vehicle travelling in the opposite direction; b) The time-window needed to comprise a single vehicle after shifting the waveforms according to a reference slowness is only 4 s, and excludes most of the second vehicle.

the motivation for which will be clarified in a moment. We then follow standard MUSIC beamforming protocol [30] by estimating the element  $(i, j)$  of the covariance matrix  $\mathbf{C}$  as:

$$C_{ij} = \frac{1}{|B|} \sum_{n \in B} Y_i' \overline{Y_j'} \quad (8)$$

where the sum extends over all frequencies  $n$  in a narrow frequency band  $B$  (here, 0.5-2 Hz). To reduce boundary effects in the Fourier transformation, we estimate  $\mathbf{C}$  using the multi-taper method [31]. The MUSIC pseudo-power spectrum is subsequently obtained as the reciprocal of the projection of the steering vectors onto the noise space of  $\mathbf{C}$ , spanned by the  $N_q - N_s$  smallest eigenvectors (with  $N_s$  denoting the presumed number of sources; in this study,  $N_s = 2$ ).

To see the advantage of performing the slowness factorisation, we consider a situation as depicted in Fig. 5a. Here, two vehicles are present in the (synthetic) DAS data, one travelling away from the interrogator (top-left to bottom-right) with a velocity  $+v_{\text{ref}}$  and one travelling in the opposite direction with a velocity  $-v_{\text{ref}}$ . The time window that is needed to fully comprise the latter vehicle is 7.5 s in duration. Consequently, almost all of the deformation induced by the former vehicle is included in this time window, and this may interfere with the overall waveform coherence. Conversely, after factoring out the reference slowness  $s_0 = -v_{\text{ref}}^{-1}$ , the vehicle travelling towards the interrogator (bottom-left to top-right) now forms a vertical trace of aligned waveforms, and a much tighter time window of 4 s can be defined that still fully comprises the vehicle's signature. The vehicle travelling in the other direction has mostly been shifted out of the time window, so that ultimately the MUSIC pseudo-power is dominated by the vehicle with negative velocity. The opposite result is achieved by setting  $s_0 = +v_{\text{ref}}^{-1}$ , highlighting the vehicle travelling with positive velocity.

Our beamforming protocol is then as follows: we begin by factoring out the reference slowness  $s_0 = v_{\text{ref}}^{-1}$  from the input data, taking  $v_{\text{ref}}$  to equal the speed limit of 80 km hr<sup>-1</sup>. Note that the input data can either be the original DAS data, or the data deconvolved by FISTA or the Deconvolution Auto-

Encoder. We then take a beamforming window of 24 DAS channels by 1.6 s in the case of original data, or 1.0 s in the case of deconvolved data. These window sizes have empirically provided the best trade-off between time-resolution on the one hand, and beamforming accuracy and robustness on the other. We then slide this window through the input data with strides of 0.2 s, estimating  $\mathbf{C}$  for each time window and computing the MUSIC pseudo-spectrum. This procedure is repeated by shifting the input data according to  $s_0 = -v_{\text{ref}}^{-1}$ . The resulting output of the algorithm for each time window is a distribution of beampower over the discretised slowness perturbation  $\Delta s_i$ , which are converted back into velocity as  $v_i = (s_0 + \Delta s_i)^{-1}$ . As such, we obtain a distribution of beampower as a function of (signed) velocity and time.

In order to identify individual vehicles in the beampower distributions, we employ a basic local peak estimation algorithm, treating the beampower obtained for each traffic direction (positive and negative  $v_i$ ) as an image with a velocity-axis and a time-axis. We calculate the maximum beampower along the velocity-axis, which gives the maximum beampower as a function of time. To these time series we apply the “find\_peaks” algorithm from the SciPy library [32], requiring a minimum distance between subsequent peaks of at least 1 s (since it is highly unlikely to find two cars trailing within 1 s of one another). The second criterion for determining peaks in the beampower is given by the topological persistence, or “prominence” threshold, of a peak, which is a measure of how well a peak stands out from its surroundings. This threshold is a hyperparameter that will be tuned to achieve an optimal trade-off between correct detections, false positives, and false negatives. Finally, for each detected peak in time the corresponding velocity is obtained as the location of maximum beampower along the velocity-axis.

## IV. QUALITATIVE RESULTS

### A. Deconvolution

As a first verification of the proposed deconvolution methods, we briefly inspect some examples of deconvolved data – see Fig. 6. Fig. 6a shows two cars travelling away from



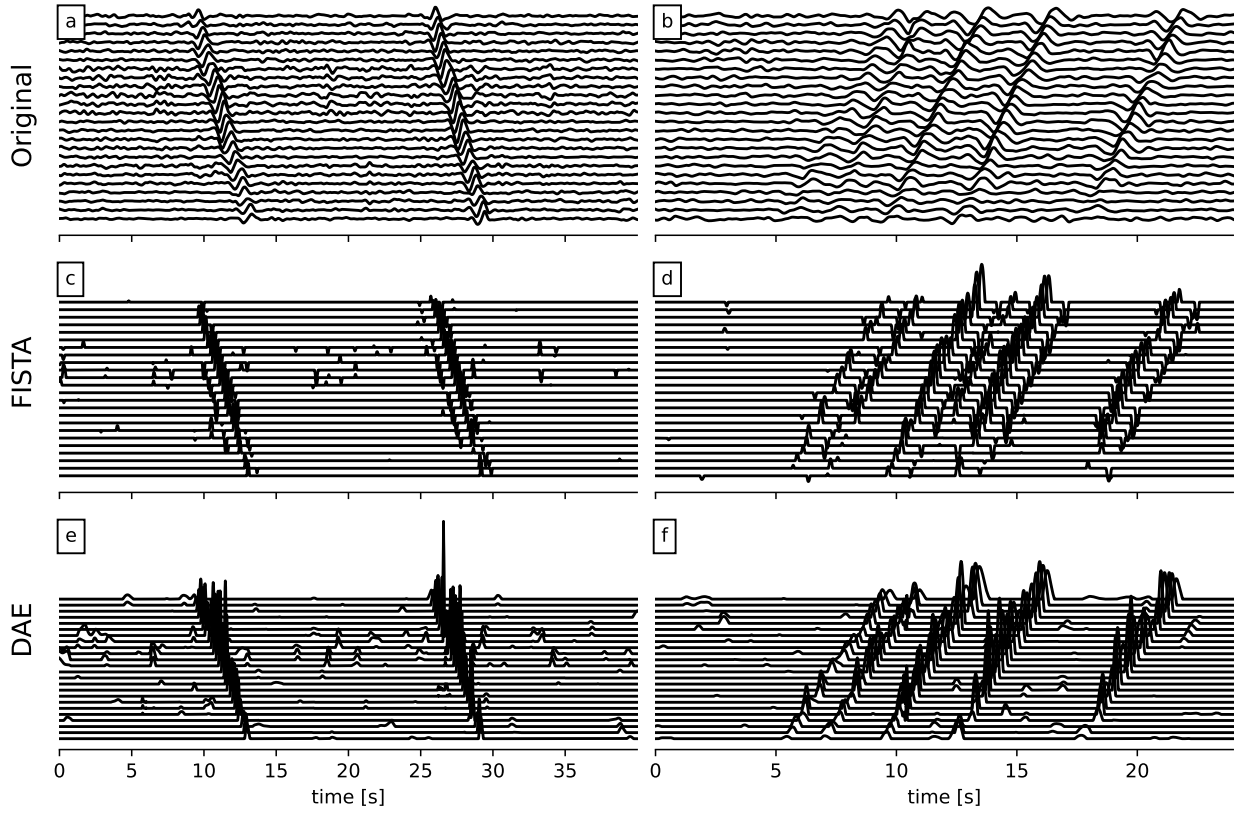


Fig. 6. Examples of the deconvolution performance of the proposed methods. a,b) Original DAS data; c,d) Deconvolved data using the FISTA method; e,f) Deconvolved data using the Deconvolution Auto-Encoder.

the interrogator, separated in time by about 15 s such that their induced deformation patterns do not interfere. For both the FISTA and Deconvolution Auto-Encoder methods (Fig. 6b and c, respectively), we see a sharp ridge of narrowly-localised impulses at the location of each car. The DAE shows some parasitic peaks at seemingly random locations, which, owing to their spatio-temporal incoherence, will not affect the beamforming performance. The FISTA deconvolution results show a minor secondary ridge of negative impulses for each car, but since these are aligned with the main ridge, they will be indistinguishable in the beamforming analysis.

A more challenging example is shown in Fig. 6b, which features a train of 5 closely-trailing cars. The closest time-separation between the first two vehicles is less than 2 s, and so this example is representative for the challenge that this study aims to address. Like in the previous example, the FISTA deconvolution approach produces multiple ridges of impulses per vehicle, while the DAE suffers from some randomly positioned parasitic peaks. But, as aforementioned, these artificial features do not necessarily affect the beamforming results.

Overall, the FISTA deconvolution results are more sparse and narrowly-localised in time, while the DAE seems to produce more robust and easily-interpretable results. However, the total computation time to process roughly 4-hours' worth of DAS data with FISTA was over 7 minutes, while passing the same amount of data through the DAE took just over 1 second. This  $> 400$  times speed-up is a significant factor

of consideration when processing large volumes of DAS data (potentially in real-time).

### B. Beamforming performance: light traffic

Next, we consider the performance of the beamforming analysis in a scenario of light traffic, such that all vehicles are clearly separated in space and time – see Fig. 7. In the first example (Fig. 7a, the same as Fig. 6a), two cars are travelling in the same direction (away from the interrogator), while in the second example (Fig. 7b) two cars are travelling away and one is travelling towards the interrogator. In panels c-h we show the (logarithm of the) MUSIC pseudo-power in red for positive velocities (away from the interrogator) and in blue for negative velocities (towards the interrogator). The estimations of the local peaks above the persistence threshold are included as cyan disks. For these relatively simple examples, the detection performance on the original and FISTA-deconvolved data is good but not perfect, with two false positives registered for the second example in the original dataset (Fig. 7d) and one false positive detection in each of the examples in the deconvolved datasets (Fig. 7e and f). The detection performance on the DAE-deconvolved data is perfect (no false or missed detections). The estimation of the velocity is consistent across all datasets, with a resolution of approximately  $\pm 5 \text{ km hr}^{-1}$ . However, beamforming on the deconvolved data clearly produces peaks in beampower that are much more localised in time.

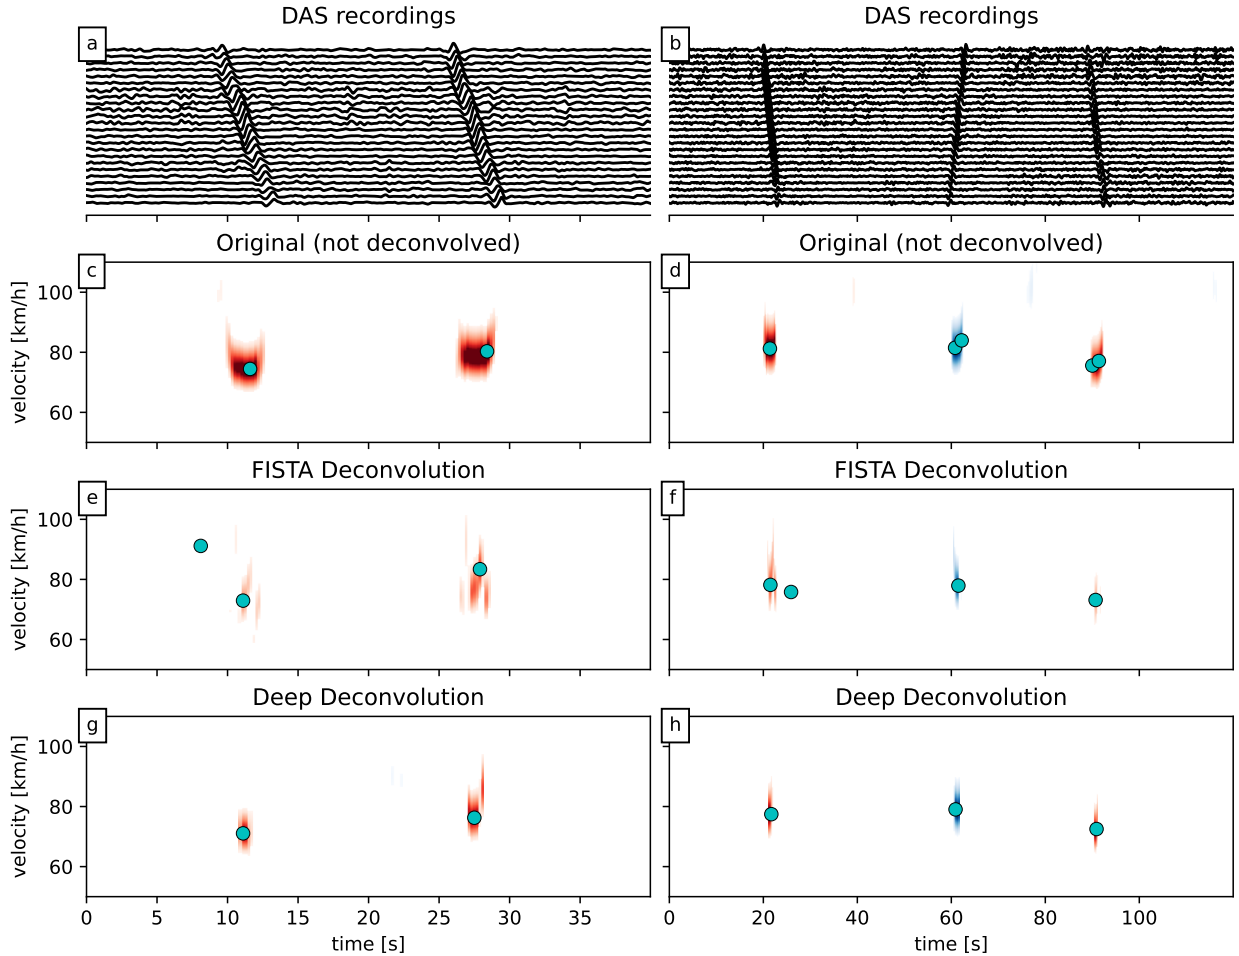


Fig. 7. Two examples of beamforming performance during light traffic conditions. The top panels (a, b) show the original DAS recordings filtered in a 0.1-2 Hz pass band. The logarithm of the MUSIC pseudo-power in panels (b-h) is indicated in red for positive velocities (vehicles travelling away from the interrogator), in blue for negative velocities (towards the interrogator). The cyan disks indicate the local peak in beampower.

One challenging aspect of the DAS data set is that it contains both personal vehicles and heavy-duty trucks/lorries, the latter being less common but exhibiting a much higher-amplitude footprint in the data. Moreover, these trucks are sufficiently heavy to generate low-frequency surface waves, which fortunately are of lower amplitude but are still visible in the data. Fig. 8a shows an example of two trucks separated about 10 s in time, with the first truck clearly generating relatively strong surface waves (e.g. the quasi-vertical line at around 10 s). When attempting to deconvolve these DAS data with the FISTA method (Fig. 8c), we see a relatively dense impulse model emerging with a temporal extent that is even larger than the original data. This is clearly undesirable, and when comparing the beamforming performance on the original with the FISTA-deconvolved data, we see that the resolution of the deconvolved data is worse (c.f. Fig. 8b with d). Lastly, the DAE manages to produce an impulse model that is rather similar to that of cars, being fairly sparse and well-localised (Fig. 8e). However, additional parasitic ridges of impulses are produced, which result in isolated clusters of beampower and are erroneously classified by the peak detector as individual vehicles (Fig. 8f).

In summary of this section, beamforming under light traffic conditions leads to similar car detection performance when considering the original DAS data, or when considering the deconvolved data (using either FISTA or the DAE). In the case of heavy trucks, beamforming the deconvolved data leads to lower-resolution detections (FISTA) or to spurious detections (FISTA and DAE). Fortunately, trucks can be easily identified based on the amplitude of the deformation they induce, and so can be dealt with in an appropriate manner, if needed.

### C. Beamforming performance: heavy traffic

While the problem of detecting cars in light traffic conditions is easily solved with standard beamforming techniques without additional treatment of the data, new challenges arise in heavy traffic conditions when vehicles are closely trailing and their signatures start to overlap – three examples of this are shown in Fig. 9. In the first example (Fig. 9a), there are two cars that are trailing within several seconds of one another, but sufficiently distant to be easily separable by eye. When performing the beamforming on these data, the peaks in beampower are overlapping owing to the size of the beamforming window (Fig. 9d). As a result, the peak detector

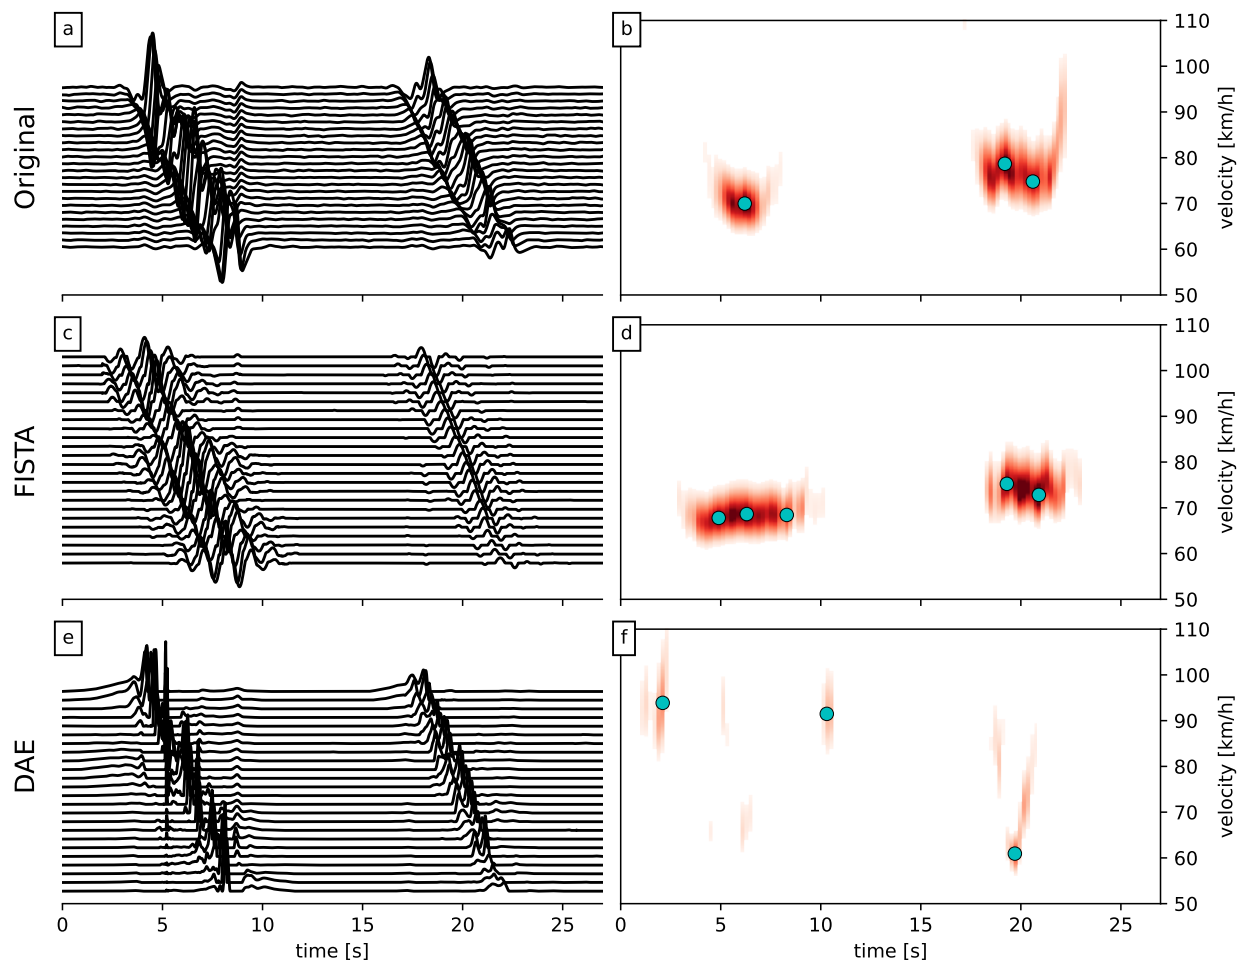


Fig. 8. Example of deconvolution and beamforming performance for heavy trucks. The left panels show a) the original DAS data, c) the deconvolved data using the FISTA method, and e) the deconvolved data using the Deconvolution Auto-Encoder. The corresponding MUSIC pseudo-power for each dataset is given in the right panels (b, d, f).

produces some false positives. Looking at the beamforming performance on the data deconvolved with FISTA (Fig. 9g) and with the DAE (Fig. 9j), we observe clearly separable peaks in beampower, facilitating more robust detections from the peak detection algorithm. In spite of this, two additional false detections are still raised for the FISTA data.

The second example (Fig. 9b; the same as in Fig. 6b) is more challenging, as the vehicles are sufficiently close to have overlapping deformation patterns. In total, 5 cars are present in this example, and the first two cars (starting at 5 s) are separated in time by less than 2 s. Performing the beamforming on the non-deconvolved data, we surprisingly recover 4 out of the 5 vehicles (Fig. 9e) with an additional false positive, even though the clusters of beampower overlap. The first car in the train did not produce any significant beampower, and is therefore undetectable. Performing the detection on the FISTA data gives a perfect score (all vehicles detected, no false positives; Fig. 9h), and the DAE data has one missed detection (Fig. 9k).

The third and last example (Fig. 9c) is the most challenging one. While it is hard to see clearly by eye, it is estimated that there are 4 vehicles travelling away from the interrogator,

followed by 4 vehicles travelling towards the interrogator. Out of these, the standard beamformer detects 3 in each direction (Fig. 9f). The detection on the FISTA data also yields 3 and 3, but for each direction at least 1 vehicle is estimated to travel very slowly ( $50 \text{ km hr}^{-1}$ ; Fig. 9i), which is much slower than anticipated and likely inaccurate. Lastly, the detection on the DAE data achieves a perfect score, facilitated by clearly localised clusters of beampower (Fig. 9l).

One might argue that the detection performance depends on the topological persistence (or prominence) threshold that is a hyperparameter for the peak detection algorithm. A higher threshold will lead to fewer peaks in beampower being classified as a detection, and so can lead to fewer false positives (or conversely, fewer missed detections for a lower threshold). Naturally there is a trade-off to be made here, which will be investigated in more detail in the next section. What is useful to mention here, is that the prominence of a peak is controlled by the “background” level of beampower, or the level of beampower in between two peaks. As can be seen in Fig. 9d, the beampower in between the first two vehicles does not fully return to the background level, which is a direct consequence of the time-window used in the beamforming. When doing the

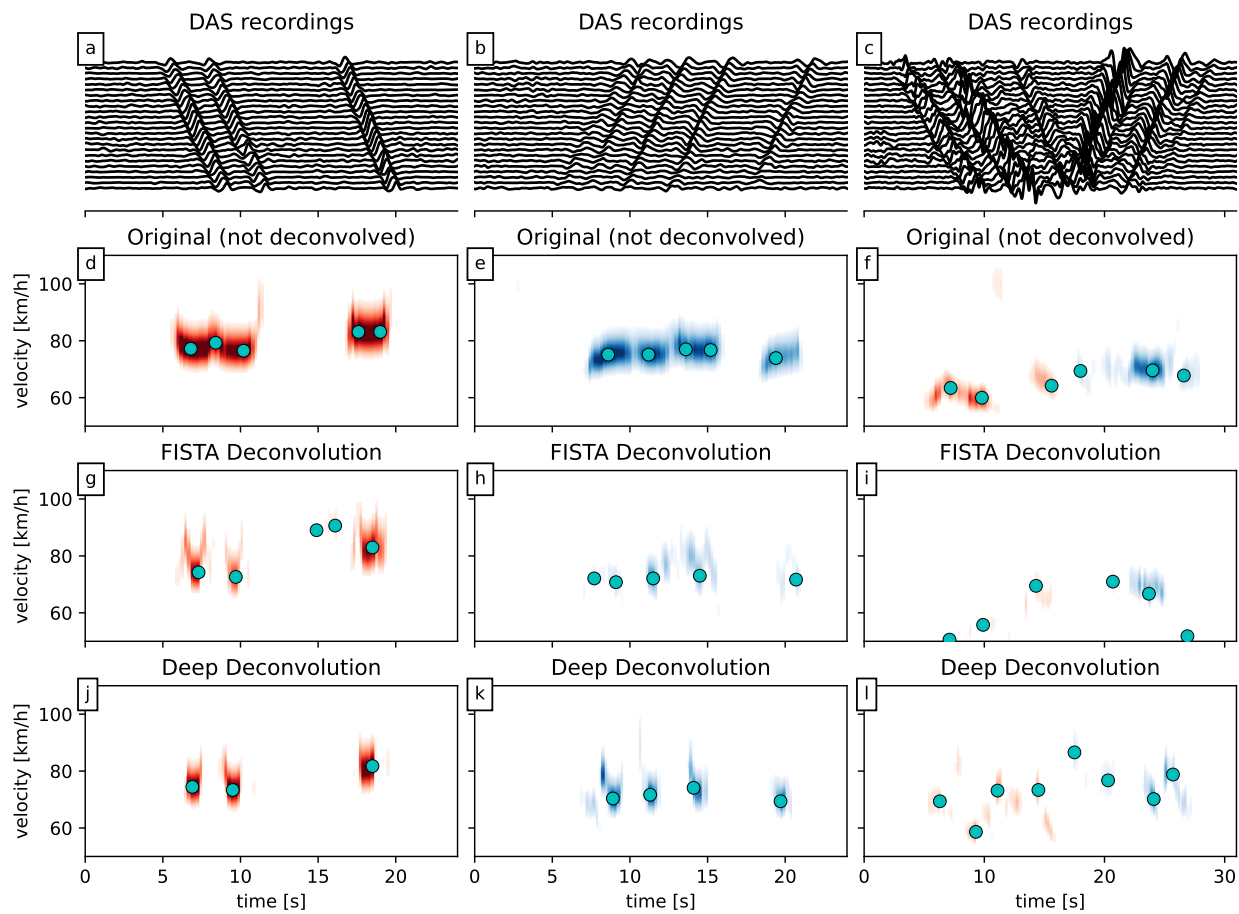


Fig. 9. Three examples of beamforming performance during heavy traffic conditions, ordered by increasing difficulty. The top panels (a, b, c) show the original DAS recordings filtered in a 0.1-2 Hz pass band. The logarithm of the MUSIC pseudo-power in panels (b-l) is indicated in red for positive velocities (vehicles travelling away from the interrogator), in blue for negative velocities (towards the interrogator). The cyan disks indicate the local peak in beampower.

beamforming on the deconvolved data, a more narrow time-window is achievable, leading to a more pronounced drop in beampower in between subsequent vehicles (Fig. 9g and j). In turn, this would permit a larger prominence threshold as each peak stands out more clearly, leading to fewer false positives. Fortunately, it is relatively straightforward to perform supplementary analyses to determine the appropriate prominent thresholds that optimally trades-off between false positives and missed detections.

## V. QUANTITATIVE RESULTS

### A. Beampower peak detection

Performing the beampower on the original or deconvolved data yields clusters of beampower, which need to be converted into detections and speed estimations of individual cars. The beampower peak detector as described in Section III-D features the peak prominence as the sole tuning parameter, which may need to vary between different datasets, as well as between the two traffic directions (away from or towards the interrogator). For each of these cases, we determine the optimal peak prominence value based on the trade-off between the true positive rate, the false positive rate, and the false discovery rate, through a comparison between the inferred

TABLE I  
DETECTION PERFORMANCE CHARACTERISTICS FOR BEAMFORMING ON THE ORIGINAL (NON-DECONVOLVED) DATA, AND DATA DECONVOLVED WITH THE FISTA METHOD AND WITH THE DAE. THE FIRST AND SECOND NUMBER IN EACH CELL CORRESPOND WITH VEHICLES TRAVELLING AWAY FROM AND TOWARDS THE INTERROGATOR, RESPECTIVELY. TP: TRUE POSITIVES; FP: FALSE POSITIVES; FN: FALSE NEGATIVES; TPR: TRUE POSITIVE RATE (ALSO KNOWN AS RECALL); FPR: FALSE POSITIVE RATE; FNR: FALSE NEGATIVE RATE.

	Original	FISTA	DAE
# of vehicles	332 forward / 97 backward		
Prominence	1.0 / 0.8	0.5 / 0.35	1.2 / 0.65
TP	245 / 72	253 / 66	266 / 76
FP	181 / 194	85 / 165	36 / 78
FN	87 / 25	79 / 31	66 / 21
TPR	0.74 / 0.74	0.76 / 0.68	0.80 / 0.78
FPR ( $\times 10^{-3}$ )	10 / 11	4.9 / 9.4	2.0 / 4.5
FNR	0.26 / 0.26	0.24 / 0.32	0.20 / 0.22
Precision	0.58 / 0.27	0.75 / 0.29	0.88 / 0.49

peaks and a manually picked dataset. This manually-labelled dataset was generated by marking the timings of vehicles travelling either towards or away from the interrogator during 1 hour of dense traffic. In total, 429 vehicles were identified during this hour, 332 of which travelling away from the interrogator, and 97 travelling towards the interrogator. For

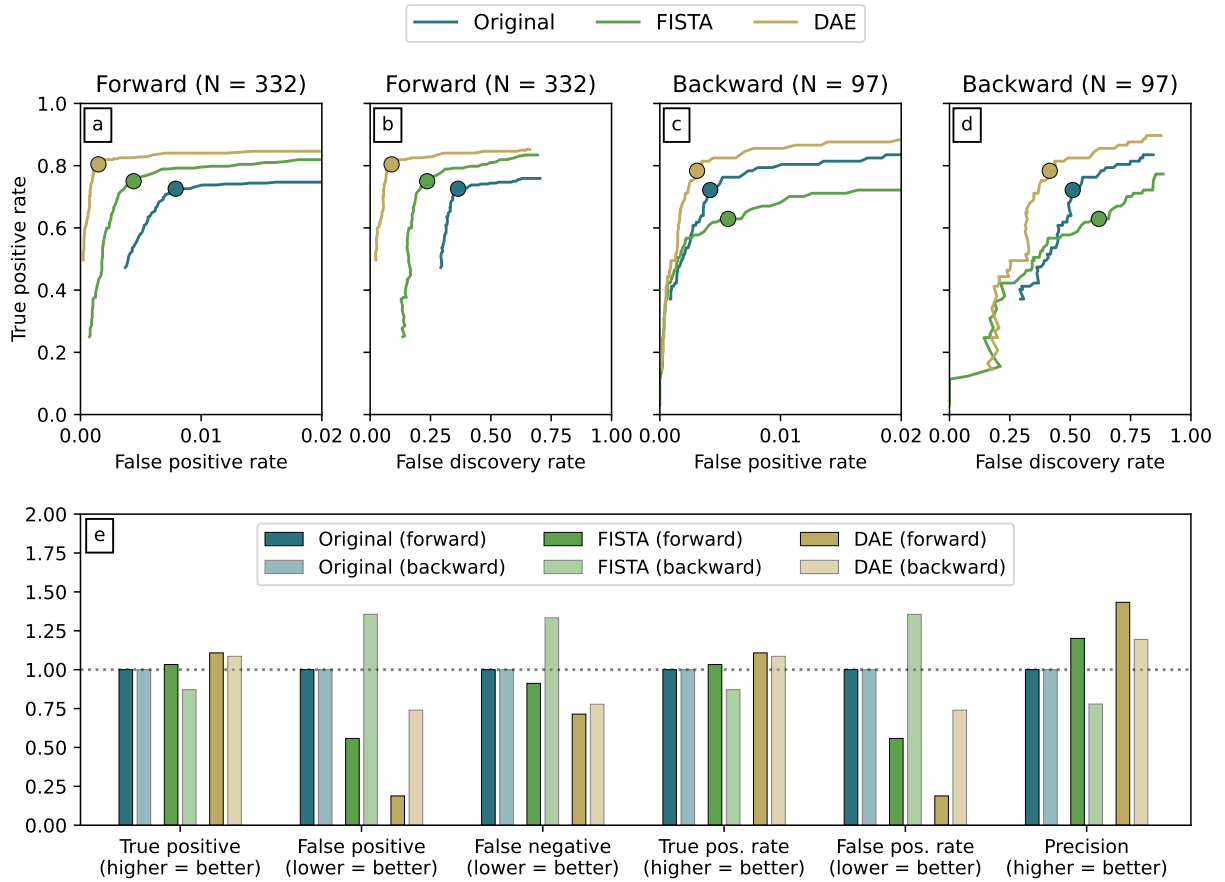


Fig. 10. Overview of the detection performance of the proposed methods. a,b) Trade-offs between the true positive rate (or recall), the false positive rate, and false discovery rate for vehicles travelling away from the interrogator (defined as the forward direction); c,d) The same as a,b, but for the vehicles travelling towards the interrogator (defined as the backward direction); The coloured dots in panels a-d indicate the optimal peak persistence threshold selected for this study. The total number of vehicles travelling in each direction are indicated in each panel title; e) Relative performance metrics measured at the optimal peak persistence threshold, normalised by the detection performance on the original (non-deconvolved) dataset.

each automatic detection from the beampower peak detector, we find the closest corresponding entry in the manual dataset and accept it as a correct detection (true positive; TP) if the time-gap between the two is less than 1 s. If the time-gap between the automatic detections and the manual picks is more than 1 s, or if multiple automatic detections are assigned to one manual pick, the detections are labelled as false positives (FP). Lastly, the number of false negatives (FN) is simply the number of manual picks that do not have a corresponding automatic detection, and the number of true negatives (TN) is the total number of analysed time windows minus the number of false positives. In these calculations, vehicle detections with an inferred speed lower than  $50 \text{ km hr}^{-1}$  or higher than  $110 \text{ km hr}^{-1}$  are discarded as outliers. The true positive rate (TPR; also known as recall), false positive rate (FPR), and false discovery rate (FDR) are then calculated as:

$$\text{TPR} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (9a)$$

$$\text{FPR} = \frac{\text{FP}}{\text{FP} + \text{TN}} \quad (9b)$$

$$\text{FDR} = \frac{\text{FP}}{\text{FP} + \text{TP}} \quad (9c)$$

The best trade-off between these three quantities is found by varying the peak prominence threshold independently for each dataset and traffic direction, with optimal values indicated by the coloured dots in Fig 10a-d. A summary of the performance characteristics is given in Table I and Fig. 10e.

Considering these performance metrics, we find that beamforming on the data deconvolved by the DAE systematically outperforms beamforming on the non-deconvolved data, or on the data deconvolved with FISTA; in all cases, the number of false positives and missed detections is lower, and consequently the number of correct detections is higher. The difference between the original and deconvolved data is not strongly apparent from the true positive (rate) because of the calibration of the prominence threshold; rather, the differences are most apparent in the false positive and false negative statistics, which clearly favour the DAE as the more reliable method for detecting vehicles in the data.

### B. Traffic flow statistics

As a final exploration of the data, we apply our vehicle detection method to the full validation dataset, which spans 24 hours of data, deconvolved with the DAE. To process

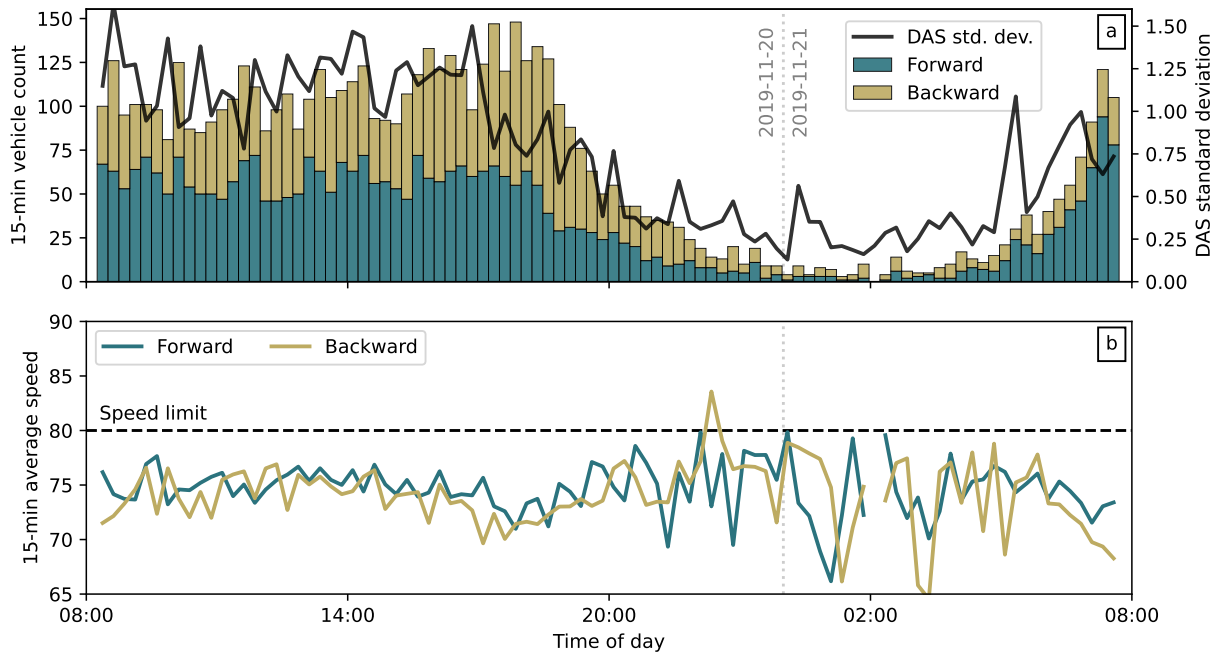


Fig. 11. Traffic statistics for 24 hours of DAS data. a) Vehicle counts in each direction for 15-minute bin intervals, alongside the 15-minute binned (normalised) standard deviation of the DAS data in the 0.1-2 Hz frequency band; b) 15-minute binned average of the speed of the detected vehicles. The road’s speed limit is as indicated.

1 day of DAS data, the DAE took just under 30 seconds. In order to deconvolve the same amount of data with FISTA, the algorithm would have taken well over 3 hours (based on the estimation cited in Section IV-A). Subsequently, the beamforming on the deconvolved data took over 18 minutes, and detecting the peaks in the beampower roughly 2 seconds. The bottleneck in the analysis is therefore the beamforming, but this was done with sliding window strides of merely 0.2 s for the purpose of high-resolution visualisation. For deployment purposes, one could consider increasing the window stride, which proportionally reduces the beamforming computation time, or to compute these in parallel instead of sequentially.

Looking at the traffic statistics shown in Fig. 11a, daytime and nighttime are directly recognisable by the traffic volume, being rather dense throughout the day with up to 150 vehicles per 15 minutes, and becoming lighter (less than 20 vehicles per 15 minutes) at night. Unsurprisingly, the standard deviation of the DAS data in a 0.1-2 Hz frequency band correlates well with these statistics. While the DAS signal variance does not reveal much beyond first-order traffic density information, variations in seismometer noise have previously been analysed in relation to the effect of COVID-19 lockdowns [33]. The DAS “noise” in itself can therefore be of interest as a heuristic in some applications, such as for investigating the effects of traffic managements strategies on traffic densities at a regional scale.

Interestingly, while for most of the day the volume of traffic travelling away from the interrogator (forward direction) is slightly less than that for traffic travelling towards the interrogator (backward direction), this balance is entirely reversed during the early-morning period (between 06:00 and 08:00 on

Thursday 21 November 2019), when the vast majority of vehicles are detected travelling in the forward direction. Manual inspection reveals that this is not a detection artefact, and that indeed many more vehicles are seen travelling in the forward direction. During the experiment, the DAS interrogator was located in the village of Alba-la-Romaine, at the northern end of the cable (see Fig. 3). The forward direction, away from the interrogator, is therefore in the direction of south. The nearest major city in the region, Montélimar, is due east of Alba-la-Romaine, and can be reached either through a minor local road, or through the major road that is the subject of this study. Normally the total travel time and distance is slightly shorter when taking the minor road. However, several days prior to the start of the DAS experiment, on 11 November 2019, an  $M_w$  4.9 earthquake occurred close to the village of Le Teil which lies on the route to Montélimar (in fact, the DAS experiment was part of a rapid response campaign to monitor aftershocks; [34]). Due to extensive structural damage to buildings in Le Teil, and to facilitate repair works, the road to Montélimar was partially blocked. It is therefore likely that many daily commuters and delivery personnel chose to take a slightly longer route to Montélimar, to circumvent traffic congestion in the village of Le Teil. As can be inferred from the average speed recorded in Fig. 11b, the additional traffic redirected to this road does not inhibit the flow of traffic (i.e., cause traffic jams), since the average speed remains fairly steady near the speed limit. Only during the evening rush hour (around 18:00) is there a small but noticeable drop in the average speed.

## VI. PERSPECTIVES FOR SMART CITIES

Having demonstrated the feasibility of using DAS to analyse dense traffic scenarios, we now turn to a brief outlook on future developments, particularly for the application in “smart cities” (*sensu lato*). It is likely that fibre-optic sensing will contribute significantly to future developments of smart cities in relation to solid media, including subsurface characterisation [35], infrastructure monitoring of railways [36] and roads [14], and structural integrity evaluations [2]. At a regional scale, fibre-optic sensing has proven useful for recording earthquakes [37] and avalanches [38], monitoring glaciers [39], and inferring hydrological conditions of the subsurface [7]. As extreme weather events like violent storms and droughts will become more likely in the near future [40], dense instrumentation and continuous observations of the subsurface will be useful (or even critical) for agriculture, water, land and forest management, preventive wildfire measures, and slope stability surveillance. Since dense fibre-optic infrastructures are already in place in urban [14] and even rural [19] and submarine [41] environments, the main technological breakthroughs needed to advance the field are those that facilitate efficient and accurate processing of the vast volumes of data generated by fibre-optic sensing.

The Deep Learning method proposed in the present study simultaneously addresses both the efficiency and accuracy of data processing workflows. As demonstrated in previous sections, deconvolution of the DAS data improves the reliability of vehicle detection methods in heavy traffic scenarios. While the use of a conventional deconvolution method like FISTA is hindered by the computation time, estimated to take over 3 hours to deconvolve a 24-hour DAS dataset, the proposed Deconvolution Auto-Encoder requires less than 30 seconds to deconvolve the same dataset while yielding higher deconvolution fidelity. The performance of the proposed method warrants its application to quasi-real time problems, which could include regional traffic flow management, incident or anomaly detection, and policy testing in (rural) localities where traffic cameras are scarce or absent.

Even in urban environments with dense traffic camera coverage, DAS-based traffic monitoring is still a valuable asset, as the deformation patterns induced by vehicles themselves encode the state of the road through effective elastic (and viscous; [42]) properties of the tarmac. Long-term variations in the deformation patterns might reveal road degradation, in a similar manner as railway track and train wheel condition is evaluated with DAS [36], [43], [44]. The feasibility of DAS for this application still needs to be tested, but undoubtedly more experiments will be conducted in the near future.

## VII. CONCLUSIONS

In this study, we consider the vehicle detection and characterisation performance of roadside Distributed Acoustic Sensing (DAS). We propose a procedure that focuses on the quasi-static deformation induced by individual vehicles, recorded by DAS as a spatio-temporal pattern of strain rate. By subsequently applying MUSIC beamforming to these DAS data, individual vehicles can be detected and characterised

in terms of their velocity, but doing so under dense traffic conditions is challenging. To improve the temporal resolution of this approach, we propose to deconvolve the characteristic impulse strain response induced by the weight of passing cars from the DAS data, which leads to a higher temporal resolution of the detections, as well as fewer false and missed detections. For the purpose of multi-channel deconvolution, we develop a self-supervised Deep Learning algorithm (a “Deconvolution Auto-Encoder”) that outperforms a commonly-used conventional deconvolution algorithm, the fast-Iterative Shrinkage Thresholding Algorithm (FISTA). Not only does the DAE lead to better vehicle detection performance, it is also  $> 400\times$  faster than FISTA, which is of practical importance for the implementation of our method in real-world settings.

## ACKNOWLEDGEMENTS

MvdE thanks J.-P. Ampuero for inspiring discussions on the deconvolution approach. The DAS interrogator was kindly provided by Febus. This research was supported by the French government through the 3IA Côte d’Azur Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-19-P3IA-0002. The Deep Learning was performed within the TensorFlow framework [45] and generic data manipulations were performed with NumPy [46] and SciPy [32]. Data visualisation was done using Matplotlib [47]. Data, code, and pre-trained models are available from [doi.org/10.6084/m9.figshare.16653163](https://doi.org/10.6084/m9.figshare.16653163).

## REFERENCES

- [1] A. H. Hartog, *An Introduction to Distributed Optical Fibre Sensors*. CRC Press, May 2017.
- [2] P. G. Hubbard, J. Xu, S. Zhang, M. Dejong, L. Luo, K. Soga, C. Papa, C. Zulberti, D. Malara, F. Fugazzotto, F. Garcia Lopez, and C. Minto, “Dynamic structural health monitoring of a model wind turbine tower using distributed acoustic sensing (DAS),” *Journal of Civil Structural Health Monitoring*, vol. 11, no. 3, pp. 833–849, Jul. 2021.
- [3] J. Mestayer, B. Cox, P. Wills, D. Kiyashchenko, J. Lopez, M. Costello, S. Bourne, G. Ugueto, R. Lupton, G. Solano, D. Hill, and A. Lewis, “Field trials of distributed acoustic sensing for geophysical monitoring,” in *SEG Technical Program Expanded Abstracts 2011*, ser. SEG Technical Program Expanded Abstracts. Society of Exploration Geophysicists, Jan. 2011, pp. 4253–4257.
- [4] 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.
- [5] A. Lellouch, S. Yuan, Z. Spica, B. Biondi, and W. L. Ellsworth, “Seismic Velocity Estimation Using Passive Downhole Distributed Acoustic Sensing Records: Examples From the San Andreas Fault Observatory at Depth,” *Journal of Geophysical Research: Solid Earth*, vol. 124, no. 7, pp. 6931–6948, 2019.
- [6] J. Ajo-Franklin, S. Dou, T. Daley, B. Freifeld, M. Robertson, C. Ulrich, T. Wood, I. Eckblaw, N. Lindsey, E. Martin, and A. Wagner, “Time-lapse surface wave monitoring of permafrost thaw using distributed acoustic sensing and a permanent automated seismic source,” in *2017 SEG International Exposition and Annual Meeting*. Society of Exploration Geophysicists, Oct. 2017.
- [7] V. R. Tribaldos and J. B. Ajo-Franklin, “Aquifer Monitoring Using Ambient Seismic Noise Recorded With Distributed Acoustic Sensing (DAS) Deployed on Dark Fiber,” *Journal of Geophysical Research: Solid Earth*, vol. 126, no. 4, p. e2020JB021004, 2021.
- [8] N. Vlachopoulos, D. Cruz, and B. Forbes, “Utilizing a novel fiber optic technology to capture the axial responses of fully grouted rock bolts,” *Journal of Rock Mechanics and Geotechnical Engineering*, vol. 10, no. 2, pp. 222–235, Apr. 2018.

- [9] L. Cheng, A. Zhang, B. Cao, J. Yang, L. Hu, and Y. Li, "An experimental study on monitoring the phreatic line of an embankment dam based on temperature detection by OFDR," *Optical Fiber Technology*, vol. 63, p. 102510, May 2021.
- [10] Z. Zhan, "Distributed Acoustic Sensing Turns Fiber-Optic Cables into Sensitive Seismic Antennas," *Seismological Research Letters*, vol. 91, no. 1, pp. 1–15, Jan. 2020.
- [11] N. J. Lindsey and E. R. Martin, "Fiber-Optic Seismology," *Annual Review of Earth and Planetary Sciences*, vol. 49, no. 1, pp. 309–336, May 2021.
- [12] E. Martin, N. Lindsey, S. Dou, J. Ajo-Franklin, T. Daley, B. Freifeld, M. Robertson, C. Ulrich, A. Wagner, and K. Bjella, "Interferometry of a roadside DAS array in Fairbanks, AK," in *SEG Technical Program Expanded Abstracts 2016*, ser. SEG Technical Program Expanded Abstracts. Society of Exploration Geophysicists, Sep. 2016, pp. 2725–2729.
- [13] 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.
- [14] S. Yuan, A. Lellouch, R. G. Clapp, and B. Biondi, "Near-surface characterization using a roadside distributed acoustic sensing array," *The Leading Edge*, vol. 39, no. 9, pp. 646–653, Sep. 2020.
- [15] N. J. Lindsey, S. Yuan, A. Lellouch, L. Gualtieri, T. Lecocq, and B. Biondi, "City-Scale Dark Fiber DAS Measurements of Infrastructure Use During the COVID-19 Pandemic," *Geophysical Research Letters*, vol. 47, no. 16, p. e2020GL089931, 2020.
- [16] K. Chambers, "Using DAS to investigate traffic patterns at Brady Hot Springs, Nevada, USA," *The Leading Edge*, vol. 39, no. 11, pp. 819–827, Nov. 2020.
- [17] X. Wang, E. F. Williams, M. Karrenbach, M. G. Herráez, H. F. Martins, and Z. Zhan, "Rose Parade Seismology: Signatures of Floats and Bands on Optical Fiber," *Seismological Research Letters*, vol. 91, no. 4, pp. 2395–2398, Jul. 2020.
- [18] H. Liu, J. Ma, W. Yan, W. Liu, X. Zhang, and C. Li, "Traffic Flow Detection Using Distributed Fiber Optic Acoustic Sensing," *IEEE Access*, vol. 6, pp. 68 968–68 980, 2018.
- [19] P. Jousset, T. Reinsch, T. Ryberg, H. Blanck, A. Clarke, R. Aghayev, G. P. Hersir, J. Henniges, M. Weber, and C. M. Krawczyk, "Dynamic strain determination using fibre-optic cables allows imaging of seismological and structural features," *Nature Communications*, vol. 9, no. 1, p. 2509, Jul. 2018.
- [20] A. Chambolle, R. De Vore, N.-Y. Lee, and B. Lucier, "Nonlinear wavelet image processing: Variational problems, compression, and noise removal through wavelet shrinkage," *IEEE Transactions on Image Processing*, vol. 7, no. 3, pp. 319–335, Mar. 1998.
- [21] I. Daubechies, M. Defrise, and C. D. Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Communications on Pure and Applied Mathematics*, vol. 57, no. 11, pp. 1413–1457, 2004.
- [22] A. Beck and M. Teboulle, "A fast Iterative Shrinkage-Thresholding Algorithm with application to wavelet-based image deblurring," in *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, Apr. 2009, pp. 693–696.
- [23] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, "JAX: Composable transformations of Python+NumPy programs," 2020.
- [24] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," *arXiv:1505.04597 [cs]*, May 2015.
- [25] P. Ramachandran, B. Zoph, and Q. V. Le, "Searching for Activation Functions," *arXiv:1710.05941 [cs]*, Oct. 2017.
- [26] H. L. Van Trees, *Optimum Array Processing*. New York, USA: John Wiley & Sons, Inc., Mar. 2002.
- [27] P. Goldstein and R. J. Archuleta, "Array analysis of seismic signals," *Geophysical Research Letters*, vol. 14, no. 1, pp. 13–16, 1987.
- [28] E. Kiser and M. Ishii, "Back-Projection Imaging of Earthquakes," *Annual Review of Earth and Planetary Sciences*, vol. 45, no. 1, pp. 271–299, 2017.
- [29] M. P. A. van den Ende and J.-P. Ampuero, "Evaluating seismic beam-forming capabilities of distributed acoustic sensing arrays," *Solid Earth*, vol. 12, no. 4, pp. 915–934, Apr. 2021.
- [30] R. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Transactions on Antennas and Propagation*, vol. 34, no. 3, pp. 276–280, Mar. 1986.
- [31] D. Thomson, "Spectrum estimation and harmonic analysis," *Proceedings of the IEEE*, vol. 70, no. 9, pp. 1055–1096, Sep. 1982.
- [32] 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.
- [33] T. Lecocq, S. P. Hicks, K. V. Noten, K. van Wijk, P. Koelemeijer, R. S. M. D. Plaen, F. Massin, G. Hillers, R. E. Anthony, M.-T. Apoloner, M. Arroyo-Solórzano, J. D. Assink, P. Büyükkapınar, A. Cannata, F. Cannavo, S. Carrasco, C. Caudron, E. J. Chaves, D. G. Cornwell, D. Craig, O. F. C. den Ouden, J. Diaz, S. Donner, C. P. Evangelidis, L. Evers, B. Fauville, G. A. Fernandez, D. Giannopoulos, S. J. Gibbons, T. Girona, B. Grecu, M. Grunberg, G. Hetényi, A. Horleston, A. Inza, J. C. E. Irving, M. Jamalrehyani, A. Kafka, M. R. Koymans, C. R. Labeledz, E. Larose, N. J. Lindsey, M. McKinnon, T. Megies, M. S. Miller, W. Minarik, L. Moresi, V. H. Márquez-Ramírez, M. Möllhoff, I. M. Nesbitt, S. Niyogi, J. Ojeda, A. Oth, S. Proud, J. Pulli, L. Retailleau, A. E. Rintamäki, C. Satriano, M. K. Savage, S. Shani-Kadmiel, R. Sleeman, E. Sokos, K. Stammler, A. E. Stott, S. Subedi, M. B. Sørensen, T. Taira, M. Tapia, F. Turhan, B. van der Pluijm, M. Vanstone, J. Vergne, T. A. T. Vuorinen, T. Warren, J. Wassermann, and H. Xiao, "Global quieting of high-frequency seismic noise due to COVID-19 pandemic lockdown measures," *Science*, vol. 369, no. 6509, pp. 1338–1343, Sep. 2020.
- [34] C. Cornou, J.-P. Ampuero, C. Aubert, L. Audin, S. Baize, J. Billant, F. Brenguier, M. Causse, M. Chlieh, A. Combey, M. de Michele, B. Delouis, A. Deschamps, M. Ferry, M. Fomelis, B. Froment, C. Gélis, R. Grandin, J.-R. Grasso, E. Hannouz, S. Hok, A. Jung, R. Jolivet, M. Langlais, P. Langlaude, C. Larroque, P. H. Leloup, K. Manchuel, L. Marconato, C. Maron, E. Mathot, E. Maufroy, D. Mercierat, M. Metois, E. Nayman, I. Pondaven, L. Provost, J. Régnier, J.-F. Ritz, D. Rivet, A. Schlupp, A. Sladen, C. Voisin, A. Walpersdorf, D. Wolyniec, P. Allemand, E. Beck, E. Bertrand, V. Bertrand, P. Briole, D. Brunel, O. Cavalié, J. Chèze, F. Courboulex, I. Douste-Bacquet, R. Dretzen, T. Giampietro, M. Godano, P. Grandjean, M. Grunberg, G. Guerin, S. Guillot, E. E. Haber, A. Hernandez, H. Jomard, C. Lasserre, C. Liang, I. Lior, X. Martin, D. Mata, M. Menager, A. Mercier, A. Mordret, E. Oral, A. Paul, F. Peix, C. Pequegnat, M. Pernoud, C. Satriano, R. Sassi, M. Schaming, V. Sellier, C. Sira, A. Socquet, C. Sue, A. Trilla, M. Vallée, M. van den Ende, P. Vernant, B. Vial, and H. Weng, "Rapid response to the M w 4.9 earthquake of November 11, 2019 in Le Teil, Lower Rhône Valley, France," *Comptes Rendus. Géoscience*, vol. 353, no. S1, pp. 1–23, Mar. 2021.
- [35] 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.
- [36] C. Wiesmeyer, A. Papp, H. Garn, G. Neunteufel, and W. Zottl, "Train tracking and train condition monitoring by Distributed Acoustic Sensing," in *Proceedings of 7th Transport Research Arena TRA 2018*. Vienna: Zenodo, Apr. 2018.
- [37] T. Zhu, J. Shen, and E. R. Martin, "Sensing Earth and environment dynamics by telecommunication fiber-optic sensors: An urban experiment in Pennsylvania, USA," *Solid Earth*, vol. 12, no. 1, pp. 219–235, Jan. 2021.
- [38] A. Fichtner, P. Edme, P. Paitz, N. Lindner, M. Hohl, P. Huguenin, B. Sovilla, P. Roig-Lafon, E. Surinach, and F. Walter, "Observing avalanche dynamics with Distributed Acoustic Sensing," *European Geosciences Union*, Other, Mar. 2021.
- [39] F. Walter, D. Gräff, F. Lindner, P. Paitz, M. Köpfl, M. Chmiel, and A. Fichtner, "Distributed acoustic sensing of microseismic sources and wave propagation in glaciated terrain," *Nature Communications*, vol. 11, no. 1, p. 2436, May 2020.
- [40] V. Masson-Delmotte, P. Zhai, S. Pirani, L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M. I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J. B. R. Matthews, T. K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou, "Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change (IPCC)," Tech. Rep., 2021.
- [41] 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.
- [42] A. Ahmed and S. Erlingsson, “Viscoelastic Response Modelling of a Pavement under Moving Load,” *Transportation Research Procedia*, vol. 14, pp. 748–757, Jan. 2016.
- [43] D. Milne, A. Masoudi, E. Ferro, G. Watson, and L. Le Pen, “An analysis of railway track behaviour based on distributed optical fibre acoustic sensing,” *Mechanical Systems and Signal Processing*, vol. 142, p. 106769, Aug. 2020.
- [44] I. Vidovic and S. Marschnig, “Optical Fibres for Condition Monitoring of Railway Infrastructure—Encouraging Data Source or Errant Effort?” *Applied Sciences*, vol. 10, no. 17, p. 6016, Jan. 2020.
- [45] 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.
- [46] 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.
- [47] J. D. Hunter, “Matplotlib: A 2D Graphics Environment,” *Computing in Science & Engineering*, vol. 9, no. 3, pp. 90–95, 2007.