Deep Deconvolution Applied to Distributed Acoustic Sensing for Traffic Analysis

Martijn van den Ende* André Ferrari*, Anthony Sladen†, and Cédric Richard*
* Université Côte d'Azur, OCA, UMR Lagrange, France
† Université Côte d'Azur, IRD, CNRS, Observatoire de la Côte d'Azur, Géoazur, France
Email: martijn.vandenende@oca.eu

***Abstract*—Distributed Acoustic Sensing (DAS) is a nascent technology that facilitates the measurement of vibrations along fibre-optic telecommunication cables, which has numerous novel applications in many domains of science and engineering. In the present study, we use DAS to analyse traffic along a fibre-optic cable deployed along a major road in Nice, France. We recently proposed to use a MUSIC beamforming algorithm for the objective of estimating the speed of individual vehicles. To greatly improve the accuracy and precision of the beamforming results, we here propose a Deep Deconvolution algorithm applied to the data prior to beamforming. The accuracy of the speed estimation is in the range of 0.14-$0.25\ \mathrm{km\ h^{-1}}$, which is at least one order of magnitude better than conventional methods. DAS therefore has great potential in urban traffic analysis applications.**

I. Introduction

Distributed Acoustic Sensing (DAS [1]) is a laser-pulsing technology that converts fibre-optic telecommunication cables into arrays of thousands of vibration sensors, positioned every few metres along the fibre-optic cable. An interrogator unit systematically sends short pulses of laser light into one end of an optical glass fibre, and measures the phase and/or amplitude of the back-scattered photons that have interacted with nanometric-scale defects along its path through the fibre. Through interferometric techniques, the (rate of) stretching of the fibre can be inferred from the back-scattered measurement at equidistant points along the fibre, at temporal sampling rates up to several kHz. This has enabled numerous applications in science and engineering, including (but not limited to) geophysics and seismology [2]–[4], structural integrity monitoring [5], and hydrology [6].

One particular application of interest of the present study, is that of vehicular traffic monitoring. As many commercial telecommunication cables are being deployed immediately adjacent to public roads, DAS has the ability to record the passage of vehicles over these roads [7], [8]. When a vehicle drives past a given location along the cable, its weight pressing down on the road causes small deflections in the subsurface, which is transferred to the fibre and subsequently measured by the interrogator (see Fig. 1). By estimating the timing at which a vehicle passes by a given sensing point, and knowing the (fixed) spacing between sensing points, one can precisely obtain the location and velocity of the vehicle. Roadside DAS therefore holds enormous potential for high-resolution traffic monitoring, and can be complementary to conventional instrumentation like traffic cameras and inductive loops embedded in the road.

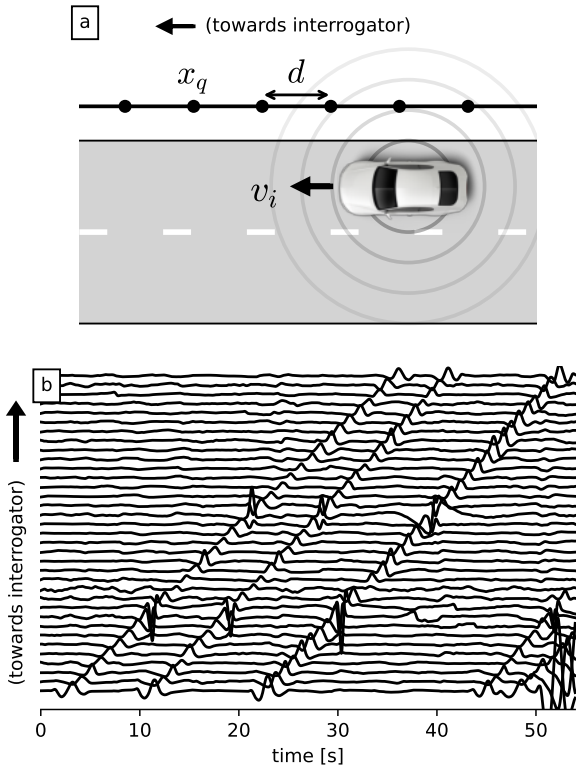

Fig. 1. a) Schematic illustration of the layout. A vehicle with velocity v_i travels along the DAS cable that was deployed parallel to the road. The spacing between the sensors is d, with the q-th sensor being located at x_q; b) Example of DAS data containing three cars. Each black line represents a time-series measurement at a given sensor. Each car traces out a diagonal line in the data, the slope of which equals its velocity.

To enable "smart city" applications facilitated by DAS, such as advanced traffic control, robust and accurate DAS analysis algorithms are needed. We recently showed in [9] that state-of-the-art speed estimation accuracy can be achieved using a MUSIC beamforming algorithm. To massively improve upon the accuracy and precision of the speed estimation, we leverage a self-supervised, non-blind Deep Learning deconvolution model that deconvolves the characteristic signal of cars and other vehicles from the DAS data.

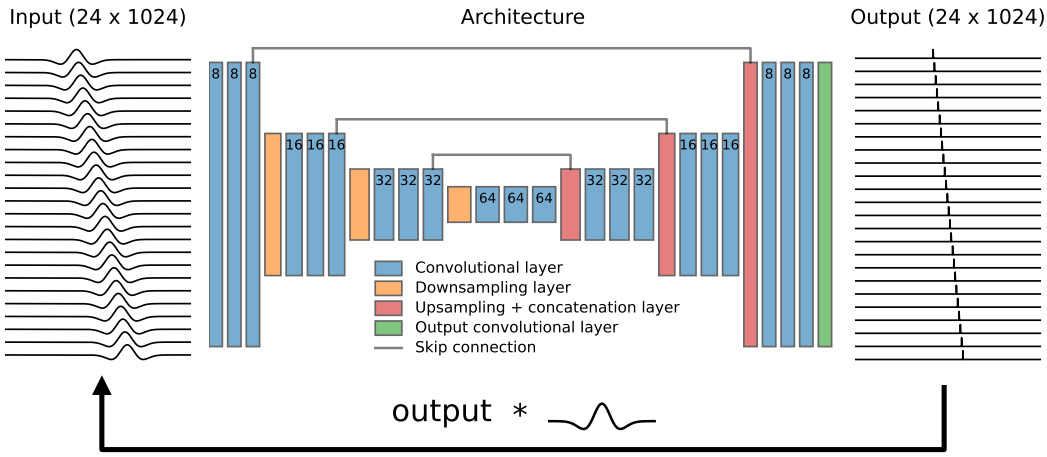

Fig. 2. Architecture of the Deconvolution Auto Encoder. We take a set of time-series of size 1024 samples, recorded by 24 consecutive DAS sensor. This input is passed to a U-Net Auto-Encoder and produces an output $\hat{M}$ of the same size as the input. Intermediate down- and upsampling is performed with a factor 2 along the spatial axis, and with a factor 4 along the temporal axis. The output is subsequently convolved with the known impulse response of a car, and the learning objective is to minimise the difference between the convolved output and the original input (subject to a sparsity constraint on M).

II. Approach

A. Signal model

The signal of a vehicle travelling along the fibre-optic cable, as recorded by DAS, is the result of deformation in the subsurface induced by the weight of the vehicle. This deformation can be precisely modelled based on the Flamant-Boussinesq approximation for a point load [2], [8]:

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

$$\varepsilon = \frac{1}{d}\left[u_x\left(x + \frac{d}{2}, y, z\right) - u_x\left(x - \frac{d}{2}, y, z\right)\right] \tag{1b}$$

In this model, a DAS channel is positioned at a coordinate point (x, y, z) (x: road-tangential; y: road-perpendicular; z: depth beneath the surface; all in units of distance) with distance $r = \sqrt{x^2 + y^2 + z^2}$ from the point load (i.e. the vehicle) centred at the origin. The weight of the car F (units: N) induces deformation u_x in a uniform elastic medium with shear modulus G (units: Pa) and Poisson's ratio ν (dimensionless), which is measured by DAS as a strain ε (dimensionless) averaged over the gauge length d. Owing to the point-load assumption, (1b) can be interpreted as the impulse response of the DAS system in response a vehicle passing by a given DAS sensor. For a given (known) position of the fibre with respect to the road (y, z), along with estimates of the elastic properties of the medium (G, ν), the only variable in (1b) is the position x of the vehicle of constant weight F. In a fixed reference frame centred at a given DAS channel q, the x-coordinate can be replaced by the vehicle's velocity v as $x = x_q + v\,(t - t_0)$ for some initial time t_0 and DAS sensor location x_q, which converts (1b) into a signal model of the DAS time series. This is illustrated in Fig. 1b, in which each of the three cars traces out a diagonal line of which the slope is equal to v. This time shift caused by translation of the vehicle makes the DAS recordings amenable to beamforming analysis [9].

B. Deconvolution Auto-Encoder

As is apparent from Fig. 1b and from the theoretical considerations laid out in the previous subsection, the characteristic signature of a car as recorded by DAS is the same for each car (up to a proportionality constant). We can therefore achieve higher resolution traffic measurements by deconvolving this characteristic impulse response from the DAS data. A first solution would be to rely on a classical deconvolution model based on the squared loss data fidelity term regularised by a sparsity-promoting ℓ_1 term, i.e.:

$$\hat{m}_q = \arg\min_{m_q} \left\{||y_q - h * m_q||_2^2 + \lambda||m_q||_1\right\} \tag{2}$$

in which y_q is a vector containing the DAS time series at channel q, $h * m_q$ represents the convolution between a characteristic impulse response h (see II-A) with an underlying impulse model m_q at DAS channel q, and λ controls the regularisation strength. However, this formulation, which can be efficiently solved by a (F)ISTA algorithm [10], [11], does not take into account the correlated nature of the DAS measurements; when a car travels along the cable, the associated DAS measurements $\{y_1, y_2, \dots\}$ are not statistically independent. As shown in a previous study [12], taking into account this statistical dependence can greatly improve the quality of the deconvolution results. To this end we employ a Deep Learning model N_θ (Fig. 2) parameterised by θ, which takes as an input a subset of the DAS data collected as a matrix Y with lines y_q and produces some output (M) of the same size as the input. We then convolve each line of this model output by the (known) impulse response of a car, i.e.

$\hat{Y} = h*M = h*N_\theta(Y)$. We then define the learning objective as:

$$\hat{\theta} = \arg\min_\theta \mathbb{E}\left\{||Y - h * N_\theta(Y)||_2^2 + \lambda||N_\theta(Y)||_1\right\} \quad (3)$$

in which the expectation is taken over the distribution of Y, which is approximated by K samples in a training set. Following the completion of the training phase, the output of the model $\hat{M} = N_{\hat{\theta}}(Y)$ is such that upon convolution with the car's impulse response, the original input Y is approximated; the Deep Learning algorithm therefore represents a deconvolution operation. Since all variables in (3) are only a function of the input Y, it is clear that this learning approach is entirely self-supervised (no "ground truth" for M needs to be provided). Similar approaches were adopted in recent self-supervised image denoising studies [13], [14], in which deconvolution was a by-product of the denoising procedure. In the present work, we focus on the deconvolution of multivariate time series to assess the effect of subsequent speed analysis.

III. Experimental setup

The data analysed for this study were acquired during a measurement campaign performed in the city of Nice, in the south of France. The fibre was nested within a bundle of cables firmly attached to the side of a multi-lane suspended road crossing the city, and was sensed with an hDAS interrogator (Aragon Photonics) with a channel spacing equal to the gauge length of $d = 10$ m, at a temporal sampling frequency of 250 Hz. To preprocess the data, we applied a bandpass filter with a 0.1-2 Hz pass band, and downsampled the data to 25 Hz. The Deconvolution Auto-Encoder was trained on 30 minutes' worth of data containing around 1000 cars (split 80-20 in a training and validation set) over a total of 5000 epochs. For each batch in the training set, input samples were randomly generated from the training set, choosing a random time and DAS channel defining the starting point of each input sample window. The regularisation strength λ was set to 10 (both for training and testing), which we qualitatively assessed to provide a good trade-off between sparsity and fidelity (higher values leading to increased sparsity). The loss function was minimised by the Adam optimisation algorithm. We assess the performance of the proposed workflow based on four typical examples of DAS recordings of cars travelling towards the interrogator (see top row of panels in Fig. 3). These example recordings were deconvolved with the Deconvolution Auto-Encoder (bottom row in Fig. 3). Both the original and the deconvolved data were subsequently analysed with MUSIC beamforming.

The beamforming analysis is performed as described in [9] (using Model #1 therein), assuming a reference velocity $v_{ref} = 70$ km h^{-1}, being the speed limit of the road under study. For visualisation purposes, all of the waveforms in Fig. 3 are shifted in accordance with this reference velocity. Under this transformation, a vehicle with a speed of exactly v_{ref} would trace out a vertical line, with faster vehicles tracing out a diagonal line from top left to bottom right (and vice versa for slower vehicles). We apply the MUSIC algorithm to a sliding window of 15 consecutive DAS sensor recordings, equivalent to 140 m distance. This sliding window traverses the data along the spatial axis, estimating the distribution of beampower as a function of distance along the cable.

For reference, we also estimate the velocity of each selected vehicle by estimating the timing of passage at a given DAS sensor as the timing of the peak strain for that sensor (i.e. by taking the peaks seen in e.g. Fig. 1b). The average velocity between consecutive DAS channels is consequently $v = d/(t_{i+1} - t_i)$. Owing to the presence of noise in the data (with spatially non-uniform signal-to-noise ratio), this method is sensitive to outliers, and manual fine-tuning with a median filter and a Savitzky-Golay filter was needed to obtain stable estimates of the vehicle's velocity $v = d\Delta q/\Delta t_{peak}$ over a distance interval $d\Delta q$. We suppose that the extensive fine-tuning renders this approach based on peaks in the DAS recordings infeasible for automated traffic analysis.

IV. Beamforming results

Corresponding with the four selected vehicles shown in Fig. 3, we plot the distribution of beampower for each vehicle as a function of distance along the cable in Fig. 4. By comparing the top row (original data) with the bottom row (deconvolved data), it becomes immediately apparent that the distribution of beampower is much more narrowly distributed for the deconvolved data than for the original data. Especially for vehicles 3 and 4 (Fig. 4c and d), the beampower distribution for the original data is very broad and multi-modal, which inhibits a precise estimation of the vehicles' velocities. By contrast, the beampower distribution for the deconvolved data of these vehicles (Fig. 4g and h) has only a single and sharp peak in beampower for each DAS channel, which could be easily detected and characterised with a basic automatic peak detector.

In comparison to the baseline estimation of the velocity (based on peaks in the strain, as described in the previous section), both data sets seem to be very accurate. We can quantify this accuracy by taking the location of the peak in beampower $\hat{v}_i$ for each channel i, and computing the absolute difference with the corresponding baseline estimation v_i^b at the same channel. The mean absolute difference averaged over all the channels, i.e. $\mathbb{E}_i\left[|\hat{v}_i - v_i^b|\right]$, is indicated in each panel in Fig. 4. Likewise, the precision of the method is estimated as the *full width at half maximum* W_i (i.e. the width of the peak at 50 % of its maximum) of the beampower peak for each channel i: precision $= \mathbb{E}_i[W_i]$. This measure is not entirely meaningful for a multi-modal distribution like seen in Fig. 4d, in which case the estimated precision is merely indicative.

Considering these quantitative performance metrics, we find that the beamforming results on the deconvolved data systematically outperform those for the original data. In the most extreme example, the accuracy and precision for the deconvolved data (0.17 and 1.32 km h^{-1}, respectively) are almost one order of magnitude better than for the original data

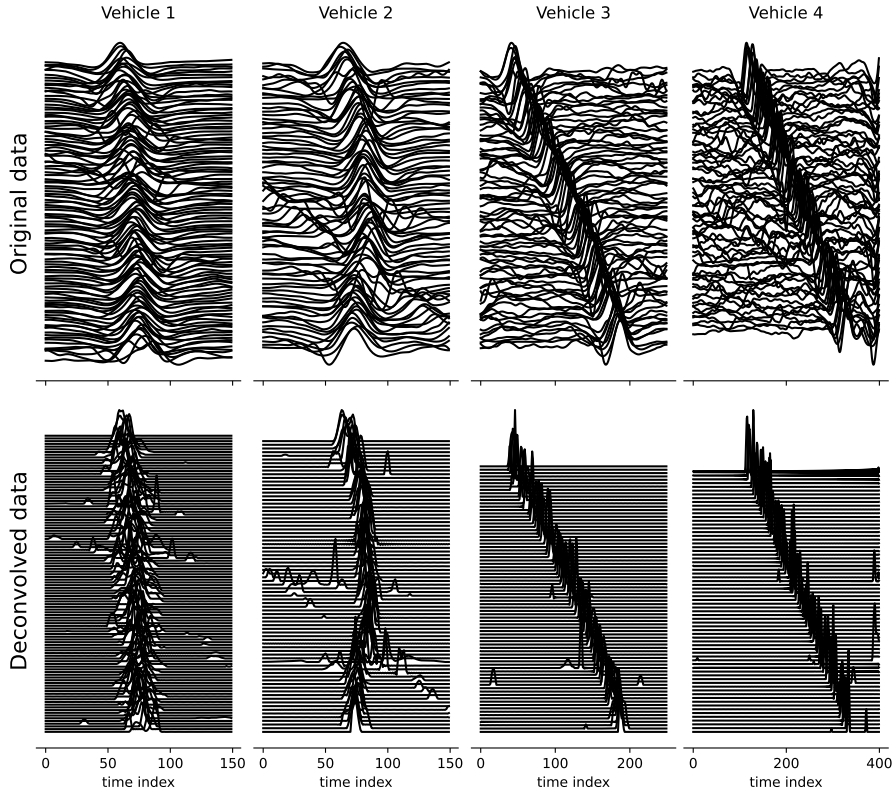

Fig. 3. Top row: four selected examples of recordings of cars, characterised by a large-amplitude wiggle near the centre of each panel. The waveforms are shifted according to a reference velocity of 70 $\mathrm{km\ h^{-1}}$ (see text); Bottom row: results of the Deconvolution Auto-Encoder algorithm, taking the waveforms from the top row as an input.

(3.55 and 7.14 $\mathrm{km\ h^{-1}}$, resp.). Moreover, the performance on the deconvolved data is much more consistent across samples, which is an important aspect to consider for real-world implementations. Previous work [12] has also shown that the proposed deconvolution method is hundreds of times faster than conventional, iterative deconvolution algorithms, which is particularly relevant for real-time DAS data analysis applications.

V. Conclusions & future perspectives

In this paper we demonstrate the efficiency of Deep Deconvolution preprocessing to improve the accuracy of MUSIC beamforming algorithms in estimating the velocity of isolated cars in Distributed Acoustic Sensing (DAS) data. Particularly when the characteristic signal of a vehicle is deconvolved from the data do the estimated velocities achieve extremely good accuracy and precision. With an accuracy in the range of 0.14-0.25 $\mathrm{km\ h^{-1}}$ and a precision in the range of 0.82-1.54 $\mathrm{km\ h^{-1}}$, DAS-based vehicle speed estimations are very competitive compared to established traffic analysis techniques. To give an example, the winning contender of the 2018 NVIDIA AI City Challenge [15] achieved an RMS accuracy of 6.6 $\mathrm{km\ h^{-1}}$ based on traffic camera data, using state-of-the-art computer vision techniques. Handheld radar guns could in principle achieve an accuracy of less than 1 $\mathrm{km\ h^{-1}}$, but have a tendency to produce outlier results [16]. Since the performance of DAS does not depend on environmental factors (weather conditions, lighting situation, etc.), we expect DAS to deliver accurate and consistent performance at all times.

Another major advantage of using DAS for traffic analysis, is that the fibre-optic infrastructure it relies on is already in place in many urban locations. While dedicated deployments could achieve better signal-to-noise ratios as a result of specific deployment protocols (e.g. improving coupling between the fibre and the road [17]), existing telecom cable deployments could be sufficient (as we demonstrate in this study). Upscaling DAS technologies can therefore be logistically more feasible than discrete sensor networks (inductive loops, cameras, roadside laser guns, etc.).

Lastly, we stress the notion that "smart city" applications facilitated by DAS are still in their infancy, and that advances in laser pulsing, fibre manufacturing, and cable deployment technologies may yield massive improvements in the signal quality that may bolster signal analysis algorithms in the near future.

Acknowledgements

This research was supported by the French government through the 3IA Côte d'Azur Investments in the Future project with the reference number ANR-19-P3IA-0002. Access to the fibre was provided by the Metropole Nice Côte d'Azur.

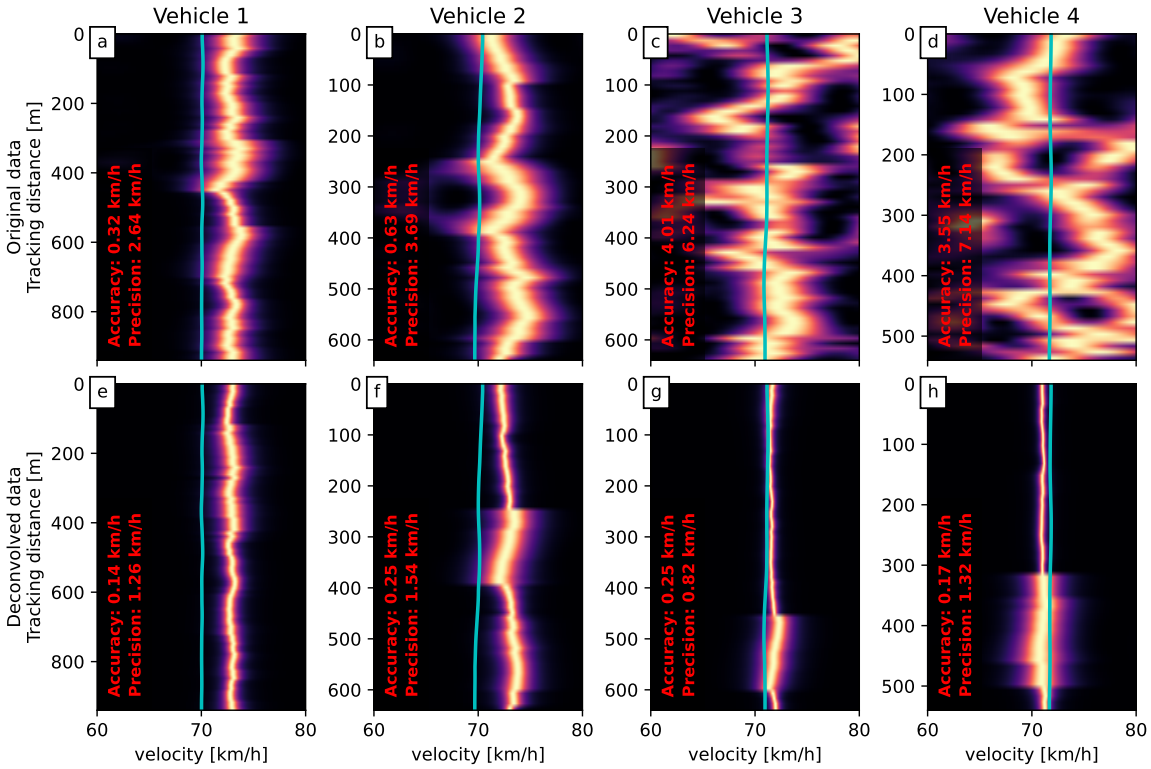

Fig. 4. Overview of beamforming performance. Top row: the beampower distributions for the four selected vehicles estimated from the original data (corresponding with the top row in Fig. 3). Bottom row: the beampower distributions estimated from the deconvolved data (corresponding with the bottom row in Fig. 3). In each panel, we indicate the velocity estimated from the peak strain as cyan lines, and the estimated accuracy and precision as red text.

References

[1] A. H. Hartog, *An Introduction to Distributed Optical Fibre Sensors*. CRC Press, May 2017.

[2] P. Jousset, T. Reinsch, T. Ryberg, H. Blanck, A. Clarke, R. Aghayev, G. P. Hersir, J. Henninges, 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.

[3] 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.

[4] 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.

[5] 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.

[6] 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.

[7] 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.

[8] 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.

[9] M. van den Ende, A. Ferrari, A. Sladen, and C. Richard, "Next-Generation Traffic Monitoring with Distributed Acoustic Sensing Arrays and Optimum Array Processing," in *Proceedings of the 2021 55th Asilomar Conference on Signals, Systems, and Computers*. IEEE, 2022.

[10] 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.

[11] A. Beck and M. Teboulle, "A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems," *SIAM Journal on Imaging Sciences*, vol. 2, no. 1, pp. 183–202, Jan. 2009.

[12] M. P. A. van den Ende, A. Ferrari, A. Sladen, and C. Richard, "Deep Deconvolution for Traffic Analysis with Distributed Acoustic Sensing Data," *EarthArxiv*, Sep. 2021.

[13] A. S. Goncharova, A. Honigmann, F. Jug, and A. Krull, "Improving Blind Spot Denoising for Microscopy," in *Computer Vision – ECCV 2020 Workshops*, ser. Lecture Notes in Computer Science, A. Bartoli and A. Fusiello, Eds. Cham: Springer International Publishing, 2020, pp. 380–393.

[14] H. Kobayashi, A. C. Solak, J. Batson, and L. A. Royer, "Image Deconvolution via Noise-Tolerant Self-Supervised Inversion," Jun. 2020.

[15] M. Naphade, M.-C. Chang, A. Sharma, D. C. Anastasiu, V. Jagarlamudi, P. Chakraborty, T. Huang, S. Wang, M.-Y. Liu, R. Chellappa, J.-N. Hwang, and S. Lyu, "The 2018 NVIDIA AI City Challenge," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. Salt Lake City, UT, USA: IEEE, Jun. 2018, pp. 53–537.

[16] M. A. Adnan, N. Sulaiman, N. I. Zainuddin, and T. B. H. T. Besar, "Vehicle speed measurement technique using various speed detection instrumentation," in *2013 IEEE Business Engineering and Industrial Applications Colloquium (BEIAC)*, Apr. 2013, pp. 668–672.

[17] P. G. Hubbard, R. Ou, T. Xu, L. Luo, H. Nonaka, M. Karrenbach, and K. Soga, "Road Deformation Monitoring and Event Detection using Asphalt-embedded Distributed Acoustic Sensing (DAS)," *Engrxiv*, Jan. 2022.