Real-Time High-Precision Detection of Vehicle Trajectories Using DAS

Huijuan WU*, Zhichao WANG, Yuanyuan SONG, Anjie SUN, Chenhui HUANG, Shibo HAN, Yu WU, and Yunjiang RAO*

Key Laboratory of Optical Fiber Sensing & Communications (Ministry of Education of the People's Republic of China), University of Electronic Science and Technology of China, Chengdu 611731, China

*Corresponding authors: Huijuan WU and Yunjiang RAO   E-mails: hjwu@uestc.edu.cn and yjrao@uestc.edu.cn

Abstract: Fiber-optic distributed acoustic sensing (DAS) offers a promising solution for continuous traffic monitoring; however, its widespread deployment is often hindered by poor signal quality, resulting in fragmented and faint vehicle trajectories. Existing techniques – including conventional signal processing and deep learning models – struggle to accurately reconstruct trajectories and estimate traffic parameters under such challenging conditions. To overcome these limitations, we propose the DAS-hierarchical vehicle estimation network (DAS-HiVENet), an end-to-end framework that fundamentally advances the state-of-the-art through three key innovations: a two-stage preprocessing pipeline for noise suppression and trajectory preservation; a novel generative adversarial network (GAN) with an enhanced U-shaped convolutional neural network (U-net) generator to reconstruct high-fidelity trajectories from degraded inputs; a rotated-you only look once (R-YOLO) detector using oriented bounding boxes to accurately detect slanted trajectories. Extensive field evaluations on multiple expressways confirm that it surpasses existing methods with breakthrough performance: a trajectory intersection over union (IoU) of 0.707 6, vehicle counting detection rate of 96.7%, and speed estimation errors as low as 1.422 km/h for the mean absolute error (MAE) and 1.796% for the mean absolute percentage error (MAPE) over 30 minutes. Even in challenging bridge scenarios with severe trajectory adhesion, DAS-HiVENet maintains an over 96% detection rate and under 4% MAPE in speed estimation – significantly outperforming alternatives.

Keywords: DAS; traffic monitoring; vehicle trajectory extraction; speed estimation

Citation: Huijuan WU, Zhichao WANG, Yuanyuan SONG, Anjie SUN, Chenhui HUANG, Shibo HAN, *et al.*, "Real-Time High-Precision Detection of Vehicle Trajectories Using DAS," *Photonic Sensors*, 2026, 16(1): 9560008.

1. Introduction

The rapid global expansion of road networks, notably exemplified by China's 180 000-km expressway system, amplifies the demand for large-scale, continuous traffic monitoring to mitigate pressing safety, and congestion issues. Real-time and high-resolution surveillance is essential to enable the intelligent transportation systems required within the evolving Internet of Things (IoT) framework [1]. However, conventional technologies – including inductive loops [2], video cameras [3–5], and radar systems [6, 7] – are

Received: 30 October 2025 / Revised: 23 December 2025

DOI: 10.26599/PhoS.2026.9560008
Article type: Research article

清华大学出版社 Tsinghua University Press | SciOpen

constrained by inherent drawbacks that limit their effectiveness. Inductive loops provide limited spatial data and are susceptible to physical damage; video systems suffer from poor lighting and adverse weather; the radar involves high cost and interpretation complexity. As a result, none of these established methods simultaneously meets the requirements for wide-area coverage, resilience in all weather conditions, minimal maintenance, and operational reliability.

Fiber-optic distributed acoustic sensing (DAS), based on phase-sensitive optical time-domain reflectometry (Φ-OTDR), has emerged as a promising alternative. By turning fiber cables into continuous vibration sensor arrays, the DAS enables real-time traffic monitoring over tens of kilometers with meter-scale spatial resolution [8–10]. Its potential has been demonstrated in several practical scenarios: repurposed geothermal fiber arrays for traffic analysis [11], beamforming-based tracking during urban lockdowns [12], and hybrid systems combining the DAS with mobile sensing [13].

Despite these advantages, field-deployed DAS systems often suffer from a low signal-to-noise ratio (SNR) and fragmented vehicle trajectories due to uneven fiber-roadbed coupling, pavement heterogeneity, and environmental noises. Current research focuses on two interrelated fronts: enhancing signal quality and extracting traffic parameters. Techniques, such as frequency-wavenumber filtering [14] and adaptive wavelet thresholding [15], have been applied to improve the SNR. For trajectory extraction and parameter estimation, both conventional and machine learning approaches are used – including the Hough transforms [16], multiple signal classification (MUSIC) algorithms [17], and more recent U-shaped convolutional neural network (U-net) [18] and you only look once (YOLO) based models [19, 20]. While these methods have shown promise, they often struggle under real-world low-SNR conditions, where trajectory fragmentation and adhesion severely limit detection accuracy and generalization.

To overcome these challenges, we propose a DAS-hierarchical vehicle estimation network (DAS-HiVENet) – an end-to-end framework integrating a two-stage preprocessing module (combining low-pass filtering and adaptive wavelet denoising), a specialized adversarial network for trajectory reconstruction, and a rotated-you only look once (R-YOLO) model for oriented detection of slanted trajectories. This approach significantly improves trajectory continuity and parameter estimation accuracy under realistic noises and variability.

The rest of this paper is organized as follows: Section 2 introduces the sensing principle and application challenges of the DAS in traffic monitoring. Section 3 details the proposed methodology. Section 4 presents field tests and discusses the results. Section 5 concludes the paper.

2. Challenges of fiber-optic DAS for traffic monitoring

The principle of a Φ-OTDR based DAS system in traffic monitoring is illustrated in Fig. 1. In essence, it transforms an optical fiber into a distributed vibration sensor by detecting coherent Rayleigh backscatter. Vehicle-induced ground vibrations modulate the phase of the backscattered light, which is demodulated to construct a spatiotemporal (S-T) matrix. In this matrix, moving vehicles manifest as diagonal trajectories, whose slopes correlate with their speeds.

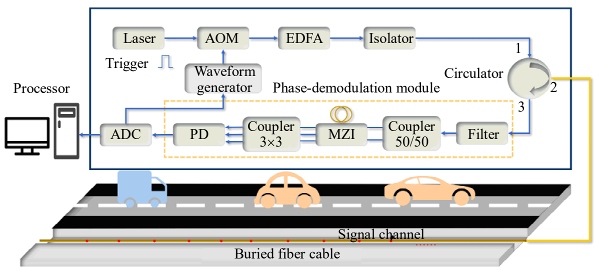

Fig. 1 Fiber-optic DAS system in traffic monitoring.

Despite its conceptual elegance, the practical application of the DAS in real-world traffic monitoring is fundamentally constrained by two pervasive and interconnected challenges: the low SNR and severe trajectory fragmentation/adhesion.

2.1 Inherent low-SNR challenge

The primary obstacle is the inherently weak target signal susceptible to being submerged in ambient noises. Vehicle-induced vibrations experience significant attenuation as they propagate through the roadbed and soil before reaching the fiber. This intrinsic signal weakness is further exacerbated by suboptimal fiber-roadbed coupling, which drastically reduces strain transfer efficiency [21]. Concurrently, the system must contend with strong ambient noises from environmental disturbances and acoustic crosstalk from adjacent lanes, often overwhelming the faint vibrational signatures of interest [22]. These factors collectively create a low-SNR regime that severely impedes reliable initial detection.

2.2 Trajectory integrity challenge

Beyond mere detection, accurately reconstructing continuous vehicle trajectories is complicated by multiple physical phenomena. The low SNR directly causes intermittent signal dropout, resulting in fragmented trajectories. More fundamentally, the spatially non-uniform sensitivity of the fiber itself, caused by variations in cable coupling [21] and inherent optical signal fading [23], creates vibration “blind spots” along the sensing path. This non-uniform response leads to discontinuous and broken trajectory patterns.

Furthermore, in multi-vehicle scenarios, the problem intensifies. The superposition of vibrational wavefields from closely spaced vehicles can cause trajectories to merge, making them inseparable in the S-T domain. A particularly challenging case arises from vehicle-bridge interactions. The bridge structure, acting as an elastic system, continues to “ring” with damped oscillations long after the excitation has passed. This prolonged vibrational response, combined with wave reflections at structural boundaries, leads to severe trajectory adhesion and superposition when multiple vehicles are present on the bridge.

These intertwined data quality issues – rooted in the physics of vibration wave propagation, road structural dynamics, and electro-optic sensing using fibers – create a complex problem landscape. Conventional signal processing algorithms, which assume relatively clean and continuous features, frequently fail under these conditions. The core limitations of the low SNR and trajectory discontinuity/coalescence fundamentally undermine the reliability of traditional methods for vehicle counting and speed estimation, necessitating a more robust and intelligent solution.

3. Vehicle detection using the DAS-HiVENet framework

As illustrated in Fig. 2, the proposed DAS-HiVENet framework detects vehicles in three stages: preprocessing, trajectory extraction, and parameter estimation. The raw DAS signals are first denoised using a combined low-pass and adaptive wavelet filter. To recover weak and fragmented trajectories, a generative adversarial network with PatchGAN supervision is introduced to enhance local structural coherence. The refined trajectories are then fed into an R-YOLO model for estimating the vehicle count and speed based on geometric interpretation. Each component is detailed in the following sections.

3.1 Preprocessing

The preprocessing module processes the one-dimensional (1D) raw signals from each spatial point in the DAS-acquired S-T matrix, serving as the foundation for all subsequent analysis. Raw vehicle-induced vibration signals are non-stationary and contaminated by two dominant noise types: the wide-band environmental/system noise that is irrelevant with vehicle signals, and the in-band low-frequency noise that spectrally overlaps with vehicle vibrations. Since a single denoising strategy is ineffective against both, we design a two-stage preprocessing pipeline as shown in Fig. 2, which systematically suppresses noises while preserves trajectory visibility in the resulting waterfall maps.

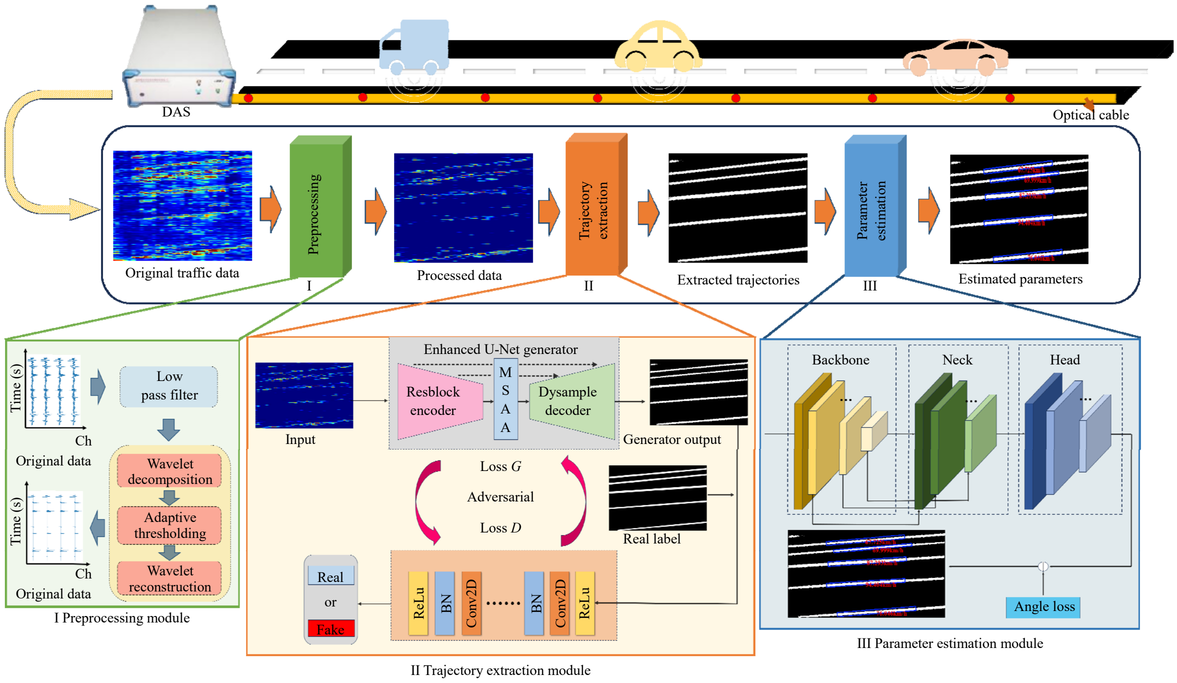

Fig. 2 End-to-end DAS-HiVENet framework for vehicle detection (MSAA: multi-scale attention aggregation; ReLU: rectified linear unit; BN: batch normalization; Conv2D: two-dimensional convolution).

Vehicle vibration energy is primarily concentrated in low-frequency bands. We first apply a 0 Hz–20 Hz Butterworth low-pass filter to remove the high-frequency noise unrelated to vehicle signatures, such as electronic interference or ambient vibration, thereby establishing a cleaner signal baseline for subsequent processing.

Despite filtering, the in-band noise persisting in the vehicle signal frequency range remains challenging. Conventional linear filters tend to distort useful signals during separation. Although wavelet transforms offer excellent time-frequency localization, fixed-threshold denoising often leads to over-smoothing of weak vehicle components or incomplete noise removal under fluctuating signal conditions.

To address this, we introduce an adaptive wavelet thresholding scheme driven by the input signal energy distribution. The procedure involves three stages: wavelet decomposition, adaptive thresholding, and signal reconstruction. First, each temporal signal in the S-T matrix is decomposed via the discrete wavelet transform (DWT). The wavelet coefficients are updated based on cumulative energy from the lowest frequency band upward:

$$W'_{j,k}=\begin{cases}\operatorname{sign}(W_{j,k})\cdot\max(|W_{j,k}|-\mathrm{th},0), & \text{if } \sum_{m=j}^{J}\left(E_m/E_{\mathrm{total}}\right)>\gamma \\ W_{j,k}, & \text{otherwise}\end{cases} \tag{1}$$

where $W_{j,k}$ denotes the coefficients at the scale j and translation k, and $W'_{j,k}$ is its updated version. At each scale, the ratio of the cumulative energy to the total energy is compared to a preset energy retention threshold, γ=0.99. If this ratio is below γ, coefficients remain unchanged to preserve low-frequency content. Otherwise, soft-thresholding is applied, removing only noise-dominated high-frequency components that constitute about 1% of total energy. The maximum noise coefficient threshold, $\mathrm{th}=\sqrt{E_{\text{total}}\cdot(1-\gamma)/J}$, is determined under the assumption of the uniform noise distribution per level. Finally, the denoised signal is reconstructed via the inverse DWT.

3.2 Trajectory reconstruction

Following preprocessing, denoised signals are reassembled into a 2D spatiotemporal map. However, real-world DAS data often contain unclear or fragmented trajectories, for which conventional U-net models [18] fail to reconstruct accurately due to limited global context capture and reliance on the pixel-wise loss.

To address these challenges, we develop a customized adversarial network, as depicted in Fig. 3, building upon the foundational framework of [24]. While our approach retains the adversarial training scheme between a generator and a PatchGAN discriminator to ensure structural realism, we have specifically redesigned the generator architecture to meet the demands of DAS trajectory reconstruction. The key customizations include the follows.

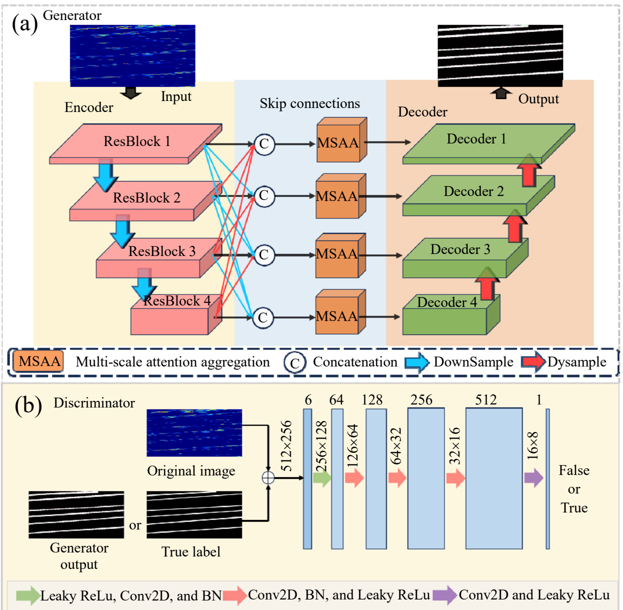

Fig. 3 Architecture of the proposed adversarial network for trajectory extraction, comprising: (a) a generator based on an enhanced U-net structure and (b) a PatchGAN-based discriminator.

1) ResBlock encoder: built on residual blocks, the encoder mitigates vanishing gradients and supports deeper networks, enabling effective extraction of multi-scale features – essential for detecting faint vehicle trajectories obscured by noises.

2) MSAA-based skip connection: replacing standard skip connections, the MSAA module aggregates multi-resolution encoder features and refines them via dual spatial and channel attention pathways. This enhances global contextual awareness and emphasizes salient features, critical for reconstructing trajectory continuity.

3) Dynamic upsampling (Dysample) decoder: replacing conventional upsampling, Dysample treats upsampling as a learnable point resampling process. It generates adaptive sampling grids via predicted offsets, enabling finer reconstruction of weak or thin trajectories. The sampling set S is derived as

$$S = G + O \tag{2}$$

where G ($G \in R^{2g\times sH\times sW}$) is the base grid and O ($O \in R^{2g\times sH\times sW}$) denotes the learned offsets, rescaled to prevent boundary overlap.

This generator is trained adversarial against a PatchGAN discriminator, which evaluates local patch authenticity to enforce structural realism. The architecture and implementation of the proposed adversarial network are detailed in S1, Supplementary files.

The overall objective for the trajectory extraction module integrates adversarial and reconstruction losses. The generator loss is defined as

$$\mathcal{L}_G = \mathcal{L}_{\text{adv}} + \lambda\mathcal{L}_{\text{pixel}} \tag{3}$$

where the adversarial loss $\mathcal{L}_{\text{adv}} = -\ln D(x, G(x))$, and the pixel-wise loss $\mathcal{L}_{\text{pixel}} = \|y - G(x)\|_1$ that is the L_1 loss between the generated image $G(x)$ and the ground truth y, weighted by λ, where λ=100 ensures that the generated trajectories are well aligned with the ground-truth spatial structure while maintaining clear and realistic boundaries.

Concurrently, we adopt a PatchGAN discriminator [25] that assesses local image patches for authenticity. The discriminator is optimized using a standard loss for distinguishing real and synthetic samples as

$$\mathcal{L}_D = \mathcal{L}_{\text{real}} + \mathcal{L}_{\text{fake}} \tag{4}$$

where

$$\mathcal{L}_{\text{real}} = -\ln D(x, y)$$
$$\mathcal{L}_{\text{fake}} = -\ln\{1 - D[x, G(x)]\}$$

which denote the losses on real and generated samples, respectively. Details about the discriminator and the evaluation metrics for trajectory reconstruction can be found in S2 and S3, Supplementary files.

3.3 Parameter estimation

The parameter estimation aims to extract vehicle count and speed accurately and in real time from the reconstructed trajectory maps. Under normal highway operating conditions, without abrupt speed changes or congestion-induced trajectory curvature, vehicle trajectories always appear as slanted lines in spatiotemporal images, where the slope encodes speed. Conventional YOLO detectors, using axis-aligned bounding boxes (HBBs), are ill-suited for this task due to poor localization of angled features, inability to resolve overlapping trajectories, and lack of a direct geometric link to speed.

To address these issues, we introduce R-YOLO borrowed from [26], an oriented object detection model based on YOLOv8 oriented bounding boxes (OBBs). As illustrated in Fig. 4, R-YOLO extends the bounding box representation to five parameters (x_c, y_c, w, h, θ) where θ is the angle between the longer box side and the x-axis. This allows the detector to align precisely with slanted trajectories. The comprehensive architecture of the R-YOLO model and its implementation details are provided in S4, Supplementary files.

The vehicle speed is directly derived from the orientation θ of the detected rotated box:

$$v = \frac{\Delta d_{\text{pixel}}}{\Delta t_{\text{pixel}} \cdot \tan\theta} \tag{5}$$

where Δd_{pixel} and Δt_{pixel} denote the spatial and temporal sampling intervals per pixel.

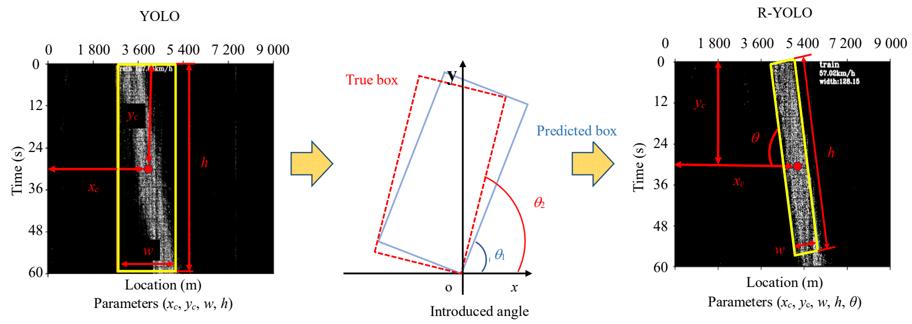

Fig. 4 Schematic diagram illustrating the angle regression principle of R-YOLO for detecting slanted vehicle trajectories.

4. Field test and results discussion

4.1 Field validation and results analysis

4.1.1 Experimental setup

Field validation was performed on a 1-km segment of the Beijing-Harbin (Jingha) Expressway (Qinhuangdao, China) from March to July 2025. As shown in Fig. 5, a commercial DAS system was deployed with a sampling rate of 2 kHz and spatial resolution of 10 m, providing single-ended detection over 40 km. A dual-core single-mode fiber (the core diameter: 0.25 mm) was installed along the road shoulder and cemented in place to ensure effective soil-fiber coupling [Figs. 5(d) and 5(e)]. The test section comprised a four-lane bidirectional roadway with embankments, cut sections, and curves, experiencing over 5 000 vehicles per day – more than 25% of which were heavy trucks – representing realistic highway conditions.

A two-stage calibration procedure was applied to georeferenced the fiber channels: initial coarse alignment based on the 10 m resolution, followed by precision validation using hammer tests at landmarks (e.g., bridges and culverts). All 100 sensing channels were registered with the Global Positioning System (GPS) coordinates in a fiber calibration database, enabling vehicle localization within ±10 m and establishing a reliable foundation for trajectory analysis.

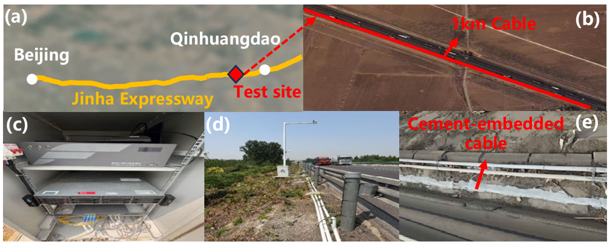

Fig. 5 Deployment of the DAS monitoring system at the Jingha Expressway field site: (a) satellite overview, (b) experimental site, and (c)–(e) DAS equipment and cement-embedded optical fiber installation.

4.1.2 Preprocessing results

Through comparative experiments, preprocessing parameters were optimized to a 20 Hz low-pass cutoff, Daubechies-22 wavelet, and 5-level decomposition. A spatiotemporal frame was defined as 500 m×120 s, under the assumption of the near-constant vehicle speed over 500 m, yielding quasi-linear trajectories.

Figure 6 presents the comparative visualization of the waterfall plots before and after preprocessing, where the horizontal axis spans 50 spatial channels with 10 m intervals and the vertical axis represents a 120 s duration, accompanied by 1D profiles illustrating the specific vibration signal extracted from the 46th channel. As observed, the raw maps [Figs. 6(a) and 6(b)] are heavily contaminated by the ambient noise that obscures the vehicle-induced diagonal streaks; in contrast, the preprocessed results [Figs. 6(c) and 6(d)] demonstrate effective background noise suppression and reduced trajectory adhesion, yielding clearly separated and enhanced vehicle tracks. This superior performance was further evaluated by additional comparisons with traditional preprocessing techniques, such as 2D median filtering and VisuShrink wavelet denoising, which are detailed in S5, Supplementary files. The results validate that our preprocessing pipeline behaves better and can ensure high-quality inputs for subsequent trajectory extraction.

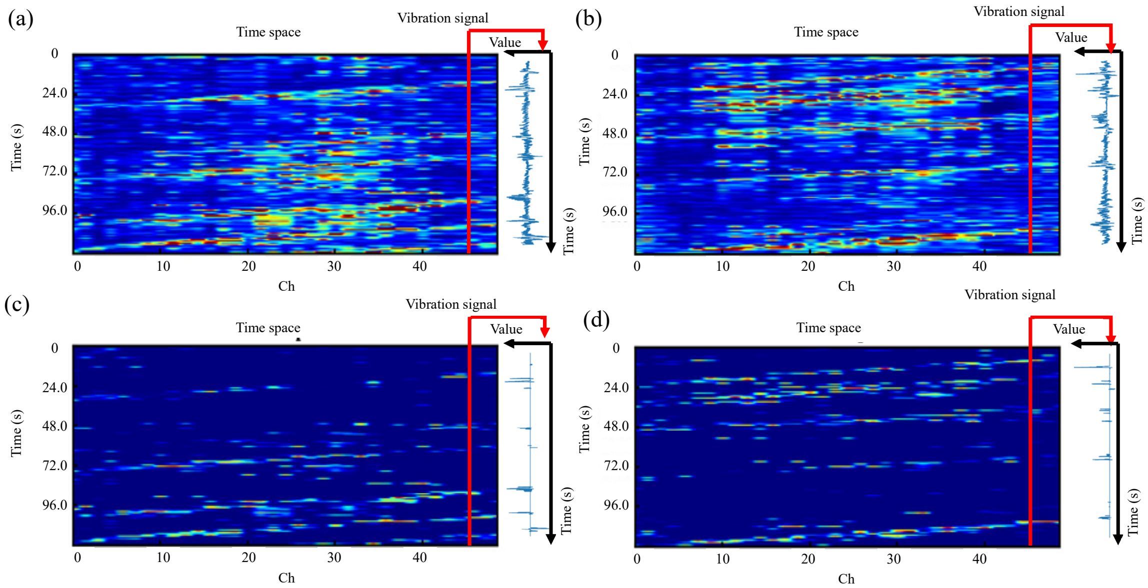

Fig. 6 Original and preprocessed data from two expressway segments: (a) and (b) original DAS monitoring data, (c) preprocessed result of (a), and (d) preprocessed result of (b).

4.1.3 Trajectory reconstruction

We constructed a dataset of 319 training and 15 validation image pairs, which included only subgrade road sections, with no bridge or culvert samples. Each pair consisted of a preprocessed waterfall map and its manually annotated trajectory mask, in which a global-context annotation strategy was used, filling missing trajectories via linear interpolation under a short-range constant-velocity assumption. In this case, the image resolution was set as (512 pixel×256 pixel) to make tradeoff between the parameter estimation precision and the calculation efficiency. Our enhanced U-net-based GAN was compared against a standard U-net [18] and a U-net-based GAN [25] under identical training conditions.

Qualitative results are shown in Figs. 7 and 8. In Fig. 7, two faint, discontinuous trajectories (marked by orange arrows) are largely missed or fragmented by the baseline models [Figs. 7(b) and 7(c)]. By contrast, our model [Fig. 7(d)] successfully reconstructed them into complete and coherent paths. Similar improvements are observed in Fig. 8, where our method recovers weak and broken segments that other models fail to reconstruct.

To complement the qualitative assessment, we employ three quantitative metrics for rigorous evaluation of the trajectory reconstruction performance: intersection over union (IoU), Dice similarity coefficient (Dice), and multi-scale structural similarity (MS-SSIM). Computation details can be seen in the S3, Supplementary files.

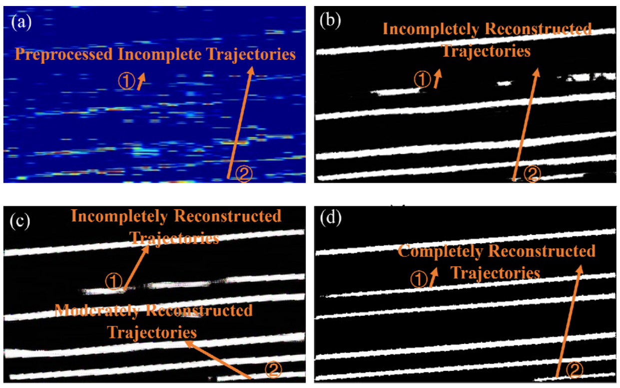

Fig. 7 Comparative outputs of trajectory reconstruction in Fig. 6(c): (a) Input data, (b) standard U-net [18], (c) U-net-based GAN [25], and (d) enhanced U-net-based GAN (ours).

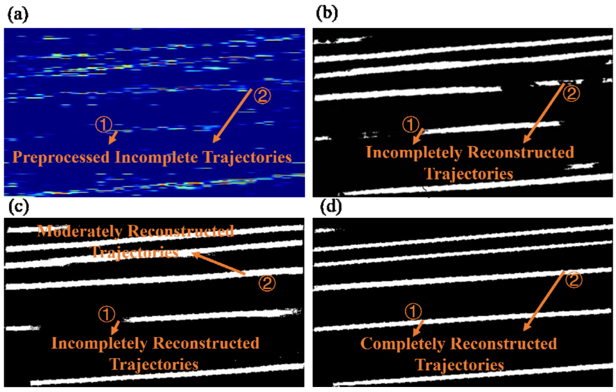

Fig. 8 Comparative outputs of trajectory reconstruction in Fig. 6(d): (a) input data, (b) standard U-net [18], (c) U-net-based GAN [25], and (d) enhanced U-net-based GAN (ours).

As summarized in Table 1, our model achieves an IoU of 0.707 6, Dice of 0.828 5, and MS-SSIM of 0.910 3, outperforming the U-net-based GAN by 17.85%, 10.61%, and 11.60%, and the standard U-net by 21.71%, 12.81%, and 15.20%, respectively. These results confirm that our adversarial training and architectural enhancements lead to more accurate and structurally realistic trajectory recovery.

Table 1 Performance comparison of IoU, Dice, and MS-SSIM metrics among different trajectory reconstruction networks.

Network	IoU	Dice	MS-SSIM
Enhanced U-net based GAN (ours)	0.707 6	0.828 5	0.910 3
U-net based GAN	0.600 4	0.749 0	0.815 7
Standard U-net	0.581 4	0.734 4	0.790 2

4.1.4 Traffic parameter estimation

Reconstructed trajectories from Figs. 7(d) and 8(d) were used to evaluate vehicle counting and speed estimation performance. As illustrated in Fig. 9, axis-aligned bounding boxes from the standard YOLO [Figs. 9(a) and 9(c)] only coarsely enclosed slanted trajectories, forcing speed estimates to rely indirectly on the bounding box diagonal. In contrast, R-YOLO [Figs. 9(b) and 9(d)] used rotated bounding boxes to align precisely with trajectory orientation, enabling direct speed estimation from the box angle.

For a 30-minute period (09:18–09:48, June 15, 2025), ground-truth manual inspection identified 90 vehicles. The R-YOLO detected 87 vehicles (96.67% detection rate), compared to 78 (86.67%) by YOLO (Fig. 10). The undercount by the R-YOLO was mainly due to occasional trajectory merging in adjacent lanes.

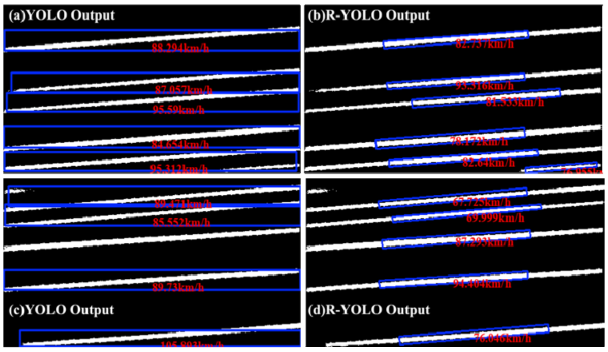

Fig. 9 Vehicle parameter estimation with the conventional YOLO (HBB) and the proposed R-YOLO (OBB): (a) and (c) outputs from YOLO; (b) and (d) outputs from R-YOLO.

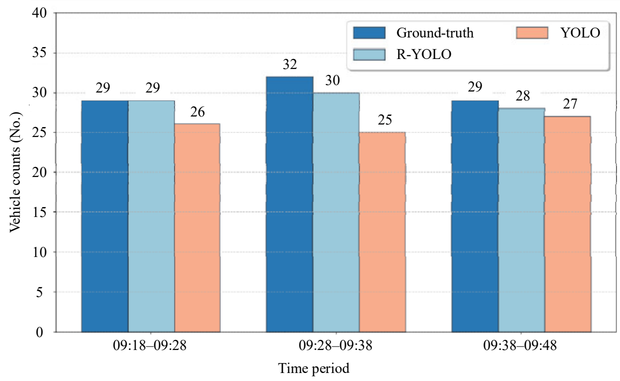

Fig. 10 Vehicle count at 10-minute intervals over 30 minutes.

Speed estimation results are shown in Fig. 11. The R-YOLO estimates [Fig. 11(a), light blue] closely matched the ground-truth distribution (dark blue), with errors tightly clustered near zero [Fig. 11(b)]. The YOLO estimates (orange) showed a consistent greater dispersion [Fig. 11(c)], stemming from its reliance on the bounding box diagonal as an inaccurate proxy for the trajectory slope.

As summarized in Table 2, the R-YOLO achieved a mean absolute error (MAE) of 1.422 km/h and mean absolute percentage error (MAPE) of 1.796% in speed estimation – nearly an order of the magnitude better than the YOLO (MAE: 11.468 km/h and MAPE: 14.641%). This confirms the fundamental advantage of direct angle-based speed measurement in the oriented object detection.

The computational efficiency of DAS-HiVENet is illustrated in Fig. 12. Using a standard workstation with an AMD Ryzen 5 5600G CPU and 32 GB RAM, the entire pipeline completed in just 1.58 s, dominated by preprocessing (0.98 s), with trajectory extraction and parameter estimation requiring only 0.44 s and 0.16 s, respectively. This efficiency allows the system to process a 37.98 km-long and two-minute DAS waterfall dataset within the acquisition time frame, demonstrating its strong capability for online real-time traffic monitoring.

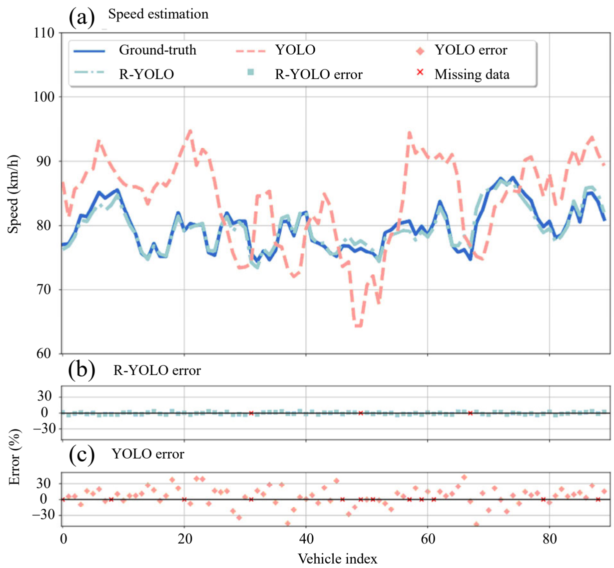

Fig. 11 Vehicle speed estimation over a 30-minute period: (a) estimated versus ground-truth speeds; (b) error distribution of the R-YOLO, and (c) error distribution of the YOLO.

Table 2 Quantitative comparison of vehicle count and speed estimation performance.

Network	Vehicle count detection rate (%)	MAE (km/h)	MAPE (%)
R-YOLO	96.67	1.422	1.796
YOLO	86.67	11.468	14.641

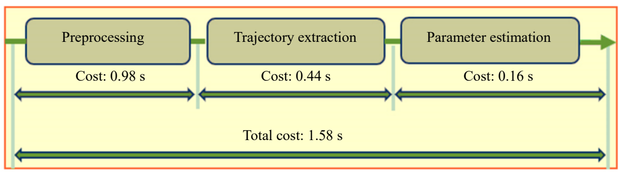

Fig. 12 Time consumption of DAS-HiVENets.

4.2 Cross-scenario validation

To assess generalization, additional tests were conducted on the Qingyin Expressway (Shijiazhuang, China), using a similar setup except for a 1 kHz sampling rate. A 2 km fiber section was deployed, including a 200 m bridge, a 10 m culvert, and standard embankment subgrade. Two representative segments were analyzed.

1) Segment 1 (flat embankment): as shown in Fig. 13, trajectories are relatively clear, though signals from opposite lanes were weak, they were successfully suppressed as noise during preprocessing.

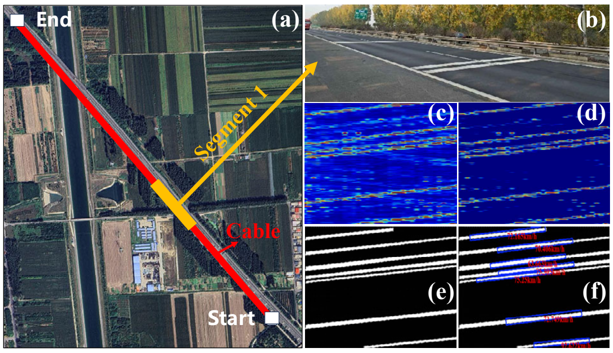

Fig. 13 Monitoring overview and processing results for Segment 1: (a) overview of the fiber layout, (b) field monitoring site, and (c)–(f) outputs from successive processing stages.

2) Segment 2 (bridge deck): in Fig. 14, bridge-induced vibration prolongation causes significant trajectory broadening and adhesion (yellow dotted region). Nonetheless, the proposed pipeline maintained robust performance.

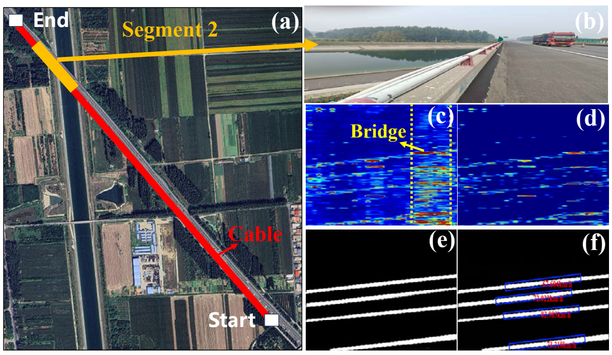

Fig. 14 Monitoring overview and processing results for Segment 2: (a) overview of the fiber layout, (b) field monitoring site, and (c)–(f) outputs from successive processing stages.

The network used in this cross-scenario validation was the same with that of the model trained on the Jingha Expressway dataset in Sub-Section 4.1. Comparative analysis of raw vibration patterns revealed distinct trajectory characteristics between subgrade types. On solid embankments [Fig. 13(c)], trajectories maintained clearer linear structures despite the moderate noise levels. In contrast, bridge sections [Fig. 14(c)] exhibited significant trajectory broadening and adhesion due to prolonged structural vibrations, substantially complicating detection. Notably, DAS-HiVENet consistently delivered robust performance in both scenarios, effectively denoising inputs, reconstructing complete trajectories, and accurately estimating parameters [Figs. 13(d)–13(f), and 14(d)–14(f)].

We further evaluated our framework against two mainstream classes of detection methods: the widely adopted deep learning-based YOLO [20] and the conventional image-processing-based Hough transform [16]. As visualized in Fig. 15 and summarized in Table 3, DAS-HiVENet significantly outperformed both representative approaches across test segments. In Segment 1, it achieved a 97.11% counting detection rate with speed estimation errors of 2.43 km/h (MAE) and 3.26% (MAPE), compared to an 81.67% detection rate and over 11 km/h MAE for the YOLO, and a below 68% detection rate for the Hough transform. Under the more challenging bridge conditions in Segment 2, it maintained a 96.43% detection rate, 3.00 km/h MAE, and 3.99% MAPE, demonstrating strong generalization. In Fig. 15, it can be also seen that the ineffective speed detection results marking with the red crosses are increasing largely, because the unreasonable detected speed values increased a lot. Besides, the two existing methods, YOLO detection [20] and the Hough transform [16], exhibited a sharp increase in ineffective speed detections (marked with red crosses), directly resulting from a rise in unreasonable values. In contrast, the proposed

DAS-HiVENet yielded results that closely match the ground truth, resulting in significantly fewer invalid readings.

Quantitative analysis further confirmed its superior trajectory retention, with vehicle tracking error rates of only 2.89% and 3.57% in Segments 1 and 2, respectively. In contrast, both mainstream methods suffered from significantly higher error rates (18.23%–26.36%), failing to detect faint trajectories or misclassifying noises as valid signals. These results validated the ability of DAS-HiVENet to enhance weak signatures and reconnect fragmented trajectories under real-world noisy and overlapping conditions, offering a clear advantage over existing representative techniques.

Despite challenges such as trajectory adhesion on bridge structures, the system consistently maintained an over 96% detection rate and under 4% MAPE in speed estimation. However, several practical factors were observed to notably influence the DAS signal quality and detection performance in operational highway environments as follows.

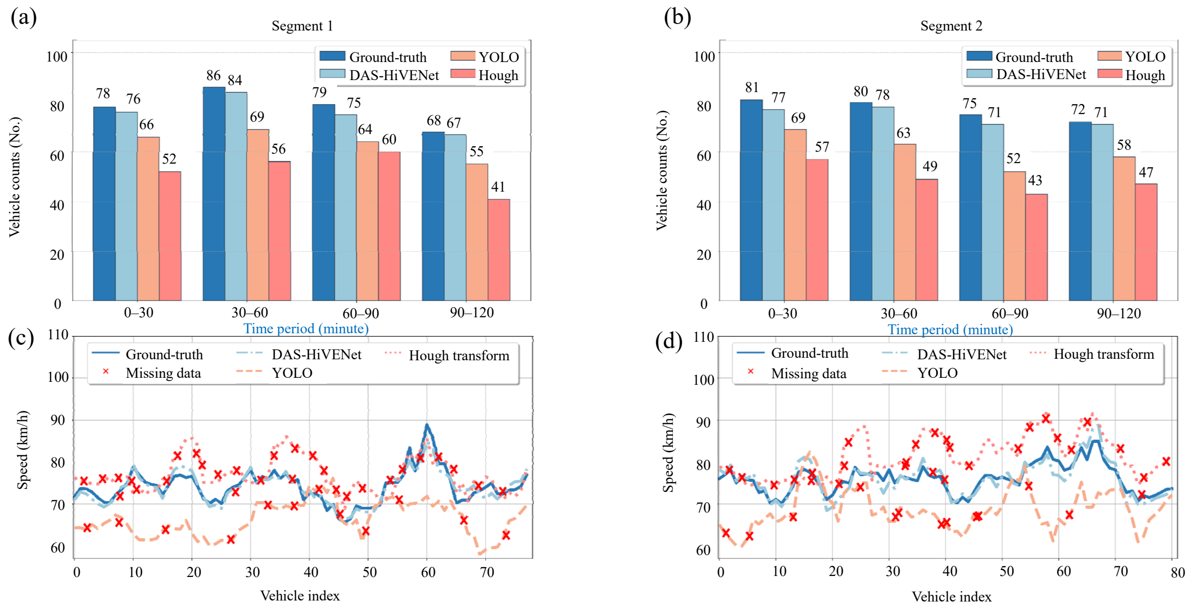

Fig. 15 Performance comparison of three vehicle detection methods, the DAS-HiVENet (proposed), YOLO [20], and Hough transform [16]: (a)–(b) two-hour continuous vehicle counts statistics in Segments 1 and 2, and (c)–(d) speed estimation over a 30-minute interval within the same period.

Table 3 Quantitative comparisons of different methods in the two segments monitoring.

	Methods	Vehicle count detection rate (%)	Speed estimation (MAE) (km/h)	Speed estimation (MAPE) (%)
Segment 1	DAS-HiVENet (proposed)	97.11	2.43	3.26
	YOLO detection [20]	81.67	11.25	14.64
	Hough transform [16]	67.20	7.98	10.64
Segment 2	DAS-HiVENet (proposed)	96.43	3.00	3.99
	YOLO detection [20]	78.57	10.92	14.36
	Hough transform [16]	63.64	8.20	11.12

1) Fiber deployment: coupling efficiency varied with installation (surface-mounted, buried, ducted, etc.). Surface-mounted and cement-fixed cabling yielded the clearest trajectories, which is currently recommended.

2) Subgrade structures: bridge decks and culverts exhibited prolonged vibrational responses, complicating trajectory separation. The method remained effective up to 200 vehicles/h under such conditions.

3) Environmental noises: wind, traffic harmonics, and instrumental noises raised the noise floor. The standard DAS hardware was proved sufficient under typical conditions. The provided trajectories remained visually discernible after preprocessing. Cases with a prohibitively low optical SNR or extreme environmental interference, rendering trajectories indiscernible visually, represent fundamental limitations for accurate algorithmic interpretation.

5. Conclusions

This study presents an end-to-end framework that transforms low-SNR DAS signals into precise and real-time traffic metrics through three key innovations: a dual-stage preprocessing pipeline for signal enhancement, a novel GAN architecture for trajectory reconstruction, and an R-YOLO detector enabling geometric-accurate parameter estimation. Field validation demonstrates exceptional performance, achieving the 0.707 6 trajectory IoU, 96.7% counting detection rate, and speed estimation errors of 1.422 km/h (MAE) and 1.796% (MAPE). The system maintains an over 96% detection rate and under 4% MAPE even under challenging bridge vibration scenarios, significantly surpassing existing approaches. This work establishes a robust and scalable solution for reliable DAS-based traffic monitoring in complicated real-world environments.

Supplementary files

The supplementary files can be downloaded from the webpage of this article in *Photonic Sensors* (https://www.sciopen.com/journal/1674-9251).

Acknowledgment

This work is supported by the National Natural Science Foundation of China (Grant Nos. 62575055 and U21A20453), and supported by the Sichuan Regional Innovation Cooperation Project (Grant No. 2026YFHZ0141).

Declarations

Conflict of Interest Yunjiang RAO is the editor-in-chief for Photonic Sensors and was not involved in the editorial review, or the decision to publish this article. All authors declare that there are no competing interests.

Permissions All the included figures, tables, or text passages that have already been published elsewhere have obtained the permission from the copyright owner(s) for both the print and online format.

Use of AI Statement None.

References

[1] C. Creβ, Z. Bing, and A. C. Knoll, "Intelligent transportation systems using roadside infrastructure: a literature survey," *IEEE Transactions on Intelligent Transportation Systems*, 2023, 25(7): 6309–6327.

[2] T. Cherrett, H. Bell, and M. McDonald, "Traffic management parameters from single inductive loop detectors," *Transportation Research Record*, 2000, 1719(1): 112–120.

[3] E. Bas, A. M. Tekalp, and F. S. Salman, "Automatic vehicle counting from video for traffic flow analysis," in *IEEE Intelligent Vehicles Symposium*, Istanbul, Turkey, Jun. 13–15, 2007, pp. 392–397.

[4] B. Li and R. Chellappa, "A generic approach to simultaneous tracking and verification in video," *IEEE Transactions on Image Processing*, 2002, 11(5): 530–544.

[5] L. Xie, G. X. Zhu, Y. Q. Wang, H. X. Xu, and Z. M. Zhang, "Real-time vehicles tracking based on Kalman filter in a video-based ITS," in *2005 International Conference on Communications, Circuits and Systems, 2005*, Hong Kong, China, May 27–30, 2005, pp. 886.

[6] J. M. Munoz-Ferreras, F. Perez-Martinez, J. Calvo-Gallego, A. Asensio-Lopez, B. P. Dorta-Naranjo, and A. Blanco-del-Campo, "Traffic surveillance system based on a high-resolution radar," *IEEE Transactions on Geoscience and Remote Sensing*, 2008, 46(6): 1624–1633.

[7] D. A. Grejner-Brzezinska, C. K. Toth, S. Moafipoor, E. Paska, and N. Csanyi, "*Vehicle classification and traffic flow estimation from airborne LiDAR/CCD data*," in Dynamic Planet. International Association

of Geodesy Symposia, P. Tregoning, C. Rizos, Eds. Berlin: Springer, 2007: 799–808.

[8] R. Juškaitis, A. M. Mamedov, V. T. Potapov, and S. V. Shatalin, "Interferometry with Rayleigh backscattering in a single-mode optical fiber," *Optics Letters*, 1994, 19(3): 225–227.

[9] K. N. Choi and H. F. Taylor, "Spectrally stable Er-fiber laser for application in phase-sensitive optical time-domain reflectometry," *IEEE Photonics Technology Letters*, 2003, 15(3): 386–388.

[10] Y. Rao, Z. Wang, H. Wu, Z. Ran, and B. Han, "Recent advances in phase-sensitive optical time domain reflectometry," *Photonic Sensors*, 2021, 11(1): 1–30.

[11] K. Chambers, "Using DAS to investigate traffic patterns at Brady Hot Springs, Nevada, USA," *The Leading Edge*, 2020, 39(11): 819–827.

[12] X. Wang, Z. W. Zhan, E. F. Williams, M. G. Herráez, H. F. Martins, and M. Karrenbach, "Ground vibrations recorded by fiber-optic cables reveal traffic response to COVID-19 lockdown measures in Pasadena, California," *Communications Earth & Environment*, 2021, 2(1): 160.

[13] S. Yuan, J. Liu, H. Y. Noh, and B. Biondi, "Urban system monitoring using combined vehicle onboard sensing and roadside distributed acoustic sensing," in *SEG International Exposition and Annual Meeting*, Denver, USA, 2021, pp. segam2021–3584136.1.

[14] H. Wang, Y. Chen, R. Min, and Y. Chen, "Urban DAS data processing and its preliminary application to city traffic monitoring," *Sensors*, 2022, 22(24): 9976.

[15] S. T. Hou, Y. J. Li, G. Wu, D. Wu, Y. X. Dong, S. Y. Zhang, *et al.*, "A method for visualizing urban road events using distributed acoustic sensing," *Urban Lifeline*, 2024, 2(1): 6.

[16] C. Wiesmeyr, C. Coronel, M. Litzenberger, H. J. Döller, H. B. Schweiger, and G. Calbris, "Distributed acoustic sensing for vehicle speed and traffic flow estimation," in *2021 IEEE International Intelligent Transportation Systems Conference*, Indianapolis, USA, 2021, pp. 2596–2601.

[17] 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 *55th Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, USA, 2021, pp. 1104–1108.

[18] F. Peng, Z. Y. Zhu, Y. J. Zhang, Q. Miao, "A deep learning image segmentation model for detection of weak vehicle-generated quasi-static strain in distributed acoustic sensing," *IEEE Transactions on Intelligent Transportation Systems*, 2025, 26(6): 8933–8944.

[19] Z. Wang, T. Y. Zhang, H. L. Chen, C. C. Zhang, and B. Shi, "Enhancing traffic monitoring with noise-robust distributed acoustic sensing and deep learning," *Journal of Applied Geophysics*, 2025, 233: 105616.

[20] Z. Ye, W. Wang, X. Wang, F. Yang, F. Peng, K. Yan, *et al.*, "Traffic flow and vehicle speed monitoring with the object detection method from the roadside distributed acoustic sensing array," *Frontiers in Earth Science*, 2023, 10: 992571.

[21] P. J. Thomas, Y. Heggelund, I. Klepsvik, J. Cook, E. Kolltveit, and T. Vaa, "The performance of distributed acoustic sensing for tracking the movement of road vehicles," *IEEE Transactions on Intelligent Transportation Systems*, 2023, 25(6): 4933–4946.

[22] J. B. Ajo-Franklin, S. Dou, N. J. Lindsey, I. Monga, C. Tracy, M. Robertson, *et al.*, "Distributed acoustic sensing using dark fiber for near-surface characterization and broadband seismic event detection," *Scientific Reports*, 9(1), 2019: 1328.

[23] F. Muñoz and M. A. Soto, "Enhancing fibre-optic distributed acoustic sensing capabilities with blind near-field array signal processing," *Nature Communications*, 2022, 13(1): 4019.

[24] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, *et al.*, "Generative adversarial nets," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Montreal, Canada, Dec. 8–13, 2014, pp. 1125–1134.

[25] P. Isola, J. Y. Zhu, T. H. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1125–1134.

[26] X. Yu, M. Lin, J. Lu, and L. L. Ou, "Oriented object detection in aerial images based on area ratio of parallelogram," *Journal of Applied Remote Sensing*, 2022, 16(3): 034510–034510.