

Graph-Enhanced Dual-Stream Feature Fusion with Pre-Trained Model for Acoustic Traffic Monitoring

Shitong Fan^{1†}, Feiyang Xiao^{1†}, Wenbo Wang², Shuhan Qi³, Qiaoxi Zhu⁴, Wenwu Wang⁵, Jian Guan^{1*}

¹Group of Intelligent Signal Processing, Harbin Engineering University, Harbin, China

²Faculty of Computing, Harbin Institute of Technology, Harbin, China

³School of Computer Science and Technology, Harbin Institute of Technology, Shenzhen, China

⁴Acoustics Lab, University of Technology Sydney, Ultimo, Australia

⁵Centre for Vision, Speech and Signal Processing, University of Surrey, Guildford, UK

Abstract—Microphone array techniques are widely used in sound source localization and smart city acoustic-based traffic monitoring, but these applications face significant challenges due to the scarcity of labeled real-world traffic audio data and the complexity and diversity of application scenarios. The DCASE Challenge’s Task 10 focuses on using multi-channel audio signals to count vehicles (cars or commercial vehicles) and identify their directions (left-to-right or vice versa). In this paper, we propose a graph-enhanced dual-stream feature fusion network (GEDF-Net) for acoustic traffic monitoring, which simultaneously considers vehicle type and direction to improve detection. We propose a graph-enhanced dual-stream feature fusion strategy which consists of a vehicle type feature extraction (VTFE) branch, a vehicle direction feature extraction (VDFE) branch, and a frame-level feature fusion module to combine the type and direction feature for enhanced performance. A pre-trained model (PANNs) is used in the VTFE branch to mitigate data scarcity and enhance the type features, followed by a graph attention mechanism to exploit temporal relationships and highlight important audio events within these features. The frame-level fusion of direction and type features enables fine-grained feature representation, resulting in better detection performance. Experiments demonstrate the effectiveness of our proposed method. GEDF-Net is our submission that achieved 1st place in the DCASE 2024 Challenge Task 10.

Index Terms—Acoustic-based traffic monitoring, transfer learning, pre-trained model, graph attention, feature fusion

I. INTRODUCTION

Acoustic-based traffic monitoring uses roadway sounds to estimate vehicle counts, speeds, and types, aiding in traffic control and anomaly detection [1]–[4]. Typically, single sensors are deployed along the roads to capture audio data for estimating speeds [4] or counting vehicles [1]–[3]. While being effective for speed estimation [4], single-sensor setups struggle with vehicle direction detection due to limited spatial resolution and challenges in distinguishing overlapping signals from multiple vehicles, which restricts the system’s ability for traffic flow monitoring. To address these limitations, microphone arrays are used to capture multi-channel signals from various positions, enabling direction detection [5]–[8]. Configurations include equidistant [6], orthogonal [7], and circular arrays

[8]. However, their effectiveness is often compromised by the scarcity of labeled real-world traffic data.

To address these challenges, the Detection and Classification of Acoustic Scenes and Events (DCASE) 2024 Challenge introduces Task 10, Acoustic-based Traffic Monitoring [3], [9], which focuses on detecting vehicle types (cars and commercial vehicles) and travel directions (left-to-right or vice versa). Due to the scarcity of real-world data, the Challenge incorporates synthetic data to evaluate its impact on system performance [9]. The Task 10 baseline [3], [10] uses a dual-branch CNN-based network to extract vehicle direction and type features from Generalized Cross-Correlation with Phase Transform (GCC-PHAT) [11] and log-Mel spectrograms, which are fused along the temporal dimension to detect vehicle events. To mitigate data scarcity, the baseline and other top-performing systems [12]–[16], including ours [17], all use pre-training on synthetic data [9] followed by fine-tuning on limited real data. The Top-3 system [13] improves the baseline by adding a matching loss [18] to improve the alignment of the predictions with ground truth, achieving slightly enhanced performance.

Other systems, such as Top-4 [14], Top-5 [15], and Top-6 [16], utilize CNN-based architectures like ResNet [19] (Top-4) and VGG11 [20] (Top-6), with Top-5 also exploring ensemble methods. However, these methods do not outperform the baseline. A common limitation among these methods, including Top-3 [13], is their limitation in capturing contextual relationships between audio events. In addition, while synthetic data supports feature learning, it does not fully resolve data scarcity and quality issues, limiting effective representation of the vehicle audio events.

The Top-2 system [12] combines vehicle direction (GCC) and type features (log-Mel spectrogram) along the temporal dimension and employs a Transformer [21] for contextual modeling. However, directly using concatenated features without refinement can introduce redundancy and degrade performance. Like above methods [3], [13]–[16], the Top-2 system [12] also adopts synthetic data to improve feature learning, which, however, can be limited by the quality of the data.

In this paper, we propose GEDF-Net, a graph-enhanced dual-stream feature fusion network with a pre-trained model, for acoustic-based traffic monitoring. The proposed GEDF-Net consists of a graph-enhanced dual-stream feature extraction

[†]These authors contributed equally to this work.

*Corresponding author.

This work was partly supported by the Natural Science Foundation of Heilongjiang Province under Grant No. LH2022F010.

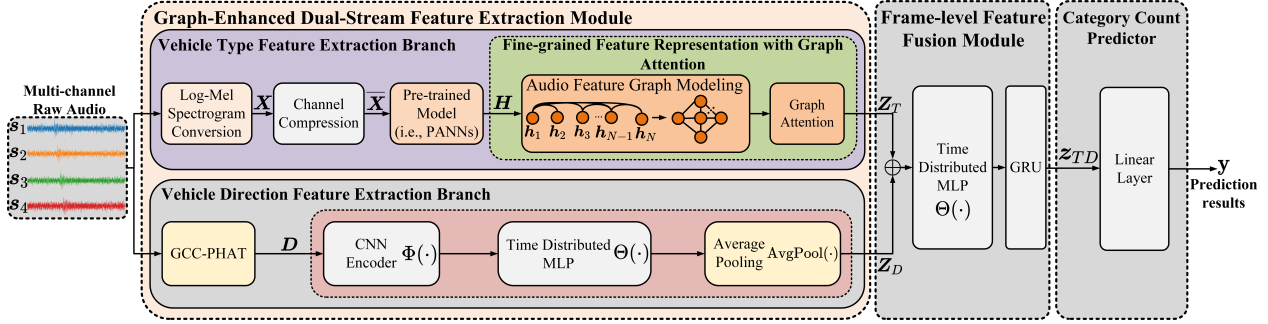


Fig. 1. Overall framework of the proposed method. The proposed GEDF-Net includes a graph-enhanced dual-stream feature extraction (GEDF) module with a vehicle type feature extraction (VTFE) branch and a vehicle direction feature extraction (VDFE) branch, a module for fusing type and direction features over time frames, and a category count predictor for vehicle prediction.

(GDFE) module with a vehicle type feature extraction (VTFE) branch and a vehicle direction feature extraction (VDFE) branch for vehicle type feature and direction feature extraction, respectively, together with a frame-level fusion module to fuse the fine-grained features for enhanced vehicle detection.

Specifically, in the VTFE branch, we use a pre-trained model (i.e. PANNs [22], pretrained on AudioSet [23]) to improve feature representation. In addition, inspired by [24], we incorporate a graph attention mechanism [25] to capture temporal relationships between audio events, treating feature frames from PANNs as nodes and their relationships as edges, thereby enhancing vehicle type feature representation. The VDFE branch extracts direction features using GCC-PHAT for time delay estimation. Then, these features are integrated along the corresponding time dimension and fused at the frame level to obtain fine-grained representation for traffic monitoring. Finally, a category count predictor counts vehicles by type and direction (e.g., cars or commercial vehicles moving left-to-right or right-to-left). Experiments are conducted on the DCASE 2024 Challenge Task 10 dataset [10] to demonstrate the effectiveness of the proposed method, which achieved 1st place in DCASE Challenge.

II. PROPOSED METHOD

Our proposed GEDF-Net, shown in Figure 1, contains a dual-stream feature extraction module including two branches, namely, the VTFE branch for extracting vehicle type features and the VDFE branch for extracting direction features, together with a module to combine these features in frame-level for fine-grained representation and a category count predictor to estimate the number of vehicles in each category.

A. Graph-Enhanced Dual-Stream Feature Extraction Module

The GEDF module is used to extract the type and direction features, both from the four-channel audio, detailed as follows.

1) *Vehicle Type Feature Extraction Branch*: Inspired by [24], the VTFE branch enhances the vehicle type feature representation by using a pre-trained model to address data scarcity and a graph attention mechanism to capture temporal relationships and emphasize important audio events for finer representation.

Feature Enhancement with Pre-trained Model: We use a pre-trained model (i.e., PANNs [22]) to extract vehicle

type features. Since PANNs is trained on AudioSet [23], which includes vehicle data, it enhances vehicle type feature representation and helps address data scarcity by incorporating external knowledge.

To achieve this, we first convert the input four-channel raw audio signals $\mathcal{S} = \{s_1, s_2, s_3, s_4\} \in \mathbb{R}^{4 \times L}$ to log-Mel spectrogram $\mathbf{X} \in \mathbb{R}^{4 \times B \times T}$ via a log-Mel spectrogram conversion operation, where L , B , and T represent the number of sampling points, mel bins, and time frames, respectively.

Then, a convolution layer, i.e., $\text{Conv}(\cdot)$, is applied to compress the four-channel log-Mel spectrogram \mathbf{X} into 1-D expression $\overline{\mathbf{X}} \in \mathbb{R}^{1 \times B \times T}$, as follows,

$$\overline{\mathbf{X}} = \text{Conv}(\mathbf{X}), \quad (1)$$

After this, a pre-trained model (i.e., PANNs [22]) is utilized to enhance the feature representation, mitigating the data scarcity for vehicle type feature extraction, as follows,

$$\mathbf{H} = \text{PANNs}(\overline{\mathbf{X}}), \quad (2)$$

where $\mathbf{H} = \{\mathbf{h}_1, \dots, \mathbf{h}_n, \dots, \mathbf{h}_N\} \in \mathbb{R}^{K \times N}$ denotes the enhanced feature, and $\mathbf{h}_n \in \mathbb{R}^{K \times 1}$ is the n -th feature frame. K and N are the dimension of the feature at each frame and the number of time frames, respectively.

Fine-grained Feature Representation with Graph Attention: Since vehicle audio events usually span multiple time frames, and feature frames of the same vehicle type should have higher correlations. To capture these contextual associations, we use audio feature graph modeling with an attention mechanism to emphasize key audio events related to vehicle traveling, achieving finer vehicle type feature representation.

Let \mathbf{h}_i and \mathbf{h}_j be the feature frames (nodes) in \mathbf{H} . The correlation between these frames is represented by the weight a_{ij} (attention coefficient) of the edge between \mathbf{h}_i and \mathbf{h}_j , calculated by a learnable linear mapping, following [24] [25],

$$a_{ij} = \text{Softmax}(\text{LeakyReLU}(\mathbf{e}^\top [\mathbf{M}\mathbf{h}_i \parallel \mathbf{M}\mathbf{h}_j])), \quad (3)$$

where $\mathbf{M} \in \mathbb{R}^{K \times K}$ is the learnable linear mapping, and $\mathbf{e} \in \mathbb{R}^{2K \times 1}$ is the learnable attention vector. \parallel denotes concatenation operation.

We can then obtain an adjacency graph $\mathcal{A} \in \mathbb{R}^{N \times N}$ from \mathbf{H} , with its element a_{ij} at i -th row and j -th column to represent the relation between feature nodes \mathbf{h}_i and \mathbf{h}_j . Then,

by aggregating the feature nodes of \mathcal{A} , we can obtain the improved type feature representation $\mathbf{Z}_T \in \mathbb{R}^{K \times N}$ as follows,

$$\mathbf{Z}_T = \mathcal{A} \mathbf{H} \mathbf{M}^\top + \mathbf{H}, \quad (4)$$

2) *Vehicle Direction Feature Extraction Branch*: We also adopt GCC-PHAT for direction feature extraction in the VDEF branch following [3]. In addition, an average pooling operation is introduced to further explore important directional information, facilitating the fusion of the vehicle type feature and direction feature over time dimension.

Specifically, the short-time Fourier transform (STFT) is employed to obtain the phase $\mathbf{P}_c \in \mathbb{R}^{F \times T}$ and $\mathbf{P}_k \in \mathbb{R}^{F \times T}$ of audio signals for each channel pair c and k , where $\{(c, k) \mid c, k \in \{1, 2, 3, 4\}, \text{ and } c \neq k\}$. F denotes the number of frequency bins. Then, the time delay $\mathbf{D}_{c,k} \in \mathbb{R}^{Q \times T}$ of the audio signals for each channel pair can be calculated as follows,

$$\mathbf{D}_{c,k} = \mathcal{F}^{-1}(\exp(j \cdot \text{Angle}(\mathbf{P}_c \odot \mathbf{P}_k^*))), \quad (5)$$

where $*$ denotes the conjugate operation, \odot represents element-wise multiplication, and $\text{Angle}(\cdot)$ computes the phase angle. $\mathcal{F}^{-1}(\cdot)$ stands for the inverse Fourier transform. Here, Q is the number of GCC-PHAT coefficients calculated from the two signals. Thus, the time delay estimation for all the four-channel audio signals can be represented as $\mathbf{D} = \{\mathbf{D}_{c,k} \mid c, k \in \{1, 2, 3, 4\}, \text{ and } c \neq k\}$.

Finally, the direction feature representation $\mathbf{Z}_D \in \mathbb{R}^{K \times N}$ can be obtained via a convolutional encoder (i.e., $\Phi(\cdot)$) and an MLP (i.e., $\Theta(\cdot)$) with the time delay estimation as,

$$\mathbf{Z}_D = \text{AvgPool}(\Phi(\Theta(\mathbf{D}))), \quad (6)$$

where $\text{AvgPool}(\cdot)$ denotes the average pooling operation.

B. Frame-level Feature Fusion Module

After extracting the vehicle type and direction features, we use a module to combine these features over each time frame to obtain a fine-grained representation accounting for both vehicle type and direction as follows,

$$\mathbf{z}_{TD} = \text{GRU}(\Theta([\mathbf{Z}_T \parallel \mathbf{Z}_D])), \quad (7)$$

where $\text{GRU}(\cdot)$ denotes the gated recurrent unit (GRU) [26], and $\mathbf{z}_{TD} \in \mathbb{R}^{1 \times d}$ is the last time step result of GRU for regressing the vehicle count, and d denotes the dimension of the result for each time frame.

C. Category Count Predictor

A linear layer is utilized as the category count predictor to estimate the counts for each category of vehicle event across vehicle type (i.e., car and commercial vehicle) and travel direction (i.e., left to right and right to left), as follows,

$$\mathbf{y} = \text{ReLU}(\mathbf{W} \mathbf{z}_{TD} + \mathbf{b}), \quad (8)$$

where $\mathbf{y} \in \mathbb{R}^{1 \times 4}$ is the prediction result for acoustic traffic monitoring, and \mathbf{W} and \mathbf{b} are the weight matrix and bias of the predictor, respectively.

TABLE I
METADATA SUMMARY FOR EACH LOCATION.

location	loc1	loc2	loc3	loc4	loc5	loc6
number of audio samples	1256	56	4129	17	96	1740
max-pass-by-speed	100	50	50	50	40	90
max-traffic-density (per minute)	1000	900	500	400	140	900
number of vehicle audio events	11515	901	21685	63	165	19116

III. EXPERIMENTS AND RESULTS

A. Experimental Setup

1) *Dataset*: Since the official evaluation set is not available, we evaluated performance on the DCASE 2024 Challenge Task 10 development dataset [10], which includes four audio event types (“car_left”, “car_right”, “cv_left”, “cv_right”) from six locations (loc1 to loc6) and synthetic data generated by an acoustic traffic simulator [3], [9]. Real data were recorded using a linear microphone array parallel to traffic flow. Sample counts per location are in Table I. Similar to the baseline method, synthetic data are also included in our training process.

2) *Evaluation Metrics*: Following the official baseline [3], we use Kendall’s Tau Rank Correlation (Kendall’s Tau Corr) and Root Mean Square Error (RMSE) as evaluation metrics. Kendall’s Tau Corr measures the ordinal association between predictions and actual results, while RMSE quantifies prediction errors. We also use a third metric, Ranking Score, as defined in the official Task 10 evaluation.

The Ranking Score evaluates systems based on average rankings across $6 \times 4 \times 2$ comparisons: 6 locations (loc1 to loc6), 4 audio event types (“car_left”, “car_right”, “cv_left”, “cv_right”), and 2 metrics (Kendall’s Tau Corr and RMSE). Each comparison is assigned a ranking score, with Rank 1 being the best and Rank N the worst. The final performance is the average of all rankings, where a lower score indicates better performance.

B. Effectiveness Validation and Analysis

We conduct ablation experiments on the development set to validate the effectiveness of our proposed GEDF-Net. We compare the full method with two variants: GEDF-Net without PANNs (w/o -P) and without graph attention (w/o -G), along with the baseline [3], [10]. Results are in Table II.

All our methods outperform the baseline, with GEDF-Net achieving the best ranking score, demonstrating its effectiveness in using pre-trained models for feature enhancement and graph attention for detailed feature representation. GEDF-Net significantly outperforms GEDF-Net (w/o -P) in locations with very few samples (e.g., loc2 with 56 samples and loc4 with 17 samples), highlighting the value of external knowledge in mitigating data scarcity.

In loc5, the traffic scenario is relatively simple, as characterized by low maximum traffic density (140 vehicles per minute) and low maximum pass-by speed (40 km/h). With a limited number of samples, GEDF-Net (w/o -P) still achieves comparable results to GEDF-Net.

Meanwhile, GEDF-Net outperforms GEDF-Net (w/o -G), showing the benefits of using graph attention modeling to

TABLE II

EFFECTIVENESS VALIDATION OF OUR PROPOSED GEDF-NET ON DCASE 2024 CHALLENGE TASK 10 DEVELOPMENT DATASET, WHERE KENDALL'S TAU RANK CORR (KENDALL), RMSE AND RANKING SCORE ARE USED FOR EVALUATION. NOTE THAT, THE RANKING SCORE REFLECTS RESULTS THAT ARE ONLY BASED ON COMPARISONS AMONG THE METHODS LISTED IN THIS TABLE.

Methods	Category	loc1		loc2		loc3		loc4		loc5		loc6		Ranking Score
		Kendall	RMSE	Kendall	RMSE	Kendall	RMSE	Kendall	RMSE	Kendall	RMSE	Kendall	RMSE	
Baseline	car_left	0.445	2.555	0.579	3.074	0.543	1.731	0.195	1.997	0.575	0.693	0.804	1.628	2.71
	car_right	0.423	2.978	0.337	2.917	0.569	1.294	0.038	1.674	0.371	0.693	0.700	1.822	
	cv_left	0.084	0.918	0.044	0.813	0.034	0.309	0.000	0.655	0.068	0.362	0.763	0.509	
	cv_right	0.076	0.882	0.051	0.604	0.322	0.212	0.000	0.463	0.257	0.252	0.641	0.530	
GEDF-Net (w/o -P)	car_left	0.410	2.654	0.630	2.510	0.555	1.716	0.049	2.295	0.582	0.629	0.805	1.646	2.52
	car_right	0.440	2.920	0.516	2.239	0.572	1.281	-0.063	2.882	0.394	0.679	0.704	1.810	
	cv_left	0.176	0.909	-0.034	0.850	0.174	0.307	0.000	0.655	0.045	0.351	0.690	0.601	
	cv_right	0.117	0.937	-0.051	0.717	0.299	0.212	0.000	0.463	0.361	0.238	0.521	0.648	
GEDF-Net (w/o -G)	car_left	0.393	2.765	0.700	2.157	0.548	1.724	0.634	1.174	0.525	0.745	0.811	1.549	2.58
	car_right	0.441	2.970	0.474	2.544	0.575	1.289	0.341	0.958	0.397	0.694	0.694	1.837	
	cv_left	0.172	0.962	0.143	0.798	0.009	0.314	0.296	0.607	-0.047	0.361	0.743	0.582	
	cv_right	0.126	0.954	0.058	0.726	-0.002	0.222	0.120	0.613	-0.083	0.300	0.611	0.576	
GEDF-Net	car_left	0.434	2.600	0.719	2.177	0.551	1.729	0.097	2.095	0.557	0.708	0.816	1.582	2.04
	car_right	0.448	2.919	0.401	2.666	0.577	1.275	0.240	1.548	0.401	0.697	0.684	1.910	
	cv_left	0.207	0.892	0.226	0.783	0.171	0.315	0.182	0.604	0.058	0.362	0.683	0.604	
	cv_right	0.126	0.861	0.171	0.677	0.377	0.195	0.445	0.428	0.357	0.208	0.570	0.594	

TABLE III

PERFORMANCE COMPARISON WITH TOP-RANKING SYSTEMS ON THE DCASE 2024 CHALLENGE TASK 10 EVALUATION SET.

Methods	Official Ranking	Ranking Score
GEDF-Net (Ours) [17]	1	3.98
Bai_JLESS_task10_1 [12]	2	4.44
Takahashi_TMU-NEE_task10_1 [13]	3	4.77
Baseline_Bosch_task10 [10]	-	5.17
Park_KT_task10_3 [14]	4	5.67
Betton-Ployon_ACSTB_task10_1 [15]	5	7.89
Cai_NCUT_task10_1 [16]	6	8.14

*Official ranking results of DCASE 2024 Challenge Task 10.

capture temporal relationships and highlight important audio events in vehicle type features. However, at loc4, GEDF-Net (w/o -G) performs better, likely due to the limited number of audio events (63) at this location, which may limit the graph attention model to learn robust representations.

C. Performance Comparison with the Top Ranking Systems

The evaluation set in DCASE 2024 Challenge Task 10 was not released to the public, and most top-ranking systems were not open-sourced either. For this reason, we could not reproduce their results for comparison. Instead, we present the official evaluation results published by DCASE Challenge organisers to showcase the advantages of our method. The official ranking scores of the top systems are shown in Table III, with detailed results available on the competition website¹.

From Table III, it can be seen that our proposed GEDF-Net as the submission system outperforms all other systems, which shows the superiority of our proposed method, demonstrating the effectiveness of using the pre-trained model to mitigate the data scarcity and the graph attention to exploit temporal relationships and highlight important audio events for acoustic traffic monitoring. Moreover, our method surpasses the Transformer-based Top-2 system (i.e., Bai_JLESS_task10_1 [12]) that simply uses log-Mel spectrogram as type feature, further illustrating the effectiveness of graph-enhanced fine-grained feature representation with the pre-trained model.

¹<https://dcase.community/challenge2024/task-acoustic-based-traffic-monitoring-results>




Fig. 2. Visualization of graph-enhanced vehicle type feature representation. The top row shows the log-Mel spectrograms of two audio samples, while the bottom row shows the learned corresponding linear interpolation adjacency graphs, with red boxes denoting the attention highlighted vehicle travel events.

D. Visualization Analysis

To demonstrate that our graph-enhanced fine-grained feature representation can capture the contextual association and highlight the important audio events related to vehicle traveling, we visualize the linear interpolation adjacency graphs of the learned vehicle type feature in Figure 2, where we can see the vehicle travel events (i.e., feature nodes) are highlighted in the interpolation adjacency graphs as indicated in the red box areas. The results further validate the effectiveness of our proposed method.

IV. CONCLUSION

In this paper, we have presented a graph-enhanced dual-stream feature fusion network with a pre-trained model for acoustic traffic monitoring, where both vehicle type and direction are taken into account for feature representation. Specifically, a pre-trained model is introduced to mitigate the data scarcity for feature enhancement and graph attention is leveraged for finer type feature representation. Experimental results demonstrate the effectiveness of our proposed method. By fusing fine-grained vehicle type feature and direction feature, our method achieved 1st place in DCASE 2024 Challenge Task 10.

REFERENCES

- [1] Slobodan Djukanović, Jiří Matas, and Tuomas Virtanen, “Robust audio-based vehicle counting in low-to-moderate traffic flow,” in *Proc. of IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2020, pp. 1608–1614.
- [2] Slobodan Djukanović, Yash Patel, Jiří Matas, and Tuomas Virtanen, “Neural network-based acoustic vehicle counting,” in *Proc. of European Signal Processing Conference (EUSIPCO)*. IEEE, 2021, pp. 561–565.
- [3] Stefano Damiano, Luca Bondi, Shabnam Ghaffarzadegan, Andre Gun-toro, and Toon van Waterschoot, “Can synthetic data boost the training of deep acoustic vehicle counting networks?,” in *Proc. of International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, April 2024.
- [4] Nikola Bulatović and Slobodan Djukanović, “Mel-spectrogram features for acoustic vehicle detection and speed estimation,” in *Proc. of International Conference on Information Technology (IT)*. IEEE, 2022, pp. 1–4.
- [5] Shigemi Ishida, Jumpei Kajimura, Masato Uchino, Shigeaki Tagashira, and Akira Fukuda, “SAVeD: Acoustic vehicle detector with speed estimation capable of sequential vehicle detection,” in *Proc. of International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2018, pp. 906–912.
- [6] A. Severdaks and M. Liepins, “Vehicle counting and motion direction detection using microphone array,” *Elektronika ir Elektrotechnika*, vol. 19, no. 8, pp. 89–92, 2013.
- [7] Grzegorz Szwoch and Józef Kotas, “Acoustic detector of road vehicles based on sound intensity,” *Sensors*, vol. 21, no. 23, pp. 7781, 2021.
- [8] Xingshui Zu, Shaojie Zhang, Feng Guo, Qin Zhao, Xin Zhang, Xing You, Huawei Liu, Baoqing Li, and Xiaobing Yuan, “Vehicle counting and moving direction identification based on small-aperture microphone array,” *Sensors*, vol. 17, no. 5, pp. 1089, 2017.
- [9] Stefano Damiano and Toon van Waterschoot, “Pyroadacoustics: A road acoustics simulator based on variable length delay lines,” in *Proc. of International Conference on Digital Audio Effects (DAFx)*, September 2022, pp. 216–223.
- [10] Shabnam Ghaffarzadegan, Luca Bondi, Wei-Cheng Lin, Abinaya Kumar, Ho-Hsiang Wu, Hans-Georg Horst, and Samarjit Das, “Sound of Traffic: A dataset for acoustic traffic identification and counting,” Tech. Rep., DCASE 2024 Challenge, June 2024.
- [11] Byoung-ho Kwon, Youngjin Park, and Youn-sik Park, “Analysis of the GCC-PHAT technique for multiple sources,” in *Proc. of International Conference on Control, Automation and Systems (ICCAS)*, 2010, pp. 2070–2073.
- [12] Dongzhe Zhang, Jisheng Bai, and Jianfeng Chen, “JLESS submission to DCASE 2024 Task10: An acoustic-based traffic monitoring solution,” Tech. Rep., DCASE 2024 Challenge, June 2024.
- [13] Tomohiro Takahashi, Natsuki Ueno, Yuma Kinoshita, Yukoh Wakabayashi, Nobutaka Ono, Makiho Sukekawa, Seishi Fukuma, and Hiroshi Nakagawa, “Neural network training with matching loss for ranking function,” Tech. Rep., DCASE 2024 Challenge, June 2024.
- [14] Yeonseok Park, TaeWoon Yeo, and Baeksan On, “Deep acoustic vehicle counting model with short-time homomorphic deconvolution,” Tech. Rep., DCASE 2024 Challenge, June 2024.
- [15] Erwann Betton-Ployon, Abbes Kacem, and Jérôme Mars, “Traffic counting system leveraged with a non-supervised counting approach,” Tech. Rep., DCASE 2024 Challenge, June 2024.
- [16] Zhilong Jiang, Xichang Cai, Ziyi Liu, and Menglong Wu, “DCASE 2024 Challenge Task10 technical report,” Tech. Rep., DCASE 2024 Challenge, June 2024.
- [17] Shitong Fan, Feiyang Xiao, Shuhan Qi, Qiaoxi Zhu, Wenwu Wang, and Jian Guan, “Fine-grained audio feature representation with pretrained model and graph attention for traffic flow monitoring,” Tech. Rep., DCASE 2024 Challenge, June 2024.
- [18] David P. Helmbold, Jyrki Kivinen, and Manfred K. Warmuth, “Relative loss bounds for single neurons,” *IEEE Transactions on Neural Networks*, vol. 10, no. 6, pp. 1291–1304, 1999.
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” in *Proc. of Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [20] Karen Simonyan and Andrew Zisserman, “Very Deep Convolutional Networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [21] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” in *Proc. of Neural Information Processing Systems (NIPS)*, 2017, vol. 30, pp. 5998–6008.
- [22] Qiuqiang Kong, Yin Cao, Turab Iqbal, Yuxuan Wang, Wenwu Wang, and Mark D. Plumbley, “PANNs: Large-scale pretrained audio neural networks for audio pattern recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2880–2894, 2020.
- [23] Jort F. Gemmeke, Daniel P. W. Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R. Channing Moore, Manoj Plakal, and Marvin Ritter, “AudioSet: An ontology and human-labeled dataset for audio events,” in *Proc. of International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 776–780.
- [24] Feiyang Xiao, Jian Guan, Qiaoxi Zhu, and Wenwu Wang, “Graph attention for automated audio captioning,” *IEEE Signal Processing Letters*, vol. 30, pp. 413–417, 2023.
- [25] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio, “Graph attention networks,” in *Proc. of International Conference on Learning Representations (ICLR)*, 2018.
- [26] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio, “Learning phrase representations using RNN encoder–decoder for statistical machine translation,” in *Proc. of Empirical Methods in Natural Language Processing (EMNLP)*, Oct. 2014, pp. 1724–1734.