

MULTI-LEVEL GRAPH LEARNING FOR AUDIO EVENT CLASSIFICATION AND HUMAN-PERCEIVED ANNOYANCE RATING PREDICTION

Yuanbo Hou¹, Qiaoqiao Ren¹, Siyang Song², Yuxin Song³, Wenwu Wang⁴, Dick Botteldooren¹

¹Ghent University, Belgium ²University of Leicester, UK ³Baidu Inc., China ⁴University of Surrey, UK

ABSTRACT

WHO’s report on environmental noise estimates that 22 M people suffer from chronic annoyance related to noise caused by audio events (AEs) from various sources. Annoyance may lead to health issues and adverse effects on metabolic and cognitive systems. In cities, monitoring noise levels does not provide insights into noticeable AEs, let alone their relations to annoyance. To create annoyance-related monitoring, this paper proposes a graph-based model to identify AEs in a soundscape, and explore relations between diverse AEs and human-perceived annoyance rating (AR). Specifically, this paper proposes a lightweight multi-level graph learning (MLGL) based on local and global semantic graphs to simultaneously perform audio event classification (AEC) and human annoyance rating prediction (ARP). Experiments show that: 1) MLGL with 4.1 M parameters improves AEC and ARP results by using semantic node information in local and global context-aware graphs; 2) MLGL captures relations between coarse- and fine-grained AEs and AR well; 3) Statistical analysis of MLGL results shows that some AEs from different sources significantly correlate with AR, which is consistent with previous research on human perception of these sound sources.

Index Terms— Noise, annoyance, audio event classification, annoyance rating prediction, multi-level graph learning

1. INTRODUCTION

World Health Organization (WHO) reports that environmental noise is the second cause of ill health in Europe after air pollution [1]. WHO hereby considers that severe annoyance can lead to cardiovascular disease [2] and reduced quality of life, with annoyance being an important pathway between noise and long-term health effects. Noticed audio events (AEs) may play a vital role in the emergence of annoyance [3]. Although most epidemiological studies guided by [4] consider a single type of sound source, different sounds occur together in real life, and a person’s annoyance rating (AR) will be affected by the combined effect of sound sources. Thus, this paper links audio event classification (AEC) [5] and annoyance rating prediction (ARP) [6]. AEC aims to determine whether there are predefined AEs in the soundscape, while ARP aims to predict the overall human-perceived AR. We are aware of the large gap between human AR in different environments, but we see our efforts as a

first step towards deploying long-term monitoring of adverse effects of environmental sound.

Most deep learning-based AEC-related studies focus on AE detection and classification, such as works in the detection and classification of acoustic scenes and events (DCASE) challenges [7]. In AEC tasks, mel spectrograms are fed into neural networks, such as CNN [8] and CRNN [9], to extract acoustic representations to classify audio clips. Recently, with the aid of large-scale datasets, e.g. AudioSet [10], and pre-trained audio models (such as PANNs [8] and AST [11]), deep learning-based models have achieved widespread success in AEC tasks. However, most AEC studies mentioned above focus only on the accuracy of AEs recognition, but have not explored the perceptual experiences that these AEs and their combinations bring to people. One of the reasons for this might be that challenges, like DCASE, do not address this aspect. For applications like monitoring for health effects and urban planning, the perception and understanding of acoustic environments become important, and AE recognition plays a vital role [12][3]. Emotions elicited directly by the sound or by the meaning given to it (the recognition) have multiple dimensions, but for the assessment of health effects, recent research mostly focuses on annoyance [13]. To simplify their approach, researchers often only use coarse-grained audio events (cAE) labels to roughly distinguish between human, natural, and mechanical sounds. However, humans use more detailed and specific fine-grained audio events (fAE) information when understanding the sound environment [3].

To extend AEC focusing on AEs to ARP on human perceptions, a hierarchical graph representation learning (HGRL) [6] has been developed recently on the DeLTA [14] dataset, the only publicly available dataset that contains both AE labels and AR. However, HGRL has shortcomings: 1) HGRL has no corresponding semantic supervision when extracting node representations to build the graph. This leads to difficulties in specifying each node’s semantics. 2) HGRL focuses on coarse-grained global information while ignoring fine-grained local information, as it uses all nodes to build a global graph. 3) The HGRL model has 92.3M parameters (*params*), which may be difficult to meet requirements in some practical applications with limited resources, e.g. edge devices. To this end, this paper proposes a lightweight attention-fused multi-level graph learning (MLGL) model

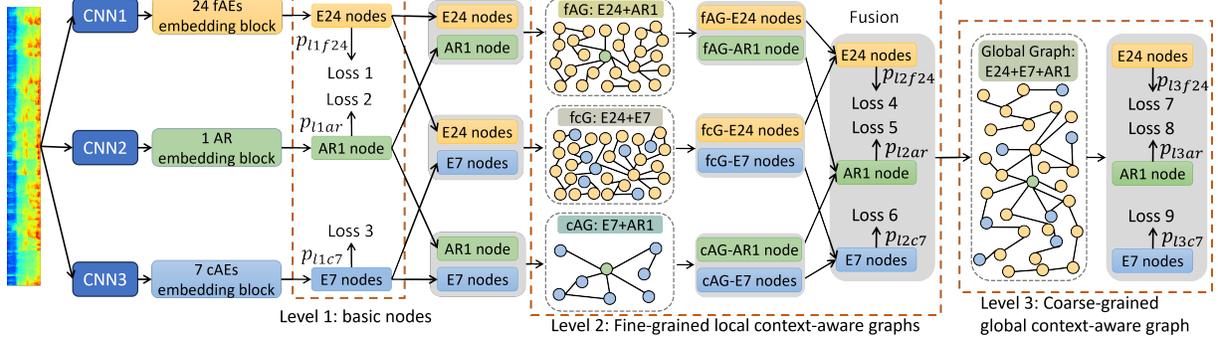


Fig. 1: The proposed lightweight attention-fused-grained multi-level graph learning (MLGL).

with local and global semantic graphs based on nodes with explicit semantic information, which facilitates the modelling of a global graph using multiple local context-aware graphs.

The contributions are as follows: 1) Compared to HGRL, the MLGL provides higher explainability for its reliance on local and global context-aware graphs. 2) The MLGL, based on graph neural networks that capture the relations between nodes well, outperforms CNN-based models. Visual analyses show that MLGL better captures relations between multi-grained AEs, and relations between AEs and AR. 3) Statistical analysis of the MLGL results shows that AEs from some sources significantly correlate with AR, which is consistent with human perception of these environmental sound sources.

2. LOCAL AND GLOBAL CONTEXT-AWARE MLGL

In this section, we first present how to extract node representations with explicit semantic information, then use semantic nodes to build local context-aware graphs (LcGs). Next, we fuse nodes from multiple LcGs by attention mechanism to enhance common objective representations of the same node in different local contexts. Finally, we construct the global context-aware graph (GcG) with attention-enhanced nodes, which captures nodes' representation and relations from the global view. We illustrate the idea with the help of labels of 24 types of fine-grained AEs (fAEs), 7 types of coarse-grained AEs (cAEs), and human-annotated AR in DeLTA [14].

2.1. Semantic node representation

In Fig. 1, MLGL uses three 3-layer VGG-like CNNs [15] to extract representations for sound-related fAEs, cAEs, and human emotion-related AR, respectively. This enables MLGL to learn semantic representations of each target, which is different from HGRL [6] where a shared 6-layer CNN is used to extract the mixed representations of 3 targets (fAEs, cAEs, and AR). Then, the extracted representations are fed into the corresponding embedding block, each consisting of multiple parallel 64-dimensional embedding (64D E_{mb}) layers, to obtain the semantic representations for each target. Unlike HGRL, which has no semantic supervision for nodes in extracting their representations, MLGL ensures that the E_{mb} layer only learns the semantic representation of its

target, which is followed by a 1-unit linear layer and a loss function to convert its output to the target prediction. Given that the labels of fAE, cAE, and AR are y_{f24} , y_{c7} , and y_{ar} , respectively, the losses at level 1 are $L_1 = BCE(p_{11f24}, y_{f24})$, $L_2 = MSE(p_{11ar}, y_{ar})$, $L_3 = BCE(p_{11c7}, y_{c7})$, where BCE is the binary cross entropy loss, MSE is the mean squared error loss, p_{11f24} , p_{11ar} , and p_{11c7} are predictions of fAE, cAE, and AR at level 1, respectively. The 64D E_{mb} of each target is viewed as the semantic representation of each node, which is then used to construct different LcGs at level 2.

2.2. Local context-aware graphs (LcGs)

Based on the obtained node representations, we construct 3 LcGs with different motives: **1) fAG:** fAEs-AR graph; **2) fcG:** fAEs-cAEs graph; **3) cAG:** cAEs-AR graph. Among them, the pure AE-based graph fcG aims to model the relations between the 24 classes of fAEs and higher-level cAEs, and use the contexts of AEs in different granularity to improve the node representations. In contrast, fAG and cAG explore the relations and complementarity between AEs and AR. Nodes in each LcG are fully connected and modelled using a 3-layer gated graph convolutional network (GCN) [16].

Attention-based node fusion. The fAG outputs 24 AR-aware fAEs (E24) node representations (marked as fAG-E24), while fcG outputs 24 cAE-aware fAEs node representations (marked as fcG-E24). Both fAG-E24 and fcG-E24 describe the same target E24 enriched by related context information in different perspectives. Hence, MLGL fuses the common information between these nodes by attention mechanism [17].

$Attention(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = softmax(\mathbf{Q}\mathbf{K}^T / \sqrt{d_k})\mathbf{V}$ (1) where $\mathbf{V}=\mathbf{K}$, and d_k is \mathbf{K} 's dimension. Here, \mathbf{Q} acts as an index to adjust \mathbf{V} . The attention output is mainly based on \mathbf{V} , so a more informative \mathbf{V} will lead to better attention results. For the E24 nodes, fAG-E24 learns E24 with 1 AR node, but fcG-E24 learns E24 with 7 cAEs nodes, so fcG-E24 will contain more context information. For fusing E24 nodes, \mathbf{Q} is fAG-E24, \mathbf{K} is fcG-E24, that is, using the AR-aware information to adjust the pure-AE information. For fusing the E7 nodes, \mathbf{Q} is cAG-E7, and \mathbf{K} is fcG-E7. For fusing the AR1 node, \mathbf{Q} is cAG-AR1, and \mathbf{K} is fAG-AR1. Similar to Section 2.1, fused nodes are fed into the following

1-unit linear layer, and a loss is used to enhance the learning of the corresponding semantic information. Losses at level 2 are $L_4=BCE(p_{l2f24}, y_{f24})$, $L_5=MSE(p_{l2ar}, y_{ar})$, and $L_6=BCE(p_{l2c7}, y_{c7})$, where p_{l2f24} , p_{l2ar} , and p_{l2c7} are the predictions of fAE, cAE, and AR at level 2, respectively.

2.3. Global context-aware graph (GcG)

The GcG comprehensively considers the semantic information of fAEs, cAEs, and AR derived from multiple LcGs. The GcG is also modelled by a 3-layer GCN to represent its nodes and global relations between nodes. Global context-aware node representations are fed into the following 1-unit linear layer, and a loss is used to complete AEC and ARP tasks. Losses at level 3 are $L_7=BCE(p_{l3f24}, y_{f24})$, $L_8=MSE(p_{l3ar}, y_{ar})$, and $L_9=BCE(p_{l3c7}, y_{c7})$, where p_{l3f24} , p_{l3ar} , and p_{l3c7} are predictions of fAE, cAE, and AR at level 3, respectively. To perform end-to-end training for MLGL, the final loss is: $L=\sum_{i=1}^9 \lambda_i L_i$, where λ_i is the parameter for weighting losses during training, fixed to 1 in our experiments. The losses at levels 1 and 2 enable the model to maintain semantic-related information while learning intermediate representations. However, the losses at level 3 correct the predictions of MLGL from a global perspective.

3. EXPERIMENTS AND RESULTS

3.1. Dataset, experiments setup, and metrics

We use the public dataset DeLTA [14] for experiments, which has 2890 15-second binaural audio clips. Each clip has labels of fAEs and the corresponding AR (continuously from 1 to 10). Following [6], the labels of 7 classes cAEs are derived from the labels of 24 classes fAEs. The training, validation, and test sets contain 2200, 245, and 445 clips, respectively.

The log-mel energy with 64 banks [8] is used as acoustic features, with a Hamming window of 46 ms with 1/3 overlap between neighbouring windows. Dropout and normalization are used to prevent over-fitting of models [18]. A batch size of 64 and AdamW optimizer [19] with a learning rate of 0.0005 are used to minimize the loss. Models are trained on a Tesla V100 GPU for 400 epochs. Dataset, code, and models are available on the *webpage* (<https://github.com/Yuanbo2020/MLGL>). Accuracy (*Acc*), *F-score*, and threshold-free *AUC* [20] are used to evaluate the classification results of AEC. Mean absolute error (*MAE*), mean squared error (*MSE*), and R2-score (*R2*) [21] are used to measure the prediction results of ARP.

3.2. Results and analysis

Performance of different levels of prediction. The proposed MLGL has 9 losses at 3 levels, responsible for different roles. In Table 1, it can be observed that good predictions from all levels have been achieved for both the 24 classes of fAEs and the 7 classes of cAEs, implying that level 1 of MLGL performs well in capturing the semantic representations of base nodes. The attention-fusion of graph-enriched

representations, further improves its performance for human-perceived ARP. This implies that representations of related nodes are enhanced in level 2 graphs. Finally, the predictions at level 3 are improved in both AEC and ARP. This reveals that node representations optimized in LcGs are beneficial for constructing a more accurate GcG. Note that the DeLTA dataset does not have labels for coarse-grained AEC, and cAE labels from [6] without human validation cannot be evaluated, so later, by default, AEC will be the fine-grained AEC.

Table 1: Multi-level predictions of MLGL for AEC and ARP.

Output Level	Fine-grained AEC		Coarse-grained AEC		ARP		
	Acc. (%)	AUC	Acc. (%)	AUC	MSE	MAE	R2
1	91.91	0.918	84.72	0.910	1.062	0.767	0.452
2	91.90	0.920	84.98	0.910	0.964	0.717	0.502
3	91.96	0.921	84.94	0.912	0.940	0.706	0.515

Model parameters and size. To simulate a realistic use case, Table 2 shows the inference time from loading the raw audio clip, extracting features, feeding them into the model, to getting predictions. Compared with HGRL, *params* of MLGL are reduced by $(92.3-4.1)/92.3 \times 100\% \approx 96\%$, and the model size is reduced by about 95%, while MLGL still improves predictions of ARP and slightly improves AEC.

Table 2: Comparison of HGRL and MLGL in detail.

Model	Params (M)	Model Size (MB)	Inference time (s)	AEC		ARP	
				Acc. (%)	AUC	MAE	R2
HGRL	92.3	353.0	0.531	91.71	0.901	0.802	0.458
MLGL	4.1	16.0	0.448	91.96	0.921	0.706	0.515

Comparison of fusion methods. At level 2, different enriched representations of the same target node are fused by attention to enhance the common information between node representations learned from different local contexts. Specifically, two embedding vectors (denoted as EV_1 , EV_2) of the same node are fused. Table 3 explores alternative fusion methods to investigate their impact on the proposed model.

Table 3: Results of different fusion ways at level 2 of MLGL.

#	Fusion Type	AEC			ARP		
		Acc. (%)	F-score (%)	AUC	MSE	MAE	R2
1	Addition	91.67	67.96	0.897	1.247	0.846	0.356
2	Concat	91.68	68.67	0.903	1.148	0.818	0.407
3	Hadamard	91.73	68.58	0.908	1.282	0.872	0.339
4	Gating	91.75	68.11	0.907	1.141	0.813	0.411
5	Attention	91.96	68.36	0.921	0.940	0.706	0.515

In Table 3, #1 refers to the addition of two vectors, while #2 refers to the concatenation of two vectors, which is then fed to a linear layer. Therefore, in #2, an extra layer is added with respect to #1. #3 performs the Hadamard product of two vectors. Similar to gated linear units [22], #4 utilizes the gated output of EV_2 to condition EV_1 , $(W_1 \times EV_1 + b_1) \odot \sigma(W_2 \times EV_2 + b_1)$, where σ is a sigmoid function, \odot is element-wise product, W and b are learnable weight and bias, respectively. Overall, the AEC performance of these fusion methods is similar, while attention-based fusion improves the score on threshold-free AUC. For ARP, the attention-based fusion significantly improved the scores on MSE, MAE, and R2. Please see the *webpage* for source code and models.

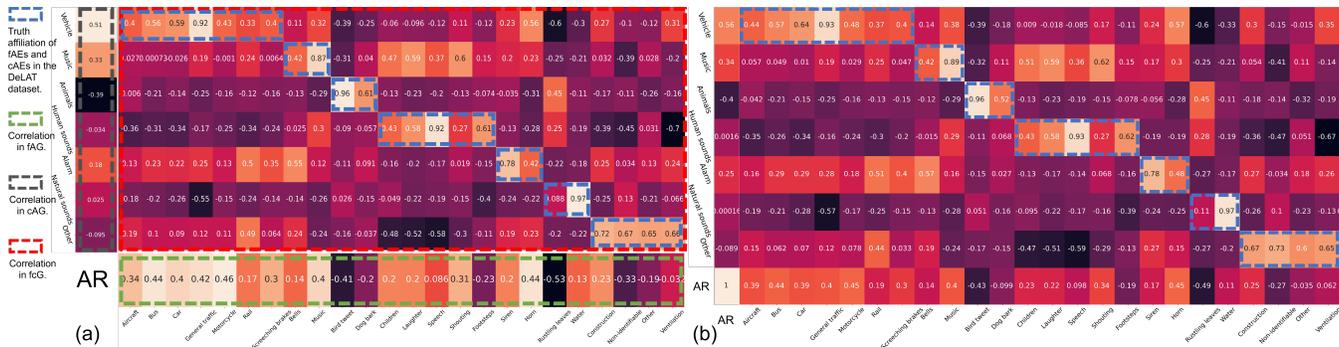


Fig. 2: (a) Correlation between node embeddings of multiple LcGs. (b) Correlation between node embeddings of the GcG.

Comparison to other methods. We compare the performance of models in Table 4, where the results of CNN, CNN-Transformer, PANNs, and HGRL are all taken from [6]. Compared with these models, MLGL, despite using much fewer parameters, achieves better results in both AEC and ARP, with MSE less than 1, and R2 greater than 0.5 for ARP.

Table 4: Comparison of different models on DeLTA dataset.

#	Model	Param. (M)	AEC		ARP		
			F-score(%)	AUC	MSE	MAE	R2
1	CNN	0.8	55.05	0.891	1.675	0.997	0.135
2	CNN-Trans.	1.6	58.66	0.851	1.445	0.966	0.254
3	PANNs	79.7	63.86	0.903	1.162	0.858	0.400
4	HGRL	92.3	67.91	0.901	1.049	0.802	0.458
5	MLGL	4.1	68.36	0.921	0.940	0.706	0.515

Explaining how the model works. To study the ability of LcGs to capture relations between nodes, and whether the GcG can model three kinds of semantic nodes in one graph, Fig. 2 visualizes the Pearson correlation coefficients [23] between node embeddings on all the test audio clips. In Fig. 2, the correlations between the nodes in fcG successfully match the associations of fAECs to cAECs in the DeLTA dataset. The correlations in the fAG indicate that *motorcycle*, *bus*, and *horn* are the most likely sounds to cause people annoyance, while *rustling leaves* and *bird tweet* sounds are the least likely to be annoying. The cAG box shows that *vehicles* sounds are more likely to annoy people, while *animals* sounds are not. These LcGs with separate motivations perform well in capturing the semantic relations between internal nodes.

In Fig. 2 (b), the GcG aligns all the semantic nodes using one graph. HGRL [6] has similar results, but the correlation of nodes in Fig. 2 (b) is significantly higher than that of HGRL. This means that MLGL can model different semantic nodes and their relations from a global view, resulting in better performance on AEC and ARP tasks, as compared with HGRL, because MLGL has nodes with explicit semantics and uses local context-aware graphs to optimize the node representations. For a clearer Fig. 2, please see the [webpage](#).

Statistical significance of correlations. For the results in Fig. 2, Table 5 further tests their significance. Due to limited space, Table 5 shows partial AEs. We first assess the distribution of MLGL’s predictions with a Shapiro–Wilk test. Since

the predictions are not normally distributed, Spearman’s rank (Spearman’s rho) [24] correlation coefficient (CC) is used. Note: ** indicates statistical significance at the 0.001 level.

Table 5: Spearman’s rho correlation coefficients for AEs with AR.

#	AE	CC w/ AR	#	AE	CC w/ AR	#	AE	CC w/ AR
1	Bus	0.504**	7	Speech	0.193**	13	General traffic	0.391**
2	Car	0.199**	8	Children	0.173**	14	Motorcycle	0.374**
3	Rail	0.329**	9	Shouting	0.443**	15	Screech brakes	0.488**
4	Bells	0.458**	10	Bird tweet	-0.522**	16	Rustling leaves	-0.585**
5	Horn	0.653**	11	Dog bark	0.029**	17	Ventilation	0.085
6	Water	-0.097**	12	Aircraft	0.415**	18	Other	0.011

In Table 5, several AEs, namely *Aircraft*, *Bus*, *Screeching brakes*, *Bells*, *Shouting*, and *Horn* exhibit strong positive correlations with AR, which indicates that an increase in the occurrences of these AEs increases the level of annoyance. In contrast, AEs like *Bird tweet* and *Rustling leaves* show strong negative correlations with AR, implying a decrease in annoyance when these are present. These statistical results from the proposed MLGL are consistent with earlier human-perception-based soundscape research [3][12].

4. CONCLUSION

We have presented the MLGL to identify audio events (AEs) generated by diverse environmental sound sources and predict human-perceived annoyance rating (AR) in real-life soundscapes. Experimental results show that: 1) Lightweight MLGL with only 4.1 M parameters performs AEC and ARP tasks as well as previous heavier models; 2) Local context-aware graphs with different motivations perform well in capturing semantic nodes, which helps the global-view graph to model the nodes and their relations; 3) Statistical analysis of MLGL results shows that some AEs are closely related to human AR, which is consistent with observations from previous soundscape research based on human perception. Future work will deploy MLGL on edge devices and measure differences in real applications to assist soundscape research.

5. ACKNOWLEDGEMENTS

This research received funding from the Flemish Government under the “Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen” programme.

6. REFERENCES

- [1] World Health Organization, *Burden of disease from environmental noise: Quantification of healthy life years lost in Europe*, WHO Regional Office for Europe, 2011.
- [2] T. Münzel, F. P. Schmidt, S. Steven, J. Herzog, A. Daiber, and M. Sørensen, “Environmental noise and the cardiovascular system,” *Journal of the American College of Cardiology*, vol. 71, no. 6, pp. 688–697, 2018.
- [3] B. De Coensel, D. Botteldooren, T. De Muer, B. Berglund, M. E. Nilsson, and P. Lercher, “A model for the perception of environmental sound based on notice-events,” *The Journal of the Acoustical Society of America*, vol. 126, no. 2, pp. 656–665, 2009.
- [4] World Health Organization, *Environmental noise guidelines for the European region*, World Health Organization. Regional Office for Europe, 2018.
- [5] H. Phan, T. N. T. Nguyen, P. Koch, and A. Mertins, “Polyphonic audio event detection: multi-label or multi-class multi-task classification problem?,” in *Proc. of ICASSP*, 2022, pp. 8877–8881.
- [6] Y. Hou, S. Song, C. Luo, A. Mitchell, Q. Ren, W. Xie, J. Kang, W. Wang, and D. Botteldooren, “Joint prediction of audio event and annoyance rating in an urban soundscape by hierarchical graph representation learning,” in *Proc. of INTERSPEECH*, 2023, pp. 331–335.
- [7] A. Politis, A. Mesaros, S. Adavanne, T. Heittola, and T. Virtanen, “Overview and evaluation of sound event localization and detection in DCASE 2019,” *IEEE/ACM TASLP*, vol. 29, pp. 684–698, 2020.
- [8] Q. Kong, Y. Cao, T. Iqbal, Y. Wang, W. Wang, and M. D. Plumbley, “PANNs: Large-scale pretrained audio neural networks for audio pattern recognition,” *IEEE/ACM TASLP*, vol. 28, pp. 2880–2894, 2020.
- [9] Y. Li, M. Liu, K. Drossos, and T. Virtanen, “Sound event detection via dilated convolutional recurrent neural networks,” in *Proc. of ICASSP*, 2020, pp. 286–290.
- [10] J. Gemmeke, D. Ellis, D. Freedman, A. Jansen, et al., “AudioSet: An ontology and human-labeled dataset for audio events,” in *Proc. of ICASSP*, 2017, pp. 776–780.
- [11] Y. Gong, Y. Chung, and J. Glass, “AST: Audio Spectrogram Transformer,” in *Proc. of INTERSPEECH*, 2021, pp. 571–575.
- [12] M. Boes, K. Filipan, B. De Coensel, and D. Botteldooren, “Machine listening for park soundscape quality assessment,” *Acta Acustica united with Acustica*, vol. 104, no. 1, pp. 121–130, 2018.
- [13] Y. Hou, Q. Ren, H. Zhang, A. Mitchell, F. Aletta, J. Kang, and D. Botteldooren, “AI-based soundscape analysis: Jointly identifying sound sources and predicting annoyance,” *The Journal of the Acoustical Society of America*, vol. 154, no. 5, pp. 3145–3157, 11 2023.
- [14] A. Mitchell, M. Erfanian, C. Soelistyo, T. Oberman, J. Kang, R. Aldridge, J. Xue, and F. Aletta, “Effects of soundscape complexity on urban noise annoyance ratings: A large-scale online listening experiment,” *International Journal of ERPH*, vol. 19, no. 22, pp. 14, 2022.
- [15] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *Proc. of ICLR*, 2015.
- [16] Y. Li, R. Zemel, M. Brockschmidt, and R. Zemel, “Gated graph sequence neural networks,” in *Proc. of ICLR*, 2016.
- [17] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, et al., “Attention is all you need,” in *Proc. of NeurIPS*, 2017, pp. 5998–6008.
- [18] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, et al., “Dropout: A simple way to prevent neural networks from overfitting,” *The journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [19] L. Ilya and H. Frank, “Decoupled weight decay regularization,” in *Proc. of ICLR*, 2019.
- [20] A. Mesaros, T. Heittola, and T. Virtanen, “Metrics for polyphonic sound event detection,” *Applied Sciences*, vol. 6, no. 6, pp. 162, 2016.
- [21] R. W. Emerson, “Regression analysis and adjusted R2,” *Journal of Visual Impairment & Blindness*, vol. 114, no. 4, pp. 332–334, 2020.
- [22] Y. N. Dauphin, A. Fan, M. Auli, and D. Grangier, “Language modeling with gated convolutional networks,” in *Proc. of ICML*, 2017, pp. 933–941.
- [23] I. Cohen, Y. Huang, J. Chen, J. Benesty, et al., “Pearson correlation coefficient,” *Noise Reduction in Speech Processing*, pp. 1–4, 2009.
- [24] F. N. David and C. L. Mallows, “The variance of spearman’s rho in normal samples,” *Biometrika*, vol. 48, no. 1/2, pp. 19–28, 1961.