



tinyCLAP: Distilling Contrastive Language-Audio Pretrained models

Francesco Paissan, Elisabetta Farella

Fondazione Bruno Kessler, Italy

{fpaissan, efarella}@fbk.eu

Abstract

Contrastive Language-Audio Pretraining (CLAP) became of crucial importance in the field of audio and speech processing. Its employment ranges from sound event detection to text-to-audio generation. However, one of the main limitations is the considerable amount of data required in the training process and the overall computational complexity during inference. This paper investigates how we can reduce the complexity of contrastive language-audio pre-trained models, yielding an efficient model that we call tinyCLAP. We derive an unimodal distillation loss from first principles and explore how the dimensionality of the shared, multimodal latent space can be reduced via pruning. tinyCLAP uses only 6% of the original Microsoft CLAP parameters with a minimal reduction (less than 5%) in zero-shot classification performance across the three sound event detection datasets on which it was tested.

Index Terms: contrastive language-audio pretraining, tinyML, sound event detection, zero-shot classification, distillation and pruning

1. Introduction

Contrastive Language-Audio Pretraining (CLAP) [1], and similarly, its image counterpart, CLIP [2], proved to be an effective technique to pretrain audio and image encoders. In particular, CLAP (Figure 1), and some of its variants [3, 4] achieved state-of-the-art performance for event detection, showcasing impressive results also in Zero-Shot (ZS) classification. To accomplish this, CLAP learns a similarity score between audio samples and text during the pretraining stage, which is then used to compute the affinity with unseen text prompts, potentially representing new classes. Due to its ability to correlate the audio and text modalities, CLAP finds applications in text-conditioned generative models as well, where a correlation between the text embedding and audio embedding is needed [5, 6].

Reducing the computational complexity of CLAP poses significant challenges. Still, it can yield many benefits in acoustic scene classification and audio generation in resource-constrained devices, such as those in IoT scenarios. The model capacity needed for learning the correlations between audio and text is high. Thus, CLAP encoders are not suited for fast and low-footprint inference. To overcome this challenge, we explore the use of two standard network compression techniques, namely knowledge distillation [7] and pruning [8], as they proved to be effective methods for learning more compact models while inheriting the representation capabilities of the teacher model. Still, standard knowledge distillation [7] is unsuitable for the CLAP audio encoder because its representations can be used in diverse scenarios, and the number of logits is not defined a priori. It means that the dimensionality of the soft

labels can change depending on the downstream task, hindering the possibility of running the vanilla knowledge distillation scheme. Additionally, text data might not be available when adapting the CLAP weights for a specific application domain. Thus, we focus on developing distillation and pruning strategies that can work with audio samples only, hindering the need to use the corresponding caption.

Previously, other approaches based on affinity mimicking and weight inheritance [9] tried to solve the problem of distilling models trained with a contrastive objective. However, these models (i) use cross-modal distillation schemes and (ii) produce networks that are too demanding for edge applications, counting up to tens of millions of parameters. To our knowledge, this manuscript is the first to (i) tackle the challenges of unimodal distillation of the CLAP audio encoder and pruning of the shared multimodal latent space, together with (ii) presenting low-footprint networks for zero-shot classification. We want to emphasize that we used the original CLAP weights¹ during the experimental evaluation of this manuscript for convenience. However, this approach can be used with all the other CLAP and CLIP variants and similarly formulated ZS classification models.

Section 2 presents the methods proposed in the paper, providing the elements needed to understand CLAP and our contribution. Then, in Section 3, we describe the experimental setup used for the evaluation of the proposed methods. Finally, we present the results in Section 4. We release the code for our submission through a companion website².

2. Methods

This section introduces CLAP’s core principles and the notation used throughout the paper. It also presents the distillation process and elaborates on how the pruning is performed. Figure 1 shows an overview of the proposed method.

2.1. Contrastive Language-Audio Pretraining

Let \mathbf{X}_a be the pre-processed input audio such that $\mathbf{X}_a \in \mathbb{R}^{F \times T}$, where F is the number of spectral components (e.g. mel bins), and T is the number of time frames. Let \mathbf{X}_t represent the corresponding text caption. Given a set of N elements, $\{\mathbf{X}_a, \mathbf{X}_t\}_{i=1}^N$ represents a batch of audio-text pairs. For convenience, we will refer to the batch as $\{\mathbf{X}_a, \mathbf{X}_t\}$ throughout the manuscript. Text and audio features are extracted via text and audio encoders. Let $f_a(\cdot)$ be the audio encoder and $f_t(\cdot)$ be the text encoder, then for a batch:

$$\hat{\mathbf{X}}_a = f_a(\mathbf{X}_a); \quad \hat{\mathbf{X}}_t = f_t(\mathbf{X}_t), \quad (1)$$

¹<https://zenodo.org/records/8378278>

²<https://fpaissan.github.io/tinyclapweb/>

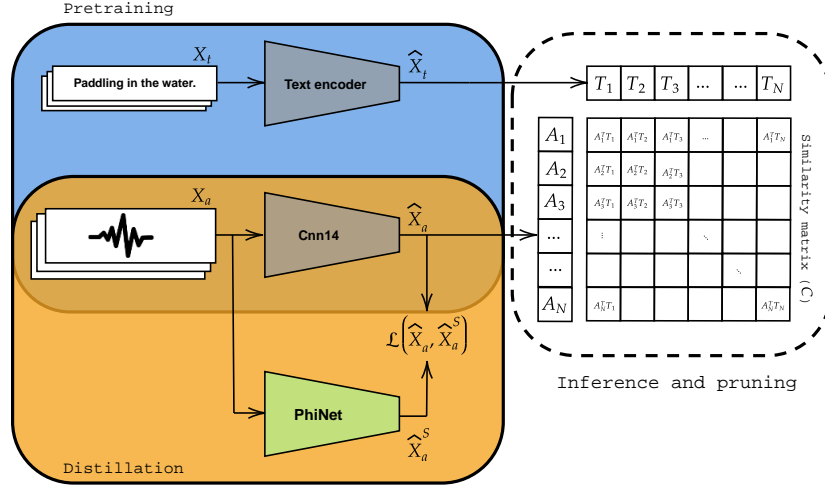


Figure 1: Diagram of the proposed distillation technique. During distillation, the student and audio encoders are aligned by minimizing the loss in Eq. 7. During the distillation stage, the encoders represented in grey are frozen, while those in green are trained. For simplicity, the image does not include the projection layers.

where $\hat{\mathbf{X}}_a \in \mathbb{R}^{N \times U}$ and $\hat{\mathbf{X}}_t \in \mathbb{R}^{N \times V}$ with U, V the audio and text latent dimensions respectively.

After encoding, CLAP brings the audio and text representations to a joint multimodal space of dimensionality d using a learned linear projection:

$$\mathbf{E}_a = \mathbf{L}_a(\hat{\mathbf{X}}_a); \quad \mathbf{E}_t = L_t(\hat{\mathbf{X}}_t), \quad (2)$$

where $\mathbf{E}_a \in \mathbb{R}^{N \times d}$, $\mathbf{E}_t \in \mathbb{R}^{N \times d}$, and L_i represent the linear projection layers.

In this shared latent space, we can measure the similarity between audio and text pairs:

$$\mathbf{C} = \tau \mathbf{E}_t \mathbf{E}_a^\top. \quad (3)$$

Note that this is a scaled cosine similarity between each other element in the original batch of N elements, with temperature τ . The CLAP loss optimizes the audio and text encoders and the linear projection layers to maximize the normalized similarity of positive pairs on the diagonal of matrix \mathbf{C} , while minimizing the similarity of the negative pairs.

2.2. Distilling Audio Representations without Text

Let $f_a^S(\cdot)$ be the distilled audio encoder, and $L_a^S(\cdot)$ be the linear projection from the student embedding to the shared multimodal latent space, then:

$$\hat{\mathbf{X}}_a^S = f_a^S(\mathbf{X}_a); \quad \mathbf{E}_a^S = L_a^S(\hat{\mathbf{X}}_a^S) \quad (4)$$

where $\hat{\mathbf{X}}_a^S \in \mathbb{R}^{N \times U}$ are the audio features extracted from the student network and $\mathbf{E}_a^S \in \mathbb{R}^{N \times d}$ is the projection in the shared latent space.

To distill the audio encoder $f_a(\cdot)$ to a more computationally efficient network (f_a^S), one option is to follow the standard knowledge distillation approach [10]. However, this method would require using both the audio and text modality to compute the soft labels for a fixed number of classes, thus limiting the generalization capabilities of the models. Instead, focusing on the structure of the latent space, we can distill the audio encoder using a single modality (e.g. only audio or text) as the

projections are in a shared latent space. We recall that the cosine similarity is maximum at convergence for positive audio-text pairs. This means that the two projections are aligned and in the same direction.

Then for a perfectly distilled model, we have:

$$\cos(\mathbf{E}_a, \mathbf{E}_t) = 1 \Leftrightarrow \cos(\mathbf{E}_a^S, \mathbf{E}_t) = 1. \quad (5)$$

Geometrically, this means that \mathbf{E}_a , \mathbf{E}_t , and \mathbf{E}_a^S are aligned and in the same direction. Thus, also

$$\cos(\mathbf{E}_a, \mathbf{E}_a^S) = 1 \quad (6)$$

holds. We want to note that the implication between Eq. 5 and Eq. 6 does not have an analytical proof as the cosine distance does not satisfy the triangle inequality [11]. However, our empirical evaluation shows that this approximation holds for the scope of this paper (see Sec. 4.1).

Eq. 6 directly defines a cost function that maximizes the cosine similarity between the new and original audio encoders, which can be used for the distillation process:

$$\mathcal{L} = -\mathbf{E}_a^S \text{sg}[\mathbf{E}_a^\top], \quad (7)$$

where $\text{sg}[\cdot]$ represents the stop-gradient operation. As mentioned, this loss function requires only the audio modality for the distillation process.

2.3. Pruning the Shared Multimodal Latent Space

To compute the loss function and distill the audio encoder using the shared multimodal latent space, it is essential that the two projections maintain the same dimensionality d . Therefore, the output of the projection layers should always be d , precluding parameter or operation reduction. After the distillation process, however, the dimensionality of the shared latent space can be reduced to $r < d$. The similarity metric measures the alignment between vectors. Therefore, zeroing out the entries associated with the smallest absolute value in the vector negligibly affects the vectors' direction and, consequently, the structure of the latent space.

Table 1: Tested networks for distillation. Hyperparameters α, β, t_0, N were derived for the different computational budgets using Hardware Aware Scaling. The last row shows the distillation results for the original CLAP encoder. Parameter count is reported for the models during the distillation process, therefore with $r = 1024$. The last row presents the results for the original CNN14 CLAP encoder without distillation.

Student model	α	β	t_0	N	Params [M]	ZS Accuracy (%)		
						TUT17	US8k	ESC50
PhiNet_1	3.00	0.75	6	7	7.0	25.2	68.3	77.4
PhiNet_2	3.00	0.75	6	9	13.0	26.4	69.7	77.2
PhiNet_3	3.00	0.75	4	7	6.2	26.1	70.3	76.5
PhiNet_4	1.50	0.75	6	7	4.4	27.5	67.9	73.0
PhiNet_5	0.75	0.75	4	7	3.5	26.7	65.2	66.1
PhiNet_6	0.75	0.75	4	4	3.2	22.1	51.8	41.9
PhiNet_7	0.75	0.75	6	4	3.3	22.3	51.6	44.1
CNN14	/	/	/	/	82.8	28.9	72.1	82.3
CNN14-CLAP	/	/	/	/	82.8	29.6	73.2	82.9

To avoid data leakage, we rank the vector entries on the training set. Therefore, we compute the average absolute value of the student projections outputs \mathbf{E}_a^S on the entire training set, such that we can rank the latent dimensions from the most important to the least, following:

$$I = \text{argsort} \left(\frac{1}{D} \sum_{i=1}^D |\mathbf{E}_{a,i}^S| \right) \quad (8)$$

where D is the number of samples in the training set, $\mathbf{E}_{a,i}^S$ is the projection of the i -th training sample in the shared, multimodal latent space, $|\cdot|$ is the element-wise absolute value operation, and the $\text{argsort}(\cdot)$ returns the indexes such that the entries of $\mathbf{E}_{a,i}^S[I]$ are sorted in a decreasing order.

2.4. Efficient Audio Encoder

To reduce the computational complexity of the original CLAP encoder, we used PhiNet [12], a scalable backbone for edge processing. PhiNet is based on inverted residual blocks. It was originally developed for computer vision applications and then benchmarked on audio processing tasks [13, 14, 15], showcasing good performance-complexity trade-offs. In the same paper, the authors proposed Hardware-Aware Scaling (HAS), a network scaling procedure that maps the computational requirements of the networks to their hyperparameters. For the purpose of this manuscript, we used HAS as a scaling technique to benchmark zero-shot classifiers with different computational budget targets.

3. Experimental setup

Pre-processing. Following the original CLAP implementation, we re-sampled all the audio files to 44.1 kHz. We computed Mel spectrograms with 64 Mel bins, a window size of 1024 samples and a hop size of 320 samples. Before feeding the samples through the neural network, we normalized the spectrograms along the frequency axis.

Encoders. As the teacher model, we used Microsoft CLAP. Its encoders are a CNN14 [16] for audio encoding and a BERT transformer [17] for text encoding. After validating the distillation loss on a self-distillation experiment, where the teacher and student network are the same CNN14, we used PhiNets

[12] with different hyperparameter configurations for the student model. The choice of hyperparameters was guided by the Hardware-Aware Scaling (HAS) principle [12] to match a variety of computational budgets. This approach allows for a proper benchmarking of how the performance of the ZS classifiers changes with the computational complexity of the encoders. Table 1 summarizes the configurations of the students and their relative ZS classification performance.

Distillation. For the distillation process, we used the audio samples from the same datasets as the original Microsoft CLAP paper [1], namely AudioCaps [18], MACS [19], FSD50k [20], and ClothoV2 [21]. Each waveform was randomly truncated during training to a continuous segment of 5 s. The captioning process was discarded, as the distillation loss only requires audio samples. For training, we used a two-stage approach. First, we used the Adam optimizer with the loss function defined in Eq. 7 and a learning rate of 3×10^{-3} to update the weights of $f_a^S(\cdot)$ and $L_a^S(\cdot)$ simultaneously for 100 epochs. Afterwards, we used the same setup, but with a learning rate of 1×10^{-3} to update the projection layer only (L_a^S). Note that the stop-gradient operator in Eq. 7 means we do not update the teacher weights during the distillation process.

Zero-Shot Evaluation. For evaluating the Zero-Shot (ZS) performance of the distilled model, we used three acoustic scene classification datasets, namely TUT17 [22], ESC50 [23] and US8k [24]. Similarly to many ZS benchmarking setups, we prepended the sentence ‘*this is the sound of*’ to the labels in order to define the caption. The text representations and projections are computed using the frozen text encoder (f_t) and projection layers (L_t) of the CLAP teacher model. As described in Section 2.3, at inference time, we use the dictionary I to only keep the top- r projection vector entries (referred to as $I^{(r)}$ from now on) of both the text and student audio encoders. As per CLAP, the class probabilities are computed with:

$$p = \text{softmax}(\mathbf{E}_a^S[I^{(r)}] \cdot \mathbf{E}_t^c[I^{(r)}]^\top) \quad (9)$$

where \mathbf{E}_t^c refers to the projection in the shared latent space of the text caption corresponding to class c .




Figure 2: a) Impact of the latent space pruning on the distilled checkpoints. The pruning procedure is described in Sec 2.3. b) Computational complexity against ZS performance of the distilled models. Note that the performance of the distilled and the original models match, considering the baseline distillation scenario. Similar plots, computed for all datasets independently, can be found on the companion website.

4. Results

The experimental results are divided into two sections. The first focuses on validating the distillation process, whilst the second one showcases how the ZS performance of distilled models drops with varying computational budgets. Due to space limitations in the manuscript, the figures showcase the average of the ZS performance for each task independently, and the tables do not explicitly mention the parameter count for all latent size combinations. We invite the reader to refer to the companion website for these additional resources.³

4.1. Distillation Sanity Checks

To validate the distillation and pruning strategies, we used the same audio encoder of the original CLAP implementation as the student encoder. This approach enables validating the distillation process itself by isolating possible contributions related to model capacity and regularization. The results of this analysis are reported in Table 1. On the three benchmarks employed in this study, the distilled CLAP encoder (CNN-14 in Table 1) achieved comparable results with the original model. In fact, the ZS performance dropped only 0.8%, averaged on the three benchmarks. Note that this slight fluctuation can also depend on the CUDA and PyTorch version or minor changes in the hyperparameters. Since it was outside the scope of this manuscript and computationally intensive, we have not performed a hyperparameter search on the training parameters of the distillation process for this network. Finally, since the distilled model performed as well as the original one, we consider the distillation process successful and Eq. 7 validated as a cost function.

4.2. Reducing Computational Complexity

As described in Sec. 2, CLAP computational complexity can be reduced with a more efficient audio encoder and a latent space pruning strategy. Hereafter, we analyze the individual contributions of these two strategies to the performance-complexity trade-off.

Audio Encoders. As shown in Table 1 and Figure 2, the networks range between 13M and 3.2M parameters, without accounting for the latent space pruning, and its consequent reduction in model size. Overall, the performance of the models drops significantly with decreasing number of parameters. For PhiNet, as expected, even the slightest increase in parameter count correlates to an increase in ZS performance. This result

aligns well with the choice of using HAS, which accommodates the biggest network for each computational budget. Quantitatively, for a reduction of approximately 4% of ZS classification performance across all benchmarks, PhiNet uses only 8% of the parameters of the original encoder without pruning the latent space. This result refers to the hyperparameters of the model denoted as PhiNet_3 in Table 1.

Latent Space Pruning. The impact of the latent space dimensionality reduction on the overall complexity of the network is related to the structure of the projection layers. Specifically, the two linear projection layers can be streamlined to output only the pertinent entries for each model, either through manual adjustments or utilizing approaches such as [25]. The impact of this complexity reduction technique is highlighted in Fig. 2a), where the average performance on the three benchmarks is showcased with respect to the size of the latent space. It is worth emphasizing that these models are not trained from scratch with a specific latent space dimensionality; rather, they are derived from the distilled checkpoints of Table 1 following the methodology described in Sect. 2.3. A clear trend emerges, indicating that for a latent space dimensionality of $r = 512$, the ZS classification performance is greater or equal to the case of $r = 1024$. Therefore, the model’s complexity is diminished by approximately 40% for the tested networks without compromising the classification performance.

5. Conclusions

In conclusion, this paper presented a distillation technique for CLAP and a pruning strategy for its latent space that only uses the audio modality. We validated the results on three public acoustic scene classification benchmarks. Our tinyCLAP model demonstrates its versatility by successfully adapting to models of varying computational complexities. We achieved considerable compression in model size (only 6% of the original parameters), with a minor drop in ZS classification accuracy (around 4%, averaged on all benchmarks). The proposed approach can be successfully applied to different CLAP variants and can extend to CLIP as well.

6. Acknowledgements

We acknowledge the support of the PNRR project FAIR - Future AI Research (PE00000013), under the NRRP MUR program funded by the NextGenerationEU.

³<https://fpaissan.github.io/tinyclapweb/>

7. References

- [1] B. Elizalde, S. Deshmukh, M. A. Ismail, and H. Wang, "Clap learning audio concepts from natural language supervision," in *2023 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2023.
- [2] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, "Learning transferable visual models from natural language supervision," in *International Conference on Machine Learning*, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:231591445>
- [3] Z. Wang, C. Subakan, K. Subramani, J. Wu, T. F. Tavares, F. Ayres, and P. Smaragdis, "Unsupervised improvement of audio-text cross-modal representations," *ArXiv*, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:258461544>
- [4] Y. Wu, K. Chen, T. Zhang, Y. Hui, T. Berg-Kirkpatrick, and S. Dubnov, "Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2023.
- [5] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, and I. Sutskever, "Zero-shot text-to-image generation," in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, M. Meila and T. Zhang, Eds. PMLR, 18–24 Jul 2021. [Online]. Available: <https://proceedings.mlr.press/v139/ramesh21a.html>
- [6] H. Liu, Z. Chen, Y. Yuan, X. Mei, X. Liu, D. P. Mandic, W. Wang, and M. Plumbley, "Audioldm: Text-to-audio generation with latent diffusion models," in *International Conference on Machine Learning*, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:256390486>
- [7] G. E. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *ArXiv*, 2015. [Online]. Available: <https://api.semanticscholar.org/CorpusID:7200347>
- [8] D. Blalock, J. J. Gonzalez Ortiz, J. Frankle, and J. Guttag, "What is the state of neural network pruning?" in *Proceedings of Machine Learning and Systems*, I. Dhillon, D. Papailiopoulos, and V. Sze, Eds., 2020. [Online]. Available: https://proceedings.mlsys.org/paper_files/paper/2020/file/6c44dc73014d66ba49b28d483a8f8b0d-Paper.pdf
- [9] K. Wu, H. Peng, Z. Zhou, B. Xiao, M. Liu, L. Yuan, H. Xuan, M. Valenzuela, X. S. Chen, X. Wang *et al.*, "Tinyclip: Clip distillation via affinity mimicking and weight inheritance," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 21 970–21 980.
- [10] J. Gou, B. Yu, S. J. Maybank, and D. Tao, "Knowledge distillation: A survey," *International Journal of Computer Vision*, 2020. [Online]. Available: <https://api.semanticscholar.org/CorpusID:219559263>
- [11] E. Schubert, "A triangle inequality for cosine similarity," in *Similarity Search and Applications: 14th International Conference, SISAP 2021, Dortmund, Germany, September 29 – October 1, 2021, Proceedings*. Berlin, Heidelberg: Springer-Verlag, 2021, p. 32–44. [Online]. Available: https://doi.org/10.1007/978-3-030-89657-7_3
- [12] F. Paissan, A. Ancilotto, and E. Farella, "Phinets: A scalable backbone for low-power ai at the edge," *ACM Transactions on Embedded Computing Systems*, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:238253101>
- [13] A. Brutti, F. Paissan, A. Ancilotto, and E. Farella, "Optimizing phinet architectures for the detection of urban sounds on low-end devices," *2022 30th European Signal Processing Conference (EUSIPCO)*, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:252999678>
- [14] F. Paissan, A. Ancilotto, A. Brutti, and E. Farella, "Scalable neural architectures for end-to-end environmental sound classification," *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:249437353>
- [15] F. Paissan, A. M. Sahabdeen, A. Ancilotto, and E. Farella, "Improving latency performance trade-off in keyword spotting applications at the edge," in *2023 9th International Workshop on Advances in Sensors and Interfaces (IWASI)*, 2023.
- [16] 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 Transactions on Audio, Speech, and Language Processing*, 2019. [Online]. Available: <https://api.semanticscholar.org/CorpusID:209444382>
- [17] J. D. M.-W. C. Kenton and L. K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of naacL-HLT*, 2019.
- [18] C. D. Kim, B. Kim, H. Lee, and G. Kim, "Audiocaps: Generating captions for audios in the wild," in *North American Chapter of the Association for Computational Linguistics*, 2019. [Online]. Available: <https://api.semanticscholar.org/CorpusID:174799768>
- [19] I. M. Morató and A. Mesáros, "Macs - multi-annotator captioned soundscapes," 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:244972894>
- [20] E. Fonseca, X. Favory, J. Pons, F. Font, and X. Serra, "Fsd50k: An open dataset of human-labeled sound events," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 829–852, 2020. [Online]. Available: <https://api.semanticscholar.org/CorpusID:222090007>
- [21] K. Drossos, S. Lipping, and T. Virtanen, "Clotho: an audio captioning dataset," *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 736–740, 2019. [Online]. Available: <https://api.semanticscholar.org/CorpusID:204800739>
- [22] K. Drossos, S. Gharib, P. Magron, and T. Virtanen, "Language modelling for sound event detection with teacher forcing and scheduled sampling," in *Workshop on Detection and Classification of Acoustic Scenes and Events*, 2019. [Online]. Available: <https://api.semanticscholar.org/CorpusID:197935289>
- [23] K. J. Piczak, "ESC: Dataset for Environmental Sound Classification," in *Proceedings of the 23rd Annual ACM Conference on Multimedia*. ACM Press, 2015. [Online]. Available: <http://dl.acm.org/citation.cfm?doi=27333373.2806390>
- [24] J. Salamon, C. Jacoby, and J. P. Bello, "A dataset and taxonomy for urban sound research," *Proceedings of the 22nd ACM international conference on Multimedia*, 2014. [Online]. Available: <https://api.semanticscholar.org/CorpusID:207217115>
- [25] G. Fang, X. Ma, M. Song, M. B. Mi, and X. Wang, "Depgraph: Towards any structural pruning," *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:256390345>