



NAS-VAD: Neural Architecture Search for Voice Activity Detection

Daniel Rho¹, Jinhyeok Park^{1,2}, and Jong Hwan Ko¹

¹Sungkyunkwan University

²Graduate School of AI, POSTECH

daniel231@skku.edu, jinhyeok1234@postech.ac.kr, jhko@skku.edu

Abstract

Various neural network-based approaches have been proposed for more robust and accurate voice activity detection (VAD). Manual design of such neural architectures is an error-prone and time-consuming process, which prompted the development of neural architecture search (NAS) that automatically designs and optimizes network architectures. While NAS has been successfully applied to improve performance in a variety of tasks, it has not yet been exploited in the VAD domain. In this paper, we present the first work that utilizes NAS approaches on the VAD task. To effectively search architectures for the VAD task, we propose a modified macro structure and a new search space with a much broader range of operations that includes attention operations. The results show that the network structures found by the proposed NAS framework outperform previous manually designed state-of-the-art VAD models in various noise-added and real-world-recorded datasets. We also show that the architectures searched on a particular dataset achieve improved generalization performance on unseen audio datasets. Our code and models are available at https://github.com/daniel03c1/NAS_VAD.

Index Terms: neural architecture search, voice activity detection

1. Introduction

Voice activity detection (VAD) is a task of detecting the existence of human voice and its onset and offset times. Early studies have utilized feature engineering and statistical signal processing methods to detect voice activity from audio signals [1, 2]. The adoption of deep learning approaches has demonstrated a significant improvement in detection performance, especially in noisy environments [3, 4, 5]. In order to further improve detection performance at low signal-to-noise ratios (SNRs), several studies have exploited contextual information in various ways and demonstrated improved performance [4, 6, 7, 8]. Zhang et al. [4], for example, proposed aggregating predictions from different temporal contexts, and other approaches employed variations of attention mechanisms to leverage contextual information [6, 7, 8]. While the existing deep learning based VAD approaches have shown remarkable progress, there is still room for further improvement of neural architectures to achieve higher VAD accuracy with a smaller model size. However, manually designing neural architectures necessitates significant engineering efforts by human designers, and it becomes increasingly difficult and time-consuming as the architecture grows in size and complexity [9, 10].

The need for a reduction of human effort in network design has led to the emergence of neural architecture search (NAS) that automatically optimizes model structures and operations. Some of the NAS methods have succeeded in creating models that outperform manually-designed models in image

classification [11, 12] and automatic speech recognition tasks [13, 14, 15]. Despite the successful adoption of NAS for various tasks, no research has explored applying NAS to the VAD task.

In this work, we propose to adopt an architecture search mechanism on the VAD task to facilitate detection performance improvement through automatic network architecture design. To the best of our knowledge, this is the first attempt to apply NAS to the VAD task. Based on a cell-based NAS approach, we propose expanding the search space to include the attention mechanism and its related operations known to be effective for the VAD task [7, 8], and modifying macro-architecture to reflect shallow depths of VAD models. The experimental results show that NAS can find VAD models that outperform the existing manually-designed network structures across a variety of audio datasets with a restricted number of parameters. In addition, the models searched by the proposed framework demonstrate enhanced generalization performance, which is an important feature of the VAD task.

2. Background and Related Work

In the field of deep learning, neural architecture search has become an important technique that can greatly reduce manual effort in neural network design. Models searched by NAS have shown comparable or even better performance than manually-designed architectures on various computer vision tasks [11, 16, 12]. In terms of the search strategy, several lines of strategies have been proposed, including evolutionary methods [17], Bayesian optimization [18, 19], reinforcement learning [20, 16], gradient-based methods [21], and random search [22]. The scope of search space is another consideration for architecture search. The search space can be the whole architecture for a given set of input and output shapes [17], or only the inner structure (operations and connections) for a fixed macro architecture [12], which acts as an exoskeleton and is not optimized during architecture search. A cell-based search space further reduces the search space by filling the inner structure with a repeating pattern, “cell” [11, 16]. As for the cell-based search space, a neural network is defined as a stack of one or two types of cells, and only the operations and connections of each cell type are optimized during architecture search. Cell-based search spaces are known to be more effective than much larger search spaces since they require fewer model evaluations to find optimal architectures without compromising performance [11].

NAS is being widely adopted in a variety of vision and natural language processing tasks [13]. Recently, audio processing tasks such as acoustic scene classification [13] and automatic speech recognition [14, 15] have begun to utilize NAS to improve model performance and design efficiency. In these speech-related tasks, several studies adjusted the search space [13] or macro structure [23] for speech processing, but

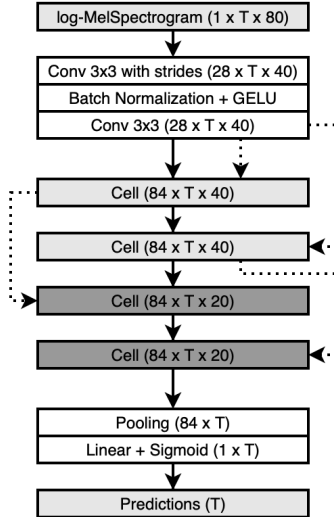


Figure 1: *The proposed macro structure. Each block denotes an operation or a module, and the parenthesis beside it shows its output shape. The first dimension is the channel dimension, and the second dimension, denoted as T , represents the time dimension, and the third number indicates the feature dimension. Dashed lines denote the second connection of each cell.*

they did not include sequential operations, which have been shown to improve VAD performance. Furthermore, there has been no study that explored the application of NAS to the VAD task. In this paper, we demonstrate the efficacy of NAS in optimizing neural architecture for enhanced voice detection performance, using a newly expanded search space and a modified macro structure.

3. NAS-VAD

In this paper, we employ NAS for the VAD task to automatically design the network architecture that can yield improved detection performance. For computational efficiency, we adopt a Bayesian optimization-based search strategy [19] on a cell-based search space. Our proposed search space includes diverse operations, including those that have been shown to be beneficial for the VAD task. We also modified the macro structure in order to search for models with fewer depths and better performance.

3.1. Search Strategy

To search for the optimal VAD model architecture, we use the cell-based Bayesian optimization approach [19], which alleviates the heavy cost of training and evaluating models using the Gaussian process as a surrogate model. Among Bayesian optimization based search methods, we adopt NAS-BOWL [19], which allows sample-efficient search compared to other Bayesian optimization based methods through the Weisfeiler-Lehman (WL) subtree graph kernel.

3.2. Search Space

We propose a new search space for the VAD task because the default search space used in NAS-BOWL was designed for image classification tasks and lacks important operations for the VAD task, including attention operations. To exploit the contextual information, we include the attention operation and its re-

lated operations (e.g., feed forward network modules) [24] into the search space. We adopt a total of 18 operations into the search space. Among them, four operations are convolutional operations, and the other six operations are attention operations. In addition, three feed forward networks [24], one squeeze-and-excitation operation [25], two gated-linear units (GLU) [26], skip connection, and zero operation (to represent a disconnection) have also been adopted.

As for convolutional operations, we used four bottleneck residual block types [27]. This includes 3x3 and 5x5 bottleneck residual blocks, each of which can have an expansion rate of 2 and 4. For bottleneck residual blocks, we use GELU [28] as an activation function.

We used three variations of multi-head attention [24]: one that just applies the attention mechanism on time, another on the feature, and the third on both time and feature dimensions. This was motivated by the spectro-temporal attention VAD [7], which improved detection performance by using two types of attention operations: one that only applies to time and the other that only applies to feature dimensions. The number of heads for each attention type was set at two or four, which resulted in a total of six attention types.

In addition, feed forward networks (FFN) [24] with an expansion rate of 0.5, 1, and 2, and gated linear units [26] with a kernel size of 3x3 and 5x5 have been adopted. We adopted pre-activation [29] in bottleneck residual blocks, multi-head attention layers, and feed-forward networks.

3.3. Macro Architecture

Our newly proposed macro architecture for NAS-VAD is shown in Fig. 1. As opposed to the basic structure [19], we use only four cells because general VAD models [30, 6, 7, 8] have a shallower depth for lower latency. Unless the input features are already in the expected shape, average pooling is applied to the input features to match the expected feature dimension size, followed by a convolutional layer with a 1x1 kernel to match the number of channels for each input port of the cells. This preprocessing operation is shown in fig. 2 as *Preprocess*. Feature reductions were applied on cells located at 1/2 of the total depth (the number of cells). In the original reduction stage used in NAS-BOWL, both height (time) and width (features) are halved, and the number of channels is doubled. However, because we need to keep the size of time dimension intact, we only reduce the width and not the height. We also employ two convolutional layers with batch normalization and GELU activation in front of the first cell, and add a linear layer right after the last cell output to generate the predictions.

4. Experiments

To evaluate the effectiveness of our NAS-VAD framework, we compared the detection performance and the parameter size of the model searched by NAS-VAD with other baseline models. We used manually designed, state-of-the-art VAD models as baselines. These baselines are built for a frame-wise voice detection framework using non-consecutive audio frames. As a result, we employed the same VAD framework for a fair comparison. Because stochastic and algorithmic VAD approaches could not work on these non-consecutive audio frames, we did not compare these models with neural network-based models.

Table 1: Comparison of model performance

trainset	model	# params	TIMIT + SoundIdeas		CV + Audioset	
			AUC	F1-score	AUC	F1-score
TIMIT + SoundIdeas	NAS-VAD	151k	0.982±0.003	0.926±0.007	0.824±0.007	0.748±0.009
	bDNN [4]	498K	0.968±0.001	0.899±0.003	0.747±0.003	0.709±0.003
	ACAM [6]	342K	0.960±0.009	0.885±0.015	0.717±0.016	0.700±0.004
	Spectro-temporal attention [7]	558K	0.977±0.003	0.915±0.007	0.781±0.017	0.724±0.007
	Self-attentive VAD [8]	155K	0.974±0.003	0.910±0.006	0.797±0.013	0.737±0.011
CV + Audioset	NAS-VAD	151k	0.965±0.001	0.894±0.001	0.976±0.001	0.915±0.002
	bDNN [4]	498K	0.941±0.000	0.861±0.002	0.950±0.000	0.871±0.000
	ACAM [6]	342K	0.940±0.002	0.860±0.003	0.948±0.002	0.869±0.003
	Spectro-temporal attention [7]	558K	0.955±0.001	0.879±0.001	0.968±0.001	0.894±0.012
	Self-attentive VAD [8]	155K	0.951±0.001	0.873±0.001	0.963±0.001	0.890±0.002

Table 2: Detection performance of VAD models on real-world datasets (AUC)

trainset	testset	Model				
		Ours	bDNN	ACAM	STA ^a	SA ^b
TIMIT	AVA	0.774	0.787	0.792	0.789	0.836
	bus stop ^c	0.980	0.876	0.924	0.962	0.954
	cons. site ^c	0.990	0.970	0.965	0.985	0.985
	park ^c	0.991	0.989	0.982	0.990	0.990
	room ^c	0.980	0.983	0.984	0.983	0.988
CV	AVA	0.905	0.868	0.858	0.892	0.892
	bus stop ^c	0.989	0.987	0.987	0.986	0.987
	cons. site ^c	0.998	0.996	0.996	0.997	0.997
	park ^c	0.991	0.989	0.989	0.990	0.989
	room ^c	0.994	0.993	0.992	0.993	0.993

^a spectro-temporal attention VAD, ^b self-attentive VAD, ^c real-world audioset from ACAM.

4.1. Experiment Settings

4.1.1. Datasets

We used noise-added speech datasets for both architecture search and model evaluation, and real-world datasets solely for performance evaluation. We used TIMIT [31] and Common Voice 7.0 [32] as clean voice sources and added noises from SoundIdeas [33] and Audiosets [34], respectively, to create noise-added speech datasets. We used the entire audio samples from the TIMIT dataset, which contains approximately 5.4 hours of English utterances. The CV dataset [32] is a publicly available multi-lingual voice dataset of more than 10,000 hours of utterances. We only used 2% of the total audio samples in the CV dataset, which is already 40 times longer than the TIMIT dataset. Because the CV dataset lacks labels for the VAD task, we used an algorithmic VAD method to extract ground truth labels from clean audio samples of the CV dataset, following VAD research conventions when ground truth labels are lacking. We used the entire SoundIdeas audio sample collection, which offers a wide range of non-speech noises such as wind and short sound effects such as a dog barking. Audioset is a publicly available, human-labeled audio dataset extracted from YouTube. We only used selected non-speech noises pre-processed by Reddy et al. [35] because the Audioset contains many audio samples containing human voices. Before adding noises,

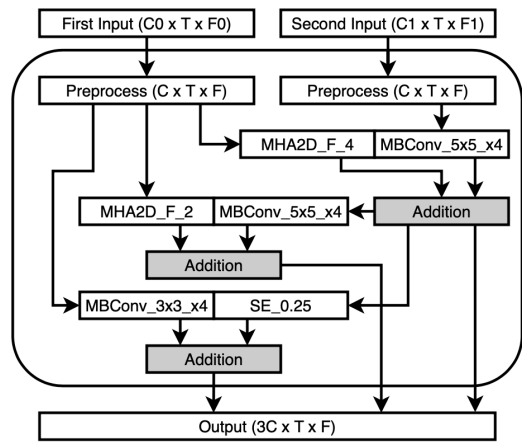


Figure 2: The cell found on TIMIT + SoundIdeas within the proposed search space. The size of the channel, time, and feature dimension are denoted by the letters C, T, and F, respectively. The Second Input is for the second connection of each cell, which is shown in Fig.1 as the dashed line. MBConv_3x3_x4 and MBConv_5x5_x4 are MobileNetV2 bottleneck residual blocks with an expansion rate of 4 and different kernel sizes (3x3 and 5x5). MHA2D_F_2 and MHA2D_F_4 are multi-head attention but only apply to the feature dimension, which has two and four heads, respectively. SE_0.25 is a squeeze-and-excitation module with a squeeze ratio of 0.25. Lastly, the three outputs resulting from the addition nodes are concatenated to construct a final result.

we padded each clean speech audio sample in order to make the number of speech and non-speech frames equal. After that, we added a randomly chosen noise to the padded speech audio sample. The SNR of noise-added audio samples was uniformly sampled from -10 to 10 dB, and SNR was calculated based on active regions of a clean speech sample and a noise audio sample. To measure the robustness of VAD models in real-world recordings, we used the AVA [36] and ACAM datasets [6] as test datasets.

We divided each audio and noise dataset into training, validation, and test groups at a ratio of 8:1:1 unless otherwise specified by the dataset. As an input feature, we used log-melspectrogram with 80 mel filterbanks, viewing each spectrogram as a single-channel image. In order to compare our

models with the baseline models, we used the same input and output settings, which are designed to detect voice activity of seven neighboring frames [t-19, t-10, t-1, t, t+1, t+10, t+19] for a given frame at time t.

4.1.2. Architecture search

We used the TIMIT-based dataset to search for an architecture. As a metric, we used Area Under the Curve (AUC), which is a commonly used metric in the VAD task performance evaluation. The default number of channels for neural networks was set to 16 in order to find VAD models with the model size similar to that of the smallest state-of-the-art VAD model [8].

4.1.3. Model Evaluation

For performance comparison, we used the following manually-designed models: bDNN [30], ACAM [6], spectro-temporal attention VAD [7], and self-attentive VAD [8]. For self-attentive VAD, we used the official model implementation, and for other models, we implemented them as described in the original papers. To train VAD models, including ours, we set the maximum training epochs to 50 for CV and 100 for TIMIT. We used the Adam optimizer with an initial learning rate of 1e-3 and the cosine annealing learning rate scheduler [37]. We used a batch size of 128 and early-stopped training if the loss did not decrease for ten consecutive epochs. For data augmentation, we used random volume scaling and frequency masking [38]. As in the baseline VAD models [4, 6], we used boosted frame-wise predictions when evaluating each model to improve voice detection performance. As for detection performance metrics, we used both AUC and f1-score, as commonly used in previous VAD studies [4, 6, 39]. We aggregated ground truth labels and corresponding predictions across multiple files to measure AUC. We binarized outputs with a threshold of 0.5 and used these binarized predictions to measure the F1-score of each model. We trained each model at least three times and reported the average test performance and its standard deviation.

To evaluate the detection performance of VAD models, we first trained and tested them on two noise-added speech datasets. After that, we also measured the performance on real-world datasets to compare the generalization performance on unseen real-world recorded audio samples.

4.2. Results

The cell structure learned on the TIMIT-based dataset in our proposed search space is depicted in Fig. 2. Table 1 shows the voice detection performance of our model and four baseline models on noise-added audio datasets. The model from our proposed search space and macro architecture shows significant performance improvement compared to the manually-designed models, outperforming every other VAD model regardless of the training and test datasets. Even though we ran architecture searches only on the TIMIT-based dataset, the resulting model showed improved detection performance also in the CV-based dataset. Furthermore, the performance gap between ours and other VAD models is markedly widened when trained on the CV-based dataset. This demonstrates that, at least in our experiments, searching for architectures in a particular dataset does not necessarily result in a decrease in generalization performance. Despite enhanced detection performance, our model has a smaller number of parameters (151K) than any other network, as demonstrated in Table 1.

Table 2 shows the generalization performance of various

Table 3: Search space comparison

search space	# params	TIMIT		CV	
		AUC	F1	AUC	F1
Ours	151k	0.982	0.926	0.824	0.748
DARTS	169K	0.981	0.923	0.807	0.732

VAD models on unseen real-world audio samples. Our model outperforms competing VAD models, achieving the best detection performance across most test datasets, as it did in the noise-added audio dataset. In particular, when trained on the CV-based dataset, our model outperformed on every real-world dataset. When trained on the TIMIT-based dataset, our model failed to outperform self-attentive VAD in room dataset [6] and in AVA. We believe this is because our model has overfitted to the noise-added training dataset when the dataset only has a small number of noise-added artificial audio samples.

In summary, using NAS, we have succeeded in finding VAD models that outperform other state-of-the-art VAD models while maintaining a similar number of parameters.

4.3. Ablation Study

As the improved VAD performance can be attributed to both the proposed search strategy and the search space, we ran an additional ablation experiment to separate the effects of the two factors. We trained two models, one from the original DARTS search space [21] and the other from our search space, on the TIMIT-based dataset and tested them on two noise-added audio datasets. As shown in Table 3, our proposed search space and macro architecture improved performance, especially in the CV dataset. Considering that we used the same search algorithm and that the resulted networks have a similar number of parameters, the performance gap can only be attributed to the modification of search space and macro structure. This demonstrates the efficacy of expanding search space by incorporating various operations, such as attention operations, along with macro structure modification.

5. Conclusion

To improve VAD performance through an automatic design process, we proposed a NAS framework with search space and macro-structure optimized for the VAD task. The results demonstrate that the models searched by NAS-VAD outperform the existing VAD models. We believe developing neural network-based VAD models in this framework might provide inspiration for more robust and accurate VAD models. We expect that the proposed framework can also be used to improve other audio processing tasks as well.

6. Acknowledgments

This work was supported by Institute of Information & communications Technology Planning Evaluation (IITP) grants (IITP-2021-0-00066, IITP-2019-0-00421, IITP-2020-0-01821, IITP-2021-0-02052) of the Korea government (MSIT).

7. References

- [1] J. Sohn, N. S. Kim, and W. Sung, "A statistical model-based voice activity detection," *IEEE Signal Processing Letters*, vol. 6, no. 1, 1999.
- [2] Z.-H. Tan, A. kr. Sarkar, and N. Dehak, "rVAD: An unsupervised

- segment-based robust voice activity detection method,” *Computer Speech & Language*, vol. 59, 2020.
- [3] X.-L. Zhang and J. Wu, “Deep belief networks based voice activity detection,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 21, no. 4, 2013.
 - [4] X.-L. Zhang and D. Wang, “Boosting contextual information for deep neural network based voice activity detection,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 2, 2015.
 - [5] T. Hughes and K. Mierle, “Recurrent neural networks for voice activity detection,” in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2013.
 - [6] J. Kim and M. Hahn, “Voice activity detection using an adaptive context attention model,” *IEEE Signal Processing Letters*, vol. 25, no. 8, 2018.
 - [7] Y. Lee, J. Min, D. K. Han, and H. Ko, “Spectro-temporal attention-based voice activity detection,” *IEEE Signal Processing Letters*, vol. 27, 2020.
 - [8] Y. R. Jo, Y. Ki Moon, W. I. Cho, and G. Sik Jo, “Self-attentive vad: Context-aware detection of voice from noise,” in *2021 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2021.
 - [9] T. Elsken, J. H. Metzen, and F. Hutter, “Neural architecture search: A survey,” *Journal of Machine Learning Research*, vol. 20, no. 55, pp. 1–21, 2019. [Online]. Available: <http://jmlr.org/papers/v20/18-598.html>
 - [10] J. Mellor, J. Turner, A. Storkey, and E. J. Crowley, “Neural architecture search without training,” in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 139. PMLR, 2021.
 - [11] B. Zoph, V. Vasudevan, J. Shlens, and Q. V. Le, “Learning transferable architectures for scalable image recognition,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
 - [12] I. Radosavovic, R. P. Kosaraju, R. Girshick, K. He, and P. Dollár, “Designing network design spaces,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
 - [13] J. Li, C. Liang, B. Zhang, Z. Wang, F. Xiang, and X. Chu, “Neural Architecture Search on Acoustic Scene Classification,” in *Proc. Interspeech 2020*, 2020.
 - [14] Y.-C. Chen, J.-Y. Hsu, C.-K. Lee, and H. yi Lee, “DARTS-ASR: Differentiable Architecture Search for Multilingual Speech Recognition and Adaptation,” in *Proc. Interspeech 2020*, 2020.
 - [15] J. Kim, J. Wang, S. Kim, and Y. Lee, “Evolved Speech-Transformer: Applying Neural Architecture Search to End-to-End Automatic Speech Recognition,” in *Proc. Interspeech 2020*, 2020.
 - [16] H. Pham, M. Guan, B. Zoph, Q. Le, and J. Dean, “Efficient neural architecture search via parameters sharing,” in *Proceedings of the 35th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 80. PMLR, 2018.
 - [17] E. Real, S. Moore, A. Selle, S. Saxena, Y. L. Suematsu, J. Tan, Q. V. Le, and A. Kurakin, “Large-scale evolution of image classifiers,” in *Proceedings of the 34th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 70. PMLR, 2017.
 - [18] H. Zhou, M. Yang, J. Wang, and W. Pan, “BayesNAS: A Bayesian approach for neural architecture search,” in *Proceedings of the 36th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 97. PMLR, 2019.
 - [19] B. X. Ru, X. Wan, X. Dong, and M. A. Osborne, “Interpretable neural architecture search via bayesian optimisation with weisfeiler-lehman kernels,” in *International Conference on Learning Representations*, 2021.
 - [20] D. P. Kingma and J. Ba, “Neural architecture search with reinforcement learning,” in *International Conference on Learning Representations*, 2017.
 - [21] H. Liu, K. Simonyan, and Y. Yang, “DARTS: Differentiable architecture search,” in *International Conference on Learning Representations*, 2019.
 - [22] L. Li and A. Talwalkar, “Random search and reproducibility for neural architecture search,” in *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, ser. Proceedings of Machine Learning Research, vol. 115. PMLR, 22–25 Jul 2020.
 - [23] S. Ding, T. Chen, X. Gong, W. Zha, and Z. Wang, “AutoSpeech: Neural Architecture Search for Speaker Recognition,” in *Proc. Interspeech 2020*, 2020.
 - [24] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
 - [25] J. Hu, L. Shen, and G. Sun, “Squeeze-and-excitation networks,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
 - [26] Y. N. Dauphin, A. Fan, M. Auli, and D. Grangier, “Language modeling with gated convolutional networks,” in *Proceedings of the 34th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 70. PMLR, 2017.
 - [27] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, “MobileNetV2: Inverted residuals and linear bottlenecks,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
 - [28] D. Hendrycks and K. Gimpel, “Gaussian error linear units (gelus),” *arXiv preprint arXiv:1606.08415*, 2016.
 - [29] K. He, X. Zhang, S. Ren, and J. Sun, “Identity mappings in deep residual networks,” in *European Conference on Computer Vision*, 2016.
 - [30] X.-L. Zhang and D. Wang, “Boosting contextual information for deep neural network based voice activity detection,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 2, 2016.
 - [31] J. Garofolo, L. Lamel, W. Fisher, J. Fiscus, D. Pallett, N. Dahlgren, and V. Zue, “Timit acoustic-phonetic continuous speech corpus,” *Linguistic Data Consortium*, 11 1992.
 - [32] R. Ardila, M. Branson, K. Davis, M. Henretty, M. Kohler, J. Meyer, R. Morais, L. Saunders, F. M. Tyers, and G. Weber, “Common voice: A massively-multilingual speech corpus,” in *Proceedings of the 12th Conference on Language Resources and Evaluation (LREC 2020)*, 2020.
 - [33] sound ideas.com, “General series 6000 combo,” [Online]. Available: <https://www.sound-ideas.com/Product/51/The-General-Series-6000-Combo>.
 - [34] J. F. Gemmeke, D. P. W. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, “Audio set: An ontology and human-labeled dataset for audio events,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2017.
 - [35] C. K. Reddy, H. Dubey, K. Koishida, A. Nair, V. Gopal, R. Cutler, S. Braun, H. Gamper, R. Aichner, and S. Srinivasan, “INTER-SPEECH 2021 Deep Noise Suppression Challenge,” in *Proc. Interspeech 2021*, 2021.
 - [36] S. Chaudhuri, J. Roth, D. P. W. Ellis, A. Gallagher, L. Kaver, R. Marvin, C. Pantofaru, N. Reale, L. Guarino Reid, K. Wilson, and Z. Xi, “AVA-Speech: A Densely Labeled Dataset of Speech Activity in Movies,” in *Proc. Interspeech 2018*, 2018.
 - [37] I. Loshchilov and F. Hutter, “SGDR: Stochastic Gradient Descent with Warm Restarts,” in *International Conference on Learning Representations*, 2017.
 - [38] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition,” in *Proc. Interspeech 2019*, 2019.
 - [39] Y. Chen, H. Dinkel, M. Wu, and K. Yu, “Voice Activity Detection in the Wild via Weakly Supervised Sound Event Detection,” in *Proc. Interspeech 2020*, 2020.