



An Efficient Speech Separation Network Based on Recurrent Fusion Dilated Convolution and Channel Attention

Junyu Wang

School of Electronic Information, Sichuan University, China

wjy446320987@163.com

Abstract

We present an efficient speech separation neural network, ARFDCN, which combines dilated convolutions, multi-scale fusion (MSF), and channel attention to overcome the limited receptive field of convolution-based networks and the high computational cost of transformer-based networks. The suggested network architecture is encoder-decoder based. By using dilated convolutions with gradually increasing dilation value to learn local and global features and fusing them at adjacent stages, the model can learn rich feature content. Meanwhile, by adding channel attention modules to the network, the model can extract channel weights, learn more important features, and thus improve its expressive power and robustness. Experimental results indicate that the model achieves a decent balance between performance and computational efficiency, making it a promising alternative to current mainstream models for practical applications.

Index Terms: speech separation, dilated convolution, integrate information, channel attention, time-domain

1. Introduction

The objective of speech separation is to separate the required sound from surrounding noise such as environmental noise and other voices. Depending on the number of channels or microphones, speech separation tasks are frequently divided into multi-channel and single-channel. Because of the challenges in encoding multi-speaker mixture speech signals into a completely separable feature space [1], speech separation remains an extremely difficult research subject, particularly in the case of single-channel. We focus on single-channel speech separation.

Numerous attempts have been undertaken throughout the years to tackle this issue, despite the fact that single-channel speech separation presents numerous difficulties. Before the age of deep learning, a number of classical approaches were used to this problem, including non-negative matrix factorization (NMF) [2, 3], computational auditory scene analysis (CASA) [4], and support vector machine (SVM) [5]. However, these models frequently struggle to represent the nonlinear connection between speech data, which severely restricts their practical utilization. With the emergence of deep learning approaches in other disciplines [6, 7], researchers have begun to employ massive amounts of data to construct models capable of separating mixed speech from unknown speakers; as a result, the performance of this task has been significantly enhanced. Currently, single-channel speech separation approaches are primarily classified into two categories: end-to-end time domain and time-frequency (T-F) domain. In the T-F domain, the input signal is first transformed into T-F features using the short-time

Fourier transform (STFT), followed by the hadamard product of the amplitude spectrum of the mixed signal and the predicted masking value to obtain the estimated amplitude spectrum of the target signal. Lastly, the enhanced amplitude spectrum and the original phase spectrum [8, 9, 10, 11] are subjected to the inverse Fourier transform (ISTFT) to separate speech.

However, due to the difficulty in reconstructing the waveform phase, the predicted source waveform is typically synthesized by mixing the initial phase, reducing the performance ceiling of the method. To overcome this issue, the paper [12] introduced the time-domain speech separation approach, which employs the codec framework based on convolutional neural networks to model the mixed waveform directly, and the separation section is comprised of multiple convolutional filters, which has achieved good results. Numerous subsequent works [13, 14, 15, 16] are devoted to constructing better separators on TASNET to better enhance speech separation performance.

Among these time-domain approaches, convolution-based multi-scale fusion (MSF) methods [17, 18] and transformer-based methods [19, 20] have achieved excellent performance in speech separation tasks by learning features at different time scales. However, transformer-based models typically require enormous computational cost and have long training time, making them less suitable for practical applications. Convolution-based models often encounter the issue of limited receptive field, which causes them to fall short of reaching the theoretical upper limit of the model. Additionally, we have observed that few methods in speech separation focus on channel attention, which may lead to difficulty capturing the channel correlations between input features, thus limiting performance.

Based on these findings, we present an ARFDCN that incorporates MSF, dilated convolutions, and channel attention, which makes full use of the time dependence of speech signals by combining exponentially growing dilated convolution [21, 22] with MSF to ensure that the network has dynamic receptive fields, so as to achieve excellent separation performance. Moreover, channel attention is added after each MSF block, allowing the network to focus on specific regions and thus enhance the ability of feature representation and discrimination of the network. Because of the application of dilated convolutions, MSF, and channel attention, the features derived from the network during the speech separation stage are more comprehensive and can make better use of the information between contexts.

2. ARFDCN

The presented ARFDCN, like Conv-Tasnet [14], involves three phases: encoder, separation layer, and decoder. First, the encoder converts the blended speech waveform into corresponding features in the separable embedding feature space. After

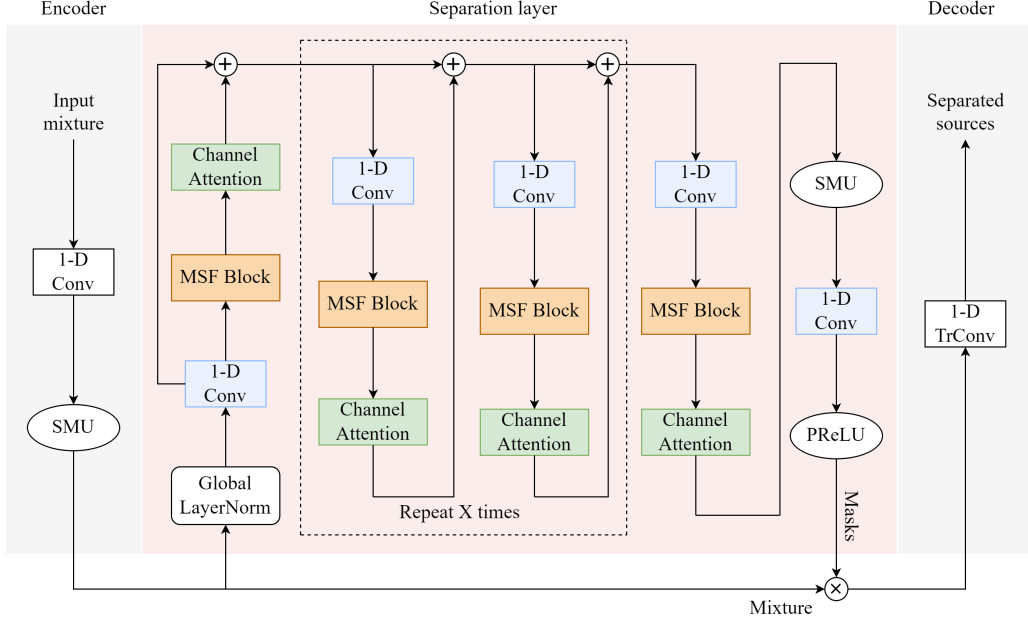


Figure 1: System flowchart of ARFDCN.

the features have been sent, a mask for each source is created by the separation layer. Finally, by transforming the mask features, the decoder reconstructs the waveforms of each source [19]. The flowchart of the model is represented in Figure 1.

2.1. Encoder

The input mixing speech is separated into overlapping segments of length H in this module, denoted by $E_s \in R^{1 \times H}$, where $s = 1, \dots, N$ indicates the segment index and N is the total number of input segments. By using a trainable 1-D convolution operation and SMU activation [23], the following E_s is encoded into a high-dimensional mixed feature $X \in R^{C \times L}$:

$$X = SMU(Conv1D(x)) \quad (1)$$

where C represents the number of channels and L is the signal length generated by 1-D convolution.

2.2. Separation layer

The separation layer is inspired by [6, 18]. It is predominantly made up of channel attention, variable dilated convolutions and MSF. After the output of the encoder is delivered to the separation layer, it first undergoes global layer normalization (GLN) [14] to boost network convergence speed, followed by 1-D convolution operation with P output channels. After the convolution output, it passes through multiple MSF blocks and channel attention modules in turn.

The specific structure of the MSF block is illustrated in Figure 2. In each MSF block, the input flows from bottom to top, obtaining different receptive field information at different stages, with a total of J stages. The lower-level dilated convolution is used to extract local detail features, while the higher-level dilated convolution is used to extract global semantic features. GLN is inserted between each dilated convolution operation to ensure more stable output data. In addition, to allow the higher-level dilated convolution layer to obtain a sufficiently large re-

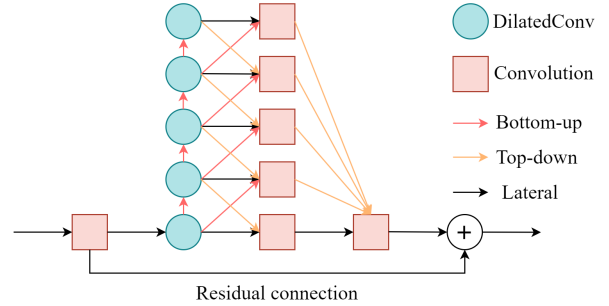


Figure 2: Multi-scale information fusion block.

ceptive field, we use dilation values that increase exponentially in powers of two from bottom to top. By fusing information from these different scales, the model can learn comprehensive features and improve its performance. Each convolution layer in the MSF block is followed by GLN and PReLU function, which helps to improve the expressiveness and robustness of the model. Finally, the residual connection is introduced between the input and output of the MSF block to enhance the representational capacity of the model and alleviate the gradient vanishing issue. Compared to the widely used dual-path network for learning local and global features, the MSF block has a shorter computation time and is more suitable for practical applications.

To further improve the generalization ability of the model and make it more focused on important features, channel attention is added after each MSF block, as shown in Figure 3. First, maximum pooling and average pooling are applied to the inputs along the last dimension, respectively. Then, a 1-D convolution with a kernel size of 5 is applied to generate a weight vector for each channel. The two weight vectors are added and multiplied with the input feature to weigh each channel and generate

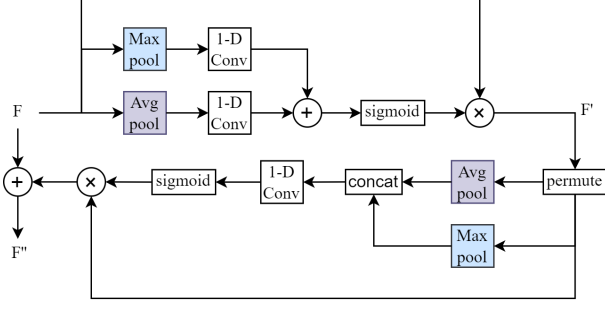


Figure 3: Channel attention module.

a new feature map F' . Next, maximum pooling and average pooling are applied to the new feature vector along its channel dimension, respectively. A 1-D convolution with a kernel size of 21 is applied to generate a weight vector for the spatial position, and the weight vector is multiplied with the input to generate the final feature. Additionally, to alleviate gradient vanishing and improve model stability, residual connection is incorporated into the channel attention module. The formula for channel attention is as follows:

$$F' = \phi(h^5(Avp(F)) + h^5(Map(F))) * F \quad (2)$$

$$F'' = \phi(h^{21}(concat[Avp(F'), Map(F')])) * F' + F \quad (3)$$

where $Avp(\cdot)$ and $Map(\cdot)$ represent average pooling and max pooling operations, respectively. h^5 represents a 1-D convolution operation with a kernel size of 5, input channel of 1, and output channel of 1. $\phi(\cdot)$ represents a sigmoid activation operation. $concat[\cdot]$ represents concatenation along the channel dimension. h^{21} represents a 1-D convolution operation with a kernel size of 21, input channel of 2, and output channel of 1.

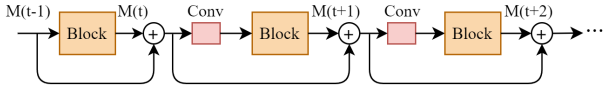


Figure 4: Integration of MSF blocks.

Moreover, to better fuse information among MSF blocks and prevent gradient vanishing. Before delivering to the subsequent block, we integrate the output of the current block with the output of all previous blocks and the input of the initial MSF block, and add a 1-D convolution before each input of the MSF block. Each convolution layer is followed by GLN and SMU activation function, as shown in Figure 4. This approach is described as follows:

$$M(t+1) = f(c(M(t) + M(t-1) + \dots + e)) \quad (4)$$

where $f(\cdot)$ represents the operation of passing through an MSF block, $c(\cdot)$ represents the 1-D convolution operation with GLN and SMU activation, $M(t)$ represents the output of the MSF block at time step t , and e represents the input of the initial MSF block.

The output of the last MSF block is passed through a 1-D convolution operation with PReLU activation to produce a mask for each source.

2.3. Decoder

The 1-D transposed convolution process of the decoder employs the same stride and convolution kernel as the encoder. The input of the decoder is obtained by element-wise multiplication between the output of the encoder X and the mask D_i of the i -th source. The transformation of the decoder can be represented as follows:

$$Y_i = ConvTranspose1d(D_i * X) \quad (5)$$

where $Y_i \in R^{1 \times T}$ is the final waveform of the i -th speaker.

3. Experimental procedures

3.1. Dataset

Using three publicly accessible English datasets, WSJ0-2Mix [8], Libri2mix [24], and WHAM! [25], we examine the separation performance of the ARFDCN model. The voice source for the WSJ0-2Mix is the WSJ0 [26]. The `si_tr_s` file generated training data for 30 hours and validation data for 10 hours. The `si_dt_05` and `si_et_05` files created 5 hours of test data. The WHAM! is a noisy speech-mixed dataset based on the WSJ0-2Mix. This dataset is created by adding a different ambient noise sample to each of the two speakers' mixed speech from the WSJ0-2Mix. In this experiment, it is required to not only separate the speaker signal but also reduce the background noise. This task is more difficult than separating the WSJ0-2Mix. The Libri2Mix is created using data from the LibriSpeech corpus [27]. In the LibriSpeech, different speaker signals are arbitrarily extracted and combined using the uniform sampling loudness unit of relative full scale (LUFS) [28] between -25 and -33 dB. Random noise samples are added to the speech mixture, uniformly distributed with loudness between -38 and -30 LUFS. The training set is the subset train-100. All speech signals are resampled at a frequency of 8kHz.

3.2. Experiment configurations

Encoder and decoder are configured with 512 channels, kernel size of 21, and stride size of 10. The number of times X and channels P in the separation layer are set to 7 and 512. The value of J in the MSF block is set to 5, and the dilation values for the dilated convolutions are set from low to high as $\{1, 2, 4, 8, 16\}$, with a kernel size of 5 and a stride size of 2.

The model is trained for 200 epochs. In WSJ0-2Mix and WHAM!, the training speech segment length is set to 4-second, whereas in Libri2Mix, the training speech segment length is set to 3-second. We utilize AdamW [29] as the optimizer and set the initial learning rate to 1×10^{-3} during training. If the validation accuracy does not increase for two successive epochs, the learning rate will be reduced to 0.9 times.

To maximize the signal fidelity of the model, this paper takes the scale-invariant source-to-distortion ratio (SI-SDR) [30, 31] as the training target. It can be described as:

$$s_{target} = \frac{\langle \tilde{y}, y \rangle y}{\|y\|^2} \quad (6)$$

$$e_{noise} = \tilde{y} - s_{target} \quad (7)$$

$$SI - SDR = 10 \log_{10} \frac{\|s_{target}\|^2}{\|e_{noise}\|^2} \quad (8)$$

where \tilde{y} and y represent estimated and clean sources, respectively. Prior to calculation, both are normalized to zero-mean.

4. Results

4.1. Comparison with other models

We begin by presenting the results for the WSJ0-2Mix dataset. Table 1 compares ARFDCN with multiple prevalent models on the WSJ0-2Mix, where the “Time” metric is computed by processing a 4-second audio sample with a sample rate of 8kHz on Intel(R) Core(TM) i9-12900KF CPU @ 3.19GHz, averaged over 100 attempts. The results indicate that Sepformer achieves only modest improvement over ARFDCN, despite significantly increasing the model size and reasoning time.

Table 1: Evaluation of different speech separation models utilizing Si-SDRi and SDRi on WSJ0-2Mix

Model	SI-SDRi	SDRi	Time(s)	Para.(M)
DPCL++ [8]	10.8	11.2	0.47	13.6
uPIT-BLSTM-ST [10]	9.8	10.0	0.97	92.7
Chimera++ [11]	11.5	11.8	0.83	32.9
BLSTM-TasNet [12]	13.2	13.6	3.66	23.6
Conv-TasNet [14]	15.3	15.6	0.53	5.6
MSGT-TasNet [15]	17.0	17.3	6.99	66.8
SuDoRM-RF 1.0x [17]	17.0	17.3	0.67	2.66
DualPathRNN [16]	18.8	19.0	5.88	2.63
DPTNet [19]	20.2	20.6	8.41	2.69
Gated DPRNN [32]	20.1	-	9.25	7.5
A-FRCNN-16 [18]	18.3	18.6	1.91	6.1
Sepformer [20]	20.4	20.5	6.58	26.0
ARFDCN	20.3	20.5	1.81	6.14

To demonstrate the generalizability of our method, relevant experiments are conducted using the WHAM! dataset. Compared with the data in the WSJ0-2Mix, random noise is added to make speech separation more difficult. Table 2 presents the findings of the ARFDCN and 10 contrasted methods on WHAM!. From these results, it can be seen that the ARFDCN and Gated DPRNN achieve almost identical performance, with significantly improved computational efficiency and reduced model size.

Table 2: Performance of different models on WHAM!

Model	SI-SDRi	SDRi	Para.(M)
Chimera++ [11]	10.0	-	32.9
BLSTM-TasNet [12]	9.8	-	23.6
Conv-TasNet [14]	12.7	-	5.6
MGST-TasNet [15]	13.1	-	66.8
SuDoRM-RF 1.0x [17]	12.9	13.3	2.66
DualPathRNN [16]	13.7	14.1	2.63
DPTNet [19]	14.9	15.3	2.69
Gated DPRNN [32]	15.2	-	7.5
A-FRCNN-16 [18]	14.5	14.8	6.1
Sepformer [20]	14.4	15.0	26.0
ARFDCN	15.2	15.5	6.14

Lastly, we verify the performance of ARFDCN on the Libri2Mix dataset. The dataset combination is more challenging to separate than the WSJ0-2Mix, and random noise has been included to get it closer to real application. Table 3 presents the outcomes of the ARFDCN and 8 compared methods on the Libri2Mix. As can be shown, the suggested method ARFDCN

Table 3: Performance of different models on Libri2Mix

Model	SI-SDRi	SDRi	Para.(M)
Chimera++ [11]	6.3	7.0	32.9
BLSTM-TasNet [12]	7.9	8.7	23.6
Conv-TasNet [14]	12.2	12.7	5.6
SuDoRM-RF 1.0x [17]	13.5	14.0	2.66
DualPathRNN [16]	16.1	16.6	2.63
DPTNet [19]	16.7	17.1	2.69
A-FRCNN-16 [18]	16.7	17.2	6.1
Sepformer [20]	16.5	17.0	26.0
ARFDCN	17.1	17.6	6.14

is superior to the compared methods, which further demonstrates its effectiveness.

4.2. Ablation study

We design four experiments to demonstrate the effects of different modules on the performance, conducted on the WSJ0-2Mix dataset. RFCN removes the channel attention module from ARFDCN and sets all dilation values in the MSF block to 1. RFDCN removes only the channel attention module from ARFDCN. ARFCN sets all dilation values in the MSF block to 1 while keeping the channel attention module. Comparing RFDCN and RFCN shows that combining dilated convolutions with MSF is more beneficial for expanding the receptive field and extracting global information. Moreover, Table 4 shows that ARFCN performs better than RFCN, and ARFDCN performs better than RFDCN, demonstrating the effectiveness of channel attention in improving model performance by focusing on important features.

Table 4: Ablation study results on WSJ0-2Mix

Model	SI-SDRi	SDRi	Time(s)
RFCN	18.9	19.2	1.83
RFDCN	19.6	19.9	1.77
ARFCN	19.7	20.0	1.87
ARFDCN	20.3	20.5	1.81

5. Conclusion

This paper presents ARFDCN, a single-channel speech separation neural network that combines dilated convolutions and MSF to learn micro and macro features at different stages and fuses them to obtain large receptive fields. Moreover, incorporating multiple channel attention modules enables the model to focus more on significant features, improving its expressive power and generalization performance. Experiments on three datasets exhibit that the network reaches a decent balance between computational cost and performance. Future research will focus on the speech separation of multiple speakers under reverberation conditions and noisy environments to better serve society.

6. References

- [1] Z. Yao, W. Pei, F. Chen, G. Lu, and D. Zhang, “Stepwise-refining speech separation network via fine-grained encoding in high-order

- latent domain,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 378–393, 2022.
- [2] M. N. Schmidt and R. K. Olsson, “Single-channel speech separation using sparse non-negative matrix factorization,” in *Proc. INTERSPEECH 2006 – 07th Annual Conference of the International Speech Communication Association*, Pittsburgh, PA, USA, 2006, pp. paper 1652–Thu2FoP.10.
 - [3] J. L. Roux, J. R. Hershey, and F. Weninger, “Deep nmf for speech separation,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 66–70.
 - [4] J. Rouat, “Computational auditory scene analysis: Principles, algorithms, and applications (wang, d. and brown, g.j., eds.; 2006) [book review],” *IEEE Transactions on Neural Networks*, vol. 19, no. 1, pp. 199–199, 2008.
 - [5] K. Han and D. Wang, “An svm based classification approach to speech separation,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2011, pp. 4632–4635.
 - [6] S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, “Cbam: Convolutional block attention module,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 3–19.
 - [7] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 770–778.
 - [8] J. R. Hershey, Z. Chen, J. L. Roux, and S. Watanabe, “Deep clustering: Discriminative embeddings for segmentation and separation,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 31–35.
 - [9] D. Yu, M. Kolbæk, Z.-H. Tan, and J. Jensen, “Permutation invariant training of deep models for speaker-independent multi-talker speech separation,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 241–245.
 - [10] M. Kolbæk, D. Yu, Z.-H. Tan, and J. Jensen, “Multitalker speech separation with utterance-level permutation invariant training of deep recurrent neural networks,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 10, pp. 1901–1913, 2017.
 - [11] Z.-Q. Wang, J. L. Roux, and J. R. Hershey, “Alternative objective functions for deep clustering,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 686–690.
 - [12] Y. Luo and N. Mesgarani, “Tasnet: Time-domain audio separation network for real-time, single-channel speech separation,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 696–700.
 - [13] D. Ditter and T. Gerkmann, “A multi-phase gammatone filterbank for speech separation via tasnet,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 36–40.
 - [14] Y. Luo and N. Mesgarani, “Conv-tasnet: Surpassing ideal time-frequency magnitude masking for speech separation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 8, pp. 1256–1266, 2019.
 - [15] Y. Zhao, C. Luo, Z.-J. Zha, and W. Zeng, “Multi-scale group transformer for long sequence modeling in speech separation,” in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2021, pp. 3251–3257.
 - [16] Y. Luo, Z. Chen, and T. Yoshioka, “Dual-path rnn: Efficient long sequence modeling for time-domain single-channel speech separation,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 46–50.
 - [17] E. Tzinis, Z. Wang, and P. Smaragdis, “Sudo rm-rf: Efficient networks for universal audio source separation,” in *2020 IEEE 30th International Workshop on Machine Learning for Signal Processing (MLSP)*, 2020, pp. 1–6.
 - [18] X. Hu, K. Li, W. Zhang, Y. Luo, J.-M. Lemerrier, and T. Gerkmann, “Speech separation using an asynchronous fully recurrent convolutional neural network,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2021, pp. 22 509–22 522.
 - [19] J. Chen, Q. Mao, and D. Liu, “Dual-path transformer network: Direct context-aware modeling for end-to-end monaural speech separation,” in *Proc. INTERSPEECH 2020 – 21st Annual Conference of the International Speech Communication Association*, Shanghai, China, 2020, pp. 2642–2646.
 - [20] S. C. M. B. C. Subakan, M. Ravanelli and J. Zhong, “Attention is all you need in speech separation,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 21–25.
 - [21] C. Lea, R. Vidal, A. Reiter, and G. D. Hager, “Temporal convolutional networks: A unified approach to action segmentation,” in *Computer Vision – ECCV 2016 Workshops*, 2016, pp. 47–54.
 - [22] C. Lea, M. D. Flynn, R. Vidal, A. Reiter, and G. D. Hager, “Temporal convolutional networks for action segmentation and detection,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 156–165.
 - [23] K. Biswas, S. Kumar, S. Banerjee, and A. K. Pandey, “Smooth maximum unit: Smooth activation function for deep networks using smoothing maximum technique,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 784–793.
 - [24] J. Cosentino, M. Pariente, S. Cornell, A. Deleforge, and E. Vincent, “Librimix: An open-source dataset for generalizable speech separation,” *arXiv preprint arXiv:2005.11262*, 2020.
 - [25] G. Wichern, J. Antognini, M. Flynn, L. R. Zhu, E. McQuinn, D. Crow, E. Manilow, and J. L. Roux, “Wham!: Extending speech separation to noisy environments,” in *Proc. INTERSPEECH 2019 – 20th Annual Conference of the International Speech Communication Association*, Graz, Austria, 2019, pp. 1368–1372.
 - [26] J. Garofolo, D. Graff, D. Paul, and D. Pallett, “Csr-i (wsj0) complete ldc93s6a,” *Web Download. Philadelphia: Linguistic Data Consortium*, vol. 83, 1993.
 - [27] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: an asr corpus based on public domain audio books,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 5206–5210.
 - [28] B. Series, “Algorithms to measure audio programme loudness and true-peak audio level,” 2011.
 - [29] I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” in *International Conference on Learning Representations (ICLR)*, 2019, pp. 1–14.
 - [30] J. L. Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, “Sdr-half-baked or well done?” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 626–630.
 - [31] E. Vincent, R. Gribonval, and C. Févotte, “Performance measurement in blind audio source separation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 4, pp. 1462–1469, 2006.
 - [32] E. Nachmani, Y. Adi, and L. Wolf, “Voice separation with an unknown number of multiple speakers,” in *International Conference on Machine Learning (ICML)*, 2020, pp. 7164–7175.