

SOFT LABEL CODING FOR END-TO-END SOUND SOURCE LOCALIZATION WITH AD-HOC MICROPHONE ARRAYS

Linfeng Feng, Yijun Gong, Xiao-Lei Zhang

School of Marine Science and Technology, Northwestern Polytechnical University, China
Research & Development Institute of Northwestern Polytechnical University in Shenzhen, China
fenglinfeng@mail.nwpu.edu.cn, gongyj@mail.nwpu.edu.cn, xiaolei.zhang@nwpu.edu.cn

ABSTRACT

Recently, an end-to-end two-dimensional sound source localization algorithm with ad-hoc microphone arrays formulates the sound source localization problem as a classification problem. The algorithm divides the target indoor space into a set of local areas, and predicts the local area where the speaker locates. However, the local areas are encoded by one-hot code, which may lose the connections between the local areas due to quantization errors. In this paper, we propose a new soft label coding method, named *label smoothing*, for the classification-based two-dimensional sound source location with ad-hoc microphone arrays. The core idea is to take the geometric connection between the classes into the label coding process. The first one is named static soft label coding (SSLC), which modifies the one-hot codes into soft codes based on the distances between the local areas. Because SSLC is handcrafted which may not be optimal, the second one, named dynamic soft label coding (DSLCL), further rectifies SSLC, by learning the soft codes according to the statistics of the predictions produced by the classification-based localization model in the training stage. Experimental results show that the proposed methods can effectively improve the localization accuracy.

Index Terms— Ad-hoc microphone arrays, sound source localization, soft label coding

1. INTRODUCTION

Sound source localization is a problem of estimating the relative positions of sound sources to the locations of microphone arrays. It has a wide range of applications, such as speech enhancement and separation [1], speech recognition [2], hearing aids [3], etc. Conventionally, the problem is simplified to the estimation of the direction of arrival (DOA) of sound sources [4]. Representative approaches include time difference of arrival (TDOA) [5], steered response power with

phase transform (SRP-PHAT) [6], and multiple signal classification (MUSIC)[7]. Recently, there is increased research interest in using deep neural network (DNN) for localizing sound sources. It takes spatial features, e.g. phase spectrogram [8], cross-correlation [9], or spatial pseudo-spectra [10], as the input of DNN to estimate the DOA. The aforementioned methods are based on a single microphone array.

If we jointly use multiple microphone arrays for sound source localization, then we might get the two-dimensional (2D) localization of sound sources directly. Ad-hoc microphone array, which collaboratively organizes a set of randomly distributed microphone arrays in space, is a solution to the problem. Early methods are mostly based conventional signal processing methods [11]. Recently, deep-learning-based methods were studied [12], which is the focus of the paper. Although the pioneering works in [13] made significant contribution to the direction, they make strong assumptions to the ad-hoc microphone arrays, such as fixed positions of microphone nodes at both the training and test stages, which do not fully utilize the advantage of the flexibility of ad-hoc arrays. Although the method in [14] can handle flexible ad-hoc arrays that consists of any number of randomly distributed ad-hoc nodes, it is a stage-wise method. Its deep models are used only at each ad-hoc node, leaving the 2-dimensional localization process an independent geometric approach.

As we know, a key advantage of deep learning is that it can address a task in an end-to-end manner. In [15], an end-to-end 2D sound source localization model with ad-hoc arrays was proposed. It formulates the 2D sound source localization problem as a classification problem, where each class represents a local area of a targeted room and is encoded by a one-hot code. Moreover, the method is effective in the situation where each ad-hoc node contains only a single microphone. However, the position of a speaker is estimated as a one-hot code, which can be problematic. For one-hot, the correct class is the same distance from any incorrect class.

To address the problem, we propose a label smoothing algorithm to remedy the weakness of the one-hot codes. Specifically, we first propose static soft label coding (SSLC) based

Xiao-Lei Zhang is the corresponding author.

This work was supported in part by the Project of the Science, Technology, and Innovation Commission of Shenzhen Municipality, China under Grant JSGG20210802152546026 and JCYJ20210324143006016, and in part by the National Science Foundation of China (NSFC) under Grant 62176211.




Fig. 1. The classification-based formulation of 2D sound source location with ad-hoc microphone arrays, where a room is partitioned into multiple local areas.

on the distances between the centers of local regions instead of the one-hot codes. Because SSLC is handcrafted that may not be accurate enough, we further propose to remedy SSLC by dynamic soft label coding (DSLCL). DSLCL is learned from the inference statistics produced by the localization model at the training stage. Experimental results demonstrate the effectiveness of the proposed method.

2. RELATED WORK

The proposed method is related to the weakness of one-hot codes. Specifically, one-hot codes set the distance between any pair of classes the same, and do not consider the inter-class similarity within classes. To address the problem, *label smoothing* was originally proposed for computer vision [16], which improves the generalization performance of classification models. However, it still ignores the inter-class similarity. An online label smoothing method based on the statistics of the predictions was proposed in [17]. The generated soft labels reflect the inter-class similarity. However, directly applying the above methods to sound source localization makes the classification model non-convergence and yields poor performance. In this paper, we propose to jointly optimize the above two soft labels for the weakness of the one-hot coding in the end-to-end sound source localization.

3. METHOD

3.1. End-to-end sound source localization

As shown in Fig. 1, the end-to-end sound source localization with ad-hoc microphone arrays partitions a targeted room space into local areas, each of which is encoded by a one-hot code. The architecture of the method, which is shown as part of Fig. 2 that does not include the module in the red box, takes the one-hot codes of the positions of the ad-hoc nodes as spatial input features, and takes the short-term Fourier transform (STFT) from the ad-hoc nodes as acoustic features. It formulates the sound source localization problem as a classification

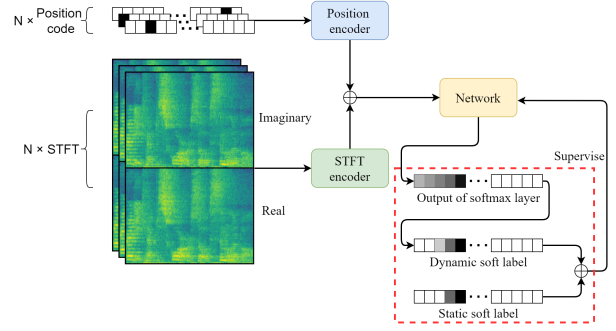


Fig. 2. The proposed end-to-end sound source localization model with soft labels. The soft label generation module is highlighted in the red box.

problem, which aims to classify a sound source into one of the grids. Because the spatial information is encoded in the one-hot codes, it could work with a special scenario where each ad-hoc node contains only a single microphone, which is the working scenario of this paper.

The original end-to-end sound source localization takes the one-hot coding to encode the room space. Suppose the set of the one-hot codes for an n -class classification problem is $\mathcal{S}^{1-hot} = \{\mathbf{S}_i^{1-hot}\}_{i=1}^n$, with the code for the i th class $\mathbf{S}_i^{1-hot} = [S_{i,1}^{1-hot}, \dots, S_{i,k}^{1-hot}, \dots, S_{i,n}^{1-hot}]$ defined as:

$$S_{i,k}^{1-hot} = \begin{cases} 1, & \text{if } k = i \\ 0, & \text{otherwise} \end{cases}, \quad \forall k = 1, \dots, n \quad (1)$$

where n is the total number of local areas.

In the following of this section, we focus on introducing the novel module highlighted in the red box in Fig. 2.

3.2. Static soft label coding

Unlike conventional classification problems where the classes may be fundamentally different and uncorrelated, the classes in our sound source localization problem are correlated. For example, we know the geometric location and similarity of the local areas, such as the distance between the centers of any two local areas. Therefore, one-hot coding may not be suitable to our problem.

Here we propose SSLC to address this issue. The SSLC for the i th class, denoted as $\mathbf{S}_i^{SSLC} = [S_{i,1}^{SSLC}, \dots, S_{i,k}^{SSLC}, \dots, S_{i,n}^{SSLC}]$, sets its k th element according to the distance between the centers of the i th and k th local areas. Formally, SSLC is defined as a Gaussian function of the distance:

$$S_{i,k}^{SSLC} = \begin{cases} \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{-d_{i,k}} e^{-\frac{\theta^2}{2\sigma^2}} d\theta, & \text{if } k \neq i \\ 1 - \sum_{\forall 1 \leq j \leq n, j \neq i} S_{j,k}, & \text{otherwise} \end{cases} \quad \forall k = 1, \dots, n \quad (2)$$

where $d_{i,k}$ is the distance between the i th and k th local areas, σ is the standard deviation that controls the spread of the Gaussian distribution. We set $\sigma = l_{ave}/\alpha_s$ where l_{ave} is the

average diagonal length of the local areas of all rooms in the training set, and α_s is a tunable hyperparameter.

Given the above formulation, we are able to train the classification model by using SSLC, i.e. $\{\mathbf{S}_i^{\text{SSLC}}\}_{i=1}^n$, as the ground-truth label instead of the one-hot coding. The training loss of a single utterance can be:

$$L^{\text{SSLC}} = - \sum_{k=1}^n S_{i,k}^{\text{SSLC}} \log \hat{y}_k \quad (3)$$

where \mathbf{x} is the input feature of the model, and \hat{y}_k is the softmax output at the k th unit.

3.3. Dynamic soft label coding

SSLC is a handcrafted label, which may not reflect the correlation between the local areas fully. To address this issue, here we propose to learn a coding book, named DSLC, via online label smoothing [17] in the training stage.

DSLCL is obtained with the guidance of the ground-truth one-hot labels in the training stage. Specifically, for the t th epoch ($t \geq 1$), suppose the softmax output of the classification model for an input utterance is $\hat{\mathbf{y}}_j = [\hat{y}_1, \dots, \hat{y}_n]$, and the ground-truth one-hot code of the speaker position of the utterance is $\mathbf{S}_i^{1\text{-hot}}$. If the speaker position is predicted correctly, i.e. $\arg \max \hat{\mathbf{y}} = i$, then we pick $\hat{\mathbf{y}}$ into a set denoted as $\mathcal{Y}_i^{(t)}$; otherwise, we discard $\hat{\mathbf{y}}$.

Then, we generate a DSLC for the i th class by averaging the elements of $\mathcal{Y}_i^{(t)}$:

$$\mathbf{S}_i^{(t)} = \frac{1}{|\mathcal{Y}_i^{(t)}|} \sum_{\hat{\mathbf{y}} \in \mathcal{Y}_i^{(t)}} \hat{\mathbf{y}}, \quad \forall i = 1, \dots, n \quad (4)$$

Finally, for the $(t+1)$ th epoch, we use DSLC generated in the t th epoch, i.e. $\{\mathbf{S}_i^{(t)}\}_{i=1}^n$, as the supervision to guide the training, where the training loss of a single utterance in this epoch is defined as:

$$L_{(t+1)}^{\text{DSLCL}} = - \sum_{k=1}^n S_{i,k}^{(t)} \log \hat{y}_k \quad (5)$$

where $S_{i,k}^{(t)}$ is the k th element of $\mathbf{S}_i^{(t)}$.

Note that, before the first epoch, i.e. $t = 0$, we use the vanilla label smoothing [16] to initialize $\{\mathbf{S}_i^{(0)}\}_{i=1}^n$.

3.4. Joint training

The reliability of DSLCL at the $(t+1)$ th epoch depends on the deep model at the t th epoch. Because the initial deep model is too weak, training DSLCL from the beginning is very difficult and sometimes unconverted.

To address this issue, we propose a joint training strategy, which uses static labels, such as SSLC or one-hot codes, to assist the training of DSLCL:

$$L_{(t)}^{\text{joint}} = \alpha_d^{(t)} L_{(t)}^{\text{DSLCL}} + (1 - \alpha_d^{(t)}) L^{\text{SSLC}} \quad (6)$$

where $\alpha_d^{(t)}$ is a tunable hyperparameter at the t th epoch.

This raises a question of how to set $\alpha_d^{(t)}$. The first choice is to set $\alpha_d^{(1)} = \alpha_d^{(2)} = \dots = \alpha_d^{(T)}$ to a constant α_d . Another choice is to adaptively adjust $\alpha_d^{(t)}$ according to the importance of $L_{(t)}^{\text{DSLCL}}$. In this paper, we set $\alpha_d^{(t)} = \text{ACC}^{(t)}$, where $\text{ACC}^{(t)}$ is the prediction accuracy of the t th model on the training set. Of course, we also can set $\alpha_d^{(t)}$ a learnable parameter of the deep model.

4. EXPERIMENTS

4.1. Experimental setup

We conducted experiments on simulated data. The source speech is from the LibriSpeech corpus [18]. The train-clean-360 subset, dev-clean subset and test-clean subset of the corpus contain 921, 40 and 40 speakers respectively. We used voice activity detection (VAD) provided by the Torchaudio module [19] to remove silent segments. The speech segments that are less than 2 seconds after VAD were discarded. For each of the remaining segments, we randomly selected a 2-second piece to generate multi-channel data with reverberation. We used the Pyroomacoustics [20] module to generate room impulse responses. The reverberation time T_{60} of all simulated speech follows a uniform distribution ranging from [0.3, 1.0] second. We further added additive noise to the reverberant speech. The additive noise was randomly selected from a large scale noise set that contains 126 hours of different types of noise segments. The signal-to-noise ratio of an utterance was selected randomly from a range of [20, 50] dB. Eventually, the number of utterances in the training, validation, and test sets are 18000, 2500, and 2500, respectively.

We generated 15 rooms of different sizes, where 10 rooms were used for training, and 5 rooms were used for testing. The length and width of the rooms range from 4m to 10m. The height of the rooms was fixed at 3m. The height of the microphone arrays and sound source was fixed at 1m. Each room was divided into 15×15 local areas in 2D. For each utterance, we first randomly picked a room, then randomly set its speech source in one of the 225 local areas, and finally randomly placed 30 ad-hoc microphone nodes in the other 224 local areas.

We used the DNN model in [15] as our backbone network. First, it uses the ad-hoc node's position code and STFT acoustic feature as the inputs for two encoders, respectively. The concatenated outputs of the two encoders are fed into a spatial-temporal attention network in the following phase. The attention network sequentially performs cross-channel attention, channel fusion, and cross-frame attention. The local region of the speaker is considered to be the average output along the time axis of the spatial-temporal attention network. ReLU was used as the activation function, and Adam was used as the optimizer to train the model.

Table 1. MAE (in meters) and ACC (%) results of the proposed five soft label coding methods and the one-hot coding.

Method	Room1		Room2		Room3		Room4		Room5		Average	
	MAE	ACC	MAE	ACC	MAE	ACC	MAE	ACC	MAE	ACC	MAE	ACC
One hot [15]	0.2135	54.6	0.2471	51.0	0.2959	46.6	0.3283	49.8	0.3207	54.0	0.2811	51.2
SSLC	0.2260	53.4	0.2566	49.2	0.3001	48.8	0.3363	43.6	0.3344	52.8	0.2907	49.6
DSLCL+One hot (α_d is a constant)	0.2226	54.2	0.2592	47.6	0.3079	43.0	0.3234	47.4	0.3225	55.4	0.2871	49.5
DSLCL+SSLC (α_d is a constant)	0.2115	57.6	0.2448	52.4	0.2873	47.0	0.3088	48.2	0.3185	56.0	0.2742	52.2
DSLCL+One hot ($\alpha_d^{(t)}$ is a variable)	0.2149	56.0	0.2689	45.0	0.3052	45.0	0.3239	49.6	0.3408	47.8	0.2907	48.7
DSLCL+SSLC ($\alpha_d^{(t)}$ is a variable)	0.2059	57.2	0.2577	47.8	0.2974	47.8	0.2998	47.4	0.3351	53.2	0.2792	50.7

Table 2. Learning error (in meters) comparison between one-hot coding and the proposed DSLCL+SSLC (α_d is a constant) method.

Room	One hot	DSLCL+SSLC	Relative reduction (%)
Room1	0.0701	0.0681	2.83
Room2	0.0942	0.0919	2.44
Room3	0.1310	0.1224	6.56
Room4	0.1484	0.1289	13.14
Room5	0.1156	0.1134	1.90
Average	0.1119	0.1049	6.26

We used mean absolute error (MAE) to evaluate the proposed method:

$$\text{MAE} = \frac{1}{I} \sum_{i=1}^I \sqrt{(x_i^{\text{spkr}} - \tilde{x}_i^{\text{spkr}})^2 + (y_i^{\text{spkr}} - \tilde{y}_i^{\text{spkr}})^2} \quad (7)$$

where I is the number of the test utterances, $(x_i^{\text{spkr}}, y_i^{\text{spkr}})$ is the ground truth coordinate of the i th speaker position, and $(\tilde{x}_i^{\text{spkr}}, \tilde{y}_i^{\text{spkr}})$ is the coordinate of the center of the local area of the predicted speaker position \tilde{y}_i .

Because the region in a local area is quantized to the center of the local area, the **quantization error** inevitably exists as an upperbound of any methods based on the classification formulation of the sound source localization. We denote the upperbound as UB-MAE. In this paper, the average UB-MAE over the 5 test rooms is 0.1693m.

4.2. Results

We trained sound source localization models with 6 label coding methods, as summarized in Table 1. From the results, we see that the proposed joint training method DSLCL+SSLC (α_d is constant) generally outperforms one-hot. In most cases, the higher the ACC, the smaller the MAE. The performance of the joint training is always better than the one-hot coding, no matter whether $\alpha_d^{(t)}$ is a constant or a variable. The schemes of adaptively adjusting α_d does not achieve better results than the schemes of setting α_d to a constant. This is probably caused by that setting $\alpha_d^{(t)} = \text{ACC}^{(t)}$ makes DSLCL overfit to the training set.

Table 3. Effect of α_s in SSLCL on performance.

α_s	2.6	2.7	2.8	2.9	3
MAE (m)	0.2793	0.2796	0.2778	0.2782	0.2822
ACC (%)	50.5	51.7	51.0	51.6	50.7

Table 4. Effect of α_d in “DSLCL+SSLC (α_d is a constant)” on performance.

α_d	0.2	0.4	0.5	0.6	0.8
MAE (m)	0.2839	0.2741	0.2726	0.2772	0.2845
ACC (%)	49.2	50.6	51.4	49.6	50.2

Because the quantization error always exist, the error in Table 1 contains two parts, one from the quantization error, and the other one from the *learning error* caused by the coding methods. We further investigated the learning error independently in Table 2 and found our proposed method achieved a relative 6.26% average reduction.

There are two tunable hyperparameters, α_s in SSLCL and α_d in “DSLCL+SSLC (α_d is a constant)”. Here we study their effects on performance. For α_s in SSLCL, when $\alpha_s < 2.6$, the probability that the speaker falls in the ground-truth local region is lower than 0.8 which results in poor performance; when $\alpha_s > 3.0$, SSLCL approaches very close to one-hot coding. Therefore, we studied $\alpha_s \in [2.6, 3.0]$ in Table 3. From the table, we see that the optimal value of α_s is 2.8. For α_d in “DSLCL+SSLC (α_d is a constant)”, we searched α_d from a range of $(0, 1)$, and find that, as shown in Table 4, the optimal value of α_d is 0.5.

5. CONCLUSION

In this paper, we propose a label smoothing algorithm against the weakness of the one-hot coding in the end-to-end sound source localization with ad-hoc microphone arrays. It includes a handcrafted SSLCL, calculated from local area center distances, and a learnable DSLCL, updated during training process using correctly classified sample statistics. Experimental results show the effectiveness of the proposed algorithm over the one-hot coding, without increasing the model complexity and training time.

6. REFERENCES

- [1] DeLiang Wang and Jitong Chen, "Supervised speech separation based on deep learning: An overview," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 10, pp. 1702–1726, 2018.
- [2] Constantin Spille, Birger Kollmeier, and Bernd T Meyer, "Comparing human and automatic speech recognition in simple and complex acoustic scenes," *Computer Speech & Language*, vol. 52, pp. 123–140, 2018.
- [3] Tim Van den Bogaert, Evelyne Carette, and Jan Wouters, "Sound source localization using hearing aids with microphones placed behind-the-ear, in-the-canal, and in-the-pinna," *International Journal of Audiology*, vol. 50, no. 3, pp. 164–176, 2011.
- [4] Pierre-Amaury Grumiaux, Sran Kitić, Laurent Girin, and Alexandre Guérin, "A survey of sound source localization with deep learning methods," *The Journal of the Acoustical Society of America*, vol. 152, no. 1, pp. 107–151, 2022.
- [5] Charles Knapp and Glifford Carter, "The generalized correlation method for estimation of time delay," *IEEE transactions on acoustics, speech, and signal processing*, vol. 24, no. 4, pp. 320–327, 1976.
- [6] Joseph H DiBiase, Harvey F Silverman, and Michael S Brandstein, "Robust localization in reverberant rooms," in *Microphone arrays*, pp. 157–180. Springer, 2001.
- [7] Ralph Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE transactions on antennas and propagation*, vol. 34, no. 3, pp. 276–280, 1986.
- [8] Soumitro Chakrabarty and Emanuël AP Habets, "Multi-speaker doa estimation using deep convolutional networks trained with noise signals," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 1, pp. 8–21, 2019.
- [9] Xiong Xiao, Shengkui Zhao, Xionghu Zhong, Douglas L Jones, Eng Siong Chng, and Haizhou Li, "A learning-based approach to direction of arrival estimation in noisy and reverberant environments," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 2814–2818.
- [10] Thi Ngoc Tho Nguyen, Woon-Seng Gan, Rishabh Ranjan, and Douglas L Jones, "Robust source counting and doa estimation using spatial pseudo-spectrum and convolutional neural network," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2626–2637, 2020.
- [11] Anthony Griffin, Anastasios Alexandridis, Despoina Pavlidi, Yiannis Mastorakis, and Athanasios Mouchtaris, "Localizing multiple audio sources in a wireless acoustic sensor network," *Signal Processing*, vol. 107, pp. 54–67, 2015.
- [12] Manuel Hahmann, Efred Fernandez-Grande, Henry Gunawan, and Peter Gerstoft, "Sound source localization using multiple ad hoc distributed microphone arrays," *JASA Express Letters*, vol. 2, no. 7, pp. 074801, 2022.
- [13] Fabio Vesperini, Paolo Vecchiotti, Emanuele Principi, Stefano Squartini, and Francesco Piazza, "Localizing speakers in multiple rooms by using deep neural networks," *Computer Speech & Language*, vol. 49, pp. 83–106, 2018.
- [14] Shupeil Liu, Yijun Gong, and Xiao-Lei Zhang, "Deep learning based two-dimensional speaker localization with large ad-hoc microphone arrays," *arXiv preprint arXiv:2210.10265*, 2022.
- [15] Yijun Gong, Shupeil Liu, and Xiao-Lei Zhang, "End-to-end two-dimensional sound source localization with ad-hoc microphone arrays," *arXiv preprint arXiv:2210.08484*, 2022.
- [16] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna, "Rethinking the inception architecture for computer vision," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2818–2826.
- [17] Chang-Bin Zhang, Peng-Tao Jiang, Qibin Hou, Yunchao Wei, Qi Han, Zhen Li, and Ming-Ming Cheng, "Delving deep into label smoothing," *IEEE Transactions on Image Processing*, vol. 30, pp. 5984–5996, 2021.
- [18] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [19] Yao-Yuan Yang, Moto Hira, Zhaoheng Ni, Artyom Astafurov, Caroline Chen, Christian Puhusch, David Pollack, Dmitriy Genzel, Donny Greenberg, Edward Z Yang, et al., "Torchaudio: Building blocks for audio and speech processing," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 6982–6986.
- [20] Robin Scheibler, Eric Bezzam, and Ivan Dokmanić, "Pyroomacoustics: A python package for audio room simulation and array processing algorithms," in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 351–355.