



Robust Training for Speaker Verification against Noisy Labels

Zhihua Fang^{1,2}, Liang He^{1,2,3,†}, Hanhan Ma^{1,2}, Xiaochen Guo^{1,2}, Lin Li⁴

¹ School of Information Science and Engineering, Xinjiang University, Urumqi 830017, China

² Xinjiang Key Laboratory of Signal Detection and Processing, Urumqi 830017, China

³ Department of Electronic Engineering, Tsinghua University, Beijing 100084, China

⁴ School of Electronic Science and Engineering, Xiamen University, Xiamen 361005, China

fangzhihua@stu.xju.edu.cn, heliang@mail.tsinghua.edu.cn

Abstract

The deep learning models used for speaker verification rely heavily on large amounts of data and correct labeling. However, noisy (incorrect) labels often occur, which degrades the performance of the system. In this paper, we propose a novel two-stage learning method to filter out noisy labels from speaker datasets. Since a DNN will first fit data with clean labels, we first train the model with all data for several epochs. Then, based on this model, the model predictions are compared with the labels using our proposed the *OR-Gate* with top- k mechanism to select the data with clean labels and the selected data is used to train the model. This process is iterated until the training is completed. We have demonstrated the effectiveness of this method in filtering noisy labels through extensive experiments and have achieved excellent performance on the VoxCeleb (1 and 2) with different added noise rates.

Index Terms: speaker verification, speaker embedding, noisy label, early learning, curriculum learning

1. Introduction

In recent years, speaker models based on deep neural networks (e.g. TDNN [1], ResNet [2], and ECAPA-TDNN [3]) have become the mainstream approaches for speaker verification, deriving many variants [4, 5, 6, 7] with excellent performance. The success of these models depends on large-scale labelled datasets [8, 9]. Unfortunately, high-quality labelled large-scale data is expensive in practice, and in addition, many large-scale datasets are collected by web crawlers [10], which inevitably contains noisy labels [11].

In the presence of noisy labels in the dataset, the loss function contains wrong labels, which leads to the model to gradient descent towards the wrong direction. Noisy labels will hinder the model to learn stable speaker features. Therefore, to learn more accurate speaker features from the data with noisy labels, we must design a robust training strategy for noisy labels.

Researchers have conducted numerous studies on image recognition with noisy labels, and these works can be roughly divided into four categories [12]: robust architecture [13, 14], robust regularization [15, 16], robust loss function [17, 18], and sample selection [19, 20]. Learning with noisy labels has been well studied in the field of images, but less in the field of speaker recognition [21, 22, 23], so it is necessary to design a speaker recognition method that is robust to noisy labels.

When the training data contains noisy labels, deep neural networks have been observed to fit the training data first with

clean labels in the early stages, and then with incorrect labels [24, 25, 26]. This approach of fitting clean data first and then noisy data is similar to the concept of curriculum learning [27, 28], which involves gradually fitting training data from easy to difficult in order to improve the performance and generalization ability of neural networks. Our work has also been inspired by these studies.

In this paper, we propose a novel and effective sample selection method for filtering the noisy labels from speaker datasets. Our approach is divided into two stages: Stage I is early learning, we use all the data to train the model for a few epochs, a speaker model with basic discriminatory ability can be obtained. Stage II is self-confident learning, which is trained on the model obtained in stage I. So-called self-confident learning is based on trusting the model and comparing the model predictions with the labels to select clean labels. Specifically, at each epoch, we compare the sample labels with the predictions of all past epoch models for these samples, and consider the label to be clean (reliable) as long as the match is successful once. Since this mechanism is similar to the or gate in logic circuits, we name this approach as *OR-Gate*. All data are divided into clean and noisy data by the *OR-Gate*, and we only use the data with clean labels to update the network parameters, while the data with noisy labels are only forward propagated to get the predictions to be used in the next epoch, and iterate the process until the training is finished. In summary, the main contributions of this work are as follows:

- We propose a two-stage learning method for filtering noisy labels in speaker datasets, and filling the gap of no sample selection method for speaker recognition with noisy labels.
- We propose a novel and effective sample selection method called the *OR-Gate*. Since the prediction results of hard samples may not match the labels, we use the top- k mechanism to compare predictions and labels.
- We conducted experiments on VoxCeleb 1 and 2 [29, 30] by manually adding various proportions of noisy labels. The experiments demonstrate that our method achieves excellent performance on both the original data and the data with various noise rates. We also evaluate the results of the *OR-Gate* dividing noisy and non-noisy labels which proves the effectiveness of the *OR-Gate* filtering noisy labels.

2. Preliminaries

Notation. For the speaker verification task, let c be the number of speakers and e be a one-hot vector with dimension of c . $D = \{(x_i, y_i)\}_{i=1}^n$ denotes the i.i.d. samples and corresponding ground-truth labels, where n is the number of utterances. $\tilde{D} = \{(x_i, \tilde{y}_i)\}_{i=1}^n$ is the dataset where the labels are corrupted and the proportion of noisy labels is η .

This work was supported by the National Key R&D Program of China under Grant No. 2022ZD0115801. Code and data available: <https://github.com/PunkMale/OR-Gate>

[†] Corresponding author.

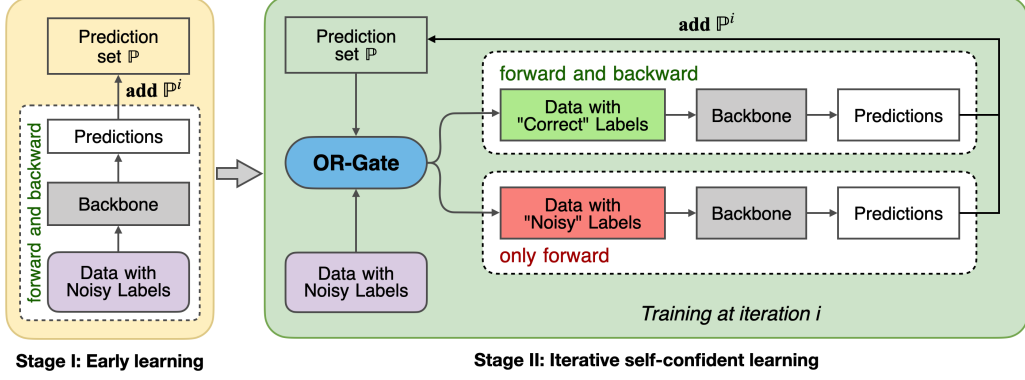


Figure 1: **The framework of the proposed two-stage learning.** **Stage I:** Train the network with all data for a few epochs, and store the model’s predictions for each sample at all epochs into the prediction set \mathbb{P} . **Stage II:** Matching all data with the prediction set \mathbb{P} , the data is divided into data with correct labels and data with noisy labels using the *OR-Gate* (see Figure 2 for details). Only the data with correct labels are used for training, and for data with noisy labels only the model predictions are output, then the predictions of all samples are added to the prediction set \mathbb{P} . Iterate the above process until the training is completed. For details, see Algorithm 1.

2.1. Problem Setup

There are two types of noisy labels: symmetric noise and asymmetric noise. Symmetric noise is generated by randomly flipping the sample labels with equal probability to incorrect labels of other classes, as shown in equation 1. Asymmetric noise is the flipping of sample labels to similar class labels. The manually added noise labels in the speaker dataset are only applicable to symmetric noise, and the existing speaker recognition with noisy labels only focused on symmetric noise [21, 22], so this work is based on symmetric noise.

$$P[\tilde{y} = e_j | y = e_i] = \frac{\eta}{c-1} \quad \forall i \neq j \quad (1)$$

Let the proportion of noisy labels be η . The aim is to learn a robust speaker recognition model without knowing the proportion of noisy labels as well as the ground-truth labels.

2.2. Learning with Noisy Label

Our work is inspired by SELF [20], which is a sample selection method of image recognition with noisy labels. The model attempts to identify correct labels progressively by self-forming ensembles of models and predictions. It uses a moving-average of ensemble models and predictions to improve filtering decisions, where the ensemble models are actually a vector \bar{p} . The moving-average is updated by the equation:

$$\bar{P}_k^j = \alpha \bar{P}_k^{j-1} + (1 - \alpha) P_k^j \quad (2)$$

whereby \bar{P}_k^j depicts the moving-average of the t -th sample at epoch j , α is a momentum that represents the confidence in the model prediction as training progresses, P_k^j is the model prediction of the t -th sample in epoch j . Briefly, P is the combined prediction of the model on the samples. By comparing the sample label y_i with the corresponding P_i , we can determine whether the label y_i is clean, and thus filter out data with clean labels for training.

The method is applied to datasets with a small number of labels, while the speaker datasets have a massive number of labels, so it is difficult to correctly classify the ground-truth label samples at the early stage of training, resulting in the inability to filter out a sufficient number of samples for learning, which affects the later training.

3. The Proposed Approach

For the problem of noisy labels in the speaker dataset, we propose a two-stage learning approach and maintain a prediction set \mathbb{P} for filtering noisy labels during the whole training process. The details are as follows.

3.1. Stage I: Early Learning

Since the number of labels in the speaker dataset is huge, it is difficult for the model to correctly predict the sample labels at the early stage of training, so the early model does not have the ability to filter the samples. Fortunately, although the neural network will eventually fit all labels, it will first fit the data with clean labels.

Therefore, we first train the model on all data for a few (hyperparameter w) epochs to get a model with basic discriminative speaker ability. This stage is called *Early Learning*, in which the sample predictions for each epoch are stored in the prediction set \mathbb{P} .

3.2. Stage II: Self-Confident Learning

The model has the basic ability to recognize the speaker after early learning, so the model predictions will be mostly close to the ground-truth labels. However, since the network parameters fluctuate during the training process, it is difficult to accurately predict the ground-truth labels of the samples in the early phase. This leads to the inability to select adequate samples for training, which further leads to the inability of the network to learn sufficiently, and this influence will be iterated.

To solve the above problem, we propose the *OR-Gate*. Similar to the or gate in logic circuits, the sample label is compared with the predictions of the model for this sample in each past epoch, and the label is considered to be clean label as long as there is one successful match. The selected data with clean labels are later used for model training, which is called self-confident learning.

Hard samples often influence the decision boundary in classification tasks, and the model predictions of hard samples generally do not match the ground-truth labels. To make the *OR-Gate* select more hard samples, we use the top- k mechanism in matching the labels to the predictions, i.e., the matching is con-




Figure 2: **The OR-Gate mechanism.** The picture represents the decision process of the t -th sample at the i -th epoch. \mathbb{I} is the indicator function and \mathbb{P}_i^k represents the set of labels for the top- k prediction of the sample x_t at the i -th epoch model. When the output $_t$ is 1, the label y_t is correct, otherwise, it is noisy.

sidered successful as long as the label y_i exists in the label set \mathbb{P}_i corresponding to the top- k probabilities of the predictions. The details are shown in Figure 2.

Algorithm 1: The proposed two-stage learning algorithm.

Input: training data $\tilde{D} = \{(x_i, \tilde{y}_i)\}_{i=1}^n$,
hyperparameters w, k

Output: speaker model

- 1 Initialize $\mathbb{P} = \emptyset, i = 0$
- // Stage I: Early Learning
- 2 **while** $i++ < w$ **do**
- 3 **for** (x_t, \tilde{y}_t) **in** \tilde{D} **do**
- 4 $P_t \leftarrow \text{train}(\text{Model}, x_t, \tilde{y}_t)$
- 5 Construct the label set \mathbb{P}_t^k containing the top- k
probabilities of P_t
- 6 $\mathbb{P} = \mathbb{P} \cup \mathbb{P}_t^k$
- // Stage II: Self-Confident Learning
- 7 **while** $i++ < \text{max_epoch}$ **do**
- 8 **for** (x_t, \tilde{y}_t) **in** \tilde{D} **do**
- 9 // The OR-Gate mechanism
- 10 **if** $\sum_{j=1}^{k-1} \mathbb{I}(\tilde{y}_t \in \mathbb{P}_t^j) > 0$ **then**
- 11 // Forward and backward
- 12 $P_t \leftarrow \text{train}(\text{Model}, x_t, \tilde{y}_t)$
- 13 **else**
- 14 // Forward only
- $P_t \leftarrow \text{Model}(x)$
- Construct the label set \mathbb{P}_t^k containing the top- k
probabilities of P_t
- $\mathbb{P} = \mathbb{P} \cup \mathbb{P}_t^k$

4. Experiments

4.1. Experimental Details

Data. We demonstrated the superiority of our proposed method by conducting comprehensive experiments on VoxCeleb1 and 2 with different proportions of noisy labels. VoxCeleb1 [29] contains 1211 speakers and 148,642 utterances, and VoxCeleb2 [30] contains 5994 speakers and 109,2009 utterances. Their labels were manually checked and can be considered clean datasets ($\eta = 0$).

To verify the noisy label robustness of the method and consider real-world scenarios, we set the noisy label ratio η to 0%, 5%, 10%, 20%, 30%, and 50%. Specifically, for a given

noise rate η , we randomly select the corresponding proportion of samples from VoxCeleb 1 or 2, flip their labels with random equal probability to other error labels, and tag these error samples (these tags are not visible when selecting the samples), and these noise tags can be used in the second stage to evaluate how correctly we partition the data.

Implementation Details. For all experiments, we use ResNet34 as the backbone, attentive statistics pooling [31] and AM-Softmax [32] respectively as the pooling layer and loss function. The input of the network is an 80-dimensional Fbank, the embedding dimension is 512 and the mini-batch size is set to 128. The network parameters are optimized by the Adam optimizer with an initial learning rate of 0.0002, which is reduced to 40% every 5 epochs. The backend uses cosine distance to score.

For VoxCeleb 1, we set $w = 5, k = 90$ and trained 80 epochs. Since the data volume of VoxCeleb 2 is much larger than that of VoxCeleb 1 and the number of labels is about 5 times that of VoxCeleb 1, therefore the early learning time of VoxCeleb 2 can be shorter and the value of k can be set larger. For VoxCeleb 2, we set $w = 4, k = 400$, and trained 60 epochs.

Performance Metrics. We compare **EER** on test data with clean labels, and since our idea is to filter noisy labels, so we also evaluate the **precision** (# clean labels selected / # of selected labels) and **recall** (# clean labels selected / # of clean labels in the data) of identifying noisy labels.

4.2. Results of comparison experiments

Performance on VoxCeleb 1. Table 1 shows the results of our experiments on VoxCeleb 1. It can be seen that the performance of the baseline deteriorates rapidly as the noise rate increases, demonstrating that noisy labels have a serious negative impact on speaker verification. Secondly, when the training data is without noisy labels, our method outperforms the baseline because our method follows the idea of curriculum learning [27, 33]. Since easy samples are easier to predict, they will be selected first by the *OR-Gate*. This results in the network learning easy samples first and then hard samples, thus outperforming the baseline. Finally, our method outperforms all other methods when training data with noisy labels.

Table 1: *EER(%) comparisons on VoxCeleb1 with different proportions of noisy labels added.*

Test Data	Vox-O (original test pairs)						
	η	0%	5%	10%	20%	30%	50%
Baseline		4.36	5.30	6.24	7.99	9.78	14.39
SELF [20] (Re-implemented)		5.52	5.94	5.90	6.34	7.50	11.90
+ Cosine		4.64	4.83	5.11	5.45	5.56	6.32
LNCL+Sub-AM [22]	+ PLDA	3.92	4.45	5.02	5.47	5.96	7.29
+ NL-PLDA		3.92	4.32	4.63	4.85	5.04	5.65
OR-Gate		4.08	4.07	4.22	4.28	4.41	5.53

Performance on VoxCeleb 2. We further conducted experiments on VoxCeleb 2, which contains more speakers and utterances and is relatively difficult to predict, as shown in Table 2. The results show that the proposed method still shows excellent performance for large-scale data containing noisy labels, and has a notable advantage at high noise rates, keeping the EER within 2% even at a noise rate η of 50%.

We also observed a special phenomenon when the training data contained few noisy labels, using our method resulted in better performance of the model. This shows that when the

Table 2: *EER(%) comparisons on VoxCeleb2 with different proportions of noisy labels added.*

Test Data		Vox-O (original test pairs)						Vox-H (hard test pairs)					
η		0%	5%	10%	20%	30%	50%	0%	5%	10%	20%	30%	50%
LNCL+Sub-AM [22]	Baseline	1.69	1.72	1.90	2.21	2.88	4.32	2.89	3.11	3.25	3.64	4.42	6.74
	+ Cosine	1.63	1.63	1.64	1.69	1.80	2.25	3.02	3.05	3.06	3.26	3.26	3.77
	+ PLDA	1.71	1.78	1.81	2.05	2.15	2.86	3.03	3.10	3.15	3.57	3.76	4.76
	+ NL-PLDA	1.70	1.72	1.73	1.75	1.77	2.19	3.02	3.07	3.10	3.20	3.33	3.58
OR-Gate		1.64	1.65	1.62	1.67	1.72	1.97	2.93	2.71	2.77	2.96	2.87	3.11

Table 3: *Precision and Recall of selecting clean labels at different periods of training.*

Metric	Precision						Recall					
η	0%	5%	10%	20%	30%	50%	0%	5%	10%	20%	30%	50%
VoxCeleb 1												
epoch ₆	1.0000	0.9999	0.9999	0.9999	0.9999	1.0000	0.9523	0.9373	0.9161	0.8298	0.7018	0.3054
epoch ₁₀	1.0000	0.9998	0.9997	0.9996	0.9992	0.9974	0.9857	0.9841	0.9800	0.9603	0.9250	0.6810
epoch ₂₀	1.0000	0.9997	0.9995	0.9989	0.9980	0.9928	0.9942	0.9940	0.9933	0.9897	0.9780	0.8123
epoch ₅₀	1.0000	0.9995	0.9990	0.9981	0.9963	0.9876	0.9975	0.9975	0.9971	0.9961	0.9916	0.8586
epoch ₈₀	1.0000	0.9995	0.9988	0.9976	0.9953	0.9847	0.9980	0.9982	0.9978	0.9969	0.9929	0.8598
VoxCeleb 2												
epoch ₅	1.0000	0.9999	0.9999	0.9998	0.9999	0.9999	0.9774	0.9743	0.9706	0.9545	0.9202	0.7093
epoch ₁₀	1.0000	0.9999	0.9997	0.9994	0.9992	0.9979	0.9842	0.9837	0.9829	0.9806	0.9762	0.9369
epoch ₂₀	1.0000	0.9998	0.9996	0.9991	0.9986	0.9961	0.9865	0.9862	0.9858	0.9845	0.9818	0.9499
epoch ₄₀	1.0000	0.9997	0.9995	0.9988	0.9980	0.9945	0.9880	0.9878	0.9876	0.9866	0.9841	0.9552
epoch ₆₀	1.0000	0.9997	0.9994	0.9986	0.9977	0.9937	0.9888	0.9886	0.9884	0.9873	0.9850	0.9569

scale of the data is not very different, containing a few noisy labels allows the model to obtain better generalization in early learning and thus to follow the guidelines of curriculum learning (selecting clean samples for training from easy to hard) in the self-confident learning stage.

4.3. Efficacy of detecting clean samples.

Although the EER results on noise-added VoxCeleb 1 and 2 demonstrate the validity of the *OR-Gate*, it is also meaningful to evaluate the precision and recall of the selected samples because of the sample selection method we used.

Precision analysis. Precision can reflect the ability of our method to filter noisy labels. As can be seen from Table 3, the highest precision (almost 1) was achieved on both datasets after early learning to divide the data, demonstrating that early learning to first fit the data with clean labels is what makes the data filtered after early learning almost free of noisy labels. Although the precision decreases during the self-confident learning period, the precision remains above 98% throughout the training period until the end, and the tiny percentage of noisy labels has minimal impact on the model.

Recall analysis. Recall reflects the ability of our method to select clean samples. As can be seen from Table 3, the *OR-Gate* can select a large portion of clean data after early learning, except for the case of the small dataset (VoxCeleb 1) with high noise rate (50%). This proves that a model with basic speaker recognition capability can be obtained after early learning, which proves that our idea is right. When the noise rate is below 50%, the *OR-Gate* selects more than 98% of clean samples by self-confident learning, and when the noise rate is as high as 50%, the *OR-Gate* also selects the majority of clean samples. The above data demonstrate the *OR-Gate* can select as many clean samples as possible while filtering noisy labels.

4.4. Ablation experiments

We also conducted ablation experiments. Table 4 shows that the performance of the model without early learning is much worse than the baseline, which indicates that early learning is essential in our approach. It is only through early learning that a speaker model with basic recognition ability can be obtained for "self-confident" learning of Stage II. In addition, we can see that adding the top-*k* mechanism to the *OR-Gate* can improve the performance more than not adding it, which indicates that the top-*k* mechanism can help the model select more hard samples to improve the classification ability of the model and get higher quality speaker embedding.

Table 4: *EER (%) comparisons of the ablation experiments on VoxCeleb 1.*

Test Data	Vox-O (original test pairs)					
η	0%	5%	10%	20%	30%	50%
Baseline	4.36	5.30	6.24	7.99	9.78	14.39
OR-Gate w/o early learning	5.99	6.27	6.87	8.37	13.05	25.83
OR-Gate w/o top- <i>k</i> mechanism	5.21	4.81	5.15	5.58	6.17	8.55
OR-Gate	4.08	4.07	4.22	4.28	4.41	5.53

5. Conclusions

In this paper, we proposed a novel and easy-to-implement framework for filtering noisy labels in speaker datasets. Specifically, a model with basic speaker discrimination ability is first obtained by early learning, and then self-confident learning is conducted based on this model, where the network is trained using our proposed the *OR-Gate* with the top-*k* mechanism to select clean data. Excellent performance was achieved on both VoxCeleb 1 and 2 with various noise ratios added.

6. References

- [1] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, “X-vectors: Robust dnn embeddings for speaker recognition,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5329–5333.
- [2] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [3] B. Desplanques, J. Thienpondt, and K. Demuyne, “ECAPA-TDNN: Emphasized Channel Attention, Propagation and Aggregation in TDNN Based Speaker Verification,” in *Proc. Interspeech 2020*, 2020, pp. 3830–3834. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2020-2650>
- [4] Z.-K. Wan, Q.-H. Ren, Y.-C. Qin, and Q.-R. Mao, “Statistical pyramid dense time delay neural network for speaker verification,” in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 7532–7536.
- [5] B. Han, Z. Chen, B. Liu, and Y. Qian, “Mlp-svnet: A multi-layer perceptrons based network for speaker verification,” in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 7522–7526.
- [6] Y. Zhang, Z. Lv, H. Wu, S. Zhang, P. Hu, Z. Wu, H. yi Lee, and H. Meng, “MFA-Conformer: Multi-scale Feature Aggregation Conformer for Automatic Speaker Verification,” in *Proc. Interspeech 2022*, 2022, pp. 306–310.
- [7] B. Han, Z. Chen, and Y. Qian, “Self-Supervised Speaker Verification Using Dynamic Loss-Gate and Label Correction,” in *Proc. Interspeech 2022*, 2022, pp. 4780–4784.
- [8] A. Nagrani, J. S. Chung, W. Xie, and A. Zisserman, “Voxceleb: Large-scale speaker verification in the wild,” *Computer Speech & Language*, vol. 60, p. 101027, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0885230819302712>
- [9] L. Li, R. Liu, J. Kang, Y. Fan, H. Cui, Y. Cai, R. Vipperla, T. F. Zheng, and D. Wang, “Cn-celeb: Multi-genre speaker recognition,” *Speech Communication*, vol. 137, pp. 77–91, 2022.
- [10] Y. Guo, L. Zhang, Y. Hu, X. He, and J. Gao, “Ms-celeb-1m: A dataset and benchmark for large-scale face recognition,” in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds. Cham: Springer International Publishing, 2016, pp. 87–102.
- [11] E. Fonseca, M. Plakal, D. P. W. Ellis, F. Font, X. Favory, and X. Serra, “Learning sound event classifiers from web audio with noisy labels,” in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 21–25.
- [12] H. Song, M. Kim, D. Park, Y. Shin, and J.-G. Lee, “Learning from noisy labels with deep neural networks: A survey,” *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–19, 2022.
- [13] J. Yao, J. Wang, I. W. Tsang, Y. Zhang, J. Sun, C. Zhang, and R. Zhang, “Deep learning from noisy image labels with quality embedding,” *IEEE Transactions on Image Processing*, vol. 28, no. 4, pp. 1909–1922, 2019.
- [14] K. Lee, S. Yun, K. Lee, H. Lee, B. Li, and J. Shin, “Robust inference via generative classifiers for handling noisy labels,” in *ICML*, 09–15 Jun 2019, pp. 3763–3772. [Online]. Available: <https://proceedings.mlr.press/v97/lee19f.html>
- [15] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, “mixup: Beyond empirical risk minimization,” in *International Conference on Learning Representations*, 2018.
- [16] H. Wei, L. Tao, R. XIE, and B. An, “Open-set label noise can improve robustness against inherent label noise,” in *Advances in Neural Information Processing Systems*, vol. 34. Curran Associates, Inc., 2021, pp. 7978–7992. [Online]. Available: <https://proceedings.neurips.cc/paper/2021/file/428fca9bc1921c25c5121f9da7815cde-Paper.pdf>
- [17] Y. Wang, X. Ma, Z. Chen, Y. Luo, J. Yi, and J. Bailey, “Symmetric cross entropy for robust learning with noisy labels,” in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 322–330.
- [18] L. Feng, S. Shu, Z. Lin, F. Lv, L. Li, and B. An, “Can cross entropy loss be robust to label noise?” in *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, 7 2020, pp. 2206–2212.
- [19] J. Li, R. Socher, and S. C. Hoi, “Dividemix: Learning with noisy labels as semi-supervised learning,” in *International Conference on Learning Representations*, 2020.
- [20] D. T. Nguyen, C. K. Mummadi, T. P. N. Ngo, T. H. P. Nguyen, L. Beggel, and T. Brox, “Self: Learning to filter noisy labels with self-ensembling,” in *International Conference on Learning Representations*, 2020.
- [21] L. Li, F. Tong, and Q. Hong, “When speaker recognition meets noisy labels: Optimizations for front-ends and back-ends,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 1586–1599, 2022.
- [22] F. Tong, Y. Liu, S. Li, J. Wang, L. Li, and Q. Hong, “Automatic Error Correction for Speaker Embedding Learning with Noisy Labels,” in *Proc. Interspeech 2021*, 2021, pp. 4628–4632.
- [23] B. J. Borgström and P. Torres-Carrasquillo, “Bayesian estimation of plda with noisy training labels, with applications to speaker verification,” in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7594–7598.
- [24] S. Liu, K. Liu, W. Zhu, Y. Shen, and C. Fernandez-Granda, “Adaptive early-learning correction for segmentation from noisy annotations,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 2606–2616.
- [25] Y. Bai, E. Yang, B. Han, Y. Yang, J. Li, Y. Mao, G. Niu, and T. Liu, “Understanding and improving early stopping for learning with noisy labels,” in *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., vol. 34. Curran Associates, Inc., 2021, pp. 24392–24403.
- [26] Y. Xue, K. Whitecross, and B. Mirzasoleiman, “Investigating why contrastive learning benefits robustness against label noise,” in *Proceedings of the 39th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, Eds., vol. 162. PMLR, 17–23 Jul 2022, pp. 24851–24871.
- [27] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, “Curriculum learning,” in *Proceedings of the 26th Annual International Conference on Machine Learning*, ser. ICML ’09. New York, NY, USA: Association for Computing Machinery, 2009, p. 41–48. [Online]. Available: <https://doi.org/10.1145/1553374.1553380>
- [28] S. Ranjan and J. H. L. Hansen, “Curriculum learning based approaches for noise robust speaker recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 1, pp. 197–210, 2018.
- [29] A. Nagrani, J. S. Chung, and A. Zisserman, “VoxCeleb: A Large-Scale Speaker Identification Dataset,” in *Proc. Interspeech 2017*, 2017, pp. 2616–2620.
- [30] J. S. Chung, A. Nagrani, and A. Zisserman, “VoxCeleb2: Deep Speaker Recognition,” in *Proc. Interspeech 2018*, 2018, pp. 1086–1090.
- [31] K. Okabe, T. Koshinaka, and K. Shinoda, “Attentive Statistics Pooling for Deep Speaker Embedding,” in *Proc. Interspeech 2018*, 2018, pp. 2252–2256.
- [32] F. Wang, J. Cheng, W. Liu, and H. Liu, “Additive margin softmax for face verification,” *IEEE Signal Processing Letters*, vol. 25, no. 7, pp. 926–930, 2018.
- [33] M. P. Kumar, B. Packer, and D. Koller, “Self-paced learning for latent variable models,” in *Proceedings of the 23rd International Conference on Neural Information Processing Systems - Volume 1*, ser. NIPS’10. Red Hook, NY, USA: Curran Associates Inc., 2010, p. 1189–1197.