Improvement of Serial Approach to Anomalous Sound Detection by Incorporating Two Binary Cross-Entropies for Outlier Exposure

Ibuki Kuroyanagi*, Tomoki Hayashi*[†], Kazuya Takeda*, Tomoki Toda*,

*Nagoya University, Nagoya, Japan

[†]Human Dataware Lab. Co., Ltd., Nagoya, Japan *{kuroyanagi.ibuki, hayashi.tomoki}@g.sp.m.is.nagoya-u.ac.jp, [†]hayashi@hdwlab.co.jp, *takeda@i.nagoya-u.ac.jp, *tomoki@icts.nagoya-u.ac.jp

Abstract-Anomalous sound detection systems must detect unknown, atypical sounds using only normal audio data. Conventional methods use the serial method, a combination of outlier exposure (OE), which classifies normal and pseudo-anomalous data and obtains embedding, and inlier modeling (IM), which models the probability distribution of the embedding. Although the serial method shows high performance due to the powerful feature extraction of OE and the robustness of IM, OE still has a problem that doesn't work well when the normal and pseudoanomalous data are too similar or too different. To explicitly distinguish these data, the proposed method uses multi-task learning of two binary cross-entropies when training OE. The first is a loss that classifies the sound of the target machine to which product it is emitted from, which deals with the case where the normal data and the pseudo-anomalous data are too similar. The second is a loss that identifies whether the sound is emitted from the target machine or not, which deals with the case where the normal data and the pseudo-anomalous data are too different. We perform our experiments with DCASE 2021 Task 2 dataset. Our proposed single-model method outperforms the top-ranked method, which combines multiple models, by 2.1 % in AUC.

Index Terms—anomalous sound detection, outlier exposure, inlier modeling, hypersphere, multi-task learning

I. INTRODUCTION

All machines in factories, plants, office buildings, etc., require regular maintenance to keep them functioning properly, and they can also break down or malfunction. It is important to quickly address such equipment problems to prevent damage to the machine, or serious accidents. In the past, skilled maintenance technicians would monitor the operation of machines and diagnose their condition by listening to them. However, with the decrease in the rapidly aging working population, providing quality maintenance services with fewer skilled workers is becoming more challenging. Furthermore, automated factories and plants are also becoming more common [1]. In response to these trends, methods for automatically detecting anomalous sounds have been developed [2], [3].

Anomalous sound detection (ASD) is the task of identifying whether the sound emitted by a target machine is normal or anomalous. However, ASD is very different from simple, binary classification problems [4] because it is difficult to collect data on every possible anomalous sound. These sounds rarely occur during the normal operation, and the possible types of anomalous sounds are very diverse. Therefore, it is more practical to detect unknown, anomalous sounds using only normal sounds [5].

Currently, two types of ASD approaches are mainly used: inlier modeling (IM) and outlier exposure (OE). IM is a method that models the probability distribution of normal data and detects data that does not correspond to the model as anomalous data. IM methods such as autoencoders [6]-[8], local outlier factor (LOF) [9], gaussian mixture models (GMM) [10], [11], normalizing flows [12], [13] have been used. IM is robust, but it is difficult to extract effective features. In contrast, OE is a method for learning the decision boundaries of normal data by classifying normal and pseudoanomalous data [14], [15]. OE methods such as deep semisupervised anomaly detection [16], [17], and deep double centroids semi-supervised anomaly detection (DDCSAD) [18] have been used. OE is easy to extract effective features, but it is not robust. It also does not work well when normal and pseudo-anomalous data are too similar or too different [19].

Recent studies have proposed methods that use a combination of IM and OE methods, and these approaches have achieved high ASD performance [19]. One such approach, called the parallel method, ensembles anomaly scores for both IM and OE to compensate each other for the weakness of IM and OE [12], [20]. However, the parallel method requires multiple models, which use different training processes, to be created to obtain combined anomaly scores, increasing the cost of system development and maintenance. Another combination approach, called the serial method, uses OE- and IM-based methods in series. The serial method consists of the following three steps: First, we train the embedding with an OE method classification task using normal and pseudo-anomalous data. Next, using the embedding extracted by the OE method in Step 1, an IM method trains a normal data distribution. Finally, data that differs from the distribution constructed in Step 2 is detected as anomalous. The Serial method solves the difficulty of IM feature extraction and the robustness of OE [21]-[23]. Still, it leaves the problem of when the normal and pseudoanomalous data of OE are too similar or too different.

To explicitly distinguish between normal and pseudoanomalous data that are too similar or too different, the proposed method uses multi-task learning of two binary crossentropies when training OE in the serial method. Here, the datasets [24], [25] contain several machine types such as fan, pump, valve, etc. Each machine type has K product IDs such as ID 0, 1, 2, etc. The first is a loss that classifies the sound of the target machine type to which product ID it is emitted from, which deals with the case where the normal data and the pseudo-anomalous data are too similar. The second is a loss that identifies whether the sound is emitted from the target machine type or not, which deals with the case where the normal data and the pseudo-anomalous data are too different. The proposed method overcomes the weaknesses of OE. Furthermore, the proposed method of OE is easy to develop and manage since one model corresponds to one machine type, and there is no need to create a model for each product ID [26]. We evaluate the performance of the proposed method by conducting experiments and showing its effectiveness. Also, we visualize the embedding by t-SNE [27] and qualitatively discuss the effect of using the serial method and the change of embedding by the proposed method.





II. OE-BASED METHOD

To improve the OE component of our serial method, we focus on DDCSAD [18], which has demonstrated high ASD performance. DDCSAD is trained using multi-task binary classification and metric learning. A binary classification model is trained using the target product ID of the target machine type as positive samples, and the other product IDs and the other machine types as negative samples. In addition, we define the centroid for positive and negative samples and perform metric learning to minimize within-class variance and maximize between-class variance for each class. During inference, we use the probability calculated using binary classification and the weighted average of the distance between the centroid of the positive samples and the embedding as the anomaly score.

Fig. 1a shows an image of an embedding obtained using DDCSAD. In this method, the pseudo-anomalous data move away from the centroid of the positive sample and move closer to the negative sample, which is considered robust when the distributions of normal and anomalous data are very different. However, there are several problems with this

method. First, when the distributions of normal and pseudoanomalous data are similar, the performance of this method is degraded because the decision boundaries of normal and pseudo-anomalous data are different from those of normal and anomalous data. Second, when using Euclidean distance, the embedding of normal data is assumed to follow a normal distribution. However, if the model is not expressive enough, the above assumption cannot be satisfied, and performance is degraded. Furthermore, since a model is created for each product ID, the performance variation of the model becomes large. In addition, the development and maintenance of the model become complicated.

III. PROPOSED METHOD

Fig. 1b shows an image of an embedding obtained by our proposed method. Fig. 2 shows an overview of the proposed method. To explicitly distinguish between normal and pseudo-anomalous data that are too similar or too different, the proposed method uses multi-task learning of two binary cross-entropies when training OE. The first is a product ID classification loss, which deals with cases where the normal and pseudo-anomalous data are too similar. The second is a machine type classification loss, which deals with cases where the normal and pseudo-anomalous data are too different.

If the audio input is represented as x_i (i = 1, 2, ..., N), the machine type is represented as t_i (i = 1, 2, ..., N), where t_i is 1 for the target machine type and 0 for the other machine types. Each machine type has K product IDs, and x_i belongs to one of them. The one-hot vector for the product ID is represented as $y_{\{i,k\}}$ (i = 1, 2, ..., N, k = 1, 2, ..., K), where $y_{\{i,k\}}$ is 1 for the k th element and 0 for the other elements when the product ID is k. Product IDs classification loss can be represented as:

$$\mathcal{L}_{\text{product}} = -\frac{1}{K \sum_{i=1}^{N} t_i} \sum_{i=1}^{N} \sum_{k=1}^{K} t_i \left\{ y_{\{i,k\}} \log \left(g_{\text{product}} \left(f(x_i) \right) \right) + (1 - y_{\{i,k\}}) \log \left(1 - g_{\text{product}} \left(f(x_i) \right) \right) \right\},$$
(1)

where f is an encoder, and g_{product} is a linear transformation. Machine type classification loss can be represented as:

$$\mathcal{L}_{\text{machine}} = -\frac{1}{N} \sum_{i=1}^{N} \left\{ t_i \log \left(g_{\text{affine}} \left(||f(x_i)||^2 \right) \right) + (1 - t_i) \log \left(1 - g_{\text{affine}} \left(||f(x_i)||^2 \right) \right) \right\},$$
(2)

where g_{affine} is an affine transformation. When creating minibatches, we use a batch sampler so that the value of t is 1:1. The final combined loss function is:

$$\mathcal{L} = \mathcal{L}_{\text{machine}} + \lambda \mathcal{L}_{\text{product}}, \qquad (3)$$

where λ is a hyperparameter. Mixup [28] is applied in minibatches to obtain intermediate features between normal and pseudo-anomalous data, and each sample applies Eq. 1 if the target machine type is included. In other words, we identify the product IDs of the target machine type but not the product IDs of the other machine types. The method for calculating the anomaly score a_i (i = 1, 2, ..., N) is:

$$a_{i} = \mathcal{A}\left(h\left(f\left(\mathcal{X}_{i}\right)\right)\right), \qquad (4)$$

where \mathcal{X}_i is a set of S segments that divide x_i into T seconds, allowing for overlap, h is a post-processing function for IM such as GMM or LOF, and \mathcal{A} is the aggregator of the anomaly scores such as max/average pooling.

By using the product IDs classification, anomalous sounds that are similar to normal sounds are distributed around each product ID. By using the machine type classification by the norm of embedding, anomalous sounds that differ significantly from normal sounds are collected near the center of the hypersphere. This approach is inspired by DDCSAD's training method, which involves collecting pseudo-anomalous data into a single point. Since the proposed method explicitly distinguishes between normal and pseudo-anomalous data that are too similar or too different, it avoids the insufficient expressiveness of the OE embedding. In other words, the proposed method obtains an embedding suitable for detecting normal or similar sounds. In addition, it is easy to develop and maintain the system because the model is created for each machine type, not for each product ID in each machine type.



Fig. 2: Overview of our proposed method.

TABLE I: Hyperparameters for each machine type.LR=learning rate, BS=batch size

	fan	gearbox	pump	valve	slider	ToyCar	ToyTrain
LR	0.001	0.001	0.001	0.0005	0.001	0.001	0.0005
BS	32	128	128	128	128	128	32
λ	0.1	10.0	10.0	10.0	10.0	10.0	0.1
h	GMM	GMM	GMM	GMM	GMM	LOF	LOF
p	16	64	2	32	2	16	8

IV. EXPERIMENTAL EVALUATION

A. Experimental conditions

To evaluate the performance of the proposed method, we conducted experiments using the data from the DCASE 2021



Fig. 3: Visualizations of various embeddings

Task 2 Challenge (MIMII Due [24], ToyADMOS2 [25]). The training and evaluation data in the same domain were used. We used sound from seven machine types: fan, gearbox, pump, valve, slider, ToyCar, and ToyTrain. Each machine type has 6 product IDs. For each product ID, 1,000 normal sound samples were used as training data, while 100 normal and 100 anomalous sound samples were used for evaluation data. ID 0, 1, and 2 of the evaluation data were used as validation data to determine the hyperparameters p of IM and λ , and ID 3, 4, and 5 were used as test data. When training Encoder f, 90% of the training data was randomly selected, and the remaining 10% was used for validation. Each recording is a single-channel, 10 sec. segment of audio sampled at 16 kHz.

Each machine's amplitude was calculated and normalized during preprocessing to have a mean of 0 and a variance of 1. For each audio input sequence, we extracted a log-compressed Mel-spectrogram with a window size of 128 ms, a hop size of 16 ms, and 224 Mel-spaced frequency bins in the range of 50–7800 Hz, in 2.0 sec. These features were passed to encoder f

TABLE II: Performance evaluation results. Values represent the harmonic mean of AUC [%] and pAUC (p = 0.1) [%] for each product ID."All/Har-mean" column values represent the harmonic mean of AUC and pAUC over all machines and product IDs.#models column values represent the number of required models for one machine type.

Method	fan	gearbox	pump	valve	slider	ToyCar	ToyTrain	All/Har-mean	#models
Parallel method [12]	56.11	62.43	85.33	66.68	74.58	68.46	69.04	67.93	5
Serial method [22]	82.37	66.22	77.21	72.06	78.02	54.72	50.88	66.78	1
DDCSAD [18]	69.21	58.13	68.55	75.66	59.56	57.71	57.71	63.12	3
Proposed method	84.35	68.42	71.60	65.03	83.97	62.08	58.83	69.42	1

TABLE III: Ablation study results. Values represent the harmonic mean of AUC [%] and pAUC (p = 0.1) [%] for each product ID. "All/Har-mean" column values represent the harmonic mean of AUC and pAUC over all machines and product IDs.

Method	fan	gearbox	pump	valve	slider	ToyCar	ToyTrain	All/Har-mean
Proposed method	84.35	68.42	71.60	65.03	83.97	62.08	58.83	69.42
w/o Mixup	86.35	71.54	66.52	67.66	81.88	52.51	59.90	67.75
w/o $h(\cdot)$	64.34	49.80	66.25	57.19	61.69	50.51	53.35	56.93
Only product IDs classification	85.39	66.07	66.88	58.16	70.50	51.68	52.78	62.79

using EfficientNet-B0 [29]. Encoder f applied global average pooling to the last convolutional layer and performs two nonlinear transformations to obtain a 128-dimensional embedding. Weights of g_{product} and g_{affine} were handled as trainable parameters. We used AdamW [30], OneCycleLR [31], and Mixup ($\alpha = 0.2$) for training. Learning rate, batch size, λ , inlier model h, and the inlier model's hyperparameters p, which the number of components for GMM and the number of neighbors for LOF are shown in Table I. We used the model with the smallest loss of validation data of the training data after 300 epochs. We trained the inlier model h using the validation data of the training data for each product ID. We determined the hyperparameters of the inlier model h using the validation data of the evaluation data. During inference, we divided 10.0 sec. clips into S = 10 segments, with overlap allowed so that each segment was T = 2.0 sec. The GMM used the negative log-likelihood as the anomaly score, while the LOF used the outlier score. The aggregator \mathcal{A} was the mean of the anomaly scores above the median for the GMM and the mean of the entire anomaly scores for the LOF.

B. Results

The parallel [12] and serial [22] methods shown in Table II are the first and second ranked methods in the DCASE 2021 Task 2 Challenge, respectively, while DDCSAD [18] is a conventional method. Table II shows that for All/Har-mean, the proposed method outperformed all of the other methods. As shown in the #models in Table II, the proposed method using a single model outperformed the parallel method ensembled different five models. The proposed method improved the performance of ToyCar and ToyTrain by more than 7% over the serial method, and we believed that the proposed method obtained more suitable embedding for ASD. We believe that the reason for the performance improvement of the proposed method compared to DDCSAD is that it is more flexible and effective in modeling the normal data by IM without assuming a normal distribution for the normal data. Focusing on h, we achieved the best performance when using a GMM with the machine types in MIMII Due dataset and a LOF with the machine types in ToyADMOS2 dataset, suggesting that it was important to use a suitable function h for the dataset.

Our ablation study results are shown in Table III. Fig. 3 shows the results when the embedding of each method is visualized using t-SNE [27]. Table III shows that using Mixup improved the pump, slider, and ToyCar. We believed that Mixup effectively trained data with a high similarity between product IDs. However, Mixup did not improve ASD performance for the other machine types. Based on these, we considered Mixup to be one of the most important hyperparameters.

We then examined the difference in performance with and without h, the IM component of our proposed method. When we did not use h, we used the average of the output probabilities of the product IDs as the anomaly score instead [19]. The experiment shows that IM was effective for ASD. We believe that even though some parts of the data distribution are less informative for the product ID classification and tend to be ignored in calculating the output probabilities of the product IDs, they are still helpful in detecting differences in the data distribution between normal and anomalous data. IM can detect such differences in distribution by GMM or LOF.

We then considered the effect of Eq. 2, which calculates the machine type classification loss. The only product IDs classification in Table III is almost the same as the serial method [22]. The performance of the proposed method outperformed that of the only product IDs classification, since Eq. 2 collected anomalous data that differed significantly from normal data near the center of the hypersphere. More anomalous data was distributed near the other machine types when using the proposed method, as shown in Fig. 3c and Fig. 3d, than when the only product IDs classification was used, as shown in Fig. 3e and Fig. 3f.

V. CONCLUSION

In this paper, we proposed using multi-task learning of two binary cross-entropies: the product ID classification loss and the machine type classification loss when training OE in the serial method. The product ID classification loss and the machine type classification loss correspond when normal and pseudo-anomalous data are too similar or too different, respectively. Even though the proposed method used a single model, it outperformed conventional ensembled methods. Our ablation study and visualization of embeddings confirmed the effectiveness of using post-processing and the effectiveness of the proposed method when detecting anomalous data that is too different from normal data. In future work, we will investigate the impact of training data on ASD performance and the use of data with different domains.

ACKNOWLEDGMENT

This paper was partly supported by a project, JPNP20006, commissioned by NEDO.

REFERENCES

- D. Huang, C. Lin, C. Chen, and J. Sze, "The Internet technology for defect detection system with deep learning method in smart factory," in 2018 4th International Conference on Information Management (ICIM), 2018, pp. 98–102.
- [2] T. Hayashi, T. Komatsu, R. Kondo, T. Toda, and K. Takeda, "Anomalous sound event detection based on wavenet," in 2018 26th European Signal Processing Conference (EUSIPCO). IEEE, 2018, pp. 2494–2498.
- [3] B. Bayram, T. B. Duman, and G. Ince, "Real time detection of acoustic anomalies in industrial processes using sequential autoencoders," *Expert Systems*, vol. 38, no. 1, p. e12564, 2021.
- [4] Y. Koizumi, Y. Kawaguchi, K. Imoto, T. Nakamura, Y. Nikaido, R. Tanabe, H. Purohit, K. Suefusa, T. Endo, M. Yasuda, and N. Harada, "Description and discussion on DCASE2020 challenge task2: Unsupervised anomalous sound detection for machine condition monitoring," in *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2020 Workshop (DCASE2020)*, November 2020, pp. 81–85.
- [5] Y. Kawaguchi, R. Tanabe, T. Endo, K. Ichige, and K. Hamada, "Anomaly Detection Based on an Ensemble of Dereverberation and Anomalous Sound Extraction," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing* (*ICASSP*), 2019, pp. 865–869.
- [6] R. Giri, S. V. Tenneti, F. Cheng, K. Helwani, U. Isik, and A. Krishnaswamy, "Self-supervised classification for detecting anomalous sounds," in *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2020 Workshop (DCASE2020)*, Tokyo, Japan, November 2020, pp. 46–50.
- [7] T. Hayashi, T. Yoshimura, and Y. Adachi, "CONFORMER-BASED ID-AWARE AUTOENCODER FOR UNSUPERVISED ANOMALOUS SOUND DETECTION," DCASE2020 Challenge, Tech. Rep., July 2020.
- [8] K. Suefusa, T. Nishida, H. Purohit, R. Tanabe, T. Endo, and Y. Kawaguchi, "Anomalous Sound Detection Based on Interpolation Deep Neural Network," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 271–275.
- [9] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander, "LOF: Identifying Density-Based Local Outliers," *SIGMOD Rec.*, vol. 29, no. 2, p. 93–104, may 2000. [Online]. Available: https://doi.org/10.1145/335191.335388
- [10] D. W. Scott, "Outlier Detection and Clustering by Partial Mixture Modeling," in COMPSTAT 2004 — Proceedings in Computational Statistics, Springer. Physica-Verlag HD, 2004, pp. 453–464.
- [11] W. Liu, D. Cui, Z. Peng, and J. Zhong, "Outlier Detection Algorithm Based on Gaussian Mixture Model," in 2019 IEEE International Conference on Power, Intelligent Computing and Systems (ICPICS), 2019, pp. 488–492.
- [12] J. A. Lopez, G. Stemmer, P. Lopez Meyer, P. Singh, J. Del Hoyo Ontiveros, and H. Cordourier, "Ensemble of complementary anomaly detectors under domain shifted conditions," in *Proceedings of the* 6th Detection and Classification of Acoustic Scenes and Events 2021 Workshop (DCASE2021), Barcelona, Spain, November 2021, pp. 11– 15.
- [13] K. Dohi, T. Endo, H. Purohit, R. Tanabe, and Y. Kawaguchi, "Flow-Based Self-Supervised Density Estimation for Anomalous Sound Detection," in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 336–340.
- [14] P. Primus, V. Haunschmid, P. Praher, and G. Widmer, "Anomalous sound detection as a simple binary classification problem with careful selection of proxy outlier examples," in *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2020 Workshop (DCASE2020)*, Tokyo, Japan, November 2020, pp. 170–174.

- [15] M. Kim, M. T. Ho, and H.-G. Kang, "Self-supervised Complex Network for Machine Sound Anomaly Detection," in 2021 29th European Signal Processing Conference (EUSIPCO). IEEE, 2021, pp. 586–590.
- [16] L. Ruff, R. A. Vandermeulen, N. Görnitz, A. Binder, E. Müller, K.-R. Müller, and M. Kloft, "Deep Semi-Supervised Anomaly Detection," in *International Conference on Learning Representations*, 2020. [Online]. Available: https://openreview.net/forum?id=HkgH0TEYwH
- [17] L. Ruff, R. A. Vandermeulen, B. J. Franks, K.-R. Müller, and M. Kloft, "Rethinking Assumptions in Deep Anomaly Detection," in *ICML 2021* Workshop on Uncertainty & Robustness in Deep Learning, 2021.
- [18] I. Kuroyanagi, T. Hayashi, K. Takeda, and T. Toda, "Anomalous Sound Detection Using a Binary Classification Model and Class Centroids," in 2021 29th European Signal Processing Conference (EUSIPCO). IEEE, 2021, pp. 1995–1999.
- [19] Y. Kawaguchi, K. Imoto, Y. Koizumi, N. Harada, D. Niizumi, K. Dohi, R. Tanabe, H. Purohit, and T. Endo, "Description and discussion on dcase 2021 challenge task 2: Unsupervised anomalous detection for machine condition monitoring under domain shifted conditions," in *Proceedings of the 6th Detection and Classification of Acoustic Scenes and Events 2021 Workshop (DCASE2021)*, Barcelona, Spain, November 2021, pp. 186–190.
- [20] I. Kuroyanagi, T. Hayashi, Y. Adachi, T. Yoshimura, K. Takeda, and T. Toda, "An ensemble approach to anomalous sound detection based on conformer-based autoencoder and binary classifier incorporated with metric learning," in *Proceedings of the 6th Detection and Classification of Acoustic Scenes and Events 2021 Workshop (DCASE2021)*, Barcelona, Spain, November 2021, pp. 110–114.
- [21] K. Sohn, C.-L. Li, J. Yoon, M. Jin, and T. Pfister, "Learning and evaluating representations for deep one-class classification," in *International Conference on Learning Representations*, 2021. [Online]. Available: https://openreview.net/forum?id=HCSgyPUfeDj
- [22] K. Morita, T. Yano, and K. Tran, "Anomalous sound detection using cnn-based features by self supervised learning," DCASE2021 Challenge, Tech. Rep., July 2021.
- [23] K. Wilkinghoff, "Combining multiple distributions based on sub-cluster adacos for anomalous sound detection under domain shifted conditions," in *Proceedings of the 6th Detection and Classification of Acoustic Scenes* and Events 2021 Workshop (DCASE2021), Barcelona, Spain, November 2021, pp. 55–59.
- [24] R. Tanabe et al., "MIMII Due: Sound Dataset for Malfunctioning Industrial Machine Investigation and Inspection with Domain Shifts Due to Changes in Operational and Environmental Conditions," in 2021 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA), 2021, pp. 21–25.
- [25] N. Harada *et al.*, "Toyadmos2: Another dataset of miniature-machine operating sounds for anomalous sound detection under domain shift conditions," in *Proceedings of the 6th Detection and Classification of Acoustic Scenes and Events 2021 Workshop (DCASE2021)*, Barcelona, Spain, November 2021, pp. 1–5.
- [26] I. Kuroyanagi, T. Hayashi, Y. Adachi, T. Yoshimura, K. Takeda, and T. Toda, "Anomalous sound detection with ensemble of autoencoder and binary classification approaches," DCASE2021 Challenge, Tech. Rep., July 2021.
- [27] L. van der Maaten and G. Hinton, "Visualizing Data using t-SNE," Journal of Machine Learning Research, vol. 9, no. 86, pp. 2579–2605, 2008. [Online]. Available: http://jmlr.org/papers/v9/vandermaaten08a.html
- [28] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond Empirical Risk Minimization," in *International Conference on Learning Representations*, 2018. [Online]. Available: https://openreview.net/forum?id=r1Ddp1-Rb
- [29] Q. Xie, M.-T. Luong, E. Hovy, and Q. V. Le, "Self-training with noisy student improves imagenet classification," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 10687–10698.
- [30] I. Loshchilov and F. Hutter, "Decoupled Weight Decay Regularization," in 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019. OpenReview.net, 2019. [Online]. Available: https://openreview.net/forum?id=Bkg6RiCqY7
- [31] L. N. Smith and N. Topin, "Super-convergence: very fast training of neural networks using large learning rates," in *Artificial Intelligence* and Machine Learning for Multi-Domain Operations Applications, T. Pham, Ed., vol. 11006, International Society for Optics and Photonics. SPIE, 2019, pp. 369 – 386. [Online]. Available: https://doi.org/10.1117/12.2520589