

# Discriminative Neighborhood Smoothing for Generative Anomalous Sound Detection

Takuya Fujimura  
Nagoya University  
Nagoya, Japan

fujimura.takuya@g.sp.m.is.nagoya-u.ac.jp

Keisuke Imoto  
Doshisha University  
Kyoto, Japan

keisuke.imoto@ieee.org

Tomoki Toda  
Nagoya University  
Nagoya, Japan

tomoki@icts.nagoya-u.ac.jp

**Abstract**—We propose discriminative neighborhood smoothing of generative anomaly scores for anomalous sound detection. While the discriminative approach is known to achieve better performance than generative approaches often, we have found that it sometimes causes significant performance degradation due to the discrepancy between the training and test data, making it less robust than the generative approach. Our proposed method aims to compensate for the disadvantages of generative and discriminative approaches by combining them. Generative anomaly scores are smoothed using multiple samples with similar discriminative features to improve the performance of the generative approach in an ensemble manner while keeping its robustness. Experimental results show that our proposed method greatly improves the original generative method, including absolute improvement of 22% in AUC and robustly works, while a discriminative method suffers from the discrepancy.

**Index Terms**—anomalous sound detection, discriminative method, generative method, ensemble

## I. INTRODUCTION

Anomalous sound detection (ASD) aims to automatically detect mechanical failures from sounds emitted by a target machine [1]–[3]. This technique is used in machine condition monitoring systems within factories. As anomalous sound data are rarely observed and unknown failures must also be detected in the ASD task, an ASD system needs to be developed without using labeled data. This is typically done by constructing an anomaly score calculator using only the normal sounds, where high scores should be assigned to anomalous sounds.

ASD methods are classified into two approaches: generative and discriminative approaches. The generative approach directly models the audio feature of the normal sounds and calculates the degree of deviation from the model as an anomaly score [1], [4]–[7]. Autoencoder (AE) [1], [4], normalizing flow (NF) [5], [6], and Gaussian mixture models (GMM) [7] are popular modeling methods. AE is trained to reconstruct normal sounds and calculates the anomaly score of an observed sound by its reconstruction error. NF and GMM are trained to maximize its likelihood for normal sounds and calculate the negative likelihood as the anomaly score. This generative approach is a proper strategy for ASD when the labeled anomalous data is not available during development; however, it is difficult to construct highly accurate generative models. For this reason, the generative approach tends to have insufficient detection performance.




Fig. 1. Illustration of the feature space without and with the discrepancy between the characteristics of the training and test data.

Another strategy is the discriminative approach [8]–[12]. This approach models the normal area in the discriminative feature space by utilizing the discriminative model trained to classify normal and *pseudo-anomalous* sounds. The pseudo-anomalous sounds are prepared by various methods such as pitch shifting, time stretching [8], mixup, changing the statistics of spectrograms [9], and utilizing sounds that are not emitted from a target machine [10], [11]. This approach assumes that anomalous sounds will be located outside of the normal area (i.e., decrease in the posterior probability of normal sound), and calculates anomaly scores using classification probabilities or deviations from the inlier model constructed with the normal training data (Fig. 1 (a)). The discriminative approach has outperformed the generative approach by turning the problem setting from a harder one (i.e., modeling a probability density function) to easier one (i.e., modeling a posterior probability function). However, when there is a discrepancy in data characteristics between the training and test data, the test data sometimes deviate from the inlier model even if they are normal sounds (Fig. 1 (b)). In this case, the inlier model does not reflect the correct relationship for the anomaly score, leading to significant performance degradation.

The goal of this work is to outperform the generative approach without the risk of performance degradation caused by the discrepancy between the training and test data. To achieve this goal, we propose the discriminative neighborhood smoothing of a generative method, based on our finding that normal and anomalous sounds tend to be distinguished in the discriminative feature space even in the presence of discrep-

ancies (Fig. 1 (b)). The proposed method improves the performance of the original generative approach by ensembling the anomaly scores of multiple samples with similar discriminative features. In the experimental evaluation, we demonstrate that (1) the proposed method greatly improves the performance of the original generative method and, (2) robustly works and has the potential to improve its performance even in a situation where the discriminative method does not work due to the discrepancy. Furthermore, (3) we conduct several analyses to clarify its behavior.

## II. RELATED WORKS

### A. Generative approach

The autoencoder (AE)-based method [1] is a popular generative method. This method trains the AE to minimize the following reconstruction loss  $\mathcal{L}_{\text{AE}}$ .

$$\mathcal{L}_{\text{AE}} = \frac{1}{MD} \sum_m \|\psi(\mathbf{x}_m^-) - f_G(\psi(\mathbf{x}_m^-); \boldsymbol{\theta})\|_2^2, \quad (1)$$

where  $M$  is a mini-batch size,  $\mathbf{x}_m^-$  is a  $m$ -th normal sound,  $\psi(\cdot)$  extracts  $D$ -dimensional audio feature such as Mel-spectrogram,  $f_G$  is a network of AE and  $\boldsymbol{\theta}$  is the set of its parameters. During inference, this method calculates the following anomaly score  $a_{\text{AE}}(\mathbf{y}_m)$  for the observed sound  $\mathbf{y}_m$ .

$$a_{\text{AE}}(\mathbf{y}_m) = \frac{1}{D} \|\psi(\mathbf{y}_m) - f_G(\psi(\mathbf{y}_m); \boldsymbol{\theta})\|_2^2. \quad (2)$$

The AE-based method calculates the reconstruction loss for the anomaly score assuming that the reconstruction quality of the anomaly sound will be poor because the network is not trained with anomaly sounds.

Thus, the generative approach, such as the AE-based method, requires highly accurate modeling of normal sounds. However, this modeling is very challenging, often resulting in insufficient performance in detecting anomalous sounds.

### B. Discriminative approach

Various discriminative approaches have been investigated [8]–[12], among which SerialOE [10] shows high performance. It uses sounds emitted from non-target machines as pseudo-anomalous sounds for the training of the discriminative feature extractor. Besides, it utilizes normal sounds emitted from different products of the target machine type to classify each of them where each product is identified by product ID. SerialOE achieves high performance by calculating the anomaly score based on the discriminative feature space, and it is realized by two-stage training. First, the discriminative feature extractor  $f_D(\cdot)$  is trained with the following multi-task loss  $\mathcal{L}_{\text{SerialOE}}$ .

$$\mathcal{L}_{\text{SerialOE}} = \mathcal{L}_{\text{machine}} + \lambda_{\text{id}} \mathcal{L}_{\text{id}}, \quad (3)$$

where  $\lambda_{\text{id}}$  is a hyperparameter for weighting,  $\mathcal{L}_{\text{machine}}$  is a loss for classifying normal and pseudo-anomalous sounds as follows:

$$\mathcal{L}_{\text{machine}} = \frac{1}{M} \sum_m \text{BCE}(t_m^{\text{mac}}, \sigma(g_{\text{mac}}(f_D(\psi(\mathbf{x}_m))))), \quad (4)$$

$$\text{BCE}(p, q) = p \cdot \log q + (1 - p) \cdot \log(1 - q), \quad (5)$$

where  $\mathbf{x}_m$  is the  $m$ -th training data, and  $t_m^{\text{mac}}$  is the label that is 1 for the normal sound and 0 for the pseudo-anomalous sound.  $g_{\text{mac}}(\cdot)$  is a function transforming the discriminative feature into a scalar, and  $\sigma(\cdot)$  is a sigmoid function.  $\mathcal{L}_{\text{id}}$  is a loss for classifying sounds of the target machine in detail using the product ID.

$$\mathcal{L}_{\text{id}} = \frac{1}{MC} \sum_m \sum_c \text{BCE}(t_{m,c}^{\text{id}}, \sigma(g_{\text{id}}(f_D(\psi(\mathbf{x}_m))))), \quad (6)$$

where the target machine has  $C$  classes for the product ID,  $t_{m,c}^{\text{id}}$  is its one-hot label. When  $\mathbf{x}_m$  belongs to the  $c$ -th class,  $t_{m,c}^{\text{id}}$  is 1 and the other elements of  $t_m^{\text{id}}$  are 0. This multi-task training manages both major and minor differences in the audio features and constructs a better feature space in which anomalous sounds are distinguished from normal sounds.

Next, an inlier model  $h$  such as a GMM is trained with the discriminative features of the normal training data. During inference, SerialOE calculates the anomaly score of the observed sound  $\mathbf{y}_m$  using its discriminative feature  $f_D(\mathbf{y}_m)$  and a trained inlier model  $h$ . For example, GMM can calculate the negative log-likelihood as the anomaly score (Fig. 1 (a)).

The discriminative approach such as SerialOE has achieved high performance by using the discriminative feature space. However, when there is a discrepancy between the characteristics of the training and test data, the test data forms a different cluster from the cluster of the training data, even if they are normal sounds. In addition, the anomalous cluster of the test data is sometimes located closer to the inlier model  $h$  than the normal cluster of the test data (Fig. 1 (b)). In this case, if we use the degree of deviation from  $h$  as the anomaly score, the score relationship of normal and anomalous sounds will be reversed, leading to significant performance degradation [13]. Thus, the discriminative approach has a risk of yielding poor performance due to the discrepancy between the training and test data.

## III. PROPOSED METHOD

We have two important observations: (1) the generative approach is less sensitive to a discrepancy between the training and test data but its detection performance is insufficient, and (2) normal and anomalous sounds in the test data can be distinguished in the discriminative feature space even in the presence of the discrepancy [13]. Based on these observations, we propose the discriminative neighborhood smoothing of the generative approach (Fig. 2).

The proposed method works as follows. First, we obtain a set of discriminative features  $\mathcal{F} = \{f_D(\mathbf{y}_i) | \mathbf{y}_i \in \mathcal{Y}\}$

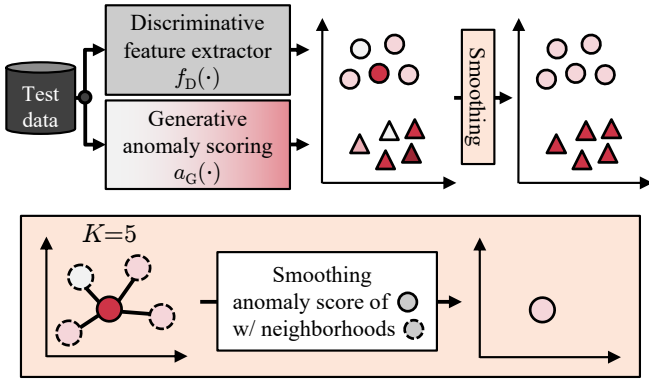


Fig. 2. Overview of the proposed method. Test data are processed with the discriminative feature extractor and generative anomaly scoring in parallel and we obtain the discriminative features to which the generative anomaly scores are assigned. Anomaly scores are smoothed in the discriminative space and are expected to be modified to the correct score using multiple neighborhood samples.

from a set of observed sounds during the operating phase  $\mathcal{Y} = \{\mathbf{y}_0, \dots, \mathbf{y}_N\}$  (i.e.,  $\mathcal{Y}$  is a test dataset). Next, the proposed method calculates the anomaly score  $a_{\text{prop}}(\mathbf{y}_i)$  for the observation  $\mathbf{y}_i$  as follows:

$$a_{\text{prop}}(\mathbf{y}_i) = \frac{1}{K} \left( a_G(\mathbf{y}_i) + \sum_{k \in \mathcal{N}_{\mathcal{F}}(\mathbf{y}_i)} a_G(\mathbf{y}_k) \right), \quad (7)$$

where  $a_G(\cdot)$  calculates the anomaly score by some generative method.  $\mathcal{N}_{\mathcal{F}}(\mathbf{y}_i)$  is the set of indices of neighborhood of  $f_D(\mathbf{y}_i)$  on the  $\mathcal{F}$  except for  $i$ ,  $K$  is the number of neighbors, and  $|\mathcal{N}_{\mathcal{F}}(\mathbf{y}_i)| = K - 1$ .

The proposed method improves the detection performance from the original generative method by ensembling anomaly scores of multiple samples with similar discriminative features. For example, a high anomaly score wrongly assigned to normal sound will be modified to a correctly low score by smoothing if most of its neighborhoods are expected to be normal sounds (Fig. 2). Also, the proposed method is robust to discrepancies because it does not use the deviation from the inlier model  $h$  as the anomaly score. Thus, the proposed method achieves robust performance by relaxing the requirements for the discriminative feature space from both (i) normal and anomalous sounds are distinguished and (ii) normal sounds of test and training data are matched, to only (i). It should be noted that the proposed method uses the test data to modify anomaly scores. Therefore, the ASD system using our proposed method works by regularly updating  $\mathcal{F}$  with the data observed during the operating phase.

#### IV. EXPERIMENTAL EVALUATION

##### A. Experimental Setups

We conducted an experimental evaluation using the DCASE 2021 Task 2 Challenge datasets (MIMII DUE [14] and Toy-ADMO2 [15]). The datasets included seven machine types: fan, gearbox, pump, slider, ToyCar, ToyTrain, and valve. It also provided two domains (source and target domains) and six sections for each machine type to handle the domain-shifted

TABLE I  
HYPERPARAMETERS FOR EACH MACHINE TYPE. “ALL” OF THE “DATA” COLUMNS MEANS THE USE OF BOTH SOURCE AND TARGET DOMAIN.

	SerialOE		Proposed	Proposed (Oracle)
	# clusters	Data	$K$ Data	$K$ Data
fan	2	Source	512 Source	1,024 Source
gearbox	32	Source	1,024 All	2,048 All
pump	8	Source	128 Source	1,024 Source
slider	4	All	64 Source	128 Source
ToyCar	4	All	128 Source	2,048 All
ToyTrain	16	Source	1,024 All	256 All
valve	2	All	128 Source	1,024 Source

problem, where the section was a subset for the evaluation and it reflected the product ID. The training data had approximately 1,003 samples of normal data for each section of each machine type, of which only three samples were in the target domain and the others were in the source domain. The test data had 100 samples of normal and anomalous data for each machine, section, and domain. Each recording was a 10-second single channel segment sampled at 16 kHz and we standardized its amplitude to have a mean of 0.0 and a variance of 1.0. Sections 0, 1, and 2 of the test data were used as the validation data to determine the hyperparameters, while Sections 3, 4, and 5 were used as the evaluation data. To train the feature extractor and the generative method, we used 85% of the source domain data and two samples of the target domain data. The remaining samples were treated as the validation set during training.

Here, we used an AE as  $a_G(\cdot)$ . The network was a ten-layer fully connected network [1]. The audio input feature  $\psi(\cdot)$  was the five adjacent frames of the log Mel power spectrogram with a window size of 64 ms, a hop size of 32 ms, and 128 Mel-spaced frequency bins in the range of 50 Hz to 7800 Hz. We trained the network for 1,500 epochs using the Adam optimizer [16] with a fixed learning rate of 0.001. The mini-batch included 16 of the 4-second segments.

We then used a feature extractor of the SerialOE as  $f_D(\cdot)$ . The network consisted of the EfficientNet B0 [17], global average pooling, and two nonlinear transformations, and it outputted a 128-dimensional feature. Both  $g_{\text{mac}}(\cdot)$  and  $g_{\text{id}}(\cdot)$  were linear transformations, and  $\lambda_{\text{id}}$  was set to 5. The audio input feature  $\psi(\cdot)$  was the 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 Hz to 7,800 Hz. We trained the network for 150 epochs using the AdamW optimizer [18], OneCycleLR scheduler [19] with a learning rate of 0.001, and mixup [20]. The mini-batch included 128 of the 4-second segments and we used a batch sampler so that the ratio of the value of  $t^{\text{mac}} \in \{0, 1\}$  was 1:1. We also used the SerialOE as the comparison method and its  $h$  was the GMM.

During inference, we divided the original 10-second signal into seven 4-second segments with 75% overlapping and calculated the final anomaly score by averaging the score of each segment. As for the hyperparameters, the GMM of SerialOE included the number of clusters and the domain of its training data, and the proposed method included the number

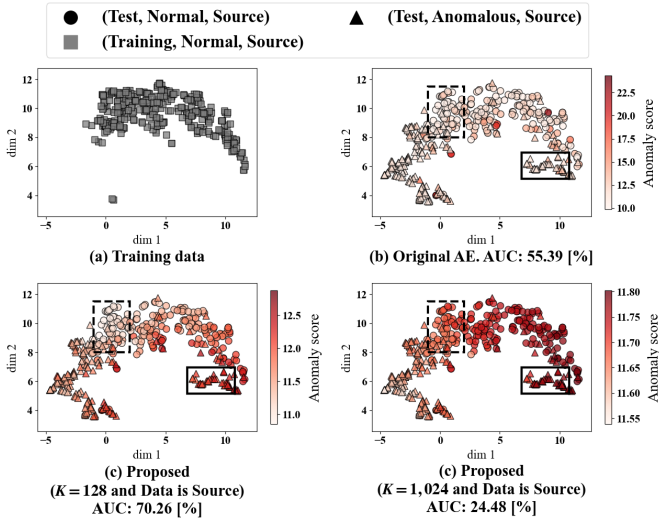


Fig. 3. Visualization of smoothing in the source domain of Section 4 of the slider. The colormap shows the anomaly score of  $a_G$  or  $a_{prop}$ . For visibility, only 30% of the entire sample was randomly selected and displayed. All figures have the same axes.

of neighbors  $K$  and the domain of the test data used for smoothing. These were selected based on the validation results and the resulting settings are shown in Table I. For reference, we also evaluated the Oracle version of the proposed method, which used the best hyperparameters on the evaluation set.

### B. Results

The results in Table II show that the proposed method improves the performance of AE in machines other than ToyCar and ToyTrain, achieving an absolute improvement of 7% in All-hmean. The proposed method makes especially significant improvements larger than 13% in the fan, pump, and slider where SerialOE achieves high performance larger than 75%. We can also see that SerialOE has an AUC less than 50% in ToyCar-5 and the discriminative approach has a risk that falls into significant performance degradation. Even in such a situation, the proposed method works robustly and even has the potential to improve the performance of original generative methods (see Proposed (Oracle)) because an AUC significantly below 50% indicates that normal and anomalous sounds are distinguished while the cluster of the anomalous sounds is closer to  $h$ . The proposed method can further improve its performance by setting proper hyperparameters (e.g., improvement of 22% in the valve). Our future work includes its adaptive setting depending on the space.

Figures 3 and 4 show the feature space using UMAP [21] with the original and smoothed anomaly scores. In the solid frames in Fig. 3, we can see that the proposed method modifies the anomaly scores of the anomalous sounds to higher scores. On the other hand, it is clear that a too-large  $K$  causes performance degradation like in the dashed frame, where normal sound is also smoothed with anomalous sounds. In the solid frames in Fig. 4, we can see that the normal sounds of the test data deviate from the cluster of the training data. The proposed method robustly works in such a situation and

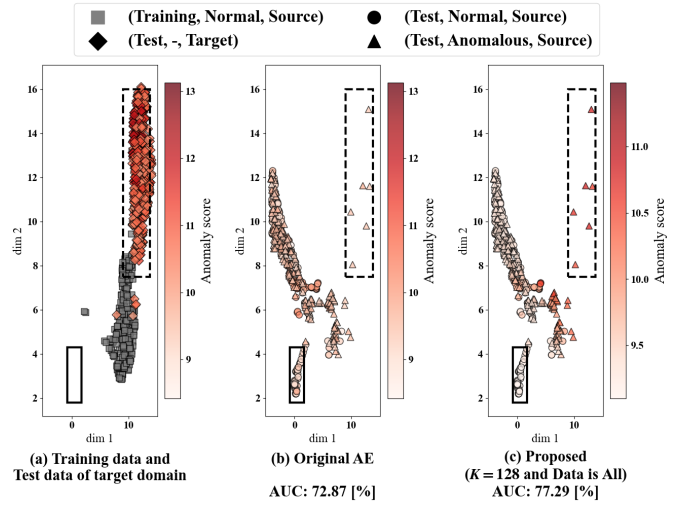


Fig. 4. Visualization of smoothing in Section 5 of ToyCar. The colormap shows the anomaly score of  $a_G$  or  $a_{prop}$ . For visibility, only 30% of the entire sample was randomly selected and displayed. All figures have the same axes.

modifies the anomaly scores of the normal sounds to lower scores. In the dashed frames, anomalous sounds are smoothed with the sounds of the target domain which have relatively high anomaly scores, resulting in performance improvements.

To investigate the behavior of the proposed method in a real situation where we can get new samples sequentially, we evaluate it by changing the number of samples used for smoothing. Figure 5 shows the evaluation result when the number of samples used for  $\mathcal{Y}$  is changed from 5% to 100% of the total evaluation data in increments of 5%, while keeping the ratio of normal to anomalous and source to target domains in each section. We randomly selected used samples and plotted the mean and standard deviation of 20 trials with different random seeds. We always gave the best hyperparameter to the proposed method. In Fig. 5, we can see an overall increasing trend in the AUC and a decreasing trend in the standard deviation corresponding to an increase in the number of samples. The result also shows that even with the small number of samples, the proposed method can improve the original generative method by giving proper hyperparameters, although there is some variation. This experiment shows that it is desirable to have regular updates of  $\mathcal{Y}$ ; however, there is no additional cost for recording sounds, since ASD systems always do this for monitoring. Rather, we believe that past data should be accumulated and utilized.

### V. CONCLUSION

In this paper, we proposed the discriminative neighborhood smoothing of generative anomaly scores to achieve high-performance ASD without the risk of degradation due to the discrepancy between the training and test data. The proposed method improved the performance of the generative method by smoothing the anomaly scores of multiple samples with similar discriminative features. This method does not suffer from the discrepancy because it does not use the inlier model to calculate the anomaly scores. Our experimental results showed that

TABLE II

EVALUATION RESULTS. VALUES OF EACH MACHINE REPRESENT THE HARMONIC MEAN OF THE AUC [%] AND PAUC ( $p = 0.1$ ) [%] OVER ALL SECTIONS. VALUES OF THE ALL-HMEAN REPRESENT THE HARMONIC MEAN OF THE AUC AND PAUC OVER ALL MACHINES AND SECTIONS. VALUES OF THE TOYCAR-5 REPRESENT THE AUC OF TOYCAR IN SECTION 5. HYPERPARAMETERS OF PROPOSED (ORACLE) ON THE TOYCAR-5 ARE 128 FOR  $K$  AND ALL FOR DATA

Method	fan	gearbox	pump	slider	ToyCar	ToyTrain	valve	All-hmean	ToyCar-5
SerialOE	<b>80.78</b>	60.77	<b>76.86</b>	<b>81.49</b>	49.37	55.06	64.02	64.72	33.58
AE	57.11	59.50	54.70	55.63	58.61	58.74	51.27	56.38	72.87
Proposed	73.22	64.42	71.84	69.13	58.18	55.91	56.95	63.52	72.13
Proposed (Oracle)	74.56	<b>64.64</b>	74.76	70.63	<b>63.09</b>	<b>70.19</b>	<b>73.28</b>	<b>69.89</b>	<b>77.29</b>




Fig. 5. Evaluation changing the number of samples used for smoothing. The evaluation metric is a harmonic mean of the AUC [%] and pAUC ( $p = 0.1$ ) [%] over all sections. The error band shows the standard deviation.

(1) the proposed method greatly improved the performance of AE, (2) it robustly worked despite discrepancies, (3) it had the potential to outperform SerialOE by removing the risk due to the discrepancy, and (4) it is desirable to accumulate and utilize the past data for its stability.

#### ACKNOWLEDGMENT

This paper was partly supported by a project, JPNP20006, commissioned by NEDO, JSPS KAKENHI Grant Number JP20H00102, and AIST KAKUSEI project (FY2023).

#### REFERENCES

- [1] Y. Koizumi, Y. Kawaguchi, K. Imoto, *et al.*, “Description and discussion on DCASE2020 challenge task2: Unsupervised anomalous sound detection for machine condition monitoring,” in *Proc. DCASE*, 2020, pp. 81–85.
- [2] Y. Kawaguchi, K. Imoto, Y. Koizumi, *et al.*, “Description and discussion on DCASE 2021 challenge task 2: Unsupervised anomalous detection for machine condition monitoring under domain shifted conditions,” in *Proc. DCASE*, 2021, pp. 186–190.
- [3] K. Dohi, K. Imoto, N. Harada, *et al.*, “Description and discussion on DCASE 2022 challenge task 2: Unsupervised anomalous sound detection for machine condition monitoring applying domain generalization techniques,” in *Proc. DCASE*, 2022, pp. 1–5.
- [4] K. Suefusa, T. Nishida, H. Purohit, R. Tanabe, T. Endo, and Y. Kawaguchi, “Anomalous sound detection based on interpolation deep neural network,” in *Proc. ICASSP*, 2020, pp. 271–275.
- [5] K. Dohi, T. Endo, H. Purohit, R. Tanabe, and Y. Kawaguchi, “Flow-based self-supervised density estimation for anomalous sound detection,” in *Proc. ICASSP*, 2021, pp. 336–340.
- [6] K. Dohi, T. Endo, and Y. Kawaguchi, “Disentangling physical parameters for anomalous sound detection under domain shifts,” in *Proc. EUSIPCO*, 2022, pp. 279–283.
- [7] J. Guan, Y. Liu, Q. Zhu, T. Zheng, J. Han, and W. Wang, “Time-weighted frequency domain audio representation with GMM estimator for anomalous sound detection,” in *Proc. ICASSP*, 2023, pp. 1–5.
- [8] T. Inoue, P. Vinayavekhin, S. Morikuni, *et al.*, “Detection of anomalous sounds for machine condition monitoring using classification confidence,” in *Proc. DCASE*, 2020, pp. 66–70.
- [9] H. Chen, Y. Song, Z. Zhuo, *et al.*, “An effective anomalous sound detection method based on representation learning with simulated anomalies,” in *Proc. ICASSP*, 2023, pp. 1–5.
- [10] I. Kuroyanagi, T. Hayashi, K. Takeda, and T. Toda, “Improvement of serial approach to anomalous sound detection by incorporating two binary cross-entropies for outlier exposure,” in *Proc. EUSIPCO*, 2022, pp. 294–298.
- [11] 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 *Proc. DCASE*, 2020, pp. 170–174.
- [12] K. Wilkinghoff, “Design choices for learning embeddings from auxiliary tasks for domain generalization in anomalous sound detection,” in *Proc. ICASSP*, 2023, pp. 1–5.
- [13] T. Fujimura, I. Kuroyanagi, T. Hayashi, and T. Toda, “Anomalous sound detection by end-to-end training of outlier exposure and normalizing flow with domain generalization techniques,” *DCASE2023 Challenge*, Tech. Rep., 2023.
- [14] R. Tanabe, H. Purohit, K. Dohi, *et al.*, “MIMII DUE: Sound dataset for malfunctioning industrial machine investigation and inspection with domain shifts due to changes in operational and environmental conditions,” *Proc. WASPAA*, pp. 21–25, 2021.
- [15] N. Harada, D. Niizumi, D. Takeuchi, Y. Ohishi, M. Yasuda, and S. Saito, “ToyADMOS2: Another dataset of miniature-machine operating sounds for anomalous sound detection under domain shift conditions,” in *Proc. DCASE*, 2021, pp. 1–5.
- [16] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Proc. ICLR*, 2015, pp. 1–15.
- [17] Q. Xie, M.-T. Luong, E. Hovy, and Q. V. Le, “Self-training with noisy student improves imagenet classification,” in *Proc. CVPR*, 2020, pp. 10 687–10 698.
- [18] I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” in *Proc. ICLR*, 2019.
- [19] 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*, vol. 11006, 2019, pp. 369–386.
- [20] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, “Mixup: Beyond empirical risk minimization,” in *Proc. ICLR*, 2018.
- [21] L. McInnes, J. Healy, N. Saul, and L. Grossberger, “UMAP: Uniform Manifold Approximation and Projection,” *The Journal of Open Source Software*, vol. 3, no. 29, 2018, 63 pages.