

Original Paper

Serial-OE: Anomalous Sound Detection Based on Serial Method with Outlier Exposure Capable of Using Small Amounts of Anomalous Data for Training

Ibuki Kuroyanagi^{1*}, Tomoki Hayashi¹, Kazuya Takeda² and Tomoki Toda³

¹*The Graduate School of Informatics, Nagoya University, Aichi, Japan*

²*The Institutes of Innovation for Future Society, Nagoya University, Aichi, Japan*

³*The Information Technology Center, Nagoya University, Aichi, Japan*

ABSTRACT

We introduce Serial-OE, a new approach to anomalous sound detection (ASD) that leverages small amounts of anomalous data to improve the performance. Conventional ASD methods rely primarily on the modeling of normal data, due to the cost of collecting anomalous data from various possible types of equipment breakdowns. Our method improves upon existing ASD systems by implementing an outlier exposure framework that utilizes normal and pseudo-anomalous data for training, with the capability to also use small amounts of real anomalous data. A comprehensive evaluation using the DCASE2020 Task2 dataset shows that our method outperforms state-of-the-art ASD models. We also investigate the impact on performance of using a small amount of anomalous data during training, of using data without machine ID information,

*Corresponding author: Ibuki Kuroyanagi, kuroyanagi.ibuki@g.sp.m.is.nagoya-u.ac.jp. This paper was partly supported by Japan's New Energy and Industrial Technology Development Organization (NEDO), project JPNP20006.

Received 24 August 2024; revised 30 December 2024; accepted 30 December 2024

ISSN 2048-7703; DOI 10.1561/116.20240060

© 2025 I. Kuroyanagi, T. Hayashi, K. Takeda and T. Toda

and of using contaminated training data. Our experimental results reveal the potential of using a very limited amount of anomalous data during training to address the limitations of existing methods using only normal data for training due to the scarcity of anomalous data. This study contributes to the field by presenting a method that can be dynamically adapted to include anomalous data during the operational phase of an ASD system, paving the way for more accurate ASD.

Keywords: Anomalous sound detection, self-supervised learning, outlier exposure, serial method.

1 Introduction

Anomalous sound detection (ASD) is used to detect unusual sounds, such as equipment malfunctions. It is often used to monitor sound inside machines, or other areas that cannot be directly monitored by cameras, such as equipment in tunnels or other dark locations. Early detection of anomalous conditions can minimize losses due to outages and damage resulting from breakdowns. Due to a decrease in the number of skilled maintenance technicians in recent years, and an increase in the number of aging facilities [2], there is a greater need for systems that can automatically detect anomalous conditions and allow timely repair after mechanical breakdowns [25, 1, 15, 23].

First, let's consider the development of an ASD system. Ideally, both normal and anomalous audio data would be used to train a binary classification model. However, the possible types of anomalous data are diverse, and intentionally causing machines to break down in various ways to collect anomalous data is expensive. Therefore, the detection model of an ASD system is often trained using only normal data [26, 28]. Next, let's consider the operational phase of ASD systems. These systems are generally installed in facilities that experience occasional breakdowns due to aging machinery, or because the facility's design makes human monitoring impractical. Conversely, facilities that do not generally experience anomalous conditions do not need an ASD system. Therefore, facilities that have installed ASD systems are likely able to capture samples of various types of anomalous audio data during breakdowns, data which contains important information about possible anomalous conditions for the machine, so it would be advantageous to utilize this data to improve the performance of the ASD system. Several methods have been proposed using active learning to improve performance in sound event detection tasks by utilising a small number of labelled data [51, 36]. However, most studies on ASD have focused on the development phase, when normal data is generally used

to create the detection model for the ASD system. Few studies have focused on continuously improving ASD performance by utilizing the small amounts of anomalous data that can be obtained during the operational phase [30].

Current ASD methods can be classified into generative and discriminative model-based methods [25, 1]. Generative model-based methods identify sounds deviating from the normal distribution as anomalous data. For example, some methods use reconstruction errors obtained using an autoencoder [54, 11, 42, 21, 37], Long-Short Term Memory networks [39], WaveNet [15], or generative adversarial networks [60, 20], as well as methods that use the likelihoods of Gaussian mixture models (GMM) [44, 32, 13] or normalizing flow [40, 24, 8]. These generative model-based methods cannot use anomalous data because they are trained while generating normal data. In methods that minimize the reconstruction error of normal data, if anomalous data is added to the training data and the reconstruction error is minimized, the anomalous data can be reconstructed, so the ASD performance is degraded. Similarly, if anomalous data is added to methods that maximize the likelihood of normal data, the likelihood of anomalous data increases, so the performance is degraded. In order to utilize anomalous data with these methods, it is necessary to add an architecture for utilizing anomalous data to the original method.

Discriminative model-based methods can also be divided into two types. The first method uses normal data to make multi-class classifications. It includes methods based on classification by machine type, using either data augmentation [19], metadata classification [12, 33, 57, 55, 56, 3, 59], or contrastive learning [18, 64, 14, 4], which have demonstrated competitive performance in recent years. These methods compute anomaly scores for multiple types of machines, and have the advantage of being able to detect anomalous conditions in multiple machines using a single model. However, because these methods do not assume the use of anomalous data, it is not easy to use them without changing the architecture, even if anomalous data is available. The second method defines normal data and pseudo-anomalous data and use a binary classification. Some methods use a single model to infer the status of a single machine [30, 50, 46, 48, 41, 29], while others use a single model to infer the status of multiple machines by considering the normal sounds of all similar machines as normal data [31]. Since these methods use pseudo-anomalous data as training data, it is possible to utilize real anomalous data without making any changes to the architecture.

In this research, we also employ a discriminative model-based method that defines normal data and pseudo-anomalous data and use a binary classification, allowing us to take advantage of any real-world anomalous data that is available as additional training data, in order to improve ASD performance. In our previous work, we compared the average performance of the proposed method when using the DCASE2021 Task2 dataset to show that it improves ASD performance [31]. However, we did not perform a

component-by-component analysis of our method, evaluate the effectiveness of using of real-world anomalous data for training, or examine the possibility of contamination of the real-world anomalous data with normal data, which may occur when collecting the anomalous data. In this paper, we compare the performance of our proposed system with other ASD methods when using the DCASE2020 Task2 dataset, the major ASD dataset available [27, 43]. Our investigation into the possibility of contamination of the real anomalous data used for model training is the first such investigation of this topic in the field of ASD. We not only compare the proposed method with state-of-the-art methods proposed in previous studies, and with similar methods, but also perform ablation studies for each component, evaluate the effect of the DCASE2020 Task2-specific information known as ‘machine ID’, which has been used in previous studies, evaluate the effect of utilizing a small amount of real anomalous data for model training, and evaluate the effect of training the model with data containing anomalous and normal data from the same type of machine.¹

The rest of this paper is organized as follows: In Section 2, we describe previously proposed ASD methods which are able to utilize anomalous data. In Section 3, we provide the details of our proposed method. In Section 4, we investigate the reasons why our proposed method improves ASD performance and the impact of utilizing anomalous data during detection model training on performance through experimental evaluation. In Section 5, we note the limitations of this study, and in Section 6 we present our conclusions.

2 ASD Methods that can Utilize Anomalous Data

Few methods which utilize anomalous data have been proposed for ASD. Here, small amounts of anomalous data refers to data that accounts for approximately 0.1% to 3% of the training data. DDCSAD+BCE [30] is a method which we have proposed previously for improving the performance of ASD systems by utilizing small amounts of anomalous data for training. A feature extractor is trained by defining normal data and pseudo-anomalous data, and two loss functions are then applied. The first loss function performs binary classification using binary cross-entropy. The second loss function uses metric learning to define the centroid of each class in the feature space, then uses Euclidean distance to cluster data in the same class while pushing data from different classes farther apart. The classification results for the normal and pseudo-anomalous classes obtained through multi-task learning using these loss functions are then used as anomaly scores. This method improves performance

¹The current implementation of our proposed Serial-OE system is available at <https://github.com/ibkuroyanagi/Serial-OE>.

by not only discriminatively learning normal and pseudo-anomalous data, but also by explicitly classifying normal data and other data in the feature space. This method also improves performance by classifying anomalous data as part of the pseudo-anomalous data class.

To explain this method in more detail, here we will describe the DCASE2020 Task2 data set [27, 43] and how it was used when evaluating this method. The Task2 dataset has a hierarchical structure of six machine types (Fan, Pump, Slider, Valve, ToyCar, and ToyConveyor), as well as machine IDs that identify the manufacturer and model of each machine belonging to each machine type. There are seven machine IDs for each machine type, except for the ToyConveyor machine type, which contains only six machine IDs, thus there are 41 machine IDs in the entire data set. DDCSAD+BCE treats specific machine IDs within a machine type as normal sound data, and the normal sounds of the other machines of the same type, as well the normal sounds of the other types of machines, are used as pseudo-anomalous data. For example, if machine ID 0 of machine type Fan is considered to be our normal data, the normal sounds of the Fans with machine IDs from 1 to 6, as well as the normal sounds of the different machine types, such as Pumps, Sliders, Valves, ToyCars and ToyConveyors are treated as pseudo-anomalous data. Therefore, 41 different models need to be trained when using the DCASE2020 Task2 data set in order to evaluate ASD performance for all of the machines.

Our DDCSAD+BCE method performed well in terms of average AUC for a single model at that time. Furthermore, ASD performance was improved by adding real anomalous data to the pseudo-anomalous class and by fine tuning the detection model. However, the performance of this method is not as high as that of more recent ASD methods, and it also suffers from large variations in performance from machine ID to machine ID. Since the proposed method trains a feature extractor for each machine ID, performance is strongly influenced by the particular characteristics of each machine’s audio data, which is defined as normal data. In addition, this method uses only the classification probability from discriminative training as the anomaly score, however discriminative classification results are not necessarily equivalent to ASD results [33]. Considering the recent success achieved by combining discriminative and generative models, our method’s lack of an anomaly detection method based on generative models, such as outlier detection, may also be a factor in its lower performance [22]. Therefore, to take advantage of the small amount of anomalous data, this paper retains the framework for classifying normal and pseudo-anomalous data but combines ASD methods based on generative models to improve performance.

From a different perspective, we investigated how anomalous detection in image detection using anomalous data can be used to improve the performance of ASD systems. Random images obtained from large data sets are called outlier exposures, and the use of such data has resulted in higher anomaly detection

performance [16, 17]. DeepSAD [48] was the first deep model to exploit a small number of anomalous cases, by generalizing unsupervised DeepSVDD [46] for use in a semi-supervised anomaly detection setting. In [47], Ruff *et al.* further modified DeepSAD using a method based on cross-entropy classification, significantly improved DeepSAD’s performance. However, methods that use only a small amount of anomalous data have a problem; their output is biased against the known anomalous data, resulting in poorer performance. DAR [7] and BGDA [62] have been proposed to solve this problem. By training these methods to acquire a feature space that explicitly distinguishes normal data from the small amount of anomalous data obtained, they improve the detection performance of the small amount of known anomalous data used during training without reducing the detection performance of unknown anomalous data not used during training. This paper also improves the detection performance for known anomalous data by explicitly classifying a small amount of anomalous data in the feature space.

These methods have achieved good performance but they are difficult to apply directly to ASD because many of the techniques used are unique to image data processing. These methods utilize CutMix [63, 7], image rotation [53] and semantic segmentation [62] in order to generate pseudo-anomalous data for training. The nature of the data used to represent images and sounds differs greatly, so many factors must be considered when applying these outlier exposure methods to ASD. In particular, sound and images differ in their approach to adding data. For example, images generated by superimposing different images on top of each are not common in the real world and thus have a high probability of being anomalous. However, sounds generated by the superimposing different sounds may be normal (*e.g.*, adding background noise to a movie scene). Therefore, it is necessary to carefully design the training process with these points in mind.

3 Proposed Method

3.1 Overview

When designing the proposed method, we decided that if it was too difficult to collect anomalous audio data for the target device during the development of an ASD system, the anomaly detection model could be trained using only normal data. But if it was possible to collect a small amount of anomalous data for the target device, the system would allow the use of that data to improve anomaly detection performance. With this in mind, we developed Serial-OE, which is an ASD system consisting of a DNN-based feature extractor and a GMM-based anomaly detector in a serial method [31]. This framework incorporates the powerful feature extraction capability of DNNs for high-dimensional data,

with the high detection performance of generative models, allowing Serial-OE to achieve performance comparable to that of other recently proposed ASD methods [39, 57, 22, 53, 38, 45]. The proposed method also incorporates outlier exposure as a binary classification method for identifying normal and pseudo-anomalous data.

An overview of the proposed method is shown in Figure 1. Incorporating outlier exposure in our proposed method makes it easier to utilize anomalous data when it is available. The proposed method trains one feature extractor DNN for each machine type, and one GMM for each machine ID. Therefore, in the case of the DCASE2020 Task2, there are six machine types, so six feature extractors are trained, and 41 GMMs are trained for use as anomaly detectors since there are 41 machine IDs. Each sample is a mixture of machine operating and environmental sounds. When training the feature extractors, normal and pseudo-anomalous data are used. The normal sound of the target machine type is used as the normal data. The normal sound of the other machine types are used as pseudo-anomalous data. When real anomalous data for a machine is available, this anomalous data is treated in the same way as pseudo-anomalous data. Note that the proposed method assumes that both pseudo-anomalous data and anomalous data are chunks of the same length as normal data. When training the anomaly detectors, only normal data is used. We discuss the validity of using normal sounds from machines other than the target machine type as pseudo-anomalous data. One of the most significant factors that degrade the performance of ASD systems is the misclassification of anomalous and environment sound noises. So, ASD models should be trained to ignore environmental sound noise to prevent misclassification. A model trained by a task that identifies the operation sounds of several different machines with the same environment sound noise will ignore the effect of the environment sound noise and focus on the operation sound to classify [59]. Using this model in ASD reduces misclassification due to noise and improve performance [59]. It has also been shown experimentally that the classification of normal and pseudo-anomalous sounds by this selection improves performance [30, 41, 31]. Therefore, it is reasonable to select data that have different operation sounds from normal data and the same environment sound noise as normal data as pseudo-anomalous data and have them classified.

The proposed method uses the following steps to infer and obtain anomaly scores:

- **Preprocessing:** The input sound is divided into S -second chunks, and each chunk is converted into a mel spectrogram.
- **Feature extractor:** The mel spectrogram is input to a DNN to obtain a D -dimensional feature vector.
- **Anomalous detector:** The D -dimensional feature vector is input to the GMM for the target machine ID to obtain the negative log-likelihood for each chunk.

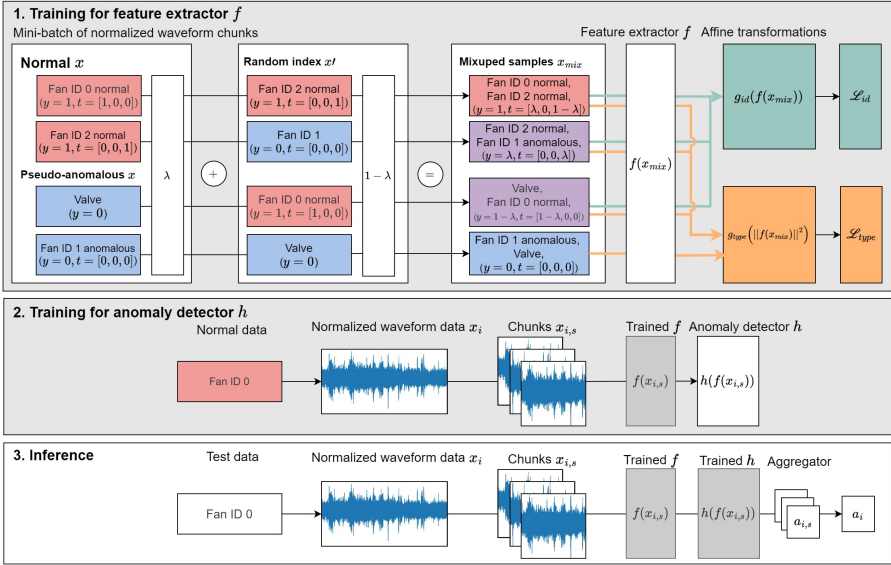


Figure 1: Overview of the proposed Serial-OE method. The gray areas represent the training phase, while the white area represents the inference phase. This example shows how to train the ASD system to detect anomalies in the sound produced by machine type Fan. Section 1 shows the training of feature extractor f , which is trained by applying two loss functions to the features extracted from a mixture of normal and pseudo-anomalous sounds, as sampled by a batch sampler in mini-batches. Section 2 shows the training of anomaly detector h . Since a different anomaly detector h is trained for each machine ID, here we train the anomaly detector using the features of normal audio from Fan ID 0 obtained from pretrained feature extractor f . Section 3 shows the inference process used to obtain the anomaly score, which is computed by dividing the test sample into multiple chunks, each of which is fed into pretrained feature extractor f and trained anomaly detector h , with the resulting anomaly scores aggregated to obtain the final score.

- **Aggregator:** Anomaly scores for all the chunks are obtained, and an aggregator is applied to these scores.

In the following subsections, we describe each step of the training and inference processes in detail.

3.2 Preprocessing

Consider N samples of training data, denoted by $x_i \mid i = 1, 2, \dots, N$. Each x_i is associated with a label $y_i \in \{0, 1\} \mid i = 1, 2, \dots, N$, where 1 is normal data and 0 is pseudo-anomalous data. Normal data refers to the normal sounds of the target machine type, and pseudo-anomalous data refers to the normal sounds of the other machine types. Anomalous data for the target machine

type is treated as $y_i = 0$. Each x_i has a class representing the machine ID, denoted by $t_i \in \{1, 2, \dots, C\} \mid i = 1, 2, \dots, N$, where C is the number of machine ID classes. When applying the loss function, t_i is converted into one-hot vector $t_{i,c} \in \{0, 1\} \mid c = 1, 2, \dots, C$. Before inputting audio data into the model, the amplitude of x is normalized by calculating the mean and variance of the amplitude across $y_i = 1$ (normal sounds). Assuming each x_i represents a sound source of L seconds, these data segments are divided into chunks of S seconds, after which a band-pass filter is applied, converting them into mel spectrograms. The resulting mel spectrogram of each chunk is denoted as $x_{i,s}$ and is input into the feature extractor.

3.3 Loss Functions for Training Feature Extractor

The proposed method uses two loss functions to train feature extractor f to obtain a D dimensional feature vector from $x_{i,s}$. We show in Figure 2 an image of the feature space obtained by the feature extractor f . In the following, we describe the training method for acquiring this feature space. The first loss function $\mathcal{L}_{\text{type}}$ is a binary cross-entropy (BCE) that classifies data as normal or pseudo-anomalous using the norm of the feature vectors. Loss function $\mathcal{L}_{\text{type}}$ is defined as follows:

$$\mathcal{L}_{\text{type}} = -\frac{1}{N} \sum_{i=1}^N \{y_i \log(g_{\text{type}}(\|f(x_{i,s})\|^2)) + (1 - y_i) \log(1 - g_{\text{type}}(\|f(x_{i,s})\|^2))\}, \quad (1)$$

where g_{type} is a function that applies the sigmoid function after performing an affine transformation. By using the norm of the feature vectors for classification, pseudo-anomalous data are distributed near the origin of the hypersphere, while normal data are distributed away from the origin. The loss function trains the model to classify pseudo-anomalous data and normal data with similar environmental sounds so that the norm of features for samples that do not contain operation sounds of normal data is 0. The model will focus on the operation sound, reducing the percentage of misclassification due to differences in environmental sound. In addition, by collecting environmental sound pseudo-anomalous data at a single point, the model prevents misclassification as normal data when data with a distribution that differs significantly from normal data is input. Although anomalous data may not always be distributed near the origin at the time of inference, this loss function is introduced to reduce the risk of treating unknown data that is different from normal data as false positives. Furthermore, adding anomalous data to the pseudo-anomalous data class will explicitly distribute data with known anomalous features near the origin. This is expected to create a feature space that makes it easier for the anomaly detectors to detect anomalies.

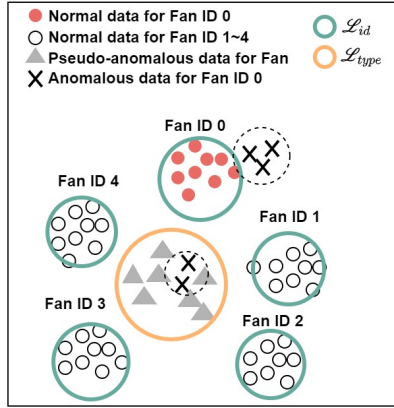


Figure 2: Feature space obtained when using a feature extractor for machine type ‘Fan’. The symbols \circ and \triangle represent the normal and pseudo-anomalous data used for training, respectively. The pseudo-anomalous data are distributed near the origin (*i.e.*, center) of the hypersphere after performing binary classification based on the norm of $\mathcal{L}_{\text{type}}$, while the normal data are distributed farther from the origin. Data for each machine ID is distributed into separate clusters using \mathcal{L}_{id} . For the purpose of illustrating our process when using the obtained feature space, normal data for Fan ID 0 are represented with a red \circ , while anomalous data for Fan ID 0 are denoted by \times . The black dashed circles represent the areas where anomalous data are distributed. It is assumed that anomalous data for Fan ID 0 will cluster around the normal data in the feature space if its characteristics are similar to the normal data (*e.g.* if anomalous data is caused by a slight scratch), and around the pseudo-anomalous data cluster if its characteristics significantly differ from the normal data, suggesting a possible breakdown. The proposed method detects anomalous data by modeling the normal data for Fan ID 0 using a GMM, in relation to such a feature space.

The second loss function \mathcal{L}_{id} is a BCE that classifies normal data by machine ID. Note that the normal data are of the same machine type. The loss function is defined as follows:

$$t'_i = y_i t_{i,c}, \quad (2)$$

$$\mathcal{L}_{\text{id}} = -\frac{1}{C \sum_{i=1}^N y_i} \sum_{i=1}^N \sum_{c=1}^C \{t'_i \log(g_{\text{id}}(f(x_{i,s}))) + (1 - t'_i) \log(1 - g_{\text{id}}(f(x_{i,s})))\}, \quad (3)$$

where g_{id} is a function that applies the sigmoid function after performing an affine transformation. \mathcal{L}_{id} is introduced with the goal of more precisely representing the characteristics of the data from the targeted machine ID, by identifying data with similar sounds. Since the characteristics of the sounds produced by each machine ID are unique, BCE is used instead of cross-entropy. Using BCE allows the classifier to be trained to focus on whether a sound

is from a particular machine, thereby becoming more robust when mixing in other sounds. This will be discussed in more detail in Section 3.4. It is important to note that by multiplying $t_{i,c}$ by y_i in Eq. 2, updates are only made for normal data. Thus, \mathcal{L}_{id} focuses training of the feature extractor only on normal data features.

The final loss function is expressed as follows:

$$\mathcal{L} = \mathcal{L}_{\text{type}} + \alpha \mathcal{L}_{\text{id}}, \quad (4)$$

where α is a hyperparameter to balance the loss functions. However, even if we train the feature extractor to update \mathcal{L} as shown, stable updating is not possible due to several problems. Therefore, we used a few additional techniques, described in Section 3.4 below, to stabilize the training of the feature extractor.

3.4 Strategies for Mini-batch Sampling

When training feature extractors, employing Mixup [65], as proposed in [57], allows the acquisition of intermediate representations between samples, with the aim of improving performance. However, when using the proposed method, two issues prevent the simple application of Mixup. The first issue is the imbalance between the normal data, which consists of data from only one machine type, and the pseudo-anomalous data, which includes data from all of the machines belonging to the five other machine types. Therefore, most of the data is pseudo-anomalous, so when Mixup is applied, Eq. 3 is hardly updated. The second issue is the difference in the scope of each loss function. Eq. 1 focuses on classification of audio data into normal and pseudo-anomalous data, while Eq. 3 focuses on the classification of normal data into machine ID classes. Although mixing up normal and pseudo-anomalous data when using in Eq. 1 is a normal Mixup, Eq. 1 does not allow machine IDs to be defined for pseudo-anomalous data, so it is necessary to consider how Mixup should be used.

To address these issues, first we defined a batch sampler for creating mini-batches in order to deal with the imbalance between normal and pseudo-anomalous data. The mini-batches created by the batch sampler contain half normal data and half pseudo-anomalous data, so if real anomalous data can be utilized, one sample of anomalous data is selected to be included in the mini-batch, reducing the number of pseudo-anomalous samples by one. One epoch is defined as a single update of the normal data. Using a batch sampler not only avoids the problem of an imbalance in data during training, but also reduces the number of iterations, making the training process more efficient. The batch sampler also plays a role in stabilizing the training process; if training is performed without the batch sampler, the proportion of pseudo-anomalous data contained in the mini-batches will increase, and the model

will continue to output values close to 0 for \mathcal{L}_{id} , as the result of an unstable training process. In my experiment conducted without the batch sampler, \mathcal{L}_{id} diverged during the training, and anomaly scores could not be obtained. Therefore, it is important to adjust the balance between the ratio of normal and pseudo-anomalous data in the mini-batches.

Next, to address the difference in the scope of the two loss functions, we devised a method for mixing multiple self-supervised labels as follows:

$$x_i^{\text{mix}} = \lambda x_i + (1 - \lambda)x_j, \quad (5)$$

$$y_i^{\text{mix}} = \lambda y_i + (1 - \lambda)y_j, \quad (6)$$

$$t_i^{\text{mix}} = \lambda y_i t_{i,c} + (1 - \lambda)y_j t_{j,c'}, \quad (7)$$

where $x_i^{\text{mix}}, y_i^{\text{mix}}, t_i^{\text{mix}}$ are samples obtained by Mixup, $x_j, y_j, t_{j,c'}$ are samples obtained by randomly shuffling the mini-batch indices, and $0 \leq \lambda \leq 1$ is the mixing coefficient, randomly sampled from Beta distribution $\mathcal{B}(\beta, \beta)$ as defined for the length of the mini-batch. When using Mixup, three types of samples are obtained:

1. a mixture of pseudo-anomalous data and pseudo-anomalous data,
2. a mixture of normal data and normal data,
3. a mixture of normal data and pseudo-anomalous data.

The mixture in the second case involves data from the same or different machine IDs. If two normal samples x_i and x_j have $t_i \neq t_j$ (different IDs within the same machine type), then mixing them yields a combination of normal data from different machine IDs. For instance, if $y_i t_{i,1} = [1, 0, \dots, 0]$ and $y_j t_{j,2} = [0, 1, \dots, 0]$, then their mixture becomes $t_i^{\text{mix}} = [\lambda, 1 - \lambda, 0, \dots, 0]$. In the third case, if one sample is normal ($y_i = 1$), such as $y_i t_{i,1} = [1, 0, \dots, 0]$, and the other is pseudo-anomalous ($y_j = 0$), such as $y_j t_{j,2} = [0, 0, \dots, 0]$, then their mixture yields $t_i^{\text{mix}} = [\lambda, 0, \dots, 0]$. The Mixup-generated samples $x_i^{\text{mix}}, y_i^{\text{mix}}, t_i^{\text{mix}}$ are applied during the training of the feature extractor by substituting them into Eqs. 1 and 3. First, let us consider Eq. 1, which can be updated in all cases since we only need to consider machine type y_i in this loss function. Next, consider Eq. 3. A sample that is a mixture of pseudo-anomalous data and pseudo-anomalous data does not affect the update of Eq. 3 since $y_i = 0$. Samples containing a mix of normal and normal data or a mix of normal and pseudo-anomalous data are used to train the feature extractor to detect the proportion of normal data, since $0 \leq y_i^{\text{mix}} \leq 1$. In other words, the model is trained to focus on whether or not the sounds subject to ASD are included. The model is expected to detect normal data when normal sounds are included, even when the SNR of the environmental sound changes.

3.5 Training Method for Anomaly Detector

A feature vector which has D dimensional features is obtained from trained feature extractor f . All of the audio samples of normal data for a particular machine ID in the training data are used to train the GMM functioning as anomaly detector h . If x represents L seconds of normal data samples, the samples used to train the GMM are divided into chunks of $M = \lceil 2L/S \rceil$, $M \in \mathbb{N}$ samples, allowing for half duplication. Feature extractor f is then applied to these chunks to obtain the features used to train the GMM. Therefore, M times more samples of data for each machine ID are used to train the GMM. Since M is assumed to be sufficiently large compared to dimensional features D in this study, the full covariance matrix is used to calculate the covariance of the GMM.

3.6 Method of Aggregating Anomaly Scores

During inference, each x_i is divided into M chunks of S seconds, similar to the training process for the GMM. As a result, M anomaly score samples are obtained for each x_i . ASD can detect both stationary anomalous sounds, such as the squealing of a belt or the sound of a broken propeller, and non-stationary anomalous sounds, such as a whistle or intermittent noise from a periodically operating machine. We explored various methods of effectively aggregating these anomaly scores. The ‘averaging’ aggregation method tends to highlight stationary sounds but may dilute the impact of non-stationary anomalies. Conversely, selecting the maximum anomaly score can lead to overlooking stationary sounds, or may bias the results toward anomalies that occur in particular data segments, especially if a sound event spans multiple segments.

To mitigate these issues, we propose an intermediate aggregation strategy. This method involves sorting all of the chunks by their anomaly scores in descending order and then calculating the average score of those exceeding the median score. Thus, our method of calculating anomaly scores, including our aggregation method, is as follows:

$$A_i = \text{sorted}(a_{i,1}, a_{i,2}, \dots, a_{i,M}) \text{ in descending order,}$$

$$a_i = \frac{1}{\lceil \frac{M}{2} \rceil} \sum_{m=1}^{\lceil \frac{M}{2} \rceil} A_i[m] \quad \text{where } A_i[m] > \text{median}(A_i), \quad (8)$$

where A_i is the list of anomaly scores $a_{i,m} \mid m = 1, 2, \dots, M$ for each chunk, and where a_i is the anomaly score. This approach presumes that both stationary and non-stationary anomalies are present in at least half of the sample’s length. For samples whose length significantly deviates from the dataset norm, we considered using quartiles to achieve a more nuanced threshold, especially

when the length of the input data sample varies widely from that of the typical dataset sample.

4 Experimental Evaluations

4.1 Dataset

We evaluated the performance of the proposed method using the DCASE2020 Task2 dataset, a representative ASD dataset which has also been used to evaluate other recently proposed ASD methods. Similar datasets include DCASE2021 Task2, DCASE2022 Task2, and DCASE2023 Task2 [22, 10, 9], however these subsequent datasets focus on domain shift control, therefore the DCASE2021 Task2 and later datasets have a hierarchical structure which is not necessarily segmented to reflect the audio data of particular target machines. This difference complicates the discussion when evaluating the proposed method and examining the factors that led to performance improvement. Therefore, we used DCASE2020 Task2 data in this study, since its focus is improving ASD performance through the use of anomalous data.

The DCASE2020 Task2 dataset consists of two datasets, MIMII [43] and ToyADMOS [27], both of which consist of machine operation and environmental sounds. The machine operation sounds were collected in a quiet environment under several different sound collection conditions, while the environmental sounds were collected in a real factory under several different conditions. The task organizers mixed these two datasets at various SNRs to generate audio samples. Each sample was 10 seconds long and was collected on one channel at 16,000 Hz. Six machine types were used: Fan, Pump, Slider, and Valve from MIMII, and ToyCar and ToyConveyor from ToyADMOS. We evaluated our proposed method using the DCASE2020 Task2 development set, which is the same evaluation setting used in most previous studies [33, 57, 3, 64, 14, 4]. DCASE2020 Task2 evaluation set was used for development purposes, to determine the hyperparameters of the model. The data for each machine ID contains about 1,000 samples of normal sounds as training data, and 200 samples each of normal and anomalous sounds as test data. There are three or four varieties of anomalous sounds for each machine type, but it is impossible to determine from the evaluation set what type of anomaly is contained in the anomalous data. When training each of the 41 anomaly detectors h , feature vectors obtained from all of the normal data for each machine ID were utilized as training data for the respective detectors.

4.2 Experimental Conditions

For the sake of specific discussion, this subsection describes the development of a model of a machine type “Fan.” Note that we have implemented the same

model for other machine types as well and used the same hyperparameters for all of them to simplify comparison. When developing an ASD model for a Fan, the normal data includes all of the normal data for all seven of the Fan machine IDs. The pseudo-anomalous data consists of the normal sounds of all of the Pump, Slider, Valve, ToyCar, and ToyConveyor machines.

First, as a pre-processing step, the mean and variance for the amplitude of all of the normal Fan data were calculated, and the amplitude of the entire training dataset was normalized. The normalized data was then randomly split into chunks of $S = 2.0$ seconds, and features were extracted using a Mel filter bank with 224 frequency bins, a window size of 128 ms, and a hop size of 16 ms, covering a frequency range from 50–7,800 Hz. We analyzed the spectrograms of the training data and found that the machine operating sounds were concentrated in the higher frequency range. We applied a band-pass filter to ensure that our model focuses on the most relevant features, the frequency range where the power was concentrated (50–7,800 Hz) across all machines. The obtained features were applied to EfficientNet-b0 [61], a powerful convolutional neural network (CNN)-based feature extractor pretrained using ImageNet [5], to obtain a $D = 128$ dimension feature vector. When training EfficientNet-b0, the learning rate was set to 0.001, the optimization method was AdamW [34], the scheduler was OneCycleLR [52], the number of epochs was 100, the size of a mini-batch was 128, α in Eq. 4 was set at 10.0, and the hyperparameters of Mixup’s beta distribution $\mathcal{B}(\beta, \beta)$ were set to $\beta = 0.2$. After training the feature extractor, the GMM of anomaly detector h was trained. Experiments using the DCASE2020 Task2 evaluation set suggested setting the number of GMM mixtures to 2. During inference, an $L = 10$ second sound source was divided into ten segments of $S = 2.0$ seconds, allowing for overlap. The hyperparameters were used, which showed the highest performance in the DCASE2020 Task2 evaluation set. Finally, an aggregator was applied to all of the segments to compute an anomaly score for the target sound source.

4.3 Evaluation of ASD Performance

The performance of the proposed method was compared with the performance of methods similar to the proposed method, as well as with the performance of recently proposed state-of-the-art methods. The methods used in our performance comparison can be briefly described as follows:

GMADE+SSC [11, 12] is the method that came first in the DCASE2020 Task2 Challenge [25]. The anomaly scores for this method are calculated by ensemble, using the anomaly scores obtained by the method that uses a masked autoencoder for density estimation for each frame (group) of the mel spectrogram (GMADE) [11] and the anomaly scores obtained by the method that performs two self-supervised classifications for machine ID classification and data augmentation (SSC) [12].

DDCSAD+BCE [30] uses the normal data of the target machine’s target machine ID as normal data. This method uses normal data other than the target machine ID of the target machine and normal data other than the target machine as pseudo-abnormal data. It is important to note that the definitions of ‘normal data’ and ‘pseudo-anomalous data’ differ from those used in the proposed method.

STgram-MFN [33] utilizes features obtained by classifying machine IDs using spectrograms and T-grams extracted using a 1D CNN.

SCAdaCos [57] is a method which trains the feature extractor by classifying the data according to the machine IDs of various machines. The feature vectors generated through this process are then utilized in a GMM. The training of the feature extractor leverages a metric known as Sub-Cluster AdaCos, which is combined with Mixup to produce a feature space optimally suited for ASD tasks. This approach uses a single model to extract feature spaces for each of the machines. A GMM is trained individually for each machine ID to calculate anomaly scores, allowing for tailored anomaly detection for individual machines. The feature extractor’s training involves using 16 subclusters, a number that, along with an equivalent number of GMM mixtures, yielded the best performance when using the development data, so we adopted it. All other training conditions are the same as those set in the original study [57].

SW-WaveNet [3] is a method that integrates spectrograms and wavegrams, the latter being derived from WaveNet, to extract additional features using WaveNet itself. This approach is unique since it is the first to employ WaveNet not as a conventional generative model but as a feature extractor, as WaveNet is able to extract more detailed features from audio data, improving the model’s sound analysis ability.

CLP-SCF [14] adopts a two-phase training methodology. In the initial phase the model is pretrained using a contrastive framework, which involves the application of contrastive learning to utilize the relationships between the sounds produced by a particular machine type and the sounds of the individual machine IDs within that machine type. In the next phase, the model is fine-tuned using CLP-SCF, which introduces self-supervised classification in order to enhance the process of learning the sound features pertinent to ASD. STgram-MFN is utilized as the feature extractor, enabling the model to identify the characteristics of sound data.

Noisy-ArcMix [4] improves the compactness of the within-class distribution by classifying samples virtually synthesized by Mixup to be close to the

normal data distribution by ArcFace [6]. In addition, a new input feature, the temporal attention log-Mel spectrogram (TAggram), derived from the temporal attention block, is introduced into STgram-MFN.

To compare the average ASD performance of each method, we used average AUC (aAUC), which is the mean of the area under the receiver operating characteristic curve (AUC) and partial AUC ($p = 0.1$). To compare stability of performance across machine IDs, which has also been examined in previous studies [33, 14, 10], we used minimum AUC (mAUC) which represents the lowest AUC of a machine ID within a particular machine type. Tables 1 and 2 compare the performance of each method. Results for the proposed Serial-OE method were obtained by the authors when using the DCASE2020 Task2 development dataset, while the performances of GMADE+SSC [11, 12], STgram-MFN [33], SW-WaveNet [3], CLP-SCF [14], and Noisy-ArcMix [4] are the results as reported in the respective studies. The performances of DDCSAD+BCE [30], SCAdaCos [57], and the proposed method were calculated five times using different seeds, and the mean and standard deviations of the results are shown.

Table 1: Average performance of compared methods when using the DCASE2020 Task2 dataset. Values represent aAUC (the mean of AUC and partial AUC ($p = 0.1$)).

Methods	Fan	Pump	Slider	Valve	ToyCar	ToyConveyor	Average
GMADE+SSC [11, 12]	80.65	83.27	93.41	91.21	92.72	73.29	86.27
DDCSAD+BCE [30]	83.42±0.26	83.01±0.57	87.27±0.97	97.90±0.83	87.72±0.42	59.72±0.97	83.17±1.58
STgram-MFN [33]	91.51	86.85	98.58	99.04	91.06	69.09	89.35
SCAdaCos [57]	86.63±0.34	90.96±0.46	99.08±0.17	93.33±1.08	95.44±0.06	70.93±0.32	89.39±0.23
SW-WaveNet [3]	94.54	84.98	96.77	98.14	92.85	74.70	90.33
CLP-SCF [14]	95.11	91.18	98.65	99.70	93.02	69.00	91.12
Noisy-ArcMix [4]	96.83	90.72	98.52	99.85	93.44	72.53	91.98
Serial-OE (Proposed)	93.29±0.07	94.87±0.13	99.56±0.17	99.19±0.10	96.54±0.12	77.77±0.49	93.54±1.00

Table 2: Stability of performance compared methods when using DCASE2020 Task2 dataset. Values represent mAUC (AUC of the lowest-performing machine ID of each machine type).

Methods	Fan	Pump	Slider	Valve	ToyCar	ToyConveyor	Average
DDCSAD+BCE [30]	69.58±1.95	71.33±0.93	79.12±1.89	95.19±3.02	74.90±1.13	54.96±1.27	74.18±1.71
STgram-MFN [33]	81.39	83.48	98.22	98.83	83.07	64.16	84.86
SCAdaCos [57]	76.55±0.93	87.83±0.38	99.09±0.18	90.31±1.16	94.24±0.26	69.25±0.71	86.88±0.32
CLP-SCF [14]	88.27	87.27	98.28	99.58	86.87	65.46	87.62
Noisy-ArcMix [4]	92.67	91.17	97.96	99.89	88.81	68.18	89.78
Serial-OE (Proposed)	84.12±0.33	94.17±0.24	99.54±0.17	99.13±0.19	94.81±0.42	70.48±1.41	90.38±1.39

Table 1 shows that the proposed method outperformed GMADE+SSC for all of the machine types. The single system performed better than the ensemble method that combined multiple systems. Table 1 also shows that the proposed method outperformed all of the conventional methods for all of the machine types except Valve and Fan. The operation sound of the Valve is a clicking sound, and the operation sound of the Fan is the sound of a propeller cutting through the wind, so there are almost no similarities between the two machines.

Compared to the best performance of the conventional method, Noisy-ArcMix, Valve’s performance difference is 0.66 points in aAUC, which is slight, and aAUC is also relatively high at 99.19%. While Fan’s performance difference is 3.54 points in aAUC, indicating a moderate difference, with an aAUC of 93.29%. The sound of the Fan is the sound of a propeller cutting through the wind, and compared to other machines, it is difficult to distinguish between the operation sound and environmental noise. Investigating the adverse effects of Fan on the proposed method is a future issue. Although compared to Noisy-ArcMix, the proposed method improves aAUC by 1.7% on average. When comparing the proposed method’s performance with that of SCAdaCos, which is similar to the proposed method, the proposed method outperformed SCAdaCos in terms of aAUC for all machine types. Noisy-ArcMix, SCAdaCos, and the proposed method are the same in that they are trained to detect trivial changes in normal data using Mixup for the classification task. The proposed method differs because it trains one model for one machine type and devises a sampling strategy where irrelevant machine types are regarded as pseudo-anomalous data. This sampling strategy is considered effective when compared to the strategy of training one model for all machine types in that it not only ignores background noise specific to the machine being monitored but also explicitly identifies mixed sounds as anomalies by being specific to the machine being monitored.

Table 2 shows the stability of the performance of each method across machine IDs. Even if the reported aAUC is high, if the mAUC is low it indicates that some anomalous data could not be detected for a specific machine ID. Therefore, if a method’s mAUC is high, its ASD performance can be evaluated as stable. Compared to Noisy-ArcMix, which was also the best performing conventional method in this stability evaluation, the proposed method achieved an average mAUC that was 0.7% higher. Compared to SCAdaCos, a similar method to the proposed method, the proposed method also achieved higher stability performance in mAUC for all machine types.

These results show that the proposed method performs well not only on average, but also across many machine IDs. All the methods shown in Tables 1 and 2, except for DDCSAD+BCE and Serial-OE, are based on the assumption that only normal data is used. Therefore, even if anomalous data can be obtained, it is necessary to incorporate additional techniques to use them. The proposed method is superior because it can easily and efficiently handle anomalous data. This point is discussed in Section 4.6.

4.4 Ablation Studies

To investigate the reasons for the superior performance achieved by the proposed method, we performed ablation studies using aAUC as the evaluation metric in order to compare the average performance of each ASD method. To allow qualitative evaluation, the feature space was visualized using t-SNE [35].

4.4.1 Evaluation of loss functions used for machine ID classification

Table 3 shows the results in terms of aAUC when the feature extractor was trained using each of the loss functions. Figure 3 shows t-SNE visualizations (perplexity= 40) of the feature space obtained when using each loss function with the Fan DCASE2020 Task2 development set. Table 3 and Figure 3 compare the most suitable loss function for \mathcal{L}_{id} in Eq. 3. GMADE+SSC was excluded from the analysis because it is an ensemble system. The proposed method and DDCSAD+BCE use BCE as the loss function for machine ID classification, and SCAdaCos uses SCAdaCos as the loss function for machine ID classification. The other methods, STgram-MFN, SW-WaveNet, CLP-SCF, and Noisy-ArcMix, use ArcFace as the loss function for machine ID classification. We compared the proposed method using these loss functions for machine ID classification previously employed in conventional methods as Eq. 3. We experimented with the proposed method by changing the BCE used for \mathcal{L}_{id} in Eq. 3 to cross-entropy, SCAdaCos, and ArcFace to evaluate the differences in the resulting feature spaces. Note that Mixup and \mathcal{L}_{type} in Eq. 1 were also applied in all versions of the system. The samples affected by changing the loss function in Eq. 3 contain two types of sounds: mixtures of normal data and normal data, and mixtures of normal data and pseudo-anomalous data. Focusing on these samples, we discuss the effects of using each loss function. BCE acquires features by focusing on whether the input sound contains the sound of the target machine ID or not, cross-entropy, SCAdaCos, and ArcFace acquire features for classification based on the class to which the input sample belongs, and ArcFace and SCAdaCos acquire features that minimize intra-class variance, and SCAdaCos uses sub-clusters, so that a machine ID which has several different features can be classified according to each feature. Cross-entropy, SCAdaCos, and ArcFace require the sum of the probabilities of each class to be one, due to the characteristics of these loss functions. Therefore, when normal and pseudo-anomalous data were mixed, the probability of the normal data was set to 1. The usual mixing ratio of each sample was used when mixing samples containing normal data and normal data.

Table 3: Average performance of the proposed method when using various loss functions for \mathcal{L}_{id} for DCASE2020 Task2 dataset. Values represent aAUC (the mean of AUC and partial AUC ($p = 0.1$)).

Loss function	Fan	Pump	Slider	Valve	ToyCar	ToyConveyor	Average
BCE (Proposed)	93.29±0.07	94.87±0.13	99.56±0.17	99.19±0.10	96.54±0.12	77.77±0.49	93.54±1.00
Cross-entropy	88.72±0.21	91.42±0.38	98.60±0.25	96.37±0.66	92.12±0.79	69.29±0.32	89.42±0.18
SCAdaCos [57]	84.83±0.52	88.72±0.59	97.81±0.68	88.97±2.68	84.11±0.79	63.38±1.01	84.64±0.46
ArcFace [6]	85.09±0.19	85.00±0.83	91.18±4.77	91.43±2.58	66.33±2.12	65.50±0.60	80.75±1.04

The results provided in Table 3 show that the best ASD performance was obtained for all machine types when BCE was used as the loss function for \mathcal{L}_{id} in the proposed method. We will first discuss the reasons for the higher

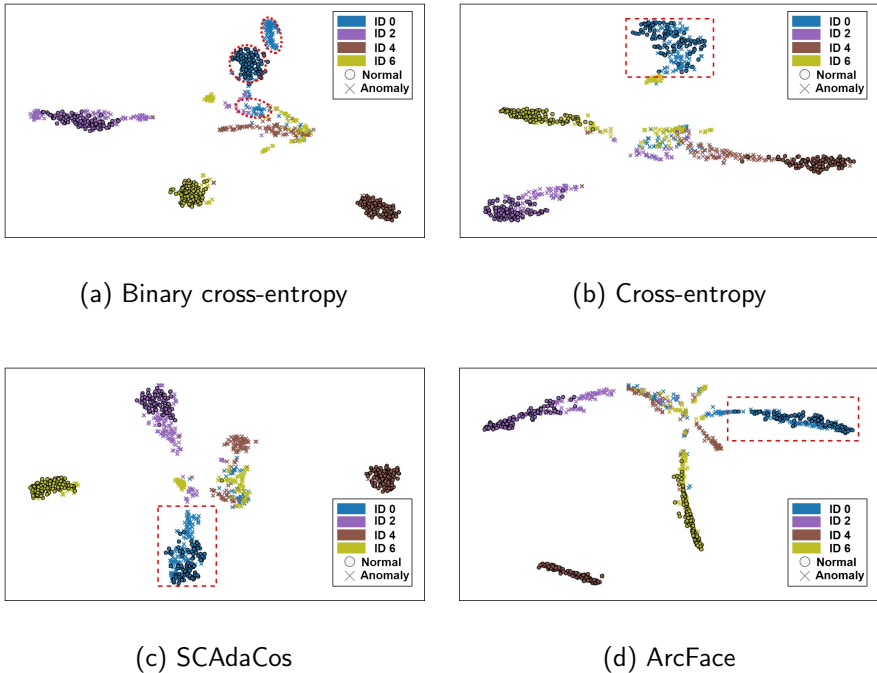


Figure 3: Visualizations of data for machine type Fan using t-SNE, when varying loss function \mathcal{L}_{id} of the proposed method. (a) Feature space distribution obtained using BCE. (b) Feature space distribution obtained using cross-entropy. (c) Feature space distribution obtained using SCAdaCos [57]. (d) Feature space distribution obtained using ArcFace [6]. The symbols \circ and \times denote normal and anomalous sounds, respectively. The normal and anomalous feature space distributions for machine ID 0 are encircled by red dotted lines.

performance of BCE compared to cross-entropy. Comparing the visualization when using BCE in Figure 3a and the visualization when using cross-entropy in Figure 3b, we can see that the classes of normal data are more tightly clustered in the BCE visualization at ID 0, shown circled with a red dotted line. Considering the characteristics of the loss function, it is thought that using features obtained by focusing on whether or not a particular machine ID is included is more effective for ASD than using features obtained by classifying machine IDs, as in the case of cross-entropy.

Next, we compare the performance of the proposed method when using SCAdaCos and ArcFace as the loss function for \mathcal{L}_{id} . Table 3 shows that ASD performance decreases when either SCAdaCos or ArcFace is used. Comparing the visualizations for BCE in Figure 3a, and for SCAdaCos in Figure 3c and ArcFace in Figure 3d, we see that both SCAdaCos and ArcFace visualizations

contain more anomalous data in the normal data cluster for machine ID 0 (red dotted line) than when using BCE. Since both SCAdaCos and ArcFace utilize $\mathcal{L}_{\text{type}}$ to perform classification, using the norm of the feature vectors, its lower performance may be due to the increased difficulty of performing angle-based classification. Therefore, we need to carefully combine a loss function using the norm of feature vectors with angle-based metric training [59].

These results suggest that which loss function is used when training the feature extractor is important for distinguishing whether the sound of the target machine is included in the input sound.

4.4.2 Evaluation of training methods for feature extractors

Table 4 shows the results of our comparison of the performance of the proposed method when each element incorporated during training of the feature extractor is deleted. First, focusing on the average, it is suggested that all three of these elements are effective because the performance of the proposed method is degraded when any of the elements are deleted compared to the proposed method, and all methods improve performance compared to only \mathcal{L}_{id} . In addition, when comparing the performance of the two elements combined and the performance of the individual elements, all combinations improved performance to the same or greater extent. These comparisons of the impact of each factor on ASD performance reveal that Mixup has the largest positive impact on performance, followed by pre-training of the weights using ImageNet, and then loss function $\mathcal{L}_{\text{type}}$, in that order.

Table 4: Average performance of ablation studies for elements other than \mathcal{L}_{id} of the proposed method when using the DCASE2020 Task2 dataset. Values represent aAUC (the mean of AUC and partial AUC ($p = 0.1$)).

Method	$\mathcal{L}_{\text{type}}$ weight	Mixup	Fan	Pump	Slider	Valve	ToyCar	ToyConveyor	Average
Serial-OE	✓	✓	93.29±0.07	94.87±0.13	99.56±0.17	99.19±0.10	96.54±0.12	77.77±0.49	93.54±1.00
w/o $\mathcal{L}_{\text{type}}$	✓	✓	90.46±0.27	92.69±0.36	98.66±0.29	98.01±0.45	94.59±0.32	74.44±1.67	91.48±0.35
w/o weight	✓	✓	87.15±0.33	87.96±0.57	98.70±0.19	95.19±0.68	91.95±0.56	72.50±0.22	88.91±0.09
w/o Mixup	✓	✓	86.52±0.37	91.98±0.76	98.40±0.32	98.83±0.58	94.30±0.44	62.40±0.74	88.74±0.22
w/ $\mathcal{L}_{\text{type}}$	✓		85.15±0.31	85.86±0.76	98.76±0.24	96.41±0.43	91.96±0.22	58.28±0.28	86.07±0.11
w/ weight		✓	83.08±0.55	84.93±0.20	97.08±1.01	99.02±0.39	92.07±0.70	61.16±0.58	86.22±0.35
w/ Mixup		✓	87.05±0.66	88.12±0.40	97.48±0.28	98.91±0.13	91.48±0.63	70.76±0.31	88.97±0.12
only \mathcal{L}_{id}			82.02±0.54	84.80±0.34	97.77±0.61	99.34±0.14	92.44±0.29	56.80±0.63	85.53±0.11

The large improvement in ASD performance when using Mixup indicates that acquiring an intermediate representation of the feature space is effective for improving performance. This allows the distribution of the anomalous data, which has features slightly different from normal data, into clusters farther away from the normal data.

The ASD performance was also improved for all machine types by initializing the feature extractor using weights from models pre-trained on the ImageNet classification task. This finding aligns with previous studies that

have reported improved performance using pre-trained models [45, 49]. Our results also show that the proposed method effectively utilizes weights from models trained on other classification tasks as initial values. However, [58] reported no performance improvement when using pre-trained models. Thus, the effect of using pre-trained weights may depend on the specific method employed.

Finally, focusing on the effect of loss function $\mathcal{L}_{\text{type}}$, it was found that ASD performance improved when $\mathcal{L}_{\text{type}}$ was used as a norm to classify normal and pseudo-anomalous data. This is because $\mathcal{L}_{\text{type}}$ concentrates anomalous data at a point near the origin, improving anomaly detection performance. $\mathcal{L}_{\text{type}}$ can also create intermediate representations between normal and pseudo-anomalous data when used with Mixup, which is also thought to improve performance.

The results of this ablation study indicate that Mixup, initialization with pretrained weights and $\mathcal{L}_{\text{type}}$ used in the proposed method were all effective for improving ASD performance.

4.4.3 Evaluation without machine ID information

The conventional and proposed methods are based on the assumption that training is performed using machine ID information. This information is used because it allows the extraction of features that reflect the target machine’s normal data in greater detail, improving ASD performance. Thus, it is desirable to use machine IDs or similar information when available, however such data is not always available. Table 5 shows ASD performance results when the feature extractor and anomaly detector parameters are trained without the use of machine ID information. When using the proposed method, we set $\lambda = 0$ in Eq. 4. Forty-one classifications were created when SCAdaCos used machine ID information, but only six were created when only machine types were used. When using each method, GMMs were trained for each machine type to create six anomaly detectors. Our results in Table V show that when machine ID information is not used, ASD performance was degraded for both models for all machine types.

Table 5: Average performance of models trained without machine ID information when using the DCASE2020 Task2 dataset. Values represent aAUC (the mean of AUC and partial AUC ($p = 0.1$)).

Method	\mathcal{L}_{id}	$\mathcal{L}_{\text{type}}$	Mixup	Fan	Pump	Slider	Valve	ToyCar	ToyConveyor	Average
Serial-OE	✓	✓	✓	93.29±0.07	94.87±0.13	99.56±0.17	99.19±0.10	96.54±0.12	77.77±0.49	93.54±1.00
SCAdaCos [57]	—	—	✓	86.63±0.34	90.96±0.46	99.08±0.17	93.33±1.08	95.44±0.06	70.93±0.32	89.39±0.23
SCAdaCos (w/o ID)	—	—	✓	75.72±0.51	77.10±0.61	87.44±0.43	76.25±0.55	76.79±0.82	54.83±0.27	74.69±0.38
w/o \mathcal{L}_{id}	—	✓	✓	77.86±0.94	81.93±1.61	88.54±1.65	71.77±1.21	69.66±1.79	54.52±0.85	74.05±0.45
only $\mathcal{L}_{\text{type}}$	—	✓	—	75.20±1.28	71.62±1.62	91.81±0.43	58.84±1.19	61.11±1.86	52.91±0.51	68.58±0.50

As was the case with machine IDs, ASD performance improved when the feature extractor was trained to obtain an intermediate representation between

normal and pseudo-anomalous data using Mixup, confirming the effectiveness of acquisition of intermediate representations.

Next, we compared the performance of the proposed method without machine IDs with that of SCAdaCos without machine IDs and found that the performance of both methods was almost the same in terms of average performance. When only normal data is used, and machine IDs cannot be used, there is no difference between the two methods. However, the proposed method has the advantage of being able to utilize anomalous data, as explained in the following subsection.

4.5 Evaluating Performance when Anomalous Data is Utilized as Pseudo-anomalous Data During Training

Since real anomalous audio data contains features useful for detecting anomalous conditions, it is likely that it can be used to improve performance. Therefore, we also investigated whether the proposed method can achieve higher ASD performance by using a small amount of anomalous data during training. We also evaluated the effectiveness of using anomalous data when information such as machine IDs is not available.

Figure 4a is a graph in which the ratio of anomalous data to training data is represented on the x-axis, while average AUC (aAUC) is represented on the vertical axis. A similar graph is shown in Figure 4b, but minimum AUC (mAUC) is shown on the vertical axis. Here, the volume of training data per machine ID is 1,000 samples, and each sample has a duration of 10 seconds, so a ratio of 1% anomalous data represents 10 samples obtained from 100 seconds of anomalous data. In the experiment, $\{2^i \mid i = 0, 1, \dots, 5\}$ samples of anomalous data were used for training. Because the evaluation data contains approximately 100 anomalous samples, we set $i = 5$ as the maximum to ensure a sufficient amount of data for evaluation. Since the proposed method defines and binarily classifies normal and pseudo-anomalous data, as described in Section 3, a small amount of real anomalous data can be used to improve ASD performance, when it is available. The same experiment was conducted when machine ID information was unavailable. DDCSAD + BCE [30], a conventional method for classifying normal and pseudo-anomalous data, was used for comparison. SCAdaCos [57] performance scores are also shown as a reference. Figures 4a and 4b show that Serial-OE achieved the highest performance in terms of both aAUC and mAUC, respectively, followed by DDCSAD+BCE. These results show that the proposed method can utilize anomalous data more effectively than DDCSAD+BCE to improve ASD performance, in terms of both average performance and performance stability. The proposed method significantly improved ASD performance even when only 0.1% of the anomalous data (1 sample of 10-second data) was added to the normal data as training data, demonstrating the usefulness of utilizing a small amount of anomalous

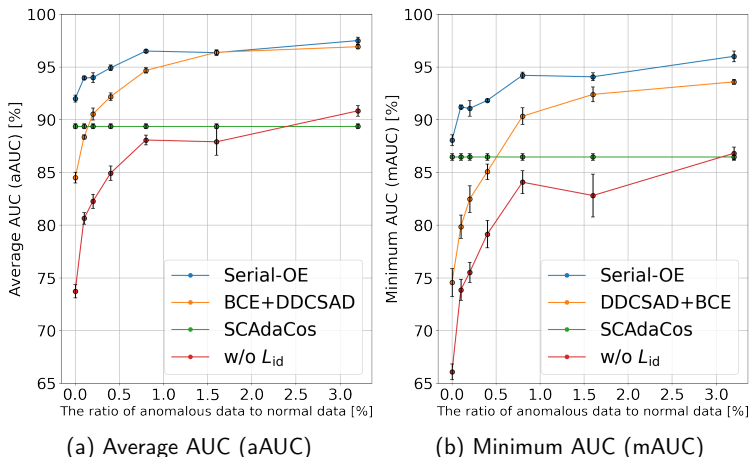


Figure 4: Relationship between the ratio of anomalous to normal data used during training and ASD performance, when using two different performance evaluation metrics: (a) aAUC [%] and (b) mAUC [%]. Error bars represent the standard error obtained from five calculations with different seeds.

data during training. This improvement was equivalent to a 2.4% improvement in performance in terms of aAUC, which is a much easier means of boosting performance than developing a new ASD method.

Next, we compared ASD performance when anomalous data is used for training, but the data does not contain machine IDs. The results in Figures 4a and 4b show that when the percentage of anomalous data is 3.2%, *i.e.*, when only 320 seconds of anomalous data is available, the proposed method’s performance is equal to or better than that of the conventional method using machine IDs (SCAdaCos). Many ASD methods are based on the assumption that machine ID information is a baseline characteristic of audio data used for ASD. However, in real environments it may not always be possible to collect data equivalent to machine IDs. On the other hand, anomalous data will likely be available, since machines monitored by ASD systems presumably experience anomalous states on a somewhat regular basis. Moreover, hours of anomalous data are not needed for training, since only a few seconds to a few minutes of data is sufficient to improve performance. Thus, the proposed method can be used to improve performance by using a small amount of anomalous data, even when machine ID information is unavailable, which is considered to be a major operational advantage. The fact that the hyperparameters do not need to be adjusted for each machine is also considered to be an operational advantage, since this allows the model to be easily adapted if many machines are being monitored, for example.

4.6 Evaluating Performance when Anomalous Data used for Training is Contaminated with Normal Data

Finally, when utilizing a small amount of anomalous data obtained during machine operation to better train the ASD system, depending on the data collection method used, this anomalous data may contaminate the normal training data. Therefore, we evaluated the performance of the proposed method when the small amount of real anomalous data added to the normal data. We randomly contaminated the normal training data with anomalous data, then trained the feature extractor and anomaly detector using this data. In Figure 5a the percentage of contaminated anomalous data used is shown on the x-axis, while performance in terms of aAUC is represented on the y-axis. Similarly, Figure 5b shows the percentage of contaminated anomalous data on the x-axis and stability in terms of mAUC on the y-axis. Figure 5 shows that the more contaminated the normal data was, the worse the performance of all the ASD methods. However, the proposed method achieved the best performance at all contamination levels compared to the other methods, thus the proposed method is more robust against contaminated anomalous data than the other methods. In particular, the proposed method was able to achieve the same performance as SCAdaCos, even when the normal data was contaminated with about 80 seconds of anomalous data. If this amount of contaminated anomalous data was successfully classified and used for training, the aAUC-based performance could be improved by nearly 5%, as shown in Figure 4a. Therefore, the management of anomalous data is very important in terms of improving model performance.

In terms of the ratio of the decrease in anomaly detection scores in relation to the ratio of anomalous data contamination, DDCSAD+BCE had the sharpest decrease in performance. DDCSAD+BCE is considered to be more vulnerable to contaminated data than other methods because it trains a single model for a single machine ID. The proposed method and SCAdaCos were less sensitive to contamination of the anomalous data than the other methods. Comparing the proposed method with SCAdaCos, the proposed method exhibited a more gradual decline in performance as the percentage of contamination with anomalous data increased, compared to SCAdaCos. The proposed method uses one model to infer one machine type, while SCAdaCos uses one model to infer all of the machine types. The influence of contaminated normal data on each machine ID may be reduced by defining the feature space for multiple types of machines. This result suggests that single models used for multiple types of machines may be more robust to contamination of the anomalous data than methods which use a different model for each machine type. On the other hand, the method that improves performance the most when utilizing anomalous data is the one which uses a different model for each machine type, so the relationship between resistance to contamination and increased performance is considered to be a trade-off.




Figure 5: Relationship between the ratio of anomalous to normal data used during training and ASD performance when the anomalous data was contaminated with various amounts of normal data, using two different evaluation metrics: (a) aAUC [%], and (b) mAUC [%]. Error bars represent the standard error obtained from five calculations with different seeds.

5 Limitations

Although this study evaluated the change in ASD performance when real anomalous data was used for training, many different types of anomalous conditions exist. Due to the limitations of the DCASE2020 Task2 dataset, this study did not investigate how effective data from various different anomalous conditions are for improving ASD performance when training the proposed method. Further research on this point would be helpful for finding the most suitable conditions when adopting the proposed Serial-OE method. Furthermore, the proposed method assumes the use of machine IDs, obtained from annotated evaluation data, which is not always available. Since the performance of both the proposed and conventional methods decreases when machine IDs are not used, it is necessary to investigate what types of data can be used as alternative features. This issue was also not investigated due to the limitations of the dataset.

Investigating the impact of domain shift on performance, and the proposed method’s applicability to edge devices were also outside the scope of this study. We believe that investigation of these issues would be advantageous.

6 Conclusion

Most previous studies on improving ASD performance have only focused on the use of normal data. In this study, we explored the use of small amounts of anomalous data for model training, based on the assumption that such data should be readily available since machines monitored by ASD systems are likely to occasionally experience anomalous states. In this study we have proposed a novel ASD method called Serial-OE, which combines the use of outlier exposure with a serial method. The proposed method improves ASD performance when using normal and pseudo-anomalous data and when utilizing a small amount of real anomalous data combined with the pseudo-anomalous data, which resulted in the best ASD performance in our evaluation experiment.

This study also evaluated the performance of each component of the proposed method, including the use of loss functions during training, visualization, impact of machine ID information, and performance when the normal data was contaminated with anomalous data, which has provided useful insights for improving ASD performance, and for applying the proposed method under various conditions.

This study focused not only on comparing the performance of the proposed method with the performance of conventional ASD methods when using normal data from the DCASE2020 Task2 dataset, but also on the practical utilization of anomalous data to further boost ASD performance. It may be obvious that utilizing anomalous data improves detection performance, however it is not obvious that a significant increase in performance can be achieved by using only a 10-second sample of the sound of a machine malfunctioning. Although we are aware of the long-held view that collecting anomalous data is not practical (an assumption which we challenge), this study was nevertheless able to demonstrate the significant benefits that can be obtained by utilizing such data, while also identifying several areas with high potential for achieving further incremental increases in ASD performance. In order to move our proposed system into the real world, various issues must still be addressed, such as countermeasures against domain shift, application to edge devices, restrictions on collectible data, automated collection of anomalous data, and data management. Therefore, we will continue our research in order to develop a highly accurate ASD system that can operate in the real world.

Biographies

Ibuki Kuroyanagi received the B.E. and M.E. degrees from Nagoya University in 2021 and 2023, respectively. He is currently pursuing a Ph.D. degree at the Graduate School of Informatics, Nagoya University. His research interests are anomalous sound detection and sound event detection.

Tomoki Hayashi received the B.E. degree in engineering, and the M.E. and Ph.D. degrees in information science from Nagoya University, Aichi, Japan, in 2014, 2016, and 2019, respectively. He is currently a Postdoctoral Researcher with Nagoya University, Nagoya, Japan and the Chief Operating Officer of Human Dataware Lab. Co., Ltd., Nagoy. He is currently the main Developer of the end-to-end speech processing toolkit ESPnet. His research interests include statistical speech and audio signal processing.

Kazuya Takeda received the B.E., M.E., and Ph.D. degrees from Nagoya University. He was with Advanced Telecommunication Research Laboratories and KDD Research and Development Laboratories. In 1995, he joined Nagoya University, where he started a research group for signal processing applications. He is currently a Professor with the Graduate School of Informatics and the Green Mobility Collaborative Research Center, Nagoya University. His research interest includes behavior signal processing, including driving behavior.

Tomoki Toda received the B.E. degree from Nagoya University, Nagoya, Japan, in 1999, and the M.E. and D.E. degrees from Nara Institute of Science and Technology (NAIST), Ikoma, Japan, in 2001 and 2003, respectively. He was a Research Fellow with the Japan Society for the Promotion of Science, from 2003 to 2005. He was then an Assistant Professor (2005–2011) and an Associate Professor (2011–2015) at NAIST. From 2015, he has been a Professor with the Information Technology Center, Nagoya University. His research interests include statistical approaches to speech processing.

References

- [1] V. Chandola, A. Banerjee, and V. Kumar, “Anomaly Detection: A Survey”, *ACM Computing Surveys*, 41(3), 2009, 58 pages, issn: 0360-0300.
- [2] B. Chen, J. Wan, L. Shu, P. Li, M. Mukherjee, and B. Yin, “Smart Factory of Industry 4.0: Key Technologies, Application Case, and Challenges”, *IEEE Access*, 6, 2018, 6505–19.
- [3] H. Chen, L. Ran, X. Sun, and C. Cai, “SW-WAVENET: Learning Representation from Spectrogram and Wavegram Using Wavenet for Anomalous Sound Detection”, in *Proc. International Conference on Acoustics, Speech and Signal Processing*, 2023, 1–5.
- [4] S. Choi and J.-W. Choi, “Noisy-Arcmix: Additive Noisy Angular Margin Loss Combined With Mixup For Anomalous Sound Detection”, in *Proc. International Conference on Acoustics, Speech and Signal Processing*, 2024, 516–20.
- [5] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “ImageNet: A large-scale hierarchical image database”, in *Proc. Computer Vision and Pattern Recognition*, 2009, 248–55.
- [6] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, “ArcFace: Additive Angular Margin Loss for Deep Face Recognition”, in *Proc. Computer Vision and Pattern Recognition*, 2019, 4685–94.

- [7] C. Ding, G. Pang, and C. Shen, “Catching Both Gray and Black Swans: Open-set Supervised Anomaly Detection”, in *Proc. Computer Vision and Pattern Recognition*, 2022, 7378–88.
- [8] K. Dohi, T. Endo, H. Purohit, R. Tanabe, and Y. Kawaguchi, “Flow-Based Self-Supervised Density Estimation for Anomalous Sound Detection”, in *Proc. International Conference on Acoustics, Speech and Signal Processing*, 2021, 336–40.
- [9] K. Dohi, K. Imoto, N. Harada, D. Niizumi, Y. Koizumi, T. Nishida, H. Purohit, R. Tanabe, T. Endo, and Y. Kawaguchi, “Description and Discussion on DCASE 2023 Challenge Task 2: First-Shot Unsupervised Anomalous Sound Detection for Machine Condition Monitoring”, in *Proc. Detection and Classification of Acoustic Scenes and Events 2023 Workshop*.
- [10] K. Dohi, T. Nishida, H. Purohit, R. Tanabe, T. Endo, M. Yamamoto, Y. Nikaido, and Y. Kawaguchi, “MIMII DG: Sound Dataset for Malfunctioning Industrial Machine Investigation and Inspection for Domain Ggeneralization Task”, in *Proc. Detection and Classification of Acoustic Scenes and Events 2022 Workshop*, 5 pages, 2022.
- [11] R. Giri, F. Cheng, K. Helwani, S. V. Tenneti, U. Isik, and A. Krishnaswamy, “Group Masked Autoencoder Based Density Estimator for Audio Anomaly Detection”, in *Proc. Detection and Classification of Acoustic Scenes and Events 2020 Workshop*, 2020, 51–5.
- [12] R. Giri, S. V. Tenneti, F. Cheng, K. Helwani, U. Isik, and A. Krishnaswamy, “Self-Supervised Classification for Detecting Anomalous Sounds”, in *Proc. Detection and Classification of Acoustic Scenes and Events 2020 Workshop*, 2020, 46–50.
- [13] 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. International Conference on Acoustics, Speech and Signal Processing*, 2023, 1–5.
- [14] J. Guan, F. Xiao, Y. Liu, Q. Zhu, and W. Wang, “Anomalous Sound Detection Using Audio Representation with Machine ID Based Contrastive Learning Pretraining”, in *Proc. International Conference on Acoustics, Speech and Signal Processing*, 2023, 1–5.
- [15] T. Hayashi, T. Komatsu, R. Kondo, T. Toda, and K. Takeda, “Anomalous sound event detection based on WavNnet”, in *Proc. European Signal Processing Conference*, 2018, 2494–8.
- [16] D. Hendrycks, M. Mazeika, and T. Dietterich, “Deep Anomaly Detection with Outlier Exposure”, in *Proc. International Conference on Learning Representations*, 18 pages, 2019.
- [17] D. Hendrycks, M. Mazeika, S. Kadavath, and D. Song, “Using Self-Supervised Learning Can Improve Model Robustness and Uncertainty”, in *Proc. Neural Information Processing Systems*, 2019, 15637–48.
- [18] H. Hojjati and N. Armanfard, “Self-Supervised Acoustic Anomaly Detection Via Contrastive Learning”, in *Proc. International Conference on Acoustics, Speech and Signal Processing*, 2022, 3253–7.
- [19] T. Inoue, P. Vinayavekhin, S. Morikuni, S. Wang, T. Hoang Trong, D. Wood, M. Tatsubori, and R. Tachibana, “Detection of Anomalous Sounds for Machine Condition Monitoring using Classification Confidence”, in *Proc. Detection and Classification of Acoustic Scenes and Events 2020 Workshop*, 2020, 66–70.
- [20] A. Jiang, W.-Q. Zhang, Y. Deng, P. Fan, and J. Liu, “Unsupervised Anomaly Detection and Localization of Machine Audio: A Gan-Based Approach”, in *Proc. International Conference on Acoustics, Speech and Signal Processing*, 2023, 1–5.
- [21] S. Kapka, “ID-Conditioned Auto-Encoder for Unsupervised Anomaly Detection”, in *Proc. Detection and Classification of Acoustic Scenes and Events 2020 Workshop*, 2020, 71–5.

- [22] 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 *Proc. Detection and Classification of Acoustic Scenes and Events 2021 Workshop*, 2021, 186–90.
- [23] Y. Kawaguchi, R. Tanabe, T. Endo, K. Ichige, and K. Hamada, “Anomaly Detection Based on an Ensemble of Dereverberation and Anomalous Sound Extraction”, in *Proc. International Conference on Acoustics, Speech and Signal Processing*, 2019, 865–9.
- [24] P. Kirichenko, P. Izmailov, and A. G. Wilson, “Why Normalizing Flows Fail to Detect Out-of-Distribution Data”, in *Proc. International Conference on Neural Information Processing Systems*, 2020, 20578–89.
- [25] 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 *Proc. Detection and Classification of Acoustic Scenes and Events 2020 Workshop*, 2020, 81–5.
- [26] Y. Koizumi, S. Saito, H. Uematsu, and N. Harada, “Optimizing acoustic feature extractor for anomalous sound detection based on Neyman-Pearson lemma”, in *Proc. European Signal Processing Conference*, 2017, 698–702.
- [27] Y. Koizumi, S. Saito, H. Uematsu, N. Harada, and K. Imoto, “ToyADMOS: A Dataset of Miniature-machine Operating Sounds for Anomalous Sound Detection”, in *Proc. Workshop on Applications of Signal Processing to Audio and Acoustics*, 2019, 308–12.
- [28] Y. Koizumi, S. Saito, H. Uematsu, Y. Kawachi, and N. Harada, “Unsupervised Detection of Anomalous Sound Based on Deep Learning and the Neyman–Pearson Lemma”, *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 27, 2019, 212–24.
- [29] 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 *Proc. Detection and Classification of Acoustic Scenes and Events 2021 Workshop*, 2021, 110–4.
- [30] I. Kuroyanagi, T. Hayashi, K. Takeda, and T. Toda, “Anomalous Sound Detection Using a Binary Classification Model and Class Centroids”, in *Proc. European Signal Processing Conference*, 2021, 1995–9.
- [31] 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. European Signal Processing Conference*, 2022, 294–8.
- [32] W. Liu, D. Cui, Z. Peng, and J. Zhong, “Outlier Detection Algorithm Based on Gaussian Mixture Model”, in *Proc. International Conference on Power, Intelligent Computing and System*, 2019, 488–92.
- [33] Y. Liu, J. Guan, Q. Zhu, and W. Wang, “Anomalous Sound Detection Using Spectral-Temporal Information Fusion”, in *Proc. International Conference on Acoustics, Speech and Signal Processing*, 2022, 816–20.
- [34] I. Loshchilov and F. Hutter, “Decoupled Weight Decay Regularization”, in *Proc. International Conference on Learning Representations*, 8 pages, 2019.
- [35] L. van der Maaten and G. Hinton, “Visualizing Data using t-SNE”, *Journal of Machine Learning Research*, 9(86), 2008, 2579–605.
- [36] M. Meire, J. Zegers, and P. Karsmakers, “Active Learning in Sound-Based Bearing Fault Detection”, in *Proc. Detection and Classification of Acoustic Scenes and Events 2023 Workshop*, 2023, 111–5.
- [37] P. Mishra, R. Verk, D. Fornasier, C. Picciarelli, and G. L. Foresti, “VT-ADL: A Vision Transformer Network for Image Anomaly Detection and Localization”, in *Proc. International Symposium on Industrial Electronics*, 2021, 1–6.

- [38] K. Morita, T. Yano, and K. Tran, “Anomalous Sound Detection Using CNN-based Features by Self Supervised Learning”, *tech. rep.*, 3 pages, DCASE2021 Challenge, 2021.
- [39] R. Müller, S. Illium, and C. Linnhoff-Popien, “Deep Recurrent Interpolation Networks for Anomalous Sound Detection”, in *Proc. International Joint Conference on Neural Networks*, 2021, 1–7.
- [40] G. Papamakarios, T. Pavlakou, and I. Murray, “Masked Autoregressive Flow for Density Estimation”, in *Proc. International Conference on Neural Information Processing Systems*, 2017, 2335–44, ISBN: 9781510860964.
- [41] 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. Detection and Classification of Acoustic Scenes and Events 2020 Workshop*, 2020, 170–4.
- [42] H. Purohit, R. Tanabe, T. Endo, K. Suefusa, Y. Nikaido, and Y. Kawaguchi, “Deep Autoencoding GMM-Based Unsupervised Anomaly Detection in Acoustic Signals and its Hyper-Parameter Optimization”, in *Proc. Detection and Classification of Acoustic Scenes and Events 2020 Workshop*, 2020, 175–9.
- [43] H. Purohit, R. Tanabe, T. Ichige, T. Endo, Y. Nikaido, K. Suefusa, and Y. Kawaguchi, “MIMII Dataset: Sound Dataset for Malfunctioning Industrial Machine Investigation and Inspection”, in *Proc. Detection and Classification of Acoustic Scenes and Events 2019 Workshop*, 2019, 209–13.
- [44] D. Reynolds, “Gaussian Mixture Models”, in *Encyclopedia of Biometrics*, 2009, 659–63, ISBN: 978-0-387-73003-5.
- [45] K. Roth, L. Pemula, J. Zepeda, B. Scholkopf, T. Brox, and P. Gehler, “Towards Total Recall in Industrial Anomaly Detection”, in *Proc. Computer Vision and Pattern Recognition*, 2022, 14298–308.
- [46] L. Ruff, R. Vandermeulen, N. Goernitz, L. Deecke, S. A. Siddiqui, A. Binder, E. Müller, and M. Kloft, “Deep One-Class Classification”, in *Proc. International Conference on Machine Learning*, 2018, 4393–402.
- [47] 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.
- [48] L. Ruff, R. A. Vandermeulen, N. Görnitz, A. Binder, E. Müller, K.-R. Müller, and M. Kloft, “Deep Semi-Supervised Anomaly Detection”, in *Proc. International Conference on Learning Representations*, 23 pages, 2020.
- [49] K. Schmid, F. Ritz, S. Illium, and R. Müller, “Analysis of Feature Representations for Anomalous Sound Detection”, in *Proc. International Conference on Agents and Artificial Intelligence*, 2021, 97–106.
- [50] B. Schölkopf, R. Williamson, A. Smola, J. Shawe-Taylor, and J. Platt, “Support Vector Method for Novelty Detection”, in *Proc. International Conference on Neural Information Processing Systems*, 1999, 582–8.
- [51] Z. Shuyang, T. Heittola, and T. Virtanen, “An Active Learning Method Using Clustering and Committee-Based Sample Selection for Sound Event Classification”, in *Proc. International Workshop on Acoustic Signal Enhancement*, 2018, 116–20.
- [52] 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, 369–86.
- [53] K. Sohn, C.-L. Li, J. Yoon, M. Jin, and T. Pfister, “Learning and Evaluating Representations for Deep One-Class Classification”, in *Proc. International Conference on Learning Representations*, 32 pages, 2021.
- [54] K. Suefusa, T. Nishida, H. Purohit, R. Tanabe, T. Endo, and Y. Kawaguchi, “Anomalous Sound Detection Based on Interpolation Deep Neural Network”, in *Proc. International Conference on Acoustics, Speech and Signal Processing*, 2020, 271–5.

- [55] K. Wilkinghoff, “Combining Multiple Distributions based on Sub-Cluster AdaCos for Anomalous Sound Detection under Domain Shifted Conditions”, in *Proc. Detection and Classification of Acoustic Scenes and Events 2021 Workshop*, 2021, 55–9, ISBN: 978-84-09-36072-7.
- [56] K. Wilkinghoff, “Design Choices for Learning Embeddings from Auxiliary Tasks for Domain Generalization in Anomalous Sound Detection”, in *Proc. International Conference on Acoustics, Speech and Signal Processing*, 2023, 1–5.
- [57] K. Wilkinghoff, “Sub-Cluster AdaCos: Learning Representations for Anomalous Sound Detection”, in *Proc. International Joint Conference on Neural Networks*, 2021, 1–8.
- [58] K. Wilkinghoff and F. Fritz, “On Using Pre-Trained Embeddings for Detecting Anomalous Sounds with Limited Training Data”, in *Proc. European Signal Processing Conference*, 2023, 186–90.
- [59] K. Wilkinghoff and F. Kurth, “Why Do Angular Margin Losses Work Well for Semi-Supervised Anomalous Sound Detection?”, *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 32, 2024, 608–22, ISSN: 2329-9304.
- [60] X. Xia, X. Pan, N. Li, X. He, L. Ma, X. Zhang, and N. Ding, “GAN-Based Anomaly Detection: A Review”, *Neurocomputing*, 493(C), 2022, 497–535, ISSN: 0925-2312.
- [61] Q. Xie, M.-T. Luong, E. Hovy, and Q. V. Le, “Self-training with noisy student improves ImagNnet classification”, in *Proc. Computer Vision and Pattern Recognition*, 2020, 10687–98.
- [62] X. Yao, R. Li, J. Zhang, J. Sun, and C. Zhang, “Explicit Boundary Guided Semi-Push-Pull Contrastive Learning for Supervised Anomaly Detection”, in *Proc. Computer Vision and Pattern Recognition*, 2023, 24490–9.
- [63] S. Yun, D. Han, S. Chun, S. Oh, Y. Yoo, and J. Choe, “CutMix: Regularization Strategy to Train Strong Classifiers With Localizable Features”, in *Proc. International Conference on Computer Vision*, 2019, 6022–31.
- [64] X.-M. Zeng, Y. Song, Z. Zhuo, Y. Zhou, Y.-H. Li, H. Xue, L.-R. Dai, and I. McLoughlin, “Joint Generative-Contrastive Representation Learning for Anomalous Sound Detection”, in *Proc. International Conference on Acoustics, Speech and Signal Processing*, 2023, 1–5.
- [65] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, “Mixup: Beyond Empirical Risk Minimization”, in *Proc. International Conference on Learning Representations*, 13 pages, 2018.