ANOMALOUS SOUND EVENT DETECTION BASED ON ONE-CLASS CLASSIFICATION USING VARIATIONAL AUTOENCODERS AND INTERVAL TYPE-2 FUZZY SETS

Zied Mnasri[a], Stefano Rovetta[b] and Francesco Masulli[b]
[a]Electrical engineering department, ENIT, University Tunis El Manar, BP 37, 1002 Tunis, Tunisia
[b]DIBRIS, Università degli studi di Genova, Via Dodecaneso 35, 16146 Genova, Italy
zied.mnasri@enit.utm.tn, {stefano.rovetta, francesco.masulli}@unige.it

***Abstract*—Audio information is increasingly being used by surveillance systems to improve their effectiveness. Thus, the present paper describes a novel method for detecting anomalous sound events in road traffic monitoring. To detect anomalies, the method combines generative variational autoencoders and interval type-2 fuzzy sets. The reconstruction error of each audio segment is computed using a baseline variational autoencoder, which offers a primary assessment of outlierness through thresholding. An interval type-2 fuzzy membership function with an optimistic/upper component and a pessimistic/lower component is employed to account for the uncertainty associated with this decision-making process. The final class attribution is made by interval comparison, based on a probabilistic technique. The evaluation results obtained after defuzzification reveal that the proposed membership function effectively enhances the performance of the baseline variational autoencoder.**

***Index Terms*—Sound event detection, anomaly detection, audio surveillance, modeling uncertainty, variational autoencoder, interval type-2 fuzzy sets**

I. Introduction

Designing an audio surveillance system depends first on the type of surveillance task, i.e. classification of all detected events, or detection of rare/anomalous/outlier events only. In case of audio event classification, several techniques, basically developed for speech/speaker recognition may be useful, such as generative models (HMM and GMM) and discriminative models (SVM and neural networks) [1]. In the latter case, several anomaly/outlier detection techniques have been applied to audio data since a few years, with different levels of efficiency. These methods can be classified into metric-based e.g. KL-divergence distance [2], reconstruction-based e.g. autoencoders [3], and domain-based e.g. one-class SVM [4].

However, in the case of audio surveillance of road traffic, a challenging issue consists in how to distinguish interesting, yet rare, sounds such as car accidents, or less hazardous events like tire skidding and harsh braking, as outliers in such a noisy environment where: i) practically all events are more or less masked by noise, ii) relevant events, such as car accidents, constitute a very small minority in comparison to normal events, and iii) from an acoustic point of view, anomalous sounds are too diverse to belong to one acoustic category, take for e.g the sound of tire skidding and that of an attacked pedestrian's scream. Therefore, one-class classification seems like an obvious choice to deal with such a problem, where training is performed only on normal samples, so that any anomalous sound would be distinguished as an outlier.

Another level of complexity is due to the vague nature of such broadly defined classes, since backgorund noise is present in all audio clips, whether normal or anomalous. Therefore, membership to any class must be considered affected by a degree of uncertainty. This calls for an explicit treatment of such uncertainty. Type-2 fuzzy sets [5] are a natural choice for this type of problem; in this work we opt for using interval type-2 fuzzy memberships to model classes because, apart from their inherent simplicity and their popularity, they minimise the need for arbitrary modeling decisions about the membership itself.

The rest of the paper is organized as follows: Section II reviews the related work, including methods and applications; Section III presents the utilized methods and the proposed approach; Section IV details the experimental protocol and the obtained results. Finally the work is summarized and commented in the conclusion.

II. Related work

The two main strategies for anomalous sound event detection are static modeling and dynamic modeling. In the first case, signals are embedded in a representation space and anomalies are detected either by distance in the latent representation or by reconstruction error. In the second case, the temporal evolution is evaluated against models of "background". The two strategies may be combined. Some examples are:

a) *One-class classification* – This method, mostly based on one-class support vector machines (OC-SVM), is a classical static anomaly detection tool, that has been applied to anomalous sound event detection in [6]–[8].

Recently, in [9], the authors have proposed an ensemble one-class SVM parallel to an MLP network to calculate the anomaly score for audio events.

b) *Autoencoders* – Several works leveraging deep/variational autoencoders, in some cases using dynamic modeling, have been recently proposed for anomalous sound event detection, e.g. [10]–[12]. In particular, the proposal of Wei et al. [13] at DCASE 2020 challenge-Task 2 is based on a reconstruction autoencoder to calculate the anomaly score through metric learning. Therefore, different types of autoencoders are tested, such as deep autoencoders, variational autoencoders, etc.

c) *Supervised learning* – Several classifiers based on recurrent/convolutional neural networks have been recently proposed for anomalous sound event detection, both statically and dynamically, e.g. [14]–[16]. For instance, in [17], a novel method for detecting road accidents through audio stream analysis is proposed.

More details about the state of the art can be found in the authors' recent survey about machine-learning-based anomalous sound event detection [1].

III. Methods

This work aims to detect anomalous events on roads, either seriously hazardous like car accident, or less hazardous such as tire skidding, harsh braking, etc. Naturally, the proportion of such events is much smaller than that of normal ones, e.g. street noise, or any other unsignificant events like pedestrians passing by, horn blowing, etc. This problem can be illsutrated by Fig. 1 that shows: a) how much small is the proportion of *anomalous* events (1) in comparison to *normal* ones (0), b) how difficult is to discriminate features of each type of events into separate clusters. The latter fact indicates the presence of a commun characteristic for both categories, i.e. background noise. For the aforementioned reasons, two unsupervised learnnig methods based on VAE's reconstrcution error are proposed in this work to split audio signals into *normal* and *anomalous* clusters [1].

A. Baseline method: Anomaly detection based on variational autoencoder with thresholding

The variational autoencoder (VAE) [18] is a special type of autoencoder, a self-supervised neural network that learns the probability distribution parameters of a latent representation of the input to reconstruct the output. The particularity of VAE consists in the encoder layer that stores the parameters of a probability distribution, e.g., mean and variance, representing the input in a latent space. Then, the decoder uses the probability distribution to generate an approximated reconstruction of the input data. A detailed description about the variational bayesian inference and probability density parameter estimation for VAE can be found in [19].

[1]The code details are available at https://github.com/zied-mnasri/Uncertainty-modeling-anomalous-SED

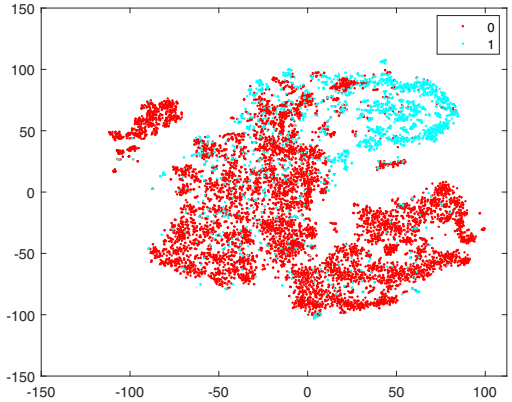

Fig. 1: t-SNE distribution of MFCC and log-Energy features for normal events (0) vs. anomalous events (1)

In the baseline method, the reconstruction error at the output of the VAE is calculated as follows:

$$\epsilon = \sqrt{\frac{\sum_{i=1}^{N}(x_i - \hat{x}_i)^2}{N}}, \tag{1}$$

where $[x_1, \ldots, x_N]$ and $[\hat{x}_1, \ldots, \hat{x}_N]$ are the input feature vector and its reconstruction through the VAE, respectively.

The final decision about outlierness is taken by simply thresholding the VAE reconstruction error, so that above the given threshold (τ) a pattern is deemed not well reconstructed by the VAE model, hence more probably anomalous:

$$\text{Event} = \begin{cases} 0 \text{ (normal)} & \text{if} \quad \epsilon \leq \tau, \\ 1 \text{ (anomalous)} & \text{otherwise.} \end{cases} \tag{2}$$

B. Proposed method: Anomaly detection based on variational autoencoders and interval type-2 fuzzy sets

To improve the performance of the baseline model, a novel method based on combining the baseline VAE to a pair of interval type-2 fuzzy membership functions is proposed.

1) Interval type-2 Fuzzy sets: The fuzzy formalism adopts continuous truth values, or set memberships, instead of a binary membership. It can be used to model partially true concepts, avoiding decisions that are weakly supported. However, according to [5], [20], when it comes to model data or membership uncertainty, the basic type-1 fuzzy sets seem unfit. Therefore, type-2 fuzzy sets have been designed, and in particular Interval-Valued type-2 fuzzy sets (IVFS) have been proposed to deal with such issues. A thourough description of a) the theory and b) the cases where ordinary fuzzy sets should be extended to IVFS, is detailed in [5] and [20], respectively. For instance, in [20], it is demonstrated that in comparison to standard type-1 fuzzy sets, IVFS allow for modeling uncertainty, making the results less specific but more credible.

2) Proposed fuzzy membership function: In this work, we employ a fuzzy membership function made of a lower/pessimistic and an upper/optimisitic components given by (3-4), respectively, where $\alpha_{j,L} = 1$ and $\alpha_{j,U} = 1$ ($j \in \{0, 1\}$)yield a piecewise linear membership, cf. Figure 2a,

and $0 < \alpha_{j,L}, \alpha_{j,U} < 1$ gives the nonlinear shape of the membership functions as in Figure 2b. The footprint of uncertainty (FOU) for each pair of membership functions is represented by the area comprised between the curves of the lower and the upper components, cf. Figure 2. It should be noted that the parameters $\{a_j, b_j\}$ ($j \in \{0,1\}$) in (3,4) are set empirically, following two main criteria: i) The performance of anomaly detection on the training set, ii) The type of uncertainty modeling, i.e. with respect to the fluctuation of a) the VAE error, or b) the membership function itself:

$$\mu_{j,L}(\epsilon) = \begin{cases} (1 - \frac{\epsilon}{a_j})^{\alpha_{j,L}} & \text{if} \quad 0 \leq \epsilon < a_j, \\ 0 & \text{if} \quad \epsilon \geq a_j, \end{cases} \tag{3}$$

and

$$\mu_{j,U}(\epsilon) = \begin{cases} 1 & \text{if} \quad 0 \leq \epsilon < a_j, \\ (\frac{b_j - \epsilon}{a_j})^{\alpha_{j,U}} & \text{if} \quad a_j \leq \epsilon < b_j, \\ 0 & \text{if} \quad \epsilon \geq b_j. \end{cases} \tag{4}$$

where $j \in \{0,1\}$ is the normal/anomalous class label, and $\{\alpha_{j,L}, \alpha_{j,U}\} \subset]0,1]$ define the shape of the membership function. It should be noted that even though $\alpha_{j,L}$ and $\alpha_{j,U}$ are empirically set, we tried to give them more sense by linking their values to the classes' weights, as shown in Table II.

3) Interval comparison method: The goal of interval comparison is to rank real-number intervals or fuzzy numbers based on their boundary values. To compare two intervals $A_i = [a_{i,1}, a_{i,2}]$ and $A_j = [a_{j,1}, a_{j,2}]$, we used the degree of preference of A_i over A_j, denoted $P(A_i > A_j)$, defined in [21] as follows:

$$P(A_i > A_j) = \frac{\max(0, a_{i,2} - a_{j,1}) - \max(0, a_{i,1} - a_{j,2})}{(a_{i,2} - a_{i,1}) + (a_{j,2} - a_{j,1})}, \tag{5}$$

such that $P(A_i > A_j) + P(A_j > A_i) = 1$.

An interval is formed for each class, using the lower and upper membership components, respectively, given by (3,4). Then, interval comparison is performed: the smaller the interval, the lesser the uncertainty, and hence the tighter the membership function. Finally, defuzzification entails matching the event class to the interval chosen as the least favored, as specified by (6), where $A_i = [\mu_{i,L}(\epsilon), \mu_{i,U}(\epsilon)]$, $A_{j \neq i} = [\mu_{j,L}(\epsilon), \mu_{j,U}(\epsilon)]$ $\forall\ j \neq i \in \{0,1\}$:

$$Event = \arg \min_{i,j \in \{0,1\}} \{P(A_i > A_{j \neq i})\}. \tag{6}$$

4) Methodology: The proposed approach proceeds as follows: i) A variational autoencoder model is trained on normal samples only, i.e. ordinary street noise without any anomalous event; ii) For each input audio clip in the test phase, the VAE reconstruction error is calculated between the input features and the reconstructed ones obtained at the output, using (1); iii) For each input audio clip, the VAE recontruction error is used to calculate membership functions as described above; iv) The intervals yielded for normal and anomalous classes, i.e. $[\mu_{0,L}(\epsilon), \mu_{0,U}(\epsilon)]$ and $[\mu_{1,L}(\epsilon), \mu_{1,U}(\epsilon)]$, respectively, are compared using the method previously described; v) Finally, the corresponding event class is yielded by (6).

IV. Experiments and results

A. Audio materials

The MIVIA dataset [17] has been designed for an audio-based road surveillance system. Recordings were realized in a real road environment at 23 locations in the province of Salerno, Italy. Three particular audio events are considered as anomalous, i.e. car crash, tire skidding and harsh braking, whereas all other street sounds are considered as normal events. The total duration of the database is one hour, divided in 57 audio clips of approximately one minute each. However, to increase the number of training samples, and especially to ensure that each audio segment contains at most one anomalous event, each audio clip was segmented into short chunks of 1 sec each with 50% overlap, so that each chunk is labeled as normal or anomalous. Therefore, we opted to make classification by chunks.

B. Experiments and results

Experiments are conducted as follows: First, each audio clip in the database is segmented into overlapping chunks. Secondly, for each chunk, the mel-spectrogram is computed using 64 mel-bands and a Hann window of length 1024 points with 50% hop rate. For each feature, i.e. MFCC, Δ-MFCC and Δ-Δ-MFCC, a matrix containing all vectors at the chunk level is built. Finally a 3D matrix composed of the so-obtained matrices is constructed.

The VAE was implemented using the architecture listed in Table I. The final model was obtained using 100 epochs and a batch size of 32. Training and validation of the autoencoders were processed on 80% of the available data, whereas test was conducted on the remaining 20%.

The evaluation results listed in Table II are expressed in terms of overall accuracy (Acc), calculated for all test samples, and class-wise precision (P_j), recall (R_j) and $F1$ scores ($F1_j$), respectively calculated for each class. Besides, since one-class classification is rather an unsupervised task, we also opted to use metrics that are more suited for this type of problems, such as the *area under the curve* (AUC) and the *partial area under the curve* (p-AUC). The formulas of the aforementioned evaluation metrics are thoroughly detailed in [1].

C. Analysis of results

In Table II, the results of the implemented methods testify the contribution of the proposed fuzzy membership functions to improve anomaly detection. The effects of using fuzzy membership can be listed as follows: a) Overall accuracy rates are enhanced, from 84% for OC-SVM, to 88% for the baseline VAE-Threshold method and up to 97% for the proposed VAE-IVFS with linear/nonlinear piecewise memberships. b) In terms of precision, recall and F1-score, the results of VAE-IVFS are not only the highest, but also the most balanced between *normal* and *anomalous* classes. c) Regarding unsupervised learning-dedicated metrics, i.e. AUC and p-AUC, the results show the improvement registered by VAE-IVFS, especially using the non-linear membership function. In particular, the p-AUC shows that notwithstanding

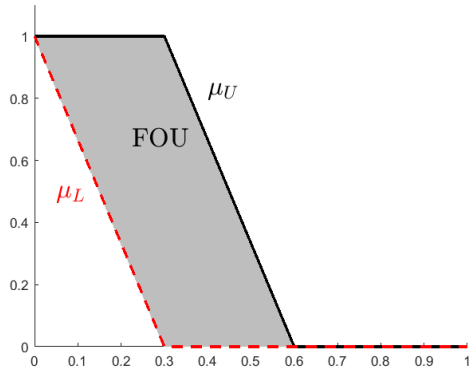

(a) Piecewise linear membership function ($\alpha_{j,L} = \alpha_{j,U} = 1$ in (3-4))

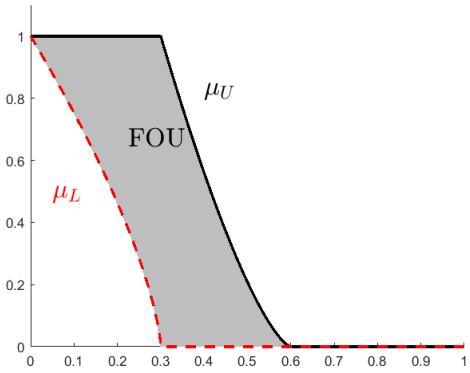

(b) Piecewise nonlinear membership function ($0 < \alpha_{j,L}, \alpha_{j,U} < 1$ in (3-4))

Fig. 2: Proposed 2-component piecewise linear and nonlinear membership functions; the dashed and the continuous lines indicate the lower membership $\mu_L(\epsilon)$ and the upper membership $\mu_U(\epsilon)$ components, respectively, cf. (3,4), (ϵ is the VAE error given by (1)); the area comprised between μ_L and μ_U (in grey) is the footprint of uncertainty (FOU)

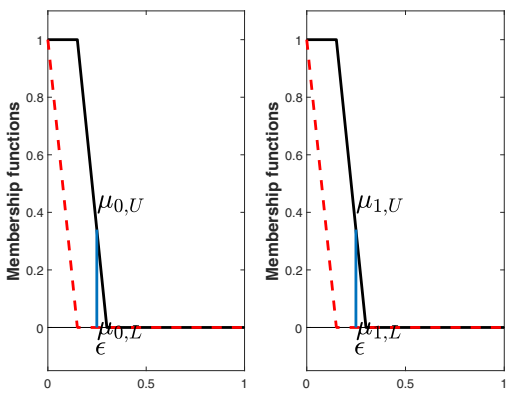

(a) Evaluating uncertainty based on VAE's error using the piecewise linear membership

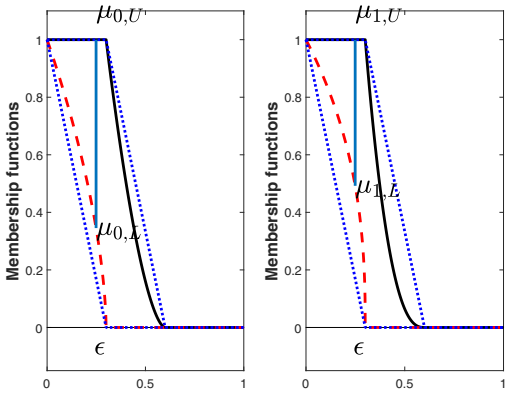

(b) Evaluating uncertainty based on primary membership using the piecewise nonlinear membership function

Fig. 3: Evaluating uncertainty using the proposed 2-component piecewise linear/nonlinear membership functions (ϵ is the VAE error given by (1)); the dashed and the continuous lines indicate the lower membership $\mu_{j,L}(\epsilon)$ and the upper membership $\mu_{j,U}(\epsilon)$ components, respectively, cf. (3,4); the straight vertical lines in all figures indicate the intervals $[\mu_{j,L}(\epsilon), \mu_{j,U}(\epsilon)]$ ($j \in \{0, 1\}$); the dotted line in (b) refers to the maximum area of the FOU region (obtained for $\alpha_{j,L} = \alpha_{j,U} = 1$ in (3,4))

TABLE I: Convolutional variational autoencoder (CVAE) architecture; (*) In the validated architecture, Bottleneck is set to 40

Part of the autoencoder	Hidden layers	Size×number of conv. filters	Transfer layers	Size of stride
Encoder network	Conv2D	(3×3)×32	Relu	(2×2)
	Conv2D	(3×3)×64	Relu	(2×2)
	Conv2D	(3×3)×128	Relu	(2×2)
Code layer	Fully connected	Bottleneck*		
Decoder network	Transposed Conv2D	(3×3)×128	Relu	(2×2)
	Transposed Conv2D	(3×3)×64	Relu	(2×2)
	Transposed Conv2D	(3×3)×32	Relu	(2×2)

the good AUC rates of OC-SVM, it hides a high false positive rate, whereas the p-AUC results for VAE-IVFS using either linear or nonlinear membership functions confirm the good AUC obtained, either for linear or non-linear membership functions. d) The contribution of introducing the interval-valued membership function to correct the performance of the baseline VAE method is clear. Notwithstanding the good rates of the baseline method for the normal class, the simple

TABLE II: Results of anomalous event detection using benchmarking OC-SVM, baseline VAE and the proposed linear membership function (cf. 3,4) using the follwing settings: for OC-SVM, the parameters ν and γ are set to 0.14 and 2.5e-5, respectively (for their high performance); for the baseline VAE and the piecewise-linear/nonlinear methods, $\tau = 0.5$ is the threshold in (2), and $\{\omega_0, \omega_1\}$ are the weights of *normal* and *anomalous* classes, respectively, such that $\omega_1 = 1 - \omega_0$; $p = 0.2$ for p-AUC; (*) NaN value is due to zero correctly estimated samples; bold characters indicate the best results obtained for each method

Method	ω_0	Acc	P_0	P_1	R_0	R_1	$F1_0$	$F1_1$	AUC	p-AUC
OC-SVM		**0.81**	**0.92**	**0.36**	**0.86**	**0.50**	**0.89**	**0.42**	**0.68**	**0.06**
VAE only (Baseline) cf. (1-2) with $\tau = 0.5$		**0.88**	**0.88**	**0.00**	**1.00**	**0.00**	**0.94**	**NaN***	**0.50**	**0.02**
VAE-IVFS with piecewise linear membership function	0.6	0.59	1.00	0.24	0.53	0.99	0.69	0.39	0.76	0.04
cf. (3-4) with $a_j = (1-\omega_j)(1-\omega_0)\tau$ and	0.7	0.79	0.99	0.39	0.76	0.94	0.86	0.55	0.85	0.08
$b_j = 2(1-\omega_j)(1-\omega_0)\tau$; $\alpha_{j,L} = \alpha_{j,U} = 1$;	0.8	0.90	0.98	0.57	0.90	0.90	0.94	0.70	0.90	0.14
$j \in \{0, 1\}$	0.9	**0.97**	**0.97**	**1.00**	**1.00**	**0.77**	**0.98**	**0.87**	**0.89**	**0.12**
VAE-IVFS with piecewise nonlinear membership function	0.6	0.59	0.99	0.24	0.54	0.95	0.69	0.39	0.74	0.04
cf. (3-4) with $a_j = (1-\omega_j)(1-\omega_0)\tau$ and	0.7	0.64	0.99	0.25	0.59	0.97	0.74	0.40	0.78	0.05
$b_j = 2(1-\omega_j)(1-\omega_0)\tau$; $\alpha_{j,L} = \omega_j$ and $\alpha_{j,U} = \frac{1}{\omega_j}$;	0.8	0.91	0.98	0.62	0.91	0.89	0.94	0.73	0.90	0.13
$j \in \{0, 1\}$	0.9	**0.97**	**0.97**	**1.00**	**1.00**	**0.80**	**0.98**	**0.89**	**0.90**	**0.12**

thresholding looks unsufficient to detect the anomalous events. e) Finally, the empirical choice and setting of the class weights $\{\omega_0, \omega_1\}$, such that $\omega_1 = 1 - \omega_0$, may help reflecting the proportion of *normal* and *anomalous* events in the dataset. This highlights the one-class classification nature of the proposed method, since the *anomalous* class has just to be treated as a minority, without needing a precise knowledge of its amount.

V. Conclusion

Based on interval-valued type 2 fuzzy sets (IVFS), this work suggested a novel approach of anomaly detection, that takes care of modeling uncertainty when input data are highly noisy. A direct application to road traffic surveillance allows detecting hazardous events such as car accidents.

The results of the various experiments allow drawing the following comments: a) IVFS appear to be more efficient than crisp one-class SVM at detecting anomaly; b) The double evaluation of uncertainty allows modeling it at different levels, due to either the variability of the input features, in case of VAE error-based uncertainty, or the ambiguity of modeling classes for audio signals in a noisy environment, in case of membership value-based uncertainty; c) Finally, the proposed method can be extended to other problems where uncertainty in audio data need to be modeled.

References

[1] Zied Mnasri, Stefano Rovetta, and Francesco Masulli, "Anomalous sound event detection: A survey of machine learning based methods and applications," *Multimedia Tools and Applications*, pp. 1–50, 2021.

[2] Nash Borges and Gerard GL Meyer, "Unsupervised distributional anomaly detection for a self-diagnostic speech activity detector," in *ACISS*. IEEE, 2008, pp. 950–955.

[3] Masataka Yamaguchi, Yuma Koizumi, and Noboru Harada, "Adaflow: Domain-adaptive density estimator with application to anomaly detection and unpaired cross-domain translation," in *ICASSP 2019*. IEEE, 2019, pp. 3647–3651.

[4] John Shawe-Taylor and Nello Cristianini, *Kernel Methods for Pattern Analysis*, Cambridge University Press, 2004.

[5] J. M. Mendel and R. I. B. John, "Type-2 fuzzy sets made simple," *IEEE Transactions on Fuzzy Systems*, vol. 10, no. 2, pp. 117–127, 2002.

[6] Francesco Aurino, Mariano Folla, Francesco Gargiulo, Vincenzo Moscato, Antonio Picariello, and Carlo Sansone, "One-class svm based approach for detecting anomalous audio events," in *ICINCS*. IEEE, 2014, pp. 145–151.

[7] Pasquale Foggia, Alessia Saggese, Nicola Strisciuglio, Mario Vento, and Nicolai Petkov, "Car crashes detection by audio analysis in crowded roads," in *AVSS*. IEEE, 2015, pp. 1–6.

[8] Sarah M Erfani, Sutharshan Rajasegarar, Shanika Karunasekera, and Christopher Leckie, "High-dimensional and large-scale anomaly detection using a linear one-class svm with deep learning," *Pattern Recognition*, vol. 58, pp. 121–134, 2016.

[9] Stefano Rovetta, Zied Mnasri, and Francesco Masulli, "Detection of hazardous road events from audio streams: An ensemble outlier detection approach," in *EAIS*. IEEE, 2020, pp. 1–6.

[10] Erik Marchi, Fabio Vesperini, Stefano Squartini, and Björn Schuller, "Deep recurrent neural network-based autoencoders for acoustic novelty detection," *Computational intelligence and neuroscience*, vol. 2017, 2017.

[11] Yohei Kawaguchi, "Anomaly detection based on feature reconstruction from subsampled audio signals," in *EUSIPCO*. IEEE, 2018, pp. 2524–2528.

[12] Tomoki Hayashi, Tatsuya Komatsu, Reishi Kondo, Tomoki Toda, and Kazuya Takeda, "Anomalous sound event detection based on wavenet," in *EUSIPCO*. IEEE, 2018, pp. 2494–2498.

[13] Qingkai Wei and Yanfang Liu, "Auto-encoder and metric-learning for anomalous sound detection task," Tech. Rep., DCASE2020 Challenge, Tech. Rep, 2020.

[14] Chieh-Chi Kao, Weiran Wang, Ming Sun, and Chao Wang, "R-crnn: Region-based convolutional recurrent neural network for audio event detection," *arXiv preprint arXiv:1808.06627*, 2018.

[15] Huy Phan, Martin Krawczyk-Becker, Timo Gerkmann, and Alfred Mertins, "Weighted and multi-task loss for rare audio event detection," in *ICASSP*. IEEE, 2018, pp. 336–340.

[16] Chieh-Chi Kao, Ming Sun, Weiran Wang, and Chao Wang, "A comparison of pooling methods on lstm models for rare acoustic event classification," in *ICASSP*. IEEE, 2020, pp. 316–320.

[17] Pasquale Foggia, Nicolai Petkov, Alessia Saggese, Nicola Strisciuglio, and Mario Vento, "Audio surveillance of roads: A system for detecting anomalous sounds," *IEEE transactions on intelligent transportation systems*, vol. 17, no. 1, pp. 279–288, 2015.

[18] Ng, Andrew, "Sparse autoencoder," https://web. stanford. edu/class/cs294a/sparseAutoencoder_2011new. pdf, 2011, Online; accessed 29 March 2020.

[19] John Paisley, David Blei, and Michael Jordan, "Variational bayesian inference with stochastic search," *arXiv preprint arXiv:1206.6430*, 2012.

[20] Humberto Bustince, "Interval-valued fuzzy sets in soft computing," *International Journal of Computational Intelligence Systems*, vol. 3, no. 2, pp. 215–222, 2010.

[21] Pavel Sevastianov, "Numerical methods for interval and fuzzy number comparison based on the probabilistic approach and dempster–shafer theory," *Information Sciences*, vol. 177, no. 21, pp. 4645–4661, 2007.