

# CURE DATASET: LADDER NETWORKS FOR AUDIO EVENT CLASSIFICATION

*Harishchandra Dubey<sup>1+</sup>, Dimitra Emmanouilidou<sup>2</sup>, Ivan J. Tashev<sup>2</sup>*

<sup>1</sup>Microsoft Corporation, Redmond, WA, USA

<sup>2</sup>Microsoft Research, Redmond, WA, USA

{hadubey, diemmano, ivantash}@microsoft.com

## ABSTRACT

Audio event classification is an important task for several applications such as surveillance, audio, video and multimedia retrieval etc. There are approximately 340 million people with hearing loss who can't perceive events happening around them. This paper establishes the CURE dataset which contains curated set of specific audio events most relevant for people with hearing loss. It is formatted as 5 sec sound recordings derived from the Freesound project. We propose a ladder network based audio event classifier. We adopted the state-of-the-art convolutional neural network (CNN) embeddings as audio features for this task. We start with signal and feature normalization that aims to reduce the mismatch between different recordings scenarios. Initially, a CNN is trained on weakly labeled Audioset data. Next, the pre-trained model is adopted as feature extractor for proposed CURE corpus. We also explore the performance of extreme learning machine (ELM) and use support vector machine (SVM) as baseline classifier. As a second evaluation set we incorporate ESC-50. Results and discussions validate the superiority of Ladder network over ELM and SVM classifier in terms of robustness and increased classification accuracy.

**Index Terms**— Audio event classification, Convolutional neural networks, Ladder network, Transfer learning, Extreme learning machine, Support vector machine.


## 1. INTRODUCTION

Audio event classification is an emerging research topic with applications in ambient sensing, low-hearing support systems, surveillance and home security monitoring, multimedia information retrieval, and many others [1]. Detecting sound events in daily-life scenarios can be a challenging task: (a) in nature, sound events of interest can exhibit similar temporal and spectral characteristics and can overlap both in time and content; (b) presence of external or inherent interference, including ambient noise, recording noise and room reverberation,

can suppress discriminatory information or alter the acoustic signature of underlying events. These challenges can lead to inevitable mismatch in training and test conditions. In recent years, deep learning advancements, publicly available datasets [2, 3, 4, 5], as well as scientific competitions [6, 7] contributed to rapid developments in the field.

In recent literature, conventional feature extraction approaches, such as Mel-frequency cepstral coefficients (MFCC), even when combined with ensemble classification models, displayed lower discriminatory capabilities than expected; thus exposing the need for more sophisticated features for audio event classification tasks [3]. Several deep neural network (DNN) based approaches have been proposed for feature extraction, among which convolutional neural networks (CNN) have shown increased performance [8, 9]. CNNs provide robust data-driven feature representations that can potentially alleviate the need for excessively complex classification models. Kumar et al. demonstrated this idea in [10] with a deep CNN network originally designed for large-scale image recognition tasks, and expanded with additional layers for the task at hand. Similar observations were made by Aytar et al. [11], and Piczak [3] where CNNs achieved superior performance than feed-forward DNN counterparts and were able to improve classification accuracy from 73% to 83%. Recent work in [12], as part of DCASE challenge, demonstrated the complications of over-fitting when using a fusion of complex feature spaces with deep neural architectures. There are 340 millions individuals with hearing loss and that shows the scope of CURE corpus. This paper presents an audio event classification scheme by exploring architectures that do not assume uniformity between training and test data, and accounts for realistic constraints in model complexity and expected system latency. The contributions of this paper are as follows: (i) establish a new data set that contains a balanced set of audio events, relevant for real-world applications, obtained under varied non-uniform acoustic conditions, and labelled having low-latency applications in mind; (ii) explore vanilla transfer learning (VTL) for extracting efficient audio embeddings from out of domain data without the need for model re-training and additional layer adjustments for in-domain data; (iii) investigate normalization procedures and propose best practices for tackling data mismatch and

<sup>+</sup>The work was done while H. Dubey was an intern at Microsoft Research Labs, Redmond, USA and PhD student at the University of Texas at Dallas, USA. H. Dubey is currently with Microsoft Corporation, Redmond, USA.



**Fig. 1.** (a-b) Proposed framework: Top sub-figure in (a-b) depicts the deep CNN architecture in the vanilla transfer learning pipeline. Audio embeddings were extracted from layer FC1 of the network, trained on out-of-domain data. Layers CC1...CC5 are double convolutional layers and C6 is a single convolutional layer. FC1 layer includes batch normalization and ReLU activation while FC2 has sigmoid activation. The global pooling layer averages segment-level embeddings to utterance-level. The soft-max layer outputs predicts the posteriors for each audio event class. Bottom sub-figure in (a-b) illustrates the decision-making pipeline that leverages audio embedding (FC1) extracted from entire audio clips in evaluation set. (c) Illustration of a two-layer LadderNet, where  $x$  and  $\hat{x}$  are input and reconstructed embeddings,  $y$  is the output label, and  $\hat{y}$  the output of noisy encoder, injected by Gaussian noise  $\mathcal{N}(0, \sigma^2)$ . Decoder paths are characterized by denoising functions  $g^{(\cdot)}$  and denoising costs  $C_d^{(\cdot)}$  at each layer.

achieving increased model-specific performance; (iv) explore Ladder Network (LadderNet) and Extreme Learning Machine (ELM) for audio event classification and propose best system for specific infrastructure needs. Section 2 summarizes the datasets used in this study. Proposed approaches are described in Section 3 and experimental results presented in Section 4. The paper concludes with a summary and discussion in Section 5.

## 2. DATA

### 2.1. AudioSet

AudioSet is a large-scale weakly labeled collection of 10 sec long audio recordings extracted from about 2 million YouTube videos, and grouped into 527 highly unbalanced classes of at least 59 recordings per class [2]. A loose hierarchy was used to categorize the sounds, and audio clips were weakly labeled containing one or more audio events with unknown start and end time. Data were organized into fixed training and validation sets by the authors and a balanced set was preferred for this work. Since many annotators label each clip, usually it is imperative to estimate inter-annotator agreement such as the correlation (or some other measures of agreement) between labels done by different annotators for same clip. AudioSet is a weakly labeled database where inter-annotator correlation was not computed. While there were no provided labels with cross-annotator agreement guarantees (a tedious process for large corpora), subsets of recordings underwent quality assessment by randomly sampling all thematically distinct classes. Depending on the application, the duration of these recordings can be too long to use for audio event classification systems that are required to respond in real-time situations. This corpus was used for out-of-domain feature learning.

### 2.2. Environmental Sound Classification (ESC-50)

ESC-50 is a labeled collection of 2000 public field audio recordings with 5 sec duration, split into 50 semantical classes under the theme of animals, natural soundscapes, human non-speech sounds, domestic sounds and exterior noises [13]. All classes are balanced containing 40 examples each. A crowd sourcing platform was used for data annotation presenting reviewers with 50 choices to choose from. ESC-50 was used in this work as part of in-domain data, with fixed train and test sets provided by the authors.

### 2.3. Proposed CURE Corpus

We established a corpus containing *CUR*ated set of specific audio *Events* (CURE) in which the audio clips were selected from a specific set of audio events. These audio events were chosen to capture most relevant events for people with hearing loss and consists of 13 classes. Unlike AudioSet and ESC-50 dataset, CURE corpus consist of 5s audio recordings that are annotated by human transcribers to generate the ground-truth audio event labels. In this way, it facilitate study of audio events classification using short 5s clip that is not possible with AudioSet and ESC-50. While ESC-50 contains 10s audio recordings for single event, AudioSet has variable duration with one or more events. Thus, CURE corpus is complementary to existing audio event datasets. Datasets used in DCASE challenges [14] target different domains excluding people with hearing loss. CURE corpus is developed for research audio event classifier for 5s recordings on smartphone. Contrary to other available datasets, CURE corpus is designed to contain practically relevant sound events, chosen for their increased sensory effect: their ability to spur strong stimulation effects on users. Audio clips were queried from Freesound online repository, noted for its diversity in data capturing modalities, and variety in acoustical room condi-

tions and environments [15]. Queried sounds were segmented into 5 sec clips without overlap, creating a pool of more than 7,000 audio clips. The choice of 5 sec duration is important and is meant to allow for enough context without blending out audio signatures of various events. It also enables system designing with realistic latency requirements, hence making it a reasonable choice. Audio clips were annotated by three judges using UHRS crowd sourcing [16] into one of 13 classes: fire and smoke alarm, baby crying and baby vocalizations, car horn, hand clapping, clock-alarm, human coughing, dog barking and dog howling, door bell ringing, door knocking, foot steps, glass breaking, ambulance or police siren, and thunderstorm sounds. Judges were restricted to 100 files to avoid bias and ground-truth was obtained when at least two judges agreed on the label. A balanced final set was achieved containing 500 sounds per class. Human accuracy was calculated by accounting for the following cases as errors: when only two judges agreed on a label (counts as 1 error) and when all judges disagreed (counts as 3 errors). CURE dataset is target domain in this paper and hence used for evaluation experiments.

### 3. METHODS

#### 3.1. Data Preparation

Data were re-sampled to 16 kHz, and further treatment was needed for the two in-domain sets. Sound data coming from a variety of sources typically has variability in recording equipment, acoustical environment and noise levels. When designing frameworks that can perform and adapt to realistic applications, it is important to avoid highly uneven conditions; this is also important for human annotators, where acoustic variations may incur unwanted bias and skew the perceived contextual information. Since this can highly affect the new dataset, CURE was pre-treated with perceptual and loudness correction using offset correction and frame-level A-weighted filtering before annotation and usage.

We further propose a feature-level normalization scheme, useful for efficient learning and classification in a variety of settings. A whitening process is applied to both training and test feature data using the training set mean and variance. Whitening was followed by length-normalization  $\hat{\mathbf{x}} = \frac{\mathbf{x}}{\|\mathbf{x}\|}$  that aims towards a compact feature representation lying on a unit hypersphere, where  $\hat{\mathbf{x}}$  is the normalized feature vector of  $\mathbf{x}$ , and  $\|\cdot\|$  represents the vector norm. Combined steps of whitening and length-normalization are simply referred to as normalization, or feature-normalization in this manuscript.

#### 3.2. Feature Extraction

Audio event detection can be done on few corpora as discussed in earlier section. For research discussed in this paper, CURE corpus is our target application i.e., in-domain data. We leveraged AudioSet data for transfer learning based on

Convolutional Neural Network (CNN). Thus, AudioSet is our out-of-domain data as it contains audio extracted from web videos. ESC-50 dataset is used as another evaluation set in addition to CURE corpus. ESC-50 is a out-of-domain evaluation set for research discussed in this paper. Transfer learning refers to the re-usage of knowledge coming from an out-of-domain task to an in-domain related setting. It is an indirect way of reusing the plurality of data and information that a related domain can offer. A deep CNN architecture inspired by the computer vision literature was adopted for vanilla transfer learning (VTL) [10, 17]. The network was trained on the large, balanced, out-of-domain corpus AudioSet (section 2.1). Log Mel-spectral features were extracted over 1.5 sec frames with 50% overlap to train the feature extracting network (Fig. 1a). Segment-level features were pooled to get utterance-level features before the soft-max output layer of dimension 527 (number of out-of-domain classes). Multi-label training loss was computed by averaging the binary cross-entropy loss per class. Network parameters were optimized using the validation set, and the resulting network was adopted for feature extraction at the FC1 layer (Fig. 1a), without the need for re-training or further adjustment to in-domain data. Extracting learned features from layer FC1 (VTL-CNN) is a great enabler for processing variable-length real-time streams of audio, by providing fixed-length embeddings. The final segment-level features were extracted using the in-domain data (CURE) or ESC-50 at layer FC1, and then pooled to utterance-level features for usage by the decision-making models.

#### 3.3. Decision Making

The audio event classification task is seen here as a classification problem. The proposed models were used in conjunction with the accompanied in-domain fixed-length data sets, but are straight-forward to be generalized as pseudo real-time predictors for continuous audio stream processing.

##### 3.3.1. Ladder Network

Ladder network (LadderNet) is a semi-supervised model designed to learn efficiently from small amounts of annotated data and has the ability to extrapolate information from unlabeled data with the assumption of low intra-class distance among feature samples [18]. It comprises of three paths: clean encoder, noisy encoder, and decoder (Fig. 1c). Training and evaluation of the network is performed using the noisy encoder, which regularizes the input by introducing zero-mean Gaussian noise  $\mathcal{N}(0, \sigma^2)$  to the noisy hidden layers and decoder path. Each layer of the decoder combines two signals, one from the layer above, and one from the corresponding layer in the noisy encoder. Lateral skip connections, a notion equivalent to hierarchical latent variable modeling [19], allow for sharing representation between the noisy encoder and corresponding decoder layers. The final prediction and testing

**Table 1.** Comparison of VTL-CNN audio embeddings with state-of-the-art acoustic features evaluated on ESC-50.

Feature	Classifier	Mean Acc.(%)
MFCC + ZCR [13]	Linear SVM	40
Proposed VTL-CNN	Linear SVM	80

is achieved via the clean encoder path. The lower layers of the network focus on accurate reconstruction between input  $x$  and reconstruction  $\hat{x}$ , while the top layers learn to predict the output label from the given input features.

### 3.3.2. Extreme Learning Machine

Extreme Learning Machine (ELM) is a single hidden layer neural network meant for simple and efficient learning [20]. Unlike traditional neural networks, ELMs can work with non-differentiable activation functions and have random hidden layer weights, that leads to analytical solution for the output layer weights. It makes them extremely fast for training, testing, and for generating predictions, compared to backpropagation-based neural networks, or conventional classification models. Despite their simplicity, their strength has been proven against various audio classification tasks [21], motivating their use in this work. ELM was used for predicting audio event labels using the VTL-CNN audio embeddings as input. Each node in the ELM output layer corresponds to the probability of an utterance belonging to the corresponding class. We used Kernel-based ELMs, including radial-basis function (RBF), linear, and polynomial kernels [22].

### 3.3.3. Support Vector Machine

Support Vector Machine (SVM) is a simple classifier that constructs a set of hyperplanes in a high-dimensional space to separate data in different classes. It is based on convex optimization meant for binary classification. A one-vs-all scheme was used for performing multi-class classification out of binary SVMs. Three kernels (linear, polynomial, and RBF) were adopted for comparison studies. Notice that SVMs can be sensitive to the choice of model hyper-parameters which were optimized for each data set and choice of kernel.

## 4. EXPERIMENTS & RESULTS

VTL-CNN audio embeddings of dimension 1024 were extracted from audio clips after pooling the outputs of layer FC1 (Fig. 1a-b). For LadderNet, 5 total layers were used with three hidden layers of dimension [2048, 1024, 256]. We used Gaussian noise of  $\mathcal{N}(0, 0.2^2)$ , with a denoising cost of [1000, 10, 0.1, 0.1, 0.1] for each layer; ReLu activations and batch normalization of size 100 was used in all hidden layers, and the network was trained over 101 epochs. The ELM network had a fixed number of hidden nodes corresponding

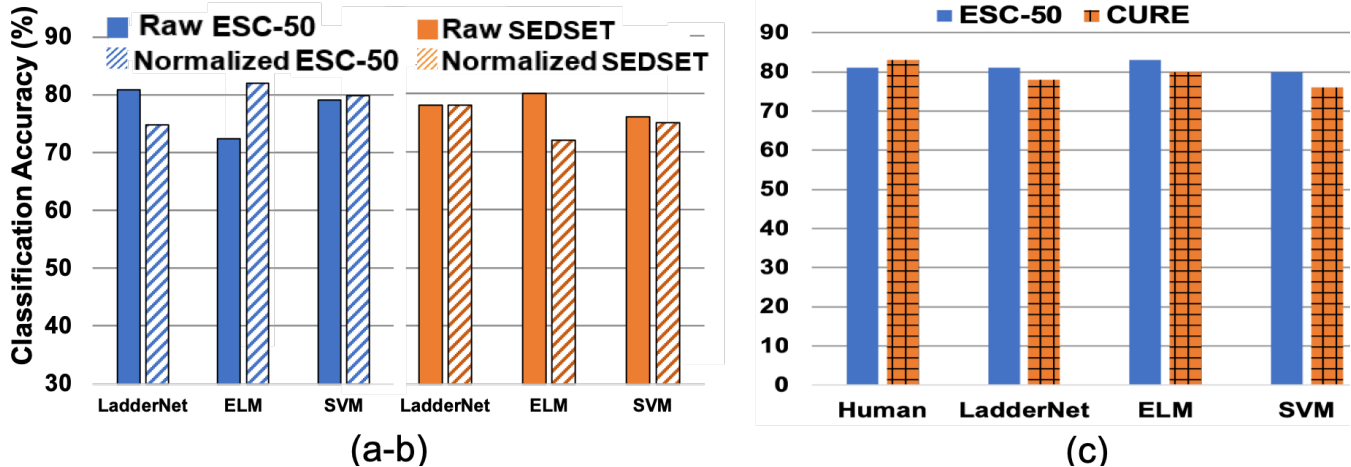
to 10x the input feature dimension, with randomly initialized hidden weights. Polynomial kernels were chosen, after optimizing over all kernels and parameters. The multi-class SVM model was also optimized over different kernels and parameters, with linear kernels outperforming all other setups, evidently due to the feature space being already adequately discriminative. For ESC-50, we used the author’s recommended 5-fold data split and report the average accuracy over all five folds. For CURE, fixed balanced training and test sets were defined and randomly split into 70-30%.

Fig. 2a-b show results on the effect of feature normalization during training and testing for both data sets. Solid bars correspond to results obtained on the original (raw) data, and bars with stripes correspond to the evaluation on feature-normalized data, for ESC-50 (a) and CURE (b) respectively, with y-axis showing (weighted) classification accuracy. The robustness of LadderNet and SVM models is evident against feature normalization; notice however that the SVM performance varied greatly during parameter optimization, which was not the case for LadderNet. In panel (a) ELM appears unable to learn the representation without proper normalization of the input data. This observation can be expected given the non-uniformity of the data. The effect of normalization was not as apparent for the CURE dataset (b), most likely due to the pre-treatment with perceptual correction. Further normalization of CURE at the feature level (striped bars) resulted in negligible changes in terms of accuracy for all models except ELM, which displayed a slight decrease, attributed to potential distortions introduced by mismatches between training and test set data. Fig. 2c shows classification accuracy for the best proposed models, where ESC-50 set was normalized at the feature level, while CURE was not.

The proposed VTL-CNN features are compared with a popular state-of-the-art feature set comprising of spectral-based MFCC and zero crossing rate (MFCC+ZCR) features, as evaluated on one of the datasets (Table 1). In accordance to previous findings [11], a significant improvement is observed, validating the efficacy of transfer-learned VTL-CNN features over domain-restricted spectral-based features.

## 5. DISCUSSION AND CONCLUSIONS

An audio event classification framework was proposed consisting of CNN-based vanilla transfer learning (VTL-CNN). The proposed approach invoked a VTL-CNN model as feature extractor trained on out-of-domain data, which eliminated the need for heavy retraining or introduction of additional and more complex network layers; Audio embeddings obtained from the pre-trained VTL-CNN were found to be more discriminative than equivalent spectral-based features; such robust features allow for efficient transfer learning and can enable the use of simpler classifiers in decision-making frameworks, when invoked with caution. The learned feature space was combined with various decision-making mod-



**Fig. 2.** (a-b) Effect of feature-based normalization on weighted classification accuracy (%) on ESC-50 and CURE dataset; Raw represents the case where audio embeddings were used as it is. Normalized refers case where audio embeddings were length-normalized leading to unit-length feature vectors. Prior to length-normalization, we also performed mean and variance normalization; (c) Best classification accuracy of proposed models for two evaluation dataset.

els, each bringing advantages under certain conditions. When tested on an existing (ESC-50) and a newly created data set (CURE), LadderNet was proven robust to data mismatches which makes it a reliable classifier for unseen acoustic conditions. This network is also flexible enough to incorporate unlabeled data, providing an advantage for practical systems where large amounts of unlabeled data can accumulate over time. The network was further found robust to parameter tuning and data normalization, making it a great candidate for applications where high data-signature mismatch is expected. Less complex classifiers such as Support Vector Machine (SVM) and Extreme Learning Machine (ELM), were found to be more sensitive to parameter tuning; furthermore, in the case of ELM, lack of proper data or feature normalization was proven highly unfavorable for the task of audio event classification under diminished data uniformity assumptions. Such less complex models are, however, better candidates for realistic applications constrained by time, latency and computational complexity, but extra steps for in-domain integration are necessary.

## 6. REFERENCES

- [1] P. K. Atrey, N. C. Maddage, and M. S. Kankanhalli, "Audio based event detection for multimedia surveillance," in *IEEE ICASSP*, 2006.
- [2] J. F. Gemmeke, D. P. W. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, "Audio set: An ontology and human-labeled dataset for audio events," in *IEEE ICASSP*, 2017, pp. 776–780.
- [3] K. J. Piczak, "Environmental sound classification with convolutional neural networks," in *IEEE MLSP*, 2015, pp. 1–6.
- [4] J. Salamon, C. Jacoby, and J. Pablo Bello, "A dataset and taxonomy for urban sound research," in *International Conference on Multimedia*, 2014, pp. 1041–1044.
- [5] A. Mesaros, T. Heittola, and T. Virtanen, "TUT database for acoustic scene classification and sound event detection," in *IEEE EUSIPCO*, 2016, pp. 1128–1132.
- [6] Annamaria Mesaros, Toni Heittola, Emmanouil Benetos, Peter Foster, Mathieu Lagrange, Tuomas Virtanen, and Mark D. Plumbley, "Detection and classification of acoustic scenes and events: Outcome of the DCASE 2016 challenge," *IEEE TASLP*, vol. 26, no. 2, pp. 379–393, 2018.
- [7] Q. Kong, T. Iqbal, Y. Xu, W. Wang, and M. D. Plumbley, "DCASE 2018 Challenge baseline with convolutional neural networks," *arXiv preprint arXiv:1808.00773*, 2018.
- [8] A. Kumar and B. Raj, "Audio event detection using weakly labeled data," in *International Conference on Multimedia*, 2016, pp. 1038–1047.
- [9] E. Cakir, G. Parascandolo, T. Heittola, H. Huttunen, and T. Virtanen, "Convolutional recurrent neural networks for polyphonic sound event detection," *IEEE TASLP*, vol. 25, no. 6, pp. 1291–1303, 2017.
- [10] A. Kumar, M. Khadkevich, and C. Fugen, "Knowledge Transfer from Weakly Labeled Audio using Convolutional Neural Network for Sound Events and Scenes," in *IEEE ICASSP*, 2018, pp. 326–330.

- [11] Y. Aytar, C. Vondrick, and A. Torralba, "Soundnet: Learning sound representations from unlabeled video," in *Advances in Neural Information Processing Systems*, 2016, pp. 892–900.
- [12] J.-W. Jung, H.-S. Heo, I. Yang, S. Yoon, H. Shim, and H. Yu, "DNN-based audio scene classification for DCASE 2017: dual input features, balancing cost, and stochastic data duplication," *System*, vol. 4, no. 5, 2017.
- [13] K. J. Piczak, "ESC: Dataset for environmental sound classification," in *ACM international conference on Multimedia*, 2015, pp. 1015–1018.
- [14] A. Mesaros, T. Heittola, A. Diment, B. Elizalde, A. Shah, E. Vincent, B. Raj, and T. Virtanen, "Dcase 2017 challenge setup: Tasks, datasets and baseline system," in *DCASE 2017-Workshop on Detection and Classification of Acoustic Scenes and Events*, 2017.
- [15] F. Font, G. Roma, and X. Serra, "Freesound technical demo," in *International Conference on Multimedia*, 2013, pp. 411–412.
- [16] "UHRS: the Universal Human Relevance System," <https://www.clickworker.com/2013/09/10/uhrs-was-wo-und-wie-ein-kleiner-guide-fur-neueinsteiger>, 2013, accessed October 30, 2018.
- [17] M. Oquab, L. Bottou, I. Laptev, and J. Sivic, "Learning and transferring mid-level image representations using convolutional neural networks," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 1717–1724.
- [18] A. Rasmus, M. Berglund, M. Honkala, H. Valpola, and T. Raiko, "Semi-supervised learning with ladder networks," in *Advances in Neural Information Processing Systems*, 2015, pp. 3546–3554.
- [19] M. Pezeshki, L. Fan, P. Brakel, A. Courville, and Y. Bengio, "Deconstructing the ladder network architecture," in *International Conference on Machine Learning*, 2016, pp. 2368–2376.
- [20] G.-B. Huang, Q. Zhu, and C. Siew, "Extreme learning machine: theory and applications," *Neurocomputing*, vol. 70, no. 1-3, pp. 489–501, 2006.
- [21] K. Han, D. Yu, and I. J. Tashev, "Speech emotion recognition using deep neural network and extreme learning machine," in *Interspeech*, 2014.
- [22] A. Iosifidis, A. Tefas, and I. Pitas, "Large-scale non-linear facial image classification based on approximate kernel extreme learning machine," in *IEEE International Conference on Image Processing (ICIP)*, 2015, pp. 2449–2453.