An Analysis of State-of-the-art Activation Functions For Supervised Deep Neural Network

Anh Nguyen¹, Khoa Pham², Dat Ngo¹, Thanh Ngo³, Lam Pham⁴

Abstract—This paper provides an analysis of state-of-the-art activation functions with respect to supervised classification of deep neural network. These activation functions comprise of Rectified Linear Units (ReLU), Exponential Linear Unit (ELU), Scaled Exponential Linear Unit (SELU), Gaussian Error Linear Unit (GELU), and the Inverse Square Root Linear Unit (ISRLU). To evaluate, experiments over two deep learning network architectures integrating these activation functions are conducted. The first model, basing on Multilayer Perceptron (MLP), is evaluated with MNIST dataset to perform these activation functions. Meanwhile, the second model, likely VGGish-based architecture, is applied for Acoustic Scene Classification (ASC) Task 1A in DCASE 2018 challenge, thus evaluate whether these activation functions work well in different datasets as well as different network architectures.

Keywords—Multilayer perceptron, convolution neural network, spectrogram, log-mel.

I. INTRODUCTION

A. The Role of Activation Function in Deep Neural Network

Fig. 1: A fully-connected layer

Consider a simple fully-connected layer as shown in Fig. 1, the output column vector $\mathbf{y}^T = [y_1, y_2, \dots, y_n]$ is computed by

$$\mathbf{y} = f(\mathbf{W}\mathbf{x} + \mathbf{b}) \quad (1)$$

where $\mathbf{x}^T = [x_1, x_2, \dots, x_m]$ is the input column vector, matrix $\mathbf{W}$ and bias $\mathbf{b}^T = [b_1, b_2, \dots, b_n]$ are trainable

parameters, and f is activation function selected. From (1), it can be seen that a simple fully-connected layer is separated into two mathematical parts in order. The first part describes a linear mapping, shown by a product between input vector $\mathbf{x}$ and trainable parameters (matrix $\mathbf{W}$), thus added by the bias vector $\mathbf{b}$. The second part, called activation function f , takes the role of non-linear mapping to output $\mathbf{y}_n$. The combination of linear and non-linear mapping mentioned helps the mathematical model operate as a real neural network [1]

B. Backward & Forward Propagation in Deep Neural Network

During training process, input is fed into the network firstly. The input thus goes through network layers, experiences various linear and non-linear mappings, eventually results final output, referred to as inference output. This process is called *Forward Propagation* that shows a direction from input to output. Next, the difference between inference output and expectation (or ground truth) is performed by a loss function (or cost function). Minimizing the loss function is conducted by applying an optimization algorithm, namely *Gradient Descent*. By using *Gradient Descent* optimization, any trainable parameter w_t at certain time t in $\mathbf{W}$ mentioned in (1) is updated by

$$w_t = w_{t-1} - \eta \nabla f(w_{t-1}) \quad (2)$$

where η is learning rate (noting that the learning rate is normally constant in certain epoch), $f(w_{t-1})$ is multi-variable function as mentioned in (1), ∇ is derivative, and $\nabla f(w_{t-1})$ is the gradient value. The subtraction means that a trainable parameter is updated in the opposite direction of gradient vector. In particular, if gradient value is negative, the trainable parameter is added with the absolute of $\nabla f(w)$. if gradient is positive, the trainable parameter is subtracted by the absolute of $\nabla f(w)$. The nearer trainable parameters are to the output layer, the sooner they are updated. Therefore, the process of updating trainable parameters, referred to as *Backward Propagation*, shows the direction from output to input.

C. State-of-the-art Activation Functions Evaluated

This section analyses state-of-the-art activation functions used by a wide range of network architectures, comprising of ReLU, ELU, SELU, GELU and ISRLU. Before going through these activation functions, we firstly analyse conventional activation functions: Tanh [2] and Sigmoid [3] activation functions. As Tanh function shown in Fig. 2a, it helps to map the product of inputs and trainable parameters

¹A. Nguyen, D. Ngo, are with Department of Electronics, Ho Chi Minh City University of Technology, Vietnam National University - Ho Chi Minh, Viet Nam.

²K. Pham is with Department of Computer and Communication Engineering, University of Technology and Education HCMC, Viet Nam.

³T. Ngo is with Department of Electrical Engineering, Da Nang University of Science and Technology, Vietnam.

⁴L. Pham is with Center for Digital Safety & Security, Austrian Institute of Technology, Austria.

Fig. 2: Tanh and Sigmoid functions

to a range from -1 to 1. This function shows saturation regions where input makes output near -1 or 1. When input makes output drop in these regions, *Back Propagation* process shows a potential problem that the gradient is too small, named *vanishing gradients*. When a network suffers from *vanishing gradients*, trainable parameters will not be adjusted, result in stopping learning. As regards Sigmoid activation function, it has the same shape of Tanh function, and shows the same problem of *vanishing gradients* when output drops in saturation regions near 0 or 1.

ReLU [4]: As formula of ReLU activation function and its derivative shown by,

$$f(x) = \begin{cases} x, & \text{if } x > 0. \\ 0, & \text{if } x \leq 0. \end{cases} \quad (3)$$

this function shows simple formula and its gradient is also easy to get, as shown by

$$f'(x) = \begin{cases} 1, & \text{if } x > 0. \\ 0, & \text{if } x \leq 0. \end{cases} \quad (4)$$

As output $f(x)$ is zero for values smaller than zero, otherwise one, it explains why ReLU prevents *vanished gradients*. However, a potential issue can occur with ReLU that is called *dead state* when it is stuck at the left side of zero. If the issue happens to many nodes inside a network, it makes the network low performance. Furthermore, since the sign of gradient for all neurons is the same with ReLU, all the weights of the layer either increase or decrease. However, the ideal gradient update is that some trainable parameters may increase and others may decrease. It can be seen that ReLU helps to reduce cost of computation, speed up training process, prevent *vanishing gradients*, but it shows the other potential problems recently mentioned.

ELU [5]: It can be seen that Tanh or Sigmoid activation function have a two-sided saturation. That means the activation function may saturate in both the positive and the negative direction. In contrast, ReLU activation function shows one-sided saturations. Up to now, it seems that there are two risks: *dead state* by dying neurons with ReLU or a risk for *vanishing gradients* with Sigmoid or Tanh. To tackle this issue, ELU activation function is developed by

$$f(x) = \begin{cases} x, & \text{if } x > 0. \\ \alpha(e^x - 1), & \text{if } x \leq 0. \end{cases} \quad (5)$$

$$f'(x) = \begin{cases} 1, & \text{if } x > 0. \\ \alpha e^x, & \text{if } x \leq 0. \end{cases} \quad (6)$$

where α is constant that normally set to 1. It can be seen that ELU is very exactly similar to ReLU with respect to positive inputs. However, when input is negative, ELU function results in a negative output ruled by an exponent function. Since, ELU produces a negative value instead of 0 as ReLU function, the problem is that weights updated only in one direction cannot occur. By using exponent function when input is negative, there is now a slope on the left side of ELU. Although the slope is usually tiny (e.g. 0.01), but it is there, makes learning happening always and the neural node cannot die. As a result, it helps to solve both *vanishing gradients* and *dead state* problems.

SELU [6]: In terms of neural networks and training process, normalization can occur in three different stages. Firstly, input is normalized before feeding into network model. Secondly, it is batch normalization where the batch of input are normalized at the certain layer of the network model. Finally, it is called internal normalization or self normalization, and SELU is developed for this purpose. If input and output of all layers in a network model follow normal distribution ($\mu = 0, \sigma = 1$), it is helpful to speed up training process by quickly converging. Inspire from this idea, SELU activation function, shown as (7), is designed to be able to make output into a normal distribution. In the other words, the function operation helps to map the mean and variance of output at any layer in the network close to normal distribution.

$$f(x) = \begin{cases} \lambda x, & \text{if } x > 0. \\ \lambda \alpha(e^x - 1), & \text{if } x \leq 0. \end{cases} \quad (7)$$

$$f'(x) = \begin{cases} \lambda, & \text{if } x > 0. \\ \lambda \alpha e^x, & \text{if } x \leq 0. \end{cases} \quad (8)$$

λ is called scale constant. λ and α is set to 1.0507 and 1.6732 that helps to make output distribution nearer normal distribution [6]. However, in order to use SELU, it requires some conditions (a) the neural network consists only a stack of dense layers, (b) all the hidden layers use SELU activation function, (c) the input features should be standardized, (d) the hidden layers weight must be initialized with LeCun normal initialization, and (e) the network must be sequential. The Lecun normal initialization is describe by

$$w_n \sim \mathcal{N}(\mu = 0, \sigma = \sqrt{1/n}) \quad (9)$$

where n is the number of trainable parameters of previous layer. In summary, SELU is similar to ReLU that enables deep neural networks to tackle *vanishing gradients*. In contrast to ReLU, SELU cannot die as they have negative side (same as ELU). In case that conditions of using SELU are satisfied, it is beneficial to network to learn faster and be better than other activation functions mentioned.

GELU [7]: As analyse mentioned above, activations such as ReLU, ELU and SELU enable network become faster and

better convergence than that in Sigmoid or Tanh. However, these cannot cover Dropout regularization [8] that randomly multiplies a few activation functions by 0 at certain nodes in layers. Inspiration from an idea that an activation function would have both these advantages, GELU activation function [9] is developed. Mathematically, if ELU is used, the output of this function is thus multiplied by a zero-one mask. This means that a dropout technique is applied after ELU function. The zero-one mask is stochastically determined, also dependent upon the input x , and simulated by Bernoulli distribution $\Theta(x)$ where $\Theta(x) = P(X \leq x)$, $X \sim \mathcal{N}(0, 1)$. With this idea, the GELU is approximated by

$$f(x) = 0.5x(1 + \tanh[\sqrt{2/\pi}(x + 0.044715x^3)]) \quad (10)$$

Thus, derivative of GELU function used for back propagation is presented by

$$\begin{aligned} f'(x) = & 0.5 \tanh(0.0356774x^3 + 0.797885x) \\ & + 0.5 + (0.0535161x^3 + 0.398942x) \\ & \times \cosh^{-2}(0.0356774x^3 + 0.797885x) \end{aligned} \quad (11)$$

ISRLU [10]: As regards ELU activation function, it is beneficial to solve *vanishing gradients*. Furthermore, as this function results in both negative and positive output for both sides of zero, it is useful for the network to shift mean to close zero, leading to a near distribution that speeds up training process and converging. However, the exponent function on the left side of zero highly cost computation. To further speed up learning, Inverse Square Root Linear Unit (ISRLU) [10] is proposed by,

$$f(x) = \begin{cases} x, & \text{if } x \geq 0. \\ x(\frac{1}{\sqrt{1+\alpha x^2}}), & \text{if } x < 0. \end{cases} \quad (12)$$

$$f'(x) = \begin{cases} 1, & \text{if } x \geq 0. \\ (\frac{1}{\sqrt{1+\alpha x^2}})^3, & \text{if } x < 0. \end{cases} \quad (13)$$

where constant α is normally ranged from 1 to 3. As can be seen that the exponent function is replaced by a polynomial that is useful to reduce cost of computation.

II. EVALUATION ACTIVATION FUNCTIONS WITH MLP AND MNIST DATASET

To evaluate these activation functions mentioned, we firstly propose a neural network, basing on Multilayer Perceptron (MLP) [11], and verify the network with MNIST dataset [12]. Although a lot of suggestions were provided to configure a MLP based network, there is no magic formula to select the optimum number of hidden layers or hidden neurons in each layer [13] in terms of the MLP architecture. Therefore, we present strategies to configure a MLP based network below.

A. Decision of the number of hidden layers

Given by MNIST dataset [12], while the input and output size are fixed to 784 (an MNIST image size is 28×28 is flatten into 784-dimensional vector) and 10 (the number of categories form 0 to 9 classified) respectively, the minimum number of hidden layers should be larger. This is explained by:

- Without hidden layers, the network only is capable of representing linearly separable functions or decisions.
- The network with only one hidden layer approximates any function that contains a continuous mapping from one finite space to another.
- Two-hidden-layer networks can represent an arbitrary decision boundary to arbitrary accuracy with rational activation functions and can approximate any smooth mapping to any accuracy.
- More-than-two-hidden-layer networks can learn complex representations (sort of automatic feature engineering).

Therefore, three hidden layers are decided to do experiments in this paper. Although three hidden layers may or may not enough to learn MNIST dataset, it is useful to speed up training process due to small network size. Selecting low volume of trainable parameters also helps network more general, thus evaluates the effect of mentioned activation functions better.

B. The number of nodes in each hidden layer

That is hard questions to indicate suitable number of nodes in each layer. This paper proposes some strategies below,

- The number of neural nodes in each hidden layer should be a power of two.
- The selected number of nodes are evaluated by conducting experiments.

C. Experimental setting

MNIST dataset is randomly separated into three subsets. The first subset is used for training. During training, the second subset, referred to as Evaluation subset (Eva.), is used to evaluate model, thus store the best model. Finally, the third subset, referred to as Test subset (Test), is used to evaluate the best model obtained from the training process.

In terms of network parameters, the initial trainable parameters are selected by Uniform distribution with scale set to 0.1 and learning rate is set to 0.001. These values are be fixed during evaluating the number of nodes in each hidden layer. Some other settings are decided below,

- Using ReLU as the activation function at hidden layers.
- Using Softmax function at the final layer to perform output probability.
- Apply Adam algorithm [14] for optimization and Cross-entropy is used for loss function as Eq. (14)

$$\text{Loss}_{\text{EN}}(\Theta) = -\frac{1}{N} \sum_{n=1}^N \mathbf{y}_n \log \{\hat{\mathbf{y}}_n(\Theta)\} + \frac{\lambda}{2} \|\Theta\|_2^2 \quad (14)$$

where Θ is all parameters, constant λ is set initially to 0.0001, both batch size N and the epoch number are set to 100, and $\mathbf{y}_n$ and $\hat{\mathbf{y}}_n$ denote expected and predicted results.

D. Experiment Results and Discussion

As shown in Table I, we evaluate 9 different MLP based network architectures and the network architecture of 784-2048-2048-512-10 shows the best performance of 97.5% and 96.4% over Evaluation (Eva.) and Test (Test) subsets, respectively. Thus, this network configuration is selected to evaluate activation functions. By using the MLP network

TABLE I: Different MLP configurations with the highest accuracy in **bold**

Input layer	Hidden layer 01	Hidden layer 02	Hidden layer 03	Output layer	Acc. % (Eva./Test)
784	2048	1024	512	10	96.6/96.0
784	1024	512	256	10	94.4/94.4
784	512	256	128	10	92.4/92.6
784	256	128	60	10	89.9/90.4
784	512	512	512	10	94.2/94.5
784	1024	1024	256	10	95.2/95.3
784	1024	1024	512	10	95.6/95.7
784	1024	1024	1024	10	9.65/9.61
784	2048	2048	512	10	97.5/96.4

TABLE II: Activation functions' performance with the highest accuracy and lowest time consumption in **bold**

Activation Function	Eva./Test Acc. (%)	Training Time/Epoch (second)
RELU	97.5/96.4	11.86
ELU	97.3/96.5	12.14
SELU	99.0/97.0	12.79
GELU	97.1/96.6	82.37
ISRLU	97.9/96.8	23.98

architecture of 784-2048-2048-512-10 selected and keeping setting parameters of learning rate, Epoch number and Batch size, ReLU function is replaced by ELU, SELU, GELU, and ISRLU functions to evaluate in the order. As results obtained shown in Table II, it can be seen that four activation functions analysed perform better than ReLU. This can be explained that these four activation functions result certain value when input is negative, prevent both the *vanishing gradients* and the *dead state*. Meanwhile ReLU may drop in the *dead state* when input is negative. In terms of training time, ELU and SELU consume little more time per training epoch compared with ReLU. GELU and ISRLU, by contrast, cost more time.

III. EVALUATE ACTIVATION FUNCTIONS WITH VGGISH NETWORK AND DCASE 2018 TASK 1A

To evaluate whether these activation functions also work well on different datasets and network architectures with respect to a certain application, we further conduct experiments over a Convolutional Neural Network (CNN) with DCASE 2018 Task 1A dataset [15].

A. DCASE 2018 Task 1A dataset

DCASE 2018 Task 1A [15] defines a Acoustic Scene Classification task, one of two major researchs of *Machine Hearing* [16]. In particular, the challenge is to classify 10-second audio recordings into 10 sound scene context as shown in Table III. As evaluation set has not published, development set is separated into two subsets, referred to as training set and test set, for training and testing processes respectively.

B. Feature Engineering Extraction

Firstly, the draw audio recordings are transformed into spectrograms where both temporal and frequency information are well presented. Next, the entire spectrograms are splitted into smaller patches of 128×128 that are suitable for back-end classification. In this paper, we apply log-mel for spectrogram transformation and use Librosa toolbox [17] to implement with setting parameters shown in Table IV.

TABLE III: DCASE2018 Task 1A dataset

Categories	Training set	Test set
Airport	599	265
Bus	622	242
Metro	603	261
Metro Station	605	259
Park	622	242
Public Station	648	216
Shopping Mall	585	279
Pedestrian Street	617	247
Traffic Street	618	246
Tram	603	261
Total files	6122	2518

TABLE IV: Spectrogram parameters

Parameters	Setting
FFT Number	2048
Hop Size	256
Window Size	1024
Log-mel Filter Number	128
Patch Size	128×128

C. Back-end CNN based Classifier

As regards back-end classification, a CNN based network architecture, likely VGG-9 [18], is proposed as shown in Table V. It can be seen that the CNN based model comprises of two main parts, CNN and DNN in the order. The CNN part is performed by convolutional layer (Cv [kernel size]), Rectify Liner Unit (ReLU), Batch normalization (BN), Average pooling (AP [kernel size], Drop out (Dr (dropped percentage))), and Global average pooling (GAP). Meanwhile, DNN part comprises of three fully connected layers (FC) followed by Rectify Linear Unit and Drop out. At the final fully connected layer, Softmax is used to perform probability of 10 categories.

D. Experimental Setting

Similar to experiments with MNIST dataset and MLP based network, the CNN based network is also built on Tensorflow framework, use Cross-entropy for loss function as mentioned in (14), and Adam for optimization.

TABLE V: CNN based network architecture

Network architecture	Output
CNN	
Input layer (image patch)	128×128×1
Cv [9×9] - ReLU - BN - AP [2×2] - Dr (10%)	64×64×32
Cv [7×7] - ReLU - BN - AP [2×2] - Dr (20%)	32×32×64
Cv [5×5] - ReLU - BN - Dr (30%)	32×32×128
Cv [5×5] - ReLU - BN - AP [2×2] - Dr (30%)	16×16×128
Cv [3×3] - ReLU - BN - Dr (40%)	16×16×256
Cv [3×3] - ReLU - BN - GAP - Dr (40%)	256
DNN	
Input layer (vector)	256
FC - ReLU - Dr (50%)	512
FC - ReLU - Dr (50%)	1024
FC - Softmax	10

E. Experimental Results and Discussion

TABLE VI: Performance over DCASE 2018 Task 1A

Network Option	Activation Func.	Acc. %
w/o BN, w/o Dr	ReLU	53.2
w/o BN, w/o Dr	ELU	54.7
w/o BN, w/o Dr	SELU	57.1
w/o BN, w/o Dr	GELU	58.5
w/o BN, w/o Dr	ISRLU	57.3
w/ BN, w/ Dr	ReLU	63.9
w/ BN, w/ Dr	ELU	61.7
w/ BN, w/ Dr	SELU	61.4
w/ BN, w/ Dr	GELU	63.1
w/ BN, w/ Dr	ISRLU	60.1
w/o BN, w/ Dr	SELU	59.1
w/ BN, w/o Dr	GELU	60.0

To evaluate activation functions with the CNN based network architecture and DCASE 2018 Task 1A dataset, we propose to conduct three main experiments. In the first experiment, we use the CNN based network proposed, but remove Batch normalization and Dropout layers. Next, ReLU is replaced by the other activation functions to evaluate. The second experiment keeps the CNN based network proposed unchanged, thus replace ReLU by the other activation functions. For the final experiment, we evaluate whether SELU and GELU can replace the roles of Batch normalization and Dropout layers as these activation functions' advantages mentioned in I-C. To evaluate SELU, we remove the Batch normalization layer in the CNN based network. Meanwhile, we remove the Dropout layer in the CNN based network to evaluate GELU.

As results shown in the top part of Table VI, without using Batch normalization and Dropout layers, activation functions such as ELU, SELU, GELU and ISRLU help to improve the classification accuracy compared with ReLU. This is similar to experiments' result over MNIST dataset and MLP based network mentioned in Section II-D. Again, using activation functions such as ELU, SELU, GELU and ISRLU to deal with *vanished gradients* issue is confirmed.

As results shown in the middle part of Table VI, when Batch normalization and Dropout layers are used, CNN based network with ReLU outperforms the other activation functions. It is interesting to indicate that Batch normalization and Dropout layers effectively help to improve performance, thus reduce the effect of *vanish gradient* from using ReLU.

Looking at the bottom part in Table VI, it can be seen that SELU and GELU cannot replace the roles of Batch normalization and Dropout layers within the experiment mentioned (CNN based network over DCASE 2018 Task 1A). Compare with DCASE 2018 Task 1A baseline [15], our best result of 63.9% improves by 4.2%. As this paper focuses on analysing activation functions rather than competing with systems submitted to DCASE 2018 Task 1A challenge, the proposed model only uses single spectrogram and compared with the DCASE 2018 Task 1A baseline (Almost submitted systems for this challenge or papers [19], [20], [21] exploring this dataset make use of multiple spectrograms).

IV. CONCLUSION

This paper has presented an analysis of using different activation functions in deep learning networks with respect to supervised learning. By conducting extensive experiments on MLP based network with MNIST dataset and CNN based network with DCASE 2018 Task 1A dataset, the roles of activation functions are confirmed, thus further evaluated over the specific application of Acoustic Scene Classification. The experiments indicate that although ReLU, ELU, SELU, GELU and ISRLU can help to prevent *vanished gradients* or/and *dead state* issue, these activation functions should combine with other statistic layers such as Batch normalization or Dropout to further improve the ASC system performance.

REFERENCES

- [1] Morteza Mohammadzadeh, Lei Chen, Ali Ghaffari, and John Willison, "A combination of linear and nonlinear activation functions in neural networks for modeling a de-superheater," *Simulation Modelling Practice and Theory*, vol. 17, no. 2, pp. 398–407, 2009.
- [2] Yang Liu, Alexandras Neophytou, Sunando Sengupta, and Eric Sommerlade, "Cross-modal spectrum transformation network for acoustic scene classification," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 830–834.
- [3] Jun Han and Claudio Moraga, "The influence of the sigmoid function parameters on the speed of backpropagation learning," in *International Workshop on Artificial Neural Networks*. Springer, 1995, pp. 195–201.
- [4] A. Mohamed and H. Raafat, "Recognition of heart sounds and murmurs for cardiac diagnosis," *Proc. Of The International Conference On Pattern Recognition*, pp. 1009–1011, 1988.
- [5] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter, "Fast and accurate deep network learning by exponential linear units (elus)," *arXiv preprint arXiv:1511.07289*, 2015.
- [6] Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter, "Self-normalizing neural networks," in *Advances in neural information processing systems*, 2017, pp. 971–980.
- [7] Ben Pinkowski, "A template-based approach for recognition of intermittent sounds," pp. 51–57, 1991.
- [8] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *The journal of machine learning research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [9] Dan Hendrycks and Kevin Gimpel, "Gaussian error linear units (gelus)," *arXiv preprint arXiv:1606.08415*, 2016.
- [10] Brad Carlile, Guy Delamarter, Paul Kinney, Akiko Marti, and Brian Whitney, "Improving deep learning by inverse square root linear units (isrlus)," *arXiv preprint arXiv:1710.09967*, 2017.
- [11] Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep Learning*. MIT Press, 2016, <http://www.deeplearningbook.org>.
- [12] Kim, Burred, and Sikora, "How efficient is mpeg-7 for general sound recognition?," Jun 2004.

- [13] Casey Michael, “Mpeg-7 sound-recognition tools,” 2001, vol. 11, pp. 737–747.
- [14] Diederik P Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [15] Annamaria Mesaros, Toni Heittola, and Tuomas Virtanen, “A multi-device dataset for urban acoustic scene classification,” in *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2018 Workshop (DCASE2018)*, November 2018, pp. 9–13.
- [16] Richard F Lyon, *Human and machine hearing*, Cambridge University Press, 2017.
- [17] Brian McFee, Colin Raffel, Dawen Liang, Daniel PW Ellis, Matt McVicar, Eric Battenberg, and Oriol Nieto, “librosa: Audio and music signal analysis in python,” in *Proceedings of the 14th python in science conference*, 2015, vol. 8, pp. 18–25.
- [18] Karen Simonyan and Andrew Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [19] Lam Pham, Huy Phan, Truc Nguyen, Ramaswamy Palaniappan, Alfred Mertins, and Ian McLoughlin, “Robust acoustic scene classification using a multi-spectrogram encoder-decoder framework,” *arXiv preprint arXiv:2002.04502*, 2020.
- [20] Lam Pham, Ian McLoughlin, Huy Phan, Ramaswamy Palaniappan, and Yue Lang, “Bag-of-features models based on c-dnn network for acoustic scene classification,” in *Audio Engineering Society Conference: 2019 AES International Conference on Audio Forensics*. Audio Engineering Society, 2019.
- [21] Dat Ngo, Hao Hoang, Anh Nguyen, Tien Ly, and Lam Pham, “Sound context classification basing on join learning model and multi-spectrogram features,” *arXiv preprint arXiv:2005.12779*, 2020.