

A LOTTERY TICKET HYPOTHESIS FRAMEWORK FOR LOW-COMPLEXITY DEVICE-ROBUST NEURAL ACOUSTIC SCENE CLASSIFICATION

Technical Report

Chao-Han Huck Yang^{†1}, Hu Hu^{†1}, Sabato Marco Siniscalchi^{1,2}, Qing Wang³, Yuyang Wang³,
Xianjun Xia⁴, Yuanjun Zhao⁴, Yuzhong Wu⁴, Yannan Wang⁴, Jun Du³, Chin-Hui Lee¹

¹School of Electrical and Computer Engineering, Georgia Institute of Technology, GA, USA

²Computer Engineering School, University of Enna Kore, Italy

³University of Science and Technology of China, HeFei, China

⁴Tencent Media Lab, Tencent Corporation, China

ABSTRACT

We propose a novel neural model compression strategy combining data augmentation, knowledge transfer, pruning, and quantization for device-robust acoustic scene classification (ASC). Specifically, we tackle the ASC task in a low-resource environment leveraging a recently proposed advanced neural network pruning mechanism, namely Lottery Ticket Hypothesis (LTH), to find a sub-network neural model associated with a small amount non-zero model parameters. The effectiveness of LTH for low-complexity acoustic modeling is assessed by investigating various data augmentation and compression schemes, and we report an efficient joint framework for low-complexity multi-device ASC, called *Acoustic Lottery*. Acoustic Lottery could compress an ASC model up to $1/10^4$ and attain a superior performance (validation accuracy of 79.4% and Log loss of 0.64) compared to its not compressed seed model. All results reported in this work are based on a joint effort of four groups, namely GT-USTC-UKE-Tencent, aiming to address the “Low-Complexity Acoustic Scene Classification (ASC) with Multiple Devices” in the DCASE 2021 Challenge Task 1a.

Index Terms— Lottery ticket hypothesis, Teacher-student learning, Acoustic scene classification, and Device-robustness

1. INTRODUCTION

Acoustic scene classification (ASC) aims to recognize a set of given environment classes (e.g., airport and urban park) from real-worlds sound examples. Analysis and learning to predict acoustic scene sounds are important topics associated with various mobile and on-device intelligent applications [1]. The Detection and Classification of Acoustic Scenes and Events (DCASE) challenges [2, 3, 4, 5] provide a comprehensive evaluation platform and benchmark data to encourage and boost sound scene research communities. DCASE 2021 Task 1a [6] focuses on developing low-complexity acoustic modeling (AM) solutions for predicting sounds recorded from multiple devices (e.g., electret binaural microphones, smartphones, and action cameras). The goal is to design a device-robust ASC system preserving generalization power over audios recorded by different devices, and highlighting the importance of low-complexity requirements.

[†]The authors have equal contribution on this work.

From previous DCASE challenges, we observed that several competitive ASC systems [7, 8, 9] benefited from large-scale convolutional neural models combined with several data augmentation schemes, but whether we can attain the generalization power of those complex models with a low-complexity architecture is the research goal to be addressed in DCASE 2021 challenge. To this end, we focus on addressing two basic questions: (i) Are some well-performed device-robust ASC models overparameterized? (ii) Can we take advantage of some overparameterized models to design a low-complexity ASC framework on multi-device data?

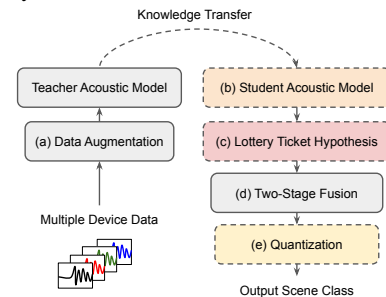


Figure 1: The proposed **Acoustic Lottery** (AL) framework.

In the quest for addressing the above questions, we deployed a novel framework, namely “Acoustic Lottery,” for DCASE 2021 Task 1a, which will be described in the following sections. As shown in Figure 1, our Acoustic Lottery system consists of (a) a data augmentation process to improve model generalization, (b) a teach-student learning mechanism to transfer knowledge, from a large teacher model to a small student model, (c) a Lottery Ticket Hypothesis [10] based pruning method to deliver a low-complexity model, (d) a two-stage fusion technique [11] to improve model prediction, and finally (e) a quantization block to deploy a final model owing less than 128 KB non-zero parameters, which is the requirement of Task 1a. A detailed presentation of each block in Figure 1 is discussed in the following sections.

2. LOW-COMPLEXITY ACOUSTIC MODELING FRAMEWORK

2.1. Data Augmentation Strategy

In some previous works [7, 11, 12, 13], data augmentation strategies played a key role to attain competitive log-loss results on the

Algorithm 1 LTH for Device-Robust Acoustic Modeling

1. **Input:** a model, G_0 ; augmented sound data, D .
 2. **Randomly Initialize Weights** (Θ_0).
 3. **Initialize Model:** $G_0(\Theta_0) \rightarrow G_1$
 4. **For** $t = 1, \dots, T$: # Pruning Searching Iterations
 5. **For** $e = 1, \dots, E$: # Gradient Training Epochs
 6. $\Theta_e \rightarrow \Theta$: TSL-train G_t with D for its final weights (Θ_t)
 7. **If** $t < T$: # LTH Pruning Strategy
 8. Mask(Θ_t) to get a pruned graph G_p from G_0
 9. Load homologous initial weights $\Theta_p \in \Theta_0$ from G_p
 10. Update target model $G_p(\Theta_p) \rightarrow G_{t+1}$
 11. **Output:** A well-trained pruned model $G_T(\Theta_T)$
-

DCASE 2020 validation set [9]. With the goal of deploying a seed model with good generalization capabilities to deal with the multiple device acoustic condition, the first module (Figure. 1 (a)) of our DCASE 2021 system thereby builds upon data augmentation methods investigated in [7]. To be specific, we use i) Mixup [14], ii) Spectrum augmentation [15], iii) Spectrum correction [16], iv) Pitch shift, v) Speed change, vi) Random noise, and vii) Mix audios. The reverberation data described in [7] is not used in our experiments.

2.2. Teacher-Student Learning (TSL)

Teacher-Student Learning (TSL), also named as Knowledge Distillation (KD), is a widely investigated approach for model compression [17, 18]. Specifically, it transfers knowledge from a large and complex deep model (teacher model) to a smaller one (student model). The main idea is to establish a framework that makes the student directly mimicking the final prediction of teacher. Formally, the soften outputs of a network can be computed by $p = \text{softmax}(\frac{\alpha}{\tau})$, where α is the vector of logits (pre-softmax activations) and τ is a temperature parameter to control the smoothness [17]. Accordingly, the distillation loss for soft logits can be written as the Kullback-Leibler divergence between the teacher and student soften outputs. In this work, we followed the approaches in [7] to build a large two-stage ASC system, serving as the teacher model. Then a teacher-student learning method is used to distill knowledge to a low-complexity student model, as shown in Figure 1 (b).

2.3. Lottery Ticket Hypothesis Pruning

Next we have investigated advanced pruning techniques to further reduce non-zero model parameters of the student. Although neural network pruning methods often negatively affect both model prediction performance and generalization power, a recent study, referred to as Lottery Ticket Hypothesis [10] (LTH), showed a quite surprising phenomenon, namely pruned neural networks (sub-networks) could be trained attaining a performance that was equal to or better than the not pruned original model if the not pruned parameters were set to the same initial random weights used for the non-pruned model. Interestingly, LTH-based low-complexity neural models had proven competitive prediction performance on several image classification tasks [10] and recently have been supported with some theoretical findings [19] related to overparameterization.

Algorithm Design: In **Algorithm 1**, we detail our approach under the Acoustic Lottery framework: In step (1), we first choose a model with its original neural architecture (e.g., Inception in our case) G_0 and record its initial weights parameters Θ_0 in step (2). In

our work, we incorporate teacher-student learning framework discussed in Section 2.2 with the goal of mimic prediction accuracy and generalization adapted of the teacher acoustic model - a complex model trained separately. At the end of each training phase, a pruning iteration is started if the current iteration t is less than T . The LTH pruning searches for a low-complexity model in steps (7) through (10).

From our empirical findings in DCASE 2021 Task 1A data, we found that the proposed Acoustic Lottery only needs one or two ($T=1$ or 2 in **Algorithm 1**) searching iteration(s) to find a good low-complexity acoustic model without a significant drop in the ASC classification accuracy compared to the high-complexity teacher model on the validation set. To select the mask function in step (8), we evaluate three major LTH strategies, namely: (i) large-final; (ii) small weights, and (iii) global small weights, which were proposed in [10]. We found the small weights strategy allows us to attain better trade-off between classification accuracy and compression rate compared to the other two mentioned methods as shown in Figure 2. Therefore, we selected “small weights” as pruning strategy to be used in our final submission. Finally, a well-trained pruned student acoustic model is deployed in step (10) of **Algorithm 1**




Figure 2: We compared empirical performance of different LTH-masking strategies [20] versus sparsity level (weights remaining).

Visualization: To better interpret weights distribution in an LTH-pruned neural acoustic model, we visualize a shallow inception model (excluding convolutional layers due to their dimensional conflicts) on Index 3 in Table 1 and its LTH-pruned results as Index 5 in in Table 1 shown in Figure 3. In Figure 3b, we can observe that the proposed Acoustic Lottery framework can discover a well-trained model using only sparse weights with up to a 149× compression rate.

2.4. Two-Stage Fusion and Multi-Task Learning

To boost ASC performance, we follow the investigation in the two-stage ASC scheme discussed in [11], where the relationship between the 3-class and 10-class ASC systems were exploited to boost the 10-class ASC system. This step is carried out in the module (d) in Figure 1. The key idea is that the labels of the two subtasks, 3-class and 10-class problems, differ in the degree of abstraction and using the two labels together could be helpful. In our setup, the 3-class classifier classifies an input scene audio into one of three broad classes: in-door, out-door, and transportation. This 3-class classification way is from our prior knowledge that scene audios can be roughly categorized into such three classes. The 10-class classifier is actually the main classifier. Each audio clip should belong to one of the three / ten classes. The final acoustic scene class is chosen by the score fusion of those two classifiers. If we let \mathbb{C}^1 and \mathbb{C}^2 denote the set of three broad classes, and ten classes, respectively, and let F^1 and F^2 indicate the output of the first and second classifier,




Figure 3: Visualized of layer-wise weights distribution by the LTH approach applied to the student neural acoustic model: (a) Shallow Inception (SIC) student and (b) LTH-Pruned (99.33%) SIC student.

respectively. The final predicted class $Class(x)$ for the input x is:

$$Class(x) = \underset{q, (p \in \mathcal{C}^1, q \in \mathcal{C}^2, p \supset q)}{\operatorname{argmax}} F_p^1(x) * F_q^2(x),$$

where $p \supset q$ means that p can be thought of a super set of q . For example, transportation class is the super set for bus, tram, and metro classes. Therefore, the probability of an input audio clip to be from the public square scene is equal to the product of the probability of out-door place, $F_p^1(x)$, and that of public square, $F_q^2(x)$.

However, the two ASC classifiers are trained separately, which means the total parameters will be doubled. In [8], the authors argued that joint training of two subtasks could be even more efficient. Specifically, the 3-class classifier and the 10-class classifier can be learned in a multi-task learning (MTL) [21] manner. The two classifiers can share some parameters, where only output layers are different. MTL is expected to perform as well as two-stage method but save parameters. We thus study that setting as an ablation module in our experimental section.

2.5. Quantization for Model Compression

As the main goal is to deploy a system with a size within 128 Kilo-bytes (KB), we further use a post-training quantization method with dynamic range quantization (DRQ), as shown in Figure 1 (e). DRQ is the simplest form of post-training quantization, which statically quantizes only weights from floating point to integer, which has 8-bits of precision. Moreover, activations are dynamically quantize based on their range to 8-bits. Leveraging upon DRQ, we thus convert our neural acoustic model from a 32-bit format to a 8-bit format, which compresses the model size to about 1/4 of its original size as our final model.

3. EXPERIMENTAL SETUP & RESULTS

3.1. Feature Extraction

We follow the same settings from DCASE 2020 Task-1a extracting acoustic features for DCASE 2021 Task-1a [6] before using the features to train low-complexity described in Section 2 and Figure 1. Log-mel filter bank (LMFB) features were used in our experiments as audio features. The input audio waveform is analyzed with a 2048 SFFT points, a window size of 2048 samples, and a

frame shift of 1024 samples. Thus the final input tensor size is thus $423 \times 128 \times 3$ for Task 1a. Before feeding the speech feature tensors into CNN classifier, we scaled each feature value into $[0,1]$.

3.2. Model Training

All ASC systems are evaluated on the DCASE 2020 task 1a development data set [5], which consists of $\sim 14K$ 10-second single-channel train audio clips and $\sim 3K$ test audio clips recorded by 9 different devices, including real devices A, B, C, and simulated device s1-s6. Only device A, B, C, s1-s3 are in the training set; whereas, devices s4-s6 are unseen in the training phase. The greatest amount of training audio clips are recorded with device A, namely over 10K audio clips. In the test set, the number of waveforms from each device is the same.

We use two different Inception [23] models as our target models, namely Shallow Inception (SIC) and Large Inception (LIC). SIC has two inception blocks whereas LIC has three inception blocks and more filters. The size computed by the way recommended in [22] of the original SIC and LIC are 503KB and 3528KB, respectively. All Inception models in this work are built with Keras based on Tensorflow2. Stochastic gradient descent (SGD) with a cosine-decay-restart learning rate scheduler is used to train all deep models. Maximum and minimum learning rates are 0.1, and $1e-5$, respectively. In our final submission, all development data is used. And due to there is no validation data, we use the output of model when learning rate hits the minimum number.

3.3. Results on Task 1a

In Table 1, we report only some of the evaluation results for low-complexity models collected on Task 1a due to space constraints. Two inception models: (i) shallow inception model (SIC) and (ii) large inception model (LTC), are investigated under the proposed Acoustic Lottery framework. By evaluating several low-complexity strategies shown in 1. From the results, Index (0) is the official baseline, which has the size of 90.3KB but very low accuracy and high log loss. Index (1) and Index (2) are results from [7], where a two-stage system is used. Although they achieve very good performance (77.6% for two-stage FCNN and 81.9% for two-stage ensemble), their size is very large, which are 132MB and 332MB, respectively. It should be noted that the reverberation augmented data is used for Index (1) and (2) but not for others.

The Index (3-11) in Table 1 are results of SIC. We here perform the ablation study for each method we propose. Index (3) is the SIC baseline, which has the size of 503KB, accuracy of 67.8%, and log loss of 0.954. With the use of TSL, shown as Index (4), we can improve the accuracy and log loss while keeping the model size unchanged. We use the two-stage FCNN model, shown as Index (1), as the teacher model. Index (5) shows the result of using LTH, where we can significantly reduce the model size (around $149\times$ compression rate. Although model parameters are reduced in a huge scale, the model performance shows much better than the SIC baseline: Index (3). This verifies our argument that the models are overparameterized a lot.

Index (6) and (7) shows the results by only using two-stage fusion or MTL. From the results, we can see the two-stage can boost the performance, but the method will double the model size. By using a compromise method, MTL, can work in the same manner but save parameters. However, it's slightly worse than using two-stage. Index (8) shows the result by only using quantization. The model parameters are quantized from float32 to float8. Although

Table 1: Experimental results on 2021 Task 1a. 'TSL' means performing teacher-student learning. 'LTH' means pruning with the Lottery Ticket Hypothesis algorithm. 'Two-stage' means using a two-stage fusion system. 'MTL' means using multi-task learning system. 'Quant' means using quantization on model parameters (float32 to float8). 'Aug' means using extra augmented data (Methods in Section 2.1). 'System size' is according to non-zero parameters [6]. All 'Y' in the table means we used that method. Acc. indicates validation accuracy.

Idx.	System	TSL	LTH	Two-stage	MTL	Quant	Aug	System size	Acc. %	Log Loss
(0)	Official Baseline [22]	-	-	-	-	-	-	90.3KB	47.7	1.473
(1)	Two-stage FCNN [11]	-	-	Y	-	-	Y	132MB	77.6	0.795
(2)	Two-stage Ensemble [11]	-	-	Y	-	-	Y	332MB	81.9	0.829
(3)	SIC	-	-	-	-	-	-	503KB	67.8	0.954
(4)	SIC	Y	-	-	-	-	-	503KB	68.9	0.919
(5)	SIC	-	Y	-	-	-	-	3.4KB	68.2	0.914
(6)	SIC	-	-	Y	-	-	-	1006 KB	68.9	0.914
(7)	SIC	-	-	-	Y	-	-	504KB	68.0	0.915
(8)	SIC	-	-	-	-	Y	-	126KB	66.9	0.972
(9)	SIC	Y	-	Y	-	-	-	1006KB	69.2	0.874
(10)	SIC	Y	-	Y	-	Y	-	252KB	68.4	0.906
(11)	SIC	-	Y	-	-	Y	-	0.9KB	67.7	0.931
(12)	LIC	-	-	-	-	-	-	3528KB	69.0	0.891
(13)	LIC	Y	-	-	-	-	-	3528KB	69.9	0.880
(14)	LIC	-	Y	-	-	-	-	23.6KB	69.2	0.878
(15)	LIC	-	-	Y	-	-	-	7056 KB	70.0	0.848
(16)	LIC	Y	Y	Y	-	-	-	47.2KB	70.8	0.833
(17)	LIC	Y	Y	-	-	-	Y	23.6KB	78.3	0.721
(18)	LIC	Y	Y	Y	-	-	Y	47.2KB	78.4	0.654
(19)	LIC	Y	Y	-	-	Y	Y	5.9KB	78.2	0.723
(20)	Ensemble of LICs	Y	Y	Y	-	Y	Y	47.2KB	78.7	0.644
(21)	Ensemble of LICs	Y	Y	Y	-	-	Y	118KB	79.0	0.637
(22)	Ensemble of LICs and SICs	Y	Y	Y	-	-	Y	122KB	79.4	0.640

it obtains a $4\times$ compression rate, the performance worsens when compared with the SIC baseline. However, according to our experiments, we find that the ensemble of 4 quantized models shows better results than an unquantized model, which shows the potential of quantization. With the combination of proposed approaches, we can further boost the performance of SIC model, as shown in Index (9) to (11) of Table 1. We can at most compress the SIC model to 0.9KB, shown as Index (11), with even better performance than SIC baseline. As for LIC models, shown in Index (12) to (19), the same conclusions as SIC can be observed. Furthermore, when training by augmented data, system robustness can be further boosted. As for LIC, we can at most compress it to 5.9KB, which a log loss of 0.723. And the best log loss can be obtained by an 47.2KB system, shown as Index (18), with an accuracy of 78.4% and log loss of 0.654. The model size limitation of DCASE 2021 task 1a is 128KB. Thus we investigate ensemble systems. As shown in Index (20) to (22), the model ensemble can further increase the performance. Index (20) is the ensemble of four quantized 10-class LICs and one unquantized 3-class LIC. Index (21) is the ensemble of three 10-class LICs and two 3-class LICs. Index (22) is further ensemble with a SIC on system of Index (21). It can obtain 79.4% accuracy and 0.640 log loss, with a model size of 122KB.

For our final submitted four systems: four “two-stage ensembles” of different LIC and SIC models with LTH pruning are selected. We obtain SICs and LICs from different training epochs by training with different combinations of data augmentation strategies and training criterion (one-hot labels or TS learning). Specifi-

cally, for system (a), we use two 3-class LICs, three 10-class LICs and one 10-class SIC. So the total non-zero parameter size of System (a) is 122KB ($23.6\text{KB} \times 5 + 3.4\text{KB} \times 1$). System (b) uses eight 3-class quantized-LICs, two 3-class quantized-SICs, ten 10-class quantized-LICs, and three 10-class quantized-SICs. So the total size of System (b) is 110KB ($5.9\text{KB} \times 18 + 0.9\text{KB} \times 5$). System (c) uses two 3-class LICs, two 3-class SIC, two 10-class LICs, four 10-class quantized-LICs, and one 10-class quantized-SIC. So the total size of System (c) is 125KB ($23.6\text{KB} \times 4 + 5.9\text{KB} \times 4 + 3.4\text{KB} \times 2 + 0.9\text{KB} \times 1$). System (d) uses two 3-class LICs, four 3-class quantized-LICs, one 10-class LIC, four 10-class quantized-LICs, and one 10-class SIC. In System (d) we give non-quantized models 4 times larger score weights than quantized models when doing ensemble. And the total size of System (d) is 122KB ($23.6\text{KB} \times 3 + 5.9\text{KB} \times 8 + 3.4\text{KB} \times 1$). The results of system (1) on development set is specified in Index (22) of Table 1.

4. DISCUSSION & CONCLUSION

As low-complexity acoustic modeling, a lottery ticket hypothesis framework, Acoustic Lottery, is proposed and provides competitive results. As the very first attempt on applying LTH for acoustic learning and modeling, our future works included theoretical analysis on the success of LTH and its relationship between knowledge distillation for different acoustic and robust speech processing tasks [24]. We will open source our proposed framework and findings to the community at https://github.com/MihawkHu/Acoustic_Lottery.

5. REFERENCES

- [1] D. Yang, H. Wang, and Y. Zou, “Unsupervised multi-target domain adaptation for acoustic scene classification,” *arXiv preprint arXiv:2105.10340*, 2021.
- [2] A. Mesaros, T. Heittola, E. Benetos, P. Foster, M. Lagrange, T. Virtanen, and M. D. Plumbley, “Detection and classification of acoustic scenes and events: Outcome of the DCASE 2016 challenge,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 2, pp. 379–393, 2018.
- [3] A. Mesaros, T. Heittola, and T. 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.
- [4] 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 *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2017 Workshop (DCASE2017)*, November 2017, pp. 85–92.
- [5] T. Heittola, A. Mesaros, and T. Virtanen, “Acoustic scene classification in dcase 2020 challenge: generalization across devices and low complexity solutions,” in *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2020 Workshop (DCASE2020)*, 2020, submitted. [Online]. Available: <https://arxiv.org/abs/2005.14623>
- [6] I. Martín-Morató, T. Heittola, A. Mesaros, and T. Virtanen, “Low-complexity acoustic scene classification for multi-device audio: analysis of dcase 2021 challenge systems,” *arXiv preprint arXiv:2105.13734*, 2021.
- [7] H. Hu, C.-H. H. Yang, X. Xia, X. Bai, X. Tang, Y. Wang, S. Niu, L. Chai, J. Li, H. Zhu, F. Bao, Y. Zhao, S. M. Siniscalchi, Y. Wang, J. Du, and C.-H. Lee, “Device-robust acoustic scene classification based on two-stage categorization and data augmentation,” 2020.
- [8] H.-j. Shim, J.-h. Kim, J.-w. Jung, and H.-J. Yu, “Attentive max feature map for acoustic scene classification with joint learning considering the abstraction of classes,” *arXiv preprint arXiv:2104.07213*, 2021.
- [9] T. Heittola, A. Mesaros, and T. Virtanen, “Acoustic scene classification in dcase 2020 challenge: generalization across devices and low complexity solutions,” *arXiv preprint arXiv:2005.14623*, 2020.
- [10] J. Frankle and M. Carbin, “The lottery ticket hypothesis: Finding sparse, trainable neural networks,” in *International Conference on Learning Representations*, 2018.
- [11] H. Hu, C.-H. H. Yang, X. Xia, X. Bai, X. Tang, Y. Wang, S. Niu, L. Chai, J. Li, H. Zhu, *et al.*, “A two-stage approach to device-robust acoustic scene classification,” *arXiv preprint arXiv:2011.01447*, 2020.
- [12] H. Chen, Z. Liu, Z. Liu, P. Zhang, and Y. Yan, “Integrating the data augmentation scheme with various classifiers for acoustic scene modeling,” DCASE2019 Challenge, Tech. Rep., June 2019.
- [13] K. Koutini, H. Eghbal-zadeh, and G. Widmer, “Acoustic scene classification and audio tagging with receptive-field-regularized CNNs,” DCASE2019 Challenge, Tech. Rep., June 2019.
- [14] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, “mixup: Beyond empirical risk minimization,” *arXiv preprint arXiv:1710.09412*, 2017.
- [15] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” *arXiv preprint arXiv:1904.08779*, 2019.
- [16] T. Nguyen, F. Pernkopf, and M. Kosmider, “Acoustic scene classification for mismatched recording devices using heated-up softmax and spectrum correction,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 126–130.
- [17] G. Hinton, O. Vinyals, and J. Dean, “Distilling the knowledge in a neural network,” *arXiv preprint arXiv:1503.02531*, 2015.
- [18] J. Li, R. Zhao, J.-T. Huang, and Y. Gong, “Learning small-size dnn with output-distribution-based criteria,” in *Fifteenth annual conference of the international speech communication association*, 2014.
- [19] E. Malach, G. Yehudai, S. Shalev-Schwartz, and O. Shamir, “Proving the lottery ticket hypothesis: Pruning is all you need,” in *International Conference on Machine Learning*. PMLR, 2020, pp. 6682–6691.
- [20] H. Zhou, J. Lan, R. Liu, and J. Yosinski, “Deconstructing lottery tickets: Zeros, signs, and the supermask,” *arXiv preprint arXiv:1905.01067*, 2019.
- [21] C.-H. H. Yang, L. Liu, A. Gandhe, Y. Gu, A. Raju, D. Filimonov, and I. Bulyko, “Multi-task language modeling for improving speech recognition of rare words,” *arXiv preprint arXiv:2011.11715*, 2020.
- [22] I. Martín-Morató, T. Heittola, A. Mesaros, and T. Virtanen, “Low-complexity acoustic scene classification for multi-device audio: analysis of dcase 2021 challenge systems,” 2021.
- [23] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, “Going deeper with convolutions,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1–9.
- [24] C.-H. Yang, J. Qi, P.-Y. Chen, X. Ma, and C.-H. Lee, “Characterizing speech adversarial examples using self-attention u-net enhancement,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 3107–3111.