



Multi-Dataset Co-Training with Sharpness-Aware Optimization for Audio Anti-spoofing

Hye-jin Shim¹, Jee-weon Jung^{2,†}, Tomi Kinnunen¹

¹University of Eastern Finland, Finland

²Carnegie Mellon University, USA

hyejin.shim@uef.fi, jeeveonj@andrew.cmu.edu, tomi.kinnunen@uef.fi

Abstract

Audio anti-spoofing for automatic speaker verification aims to safeguard users' identities from spoofing attacks. Although state-of-the-art spoofing countermeasure(CM) models perform well on specific datasets, they lack generalization when evaluated with different datasets. To address this limitation, previous studies have explored large pre-trained models, which require significant resources and time. We aim to develop a compact but well-generalizing CM model that can compete with large pre-trained models. Our approach involves multi-dataset co-training and sharpness-aware minimization, which has not been investigated in this domain. Extensive experiments reveal that proposed method yield competitive results across various datasets while utilizing 4,000 times less parameters than the large pre-trained models.

Index Terms: audio spoofing, spoofing detection, sharpness aware minimization, generalization, multi-dataset training

1. Introduction

Automatic speaker verification (ASV) systems [1], even state-of-the-art, have been reported to be easily deceived by *spoofing attacks* including speech synthesis (text-to-speech, TTS), voice conversion (VC). For the reliable ASV systems, *audio anti-spoofing* has emerged which aims to distinguish the utterance from a real human (*bona fide*) or spoofing attacks (*spoofed*). To develop spoofing countermeasure (CM) models, various studies have conducted focusing on feature [2–4], model architecture [5–8], and other techniques (e.g. loss function and data augmentation) [9–11].

While most state-of-the-art CMs perform well on specific evaluation datasets, they do not generalize well when cross-evaluated on different datasets [12–18]. Several studies have explored improving generalization capability and literature can be divided into two strands: The first strand exploits techniques to develop a model with feature adjustment, gradient-based methods, and adversarial learning [12, 13, 15]. The other strand develops model using domain adaptation, continual learning, and self-supervised learning to leverage large-scale datasets [16–21]. Even though the latter demonstrated promising results with a large gap, it typically requires additional learning steps with a large, heavyweight model such as wav2Vec2.0 [22] or HuBERT [23] that typically contains billions of parameters.

Dating back to the base assumption of large-scale pre-trained models, there is a core premise of “*more data leads to better (generalization) performance*” that lies on deep learning as well as statistics. It expects that exposure of large amounts

Table 1: Performance comparison of using a single dataset and multiple datasets in audio anti-spoofing. Evaluation is conducted with the ASVspoof 2019 LA dataset.

Train Dataset(s)	EER(%)
ASVspoof 2015	38.83
ASVspoof 2019	1.38
ASVspoof 2019 + ASVspoof 2015	1.56
ASVspoof 2019 + ASVspoof 2015 + WaveFake	1.76

of data to the model can lead to better generalization. Following this idea, abundant studies have demonstrated the effectiveness of training a model on diverse datasets to enhance their robustness and adaptability, namely unsupervised/semi-supervised pre-training and self-supervised learning. Otherwise, training a model using multiple datasets at once has been well-known as a challenging and unsolved problem because of different characteristics of datasets that may interfere with the target task. Nevertheless, a few studies have been conducted in this direction [24–27]. There is a potential to be improved, so further inspection and exploration yet remain.

In this study, we aim to develop a compact and well-generalized CM model leveraging multi-dataset co-training. At an early stage of the present study, we conducted pilot experiments with gradually enlarging datasets for training the model at once. The result is shown in Table 1 and its outcome indicates that merely combining datasets that span different domains does not guarantee generalization. These observations motivated us to address more elaborate ways to optimize the model when using multiple training datasets. To mitigate this problem, we explore the way to reduce the perturbation which can be caused by domain mismatch across different datasets. Recent studies have shown gradient-based methods, especially sharpness-aware related works [28–30], demonstrate to avoid a severe perturbation during the training process and enhance generalization capability. Here, the term *sharpness* in this context refers to the curvature of neighborhoods in the loss surface.

To this end, we exploit two recently proposed optimization techniques: *sharpness-aware minimization* (SAM) [28] and *Adaptive sharpness-aware minimization* (ASAM) [29]. SAM and ASAM are both designed to find flat minimas by taking into account the sharpness in the loss surface in addition to the gradients. Hence, we hypothesize that combining sharpness-aware training — an approach designed to avoid sharp loss minima — with multiple dataset co-training (to handle diverse data) has the potential to lead to improving the generalization of CM.

Our study first attempts to optimize multi-dataset co-training and also practical effectiveness of sharpness-aware training remains presently unknown in the CM task. Provid-

[†] Work done while author was in Naver Corporation, South Korea.

ing initial answers to this question forms the main novelty of our work; we implement our proposed method using the state-of-the-art graph neural network-based AASIST [8] model with various evaluation data. Our comprehensive experimental validation reveals that both approaches are effective in that the proposed model shows competitive results throughout various datasets using a number of parameters more than 4000 times less than the large pre-trained models and leads to better generalization.

2. Multi-dataset training

2.1. Related works

Several previous works have focused on enlarging the amount of training data for the training based on “*more data leads to better (generalization) performance*” in line with fitting the general distribution with large amounts of data. Both unsupervised learning and self-supervised learning align with this principle aimed to enhance model generalization with more data. However, above mentioned studies are taking into account the conditions that are hard to get labeled data for the same task. It means if labeled data is available, it is basically helpful for training, however, multi-dataset co-training is even unveiled yet compared to the methodologies for exploiting unlabeled data. Few studies [27, 31] have conducted multi-dataset co-training in other domains.

There exists a few preliminary works in audio spoofing utilizing multiple datasets [13, 32]. They concentrated on developing a single model that can detect diverse types of attacks as audio spoofing attacks can be divided into two categories: logical access (LA) and physical access (PA). The former includes TTS and VC attacks, whereas the latter refers to replay attacks only. However, no research has explored multi-dataset training to deal with one category of attack (either LA or PA). The potential effectiveness of training the model with multiple datasets simultaneously for the same task has yet to be explored in depth, but it would be worthwhile to investigate further.

2.2. Summary of datasets used in this study

We use three datasets concurrently to train a *single model* in a *single phase*: ASVspoof 2015 [33], ASVspoof 2019 LA [34], and WaveFake [35]. The latest ASVspoof edition in 2021 additionally introduced DeepFake (DF) scenario which includes lossy codecs used for media storage. In this study, we *only* deal with LA spoofing attacks for training a model. To address generalization to an unseen domain, we use the ASVspoof 2021 LA and DF tasks for the evaluation. Note that ASVspoof 2015, ASVspoof 2019, and a part of ASVspoof 2021 are based upon the Voice Cloning Toolkit (VCTK) corpus [36]; however, they cover different attacks. Supporting this basis, it is assumed that ASVspoof 2015 evaluation can be theoretically easy when a CM is trained on the more diverse LA 2019 train set. However, it has empirically confirmed that this is not the case [13, 18]. An overview of the selected datasets is shown in Table 2.

ASVspoof 2015 [33] is the earliest and smallest database among the four existing ASVspoof editions. The evaluation set consists of five known and five unknown attacks composed of different TTS and VC systems. In this context, the term *known* attack indicates an attack in the train and test set is overlapped, while *unknown* attack indicates to scenarios where the test set includes attacks that were not encountered during the training phase.

Table 2: Summary of data statistics used in this study. #Spks, # Utts, and # Conds refer to the number of speakers, utterances, and spoofing conditions, respectively. Division of train and test set is indicated by /.

Dataset	# Spks	# Utts	# Conds
ASVspoof 2015	25 / 46	16375 / 193404	5 / 10
ASVspoof 2019 LA	20 / 48	25380 / 108978	6 / 13
WaveFake	2	117985	6
ASVspoof 2021 LA	- / 48	181566	- / 13
ASVspoof 2021 DF	- / 48	611829	- / 13

ASVspoof 2019 [34] is a large-scale dataset that covers advanced technologies developed during the four years following ASVspoof 2015. It includes 6 and 13 types of spoofing attacks in train and test, respectively. There are two known attacks, four partially known attacks, and seven unknown attacks between train and test sets. Here, *partially known* attack denotes a scenario where some of the attacks are present in both the train set and test set, but some attacks are not present in the train set.

WaveFake [35] is collected using six different state-of-the-art TTS methods. It considers the created samples to resemble the training distributions. All spoofed data has been generated using the last, competitive VC and TTS models. Note that we utilize whole spoofed speech for WaveFake since no standardized test protocol exists. WaveFake contains spoofed utterances from two speakers, originating from LJSPEECH [37] and JSUT [38] datasets, respectively.

ASVspoof 2021 [39] is the latest and hardest edition of the ASVspoof challenge series. It only contains a test set and introduces real telephony systems both applied to encoding and transmission artifacts in the LA scenario of the ASVspoof 2019 dataset. The DF scenario additionally consists of bona fide and spoofed data processed through various audio compressors containing data from two additional datasets, VCC 2018 [40] and VCC 2020 [41].

3. Sharpness-aware optimizations

When working with multiple datasets simultaneously, a model may easily struggle with domain information between the main task, despite having access to explicit class labels. Domain information can distract the model from converging, although it may generalize better once the train loss has well converged. We thus seek methods that can prevent the model from being distracted by domain discrepancies between different datasets. Besides, this direction can also cast away the necessity of the additional pre-training and fine-tuning steps in a straight way.

In particular, *sharpness-aware minimization* (SAM) [28] recently demonstrates state-of-the-art performance in various tasks [42–45]. SAM aims to find a *flat* region in the parameter space with both low loss itself and neighborhoods, seeking flat minima. It uses worst-case perturbation of the model parameters on every iteration in the training phase and can be easily implemented on top of existing optimizers such as Adam [46]. Moreover, there are several follow-up studies related to sharpness. For instance, the authors in [29] proposed a scale-invariant version, *adaptive* variant of SAM, adaptive SAM (ASAM). It solves the scale dependency problem by removing the effect of scaling and helps build a solid correlation with the generalization gap. In the following subsections, we detail both SAM and ASAM.

Table 3: *In-domain and out-domain evaluation results for ASVspoof 2015, ASVspoof 2019 LA, ASVspoof 2021 LA & DF. For model training, ASVspoof 2015, ASVspoof 2019 LA, and WaveFake are used. The first row and column refer to training and evaluation datasets, respectively. The boldface in the last row and the last column indicate the best results between training objectives and training datasets, respectively. On the other hand, the underline results refer to the best results in each evaluation set.*

	ASVspoof 2015			ASVspoof 2019 LA			ASVspoof 2021 LA			ASVspoof 2021 DF			Average
	w/o SAM	SAM	ASAM	w/o SAM	SAM	ASAM	w/o SAM	SAM	ASAM	w/o SAM	SAM	ASAM	
2015	8.25	6.50	5.83	38.83	29.50	30.70	39.87	32.27	31.09	33.54	28.40	21.80	25.55
2019 LA	5.98	4.32	3.53	1.38	1.06	1.48	12.18	<u>7.08</u>	10.18	<u>18.20</u>	21.16	19.58	8.84
2015 + 2019 LA	0.69	0.80	0.83	1.56	1.07	<u>0.99</u>	12.25	12.52	9.37	21.02	26.32	22.14	9.13
2015 + 2019 LA + WaveFake	0.82	<u>0.66</u>	1.82	1.76	1.78	1.27	11.05	13.08	12.41	19.76	19.49	19.84	8.65
<i>Average</i>	3.94	3.07	3.00	10.88	8.35	8.61	18.84	16.24	15.76	23.13	23.84	20.84	

3.1. Sharpness-Aware Minimization

Given a labeled training set $S = \{(x_1, y_1), \dots, (x_n, y_n)\}$ drawn from i.i.d an unknown data distribution D , the training loss function with model parameter \mathbf{w} and population loss are defined as L_S and $L_D(\mathbf{w}) \triangleq \mathbb{E}_{(x,y) \sim D} [l(\mathbf{w}, x, y)]$, respectively. While the training loss is the empirical/sample-based estimator, the population loss refers to the corresponding theoretical quantity when the knowledge of the actual joint distribution (x, y) is fully known. So, population loss can be thought of as the empirical training loss, applied to an infinitely large training set ($n \rightarrow \infty$).

Our goal is to select \mathbf{w} not only for having low training loss $L_S(\mathbf{w})$ but also for low population loss $L_D(\mathbf{w})$, for improved generalization. To achieve such a goal, SAM is designed to minimize the following PAC-Bayesian generalization upper bound, where the sharpness term appears explicitly in another following equation. (The term in brackets indicates the sharpness of L_S at \mathbf{w}):

$$L_D(\mathbf{w}) \leq \max_{\|\varepsilon\|_2 \leq \rho} L_S(\mathbf{w} + \varepsilon) + h(\|\mathbf{w}\|_2^2 / \rho^2)$$

$$= [\max_{\|\varepsilon\|_2 \leq \rho} L_S(\mathbf{w} + \varepsilon) - L_S(\mathbf{w})] + L_S(\mathbf{w}) + h(\|\mathbf{w}\|_2^2 / \rho^2)$$

Here, h is a strictly increasing function under conditions on $L_D(\mathbf{w})$ and ρ is a predefined constant controlling the radius of a neighborhood in l^p ball ($p \in [1, \infty]$, [28] revealed that $p = 2$ is optimal). A detailed explanation of the PAC-Bayesian generalization bound is omitted for the limited space. Refer to Appendix A.1 of [28]) for full details.

Finally, for any $\rho > 0$ and $\varepsilon \approx 0$ (to avoid division by 0), the model loss is defined as:

$$\min_{\mathbf{w}} L_S^{\text{SAM}}(\mathbf{w}) + \lambda \|\mathbf{w}\|_2^2,$$

where $L_S^{\text{SAM}}(\mathbf{w}) \triangleq \max_{\|\varepsilon\|_p \leq \rho} L_S(\mathbf{w} + \varepsilon)$

3.2. Adaptive Sharpness-Aware Minimization

Even if the vanilla SAM performs usually well, it is easily affected by parameter re-scaling as a sharpness term in SAM defined on a rigid region with a fixed radius. This may disturb the generalization of SAM.

To address this shortcoming, ASAM [29] uses adaptive sharpness that removes the effect of scaling and adjusting the

Table 4: *Per-attack results of ASVspoof 2021 DF evaluation. The best results are selected between w/o SAM, SAM, and ASAM. The best results in the same row are represented in boldface. ((a) ASVspoof 2015, (b) ASVspoof 2019, (c) WaveFake.)*

Attack	(a)	(b)	(a)+(b)	(a)+(b)+(c)
Traditional	21.54	12.18	14.18	10.77
Wav.Concat.	55.22	12.07	20.32	13.09
Neural AR	46.82	23.10	27.70	24.87
Neural non-AR	40.23	20.47	25.05	23.21
Unknown	24.45	20.38	17.94	16.05
Pooled	33.54	18.20	22.14	19.49

maximization region, leading to an improved training path. It utilizes the normalization operator $T_{\mathbf{w}}^{-1}$ and achieves better generalization compared to SAM. Firstly, the normalization operator of weight \mathbf{w} can be defined, if $T_{\mathbf{w}}$ is a family of invertible linear operators and $T_{\mathbf{A}}^{-1} \mathbf{w} = T_{\mathbf{w}}^{-1}$ for any invertible scaling operator \mathbf{A} , which does not alter the loss function. With this normalization operator, adaptive sharpness objective function is defined as:

$$L_S^{\text{ASAM}}(\mathbf{w}) \triangleq \max_{\|T_{\mathbf{w}}^{-1} \varepsilon\|_p \leq \rho} L_S(\mathbf{w} + \varepsilon)$$

4. Experimental settings

For experiments, we deploy the ‘‘light’’ version of the recent AASIST model [8], referred to as AASIST-L. It includes a graph attention layer to capture information both in spectral and temporal domains and max graph operations to select features in a competitive manner. The main difference between AASIST and AASIST-L is the number of parameters for practical purposes including 297K and 85K numbers of parameters, respectively. We use Adam [46] as our base optimizer. When exploiting SAM and ASAM, optimization proceeds similarly, but with the additional sharpness term added to the training loss as explained above. As our aim is to focus on the generalization throughout corpora rather than dealing with architectural details, we did not adjust parameters (e.g. learning rate, pooling ratio). Full details of AASIST model can be referred to [8]. All models were implemented using PyTorch and trained for 100 epochs. Performance evaluation is done based on equal error rate (EER) and we selected the best-performing model in terms of EER on the development set. Code for experiments of this study is available in: <https://github.com/ASIST-L>

Table 5: The comparison of mini-batch composition strategies using ASVspoof 2019 LA evaluation. Pooled mini-batch refers to the condition which ignores the balance between datasets, and balanced mini-batch refers to the condition which considers the balance between multiple datasets.

Training datasets	Mini-batch	Loss	EER
2015 + 2019 LA	pooled	w/o SAM	1.56
2015 + 2019 LA	balanced	w/o SAM	1.49
2015 + 2019 LA + WaveFake	pooled	ASAM	1.27
2015 + 2019 LA + WaveFake	balanced	ASAM	1.09

//github.com/shimhz/MDL_sharpness.

5. Results and Analyses

Main results In Table 3, we validate our proposed model using both in-domain datasets (ASVspoof 2015 and ASVspoof 2019 LA) and out-of-domain datasets (ASVspoof 2021 LA and ASVspoof 2021 DF). Firstly, as for the multi-dataset co-training, the pooled results are shown in the last column. We confirm that the best result achieved when all three datasets are all utilized. Even though the full usage of three datasets did not show a consistent result as the best result, the effectiveness of using multiple datasets was demonstrated.

Secondly, the results of sharpness-aware optimizations can be easily comparable in the last row. Sharpness-aware optimization methods improve the performance in most cases, regardless of whether using a single dataset or multiple datasets. Except for the ASVspoof 2019 LA evaluation case that gets the lowest EER using SAM, ASAM shows the best results. Through results, we observe that SAM and ASAM substantially benefit the model optimization.

Per-attack results on ASVspoof 2021 DF For further analysis, Table 4 shows the results for each attack in ASVspoof 2021 DF evaluation result. We depict the best results for each training dataset combination : each column refers to ASVspoof 2015 (w/ ASAM), ASVspoof 2019 (w/o SAM), ASVspoof 2015 + ASVspoof 2019 (w/o SAM), and ASVspoof2015 + ASVspoof 2019 + WaveFake (w/ SAM). The interesting thing to be noted is that (a)+(b)+(c) which includes all datasets show superior in Traditional and Unknown attack in a large gap compared to the best result which is from (b) only. In terms of generalization, *unknown* is the most critical subset; thus we interpret that the lowest EER in unknown signifies better generalization. Hence, these results back up the effectiveness of the proposed method for the generalization performance.

Mini-batch composition When utilizing multiple datasets, an imbalance exists between different datasets. We thus further explore whether balancing the number of samples drawn from each dataset can be advantageous in terms of performance. Table 5 describes the results, which confirm that simply balancing the samples between different datasets within a mini-batch are helpful. We confirm improvement in both reported train dataset configurations, where the performance of the most extensive setting is further improved by 14% relative.

Comparison with other studies In Table 6, we compare our results with other state-of-the-art systems including the studies utilizing a large pre-trained model. Among four evaluation protocols, our model demonstrates competitive performance

Table 6: Comparison with other state-of-the-art results including the research which utilized pre-trained model with large datasets.

Method	# Params	EER(%)
ASVspoof 2015		
Primary result [47]	-	1.21
Ours	85K	0.66
ASVspoof 2019 LA		
SSAD+LCNN big [48]	-	5.31
Imag-pre + J&S [16]	-	0.87
wav2Vec2.0 [18]	317M + (290±30K)	1.28
HuBERT-XL [18]	317M + (290±30K)	3.55
Ours	85K	0.99
ASVspoof 2021 LA		
Img-pre+RawBoost [16]	-	7.71
wav2Vec2.0-XLSR [17]	317M + 297K	6.15
wav2Vec2.0-XLSR [18]	317M + 297K	9.66
HuBERT-XL [18]	317M + (290±30K)	9.55
Ours	85K	7.08
ASVspoof 2021 DF		
Img-pre+RawBoost [16]	-	19.11
wav2Vec2.0-XLSR [17]	317M + 297K	7.69
wav2Vec2.0-XLSR [18]	317M + 297K	4.75
HuBERT-XL [18]	317M + (290±30K)	13.07
Ours	85K	18.20

with only 85K parameters in two protocols: ASVspoof 2015 and ASVspoof 2019 LA. In the other two remaining protocols, our model underperforms; nonetheless, taking into account the number of parameters and training time, we argue that our approach remains competitive. Given the fact that the purpose of CM models is to aid ASV systems, lightweight yet well-generalizing models are worth further investigation.

6. Conclusions

Recent studies have widely exploited large pre-trained models to leverage as much data as possible to develop a well-generalized model. While training a single model with multiple datasets is a straightforward way to utilize diverse data without additional training, it is well-known that handling domain differences is challenging. Given the nature that CM models are inherently built to support ASV systems, these enormous systems can be potentially not applicable because of their size. In this paper, we explore a case study in the audio anti-spoofing field which lacks a large amount of data compared to other research domains. To optimize the model to handle multiple datasets simultaneously, we utilize sharpness-aware methodologies, which include a curvature-based term in the objective function to reduce the gap between the variance of the data. Using a number of parameters more than 4000 times less than the large pre-trained models, our proposed method demonstrates effectiveness in both in-domain evaluations on unknown attacks and out-of-domain evaluations.

7. Acknowledgements

This work was supported by the Academy of Finland (Decision No. 349605, project ‘‘SPEECHFAKES’’)

8. References

- [1] D. A. Reynolds, "An overview of automatic speaker recognition technology," in *Proc. ICASSP*, 2002.
- [2] R. K. Das, J. Yang and H. Li, "Long range acoustic and deep features perspective on asvspoof 2019," in *Proc. ASRU*, 2019.
- [3] J. Yang, R. K. Das and H. Li, "Significance of subband features for synthetic speech detection," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 2160–2170, 2019.
- [4] J. Yang and R. K. Das, "Long-term high frequency features for synthetic speech detection," *Digital Signal Processing*, vol. 97, pp. 102622, 2020.
- [5] C.-I. Lai, N. Chen, J. Villalba and N. Dehak, "Assert: Anti-spoofing with squeeze-excitation and residual networks," in *Proc. Interspeech*, 2019.
- [6] X. Li, N. Li, C. Weng et al., "Replay and synthetic speech detection with res2net architecture," in *Proc. ICASSP*, 2021.
- [7] H. Tak, J.-w. Jung, J. Patino et al., "End-to-end spectro-temporal graph attention networks for speaker verification anti-spoofing and speech deepfake detection," in *Proc. ASVspoof 2021 Workshop (Interspeech satellite)*, 2021.
- [8] J.-w. Jung, H.-S. Heo, H. Tak et al., "AASIST: Audio Anti-Spoofing using Integrated Spectro-Temporal Graph Attention Networks," in *Proc. ICASSP*, 2022.
- [9] Y. Zhang, F. Jiang and Z. Duan, "One-class learning towards synthetic voice spoofing detection," *IEEE Signal Processing Letters*, vol. 28, pp. 937–941, 2021.
- [10] A. Cohen, I. Rimon, E. Afalo and H. H. Permuter, "A study on data augmentation in voice anti-spoofing," *Speech Communication*, vol. 141, pp. 56–67, 2022.
- [11] H. Tak, M. Kamble, J. Patino et al., "Rawboost: A raw data boosting and augmentation method applied to automatic speaker verification anti-spoofing," in *Proc. ICASSP*, 2022.
- [12] H. Wang, H. Dinkel, S. Wang et al., "Cross-domain replay spoofing attack detection using domain adversarial training," in *Proc. Interspeech*, 2019.
- [13] R. K. Das, J. Yang and H. Li, "Assessing the scope of generalized countermeasures for anti-spoofing," in *Proc. ICASSP*, 2020.
- [14] B. Chettri, R. G. Hautamäki, M. Sahidullah and T. Kinnunen, "Data quality as predictor of voice anti-spoofing generalization," in *Proc. Interspeech*, 2021.
- [15] Y. Zhang, G. Zhu, F. Jiang and Z. Duan, "An empirical study on channel effects for synthetic voice spoofing countermeasure systems," in *Proc. Interspeech*, 2021.
- [16] J. Lu, Z. Li, Y. Zhang et al., "Acoustic or pattern? speech spoofing countermeasure based on image pre-training models," in *Proceedings of the 1st International Workshop on Deepfake Detection for Audio Multimedia*, 2022, pp. 77–84.
- [17] H. Tak, M. Todisco, X. Wang et al., "Automatic speaker verification spoofing and deepfake detection using wav2vec 2.0 and data augmentation," *Proc. Interspeech*, 2022.
- [18] X. Wang and J. Yamagishi, "Investigating active-learning-based training data selection for speech spoofing countermeasure," *Proc. SLT*, 2022.
- [19] H. Ma, J. Yi, J. Tao et al., "Continual learning for fake audio detection," *Proc. Interspeech*, 2021.
- [20] H.-j. Shim, H.-S. Heo, J.-w. Jung and H.-J. Yu, "Self-supervised pre-training with acoustic configurations for replay spoofing detection," *Proc. Interspeech*, 2019.
- [21] Z. Lv, S. Zhang, K. Tang and P. Hu, "Fake audio detection based on unsupervised pretraining models," in *Proc. ICASSP*, 2022.
- [22] A. Baeovski, Y. Zhou, A. Mohamed and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *NeurIPS*, 2020.
- [23] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai et al., "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM TASLP*, vol. 29, 2021.
- [24] P. Meletis and G. Dubbelman, "Training of convolutional networks on multiple heterogeneous datasets for street scene semantic segmentation," in *2018 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2018, pp. 1045–1050.
- [25] G. Kapidis, R. Poppe and R. C. Veltkamp, "Multi-dataset, multi-task learning of egocentric vision tasks," *IEEE PAMI*, 2021.
- [26] X. Zhou, V. Koltun and P. Krähenbühl, "Simple multi-dataset detection," in *Proc. CVPR*, 2022.
- [27] J. Liang, E. Zhang, J. Zhang and C. Shen, "Multi-dataset training of transformers for robust action recognition," *NeurIPS*, 2022.
- [28] P. Foret, A. Kleiner, H. Mobahi and B. Neyshabur, "Sharpness-aware minimization for efficiently improving generalization," *Proc. ICLR*, 2020.
- [29] J. Kwon, J. Kim, H. Park and I. K. Choi, "Asam: Adaptive sharpness-aware minimization for scale-invariant learning of deep neural networks," in *Proc. ICML*. PMLR, 2021.
- [30] Y. Liu, S. Mai, X. Chen et al., "Towards efficient and scalable sharpness-aware minimization," in *Proc. CVPR*, 2022.
- [31] X. Zhu, C. Vondrick, C. C. Fowlkes and D. Ramanan, "Do we need more training data?," *IJCV*, vol. 119, 2016.
- [32] D. Paul, M. Sahidullah and G. Saha, "Generalization of spoofing countermeasures: A case study with asvspoof 2015 and btas 2016 corpora," in *Proc. ICASSP*, 2017.
- [33] Z. Wu, T. Kinnunen, N. Evans et al., "Asvspoof 2015: the first automatic speaker verification spoofing and countermeasures challenge," in *Proc. Interspeech*, 2015.
- [34] M. Todisco, X. Wang, V. Vestman et al., "ASVspoof 2019: Future horizons in spoofed and fake audio detection," in *Proc. Interspeech*, 2019.
- [35] J. Frank and L. Schönherr, "WaveFake: A Data Set to Facilitate Audio Deepfake Detection," in *NeurIPS*, 2021.
- [36] J. Yamagishi, C. Veaux, K. MacDonald et al., "CSTR VCTK corpus: English multi-speaker corpus for cstr voice cloning toolkit," 2019, <http://dx.doi.org/10.7488/ds/1994>.
- [37] K. Ito and L. Johnson, "The lj speech dataset," <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- [38] R. Sonobe, S. Takamichi and H. Saruwatari, "JSUT corpus: free large-scale japanese speech corpus for end-to-end speech synthesis," *CoRR*, vol. abs/1711.00354, 2017.
- [39] J. Yamagishi, X. Wang et al., "ASVspoof2021: Accelerating progress in spoofed and deep fake speech detection," in *Proc. ASVspoof 2021 Workshop (Interspeech satellite)*, 2021.
- [40] J. Lorenzo-Trueba, J. Yamagishi, T. Toda et al., "The voice conversion challenge 2018: Promoting development of parallel and nonparallel methods," *Proc. Speaker Odyssey Workshop*, 2018.
- [41] Y. Zhao, W.-C. Huang, X. Tian et al., "Voice conversion challenge 2020: Intra-lingual semi-parallel and cross-lingual voice conversion," *ISCA Joint Workshop for the Blizzard Challenge and Voice Conversion Challenge 2020*, 2020.
- [42] D. Bahri, H. Mobahi and Y. Tay, "Sharpness-aware minimization improves language model generalization," *ACL*, 2022.
- [43] M. Andriushchenko and N. Flammarion, "Towards understanding sharpness-aware minimization," in *Proc. ICML*, 2022.
- [44] J. M. Coldenhoff, C. Li and Y. Zhu, "Model generalization: A sharpness aware optimization perspective," *Proc. ICLR*, 2022.
- [45] D. Caldarola, B. Caputo and M. Ciccone, "Improving generalization in federated learning by seeking flat minima," in *Proc. ECCV*, 2022.
- [46] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [47] Z. Wu, T. Kinnunen, N. Evans et al., "Asvspoof 2015: the first automatic speaker verification spoofing and countermeasures challenge," in *Proc. Interspeech*, 2015.
- [48] Z. Jiang, H. Zhu, L. Peng et al., "Self-supervised spoofing audio detection scheme," in *Proc. Interspeech*, 2020.