Improving Small Footprint Few-shot Keyword Spotting with Supervision on Auxiliary Data

Seunghan Yang, Byeonggeun Kim, Kyuhong Shim, Simyung Chang

Qualcomm AI Research[†], Qualcomm Korea YH, Seoul, Republic of Korea

{seunghan, kshim, simychan}@qti.qualcomm.com

Abstract

Few-shot keyword spotting (FS-KWS) models usually require large-scale annotated datasets to generalize to unseen target keywords. However, existing KWS datasets are limited in scale and gathering keyword-like labeled data is costly undertaking. To mitigate this issue, we propose a framework that uses easily collectible, unlabeled reading speech data as an auxiliary source. Self-supervised learning has been widely adopted for learning representations from unlabeled data; however, it is known to be suitable for large models with enough capacity and is not practical for training a small footprint FS-KWS model. Instead, we automatically annotate and filter the data to construct a keyword-like dataset, LibriWord, enabling supervision on auxiliary data. We then adopt multi-task learning that helps the model to enhance the representation power from out-of-domain auxiliary data. Our method notably improves the performance over competitive methods in the FS-KWS benchmark.
Index Terms: few-shot learning, keyword spotting

1. Introduction

Few-shot keyword spotting (FS-KWS) refers to the task of recognizing specific keywords in an audio signal, where only a limited number of examples are available for each keyword. This task has gained significant attention in recent years due to its practical importance in various applications such as voice assistants requiring user-defined keywords [1, 2, 3, 4, 5, 6, 7]. To solve this task, few-shot learning approaches have been proposed to learn a discriminative embedding space that can generate general representations for novel classes only given few examples. They mainly leverage prior knowledge from related tasks or similar data distributions using large-scale annotated training datasets [8, 9, 10, 11].

However, the keyword spotting (KWS) datasets [12, 13, 14] that are accessible to the public are typically limited in size and have a small number of categories. Acquiring a comprehensive annotated dataset suitable for KWS is a costly work. Given the existing datasets, the performance of FS-KWS is limited, as reported in Fig 1(a), even with the use of an advanced few-shot learning technique [15]. To this end, we propose to exploit auxiliary data, especially reading speech data from audiobooks, such as LibriSpeech [16] and MLS [17], which are readily publicly available in large quantities and easily collectable. These data are not built for the keyword spotting task (*e.g.*, LibriSpeech [16] for speech recognition), but they are expected to help the model learn a more robust embedding space.

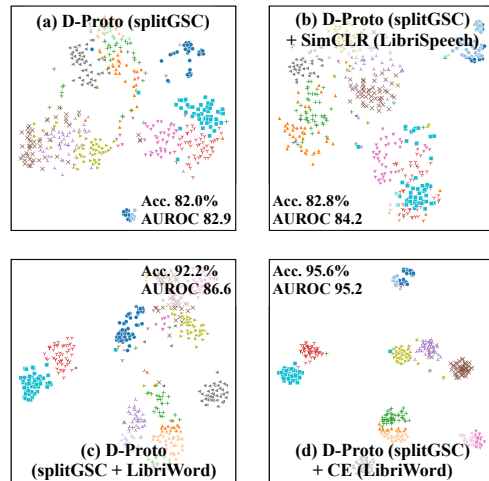One way to leverage knowledge from a large-scale unlabeled dataset is through self-supervised learning (SSL) [18, 19,

Figure 1: *t-SNE Visualization of Unseen Keywords Embeddings Using Various Learning Methods on the FS-KWS Benchmark [15]. We train four models using in-domain split-GSC and out-of-domain reading speech datasets, including LibriSpeech and our proposed word-level labeled dataset, LibriWord:* **(a)** *metric learning [15] on splitGSC,* **(b)** *(a) + self-supervised learning [18] on LibriSpeech,* **(c)** *metric learning on both datasets (splitGSC + LibriWord), and* **(d)** *our proposed AuxSL on splitGSC and LibriWord. The closed-set classification accuracy (Acc.) and open-set detection rate (AUROC) indicate the performance of the models.*

20, 21]. SSL has gained widespread usage in various applications due to its many advantages, particularly its ability to produce robust feature representations without the need for labeled data. SSL has been successfully employed in the pre-training stage of the model [19, 21] or in combination with a target objective for fine-tuning the model [22, 23]. However, prior literature has established that SSL is only beneficial to large models, as evidenced by [24, 25]. Therefore, the lightweight nature of FS-KWS models presents a challenge for the use of SSL. In this paper, we empirically derive that the existing SSL methods are also not suitable for training lightweight models in FS-KWS, as shown in Fig. 1(b). Instead, we automatically annotate the reading speech data with a word extraction technique [26] and construct a well-organized keyword-like dataset called *LibriWord* by filtering and balancing the data, effectively enabling the use of auxiliary data with supervision on lightweight models.

To effectively leverage both in-domain command-like data and out-of-domain auxiliary data, we propose a simple but effective multi-task learning (MTL) framework, called AuxSL, that incorporates an additional classifier with a supervised loss function for the out-of-domain data. AuxSL enables the model

to enhance the representation power from out-of-domain auxiliary data while primarily learning keyword representations on in-domain data. The naive approach of combining the two datasets and applying a single loss function to the entire dataset, as shown in Fig. 1(c), results in limited performance improvement. Our proposed framework is evaluated through extensive experiments on the FS-KWS benchmark, demonstrating its superior discriminative capabilities. In Fig. 1(d), the results show a relative improvement of 16% in both closed- and open-set settings when using 5-shot samples, compared to the baseline.

2. Related Works

2.1. Small Footprint Keyword Spotting

The aim of conventional keyword spotting is to detect a small set of pre-defined speech signals, such as the activation phrases "Alexa" and "OK Google". Recent research has primarily focused on enhancing accuracy while minimizing memory requirements and reducing power consumption, allowing for the deployment to an always-on system [27, 28, 29, 30]. Some studies have centered around the development of optimal architectures for keyword spotting [27, 28], while others have aimed to improve the loss function [29] or training approach [30].

2.2. Few-shot Keyword Spotting

Few-shot keyword spotting (FS-KWS) aims to support systems that require user-defined keywords. Unlike conventional keyword spotting, FS-KWS involves a scenario in which new keywords that were not seen during training are enrolled and tested on the device. To achieve FS-KWS, it is crucial to learn robust representations in the pre-training stage using a large-scale dataset. To further improve performance, various loss functions [2, 3, 4], architectures [31], and learning strategies [5, 6, 7] have been proposed. Modifying the learning strategies is closely related to our work. Kao et al. [5] present a two-stage framework that compares several popular SSL models to determine the best model for FS-KWS. Lee et al. [6] incorporate auxiliary synthetic data generated by TTS, while Shin et al. [7] propose a novel approach that leverages the linguistically corresponding patterns between speech and text sequences. Our approach is distinct from prior studies as we propose utilizing auxiliary data from reading speech domain along with a supervised multi-task learning strategy, with the objective of training lightweight models. Recently, D-ProtoNets [15] has been proposed for few-shot open-set keyword spotting, which involves open-set classes during test time. Our proposed method is expected to improve open-set performance in addition to closed-set performance, as it helps to create robust representations.

3. Proposed Methods

3.1. Problem Definition

Training objectives for few-shot learning: In the domain of few-shot learning, metric learning based techniques have been predominantly investigated for constructing a robust embedding space that enables models to extract discriminative embeddings for unseen keywords and perform keyword detection using distance metrics. One of the most popular metric learning-based methods is ProtoNets [8]. At each training iteration, ProtoNets create a classification scenario by constructing a N-way and K-shot problem from training data, where N represents the number of classes and K denotes the number of support

Table 1: *Metadata of LibriSpeech and proposed LibriWord.*

Info.	LibriSpeech	LibriWord
# of samples	4,806,675	300,000
# of keywords	58,997	1,000
Audio length	[0.0s, 11.1s]	[0.3s, 2.2s]
Annotations	Word-level labels	
Annotator	Montreal Forced Aligner [26, 32]	

samples per class. ProtoNets first construct the prototype of each class by averaging the embeddings of support samples, $c_n = \frac{1}{K} \sum_{i=1}^{K} F_\theta(x_{n,i}^s)$, where $F_\theta(\cdot)$ is a feature extractor. Subsequently, the prototypical loss function forces the query samples to minimize the distance from the corresponding prototypes as follows:

$$L_{FSL} = -\frac{1}{M} \sum_{i=1}^{M} \log p_\theta(y = n | x_{n,i}^q), \quad \text{where}$$

$$p_\theta(y = n | x_{n,i}^q) = \frac{\exp(-d(F_\theta(x_{n,i}^q), c_n))}{\sum_{n'=1}^{N} \exp(-d(F_\theta(x_{n,i}^q), c_{n'}))}. \quad (1)$$

$d(\cdot)$ can be any distance metric, $x_{n,i}^q$ is i-th query sample of class n, and M indicates the number of query samples of class n in the N-way and K-shot problem. At the inference stage, prototypes of new classes are constructed by averaging the embeddings of enrolled samples. The test samples are then classified based on the classification probability in Eq. 1.

Limitation of few-shot keyword spotting: Despite the success of few-shot learning, a large-scale dataset is required to learn a robust embedding space. For keyword spotting, the datasets usually have limited sizes, making it difficult to learn general keyword representations. While collecting in-domain command-like data is an effective way to create a robust embedding space, it is challenging to obtain such datasets. To address this issue, we propose to use easily obtainable auxiliary data, especially reading speech data.

3.2. LibriWord

Instead of collecting human-labeled keyword spotting data, we create a dataset named LibriWord containing segmented utterances and corresponding word-level labels. The samples are obtained from the LibriSpeech corpus [16], which comprises roughly 1,000 hours of read English speech at 16kHz. It contains numerous words, but lacks word-level alignments and only utterance-level transcriptions are provided. To obtain word-level segmented samples, we employ the Montreal Forced Aligner [32, 26], a word extraction technique previously used in [7, 33]. We then form a balanced dataset by organizing the extracted words based on the number of samples and eliminating similar keywords. Specifically, when a word partially overlapped with another word, such as with past tense, plural forms, or negative forms, we randomly remain one of them to simplify the representation learning process for small models. As a result, LibriWord includes 300 samples for each top 1,000 frequently appearing keywords. Table 1 presents the metadata before and after the refinement process, respectively. Learning with LibriWord, a comparatively smaller dataset than LibriSpeech, yields the benefits of reducing the burden of collecting auxiliary data and saving storage space, while achieving better performance (please see Section 4.2).

3.3. FS-KWS with Auxiliary Supervision on LibriWord

We propose a simple but effective multi-task learning framework, called AuxSL, which utilizes an additional classifier and
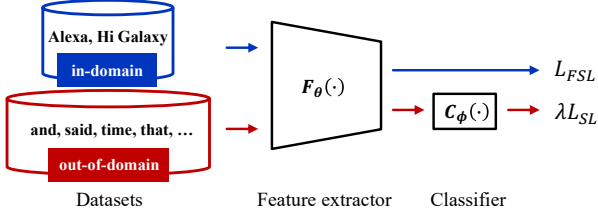
Figure 2: *Our Proposed AuxSL Framework.* $F_\theta(\cdot)$ is trained with both in-domain data for keyword spotting and out-of-domain auxiliary data from the reading speech source, and $C_\phi(\cdot)$ is trained to classify auxiliary words for helping to learn robust representations.

supervised loss function for the out-of-domain auxiliary data to mitigate potential representation skew that may arise when directly sharing the feature extractor between in-domain and out-of-domain data, motivated by [34]. Our multi-task learning objectives are calculated as follows:

$$L_{AuxSL} = L_{FSL} + \lambda L_{SL}, \qquad (2)$$

where L_{FSL} represents any few-shot learning loss function on in-domain training data and L_{SL} represents supervised loss function on out-of-domain auxiliary data. λ is a balancing parameter for auxiliary loss. Here, we maintain metric learning loss on in-domain data and use conventional cross-entropy loss, as depicted in Figure 2. This approach does not incur additional cost during inference since the classifier is not used.

For our experiment, we employ a dummy prototypical loss function of D-ProtoNets [15] for L_{FSL} to efficiently handle both closed- and open-set query samples. D-ProtoNets [15] are trained using the dummy prototypical loss, which incorporates a learnable open-set prototype specifically designed to represent the open-set class. The open-set prototype is trained jointly with the class-wise prototypes to enable query samples of both closed- and open-set to be closely associated with their corresponding prototypes in the $N + 1$ classification task. During inference, if the probability of a query test sample x_t^q belonging to the open-set class $N + 1$, *i.e.*, $p_\theta(y = N + 1|x_t^q)$, exceeds a pre-defined threshold, then it is verified as the open-set class.

4. Experiments

4.1. Experimental Settings

Dataset. We use a standard benchmark splitGSC [15] on the Google Speech Commands (GSC) dataset [12] for keyword spotting task. The splitGSC contains the train, validation, and test split designed for few-shot closed- and open-set keyword spotting. The split includes 15, 10, and 10 keywords for training, validation, and testing, respectively, with 24,400, 4,007, and 4,482 samples, respectively. Note that a silence category is included in all sets. The official background noise provided by GSC with a probability of 0.8 is used. See [15] for more details.
Implementation details. We use three different small footprint backbone models, BC-ResNet8 [27] and Res12 [35], both of which take 40-dimensional log Mel spectrograms as input with a window length of 30 ms and frame shift of 10 ms, and DS-ResNet18 [28], which takes 40-dimensional Mel-frequency cepstrum coefficient features as input. The network size of BC-ResNet8, Res12, and DS-ResNet18 is 321k, 8M, and 72k, respectively. Each model is trained for 100 epochs using the Adam optimizer with an initial learning rate of 0.001, which is step decayed by a factor of 0.5 every 20 epochs. Each epoch consists of 500 episodes, each containing 5 closed- and 5 open-
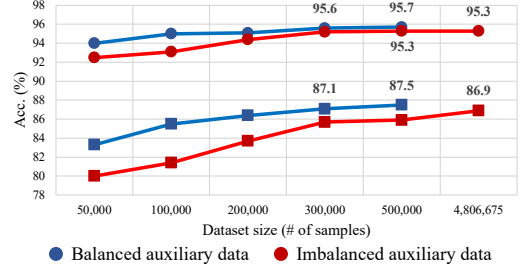

Figure 3: *The accuracy of models trained with balanced and imbalanced auxiliary data on BC-ResNet8 using the MTL path (our optimal framework). The 5-shot and 1-shot performance are denoted by circle and square, respectively.*

set classes, with 5 support samples for prototypes and 5 query samples for each class. For multi-task learning, we use a batch size of 64 with parallel sampling as in the episodic iterations, and set λ to 1.0. As an additional module, we used a 1-FC layer as the classifier. To evaluate the trained model, we used $1,000$ episodes at test time, with 1 or 5 support samples and 15 query samples per class including open-set classes. We reported the average 1-shot and 5-shot accuracy and threshold-free area under the receiver-operating characteristics (AUROC) for closed- and open-set performance with 3 different seeds.

4.2. Analysis of LibriWord Dataset

In Figure 3, we present the results obtained by training with different dataset constructions after extracting words from the imbalanced Librispeech corpus, which has highly skewed distribution of samples among its words. Notably, in this experiment, we employ our final MTL architecture solely to investigate the impact of LibriWord's dataset configuration. Our findings indicate that constructing a dataset with a balanced number of samples per keyword leads to superior performance compared to using an equally imbalanced dataset or the entire Librispeech dataset, which is roughly 16 times larger than LibriWord. Using a balanced number of data per keyword during training aids in building robust feature embeddings, a phenomenon that has also been observed in ImageNet, where research on handling imbalanced datasets is actively being pursued [36, 37]. To address the drawbacks of imbalanced datasets, we constructed LibriWord in a balanced manner and demonstrate its empirical efficacy.

4.3. Comparison of SSL and SL Methods on Auxiliary Data

In Table 2, we compared the performance of few-shot keyword spotting models trained using self-supervised learning (SSL) and supervised learning (SL) on auxiliary data. The baseline model, D-Proto [15], was trained with SL on splitGSC without using auxiliary data. We evaluated three SSL methods: (1) PreT-Big, which uses a large-scale pre-trained feature extractor from LibriSpeech and fine-tunes only the classifier for keyword spotting, (2) PreT, which pre-trains a keyword spotting model on LibriSpeech using SimCLR and BYOL, (3) MTL with SSL, which uses SSL on LibriSpeech and SL on splitGSC, and (4) MTL with knowledge distillation (KD), which uses feature distillation [38] from the large-scale pre-trained feature extractor and SL on splitGSC. We also evaluated three SL methods: (1) PreT, which pre-trains the keyword spotting model using cross-entropy loss (CE), (2) ALL, which uses all data together to train the model, and (3) AuxSL, our proposed method that uses a different path for each dataset. Balancing parameters for MTL methods are set to 1 and all hyper-parameters are chosen based

Learning scheme on auxiliary data	Training strategy	Method	5-shot Acc. (%)	5-shot AUROC	1-shot Acc. (%)	1-shot AUROC
X	Baseline (321k)	D-Proto (splitGSC)	82.0 (0.7)	82.9 (0.1)	65.8 (0.3)	75.7 (0.1)
Self-supervised learning	PreT-Big (300M)	Wav2vec (LibriSpeech) → D-Proto (splitGSC)	85.3 (0.9)	76.7 (1.2)	66.6 (1.5)	70.5 (0.4)
		Hubert (LibriSpeech) → D-Proto (splitGSC)	90.4 (0.3)	85.2 (0.6)	76.2 (0.8)	77.3 (0.9)
	PreT	SimCLR (LibriSpeech) → D-Proto (splitGSC)	79.9 (1.0)	81.2 (1.0)	64.7 (1.0)	75.0 (0.8)
		BYOL (LibriSpeech) → D-Proto (splitGSC)	82.0 (1.5)	82.7 (1.0)	65.7 (1.6)	76.2 (1.0)
	MTL	SimCLR (LibriSpeech) + D-Proto (splitGSC)	83.0 (0.8)	83.1 (0.9)	66.6 (1.6)	76.0 (1.0)
		BYOL (LibriSpeech) + D-Proto (splitGSC)	81.7 (0.0)	82.7 (0.0)	65.2 (0.4)	76.0 (0.5)
		KD w/ Wav2vec (splitGSC) + D-Proto (splitGSC)	82.4 (0.7)	83.3 (0.1)	66.4 (0.6)	76.4 (0.2)
		KD w/ Hubert (splitGSC) + D-Proto (splitGSC)	83.7 (0.5)	84.4 (0.5)	67.8 (0.6)	77.3 (0.9)
Supervised learning	PreT	CE (LibriWord) → D-Proto (splitGSC)	82.8 (1.6)	84.2 (1.1)	68.7 (1.5)	77.3 (0.8)
	ALL	D-Proto (LibriWord + splitGSC)	92.2 (0.9)	86.6 (1.3)	77.4 (2.1)	75.3 (1.6)
	Ours: AuxSL	**CE (LibriWord) + D-Proto (splitGSC)**	**95.6 (0.5)**	**95.2 (0.4)**	**87.1 (0.8)**	**89.4 (0.3)**

Table 3: *FS-KWS performance on various backbones.*

Training strategy	Backbones	5-shot Acc. (%)	5-shot AUROC	1-shot Acc. (%)	1-shot AUROC
Baseline	Res12	86.8 (0.2)	85.9 (0.3)	70.4 (0.2)	78.5 (0.1)
	DS-ResNet18	69.2 (2.9)	74.8 (1.7)	53.7 (2.2)	70.3 (1.2)
PreT (SimCLR)	Res12	83.3 (1.1)	83.8 (0.9)	68.7 (1.5)	77.1 (0.5)
	DS-ResNet18	59.0 (2.7)	69.3 (0.9)	45.8 (2.0)	65.7 (0.3)
KD (Hubert)	Res12	87.0 (0.5)	85.3 (0.1)	70.3 (0.8)	78.1 (0.4)
KD (Wav2vec)	DS-ResNet18	70.8 (0.5)	75.5 (0.6)	55.0 (0.3)	70.6 (0.3)
ALL	Res12	93.1 (0.2)	86.9 (1.7)	78.3 (0.4)	75.4 (1.0)
	DS-ResNet18	88.0 (1.3)	82.1 (2.2)	71.7 (1.2)	72.1 (2.3)
Ours: AuxSL	Res12	**95.6 (0.3)**	**94.1 (0.6)**	**85.5 (0.2)**	**84.3 (1.4)**
	DS-ResNet18	**90.8 (0.3)**	**89.2 (0.9)**	**77.4 (1.0)**	**80.2 (1.8)**

Figure 4: *A performance comparison on various model size.*

on validation performance.

Self-supervised learning on auxiliary data. It is well-known that self-supervised learning is challenging for small models [24, 25], and we empirically demonstrate that SSL is also ineffective for lightweight keyword spotting models. Despite extensive data augmentation and hyperparameter tuning, MTL with SSL and KD achieve limited performance improvement. Moreover, PreT degrades performance compared to the baseline, which trains the model on randomly initialized parameters. In small-footprint keyword spotting models, learning invariant information with SSL-based pre-training hinders the creation of keyword representations. Among the SSL methods, PreT-Big shows significantly better performance than the baseline. However, PreT-Big utilizes large feature extractors of size 300M, while other models use BC-ResNet8 that size is around 300K.

Supervised learning on auxiliary data. Different from PreT with SSL, pre-training with SL boosts the performance of the baseline, which helps to learn keyword representations. We observe that the ALL method, which uses both datasets together for FS-KWS, is more beneficial for representation learning of small models than SSL, even in the presence of a domain gap between the two datasets. This naive approach benefits from well-organized word-like speech data for keyword representation learning. Our AuxSL adopts a different path for the auxiliary data through the MTL framework, which results in significant performance improvement as the model learns useful information for the target task from the auxiliary data.

4.4. Evaluation on Various Architecture and Model Size

To evaluate the effectiveness of our proposed method on various architectures and model sizes, we conducted experiments using three different backbone architectures: BC-ResNet, Res12, and DS-ResNet18. We applied the same training and evaluation settings as in the first 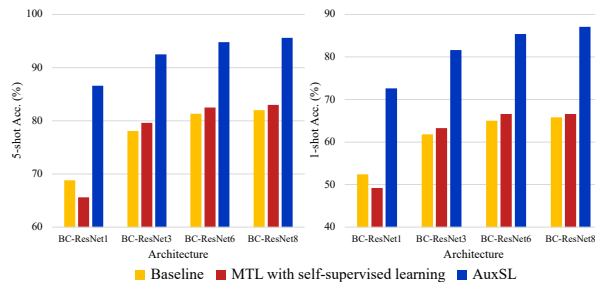sub-section, except for changing the back-bone architecture. As shown in Table 3, we apply the competitive methods to other backbones, Res12 and DS-ResNet18, but the performance improvements are marginal, similar to the results obtained with BC-ResNet8. Our proposed MTL approach outperforms other training strategies.

In Fig. 4, we observe that the performance of all methods, including the baseline, SSL, and SL, improves as the feature extractor size increases. However, SSL decreases performance for tiny models, such as BC-ResNet1 (9.2K), compared to the baseline. Remarkably, our proposed AuxSL approach applied to the smallest model (BC-ResNet1) outperforms other methods, including the largest model (BC-ResNet8). This result demonstrates the effectiveness of our method in learning discriminative representations even in extremely low-resource settings.

5. Discussion and Conclusion

In this paper, we addressed the challenge of few-shot keyword spotting (FS-KWS) by proposing a framework that leverages commonly available reading speech data as auxiliary data. Our approach has two main contributions: (1) the creation of a well-organized and balanced keyword dataset, LibriWord, and (2) AuxSL: multi-task learning (MTL) with an additional classifier to minimize the domain gap between in-domain and auxiliary data. Our results have shown that creating a keyword-balanced dataset is a practical approach for training lightweight keyword-spotting models. Moreover, we demonstrated the superiority of the proposed learning approach through extensive experiments, as evidenced by the improved performance in the FS-KWS benchmark. While our approach has achieved promising results, we recognize that there may be other techniques to mitigate domain differences between datasets, such as RFN [39] and DSBN [40]. We leave it as future work to explore and analyze the effectiveness of these methods in a scenario where there is a large domain gap between in-domain and auxiliary data.

6. References

[1] B. Kim, M. Lee, J. Lee, Y. Kim, and K. Hwang, "Query-by-example on-device keyword spotting," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2019.

[2] J. Huang, W. Gharbieh, H. S. Shim, and E. Kim, "Query-by-example keyword spotting system using multi-head attention and soft-triple loss," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021.

[3] Y. Chen, T. Ko, L. Shang, X. Chen, X. Jiang, and Q. Li, "An investigation of few-shot learning in spoken term classification," in *Proceedings of the INTERSPEECH*, 2020.

[4] A. Parnami and M. Lee, "Few-shot keyword spotting with prototypical networks," in *2022 7th International Conference on Machine Learning Technologies (ICMLT)*, 2022.

[5] W.-T. Kao, Y.-K. Wu, C.-P. Chen, Z.-S. Chen, Y.-P. Tsai, and H.-Y. Lee, "On the efficiency of integrating self-supervised learning and meta-learning for user-defined few-shot keyword spotting," in *2022 IEEE Spoken Language Technology Workshop (SLT)*, 2023.

[6] D. Lee, M. Kim, S. H. Mun, M. H. Han, and N. S. Kim, "Fully unsupervised training of few-shot keyword spotting," in *2022 IEEE Spoken Language Technology Workshop (SLT)*, 2023.

[7] H.-K. Shin, H. Han, D. Kim, S.-W. Chung, and H.-G. Kang, "Learning audio-text agreement for open-vocabulary keyword spotting," in *Proceedings of the INTERSPEECH*, 2022.

[8] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

[9] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *International Conference on Machine Learning (ICML)*, 2017.

[10] Q. Sun, Y. Liu, T.-S. Chua, and B. Schiele, "Meta-transfer learning for few-shot learning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

[11] A. Li, T. Luo, Z. Lu, T. Xiang, and L. Wang, "Large-scale few-shot learning: Knowledge transfer with class hierarchy," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

[12] P. Warden, "Speech commands: A dataset for limited-vocabulary speech recognition," *arXiv preprint arXiv:1804.03209*, 2018.

[13] A. Coucke, M. Chlieh, T. Gisselbrecht, D. Leroy, M. Poumeyrol, and T. Lavril, "Efficient keyword spotting using dilated convolutions and gating," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019.

[14] "Mandarin chinese scripted speech corpus," https://magichub.com/datasets/mandarin-chinese-scripted-speech-corpus-keyword-spotting-2/.

[15] B. Kim, S. Yang, I. Chung, and S. Chang, "Dummy prototypical networks for few-shot open-set keyword spotting," in *Proceedings of the INTERSPEECH*, 2022.

[16] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015.

[17] V. Pratap, Q. Xu, A. Sriram, G. Synnaeve, and R. Collobert, "Mls: A large-scale multilingual dataset for speech research," in *Proceedings of the INTERSPEECH*, 2020.

[18] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *International Conference on Machine Learning (ICML)*, 2020.

[19] X. Chen and K. He, "Exploring simple siamese representation learning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.

[20] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

[21] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," in *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2021.

[22] X. Zhai, A. Oliver, A. Kolesnikov, and L. Beyer, "S4l: Self-supervised semi-supervised learning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

[23] F. Cui, L. Guo, Q. Wang, P. Gao, and Y. Wang, "Exploring representation learning for small-footprint keyword spotting," in *Proceedings of the INTERSPEECH*, 2022.

[24] T. Chen, S. Kornblith, K. Swersky, M. Norouzi, and G. E. Hinton, "Big self-supervised models are strong semi-supervised learners," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

[25] Z. Fang, J. Wang, L. Wang, L. Zhang, Y. Yang, and Z. Liu, "Seed: Self-supervised distillation for visual representation," in *International Conference on Learning Representations (ICLR)*, 2021.

[26] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, "Montreal forced aligner: Trainable text-speech alignment using kaldi." in *Proceedings of the INTERSPEECH*, 2017.

[27] B. Kim, S. Chang, J. Lee, and D. Sung, "Broadcasted Residual Learning for Efficient Keyword Spotting," in *Proceedings of the INTERSPEECH*, 2021.

[28] M. Xu and X.-L. Zhang, "Depthwise separable convolutional resnet with squeeze-and-excitation blocks for small-footprint keyword spotting," in *Proceedings of the INTERSPEECH*, 2020.

[29] I. López-Espejo, Z.-H. Tan, and J. Jensen, "A novel loss function and training strategy for noise-robust keyword spotting," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2021.

[30] S. Yang, B. Kim, I. Chung, and S. Chang, "Personalized keyword spotting through multi-task learning," in *Proceedings of the INTERSPEECH*, 2022.

[31] Z. Hou, A. Walid, and S.-Y. Kung, "Meta-learning with attention for improved few-shot learning," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021.

[32] L. Lugosch, M. Ravanelli, P. Ignoto, V. S. Tomar, and Y. Bengio, "Speech model pre-training for end-to-end spoken language understanding," in *Proceedings of the INTERSPEECH*, 2019.

[33] L. Lugosch, S. Myer, and V. S. Tomar, "Donut: Ctc-based query-by-example keyword spotting," in *Advances in Neural Information Processing Systems (NeurIPS) workshop*, 2018.

[34] F. Zhou, B. Wu, and Z. Li, "Deep meta-learning: Learning to learn in the concept space," *arXiv preprint arXiv:1802.03596*, 2018.

[35] B. Oreshkin, P. Rodríguez López, and A. Lacoste, "Tadam: Task dependent adaptive metric for improved few-shot learning," *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

[36] B. Kang, S. Xie, M. Rohrbach, Z. Yan, A. Gordo, J. Feng, and Y. Kalantidis, "Decoupling representation and classifier for long-tailed recognition," in *International Conference on Learning Representations (ICLR)*, 2020.

[37] J. Ren, C. Yu, X. Ma, H. Zhao, S. Yi *et al.*, "Balanced meta-softmax for long-tailed visual recognition," *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

[38] R. Adriana, B. Nicolas, K. S. Ebrahimi, C. Antoine, G. Carlo, and B. Yoshua, "Fitnets: Hints for thin deep nets," in *International Conference on Learning Representations (ICLR)*, 2015.

[39] B. Kim, S. Yang, J. Kim, H. Park, J. Lee, and S. Chang, "Domain generalization with relaxed instance frequency-wise normalization for multi-device acoustic scene classification," in *Proceedings of the INTERSPEECH*, 2022.

[40] W.-G. Chang, T. You, S. Seo, S. Kwak, and B. Han, "Domain-specific batch normalization for unsupervised domain adaptation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.