



Few-Shot Open-Set Learning for On-Device Customization of KeyWord Spotting Systems

Manuele Rusci¹, Tinne Tuytelaars¹

¹PSI, KU Leuven, Belgium

{manuele.rusci,tinne.tuytelaars}@esat.kuleuven.be

Abstract

A personalized KeyWord Spotting (KWS) pipeline typically requires the training of a Deep Learning model on a large set of user-defined speech utterances, preventing fast customization directly applied on-device. To fill this gap, this paper investigates few-shot learning methods for open-set KWS classification by combining a deep feature encoder with a prototype-based classifier. With user-defined keywords from 10 classes of the Google Speech Command dataset, our study reports an accuracy of up to 76% in a 10-shot scenario while the false acceptance rate of unknown data is kept to 5%. In the analyzed settings, the usage of the triplet loss to train an encoder with normalized output features performs better than the prototypical networks jointly trained with a generator of dummy unknown-class prototypes. This design is also more effective than encoders trained on a classification problem and features fewer parameters than other iso-accuracy approaches.

Index Terms: Keyword Spotting Systems, Few-shot Learning, Open-Set Classification, On-Device Customization

1. Introduction

Keyword Spotting, which is the ability to recognize speech commands or *wake-words*, is getting popular among battery-powered smart audio sensors [1]. Because a KWS detection pipeline is intended to run continuously on the target system, the reduction of computation and memory costs of Deep Learning based KWS algorithms have been extensively investigated over the last years [2, 3]. The design of a custom KWS algorithm typically demands the training of a model on a dataset of collected *user-defined keywords* [4]. Despite its effectiveness, such a design process is subject to the availability of computing resources and speech recordings, preventing users from obtaining a custom solution in a short time, e.g., on-device.

Few-Shot Learning (FSL) provides a viable solution to deal with the scarcity of abundant user-defined keywords data. In the field of KWS, there are several recent FSL methods that rely on the Prototypical Network (ProtoNet) concept [5]. During system setup in the target scenario, users are asked to provide a few enrollment samples for each keyword. These reference speech data are then processed through a trained feature encoder to produce a set of feature vectors. A *class prototype* is then computed as the mean of the feature vectors for each user-defined keyword. When it comes to inference, the distances between the output feature vector of a test sample, or *embedding*, and the class prototypes are calculated. The classification output is determined by the shortest distance. To gain high accuracy, the ProtoNet's feature encoder is trained to cluster the embeddings of speech samples belonging to the same class. At the same time, feature vectors of different classes are forced to be distant

according to a given distance metric, e.g. Euclidean.

With respect to FSL techniques involving fine-tuning on a few labeled data [6], the fit of a prototype-based classifier is a low-cost option that can be easily implemented on-device without the need for backpropagation. At present, however, the ProtoNet-based FSL approaches have been primarily assessed for closed-set classification, i.e. the test categories match those of the training set [7, 8, 9]. In contrast, we argue that a customized KWS method shall work in an open-set setting, to distinguish user-defined keywords from *unknown* speech utterances. To this aim, the ProtoNet approach has been recently extended by jointly training a generator of dummy prototypes for the unknown class [10]. Unfortunately, this work makes use of training data sampled from the same distribution of the target data, i.e. different class subsets of the Google Speech Command dataset. In addition, other works showed angular variants of the prototypical loss to achieve the highest accuracy for closed-set FSL classification [9] or efficiently learning the feature encoder using the triplet loss but not in a few-shot setting [11]. Hence, we denoted the present literature on few-shot open-set learning KWS solutions to be highly fragmented and it is missing clear design guidelines for on-device KWS customization.

To bridge this gap, this paper contributes an evaluation framework for FSL architectures composed by a feature encoder and a prototype-based open-set classifier initialized with few-shot samples. More in detail, we leverage the recent Multilingual Spoken Words Corpus (MSWC) dataset [12] to train a feature extractor using the prototypical loss, its angular variant or the triplet loss. The evaluation is performed on the Google Speech Command (GSC) dataset, which is partitioned between a collection of target keywords, i.e. the positive set, and a negative set of *unknown* keywords. In our analysis we compare the open-set classifier featuring a *dummy proto generator* with either a simple variant that computes the unknown-class prototype using few random words or an alternative based on OpenMAX [13], which statistically models the distance of data samples from the class prototypes to estimate if a test sample can fit any of the known classes. When considering Depthwise-Separable Convolutional Neural Network (DSCNN) encoders tailored for low-power embedded systems [4], we show that a training process using the triplet loss and normalized features brings a superior accuracy than a ProtoNet-based method for open-set FSL classification under a fixed training epoch budget. Our code is available at: <https://github.com/mrusci/ondevice-fewshot-kws>.

2. Related Work

A set of works denoted as "Query-by-Example" investigated methods to personalize the wake-word of a KWS system after

providing a reference utterance. [14] trained a 2-layer LSTM and proposed to use a similarity score to compare the last embedding state with a reference vector. An LSTM encoder was also leveraged by [10] to output a hypothesis graph based on the phonetic-based posteriorgram for the comparison. This method is however expensive at inference time because the memory-bound LSTM workload occurs at any window step and may lead to a runtime 10-100× slower than our considered convolutional method applied on a large time window (1 s).

Differently from [15] that proposed to use MAML to train an effective representation for FSL on KWS, authors of [6] leveraged transfer learning on a large encoder model (EfficientNet) to train a classifier on the target scenario. Conversely, many recent works rely on the Prototypical Network approach [5], which does not require any training to encapsulate the information of the few-shot samples. Several studies focused only on closed-set problems and used the data from the same distribution (typically the GSC dataset) at training and test time [7, 8]. Jung *et al.* [9] showed the higher effectiveness of the angular prototypical loss on a KWS FSL scenario compared to classification-based encoders. After training a Resnet15 using data from 1000 classes of Librispeech, the authors proposed to fine-tune the feature extractor on the categories of the test dataset to gain an accuracy of up to 96% for a 10-shot closed-set problem. A personalized KWS system [16] also leveraged the angular loss to train an additional model for keyword adaptation. In parallel, [11] proposed the triplet loss to train a ResNet15-based feature extractor on LibriSpeech, which showed a top score of 97% on GSC using a kNN with all the training data, i.e. not FSL. After, [17, 18] exploited a soft-triplet loss, which combines triplet loss and softmax, for a Query-by-Example custom wake-word solution. This class of works focused on FSL training methodologies but lacks clear design guidelines and evaluation in an open-set scenario for customized KWS solutions that we address.

In the open-set context, [19] used triplet or angular prototypical losses for KWS detection combined with a memory-expensive SVM classifier. However, this study do not consider few-shot evaluations and includes the user-defined keywords in the training dataset. More related to our study is the work of Kim *et al.* [10] that we also reproduce in this paper. The authors used a ProtoNet setting and jointly trained a generator of dummy prototypes for the unknown class. Unfortunately, this study is limited by the usage of GSC for training and testing, and the evaluation is restricted to a 5-shot-5-way, in which this method scores up to 86.9% in accuracy. PEELER [20] proposed to include an extra model to estimate the variance associated with every class prototype and rely on the Mahalanobis distance at inference time for open-set classification. This method is however memory-expensive because of the overhead to store the spatial variance tensors and the estimator.

3. Method

Figure 1 depicts the design flow considered for On-Device KWS customization. The core algorithm is composed by a Deep Learning based feature encoder and an open-set classifier. Firstly, the feature encoder is trained offline (Fig. 1A), i.e. using a server machine before deployment on the target HW, on a large source dataset of labelled spoken keywords. At test time, an open-set classifier is plugged on top of the feature extractor and initialized with few-shot utterances taken from the target scenario (Fig. 1B). We consider the target dataset to be disjointed from the training data; test utterances belong to classes

not represented in the source dataset. Train and test data are also sampled from different distributions. As shown in Fig. 1C, the inference is based on a distance score accounting for the unknown class as explained in Sec. 3.2.

3.1. Training the Feature Encoder

The feature encoder $f(\cdot)$ is trained in a supervised fashion following an episode-based protocol. At every episode, the dataloader feeds the model with a batch of training data fetched from the source dataset. Every batch features a balanced number of samples from M randomly chosen classes. In the following, we detail the considered training strategies.

Prototypical Network (PN). Let us consider an episodic batch that includes S support samples $\{x_{i,j}^S\}_{i=1}^S$ and Q query samples $\{x_{i,j}^Q\}_{i=1}^Q$ for every class $j = 1, \dots, M$, for a total of $(S + Q) \times M$ utterances per batch. At every episode, a set of prototypes $c = \{c_j\}_{j=1}^M$ is firstly computed as:

$$c_j = \frac{1}{S} \sum_{i=1}^S f(x_{i,j}^S) \quad (1)$$

Then, the loss function minimizes the negative log probability of correctly predicting the category of the query samples:

$$L_{PN} = -\frac{1}{Q \cdot M} \sum_{i=1}^Q \sum_{j=1}^M \log \frac{\exp(\mathbf{s}(x_{i,j}^Q, j))}{\sum_{k=1}^M \exp(\mathbf{s}(x_{i,j}^Q, k))} \quad (2)$$

where

$$\mathbf{s}(x, j) = -d_{L2}(f(x), c_j) \quad (3)$$

and d_{L2} is the euclidean distance.

Angular Prototypical (AP). Following [21, 9], we replace the euclidean distance in Eq. 3 with the cosine similarity within the same formulation of the Prototypical Network (Eq. 2) as:

$$\mathbf{s}(x, j) = w \cdot (\cos(f(x), c_j) - m) + b \quad (4)$$

where w and b are learnable scalar parameters, c_j is computed from Eq. 1 and m is a margin different from zero for the score of the true class.

Triplet Loss (TL). The triplet loss is computed based on N_t triplets $\{x_i, x_i^+, x_i^-\}_{i=1}^{N_t}$ sampled from an episodic batch:

$$L_{TL} = \frac{1}{N_t} \sum_{i=1}^{N_t} \max(0, d_{L2}(x_i, x_i^+) - d_{L2}(x_i, x_i^-) + m) \quad (5)$$

where x_i^+ belongs to the same category of x_i and x_i^- is selected from a different random class. m is the margin.

3.1.1. Feature Encoder Architecture

Without loss of generality, this study considers the DSCNN model as the feature encoder, which is largely used for on-device KWS [4]. This deep model is composed of a stacked sequence of depthwise and pointwise convolution blocks. Every block includes a BatchNorm and a ReLU activation after the convolution. We experiment by passing to the final average pooling layer the feature maps produced by the last convolution or the ReLU layers. These feature extractor versions are denoted respectively as *DSCNN-CONV* or *DSCNN-RELU* (see Fig. 1A). Note that the latter includes only non-negative values in the high-dimensional space. In addition, we also consider the case of normalized output features [11]. Following [22], the *DSCNN-NORM* replaces the last BatchNorm layer with a LayerNorm and applies L2 normalization after the average pooling. In this last setting, the feature encoder is forced to discriminate classes based on the phase of the embeddings.




Figure 1: Overview of the approach considered for KWS open-set FSL.

3.2. Open-Set Classification

At inference time, the KWS pipeline receives K enrollment samples for every *user-defined keyword*. We consider N new keywords in the target domain and, therefore, a K -shot N -way classification problem. After computing the embeddings of the K -shot samples using the trained encoder, a classifier computes (Eq. 1) and stores the class prototypes.

When a new test sample is fed to the pipeline, the open-set classifier returns a probability score vector $P = \{p_i\}_{i=0}^N$ based on the current prototype set. Specifically, p_0 is the prediction score for the *unknown class* and $p_{i \neq 0}$ is the probability of the i -th keyword, which assumes the highest value if the distance score of Eq. 3 is the lowest. The final class prediction y is:

$$y = \begin{cases} \arg \max p_i, & \text{if } p_i \geq \gamma \\ 0, & \text{otherwise} \end{cases} \quad (6)$$

where $y = 0$ denotes the unknown class and γ is a manually-tunable parameter to tradeoff between the classifier’s precision and recall. In the following, we detail the considered variants for the open-set classifiers.

Open Nearest Class Mean (openNCM) [23]. The Nearest Class Mean classifier is typically adopted by Prototypical Networks. A simple open-set variant estimates the c_0 prototype for the *unknown class* using K random samples taken from the target domain but not belonging to the user classes. The probability score is computed by applying the SoftMax on the $(N+1)$ -sized distance vector obtained from Eq. 3. If the feature encoder is trained using the AP loss, the classifier applies L2 normalization before computing the euclidean distance.

OpenMAX [13]. Following [24], for every known class the distance from the prototype vector is statistically modeled using a Weibull distribution (we feed the 5 largest distances from the K enrollment samples to the `fit_high` function). At test time, the distances of a test sample from the known classes (Eq. 3) are scaled by a factor $w_i^{wb} \in [0, 1]$, which tends to 1 if the current sample is on the tail of the i -th Weibull class distribution. The score for the unknown category is then estimated with Eq. 7 before applying the Softmax function.

$$s(x, j = 0) = - \sum_{i=1}^N (1 - w_i^{wb}) \cdot d_{L2}(f(x), c_i) \quad (7)$$

Dummy Proto (DProto) [10]. A generator $g(\cdot)$ of dummy prototypes for the unknown category is jointly trained with a ProteNet-based feature encoder $f(\cdot)$. Thanks to this, the unknown prototype is estimated at inference time as $c_0 = g(\{c_j\}_{j=1}^N)$. Driven by the results of the original paper [10], $g(\cdot)$ is trained to produce 3 unknown prototypes¹; only the clos-

¹We rely on our implementation because not any code is available.

est to the test samples is selected at inference time to compute the final probability score, in a similar way to openNCM. At training time, data samples from a subset of categories in the episodic batch are assigned to the unknown class to jointly learn the feature extractor and the generator g . We remind the reader to the original paper for more details.

3.3. Evaluation Framework

We use the *English* partition of the Multilingual Spoken Words Corpus (MSWC) dataset [12] to train the feature encoder. We restrict the data sampling to the 500 classes with the highest number of utterances after excluding the categories of the test dataset (the GSC+ partition described below). The train samples are augmented with additive background noises taken from the DEMAND dataset [25]. Noise is applied according to a uniform probability of 0.95 and a random SNR between 0 and 5.

For the open-set testing, we refer to the Google Speech Command (GSC) dataset [26]. The GSC dataset is composed of 1 s long speech utterances from 35 categories. For our experiments, we define a positive (GSC+) and a negative (GSC-) partition. The GSC+ includes only samples belonging to 10 classes (*yes, no, up, down, left, right, on, off, stop, go*), which constitute the target categories for few-shot classification. The K enrollments to setup the classifier within a K -shot N -way test are sampled from the original *train* split of the GSC dataset. In the case of the openNCM, the prototype of the *unknown class* is computed based on data randomly fetched from 5 classes: *backward, forward, visual, follow, learn*. The GSC- partition includes the test samples from the remaining 20 classes.

To assess the performance of every tested configuration we report: (i) the classification accuracy on the GSC+ partition when γ is tuned to achieve a False Acceptance Rate (FAR)² of 5% on the GSC- dataset ($ACC_{5\%}^+$), (ii) the False Rejection Rate at FAR of 5% ($FR_{5\%}^+$), and (iii) the *AUROC*. The classification metrics are measured over the data samples fetched from the *test* split of GSC and we provide the mean statistics over 10 test repetitions.

4. Experimental Result

In this study we consider the Large and Small versions of the DSCNN, respectively indicated as DSCNN-L and DSCNN-S and featuring 407k and 22k parameters [4]. The DSCNN model is fed with a MFCC feature map computed with a window size of 40 ms and a stride of 50%. We rely on power spectrograms to avoid the costly square root operations and select the first 10 MFCC features to obtain 49×10 feature maps, as done in [4]. The size of the embedding vectors is 64 and 256 for DSCNN-S and DSCNN-L, respectively.

²Ratio of negative samples classified as positives

Table 1: Performance on the GSC testset under $\{5, 10\}$ -shot 10-way open-set classification settings.

Loss	Feature Extractor	openNCM 5-shot			OpenMAX 5-shot			Dproto 5-shot			openNCM 10-shot			OpenMAX 10-shot			Dproto 10-shot		
		$ACC_{5\%}^+$	AUROC	$FRR_{5\%}^+$	$ACC_{5\%}^+$	AUROC	$FRR_{5\%}^+$	$ACC_{5\%}^+$	AUROC	$FRR_{5\%}^+$	$ACC_{5\%}^+$	AUROC	$FRR_{5\%}^+$	$ACC_{5\%}^+$	AUROC	$FRR_{5\%}^+$	$ACC_{5\%}^+$	AUROC	$FRR_{5\%}^+$
PN	DSCNN-L-NORM	0.21	0.66	0.78	0.23	0.79	0.77	0.21	0.64	0.79	0.22	0.68	0.78	0.23	0.75	0.77	0.22	0.67	0.78
	DSCNN-L-CONV	0.54	0.86	0.46	0.12	0.87	0.87	0.64	0.91	0.35	0.62	0.89	0.37	0.48	0.89	0.50	0.71	0.93	0.28
	DSCNN-L-RELU	0.56	0.87	0.43	0.14	0.91	0.85	0.66	0.92	0.32	0.63	0.89	0.37	0.56	0.92	0.40	0.71	0.93	0.28
AP	DSCNN-L-NORM	0.66	0.92	0.29	0.44	0.94	0.54	0.65	0.93	0.30	0.71	0.93	0.25	0.66	0.93	0.30	0.70	0.94	0.25
	DSCNN-L-NORM	0.71	0.93	0.26	0.37	0.94	0.62				0.76	0.94	0.21	0.71	0.94	0.24			
	DSCNN-L-CONV	0.58	0.88	0.41	0.25	0.95	0.74				0.63	0.89	0.36	0.67	0.95	0.29			
TL	DSCNN-L-RELU	0.66	0.90	0.33	0.20	0.96	0.80				0.71	0.91	0.28	0.64	0.96	0.32			
	DSCNN-S-NORM	0.14	0.57	0.85	0.14	0.72	0.85	0.14	0.54	0.85	0.17	0.59	0.83	0.17	0.69	0.83	0.15	0.55	0.84
	DSCNN-S-CONV	0.40	0.81	0.60	0.14	0.83	0.84	0.40	0.81	0.59	0.48	0.85	0.51	0.38	0.86	0.60	0.43	0.83	0.57
AP	DSCNN-S-RELU	0.39	0.80	0.60	0.20	0.86	0.77	0.39	0.80	0.60	0.45	0.84	0.54	0.44	0.87	0.54	0.44	0.81	0.56
	DSCNN-S-NORM	0.39	0.83	0.60	0.34	0.87	0.64	0.31	0.81	0.68	0.41	0.84	0.57	0.36	0.86	0.63	0.33	0.82	0.66
	DSCNN-S-NORM	0.51	0.87	0.46	0.38	0.91	0.59				0.56	0.89	0.42	0.54	0.91	0.42			
TL	DSCNN-S-CONV	0.39	0.80	0.60	0.26	0.92	0.70				0.42	0.82	0.57	0.56	0.92	0.39			
	DSCNN-S-RELU	0.42	0.82	0.57	0.28	0.92	0.69				0.49	0.85	0.50	0.58	0.93	0.37			

Feature extractors are trained for a fixed duration of 40 epochs of 400 episodes. Adam optimizer is adopted with an initial learning rate of 0.001 and decayed by 0.1 after 20 epochs. The Triplet Loss is computed over episodic batches of 20 samples \times 80 classes; triplet negatives are randomly selected among the batch data of a different category. A margin m of 0.5 is used. In the case of PN or AP, the dataloader samples 10 support samples and 30 query samples from 40 classes at every episode. For AP, m is set to 0.5. To train the DProto feature extractor jointly with the dummy prototype generator, 16 of the 40 classes are marked as unknown.

Table 1 includes the results of the experiments, which are averaged over 3 runs to account uncertainty from different training seeds. Firstly, we notice the PN loss to achieve the best results when a large DSCNN-RELU is used. This configuration is slightly better than DSCNN-CONV while the NORM counterpart is leading to a low score, denoting a low-class separability. A $<2\%$ accuracy gap between the RELU and CONV configurations is also observed for the small DSCNN, with the latter scoring the best. More in detail, openNCM achieves the highest accuracies of 56% and 63% for, respectively, 5- and 10-shot 10-way KWS classification. In the 5-shot scenario, OpenMAX leads to a low score, meaning a correct statistical model is not inferred from few samples, as also demonstrated from all other analyzed cases. Conversely, the accuracy of OpenMAX is close to openNCM with 10 samples per-class. These observations are preserved when the model is trained using the DProto approach, achieving up to +8% accuracy and -9% FRR compared to openNCM for DSCNN-L in 10-shot configuration.

Overall, the triplet loss achieves the highest accuracies of 71% and 76% with openNCM and DSCNN-L-NORM for 5-shot and 10-shot, respectively. Differently from PN, the NORM feature extractor leads to superior performance than CONV and RELU in the case of DSCNN-L with a 4-5% accuracy improvement. We hypothesize the NORM feature extractor to foster

the class separation process based on the phase of the embeddings. The performance difference between NORM and RELU is reduced on the small model up to inverting the trend when a 10-shot OpenMAX classifier is plugged on the DSCNN-S-RELU (+4%). On the other side, the AP loss achieves a slightly lower accuracy than the TL but still gains 8-10% accuracy with respect to the PN loss on DSCNN-L, either with openNCM or DProto classifiers. In the case of DSCNN-S, AP leads to a similar accuracy level of the PN losses.

Lastly, Table 2 shows the performance of the best configurations of TL+openNCM and Dproto on a 10-shot open-set problem along with other methods. More in detail, we consider a DSCNN feature extractor trained on a classification problem following [22] and PEELER [20]. Both methods are trained on MSWC and tested on GSC, according to our evaluation framework. The usage of the TL leads to the same accuracy level as PEELER on DSCNN-L but without paying the overhead of the additional network that estimates the class variances (+6.3M parameters), while PEELER scores +5% accuracy on DSCNN-S. On the other side, the classifier-based feature extractor performs poorly in open-set because it is overconfident in the case of wrong predictions. For comparison purposes, we also trained the DSCNN models end-to-end on GSC+ with an *unknown* class composed by the 5 keywords listed in Sec. 3.3. Also in this case, we observe that a TL+openNCM achieves a similar $ACC_{5\%}^+$ in open-set. Despite the high accuracy on known classes, i.e. up to 95% if $\gamma = 0$, the end-to-end classifier fails to predict unknown classes not seen during training, i.e. low accuracy on the negative set. Conversely, the lower model capacity of DSCNN-S brings the end-to-end approach to perform best if trained on the full target dataset.

5. Conclusion

This paper investigated few-shot open-set methods to again on-device KWS customization. We built a framework to evaluate multiple open-set classifiers placed on top of feature extraction modules and initialized with few enrollments. Among the analyzed feature extractor variants trained on the MSWC dataset, we showed the triplet loss applied on normalized output features of a DSCNN-L model leads to the highest accuracy when coupled with an openNCM classifier, surpassing solutions based on prototypical networks. In the considered scenario, this solution is also comparable with the more complex PEELER method or an end-to-end model trained on the target dataset.

6. Acknowledgements

This work is partly supported by the European Horizon Europe program under grant agreement 101067475.

Table 2: 10-shot 10-way open-set Classification Comparison.

DSCNN-L — params: 407k	$ACC_{5\%}^+$	AUROC	Train Data	Extra Params
openNCM+Classif [22]+ NORM	0.52	0.89	source	-
openNCM+TL+NORM	0.76	0.94	source	-
dProto [10]+RELU	0.71	0.93	source	-
PEELER [20]	0.76	0.94	source	+6.3M
end-to-end [4]	0.76	0.93	target	-
DSCNN-S — params: 22k	$ACC_{5\%}^+$	AUROC	Train Data	Extra Params
openNCM+Classif [22]+ NORM	0.47	0.85	source	-
openNCM+TL+NORM	0.56	0.89	source	-
dProto [10]+RELU	0.44	0.82	source	-
PEELER [20]	0.60	0.88	source	+341k
end-to-end [4]	0.72	0.93	target	-

7. References

- [1] I. López-Espejo, Z.-H. Tan, J. H. Hansen, and J. Jensen, “Deep spoken keyword spotting: An overview,” *IEEE Access*, vol. 10, pp. 4169–4199, 2021.
- [2] O. Rybakov, N. Kononenko, N. Subrahmanya, M. Visontai, and S. Laurenzo, “Streaming keyword spotting on mobile devices,” *Proc. Interspeech*, pp. 2277–2281, 2020.
- [3] C. Banbury, C. Zhou, I. Fedorov, R. Matas, U. Thakker, D. Gope, V. Janapa Reddi, M. Mattina, and P. Whatmough, “Micronets: Neural network architectures for deploying tinymt applications on commodity microcontrollers,” *Proceedings of Machine Learning and Systems*, vol. 3, pp. 517–532, 2021.
- [4] Y. Zhang, N. Suda, L. Lai, and V. Chandra, “Hello edge: Keyword spotting on microcontrollers,” *arXiv preprint arXiv:1711.07128*, 2017.
- [5] J. Snell, K. Swersky, and R. Zemel, “Prototypical networks for few-shot learning,” *Advances in neural information processing systems*, vol. 30, 2017.
- [6] M. Mazumder, C. Banbury, J. Meyer, P. Warden, and V. J. Reddi, “Few-Shot Keyword Spotting in Any Language,” in *Proc. Interspeech*, 2021, pp. 4214–4218.
- [7] Y. Chen, T. Ko, and J. Wang, “A Meta-Learning Approach for User-Defined Spoken Term Classification with Varying Classes and Examples,” in *Proc. Interspeech*, 2021, pp. 4224–4228.
- [8] A. Parnami and M. Lee, “Few-shot keyword spotting with prototypical networks,” in *2022 7th International Conference on Machine Learning Technologies (ICMLT)*. New York, NY, USA: Association for Computing Machinery, 2022, p. 277–283. [Online]. Available: <https://doi.org/10.1145/3529399.3529443>
- [9] J. Jung, Y. Kim, J. Park, Y. Lim, B.-Y. Kim, Y. Jang, and J. S. Chung, “Metric learning for user-defined keyword spotting,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [10] B. Kim, S. Yang, I. Chung, and S. Chang, “Dummy Prototypical Networks for Few-Shot Open-Set Keyword Spotting,” in *Proc. Interspeech*, 2022, pp. 4621–4625.
- [11] R. Vygon and N. Mikheylovskiy, “Learning efficient representations for keyword spotting with triplet loss,” in *Speech and Computer: 23rd International Conference, SPECOM 2021, St. Petersburg, Russia, September 27–30, 2021, Proceedings 23*. Springer, 2021, pp. 773–785.
- [12] M. Mazumder, S. Chitlangia, C. Banbury, Y. Kang, J. M. Ciro, K. Achorn, D. Galvez, M. Sabini, P. Mattson, D. Kanter *et al.*, “Multilingual spoken words corpus,” in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- [13] A. Bendale and T. E. Boulton, “Towards open set deep networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 1563–1572.
- [14] G. Chen, C. Parada, and T. N. Sainath, “Query-by-example keyword spotting using long short-term memory networks,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 5236–5240.
- [15] Y. Chen, T. Ko, L. Shang, X. Chen, X. Jiang, and Q. Li, “An Investigation of Few-Shot Learning in Spoken Term Classification,” in *Proc. Interspeech*, 2020, pp. 2582–2586.
- [16] S. Yang, B. Kim, I. Chung, and S. Chang, “Personalized Keyword Spotting through Multi-task Learning,” in *Proc. Interspeech*, 2022, pp. 1881–1885.
- [17] 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 *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 6858–6862.
- [18] J. Huang, W. Gharbieh, Q. Wan, H. S. Shim, and H. C. Lee, “QbyE-MLPMixer: Query-by-Example Open-Vocabulary Keyword Spotting using MLPMixer,” in *Proc. Interspeech*, 2022, pp. 5200–5204.
- [19] J. Huh, M. Lee, H. Heo, S. Mun, and J. S. Chung, “Metric learning for keyword spotting,” in *2021 IEEE Spoken Language Technology Workshop (SLT)*, 2021, pp. 133–140.
- [20] B. Liu, H. Kang, H. Li, G. Hua, and N. Vasconcelos, “Few-shot open-set recognition using meta-learning,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 8798–8807.
- [21] J. S. Chung, J. Huh, S. Mun, M. Lee, H.-S. Heo, S. Choe, C. Ham, S. Jung, B.-J. Lee, and I. Han, “In Defence of Metric Learning for Speaker Recognition,” in *Proc. Interspeech*, 2020, pp. 2977–2981.
- [22] A. Zhai and H.-Y. Wu, “Classification is a strong baseline for deep metric learning,” *arXiv preprint arXiv:1811.12649*, 2018.
- [23] T. L. Hayes and C. Kanan, “Online continual learning for embedded devices,” in *Conference on Lifelong Learning Agents*. PMLR, 2022, pp. 744–766.
- [24] M. Mundt, Y. Hong, I. Pliushch, and V. Ramesh, “A wholistic view of continual learning with deep neural networks: Forgotten lessons and the bridge to active and open world learning,” *Neural Networks*, 2023.
- [25] J. Thiemann, N. Ito, and E. Vincent, “The diverse environments multi-channel acoustic noise database (demand): A database of multichannel environmental noise recordings,” in *Proceedings of Meetings on Acoustics ICA2013*, vol. 19, no. 1. Acoustical Society of America, 2013, p. 035081.
- [26] P. Warden, “Speech commands: A dataset for limited-vocabulary speech recognition,” *arXiv preprint arXiv:1804.03209*, 2018.