



Representation based meta-learning for few-shot spoken intent recognition

Ashish Mittal, Samarth Bharadwaj, Shreya Khare, Saneem Chemmengath,
Karthik Sankaranarayanan, Brian Kingsbury

IBM Research AI

{arakeshk, samarhb, skhare34, saneem.cg, kartsank}@in.ibm.com, bedk@us.ibm.com

Abstract

Spoken intent detection has become a popular approach to interface with various smart devices with ease. However, such systems are limited to the preset list of *intents-terms* or *commands*, which restricts the quick customization of personal devices to new intents. This paper presents a few-shot spoken intent classification approach with task-agnostic representations via meta-learning paradigm. Specifically, we leverage the popular representation based meta-learning learning to build a task-agnostic representation of utterances, that then use a linear classifier for prediction. We evaluate three such approaches on our novel experimental protocol developed on two popular spoken intent classification datasets: Google Commands and the Fluent Speech Commands dataset. For a 5-shot (1-shot) classification of novel classes, the proposed framework provides an average classification accuracy of 88.6% (76.3%) on the Google Commands dataset, and 78.5% (64.2%) on the Fluent Speech Commands dataset. The performance is comparable to traditionally supervised classification models with abundant training samples.

Index Terms: speech recognition, meta learning, intent classification

1. Introduction

Recent advances in AI powered consumer devices has brought forth huge advances in direct speech to intent or command detection. Unlike traditional speech problems, such as speech to text, which are designed to weigh all words equally, the correct prediction of intent requires identifying the essential tokens or phrases correctly. Spoken command detection is a challenging problem since the utterance may be unclear, intent may be implicit, and the response time must be low. In such a scenario, advances made in natural language understanding can be rarely leveraged and the intent or command must be identified directly from speech.

Kao *et al*[1] showcase the low memory footprint and computational time of utilizing convolution network for the task of keyword spotting, reporting over 97% accuracy on the Speech Commands dataset [2]. Recently, Poncelet *et al*[3] leverage the advances in capsules network architecture for multi-task prediction of speaker and label of command, showing the robustness of the model. Further, some approaches leverage additional text data to improve performance of speech to intent classification in low resource settings [4].

Similarly, several other end-to-end learning methods have been proposed for spoken language understanding (SLU) task [5, 6, 7, 8, 9, 10, 4, 11, 12, 13, 14] without requiring intermediate text and show promising results on multiple tasks. However, their success heavily depends on a large amount of labeled training data. Such data is usually scarce and thus limiting the performance. Furthermore, such approaches limit personal de-

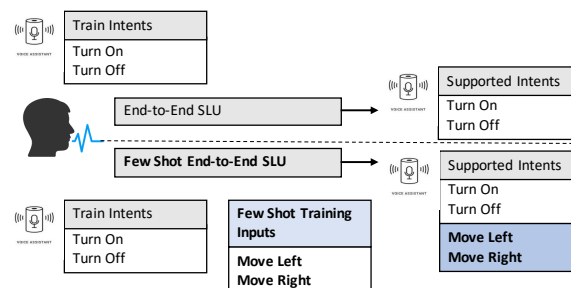


Figure 1: An illustration of the proposed few-shot approach to End-to-End SLU. A user adds new commands ‘move’ to the voice assistant.

vices to a specific list of commands or intents, restricting customization. This is illustrated in Figure 1. This paper brings the meta-learning paradigm to few-shot speech to intent classification task.

We show on two popular spoken intent detection datasets: the Google Speech Commands [2] and Fluent Speech Commands [13] datasets, that 1-shot and 5-shot classification to predict spoken commands and intents can be performed with surprisingly high accuracy. Specifically, we show that a representation based meta-learning paradigm that first learns task-agnostic representation of input utterance followed by task-specific base learners are ideal for spoken-command feature that is popular in several consumer devices. Representation based meta-learning (R-ML) simultaneously leverage the representational power of neural networks, with the low-resource requirements of linear classifiers. Further, the models are trained in an end-to-end fashion thereby passing gradients computed from classification error through the linear classifier and the representational neural network simultaneously. The end-to-end nature of these architectures makes them powerful *generalizers*, in which the small foot-print linear classifier can be easily retrained with the few-shot examples available of novel classes.

To ensure that such a classifier ensemble does not overfit to few-shot available per class, we utilize a speech representation known as PASE+ [15, 16], that is trained with only self-supervised tasks, i.e., classification tasks generated from the data itself, such as reconstruction of static-engineered features. We hypothesize that the additional strong regularization of self-supervision ensures the representation remains task-agnostic. To the best of our knowledge, this is the first study to evaluate three popular representation based meta-learning approaches for speech classification tasks as well as self-supervised regularization of meta-learning.

Similar to our proposed approach, [17] present a meta-learning formulation of intent classification problem on the Google commands dataset. Different from R-ML, the per-




Figure 2: The representation based few shot encoding learns a task agnostic feature representation. Simultaneously, the accompanying linear classifier (nearest-neighbour, SVM) learns tasks-specific class discrimination boundaries. Therefore, the proposed approach can adapt to a new task with few samples (few-shot) by updating just the linear classifier.

formance of a popular optimization based approach to meta-learning, known as model-agnostic-meta-learning (MAML) [18] is evaluated on small subsets of the dataset, though MAML has shown to be difficult to train in practice [19]. On the other hand, Koluguri *et al*[20] show that meta-learning with ProtoNet [21] is effective for child-adult speaker diarization by modeling each speaker session as novel task. Neither works perform a thorough evaluation of multiple meta-learning algorithms on standard protocols of command or intent classification problems.

2. Proposed approach

As shown in Figure 2, the proposed representation based few-shot learning paradigm is trained to be able to produce a task-agnostic representation of an utterance using the popular meta-learning paradigm. In this work, we define a *task* or *episode* as a n -way, m -shot classification problem, i.e., n classes are predicted with only (small) m audio samples per class available for training. In this section, we describe the meta-learning pipeline for few-shot learning.

2.1. Meta-learning

Task independent learning, or learning to learn is a popular paradigm of machine learning termed meta-learning [22, 18, 23, 24, 21, 25]. We consider three popular approaches from the *robust-representation* view of the meta-learning paradigm, termed ProtoNet [21], Ridge [24] and MetaOptNet [23]. The representation view achieves the goal of few-shot learning by decoupling the learning process into: task-independent encoding function \mathcal{F}_θ parameterized by θ , and task-specific classifier \mathcal{A}_ϕ parameterized by ϕ . In general, \mathcal{F}_θ is a complex neural network model and \mathcal{A}_ϕ a simple linear classifier that project the embedding from \mathcal{F}_θ to space of class labels. The optimizer minimizes the expectation of *meta-loss* \mathcal{L}^{meta} (negative log-likelihood) over T tasks to obtain task-independent optimal parameters θ (as given by Equation 1).

$$\min_{\theta} \frac{1}{T} \sum_{t=1}^T \mathcal{L}^{meta}(\mathcal{D}_q^t; \mathcal{A}_\phi, \mathcal{F}_\theta) + \mathcal{R}(\theta) \quad (1)$$

where $\phi = \arg \min_{\phi} \ell(\mathcal{D}_s^t; \mathcal{A}_\phi, \mathcal{F}_\theta)$

Algorithm 1: Meta-training routine for single batch of n -way, m -shot classification

```

Input:  $\mathcal{F}()$ ,  $\mathcal{A}$ ,  $crossEntropy()$ ,  $generateEpisode()$  [22]
Output:  $\theta^*$ 
Data:  $\mathcal{D}_{tr}$ 
/*  $n$  classes are chosen from  $\mathbf{y}$ , for which
    $m$  samples are chosen as support,
   while  $m$  from remaining samples are
   chosen as query */
i=0 loss=list()
for  $i \leq episodes\_per\_batch$  do
  support, supportLabels, query, queryLabels =
    generateEpisode( $\mathcal{D}_{tr}$ ,  $m$ ,  $n$ ) /* Sample  $n$ -way,
     $m$ -shot episode from  $\mathcal{D}_{tr}$  */
   $E_{support} = \mathcal{F}(x_{support})$ 
   $E_{query} = \mathcal{F}(x_{query})$ 
   $\mathcal{A}.train(E_{support}, supportLabels)$ 
   $logits = \mathcal{A}.predict(E_{query})$ 
   $\mathcal{L}^{meta} = crossEntropy(logits, queryLabels)$ 
   $loss.append(\mathcal{L}^{meta})$ 
 $\theta^* = optimizer.backward(mean(loss))$  /* mean
  classification loss across all
  episodes of a batch is backprop */

```

where $[(\mathcal{D}_s^t, \mathcal{D}_q^t)]_{t=1}^T$ is the few-shot episode generated from training set which are n -way, m -shot classification tasks, following the strategy proposed in [22]. Every episode has a support set \mathcal{D}_s^t which is used to train the task classifier \mathcal{A}_ϕ and a query set \mathcal{D}_q^t to train \mathcal{F}_θ . Further, ℓ is the task-classification loss which depends on the classifier \mathcal{A}_ϕ used. Task classifier could be a nearest neighbour (ProtoNet [21]), linear regression (Ridge [24]) or linear SVM (MetaOptNet [23]). In order to train the entire setup in an end-to-end fashion, in case of SVM, implicit functions are utilized to capture *approximate* gradients to train the embedding function [23]. During the validation and test routines, similar episodes are sampled from a *novel* set of classes that are previously unseen by the classifier. However, only the task classifier (\mathcal{A}) is updated from the validation or test episode samples. The meta-learning loss (\mathcal{L}^{meta}) is weight regularized ($\mathcal{R}(\theta)$) to ensure generalization and preventing over-fitting of the network. Hence, the batch optimization is performed over multiple episodes to compute the average \mathcal{L}^{meta} . The details of the encoder embedding function is described next.

2.2. Self-supervised encoder

For the input utterances encoder ($\mathcal{F}(\cdot)$), we use a convolutional neural network as the embedding function. Specifically, we utilize the PASE+ architecture [15, 16] pre-trained on LibriSpeech dataset[26] with self-supervised tasks obtained from the original signal. The convolution based architecture first encodes the input with SincNet [27], that learns custom band-pass filters for the input signal. Next, seven residual network blocks are applied with skip connections [28] that are pooled together with the QRNN temporal pooling [29] method. Additionally, in this work, we propose to continue to regularize the embedding learner obtained from PASE+ with some of the self-supervision tasks that have been shown to improve representation power of the embedding. We utilize classification *workers* that reconstruct static-features engineered from the input speech signal from the same base encoding used for meta-learning. To the best of our knowledge, our work is first to explore the benefits of self-supervision in conjunction with meta-learning task. We also hypothesize a weighted combination of the reconstruction loss from the self-supervision (with a controlling parameter α) together with the cross-entropy loss from the meta-

learning few-shot classification, as $\mathcal{L}^{total} = \alpha\mathcal{L}^{meta} + (1 - \alpha) \sum_{w=0}^W \mathcal{L}_w^{self}$, where \mathcal{L}^{self} is the combined loss of all workers of PASE+. The overall routine of the proposed approach for a single batch is shown in Algorithm 1. As mentioned, PASE+ which is pre-trained with only self-supervision is used as the embedding function. Next, we show the empirical evaluation of the proposed architecture for few-shot learning. We showcase 5-shot an 1-shot classification of intent on two popular datasets.

3. Experiments

3.1. Dataset

Experiments are performed on two popular datasets: Google Speech Commands dataset [2] and Fluent Speech Commands[13]. This work presents a speech to intent classification task with few-shots per label, the number of samples per class and classes seen in train and test is controlled via a novel protocol first proposed in [22]. First, we split the datasets into train, validation and test such that all labels are novel (completely different class labels in each split). The details of distribution in the revised protocol resulting from the new splits are shown in Table 1. Additionally, we prune the splits by utilizing the available speaker labels to create two versions of the meta-learning protocol, with and without speaker overlap, referred to as SPO and No-SPO, respectively.¹. A similar protocol is proposed by [17], however, the authors ignore speaker overlap that may inadvertently boost the reported performance.

Table 1: *Dataset statistics on Google Commands and Fluent Speech Commands of the proposed meta-learning protocol*

		Google Commands			Fluent Speech Commands		
		Train	Validation	Test	Train	Validation	Test
SPO	#Classes	18	5	5	15	8	8
	#Audio Files	42,954	10,560	11,207	13,030	7,632	8,572
	#Speakers	1881	1881	1881	97	97	97
	Duration (hours)	8.28	4.85	5.27	10.44	3.54	3.72
No-SPO	#Audio Files	30,181	1,389	1,436	10,115	811	982
	#Speakers	1503	189	189	77	10	10
	Duration (hours)	6.39	0.5	0.65	8.24	0.38	0.39

3.2. Design choices

Experiment design: We show two scenarios of few-shot learning - (i) 5-way 5-shot and (ii) 5-way 1-shot classification for both validation and test scenarios. Further, we observe empirically that the performance on the test set is better when the number of shots in training episodes are higher, hence we generate 15-shot episodes in train. All experiments are conducted for a maximum of 30 epochs, where for each epoch, one batch consists of 1000 train episodes. The episode generation strategy is performed with replacement, hence an infinite number of episodes may be generated from the training set. We cap the number of episodes per batch based on the capacity of the GPU memory. We sample 2000 episodes from validation and test and report mean and standard deviation of the total n -way accuracy [22]. All the models are optimized using SGD with Nesterov momentum[30], using an initial learning rate of 0.01 and momentum of 0.9 with step decay at every 20 epochs. We use mini-batches of 75 waveforms of 1s for Google Commands and 4s for Fluent Speech Commands dataset.

¹The dataset splits used in our paper are available at <https://github.com/AshishMittal/RMLIntent>

Table 2: *Test Intent Accuracy comparison of the ProtoNet[21] model with baseline (supervised n -way m -shot) and skyline (supervised n -way all-shot).*

Model	Google Commands		Fluent Speech Commands	
	5-shot	1-shot	5-shot	1-shot
Baseline	26.22 ± 7	22.39 ± 13	27.42 ± 7	21.32 ± 13
Skyline	95.61		88.98	
ProtoNet [21]	78.86 ± 0.40	69.30 ± 0.54	78.00 ± 0.39	64.24 ± 0.52

Embedding: As mentioned before, we use the PASE+ architecture as the encoder for the input utterances. To ensure the fast performance of the base learners, the original PASE+ encoding of 256 per frame is reduced to a vector size of 50 per frame. Further, a single vector representation is obtained by computing mean across frames. Importantly, our empirical analysis shows that retaining the self-supervised workers does not significantly affect the performance of the downstream task of few-shot classification. We experimented with several loss weighting strategies (by tuning α) and found that the self-supervised workers deter performance. We hypothesize that since the meta-learning framework gathers meta-loss \mathcal{L}^{meta} as the average of classification loss across multiple episodes, the model is already sufficiently regularized without the aid of self-supervised workers. However, the PASE+ encoder is pretrained with self-supervision workers and significantly affect the performance, showing that the workers helped to generalize the representation during pretraining and facilitating transfer to the new classes.

Benchmark: To the best of our knowledge, this is the first work to evaluate any approach to few-shot speech to intent classification tasks in literature. Hence, we also set a baseline and a hypothetical skyline on the same experimental protocol: **Baseline:** For the baseline, we create a 5-way 1-shot and 5-way 5-shot intent classification task, where we train the classification model (using the same pretrained PASE+ embedding) for the same 5 classes as the proposed protocol. **Skyline:** For the skyline, we create a 5-way intent classification task, using *all* the training data available for only those classes in the original split of both datasets (on average, n is 1786 in Google commands and 845 in Fluent.ai). Here, the model is tested on the same classes to give us the *maximum* ceiling that can be reached in a traditional classification task.

In both cases, the architecture is a traditional multi-layer perceptron network with embedding obtained from the PASE encoder trained with categorical cross-entropy loss. For both the baseline and the skyline we train the encoder, as well as the MLP neural network using SGD.

3.3. Evaluation

In this section we evaluate the performance of the proposed model for both the datasets. As shown in Table 2, we report accuracy scores as the mean and standard deviation of the accuracy of 2000 randomly sampled episodes from validation and test splits [22]. Table 3 shows a more detailed analysis of various base learners and their performance on SPO protocol and the no-SPO subset. In all cases, the test accuracy is reported at best validation accuracy obtained during the training process.

As expected, the baseline model performance is close to chance when trained on the proposed few-shot protocol for both the datasets. Note that the baseline is unable to leverage the pre-trained PASE+ vectors with such few training samples. On the other hand, the fully supervised skyline (or all-shot experiment)

Table 3: Detailed evaluation on both datasets with three base learners. The no-SPO is the subset of protocol where the speakers do not overlap.

Model		Google Commands				Fluent Speech Commands			
		5-Shot		1-Shot		5-Shot		1-Shot	
		SPO	No-SPO	SPO	No-SPO	SPO	No-SPO	SPO	No-SPO
ProtoNet [21]	Val	85.45 ± 0.32	83.48 ± 0.35	71.07 ± 0.52	70.61 ± 0.53	83.55 ± 0.38	68.57 ± 0.45	70.42 ± 0.55	56.14 ± 0.56
	Test	89.63 ± 0.27	78.86 ± 0.40	74.35 ± 0.50	69.30 ± 0.54	78.86 ± 0.40	78.00 ± 0.39	65.61 ± 0.52	64.24 ± 0.52
Ridge [24]	Val	82.38 ± 0.34	81.33 ± 0.35	68.49 ± 0.51	68.03 ± 0.51	81.47 ± 0.39	67.75 ± 0.42	70.56 ± 0.51	56.69 ± 0.55
	Test	89.11 ± 0.29	86.17 ± 0.31	75.05 ± 0.48	76.34 ± 0.49	75.75 ± 0.40	78.51 ± 0.39	61.64 ± 0.53	60.84 ± 0.51
MetaOptNet [23]	Val	76.31 ± 0.39	83.48 ± 0.35	62.90 ± 0.51	60.50 ± 0.51	78.52 ± 0.43	65.90 ± 0.46	67.68 ± 0.53	52.77 ± 0.53
	Test	84.37 ± 0.32	88.64 ± 0.29	66.61 ± 0.53	70.82 ± 0.50	70.16 ± 0.43	69.84 ± 0.44	58.71 ± 0.50	59.44 ± 0.51


is trained with all the available training data in an end-to-end fashion. While the performance of the model is high, the fine-tuning of the proposed approach only requires fine-tuning of only the base learner, which is a fraction of the computation cost of the skyline training process.

As shown in Table 2, the proposed approach achieves $76.34 \pm 0.49\%$ and $64.24 \pm 0.52\%$ on the 5-way 1-shot experiment, i.e., on a previously unseen class, the base learner fine-tunes with 5 samples (1 example per class) and provides the mentioned average performance. The low variance across 2000 randomly selected episodes from validation and from test show that the performance is fairly stable across different possible combinations of classes and instances in the test set.

Next, we show detailed performance analysis of the proposed model with various base learners used in the representational meta-learning framework in Table 3. We specifically run each experiment setting with the three popular representation learning protocols - ProtoNet [21], Ridge [24] and MetaOptNet (SVM) [23]. The best performing model on the Google Commands dataset for 5-shot classification achieves the accuracy of 89.63% at the test time, which is only few percent short of the skyline accuracy of 95.61% using only 5 examples in the training for each class. The best accuracy of 88.64% on the harder No-SPO Google Commands dataset is achieved using SVM base learner. This result signifies that our model performs equally well when there are new speakers and classes at the test time. Self-supervised learning of the PASE+ encoder and the subsequent representation learning of MetaOptNet, makes sure that the model do not overfit on speaker traits, but instead learn robust representation that is useful for the classification task. The best results for the 1-shot classification on the Google Commands dataset for both SPO and No-SPO variants are $\sim 75\%$ and $\sim 76\%$, which demonstrates the efficacy of the proposed model.

The best accuracy on the Fluent Speech Commands dataset for 5-shot classification task is around $\sim 79\%$ for the SPO variant, which is approximately close to the No-SPO variant. The best accuracy for the 1-shot classification task for the Fluent Speech Commands dataset is around $\sim 65\%$ for both the variant of the dataset. There is a drop in the accuracy for the Fluent Speech Commands dataset as compared to Google Commands, because the former has longer utterances ($\sim 3s$) as compared to the latter ($\sim 1s$). Another reason being that the classes in Fluent Speech Commands overlap significantly (e.g. deactivate lights, deactivate lights kitchen, deactivate lights washroom). Finally, we note that for the SPO variant of the dataset, ProtoNet base learner which is a nearest neighbor based works well because there is overlap in the speakers, but for the harder No-SPO variant, discriminating base learners SVM work better.

Generally, max-margin classifiers such as SVM tend to outperform nearest neighbor distance based approaches, we find



(a) Fluent Speech Commands (b) Google Commands

Figure 3: The confusion matrix is mean computed for 5-shot classification task over 1000 test episodes.

that the performance of ProtoNet is often superior to SVM. Similarly, a closed-form solution of Ridge may be more sensitive to noise than ProtoNet leading to lower performance. Specifically, during the fine-tune process the linear SVM hyper parameters are fixed as ($C=0.1$, $\text{maxIter}=15$). The robustness of nearest neighbor, particularly in 1-shot makes the case for simpler models. Despite our exhaustive experimentation, we acknowledge that further hyper-parameter tuning may improve MetaOptNet. Finally, Figure 3 shows the confusion matrix for both the datasets on all their test classes, i.e. 5-shot n -way with $n = 6, 8$ respectively. Notice that the model predictably fails on those intent classes that have significant overlap in their phrasings. In Fluent Speech Commands dataset, the class *increase,heat,bedroom* is confused with the class *increase,heat,kitchen* 19% of the times and 17% of the times with the class *decrease,heat,bedroom*, all of which share lexicons with the each other. On the Google Commands dataset there is little overlap in lexicon thus the model learns robust discriminating features using few-shots as evident by the accuracy scores.

4. Conclusions

In this work, we show the classification and generalization of our few-shot learning with the representation based meta-learning paradigm. On two benchmarks for spoken intent classification, we show that the proposed approach can fine-tune with limited examples (we tested one and five) of novel classes and perform classification task with sufficient accuracy. The proposed approach also benefits from the self-supervised pre-trained embedding function trained on an additional dataset. Intent classification and command detection have been prevalent in several consumer devices. However, the list of commands have as yet been static or pre-programmed. Our results indicates the possibility of a light weight classification model (such as a nearest neighbor or linear SVM) that can be easily fine-tuned for new commands with limited training examples.

5. References

- [1] C.-C. Kao, M. Sun, Y. Gao, S. Vitaladevuni, and C. Wang, "Sub-band convolutional neural networks for small-footprint spoken term classification," *Interspeech 2019*, Sep 2019. [Online]. Available: <http://dx.doi.org/10.21437/interspeech.2019-1766>
- [2] P. Warden, "Speech commands: A dataset for limited-vocabulary speech recognition," *CoRR*, vol. abs/1804.03209, 2018.
- [3] J. Poncelet and H. V. hamme, "Multitask learning with capsule networks for speech-to-intent applications," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2020. [Online]. Available: <http://dx.doi.org/10.1109/icassp40776.2020.9053832>
- [4] Y. Huang, H.-K. Kuo, S. Thomas, Z. Kons, K. Audhkhasi, B. Kingsbury, R. Hoory, and M. Picheny, "Leveraging unpaired text data for training end-to-end speech-to-intent systems," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7984–7988.
- [5] Y. Qian, R. Ubale, V. Ramanarayanan, P. L. Lange, D. Suendermann-Oeft, K. Evanini, and E. Tsuprun, "Exploring asr-free end-to-end modeling to improve spoken language understanding in a cloud-based dialog system," *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pp. 569–576, 2017.
- [6] P. Haghani, A. Narayanan, M. Bacchiani, G. Chuang, N. Gaur, P. Moreno, R. Prabhavalkar, Z. Qu, and A. Waters, "From audio to semantics: Approaches to end-to-end spoken language understanding," in *2018 IEEE Spoken Language Technology Workshop (SLT)*, 2018, pp. 720–726.
- [7] D. Serdyuk, Y. Wang, C. Fuegen, A. Kumar, B. Liu, and Y. Bengio, "Towards end-to-end spoken language understanding," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5754–5758.
- [8] S. Ghannay, A. Caubrière, Y. Estève, A. Laurent, and E. Morin, "End-to-end named entity extraction from speech," *CoRR*, vol. abs/1805.12045, 2018. [Online]. Available: <http://arxiv.org/abs/1805.12045>
- [9] Y. Chen, R. Price, and S. Bangalore, "Spoken language understanding without speech recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 6189–6193.
- [10] S. Ghannay, A. Caubrière, Y. Estève, N. Camelin, E. Simonnet, A. Laurent, and E. Morin, "End-to-end named entity and semantic concept extraction from speech," in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 692–699.
- [11] Y. Qian, R. Ubale, P. Lange, K. Evanini, and F. Soong, "From speech signals to semantics — tagging performance at acoustic, phonetic and word levels," in *2018 11th International Symposium on Chinese Spoken Language Processing (ISCSLP)*, 2018, pp. 280–284.
- [12] A. Caubrière, N. Tomashenko, A. Laurent, E. Morin, N. Camelin, and Y. Estève, "Curriculum-based transfer learning for an effective end-to-end spoken language understanding and domain portability," *arXiv preprint arXiv:1906.07601*, 2019.
- [13] L. Lugosch, M. Ravanelli, P. Ignoto, V. S. Tomar, and Y. Bengio, "Speech model pre-training for end-to-end spoken language understanding," *Proc. Interspeech 2019*, pp. 814–818, 2019.
- [14] S. Bhosale, I. Sheikh, S. H. Dumpala, and S. K. Kopparapu, "End-to-end spoken language understanding: Bootstrapping in low resource scenarios," in *INTERSPEECH*, 2019.
- [15] S. Pascual, M. Ravanelli, J. Serrà, A. Bonafonte, and Y. Bengio, "Learning problem-agnostic speech representations from multiple self-supervised tasks," *Proc. Interspeech 2019*, pp. 161–165, 2019.
- [16] M. Ravanelli, J. Zhong, S. Pascual, P. Swietojanski, J. Monteiro, J. Trmal, and Y. Bengio, "Multi-task self-supervised learning for robust speech recognition," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6989–6993.
- [17] Y. Chen, T. Ko, L. Shang, X. Chen, X. Jiang, and Q. Li, "An investigation of few-shot learning in spoken term classification," *arXiv preprint arXiv:1812.10233*, 2018.
- [18] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR.org, 2017, pp. 1126–1135.
- [19] A. Rajeswaran, C. Finn, S. M. Kakade, and S. Levine, "Meta-learning with implicit gradients," in *Advances in Neural Information Processing Systems*, 2019, pp. 113–124.
- [20] N. R. Koluguri, M. Kumar, S. H. Kim, C. Lord, and S. Narayanan, "Meta-learning for robust child-adult classification from speech," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 8094–8098.
- [21] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *Advances in Neural Information Processing Systems*, 2017.
- [22] O. Vinyals, C. Blundell, T. Lillicrap, and D. Wierstra, "Matching networks for one shot learning," in *Advances in neural information processing systems*, 2016, pp. 3630–3638.
- [23] K. Lee, S. Maji, A. Ravichandran, and S. Soatto, "Meta-learning with differentiable convex optimization," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10 657–10 665.
- [24] L. Bertinetto, J. F. Henriques, P. Torr, and A. Vedaldi, "Meta-learning with differentiable closed-form solvers," in *International Conference on Learning Representations*, 2019.
- [25] S. Gidaris and N. Komodakis, "Dynamic few-shot visual learning without forgetting," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4367–4375.
- [26] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [27] M. Ravanelli and Y. Bengio, "Speaker recognition from raw waveform with sinenet," in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 1021–1028.
- [28] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [29] J. Bradbury, S. Merity, C. Xiong, and R. Socher, "Quasi-Recurrent Neural Networks," *International Conference on Learning Representations (ICLR 2017)*, 2017.
- [30] Y. Nesterov, "A method for unconstrained convex minimization problem with the rate of convergence $O(1/k^2)$," in *Doklady an ussr*, vol. 269, 1983, pp. 543–547.