

# On the cross-lingual transferability of multilingual prototypical models across NLU tasks

Oralie Cattan<sup>1,2</sup>

Christophe Servan<sup>1</sup>

Sophie Rosset<sup>2</sup>

<sup>1</sup>QWANT  
61 rue de Villier,  
92200 Neuilly-sur-Seine, France  
inital.lastname@qwant.com

<sup>2</sup>Université Paris-Saclay,  
CNRS, LISN,  
91405, Orsay, France  
lastname@lisn.fr

## Abstract

Supervised deep learning-based approaches have been applied to task-oriented dialog and have proven to be effective for limited domain and language applications when a sufficient number of training examples are available. In practice, these approaches suffer from the drawbacks of domain-driven design and under-resourced languages. Domain and language models are supposed to grow and change as the problem space evolves. On one hand, research on transfer learning has demonstrated the cross-lingual ability of multilingual Transformers-based models to learn semantically rich representations. On the other, in addition to the above approaches, meta-learning have enabled the development of task and language learning algorithms capable of far generalization. Through this context, this article proposes to investigate the cross-lingual transferability of using synergistically few-shot learning with prototypical neural networks and multilingual Transformers-based models. Experiments in natural language understanding tasks on MultiATIS++ corpus shows that our approach substantially improves the observed transfer learning performances between the low and the high resource languages. More generally our approach confirms that the meaningful latent space learned in a given language can be generalized to unseen and under-resourced ones using meta-learning.

## 1 Introduction

Traditionally, Natural Language Understanding (NLU) is an intermediate module between the user interface and the dialogue management module in a dialogue system. It aims to extract semantic information from a user’s query or utterance to fill slots in a domain specific semantic frame. Domain classification, intent detection and slot filling are three core components belonging to the NLU. They are in charge of determining the domain or service of a users query, its underlying goal or intent and asso-

ciating utterance segments with conceptual labels, called slots, similar to named entity recognition.

NLU is usually defined as a supervised learning problem, involving conventional machine learning models on massive amount of annotated training data, which are language dependent. This prerequisite has prevented its widespread adoption for poorly endowed languages and for small technology companies that do not benefit from millions of users to gather data. Besides the requirement of a large amount of annotated data being available, domains, intents and slots are language dependent. Consequently, in practice, the resulting systems are hardly adaptable to expand to new languages.

As a solution to this problem, cross-lingual transfer approaches were developed to leverage the knowledge from well-resourced languages, with task specific data available to under-resourced languages with little or no data. Recent efforts focused on training Transformer models multilingually such as the multilingual version of BERT (Devlin et al., 2019). While earlier work demonstrated the effectiveness of multilingual models to learn representations which are transferable across languages, they show limitations when applied to low-resource languages (Pires et al., 2019; Conneau et al., 2020). From another perspective low-shot learning such as few-shot and zero-shot, aims to transfer knowledge learned from one language to another when the training data is limited or is missing some task labels.

As a core contribution, we explore the potential for cross-lingual transferability of multilingual Transformer-based model (Vaswani et al., 2017) (mBERT) combined with a few-shot learning algorithm based on prototypical representations. We also introduce a zero-shot scenario, where models are trained on multiple languages and evaluated on another. Our proposed approach relies on appending a mBERT encoder module to the prototypical neural network, which is a proven few-shot model,

originally designed for image classification. Our experimental results show that the generated model trained with a limited number of annotated training examples outperforms the transfer learning based approach on MultiATIS++ dataset (Xu et al., 2020; Upadhyay et al., 2018) and can be applied to unseen languages directly with decent performance.

## 2 Related work

The availability of large datasets has enabled deep learning methods to achieve great success in a variety of fields. However, most of these successes are based on supervised learning approaches, which require lots of labeled data to train. Most datasets are only available in English. Only a few other languages are supported, and most of them are considered as under-resourced languages.

Recently, meta-learning approaches have enabled the development of task-agnostic learning algorithms capable of far generalizations (cross-domain or cross-lingual) in the context of having a low-data regime. Because literature on low-shot learning is vast and diverse, only the most relevant approaches to this work are presented and we refer the reader to Vanschoren (2019) and Wang et al. (2019) for a surveys of earlier work.

### 2.1 Low-shot learning

Humans manifest a capacity of learning new concepts from few stimuli quickly and efficiently by utilizing prior knowledge and experience. Inspired by this ability, there has been a resurgence of interest in designing specialized models to perform low-shot learning. An example of this form of learning is metric-based approaches founded on the simple idea of learning a discriminative metric space in which similar samples are mapped close to each other and dissimilar ones distant. Siamese (Koch, 2015), Matching (Vinyals et al., 2016) or Prototypical (Snell et al., 2017) networks belong to this category.

#### 2.1.1 Supervised generalization

In recent years, several approaches have been introduced and refined to overcome the issue of data-limited regime. As an example, the Prototypical Neural Networks (PNNs), developed by Snell et al. (2017) originally for image classification, were used to extract representative characteristics of the data by mapping data points into an embedding space where each sample will cluster around their respective prototype representation. Fort (2017)

proposed to extend their work by adding a confidence region around prototypes with the help of Gaussian covariance models. With the aim of improving the generalization capacity of metric-based methods, Wang et al. (2018) proposed to enforce a large margin between the class prototypes by modifying the standard softmax loss function.

#### 2.1.2 Semi-supervised generalization

Other approaches, closely related to the aforementioned ones, proposed to take advantage of labeled and unlabeled data. Among them, Boney and Ilin (2017) extended PNNs to address semi-supervised image classification problems. They applied a hard clustering to assign the class for the unlabeled examples within the latent space learned by the PNNs. A close method was developed by Ren et al. (2018) to refine the prototype generation process with clustering. The authors introduced distractor classes with the aim of handling unlabeled samples not belonging to any of the known classes.

Most of these approaches have mainly been explored in the field of computer vision, and a few of them were applied to NLP fields, such like Natural Language Understanding (NLU).

### 2.2 NLU using low-shot learning

A number of different deep learning approaches have been applied to the problem of language understanding in recent years. For a thorough overview of deep learning methods in conversational language understanding, we refer the readers to Gao et al. (2018). In the context of relying on limited training resources, few-shot learning has been used for NLU tasks. Yazdani and Henderson (2015) proposes a method to leverage unlabeled data in order to find the separating hyperplanes that divide the utterances with the same label from those with different labels. Sun et al. (2019) extended PNNs for intent classification using hierarchical attention mechanisms when generating the prototype representations.

Slot filling using few-shot models has also been explored. Ferreira et al. (2015) presented a zero-shot approach based on a knowledge base and on word representations learned from unlabeled data. Fritzler et al. (2019) applied PNNs to few-shot named entity recognition by training a separate model for each entity type and Hou et al. (2019) proposed a conditional random forest-based approach enhanced with transfer mechanisms that implicitly incorporate label dependencies and sim-

ilarities. More recently, Dou et al. (2019), Bansal et al. (2020a) and Bansal et al. (2020b) applied various meta-learned models to few-shot NLU across domains and tasks.

Finally, besides the approaches of Gu et al. (2018) and Zhang et al. (2020) that focus on handling new and low-resource languages for machine translation, to the best of our knowledge, there are no approaches that combine cross-lingual transfer and meta-learning methods for NLU tasks.

### 3 Approach

In this section, we present the design of a Prototypical Neural Network and its episodic training procedure before introducing our approach.

#### 3.1 Prototypical Neural Networks

Prototypical Neural Networks (Snell et al., 2017) or PNNs are based on the computation of distance measures between seen-class prototypes to unseen ones. More specifically, a  $D$ -dimensional embedding is generated for each example  $x \in \mathbb{R}^D$  using a neural network based function  $f(\cdot)$  parameterized by  $\Theta$ . This function enhances the encoding process with better separability properties through a non-linear mapping  $f_\Theta : \mathbb{R}^D \rightarrow \mathbb{R}^M$ . The  $M$ -dimensional prototype of each class is formed as the centroid  $c_i$  of their embedded support points as seen in Equation (1):

$$c_i = \frac{1}{|S_i|} \sum_{(x_j, y_j) \in S_i} f_\Theta(x_j), \quad (1)$$

where  $S_i$  represents the set of examples labeled with class  $i$  and  $y_j$  the corresponding label of  $x_j$ . Equation (2) shows how, given a query (that is, a new and an unlabeled sample)  $q_i$ , the probability distribution over the prototypes is computed from  $d(\cdot, \cdot)$ , an arbitrary similarity measures function such as the squared euclidean distance or cosine similarity.

$$p_\Theta(y_i|q_i) = \frac{\exp(-d(f_\Theta(q_i), c_i))}{\sum_{i'} \exp(-d(f_\Theta(q_i), c_{i'}))} \quad (2)$$

Finally, the class with the highest probability is chosen by a softmax over the distances and at optimization time, the negative log-probability  $J(\Theta) = -\log p_\Theta(y_i|q_i)$  of the true class of each query point is minimized by stochastic gradient descent during an episodic learning process described in the next subsection.

#### 3.2 Episodic learning

With the aim of generalizing unseen classes from zero to few training examples per class, PNNs is trained from a collection of  $N$ -way,  $k$ -shot classification tasks through an episodic training procedure (Vinyals et al., 2016). Specifically, each episode is one mini-batch consisting of  $k$  examples from each of the  $N$  classes (both randomly sampled), used to form a labeled (support  $S$ ) and an unlabeled set of examples (query  $Q$ ). The parameter  $k$  often takes a very small value, meaning we have zero-to- $k$  labeled samples. During training, the model is fed with  $S$  to construct the class prototypes using Equation (1). Its parameters are learned in order to minimize the prototypical loss of its predictions for the examples in the given  $Q$  according to Equation (3) of Section 3.1. The evaluation is done by averaging the classification performances on query sets of many testing episodes.

#### 3.3 Transformer-based PNNs

Studies have demonstrated that contextualized representations produced by language models such as ELMo (Peters et al., 2018) or BERT (Devlin et al., 2019) gave neural networks a better training initializations. Rather than training the initialized encoder of PNNs with feature extractors such as convolution or recurrent networks we propose to induce robustness of the pre-trained multilingual BERT (mBERT) to test the distinctiveness of the representation of each class accross languages. The embedding layer is initialized with the pre-trained mBERT embeddings and fine-tuned together with a dense linear layer that defines the embedding space where the prototype-based classifier operates. This latent space is used to learn prototypes of each class by estimating their mean and the chosen class is derived from the output layer of the network based on a softmax over distance to the class prototypes. The motivation behind fine-tuning the encoder with prototypical loss is to induce better generalization properties at test-time to new class labels not seen during training given only a few examples.

#### 3.4 The cross-lingual way

As introduced earlier, even though recent works demonstrate strong cross-lingual transfer capability of multilingual pretrained BERT, they exhibit limitations when applied to low-resource languages (Pires et al., 2019; Conneau et al., 2020). To enable cross-lingual transfer according to our few-shot

Language	# utterances			# tokens			# intents	# slot types
	train	dev	test	train	dev	test		
English	} 4488	490	893	50755	5445	9164	} 18	84
Spanish				55197	5927	10338		
Portuguese				55052	5909	10228		
German				51111	5517	9383		
French				55909	5769	10511		
Chinese				88194	9652	16710		
Japanese				133890	14416	25939		
Hindi	1440	160	893	16422	1753	9755	17	75
Turkish	578	60	715	6132	686	7683	17	71

Table 1: Details of the MultiATIS++ corpus.

scenario, we construct multiple episodic batches  $E$ . From the available data, we draw the task sets by sampling a subset of labels to form a support set from data in the high-resources languages and a query set from data in the low-resource languages to be evaluated. NLU data consists of utterances composed of sentence-level intent labels and sequences of slot labels annotated in BIO format (Ramshaw and Marcus, 1995) to define the boundary of slots. The  $N$ -way  $k$ -shot NLU task is then defined as follows: given an input query utterance in a new language  $q_i$  and a  $k$ -shot support set  $S$  as references, find the most appropriate intent label or slot label sequence  $y$ :

$$\underset{\theta}{\operatorname{argmax}}_E \sum_{(q_i, y_i) \in Q} \log p_{\theta}(y_i | q_i, S). \quad (3)$$

## 4 Experiments

Our NLU experiments in cross-lingual and few-shot learning for under-resources languages are conducted on MultiATIS++ (Xu et al., 2020; Upadhyay et al., 2018) corpus, whose description follows.

### 4.1 The MultiATIS++ corpus

MultiATIS++ (Upadhyay et al., 2018; Xu et al., 2020) is the multilingual extension of the ATIS corpus (Hemphill et al., 1990), which belongs to the air travel planning domain. Originally in English (en), it has been human translated to 8 different other (distant and close) languages i.e., Spanish (es), German (de), French (fr), Portuguese (pt), Hindi (hi), Chinese (zh), Japanese (ja), and Turkish (tr). It contains 37,084 training examples and 7,859 test examples. Details of subsets statistics in terms

of the number of utterances, intents and slots are shown in Table 1. Our main concerns about this corpus are the Hindi and Turkish portions of the data, which are smaller than the other languages, covering only a subset of intents and slots and containing extremely few labeled examples.

### 4.2 Models

We use the fine-tuning procedure (Devlin et al., 2019) of the original mBERT model as our baseline. In sequence-level and token-level classification tasks, it takes the final hidden states (the last layer output of the multi-head Transformer) of the first [CLS] sequence token or each individual token representation as input of the prediction layer to compute classification scores. Since we plan to use transfer learning in the context of PNNs, we fine-tune the pre-trained mBERT model together with a dense linear layer that defines the embedding space (Section 3.3).

### 4.3 Training configurations

We perform three sets of experiments: *target only*, *multilingual* and *multilingual zero-shot*.

- **target only**: this configuration consists of using only the target language data.

We also considered two cross-lingual classification tasks with a varying quantity of data between source and target languages to investigate the behaviour of different types of knowledge transfer.

- **multilingual**: where the training strategy aims to train a network on the concatenation of all of the nine languages and testing the model for each target language.
- **multilingual zero-shot**: where the training relies on the concatenation of all training

config.	encoder	en	es	de	zh	ja	pt	fr	hi	tr
<b>target only</b>	mBERT	98.54	<b>97.31</b>	98.43	<b>97.09</b>	<b>97.20</b>	<b>97.54</b>	<b>98.88</b>	90.93	83.36
	mBERT + PNN (5w1s)	97.46	95.14	97.18	96.35	95.53	96.80	97.11	84.95	85.17
	mBERT + PNN (5w10s)	<b>98.77</b>	96.97	<b>98.54</b>	97.0	96.64	97.42	97.98	<b>91.33</b>	<b>89.33</b>
<b>multilingual</b>	mBERT	98.42	97.98	98.59	97.65	97.45	<b>98.3</b>	98.46	95.33	<b>93.93</b>
	mBERT + PNN (5w1s)	95.33	93.71	95.93	95.89	94.42	94.00	94.78	91.4	90.91
	mBERT + PNN (5w10s)	<b>99.87</b>	<b>98.54</b>	<b>98.60</b>	<b>98.67</b>	<b>98.54</b>	<b>98.32</b>	<b>98.66</b>	<b>95.49</b>	92.61
<b>multilingual (zero shot)</b>	mBERT	96.42	<b>97.98</b>	<b>97.54</b>	96.71	<b>97.45</b>	97.42	<b>97.87</b>	<b>94.37</b>	<b>91.61</b>
	mBERT + PNN (5w1s)	93.73	92.02	93.27	95.62	91.73	93.51	93.28	90.51	89.92
	mBERT + PNN (5w10s)	<b>96.47</b>	97.87	96.86	<b>97.65</b>	96.64	<b>98.10</b>	97.45	93.17	90.67

Table 2: Averaged intent accuracies obtained with PNNs on 5-way  $k$ -shot classification  $k \in [1, 10]$  (best scores are marked in bold) and baseline results.

datasets from all languages except the one we want to test.

This works only for the baseline approach (*mBERT*), but with our PNNs approach (*mBERT+PNN*), we performs few-shot learning. This means we use only a few training data in the considered language (*target only* and *multilingual* configurations).

For instance, when we evaluate our approach in the English task, we consider only a fraction of the English training dataset to train our *mBERT+PNN* model in the *target only*. In the *multilingual* configuration, our few-shot approach (*mBERT+PNN*) is trained using only a fraction of all the examples provided for each language.

#### 4.4 Training details

For all the baseline models built, we use the publicly available mBERT models pre-trained on over a hundred different languages (Devlin et al., 2019). We trained it using 20 epochs like Xu et al. (2020).

PNNs training was done using a number of 1000 episodes using Euclidean distance as suggested by the original authors (Snell et al., 2017). We consider a configuration parameter and tried a 5-way  $k$ -shot intent classification with  $k \in [1, 10]$  (5w1s and 5w10s) and 5-way 10-shots slot filling.

For all approaches we use AdamW optimizer (Loshchilov and Hutter, 2017) using a learning rate of  $5e-5$  to apply gradients with respect to the loss and weight decay.

All results are reported using the average performances of over 30 runs for intent classification and over 5 runs for slot filling (fewer amount of runs because of higher training time).

#### 4.5 Results

Our experimental findings are summarized in Tables 2 and 3 for the intent classification and the

slot-filling tasks, respectively.

##### 4.5.1 Intent classification results

Using the *target only* configuration, the baseline obtains optimal scores when applied to high resource languages, e.g. *English* (en), *French* (fr) or *German* (de) reaching nearly identical high scores. We obtain the highest baseline scores with an accuracy of 98.8 on the French model, followed by the English model with an accuracy of 98.5. Unlike other mainstream languages, the baseline is less accurate on under-resourced languages, with a loss of 7 to 15 points for intent classification on *Hindi* (hi) and *Turkish* (tr) respectively.

In *multilingual* configuration, baseline models perform reasonably well over all the high-resource languages with a significant performance boost due to the availability of additional data. The *mBERT + PNN (5w10s)* models outperformed the baseline for all languages, except for the Turkish (tr) language.

When transferring from all languages to an unseen one (*multilingual zero-shot* configuration) we observe the best results for the *mBERT* model, except *Portuguese* (pt) and *English* (en) languages, in which the *mBERT + PNN (5w10s)* is 0.5 points better.

Finally, within the framework of the intent classification task, the *mBERT + PNN (5w10s)* model achieves better overall performances in the *multilingual* configuration, especially in the case of under-resourced languages with a gain up to 9 points of accuracy, compared to the *target-only* configuration and an average of one point compared to the best model in the *multilingual zero-shot* configuration.

##### 4.5.2 Slot-filling results

Slot-filling result trends in the *target only* configuration are about one point better of F1 score for the *mBERT + PNN (5w10s)* model compared to the baseline model (*mBERT*). The *mBERT + PNN*

config.	encoder	en	es	de	zh	ja	pt	fr	hi	tr
target only	mBERT	95.64	85.52	94.88	92.93	93.13	91.71	92.78	85.12	78.22
	mBERT + PNN (5w10s)	<b>95.76</b>	<b>87.40</b>	<b>95.63</b>	<b>93.45</b>	<b>93.93</b>	<b>92.22</b>	<b>93.13</b>	<b>85.70</b>	<b>82.67</b>
multilingual	mBERT	96.02	88.03	95.03	93.63	93.01	92.31	91.18	87.39	86.83
	mBERT + PNN (5w10s)	<b>98.40</b>	<b>92.09</b>	<b>97.12</b>	<b>95.50</b>	<b>97.24</b>	<b>95.81</b>	<b>96.80</b>	<b>89.59</b>	<b>88.39</b>
multilingual (zero shot)	mBERT	<b>94.10</b>	<b>87.14</b>	<b>94.23</b>	<b>92.17</b>	<b>92.61</b>	<b>91.59</b>	<b>90.79</b>	86.14	85.86
	mBERT + PNN (5w10s)	93.25	86.99	93.57	91.82	92.38	91.19	90.39	<b>87.49</b>	<b>86.83</b>

Table 3: Averaged slot F1s obtained with PNNs on 5-way 10-shot and baseline results (highest scores are marked in bold).

(5w10s) model even outperformed the baseline by more than 4 points of F1 in the Turkish task (tr).

We can observe the same trend in the *multilingual* configuration: our approach outperformed the baseline in all languages.

On the contrary, the *mBERT + PNN (5w10s)* fails in most of language tasks in the *multilingual zero-shot* configuration, except for the *Hindi* (hi) and the *Turkish* (tr) languages.

Finally, like the intent classification task, the *mBERT + PNN (5w10s)* model achieves better overall performance in the *multilingual* configuration for all languages.

#### 4.6 Result analysis

First, our baseline results are on par with those obtained by Qin et al. (2019) and Xu et al. (2020) when they trained BERT-based models using only English training data (en) with intent accuracy scores of 97.5% and 96.08% while we obtain 98.5%. This is the same in our slot-filling experiment in which they report 94.7 F1 points while we obtain 95.6. This difference comes from our results averaging between 30 and 5 runs for intent classification and slot filling, while previous works only performed 5 runs. We also observe that, just like Xu et al. (2020), slot filling on *Spanish* (es) leads to lower results, similar to those obtained in our few-shot setting.

When transferring from all languages to an unseen one (*multilingual zero-shot configuration* in both tables 2 and 3) we obtained lower scores than the *multilingual* configurations. This means the multilingual representation captured in mBERT is efficient enough when data is available in several languages and none are available in the target considered language. But, in both cases, the combination of mBERT+PNN performs better when fewer data is available using the few-shot approach (the *multilingual* configuration). This means that our approach quickly adapts to the considered target language with only a few examples available and

enhances the mBERT multilingual transfer learning capabilities. This is especially true in the case of slot filling with gains in terms of F1-scores ranging from 2 to 5 points.

Finally, using the mBERT baseline model, transfer learning to *French* or *German* has performance scores similar to *English* while using the *Turkish* (tr) or *Hindi* (hi) yielded significant loss. This leads us to the same conclusion as Xu et al. (2020): exploiting language interrelationships learnt with transfer learning improve the model performances. This may come from the fact that French, English and German are similar and share some vocabulary while Turkish or Hindi are dissimilar to European languages (Hock and Joseph, 2019).

A detailed inspection of the PNNs results shows that in the *target only* and in the *multilingual* configurations, there is an overall and important reduction in recall values, which is balanced by an improvement of the precision values. If we analyze deeper the mislabeled examples we can observe that applying PNNs help to prevent overlapping and annotation mismatch cases that occur in the data.

We observed that MultiAtis++ corpus seems to be a highly unbalanced labeled dataset with the number of training examples per class varying from 1 to 3300. This impacts the model performance, and it could explain why we observe a lower recall and an improvement in precision using our approach, since it is based on the reduction of the amount of data.

## 5 Conclusions

In this paper, we demonstrate the opportunities in leveraging mBERT-based modeling using few-shot learning for both intent classification and slot filling tasks on under-resource languages. We found that our approach model is a highly effective technique for training models for low-resource languages. This illustrates the performance gains that can be achieved by exploiting language interrelationships

learnt with transfer learning, a conclusion further emphasised by the fact that multilingual results outperformed other configuration models (target only and specifically multilingual zero-shot) regardless of the approach. Overall, PNNs models outperform mBERT-based transfer learning approach, enabling us to train competitive NLU systems for under-resources languages with only a fraction of training examples.

From this work a new challenge naturally comes up and a possible direction is to adapt a few-shot setting to a joint approach of intent detection and slot filling, like in Zhang and Wang (2016), Liu and Lane (2016) and Zhang et al. (2019), which demonstrates that performing these two tasks jointly improves the performance of both of them.

## References

- Trapit Bansal, Rishikesh Jha, and Andrew McCallum. 2020a. [Learning to few-shot learn across diverse natural language classification tasks](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5108–5123, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Trapit Bansal, Rishikesh Jha, Tsendsuren Munkhdalai, and Andrew McCallum. 2020b. [Self-supervised meta-learning for few-shot natural language classification tasks](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 522–534, Online. Association for Computational Linguistics.
- Rinu Boney and Alexander Ilin. 2017. [Semi-supervised few-shot learning with prototypical networks](#). In *Workshop on Meta-Learning 2017 (NIPS 2017)*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Zi-Yi Dou, Keyi Yu, and Antonios Anastasopoulos. 2019. [Investigating meta-learning algorithms for low-resource natural language understanding tasks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1192–1197, Hong Kong, China. Association for Computational Linguistics.
- Emmanuel Ferreira, Bassam Jabaian, and Fabrice Lefèvre. 2015. [Online adaptative zero-shot learning spoken language understanding using word-embedding](#). *IEEE International Conference on Acoustics, Speech and SP*.
- Stanislav Fort. 2017. [Gaussian prototypical networks for few-shot learning on omniglot](#). In *Workshop on Bayesian Deep Learning (NIPS 2017)*.
- Alexander Fritzier, Varvara Logacheva, and Maksim KretoV. 2019. [Few-shot classification in named entity recognition task](#). In *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing, SAC '19*, page 993–1000, New York, NY, USA. Association for Computing Machinery.
- Jianfeng Gao, Michel Galley, and Lihong Li. 2018. [Neural approaches to conversational ai](#). In *The 41st International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '18*, page 1371–1374, New York, NY, USA. Association for Computing Machinery.
- Jiatao Gu, Yong Wang, Yun Chen, Victor O. K. Li, and Kyunghyun Cho. 2018. [Meta-learning for low-resource neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3622–3631, Brussels, Belgium. Association for Computational Linguistics.
- Charles T Hemphill, John J Godfrey, and George R Doddington. 1990. [The atis spoken language systems pilot corpus](#). In *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27, 1990*.
- Hans Henrich Hock and Brian D Joseph. 2019. [Language history, language change, and language relationship: An introduction to historical and comparative linguistics](#). Walter de Gruyter GmbH & Co KG.
- Yutai Hou, Zhihan Zhou, Yijia Liu, Ning Wang, Wanxiang Che, Han Liu, and Ting Liu. 2019. [Few-shot sequence labeling with label dependency transfer](#). *CoRR*, abs/1906.08711.
- Gregory R. Koch. 2015. [Siamese neural networks for one-shot image recognition](#). In *ICML Deep Learning Workshop*.
- Bing Liu and Ian Lane. 2016. [Attention-based recurrent neural network models for joint intent detection and slot filling](#). *INTERSPEECH*, pages 685–689.

- Ilya Loshchilov and Frank Hutter. 2017. [Fixing weight decay regularization in adam](#). *CoRR*, abs/1711.05101.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proc. of NAACL*.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Libo Qin, Wanxiang Che, Yangming Li, Haoyang Wen, and Ting Liu. 2019. [A stack-propagation framework with token-level intent detection for spoken language understanding](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2078–2087, Hong Kong, China. Association for Computational Linguistics.
- Lance Ramshaw and Mitch Marcus. 1995. [Text chunking using transformation-based learning](#). In *Third Workshop on Very Large Corpora*.
- Mengye Ren, Sachin Ravi, Eleni Triantafillou, Jake Snell, Kevin Swersky, Josh B. Tenenbaum, Hugo Larochelle, and Richard S. Zemel. 2018. [Meta-learning for semi-supervised few-shot classification](#). In *International Conference on Learning Representations*.
- Jake Snell, Kevin Swersky, and Richard Zemel. 2017. [Prototypical networks for few-shot learning](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4077–4087. Curran Associates, Inc.
- Shengli Sun, Qingfeng Sun, Kevin Zhou, and Tengchao Lv. 2019. [Hierarchical attention prototypical networks for few-shot text classification](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 476–485, Hong Kong, China. Association for Computational Linguistics.
- Shyam Upadhyay, Manaal Faruqui, Gokhan Tür, Hakkani-Tür Dilek, and Larry Heck. 2018. (almost) zero-shot cross-lingual spoken language understanding. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6034–6038. IEEE.
- Joaquin Vanschoren. 2019. *Meta-Learning*, pages 35–61. Springer International Publishing, Cham.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, undefin dukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 6000–6010, Red Hook, NY, USA. Curran Associates Inc.
- Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. 2016. Matching networks for one shot learning. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS’16*, page 3637–3645, Red Hook, NY, USA. Curran Associates Inc.
- Wei Wang, Vincent W. Zheng, Han Yu, and Chunyan Miao. 2019. [A survey of zero-shot learning: Settings, methods, and applications](#). *ACM Trans. Intell. Syst. Technol.*, 10(2).
- Yong Wang, Xiao-Ming Wu, Qimai Li, Jiatao Gu, Wangmeng Xiang, Lei Zhang, and Victor O. K. Li. 2018. [Large margin few-shot learning](#). *CoRR*, abs/1807.02872.
- Weijia Xu, Batoool Haider, and Saab Mansour. 2020. [End-to-end slot alignment and recognition for cross-lingual NLU](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5052–5063, Online. Association for Computational Linguistics.
- Majid Yazdani and James Henderson. 2015. [A model of zero-shot learning of spoken language understanding](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 244–249, Lisbon, Portugal. Association for Computational Linguistics.
- Biao Zhang, Philip Williams, Ivan Titov, and Rico Senrich. 2020. [Improving massively multilingual neural machine translation and zero-shot translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1628–1639, Online. Association for Computational Linguistics.
- Chenwei Zhang, Yaliang Li, Nan Du, Wei Fan, and Philip Yu. 2019. [Joint slot filling and intent detection via capsule neural networks](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5259–5267, Florence, Italy. Association for Computational Linguistics.
- Xiaodong Zhang and Houfeng Wang. 2016. A joint model of intent determination and slot filling for spoken language understanding. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI’16*, page 2993–2999. AAAI Press.