

META LEARNING TO CLASSIFY INTENT AND SLOT LABELS WITH NOISY FEW SHOT EXAMPLES

Shang-Wen Li, Jason Krone, Shuyan Dong, Yi Zhang, Yaser Al-onaihan

Amazon AI

{shangwel, kronej, shuyand, yizhngn, onaihan}@amazon.com

ABSTRACT

Recently deep learning has dominated many machine learning areas, including spoken language understanding (SLU). However, deep learning models are notorious for being data-hungry, and the heavily optimized models are usually sensitive to the quality of the training examples provided and the consistency between training and inference conditions. To improve the performance of SLU models on tasks with noisy and low training resources, we propose a new SLU benchmarking task: few-shot robust SLU, where SLU comprises two core problems, intent classification (IC) and slot labeling (SL). We establish the task by defining few-shot splits on three public IC/SL datasets, ATIS, SNIPS, and TOP, and adding two types of natural noises (adaptation example missing/replacing and modality mismatch) to the splits. We further propose a novel noise-robust few-shot SLU model based on prototypical networks. We show the model consistently outperforms the conventional fine-tuning baseline and another popular meta-learning method, Model-Agnostic Meta-Learning (MAML), in terms of achieving better IC accuracy and SL F1, and yielding smaller performance variation when noises are present.

Index Terms: meta-learning, few-shot learning, spoken language understanding (SLU)

1. INTRODUCTION AND RELATED WORKS

Goal-oriented dialogue systems is a hot topic in machine learning research. The systems have widespread applications in the industry and are the foundation of many successful products, including Alexa, Siri, Google Assistant, and Cortana. One core component of a dialog system is spoken language understanding (SLU), which consists of two main problems, intent classification (IC) and slot labeling (SL) [1, 2]. In IC, we attempt to classify the goal of a user query, usually input in text or transcribed by automatic speech recognition (ASR) system from audio. SL, similar to the named-entity recognition (NER) problem, aims to label each token in a query an entity type. The only difference is that entity types in SL are domain-specific and based upon dialog ontology. Recent advances in neural models have enabled greatly improved SLU [3, 4, 5, 6, 7].

However, two significant challenges hinder the broad application and expansion of the SLU models in industrial settings. First of all, neural methods require a large amount of labeled data for training [8]. SLU is often coupled with the ontology of the underlying dialog system and thus domain-dependent. Collecting a large number of in-domain labeled data for neural models is prohibitively expensive and time-consuming. Secondly, the performance of SLU models in practice often suffers from fluctuations due to various types of noises. One common noise is adaptation data perturbation. In many industrial applications such as cloud services¹, the SLU model is built by fine-tuning (or adapting) a pre-trained, shared network to the target domain with data provided by developers. The developers often have a limited background in SLU and machine learning. Thus the data provided varies in quality and is subject to different types of perturbations, such as missing or replaced data samples (e.g., a subset of an optimal example set is missing or replaced by redundant examples in the adaptation data) and typos. Another common noise comes from the mismatch of input modalities between adaptation and inference stages. For instance, the model is adapted with human transcription yet deployed to understand ASR decoded text, or the input at adaptation and inference stages relies on the recognition of different versions of ASR models. Given that most neural methods comprise a large number of parameters and are heavily optimized for the training (i.e., adaptation in the context of cloud service) data provided, the resulting model is usually sensitive to these noises. The requirement of noise-free adaptation and inference conditions also prohibits the use of neural SLU techniques because it is often infeasible to achieve such conditions.

Transfer learning and meta-learning are two conventional techniques that have been applied to address the challenge of data scarcity. Transfer learning usually refers to pre-training initial models using mismatched domains with rich human annotations and then adapting the models with limited labels in targeted domains. Previous works [9, 10, 11, 12, 13] have shown promising results in applying transfer learning to SLU. Note that pre-training discussed here covers meth-

¹Alexa ASK: <https://developer.amazon.com/en-US/alexa/alexa-skills-kit>; Google DialogFlow: <https://dialogflow.com/>

ods including using a pre-trained language model like BERT [14] directly and further training downstream tasks on data in mismatched domains with the pre-trained model. In the following, we focus on the latter due to utilizing data from other domains better and yielding higher accuracy. In recent years, meta-learning has gained growing interest among the machine learning fields for tackling few-shot learning (i.e., data scarcity) scenarios. Model-Agnostic Meta-Learning (MAML) [15] focuses on learning parameter initialization from multiple subtasks, such that the initialization can be fine-tuned with few labels and yield good performance in targeted tasks. Metric-based meta-learning, including prototypical networks (ProtoNets) [16, 17] and matching networks [18], aim to learn embedding or metric space which can be generalized to domains unseen in the training set after adaptation with a small number of examples from the unseen domains. Recent work unveils excellent potential in applying meta-learning techniques to SLU in the few-shot learning context [19].

As compared to data scarcity, another challenge for SLU, the robustness against noises, is also gaining attention. Simulated ASR errors are used to augment training data for SLU models [20]. Researchers also leverage information from confusion networks or lattices [2, 21, 22, 23, 24, 25], and adversarial training techniques [26, 27] for models to learn query embeddings that are robust against ASR errors. For text input, methods have also been explored on model robustness against noises from misspelling and acronym [28]. In contrast to these noise types that have gained attention, to our best knowledge, there is no prior work investigating the impact of missing or replaced examples in adaptation data. Moreover, the intersection of data scarcity and noise robustness is unexplored. Since the scarcity of labeled data and data noisiness usually co-occur in SLU applications (both reflect the difficulty of acquiring annotated data), the lack of studies in the intersectional areas hinders the use of neural SLU models and its expansion to broader use cases.

Given the deficiency, we establish a novel few-shot noisy SLU task by introducing two common types of natural noise, adaptation example missing/replacing and modality mismatch, to the previously defined few-shot IC/SL splits [19]. The task is built upon three public datasets, ATIS [29], SNIPS [30], and TOP [31]. We further propose a noise-robust few-shot SLU model based on ProtoNets for the established task. In summary, our primary contributions are 3-fold: 1) formulating the first few-shot noisy SLU task and evaluation framework, 2) proposing the first working solution for the few-shot noisy SLU with the existing ProtoNet algorithm, and 3) in the context of noisy and scarce learning examples, comparing the performance of the proposed method with conventional techniques, including MAML and fine-tuning based adaptation.

2. APPROACHES

In this section, we dive deep into the formulation of few-shot noisy SLU tasks. We also elaborate on the method we propose to overcome this challenging task.

2.1. Problem formulation

2.1.1. Few-shot SLU

The goal of the few-shot SLU is to adapt an IC/SL classifier f_θ pre-trained on data, $D^{pretrain}$, in mismatch domains with rich annotation to new domains using data, D^{test} , which comprises few labeled examples per class. In this setting, pre-training and test splits are two disjoint class sets $L^{pretrain}$ and L^{test} from $D^{pretrain}$ and D^{test} respectively. Classes in $L^{pretrain}$ are used for model pre-training while those in L^{test} are held out for test time to adapt the pre-trained model and evaluate generalizability. The test is done episodically, where each episode is a mini adaptation task containing a support set S and a query set Q . Adaptation of f_θ to new domains is achieved with a few labeled examples provided by the support set, whereas model performance is evaluated by averaging metrics measured in each episode’s query set. In our implementation, we also pre-train meta-learning methods on episodes, since a previous study showed that pre-training with matched condition yields better performance [19].

We follow the setup in [19] to build support and query sets. We define $S^m = \cup_{l \in L^m} S_l$, where $S_l = \{(x_i^j, l, t_l^i), i \in (1, \dots, k_s)\}$ and m is either the pre-train or test split. Here x_i^j is the input utterance of the i -th example for intent (i.e., class) l , t_l^i is the slot label of x_i^j , and k_s is the number of examples per class in the support set. Similarly Q^m is defined as $\cup_{l \in L^m} Q_l$, where $Q_l = \{(x_j^l, l, t_l^j), j \in (1, \dots, k_q)\}$ and $Q_l \cap S_l = \emptyset$. In this way, we construct an episode by sampling k_s and k_q examples per intent l in $L^{pretrain}$ and L^{test} from $D^{pretrain}$ and D^{test} respectively. Note that in practice D^{test} consists of more than a few examples per class, and thus we downsample D^{test} to resemble a few-shot learning task in each episode. Besides, for fair evaluation, we map slot labels in the query set but not in the support one to **Other**, which is excluded when we evaluate SL performance. The mapping also guarantees not updating gradient for slot labels unseen in support set during pre-training episodically.

2.1.2. Few-shot noisy SLU

Noises usually co-occur with data scarcity in the industrial settings of SLU. Thus we further formulate a few-shot noisy SLU task for benchmark and model development. The task is built upon the few-shot SLU described above, and the goal is to adapt the IC/SL classifier with few examples such that the resulting model can perform well and robust in new domains when noise exists. Specifically, we investigate two types of

noise, adaptation example missing/replacing and modality mismatch².

For adaptation example missing/replacing, $S^{pretrain}$, $Q^{pretrain}$, and Q^{test} are kept identical to the ones in few-shot SLU, while the support set in test split, \bar{S}^{test} , is the perturbed S^{test} . In the following experiment, we perturb S^{test} by either removing c examples per intent from S^{test} , or replacing c examples per intent with ones sampled from D^{test} but excluding S^{test} and Q^{test} . We choose to study these two types of perturbation, for they are the basic operation on the example set at the utterance level. A more complicated variation can be built by combining the two. To quantify the model robustness against the missing and replacing operation, in each episode at the test stage, we adapt the pre-trained classifier with S^{test} and \bar{S}^{test} separately, and evaluate on Q^{test} to measure the performance difference between models adapted with the original and perturbed support set. With this setup, we estimate how well and robust classifiers can perform with a network pre-trained on mismatched but rich-annotated domains as well as a small and perturbed adaptation set. Good and robust performance in such a setting is especially useful in the context of cloud services.

In modality mismatch, the goal is to benchmark performance impact when the preprocessing pipeline of data used in pre-training and adaptation is different from the one in inference. We simulate this noise type by replacing examples in Q^{test} with ASR hypotheses and IC/SL annotation on them. Human transcription is still used in pre-training split and S^{test} for adaptation. Since most SLU benchmarking datasets only provide IC/SL annotation on human transcription, further data processing is required. Here we adopt a common technique, noise-corrupted synthesized speech [32, 33], to obtain ASR hypotheses and corresponding annotation. We first apply TTS on text input and feed the synthesized audio into ASR to generate the ASR hypothesis. Levenshtein alignment between tokens in the hypothesis and original text is then adopted to project SL annotation onto the hypothesis (note that intent labeling is not affected by ASR results). Projected SL labels are reviewed and corrected by human annotators for the experiment. In such a manner, we generate examples to measure model performance and robustness when the data preprocess pipeline, i.e., text input from human vs. audio input recognized by ASR, is mismatched.

2.2. Learning frameworks

To build robust models for the few-shot noisy SLU, we propose ProtoNets [16] based SLU. ProtoNets is a popular meta-learning framework for the few-shot learning scenario. We apply ProtoNets to IC and SL problems by first representing

²We select the two types of noises as they are common in cloud services, where the input modality at deployment can be different from development; the provided adaptation data and its quality can fluctuate due to developers' limited background or deletion per user privacy concerns. We plan to explore other prevalent noises, including typos and acronyms, in future work.

the input utterance and tokens with encoder $Enc(x) : \rightarrow \mathbb{R}^D$. Then we compute prototypes, i.e., centroids of examples, of each intent and slot class with the embedding. That is

$$c_l = \frac{1}{|S_l|} \sum_{(x_l^i, l, t_l^i) \in S_l} Enc(x_l^i) \quad (1)$$

$$c_a = \frac{1}{|S_a|} \sum_{(x_{l_w}^i, l, t_{l_w}^i) \in S_a} Enc(x_{l_w}^i | x_l^i) \quad (2)$$

Here c_l and c_a are the prototype for intent l and slot a ; $Enc(x_l^i)$ and $Enc(x_{l_w}^i | x_l^i)$ are the embeddings of utterance x_l^i and token $x_{l_w}^i$ in utterance x_l^i respectively; w is the token index;

$$S_a = \{(x_{l_w}^i, l, t_{l_w}^i) | i \in (1 \dots k_s), l \in L, w \in (1 \dots |x_l^i|), t_{l_w}^i = a\} \quad (3)$$

is the token level example set for slot a . Given an example $(x^*, l^*, t^*) \in Q$, we predict the intent and each token's slot by computing the softmax of distance from the example to prototypes. Specifically,

$$p(l^* = l | x^*, S) = \frac{\exp(-|Enc(x^*) - c_l|^2)}{\sum_{l'} \exp(-|Enc(x^*) - c_{l'}|^2)} \quad (4)$$

$$p(t_w^* = a | x^*, S) = \frac{\exp(-|Enc(t_w^* | x^*) - c_a|^2)}{\sum_{a'} \exp(-|Enc(t_w^* | x^*) - c_{a'}|^2)} \quad (5)$$

We denote the approach of building IC and SL classifiers with this framework as **Proto**.

In the implementation, we jointly pre-train IC and SL. Thus the loss function is defined as the sum of IC and SL negative log-likelihood averaged over instances in $Q^{pretrain}$ given prototypes computed from $S^{pretrain}$. We backward propagate the gradient of loss at pre-training episodically to tune the encoder. At the testing, the learned encoder along with S^{test} is used for calculating class prototype and predict examples in Q^{test} .

For comparison, we also build two baselines, one based on MAML (denoted as **MAML**) and another based on finetuning (denoted as **Finetune**). MAML is another popular meta-learning framework. We utilize MAML to optimize parameters θ of the IC-SL classifier $f(x|\theta)$ on $S^{pretrain}$ and $Q^{pretrain}$, such that after θ is adapted with the S^{test} for d steps (i.e., backward update gradient by d epochs), the resulting classifier $f(x|\theta')$ can generalize well in Q^{test} . Concretely, we perform the following two-step optimization at pre-training:

$$\begin{aligned} 1. & \theta \leftarrow finetune(\theta, d|S) \\ 2. & \theta \leftarrow \theta - \nabla_{\theta} loss(f(\cdot|\theta'), Q) \end{aligned} \quad (6)$$

Then the learned classifier $f(\cdot|\theta)$ is adapted and evaluated with S^{test} and Q^{test} respectively. Besides, for computational efficiency, in implementation, we adopt first-order approximation of MAML, foMAML, which has been shown to achieve similar performance with less computation [15, 19].

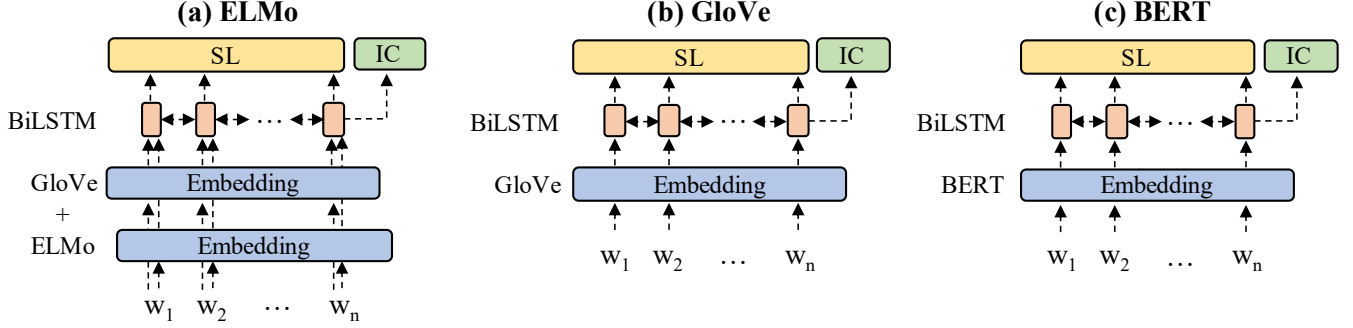


Fig. 1. The architectures for the IC-SL joint classifiers built with various learning frameworks. Three popular architectures for SLU classifiers, **ELMo**, **GloVe**, and **BERT**, are shown.

On the other hand, **Finetune** is a common supervised-learning framework for low resource SLU [34]. In **Finetune**, we pre-train models with examples from $S^{pretrain}$, $Q^{pretrain}$ in batch. Adaptation and evaluation are also conducted in episodes, where the output layers of pre-trained models are first adapted with S^{test} , and then the resulting IC-SL models are evaluated with Q^{test} .

2.3. Model architecture

Figure 1 visualizes the architectures for the IC-SL joint classifiers built with these frameworks. Three popular architectures for SLU, **ELMo**, **GloVe**, and **BERT**, are investigated [19]. The architectures share the same design of using a bi-LSTM [35] layer to encode embeddings and then fully connected IC and SL prediction layers with bi-LSTM hidden states as input. These classifiers differ in the utilized embeddings, where **ELMo** adopts a concatenation of GloVe [36] and ELMo [37], **GloVe** uses GloVe only and **BERT** employs a pre-trained BERT for encoding tokens. In our experiments, we keep these pre-trained encoders frozen in adaptation since our previous study shows that it is insufficient to adapt these encoders with few-shot examples [19]. When applying these architectures to learning framework **Proto**, the distance between the output and prototypes is further computed, and the probability of an IC or SL class is the softmax of the distance. As for **MAML** and **Finetune**, the output of IC and SL layers is used directly as the prediction logits. Note that our learning frameworks discussed above are architecture agnostic. For the generalizability of our experiment, we choose three popular architectures in SLU to explore. Other model backbones are also viable for these frameworks.

3. EXPERIMENTS

3.1. Datasets and experiment setup

We build the few-shot learning task to evaluate the proposed approach based on three public SLU datasets: ATIS [29], SNIPS [30], and TOP [31]. ATIS is a dataset in the airline

Task\Splits	Pre-train	Validation	Test
SNIPS-fs	(20345, 7, TOP) (4373, 5, ATIS)	(4333, 5, TOP) (662, 7, ATIS)	(6254, 3, SNIPS)
ATIS-fs	(20345, 7, TOP) (8230, 5, SNIPS)	(4333, 5, TOP)	(829, 7, ATIS)
TOP-fs	(4373, 5, ATIS) (8230, 5, SNIPS)	(662, 7, ATIS)	(4426, 6, TOP)

Table 1. Statistics of pre-train, validation, and test splits for established few-shot datasets, shown in the form of (utterance counts, intent counts, datasets from which the intents were selected).

domain, while SNIPS comprises utterances in inquiring music, media, and weather. TOP pertains to navigation and event search with nested and flat intent labels. As discussed in previous work [19], we only utilize the non-hierarchical intents in experiments for comparable results. In the context of few-shot learning, data for pre-training and adaptation is often in mismatched domains. Hence, we build the few-shot learning datasets, SNIPS-fs, ATIS-fs, and TOP-fs, by manually selecting intents from SNIPS, ATIS, and TOP respectively for pre-training and validation splits. In Table 3.1, we provide detailed statistics of these few-shot datasets³.

With the established splits, we extract episodes. At the pre-training stage, the number of intents in each episode, i.e., the n -way, is sampled from $[3, |L^{pretrain}|]$ uniformly. After that, we sample n intents from $L^{pretrain}$, and sample k_s and k_q utterances respectively for each sampled intent as the support and query set. At the validation and test stage, n is set to $|L^{validation}|$ and $|L^{test}|$, and the remaining settings are the same. Additionally, with the same rationale as many meta-learning studies where k_s and k_q are set to some small num-

³The splits of intents are selected by maximizing the distance between intents belonging to different splits, where each intent is represented by the average over the BERT-CLS embeddings of its utterances. We also investigated different split methods, such as random sampling, and observed no significant difference. We will release the splits and resulting datasets for reproducibility.

bers arbitrarily [15, 16], in the experiment here, we let both k_s and k_q equal to 10.

Additional steps are required to introduce noises to these few-shot datasets. For adaptation example missing/replacing, we keep the pre-training and validation episodes the same, while perturbing the test episodes by sampling c (set to arbitrarily small numbers, 1 to 5, in the experiment) utterances from each intent in the support set, and *remove* the sampled ones or *replace* with others. For modality mismatch, since SNIPS and TOP only contain human text input, we use commercial TTS⁴ and ASR services⁵, to synthesize the audio and decode the audio back to text. Slot labels on the human text are projected onto ASR hypotheses with Levenshtein alignment. We adopt a similar process for ATIS but skip TTS since audio recordings are available in ATIS. The word error rate (WER) of decoded audio for ATIS, SNIPS, and TOP is 18.4%, 16.2%, and 14.7%, respectively.

In the experiment here, we pre-train and adapt all the models using Adam optimizer. The learning rate is set to 0.001 in **Finetune** and **Proto**. For **MAML**, the rate is 0.003 at the pre-training step and 0.01 at the adaptation with d set to 8. We pre-train models for 30 epochs either episodically (**MAML** or **Proto**) or in batch (**Finetune**) with size 512. At the testing, we adapt these models for ten steps on the support set. We select all the hyperparameters for each approach separately, such that the hyperparameters yield the minimum IC and SL joint loss averaged over the three validation sets without perturbation. As for the model architecture, we use a 2-layer BiLSTM with 256 hidden units for contextual encoding. GloVe with 300 dimensions, ELMo with 512 dimensions, and BERT-medium with 512 dimensions are selected for the token embeddings as these settings are commonly adopted in SLU experiments. To assess the model performance, we report the average IC accuracy and SL F1 over 100 episodes and three random seeds, a typical setting in few-shot learning to avoid performance fluctuation [19].

3.2. Experiment results

Then we evaluate the performance and robustness of the proposed method (**Proto**) as well as baselines (**MAML** and **Finetune**) on established few-shot noisy tasks. Table 3.2 shows the IC accuracy and SL F1 for the methods under three settings of adaptation example perturbation (i.e., no perturbation, *removing* or *replacing* c utterances). Here, we start from minimal perturbation, $c = 1$. For evaluating robustness, we also report in parentheses the absolute accuracy and F1 difference between each perturbed setting and no perturbation counterpart⁶. Similar to previous observation [19], we find

⁴Amazon Polly: <https://aws.amazon.com/polly/>

⁵Amazon Transcribe: <https://aws.amazon.com/transcribe/>

⁶Note that for measuring the absolute difference, we first calculate the difference in IC/SL performance between the noisy and no perturbation setting for each episode. The average and standard deviation of the difference over episodes are reported in the table. The reported metrics are different from

Proto outperforms **MAML** and **Finetune** consistently for IC and SL problems in all the few-shot datasets; also the relatively low SL F1 from baselines suggests that few-shot SL is challenging. In addition to the three learning frameworks, the impact of various model architectures is investigated as well. We observe that **BERT** consistently yields better performance than **GloVe** and **ELMo** (c.f., row 1 to 9), presumably because BERT encodes more knowledge for SLU. Thus, in the following, we choose **BERT** as the backbone model architecture.

Results in Table 3.2 also suggest that **Proto** yields the most robust performance against adaptation example missing/replacing, with absolute differences of IC accuracy and SL F1 between 0.2 and 1.7. We surmise the robustness results from the model’s ensemble nature, where the inference can be viewed as an aggregation of classifier prediction based on the distance to examples. **Finetune** is comparably robust for IC in *replacement* but worse in *removal*, presumably because the latter leaves fewer adaptation utterances, which can be consequential at few-shot learning. The SL performance from **Finetune** is too low to measure the robustness. **MAML**, on the other hand, exhibits a large variation in performance. We believe the reason is that the adaptation in MAML, which decides where to evaluate the gradient, amplifies perturbation.

We further vary c , the number of perturbed examples, and quantify its performance impact. In Figure 2, we present the absolute difference in IC accuracy between perturbed and no perturbation setting over various c (we decided not to look into SL robustness since only **Proto** yields satisfactory results even before perturbation). *Remove* and *replace* perturbation is shown in the left and right panel respectively. Again **Proto** yields the most robust results (in terms of least difference) compared to baselines in explored perturbation settings. Also, utterance *removing* is more challenging for model robustness. Both findings strengthen our argument above.

Lastly, we measure the model robustness against modality mismatch. In Table 3.2, we report the IC accuracy and SL F1 when models are pre-trained and adapted in human transcription while evaluated with ASR hypotheses. Similar to what we found above, **Proto** yields the best performance in both IC and SL. By comparing the results in the condition of mismatched modality reported here with the matched modality counterpart (i.e., no perturbation in Table 3.2), we observe that **Proto** is again the most robust approach in IC (accuracy drop ranging from 0.3 to 2.0 for **Proto**, 3.0 to 4.4 for **Finetune**, and 3.5 to 4.3 for **MAML**). SL result from **Proto** is stable as well (2.2 to 4.3 F1 drop), while **Finetune** and **MAML** yield relatively low F1 scores. Findings here agree with the observation made above for adaptation example missing/replacing, and further support our discussion about the robustness of different learning frameworks.

those obtained by directly computing the difference between the averaged performance reported in the table for the noisy and no perturbation settings.

		IC acc. (absolute acc. difference)			SL F1 (absolute F1 difference)		
		SNIPS-fs	ATIS-fs	TOP-fs	SNIPS-fs	ATIS-fs	TOP-fs
No perturbation	FINETUNE	83.6±0.8	69.9±1.6	57.7±0.7	19.6±0.7	20.1±0.6	15.7±0.6
	FINETUNE-GloVe	81.0±0.7	65.6±1.1	53.1±0.6	18.0±0.4	18.2±0.5	14.7±0.4
	FINETUNE-ELMo	82.1±1.2	68.7±1.9	58.0±1.0	19.3±0.6	19.4±0.8	14.9±0.4
	MAML	87.4±0.8	71.1±1.1	57.6±0.5	33.0±1.3	26.3±0.9	22.9±0.8
	MAML-GloVe	79.4±0.9	65.9±0.9	53.2±0.6	30.1±1.0	24.9±0.8	21.3±0.6
	MAML-ELMo	83.5±0.9	69.6±1.3	54.5±0.8	31.9±1.1	25.4±1.2	23.2±0.9
	Proto	90.9±0.3	75.3±0.7	61.9±1.1	45.4±0.5	42.7±1.6	35.4±0.5
	Proto-GloVe	75.3±0.4	70.1±0.6	52.7±1.3	32.2±0.7	41.5±1.5	30.6±0.6
	Proto-ELMo	87.1±0.5	76.0±0.8	59.8±1.2	43.1±0.6	41.2±1.8	35.7±0.8
Remove 1 utterance per intent	FINETUNE	83.4±0.6 (3.1±3.0)	69.3±1.4 (4.1±3.3)	57.4±0.6 (4.0±2.9)	18.9±0.7 (1.1±1.8)	19.8±0.7 (1.1±1.4)	15.3±0.4 (0.9±0.6)
	MAML	87.3±0.5 (4.3±4.2)	71.3±1.6 (11.5±8.3)	58.8±1.4 (10.7±8.8)	32.6±1.4 (5.2±4.0)	25.9±0.8 (3.8±2.8)	22.3±0.7 (4.4±3.5)
	Proto	90.7±0.3 (0.8±1.6)	75.0±0.6 (0.3±0.6)	61.3±0.9 (1.3±1.4)	44.8±0.6 (1.5±1.5)	42.5±1.5 (0.4±0.5)	35.2±0.5 (0.2±0.3)
	FINETUNE	83.2±0.9 (1.9±2.4)	69.4±1.5 (0.6±1.0)	57.1±0.8 (1.6±1.5)	18.8±0.5 (1.7±1.5)	20.0±0.6 (1.1±0.5)	15.5±0.5 (1.0±0.2)
	MAML	87.5±0.7 (2.0±2.7)	71.0±1.7 (5.3±5.0)	57.5±1.3 (4.9±5.4)	32.8±1.3 (3.2±2.6)	25.6±0.1 (1.8±1.7)	22.8±0.6 (1.5±2.0)
	Proto	90.9±0.4 (0.9±1.8)	75.2±0.6 (0.4±0.7)	62.1±1.0 (1.7±1.7)	45.5±0.6 (1.7±1.6)	42.6±1.5 (0.6±0.6)	35.1±0.5 (0.3±0.4)

Table 2. Average IC accuracy and SL F1 over 100 test episodes for three few-shot datasets in the form of mean \pm standard deviation, computed over three random seeds. We show results for three methods and perturbation settings. In the parentheses, the absolute accuracy and F1 difference on test episodes between *replace* or *remove* and no perturbation counterpart is reported. Model architecture **BERT** is used by default if not specified.




Fig. 2. The absolute IC accuracy difference between no perturbation and perturbed settings with various c , the number of utterances *removed* or *replaced* from the adaptation set.

4. CONCLUSIONS

In this paper, we establish a novel SLU task, the few-shot noisy SLU, with existing public datasets. We further propose a ProtoNets based approach, **Proto**, to build IC and SL classifiers with few noisy examples. When there is no noise in few-

		SNIPS-fs	ATIS-fs	TOP-fs
IC acc.	FINETUNE	79.4±1.8	66.9±1.9	53.3±1.1
	MAML	83.3±0.9	68.1±0.9	54.0±1.0
	Proto	89.0±0.7	75.0±0.7	59.9±0.6
SL F1	FINETUNE	17.5±0.4	19.2±0.2	14.5±0.3
	MAML	26.7±1.6	20.7±0.5	18.8±0.5
	Proto	41.1±0.6	39.7±1.3	33.2±0.6

Table 3. IC accuracy and SL F1 from models pre-trained and adapted in human transcription while evaluated with ASR hypotheses.

shot examples, **Proto** yields better performance than other approaches utilizing MAML and fine-tuning frameworks. **Proto** also achieves the highest and most robust IC accuracy and SL F1 when two types of noise, adaptation example missing/replacing and modality mismatch, are injected in adaption and evaluation set respectively. We believe the ensemble nature of ProtoNets benefits the model robustness, and the simplicity of **Proto**'s model architecture is also helpful in the few-shot noisy scenario. Our contribution here is a step toward the efficient and robust deployment of SLU models. While our results are promising, there is still substantial work, from the creation of few-shot SLU datasets covering more noises to studies of faster and stabler learning algorithms, in pursuit of the goal.

5. REFERENCES

- [1] Gökhan Tür, Jerry Wright, Allen Gorin, Giuseppe Riccardi, and Dilek Hakkani-Tür, “Improving spoken language understanding using word confusion networks,” in *Seventh International Conference on Spoken Language Processing*, 2002.
- [2] Dilek Hakkani-Tür, Frédéric Béchet, Giuseppe Riccardi, and Gokhan Tür, “Beyond asr 1-best: Using word confusion networks in spoken language understanding,” *Computer Speech & Language*, vol. 20, no. 4, pp. 495–514, 2006.
- [3] Kaisheng Yao, Baolin Peng, Yu Zhang, Dong Yu, Geoffrey Zweig, and Yangyang Shi, “Spoken language understanding using long short-term memory neural networks,” in *2014 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2014, pp. 189–194.
- [4] Daniel Guo, Gökhan Tür, Wen-tau Yih, and Geoffrey Zweig, “Joint semantic utterance classification and slot filling with recursive neural networks,” in *2014 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2014, pp. 554–559.
- [5] Grégoire Mesnil, Yann Dauphin, Kaisheng Yao, Yoshua Bengio, Li Deng, Dilek Hakkani-Tür, Xiaodong He, Larry Heck, Gökhan Tür, Dong Yu, et al., “Using recurrent neural networks for slot filling in spoken language understanding,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 3, pp. 530–539, 2014.
- [6] Jin Cao, Jun Wang, Wael Hamza, Kelly Vanee, and Shang-Wen Li, “Style attuned pre-training and parameter efficient fine-tuning for spoken language understanding,” *Proc. Interspeech 2020*, pp. 1570–1574, 2020.
- [7] Cheng-I Lai, Jin Cao, Sravan Bodapati, and Shang-Wen Li, “Towards semi-supervised semantics understanding from speech,” in *Workshop on Self-Supervised Learning for Speech and Audio Processing, NeurIPS*, 2020.
- [8] Andy T Liu, Shang-Wen Li, and Hung-yi Lee, “Tera: Self-supervised learning of transformer encoder representation for speech,” *arXiv preprint arXiv:2007.06028*, 2020.
- [9] Zhilin Yang, Ruslan Salakhutdinov, and William W Cohen, “Transfer learning for sequence tagging with hierarchical recurrent networks,” *arXiv preprint arXiv:1703.06345*, 2017.
- [10] Anjishnu Kumar, Arpit Gupta, Julian Chan, Sam Tucker, Bjorn Hoffmeister, Markus Dreyer, Stanislav Peshterliev, Ankur Gandhe, Denis Filiminov, Ariya Rastrow, et al., “Just ask: building an architecture for extensible self-service spoken language understanding,” *arXiv preprint arXiv:1711.00549*, 2017.
- [11] Sebastian Schuster, Sonal Gupta, Rushin Shah, and Mike Lewis, “Cross-lingual transfer learning for multilingual task oriented dialog,” in *Proceedings of NAACL-HLT*, 2019, pp. 3795–3805.
- [12] Qian Chen, Zhu Zhuo, and Wen Wang, “Bert for joint intent classification and slot filling,” *arXiv preprint arXiv:1902.10909*, 2019.
- [13] Cheng-I Lai, Yung-Sung Chuang, Hung-Yi Lee, Shang-Wen Li, and James Glass, “Semi-supervised spoken language understanding via self-supervised speech and language model pretraining,” *arXiv preprint arXiv:2010.13826*, 2020.
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [15] Chelsea Finn, Pieter Abbeel, and Sergey Levine, “Model-agnostic meta-learning for fast adaptation of deep networks,” in *Proceedings of the 34th International Conference on Machine Learning, JMLR. org*, 2017, vol. 70, pp. 1126–1135.
- [16] Jake Snell, Kevin Swersky, and Richard Zemel, “Prototypical networks for few-shot learning,” in *Advances in neural information processing systems*, 2017, pp. 4077–4087.
- [17] Hongyin Luo, Shang-Wen Li, and James Glass, “Prototypical Q Networks for Automatic Conversational Diagnosis and Few-Shot New Disease Adaption,” in *Proc. Interspeech 2020*, 2020, pp. 3895–3899.
- [18] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al., “Matching networks for one shot learning,” in *Advances in neural information processing systems*, 2016, pp. 3630–3638.
- [19] Jason Krone, Yi Zhang, and Mona Diab, “Learning to classify intents and slot labels given a handful of examples,” in *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI, ACL*, 2020.
- [20] Edwin Simonnet, Sahar Ghannay, Nathalie Camelin, and Yannick Estève, “Simulating ASR errors for training SLU systems,” in *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC)*, 2018.
- [21] Gökhan Tür, Anoop Deoras, and Dilek Hakkani-Tür, “Semantic parsing using word confusion networks with conditional random fields,” in *INTERSPEECH*, 2013.

- [22] Faisal Ladhak, Ankur Gandhe, Markus Dreyer, Lambert Mathias, Ariya Rastrow, and Björn Hoffmeister, “Latticernn: Recurrent neural networks over lattices,” in *Interspeech*, 2016.
- [23] Chao-Wei Huang and Yun-Nung Chen, “Adapting pre-trained transformer to lattices for spoken language understanding,” in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2019.
- [24] Ryo Masumura, Yusuke Ijima, Taichi Asami, Hirokazu Masataki, and Ryuichiro Higashinaka, “Neural confnet classification: Fully neural network based spoken utterance classification using word confusion networks,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 6039–6043.
- [25] Sibel Yaman, Li Deng, Dong Yu, Ye-Yi Wang, and Alex Acero, “An integrative and discriminative technique for spoken utterance classification,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 16, no. 6, pp. 1207–1214, 2008.
- [26] Su Zhu, Ouyu Lan, and Kai Yu, “Robust spoken language understanding with unsupervised asr-error adaptation,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018.
- [27] Chia-Hsuan Lee, Yun-Nung Chen, and Hung-Yi Lee, “Mitigating the impact of speech recognition errors on spoken question answering by adversarial domain adaptation,” in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 7300–7304.
- [28] Chen Li and Yang Liu, “Improving named entity recognition in tweets via detecting non-standard words,” in *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2015, pp. 929–938.
- [29] Charles T Hemphill, John J Godfrey, and George R Doddington, “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*.
- [30] Alice Coucke, Alaa Saade, Adrien Ball, Théodore Bluche, Alexandre Caulier, David Leroy, Clément Doumouro, Thibault Gisselbrecht, Francesco Caltagirone, Thibaut Lavril, et al., “Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces,” *arXiv preprint arXiv:1805.10190*, 2018.
- [31] Sonal Gupta, Rushin Shah, Mrinal Mohit, Anuj Kumar, and Mike Lewis, “Semantic parsing for task oriented dialog using hierarchical representations,” in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018, pp. 2787–2792.
- [32] Sihui Wang, Tom Gunter, and David VanDyke, “On modelling uncertainty in neural language generation for policy optimisation in voice-triggered dialog assistants,” in *2nd Workshop on Conversational AI: Today’s Practice and Tomorrow’s Potential, NeurIPS*, 2018.
- [33] Maryam Fazel-Zarandi, Shang-Wen Li, Jin Cao, Jared Casale, Peter Henderson, David Whitney, and Alborz Geramifard, “Learning robust dialog policies in noisy environments,” in *Workshop on Conversational AI: Today’s Practice and Tomorrow’s Potential, NeurIPS*, 2017.
- [34] Anuj Goyal, Angeliki Metallinou, and Spyros Matsoukas, “Fast and scalable expansion of natural language understanding functionality for intelligent agents,” *arXiv preprint arXiv:1805.01542*, 2018.
- [35] Sepp Hochreiter and Jürgen Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [36] Jeffrey Pennington, Richard Socher, and Christopher D Manning, “Glove: Global vectors for word representation,” in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 1532–1543.
- [37] Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer, “Deep contextualized word representations,” in *Proceedings of NAACL-HLT*, 2018, pp. 2227–2237.