

Novel Slot Detection: A Benchmark for Discovering Unknown Slot Types in the Task-Oriented Dialogue System

Yanan Wu^{1*}, Zhiyuan Zeng^{1*}, Keqing He^{2*}, Hong Xu¹
Yuanmeng Yan¹, Huixing Jiang², Weiran Xu^{1*}

¹Pattern Recognition & Intelligent System Laboratory

¹Beijing University of Posts and Telecommunications, Beijing, China

²Meituan Group, Beijing, China

{yanan.wu, zengzhiyuan, xuhong, yanyuanmeng, xuweiran}@bupt.edu.cn
{hekeqing, jianghuixing}@meituan.com

Abstract

Existing slot filling models can only recognize pre-defined in-domain slot types from a limited slot set. In the practical application, a reliable dialogue system should know what it does not know. In this paper, we introduce a new task, Novel Slot Detection (NSD), in the task-oriented dialogue system. NSD aims to discover unknown or out-of-domain slot types to strengthen the capability of a dialogue system based on in-domain training data. Besides, we construct two public NSD datasets, propose several strong NSD baselines, and establish a benchmark for future work. Finally, we conduct exhaustive experiments and qualitative analysis to comprehend key challenges and provide new guidance for future directions¹.

1 Introduction

Slot filling plays a vital role to understand user queries in personal assistants such as Amazon Alexa, Apple Siri, Google Assistant, etc. It aims at identifying a sequence of tokens and extracting semantic constituents from the user queries. Given a large scale pre-collected training corpus, existing neural-based models (Mesnil et al., 2015; Liu and Lane, 2015, 2016; Goo et al., 2018; Haihong et al., 2019; Chen et al., 2019; He et al., 2020b,d; Yan et al., 2020; Louvan and Magnini, 2020; He et al., 2020a) have been actively applied to slot filling and achieved promising results.

Existing slot filling models can only recognize pre-defined entity types from a limited slot set, which is insufficient in the practical application scenario. A reliable slot filling model should not only predict the pre-defined slots but also detect potential unknown slot types to know what it doesn't

*The first three authors contribute equally. Weiran Xu is the corresponding author.

¹<https://github.com/ChestnutWYN/ACL2021-Novel-Slot-Detection>

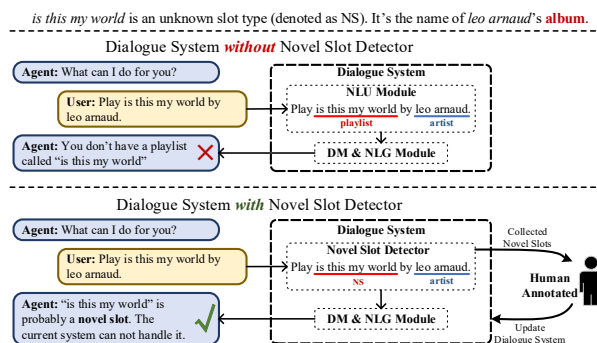


Figure 1: An example of Novel Slot Detection in the task-oriented dialogue system. Without NSD, the dialogue system gives the wrong response since it misunderstands the unknown slot “is this my world” as the in-domain *playlist* type. In contrast, NSD recognizes “is this my world” as *NS* and the system gives a fallback response. Meanwhile, with human-in-the-loop annotation, the system can increase its functions or skills.

know, which we call Novel Slot Detection (NSD) in this paper. NSD is particularly crucial in deployed systems—both to avoid performing the wrong action and to discover potential new entity types for future development and improvement. We display an example as Fig 1 shows.

In this paper, we define *Novel Slot* (NS) as new slot types that are not included in the pre-defined slot set. NSD aims to discover potential new or out-of-domain entity types to strengthen the capability of a dialogue system based on in-domain pre-collected training data. There are two aspects in the previous work related to NSD, out-of-vocabulary (OOV) recognition (Liang et al., 2017a; Zhao and Feng, 2018; Hu et al., 2019; He et al., 2020c,d; Yan et al., 2020; He et al., 2020e) and out-of-domain (OOD) intent detection (Lin and Xu, 2019; Larson et al., 2019; Xu et al., 2020a; Zeng et al., 2021b,a). OOV means many slot types can have a large number of new slot values while the training set only obtains a tiny part of slot values. OOV aims to recognize unseen slot values in training set

Utterance	play	is	this	my	world	by	leo	arnaud
Slot Filling Labels	O	B-album	I-album	I-album	I-album	O	B-artist	I-artist
Novel Slot Detection Labels	O	NS	NS	NS	NS	O	B-artist	I-artist

Table 1: Comparison between slot filling and novel slot detection. In the novel slot detection labels, we consider “album” as an unknown slot type that is out of the scope of the pre-defined slot set. Meanwhile, “artist” belonging to in-domain slot types still needs to be recognized as the original slot filling task.

for pre-defined slot types, using character embedding (Liang et al., 2017a), copy mechanism (Zhao and Feng, 2018), few/zero-shot learning (Hu et al., 2019; He et al., 2020e; Shah et al., 2019), transfer learning (Chen and Moschitti, 2019; He et al., 2020c,b) and background knowledge (Yang and Mitchell, 2017; He et al., 2020d), etc. Compared to OOV recognition, our proposed novel slot detection task focuses on detecting unknown slot types, not just unseen values. NSD faces the challenges of both OOV and no sufficient context semantics (see analysis in Section 6.2), greatly increasing the complexity of the task. Another line of related work is OOD intent detection (Hendrycks and Gimpel, 2017; Lee et al., 2018; Lin and Xu, 2019; Ren et al., 2019; Zheng et al., 2020; Xu et al., 2020a) which aims to know when a query falls outside the range of predefined supported intents. The main difference is that NSD detects unknown slot types in the token level while OOD intent detection identifies out-of-domain intent queries. NSD requires a deep understanding of the query context and is prone to label bias of O (see analysis in Section 5.3.1), making it challenging to identify unknown slot types in the task-oriented dialog system.

In this paper, we first introduce a new and important task, Novel Slot Detection (NSD), in the task-oriented dialogue system (Section 2.2). NSD plays a vital role in avoiding performing the wrong action and discovering potential new entity types for the future development of dialogue systems. Then, we construct two public NSD datasets, Snips-NSD and ATIS-NSD, based on the original slot filling datasets, Snips (Coucke et al., 2018) and ATIS (Hemphill et al., 1990) (Section 2.2). From the perspective of practical application, we consider three kinds of dataset construction strategies, Replace, Mask and Remove. Replace denotes we label the novel slot values with all O in the training set. Mask is to label with all O and mask the novel slot values. Remove is the most strict strategy where all the queries containing novel slots are removed. We dive into the details of the three different construction strategies in Section 3.2 and perform a qualitative analysis in Section 5.3.1. Be-

sides, we propose two kinds of evaluation metrics, span-level F1 and token-level F1 in Section 3.4, following the slot filling task. Span F1 considers the exact matching of a novel slot span while Token F1 focuses on prediction accuracy on each word of a novel slot span. We discuss performance comparison between the two metrics and propose a new metric, restriction-oriented span evaluation (ROSE), to combine the advantages of both in Section 5.3.3. Then, we establish a fair benchmark and propose extensive strong baselines for NSD in Section 4. Finally, we perform exhaustive experiments and qualitative analysis to shed light on the challenges that current approaches faced with NSD in Section 5.3 and 6.

Our contributions are three-fold: (1) We introduce a Novel Slot Detection (NSD) task in the task-oriented dialogue system. NSD helps avoid performing the wrong action and discovering potential new entity types for increasing functions of dialogue systems. (2) We construct two public NSD datasets and establish a benchmark for future work. (3) We conduct exhaustive experiments and qualitative analysis to comprehend key challenges and provide new guidance for future NSD work.

2 Problem Formulation

2.1 Slot Filling

Given a sentence $X = \{x_1, \dots, x_n\}$ with n tokens, the slot filling task is to predict a corresponding tag sequence $Y = \{y_1, \dots, y_n\}$ in BIO format, where each y_i can take three types of values: B-slot_type, I-slot_type and O , where “B” and “I” stand for the beginning and intermediate word of a slot and “O” means the word does not belong to any slot. Here, slot filling assumes $y_i \in \mathbf{y}$, where \mathbf{y} denotes a pre-defined slot set of size \mathcal{M} . Current approaches typically model slot filling as a sequence labeling problem using RNN (Liu and Lane, 2015, 2016; Goo et al., 2018) or pre-trained language models (Chen et al., 2019).

2.2 Novel Slot Detection

We refer to the above training data D as in-domain (IND) data. Novel slot detection aims to identify

Original Utterance		play	is	this	my	world	by	leo	arnaud
Original Slot Filling Labels		O	B-album	I-album	I-album	I-album	O	B-artist	I-artist
Strategy	Replace	play	is	this	my	world	by	leo	arnaud
		O	O	O	O	O	O	B-artist	I-artist
	Mask	play	MASK	MASK	MASK	MASK	by	leo	arnaud
		O	O	O	O	O	O	B-artist	I-artist
	Remove	-	-	-	-	-	-	-	-
		-	-	-	-	-	-	-	-

Table 2: Comparison between three processing strategies in the training set. We consider “album” as an unknown slot type and “-” denotes the sentence is removed from the training data.

unknown or out-of-domain (OOD) slot types via IND data while correctly labeling in-domain data. We denote unknown slot type as *NS* and in-domain slot types as IND in the following sections. Note that we don’t distinguish between B-NS and I-NS and unify them as *NS* because we empirically find existing models hardly discriminate B and I for an unknown slot type. We provide a detailed analysis in Section 5.3.3. We show an example of NSD in Table 1. The challenges of recognizing NSD come from two aspects, *O* tags and in-domain slots. On the one hand, models need to learn entity information for distinguishing *NS* from *O* tags. On the other hand, they require discriminating *NS* from other slot types in the pre-defined slot set. We provide a detailed error analysis in Section 6.1.

3 Dataset

Since there are not existing NSD datasets, we construct two new datasets based on the two widely used slot filling datasets, Snips (Coucke et al., 2018) and ATIS (Hemphill et al., 1990). We first briefly introduce Snips and ATIS, then elaborate on data construction and processing in detail, and display the statistic of our NSD datasets, Snips-NSD and ATIS-NSD. Finally, we define two evaluation metrics for the NSD task, Span F1 and Token F1.

3.1 Original Slot Filling Datasets

Snips² is a custom intent engine dataset. It originally has 13,084 train utterances, 700 and 700 test utterances. ATIS³ contains audio recordings of people making flight reservations. It originally has 4,478 train utterances, 500 dev and 893 test utterances. The full statistic is shown in Table 3. Note that the vocabulary only contains words in the training set, and test set words that do not exist in the vocabulary are referred to OOV words. The percentage of OOV words represents the portion of OOV words in the test set.

²<https://github.com/sonos/nlu-benchmark/tree/master/2017-06-custom-intent-engines>

³<https://github.com/yvchen/JointSLU/tree/master/data>

	Snips	ATIS
Vocabulary Size	11,241	722
Percentage of OOV words	5.95%	0.77%
Number of Slots	39	79
Training Set Size	13,084	4,478
Development Set Size	700	500
Testing Set Size	700	893

Table 3: Statistics of ATIS and Snips datasets.

3.2 Data Construction and Processing

For Snips and ATIS datasets, we keep some slot classes in training as unknown and integrate them back during testing, following (Fei and Liu, 2016; Shu et al., 2017; Lin and Xu, 2019). We randomly select part of slot types in Snips and ATIS as unknown slots (5%, 15%, and 30% in this paper). Note that the original train/val/test split is fixed. Considering class imbalance, we perform weighted sampling where the chosen probability is relevant to the number of class examples similar to (Lin and Xu, 2019). To avoid randomness of experiment results, we report the average result over 10 runs.

After we choose the unknown slot types, a critical problem is how to handle sentences including these unknown slot types in training set. For OOD intent detection, we just need to remove these sentences in training and validation set. However, for Novel Slot Detection, a sentence perhaps contains both in-domain slots and unknown slots, which is nontrivial for tackling unknown slots at the token level. We need to balance the performance of recognizing unknown slots and in-domain slots. Therefore, we propose three different processing strategies as follows: (1) **Replace**: We label the unknown slot values with all *O* in the training set while the original values remain unchanged. (2) **Mask**: We label the unknown slot values with all *O* and mask these slot values with a special token *MASK*. (3) **Remove**: All the sentences containing unknown slots are directly removed.

We display examples of the above three strategies in Table 2. For the val and test set, we just label the unknown slot values with all *NS* while keeping the in-domain labeling fixed. Note that *NS*

Snips-NSD-15%	Train	Val	Test
number of in-domain slots	33	33	33
number of unknown slots	6	6	6
percentage of OOV words	-	-	8.51%
number of queries	9,329	700	700
number of queries including unknown slots	0	192	202
number of slot values	23,176	1,794	1,790
number of unknown slot values	0	210	220

Table 4: The detailed statistics of Snips-NSD-15%.

tags only exist in the val and test set, not in the training set. Besides, we keep original in-domain slots fixed to evaluate the performance of both *NS* and in-domain slots. We aim to simulate the practical scenario where we can hardly know what unknown slots are. These three strategies all have its practical significance. Compared with others, Remove is the most suitable strategies for real-world scenarios. In practical scenario, dialog systems first train in the data set labeled by human annotators, and then applied to the actual application. In the process of interaction with the real users, novel slot types appear gradually. Therefore, we consider that the training set doesn't contain potential novel slots sentences. In other words, Remove is the most suitable strategy for NSD in real applications. What's more, Section 5.3.1 demonstrates Remove performs best while the others suffer from severe model bias by *O* tags. Therefore, we adopt Remove as the main strategy in this paper.

3.3 Statistic of New NSD Datasets

Table 4 shows the detailed statistics of Snips-NSD-15% constructed by Remove strategy, where we choose 15% classes in the training data as unknown slots.⁴ Combining Table 3 and Table 4, we can find Remove strategy removes 28.70% of queries in the original Snips training set, hence increases the percentage of OOV word from 5.95% to 8.51%. And unknown slot values account for 12.29% of total slot values in the test set.

3.4 Metrics

The traditional slot filling task uses Span F1⁵ for evaluation. Span F1 considers the exact span matching of an unknown slot span. However, we find in Section 5.3.3 that this metric is too strict to NSD

⁴Since different proportions of unknown slots have different statistics, here we only display the results of Snips-NSD-15% for brevity.

⁵<https://www.clips.uantwerpen.be/conl12000/chunking/conlleval.txt>

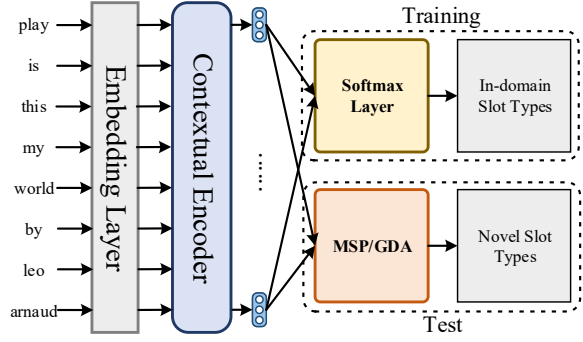


Figure 2: The overall architecture of our approach.

models. In the practical application, we only need to coarsely mine parts of words of unknown slots, then send these queries containing potential unknown slot tokens to human annotators, which has effectively reduced extensive labor and improved efficiency. Therefore, we define a more reasonable metric, Token F1 which focuses on the word-level matching of a novel slot span. We also propose a new metric, Restriction-Oriented Span Evaluation (ROSE), for a fair comparison in Section 5.3.3.

4 Methodology

In this section, we introduce the NSD models proposed in this paper and illustrate the differences between the various parallel approaches during the training and test stage.

4.1 Overall Framework

The overall structure of model is shown in Fig 2. In the training stage, we either train a multiple-class classifier or binary classifier using different training objectives. We use public BERT-large (Devlin et al., 2019) embedding layer and BiLSTM-CRF (Huang et al., 2015) for token level feature extraction. Then, in the test stage, we use the typical neural multiple classifier to predict the in-domain slot labels. Meanwhile, we use the detection algorithm, MSP or GDA to figure out novel slot tokens. Finally, we override the slot token labels which are detected as NS. In terms of training objectives, detection algorithms, and distance strategies, we compare different variants as follows.

Training objective. For in-domain slots, we propose two training objectives. Multiple classifier refers to the traditional slot filling objective setting, which performs token-level multiple classifications on the BIO tags (Ratinov and Roth, 2009) combined with different slots. Binary classifier unifies all non-O tags into one class, and the model makes

Models			5%			15%			30%		
			IND	NSD		IND	NSD		IND	NSD	
detection method	objective	distance strategy	Span F1	Span F1	Token F1	Span F1	Span F1	Token F1	Span F1	Span F1	Token F1
MSP	binary	-	87.21	12.34	25.16	71.44	12.31	39.50	58.88	8.73	40.38
	multiple	-	88.05	14.04	30.50	79.71	20.97	40.02	78.52	25.26	46.91
	binary+multiple	-	89.59	23.58	37.55	83.72	24.70	45.32	79.08	30.66	52.10
GDA	binary	difference	87.95	23.83	35.83	83.65	22.06	43.99	78.72	32.50	44.13
	binary	minimum	61.29	10.36	17.08	49.11	16.91	31.10	48.07	15.56	33.78
	multiple	difference	93.14	29.73	45.99	90.07	31.96	53.02	85.56	36.16	54.55
	multiple	minimum	93.10	31.67*	46.97*	90.18	32.19	53.75*	86.26*	38.64*	55.24*

Table 5: IND and NSD results with different proportions (5%, 15% and 30%) of classes are treated as unknown slots on Snips-NSD. * indicates the significant improvement over all baselines ($p < 0.05$).

Models			5%			15%			30%		
			IND	NSD		IND	NSD		IND	NSD	
detection method	objective	distance strategy	Span F1	Span F1	Token F1	Span F1	Span F1	Token F1	Span F1	Span F1	Token F1
MSP	binary	-	92.04	19.73	29.63	91.74	23.40	33.89	80.49	21.88	39.17
	multiple	-	94.33	27.15	31.16	92.54	39.88	42.29	87.63	40.42	47.64
	binary+multiple	-	94.41	32.49	43.48	93.29	41.23	43.13	90.14	41.76	51.87
GDA	binary	difference	93.69	27.02	34.21	92.13	30.51	36.30	88.73	30.91	45.64
	binary	minimum	93.57	15.90	20.96	90.98	24.53	27.26	88.21	26.40	39.83
	multiple	difference	95.20	47.78*	51.54*	93.92	50.92*	52.24*	92.02	51.26*	56.59*
	multiple	minimum	95.31*	41.74	45.91	93.88	43.78	46.18	91.67	45.44	52.37

Table 6: IND and NSD results with different proportions (5%, 15% and 30%) of classes are treated as unknown slots on ATIS-NSD. * indicates the significant improvement over all baselines ($p < 0.05$).

a token-level binary classification of O or non-O on the sequence. Note that in the test stage, for in-domain prediction, we both use the multiple classifier. While, for novel slot detection, we use the multiple classifier, or the binary classifier, or both of them. In Table 5 and Table 6, binary+multiple means the token will be labeled as NS only if both classifiers predict it as NS.

Detection algorithm. MSP and GDA are detection algorithms in the test stage. MSP (Maximum Softmax Probability) (Hendrycks and Gimpel, 2017) applies a threshold on the maximum softmax probability, if the maximum falls below the threshold, the token will be predicted to be a novel slot token. GDA (Gaussian Discriminant Analysis) (Xu et al., 2020a) is a generative distance-based classifier for out-of-domain detection with Euclidean space. We treat tokens not belonging to any in-domain slots (including O) as novel slot tokens for both methods. For example, with a binary classifier, if the softmax probabilities belonging to O or non-O are both lower than an MSP threshold, then the token is labeled as NS.

Distance strategy. The GDA detection is based on the distances between a target and each slot representation cluster. In original GDA, when the minimum distance is greater than a certain threshold, it is predicted to be novel slots. We propose a novel strategy named Difference, which uses the maximum distance minus the minimum distance, when the difference value of a target is less than a threshold, it is predicted as novel slots. Both

of their thresholds are obtained by optimizing the NSD metrics on the validation set.

5 Experiment and Analysis

5.1 Implementation Details

We use the public pre-trained Bert-large-uncased model to embed tokens which has 24 layers, 1024 hidden states, 16 heads and 336M parameters. The hidden size for the BiLSTM layer is set to 128. Adam is used for optimization with an initial learning rate of $2e-5$. The dropout value is fixed as 0.5, and the batch size is 64. We train the model only on in-domain labeled data. The training stage has an early stopping setting with patience equal to 10. We use the best F1 scores on the validation set to calculate the MSP and GDA thresholds adaptively. Each result of the experiments is tested for 10 times under the same setting and reports the average value. The training stage of our model lasts about 28 minutes on single Tesla T4 GPU(16 GB of memory).

5.2 Main Results

Table 5 and 6 show the experiment results with seven different models on two benchmark slot filling datasets Snips-NSD and ATIS-NSD constructed by Remove strategy. We both report NSD and IND results using Span F1 and Token F1. We compare these models from three perspectives, detection method, objective and distance strategy in the following. The analysis of effect of the propor-

Strategy	5%			15%			30%		
	IND	NSD		IND	NSD		IND	NSD	
	Span	Span	Token	Span	Span	Token	Span	Span	Token
Replace	94.52	1.93	5.27	94.33	0.66	2.29	94.02	0.27	0.82
Mask	90.08	23.10	37.91	86.52	25.07	45.92	83.37	32.14	50.68
Remove	93.10	31.67	46.97	90.18	32.19	53.75	86.26	38.64	55.24

Table 7: Comparison between different data processing strategies on Snips-NSD using GDA+Multiple+Minimum.




Figure 3: Effect of the proportion of unknown slot types.

tion of unknown slot types is described in 5.3.2.

Detection Method: MSP vs GDA. Under the same setting of objective, GDA performs better than MSP in both IND and NSD, especially in NSD. We argue that GDA models the posterior distribution on representation spaces of the feature extractor and avoids the issue of overconfident predictions (Guo et al., 2017; Liang et al., 2017b, 2018). Besides, comparing Snips-NSD and ATIS-NSD, NSD Token F1 scores on ATIS-NSD are much higher than Snips-NSD but no significant difference exists for NSD Span F1 scores. The reason is that Snips-NSD has a higher average entity length (1.83) than ATIS-NSD (1.29), making it harder to detect the exact *NS* span.

Objective: Binary vs Multiple. Under all settings, Multiple outperforms Binary with a large margin on two datasets in both IND and NSD metrics. For MSP, combining Multiple and Binary get higher F1 scores. Specifically, the Binary classifier is used to calculate the confidence of a token belonging to non-*O* type, which can judge whether the token belongs to entities and distinguish *NS* from type *O*. On the other hand, we use the Multiple classifier to calculate the confidence for tokens that are of type *NS*, to distinguish *NS* from all predefined non-*O* slot types. For GDA, we do not combine Multiple and Binary because of poor performance. Multiple achieves the best results for all the IND and NSD F1 scores. We suppose multi-class classification can better capture semantic features than binary classification.

Distance Strategy: Minimum vs Difference. We

find under the same setting of Binary, Difference strategy outperforms Minimum on both datasets for NSD metrics. But under the same setting of Multiple, there is no consistent superiority between the two distance strategies. For example, Difference outperforms Minimum for NSD metrics on ATIS-NSD, opposite to the results on Snips-NSD. We argue different distance strategies are closely related to objective settings and dataset complexity. We will leave the theoretical analysis to the future.

5.3 Qualitative Analysis

5.3.1 Effect of Different Data Processing Strategies

Table 7 displays IND and NSD metrics of three different dataset processing strategies on Snips-NSD using the same model GDA+Multiple+Minimum. In this section, we will dive into the analysis of the effects of different data processing strategies. Results show the Replace strategy gets poor performance in NSD, which proves labeling unknown slots as *O* tags will severely mislead the model. The Mask and Remove strategies are more reasonable since they remove unknown slots from the training data. Their main difference is that Mask only deletes token-level information, while Remove even eliminates the contextual information. For NSD in all datasets, Remove gains significantly better performance on both Token F1 and Span F1 than Mask by 9.06%(5%), 7.83%(15%) and 4.56%(30%) on Token F1, and 8.57%(5%), 7.12%(15%) and 6.5%(30%) on Span F1. We argue the remaining context is still misleading even if the novel slot tokens are not directly trained in the Mask strategy. Besides, Mask does not conform to the real NSD scenario. Generally, Remove is the most suitable strategy for NSD in real applications and can achieve the best performance.

5.3.2 Effect of the Proportion of Unknown Slot Types

Fig 3 displays the effect of the proportion of unknown slot types using the Remove strategy in GDA+Multiple+Minimum. Results show that with the increase of the proportion of unknown slot types, the NSD F1 scores get improvements while IND F1 scores decrease. We suppose fewer in-domain slot types help the model distinguish unknown slots from IND slots, thus NSD F1 scores get improvements. However, for in-domain slot detection, since Remove deletes all the sentences containing unknown slots in the training data, our




Figure 4: Effect of varying degrees of restrictions

	GDA+mul.+min.	MSP+bin.+mul.
ROSE-mean	40.73	34.71
ROSE-100%	40.39	33.74
ROSE-50%	41.00	35.46

Table 8: ROSE metrics on Snips-NSD using GDA+Multiple+Minimum and MSP+Binary+Multiple

models suffer from the lack of sufficient context to recognize IND slots so IND F1 scores decrease.

5.3.3 New Metric: ROSE

The previous results have shown Span F1 is much lower than the token F1. The reason is that Span F1 is a strict metric, where the model needs to correctly predict all *NS* tokens and the correct boundary. This is difficult for NSD models due to the lack of supervised information. In fact, NSD models only need to mark some tokens in the span of novel slots and send the total sequence containing the *NS* tokens back to the humans. A small number of token omissions or misjudgments are acceptable. Therefore, to meet a reasonable NSD scenario, we propose a new metric, restriction-oriented span evaluation (ROSE), to evaluate the span prediction performance under different restrictions. First, we do not punish the situation where tokens prediction exceeds the span. Then, we consider a span is correct when the number of correctly predicted tokens is greater than a settable proportion p of the span length. We take the average of the ROSE score and the original span F1 to avoid the model obtaining an outstanding result through over-long prediction. The results using Snips with 15% of novel slots are shown in Figure 4. As the degree of restriction increases, the metrics tend to decline. It indicates that the model can mostly identify more than half

Type	Proportion(%)	Span Length	Token F1	Span F1	
top 5	Object_name	21.42	3.71	55.64	20.82
	TimeRange	15.29	2.35	53.65	30.15
	Entity_name	23.14	3.09	48.56	22.83
	Music_item	14.86	1.05	46.23	34.59
	Artist	15.29	2.05	45.26	26.36
bottom 5	City	8.57	1.32	18.72	15.85
	Country	6.29	1.57	14.19	11.11
	State	5.54	1.10	13.55	10.83
	Best_rating	6.14	1.00	11.04	11.04
	Year	3.43	1.00	10.24	10.24

Table 9: Results of single unknown slot.

Type 1	Type 2	Token F1	Span F1
Object_name	-	55.64	20.82
TimeRange	-	53.65	30.15
Party_size_number	-	33.44	28.57
City	-	18.72	15.85
State	-	13.55	10.83
Object_name	TimeRange	53.88	23.37
Object_name	Party_size_number	52.81	22.35
Object_name	City	57.92	21.42
Object_name	State	56.32	19.27
TimeRange	Party_size_number	71.27*	51.03*
City	State	29.33*	27.14*

Table 10: Results of combining multiple unknown slots. * denotes that NSD performance of the combination of two unknown slots is significantly better than each single slot.

of the tokens in spans. To make a comprehensive evaluation, we defined the ROSE-mean, namely the mean of ROSE-25%, ROSE-50%, ROSE-75%, and ROSE-100%. We present results on part of proposed models in Table 8.

5.3.4 Analysis of Single Unknown Slot

To analyze the relationship between NSD performance and a single specific slot, we calculate the token and span metrics treating each single slot type as an unknown slot and show the results of the top five and bottom five for Token F1 scores in Table 9. We find that the slots with better performance often account for a larger percentage of the data set, such as Object_name or Entity_name. They also tend to have a larger value space, such as TimeRange, Music_item, or Artist. These characteristics allow the semantic representation of these slots to be distributed over a large area rather than clustered tightly together. We consider that this distribution is more reasonable because in a real application scenario, novel slots are diverse and its distribution tends to be diffuse. Performance on these types also proves that the NSD models we propose can be better generalized to a reasonable data setting.

NSD error proportion(%)	O	Open vocabulary slots	Other slots	Sum
Prediction is NS	17.79	18.84	9.07	45.70
Target is NS	18.47	7.54	28.29	54.30
Sum	36.26	26.38	37.36	100.00

Table 11: Relative proportions of several types of errors.

Error type	NS	Example
NS to O	movie_name (m_name)	text: when will paris by night aired true: O O B-m_name I-m_name I-m_name O predict: O O NS O NS O
NS to open slot	album	text: play the insoc ep true: O B-album I-album I-album predict: O B-object_name I-object_name NS
NS to other slot	artist	text: play kurt cobain ballad tunes true: O B-artist I-artist B-music_item O predict: O B-genre I-genre B-music_item O
O to NS	artist	text: the workout playlist needs more chris cross true: O B-playlist O O O B-artist I-artist predict: O B-playlist O O NS NS NS
open slots to NS	object_type	text: tell me the actors of the saga awards true: O O O B-object_name O O B-object_type O predict: O O O NS O O NS O
other slots to NS	city	text: what is the weather of east portal ks true: O O O O O B-city I-city B-state predict: O O O O O NS NS NS

Table 12: Error case from NSD prediction.

5.3.5 Analysis for Relationship of Multiple Unknown Slots

In order to explore the effect of inter-slot relationships on NSD, we conducted experiments in which two types are mixed as novel slots. Some of the results are shown in Table 10. In the five types shown in the table, Object_name is an open vocabulary slot with a wide range of values and contains many OOV tokens, TimeRange and Party_size_number often contain numbers, City and State are usually similar in semantics and context. We found that when the other types combined with Object_name, NSD performance is often maintained close to treat Object_name as a novel slot alone. The reason, on the one hand, is that the proportion of other types in the dataset is relatively small, so the overall impact on the metrics is smaller. On the other hand, due to the large semantic distribution range of the open vocabulary slot, there is a latent inclusion relationship for other types, so the mixing of a single type tends to have a slight impact on the NSD performance. We also found that the appropriate combination can significantly improve the efficiency of NSD. Such as TimeRange with Party_size_number, or City with State. This indicates that when the novel slot is similar to the in-domain slot, the model tends to predict the novel slot as a similar slot, which leads to errors. When both are treated as novel slots, these errors can be mitigated.

6 Discussion

In this section, we empirically divide all the error samples into three categories. Each type of problem contains two aspects, corresponding to NSD precision and recall, respectively. We present the relative proportions of several types of errors in Table 11, which using Snips dataset with 5% novel slots on GDA+multiple+minimum model. For each error type, we present an example in Table 12 to describe the characteristics and analyze the causes. Then, we dive into identifying the key challenges and finally proposed possible solutions for future work.

6.1 Error Analysis

Tag O. Tag *O* is the largest and most widely distributed type in the dataset, and it generally refers to the independent function tokens. Therefore, when identifying, it is easy to be confused with other types, and the confusion is more serious for novel slots without supervised learning. We observed that tokens with *O* label detected as novel slots usually exist near spans, and the function words in the span labeled as a novel slot have a probability of being predicted as *O*. We consider that this kind of problem is related to the context. Although the processing strategy of Remove can effectively reduce the misleading of *O* for the novel slots, tag *O* will still be affected by context information of other in-domain slots.

Open Vocabulary Slots. We observe that a large number of novel slot tokens are mispredicted as open vocabulary slots, while the reverse situation is much less likely to happen. This indicates that in Snips, open vocabulary slots tend to overlap or contain most other slots semantically. Even in traditional slot filling tasks, open vocabulary slots are often confused with other slots. We demonstrate this hypothesis in the analysis. Section 5.3.5 shows that NSD performs better when open vocabulary slots are treated as novel slots, and Section 5.3.4 shows that there is no significant performance change when open vocabulary slots are mixed with some semantically concentrated slots. The reason for this problem is that the definition of the dataset is not reasonable. Slots with a large value range can hardly help the personal assistant to give an appropriate reply, and the supervised information of these slots is usually incomplete.

Similar Slots. Except for the two cases mentioned above, predicting novel slots as other in-domain

slots is the most common type of error, in which similar slots account for a large part of it. Due to the overlap between vocabulary or shared similar context, the model often tend to be overconfident to predict similar slot labels, we analyze the phenomenon in Table 10, when similar types is treated as a new slot at the same time, NSD efficiency will rise significantly. We employ a generative classification method GDA, compared with the traditional MSP method, to make full use of data features and alleviate the problem.

6.2 Challenges

Based on the above analysis, we summarize the current challenges faced by the NSD task:

Function tokens. Articles, prepositions, and so on that act as connective words in a sequence. It is usually labeled with type *O*, but also found in some long-span slots, such as *Movie_name*. It can lead to confusion between *O* and novel slot when this kind of slot is the target of NSD.

Insufficient context. Correct slot detection often depends on the context, and this supervised information is missing for novel slots. Models can only conduct NSD to tokens using the original embeddings or representations trained in other contexts, which can lead to bias in the semantic modeling of the novel slot.

Dependencies between slots. There are some semantic overlaps or inclusion relationships in the slot definition of the current benchmark slot filling datasets. As a result, the semantic features are not sufficiently discriminative, and thus some outliers tokens in in-domain slots are easily confused with the novel slots.

Open vocabulary slots. Open vocabulary slots is a special kind of slot, its definition is usually macroscopic and can be further divided, the value range is broad. The representation distribution for Open vocabulary slots tends to be diffuse and uneven, which can be misleading to NSD.

6.3 Future Directions

For tag *O*, a possible solution is to use a binary model to assist identification between *O* and non-*O* function tokens, we provide a simple method in this paper and leave further optimizing to future work. Then, to decouple the dependencies between slots, it is critical to learn more discriminative features for in-domain data, using contrastive learning or prototypical network is expected to help. Besides, in the traditional slot filling task, the open

vocabulary slot problem has been researched for a long time, and accumulate many achievements. Adaptive combination and improvement of relevant methods with NSD tasks is also an important direction of our future research.

7 Related Work

OOV Recognition OOV aims to recognize unseen slot values in training set for pre-defined slot types, using character embedding (Liang et al., 2017a), copy mechanism (Zhao and Feng, 2018), few/zero-shot learning (Hu et al., 2019; Shah et al., 2019), transfer learning (Chen and Moschitti, 2019; He et al., 2020c) and background knowledge (Yang and Mitchell, 2017; He et al., 2020d), etc. Our proposed NSD task focuses on detecting unknown slot types, not just unseen values.

OOD Intent Detection Lee et al. (2018); Lin and Xu (2019); Xu et al. (2020a) aim to know when a query falls outside the range of predefined supported intents. Generally, they first learn discriminative intent representations via in-domain (IND) data, then employs detecting algorithms, such as Maximum Softmax Probability (MSP) (Hendrycks and Gimpel, 2017), Local Outlier Factor (LOF) (Lin and Xu, 2019), Gaussian Discriminant Analysis (GDA) (Xu et al., 2020b) to compute the similarity of features between OOD samples and IND samples. Compared to our proposed NSD, the main difference is that NSD detects unknown slot types in the token level while OOD intent detection identifies sentence-level OOD intent queries.

8 Conclusion

In this paper, we defined a new task, Novel Slot Detection(NSD), then provide two public datasets and establish a benchmark for it. Further, we analyze the problems of NSD through multi-angle experiments and extract the key challenges of the task. We provide some strong models for these problems and offer possible solutions for future work.

Acknowledgements

This work was partially supported by National Key R&D Program of China No. 2019YFF0303300 and Subject II No. 2019YFF0303302, DOCOMO Beijing Communications Laboratories Co., Ltd, MoE-CMCC "Artificial Intelligence" Project No. MCM20190701.

Broader Impact

Dialog systems have demonstrated remarkable performance across a wide range of applications, with the promise of a significant positive impact on human production mode and lifeway. The first step of the dialog system is to identify users' key points. In practical industrial scenario, users may make unreasonable queries which fall outside of the scope of the system-supported slot types. Previous dialogue systems will ignore this problem, which will lead to wrong operations and limit the system's development. In this paper, we firstly propose to detect not only pre-defined slot types but also potential unknown or out-of-domain slot types using MSP and GDA methods. According to exhaustive experiments and qualitative analysis, we also discuss several major challenges in Novel Slot Detection for future work. The effectiveness and robustness of the model are significantly improved by adding Novel Slot Detection, which takes a step towards the ultimate goal of enabling the safe real-world deployment of dialog systems in safety-critical domains. The experimental results have been reported on standard benchmark datasets for considerations of reproducible research.

References

- Lingzhen Chen and Alessandro Moschitti. 2019. Transfer learning for sequence labeling using source model and target data. *ArXiv*, abs/1902.05309.
- Qian Chen, Zhu Zhuo, and Wen Wang. 2019. Bert for joint intent classification and slot filling. *arXiv preprint arXiv:1902.10909*.
- A. Coucke, A. Saade, Adrien Ball, Théodore Bluche, A. Caulier, D. Leroy, Clément Doumouro, Thibault Gisselbrecht, F. Caltagirone, Thibaut Lavril, Maël Primet, and J. Dureau. 2018. Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces. *ArXiv*, abs/1805.10190.
- J. Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*.
- Geli Fei and B. Liu. 2016. Breaking the closed world assumption in text classification. In *HLT-NAACL*.
- Chih-Wen Goo, Guang Gao, Yun-Kai Hsu, Chih-Li Huo, Tsung-Chieh Chen, Keng-Wei Hsu, and Yun-Nung Chen. 2018. Slot-gated modeling for joint slot filling and intent prediction. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 753–757.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. On calibration of modern neural networks. In *ICML*.
- E Haihong, Peiqing Niu, Zhongfu Chen, and Meina Song. 2019. A novel bi-directional interrelated model for joint intent detection and slot filling. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5467–5471.
- Keqing He, Shuyu Lei, Yushu Yang, Huixing Jiang, and Zhongyuan Wang. 2020a. **Syntactic graph convolutional network for spoken language understanding**. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2728–2738, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Keqing He, Weiran Xu, and Yuanmeng Yan. 2020b. Multi-level cross-lingual transfer learning with language shared and specific knowledge for spoken language understanding. *IEEE Access*, 8:29407–29416.
- Keqing He, Yuanmeng Yan, Si hong Liu, Z. Liu, and Weiran Xu. 2020c. Learning label-relational output structure for adaptive sequence labeling. *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8.
- Keqing He, Yuanmeng Yan, and Weiran Xu. 2020d. **Learning to tag OOV tokens by integrating contextual representation and background knowledge**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 619–624, Online. Association for Computational Linguistics.
- Keqing He, Jinchao Zhang, Yuanmeng Yan, Weiran Xu, Cheng Niu, and Jie Zhou. 2020e. Contrastive zero-shot learning for cross-domain slot filling with adversarial attack. In *COLING*.
- C. T. Hemphill, J. J. Godfrey, and G. Doddington. 1990. The atis spoken language systems pilot corpus. In *HLT*.
- Dan Hendrycks and Kevin Gimpel. 2017. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *ArXiv*, abs/1610.02136.
- Ziniu Hu, Ting Chen, Kai-Wei Chang, and Yizhou Sun. 2019. **Few-shot representation learning for out-of-vocabulary words**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4102–4112, Florence, Italy. Association for Computational Linguistics.
- Zhiheng Huang, W. Xu, and Kai Yu. 2015. Bidirectional lstm-crf models for sequence tagging. *ArXiv*, abs/1508.01991.

- Stefan Larson, Anish Mahendran, Joseph Peper, Christopher Clarke, Andrew Lee, P. Hill, Jonathan K. Kummerfeld, Kevin Leach, M. Laurenzano, L. Tang, and J. Mars. 2019. An evaluation dataset for intent classification and out-of-scope prediction. *ArXiv*, abs/1909.02027.
- Kimin Lee, Kibok Lee, H. Lee, and Jinwoo Shin. 2018. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In *NeurIPS*.
- Dongyun Liang, Weiran Xu, and Yingge Zhao. 2017a. [Combining word-level and character-level representations for relation classification of informal text](#). In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 43–47, Vancouver, Canada. Association for Computational Linguistics.
- Shiyu Liang, Yixuan Li, and R. Srikant. 2017b. Principled detection of out-of-distribution examples in neural networks. *ArXiv*, abs/1706.02690.
- Shiyu Liang, Yixuan Li, and R. Srikant. 2018. Enhancing the reliability of out-of-distribution image detection in neural networks. *arXiv: Learning*.
- Ting-En Lin and H. Xu. 2019. Deep unknown intent detection with margin loss. *ArXiv*, abs/1906.00434.
- Bing Liu and Ian Lane. 2015. Recurrent neural network structured output prediction for spoken language understanding. In *Proc. NIPS Workshop on Machine Learning for Spoken Language Understanding and Interactions*.
- Bing Liu and Ian Lane. 2016. Attention-based recurrent neural network models for joint intent detection and slot filling. *arXiv preprint arXiv:1609.01454*.
- Samuel Louvan and B. Magnini. 2020. Recent neural methods on slot filling and intent classification for task-oriented dialogue systems: A survey. In *COLING*.
- Grégoire Mesnil, Yann Dauphin, Kaisheng Yao, Yoshua Bengio, Li Deng, Dilek Z. Hakkani-Tur, Xiaodong He, Larry Heck, Gokhan Tur, Dong Yu, and Geoffrey Zweig. 2015. Using recurrent neural networks for slot filling in spoken language understanding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23:530–539.
- Lev Ratinov and Dan Roth. 2009. Design challenges and misconceptions in named entity recognition. In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning (CoNLL-2009)*, pages 147–155.
- J. Ren, Peter J. Liu, E. Fertig, Jasper Snoek, Ryan Poplin, Mark A. DePristo, Joshua V. Dillon, and Balaji Lakshminarayanan. 2019. Likelihood ratios for out-of-distribution detection. In *NeurIPS*.
- Darsh J. Shah, Raghav Gupta, A. Fayazi, and Dilek Z. Hakkani-Tür. 2019. Robust zero-shot cross-domain slot filling with example values. *ArXiv*, abs/1906.06870.
- Lei Shu, Hu Xu, and Bing Liu. 2017. Doc: Deep open classification of text documents. *ArXiv*, abs/1709.08716.
- H. Xu, Keqing He, Yuanmeng Yan, Si hong Liu, Z. Liu, and Weiran Xu. 2020a. A deep generative distance-based classifier for out-of-domain detection with mahalanobis space. In *COLING*.
- Hong Xu, Keqing He, Yuanmeng Yan, Sihong Liu, Zijun Liu, and Weiran Xu. 2020b. [A deep generative distance-based classifier for out-of-domain detection with mahalanobis space](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1452–1460, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Yuanmeng Yan, Keqing He, Hong Xu, Sihong Liu, Fanyu Meng, Min Hu, and Weiran Xu. 2020. [Adversarial semantic decoupling for recognizing open-vocabulary slots](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6070–6075, Online. Association for Computational Linguistics.
- B. Yang and Tom Michael Mitchell. 2017. Leveraging knowledge bases in lstms for improving machine reading. In *ACL*.
- Zhiyuan Zeng, Keqing He, Yuanmeng Yan, Hong Xu, and Weiran Xu. 2021a. Adversarial self-supervised learning for out-of-domain detection. In *NAACL*.
- Zhiyuan Zeng, Hong Xu, Keqing He, Yuanmeng Yan, Sihong Liu, Zijun Liu, and Weiran Xu. 2021b. [Adversarial generative distance-based classifier for robust out-of-domain detection](#). In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7658–7662.
- Lin Zhao and Zhe Feng. 2018. [Improving slot filling in spoken language understanding with joint pointer and attention](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 426–431, Melbourne, Australia. Association for Computational Linguistics.
- Yinhe Zheng, Guanyi Chen, and Minlie Huang. 2020. Out-of-domain detection for natural language understanding in dialog systems. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:1198–1209.