

# Beyond the Known: Investigating LLMs Performance on Out-of-Domain Intent Detection

Pei Wang<sup>1\*</sup>, Keqing He<sup>2\*</sup>, Yejie Wang<sup>1\*</sup>, Xiaoshuai Song<sup>1</sup>, Yutao Mou<sup>1</sup>,  
Jingang Wang<sup>2</sup>, Yunsen Xian<sup>2</sup>, Xunliang Cai<sup>2</sup>, Weiran Xu<sup>1\*</sup>

{wangpei, wangyejie, songxiaoshuai, myt, xuweiran}@bupt.edu.cn  
{hekeqing, wangjingang, xianyunsen, caixunliang}@meituan.com

## Abstract

Out-of-domain (OOD) intent detection aims to examine whether the user's query falls outside the predefined domain of the system, which is crucial for the proper functioning of task-oriented dialogue (TOD) systems. Previous methods address it by fine-tuning discriminative models. Recently, some studies have been exploring the application of large language models (LLMs) represented by ChatGPT to various downstream tasks, but it is still unclear for their ability on OOD detection task. This paper conducts a comprehensive evaluation of LLMs under various experimental settings, and then outline the strengths and weaknesses of LLMs. We find that LLMs exhibit strong zero-shot and few-shot capabilities, but is still at a disadvantage compared to models fine-tuned with full resource. More deeply, through a series of additional analysis experiments, we discuss and summarize the challenges faced by LLMs and provide guidance for future work including injecting domain knowledge, strengthening knowledge transfer from IND(In-domain) to OOD, and understanding long instructions.

**Keywords:** OOD, ChatGPT, LLM

## 1. Introduction

Traditional TOD systems are based on the closed-set hypothesis (Chen et al., 2019; Yang et al., 2021; Zeng et al., 2022) and can only handle queries within a limited scope of in-domain(IND) intents. However, users may input queries with out-of-domain(OOD) intents in the real open world, which poses new challenges for TOD systems. As shown in Figure 1, OOD intent detection task aims to determine whether the intent of user queries exceeds the predefined intents, making it an essential component of TOD systems. (Tulshan and Dhage, 2018; Lin and Xu, 2019; Zeng et al., 2021; Wu et al., 2022a,b; Mou et al., 2022, 2023; Song et al., 2023b).

Previous work on OOD detection rely on fine-tuning pre-training language model (PLM), extracting the output representation of PLMs' final layer as the intent feature, and employing scoring functions based on density, distance, or energy to detect OOD samples, as shown in Figure 2 (Zeng et al., 2021; Zhou et al., 2022; Mou et al., 2022; Cho et al., 2023; Wang et al., 2023b). Recently, the emergence of large language models(LLMs) like ChatGPT<sup>1</sup> has injected new vitality into natural language process(NLP) tasks. Their superior zero-shot learning capability enables a new paradigm of NLP research and applications by prompting LLMs without finetuning (Ouyang et al., 2022; Touvron et al., 2023; Jiao et al., 2023; Wei et al., 2023; Yang et al., 2023). Given the LLMs' training on broad

text corpora and their impressive generalization skills, it's worth considering the benefits and potential challenges they may face in open-scenario intent identification. Specifically, we have raised the following questions:

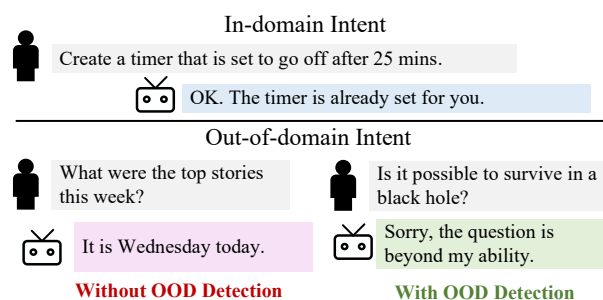


Figure 1: Explanation of the role of OOD intent detection in the TOD system. When the system encounters an intent that is beyond its supported intents, it can detect and friendly prompt the user.

1. What are the potential positive and negative effects of large language models on the Out-of-Domain (OOD) detection task?
2. What are the strengths and weaknesses of large language models, compared with traditional fine-tuned models?
3. Why do large language models exhibit certain strengths and weaknesses?
4. How can we potentially address and improve these weaknesses?

In this work, we introduce two LLM-based OOD framework, ZSD-LLM and FSD-LLM which based on different IND prior to instruct LLM to conduct

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

<sup>1</sup><https://openai.com/blog/ChatGPT>

intent detection (Section 3). Then we conduct comparative experiments between ChatGPT and discriminative methods (Section 4). In order to further explore the underlying reasons behind the experiments, we conduct a series of analytical experiments including IND intent number effect, different data split, comparison of different LLMs and different prompts effect (Section 5). Finally, we summarize the strengths and weaknesses of ChatGPT in OOD detection tasks and future improvement directions (Section 6). To the best of our knowledge, we are the first to comprehensively evaluate the performance of LLMs on OOD intent detection.

The key findings of this paper can be summarized as follows:

**What ChatGPT does well:**

- ChatGPT can achieve good zero-shot performance without providing any IND intent priors, demonstrating his powerful NLU capabilities.
- When the number of IND intents is small, ChatGPT can achieve better accuracy in few-shot settings than discriminative models.
- ChatGPT can not only perform OOD detection but also output the intent of the OOD samples, which is something that current methods based on discriminative models cannot achieve.

**What ChatGPT does not do well:**

- ChatGPT performs significantly worse than baselines with a large number of IND intents. It's manifested by an increase in misclassifications among IND intents and a substantial number of OOD samples being detected as IND when there is a higher number of IND intents.
- In rare instances, ChatGPT does not output according to our designed instructions. Particularly when the increase in intents leads to longer instructions, ChatGPT may overlook key information in the prompts, resulting in task failure.
- Compared to discriminative models, the performance of ChatGPT is affected to some extent by the number of intents. This is primarily manifested when the number of IND intents increases, resulting in a significant decline in its performance.
- ChatGPT struggles with fine-grained semantic distinctions which indicating the comprehension of ChatGPT in fine-grained intent labels is insufficient, exhibiting misalignment with human-level understanding.
- It's challenging for ChatGPT to acquire knowledge from IND demonstrations that could assist with OOD tasks. It might even perceive IND demonstrations as noisy, which could potentially harm the performance of OOD tasks.

We further summarize future LLM improvement directions which includes the following aspects:

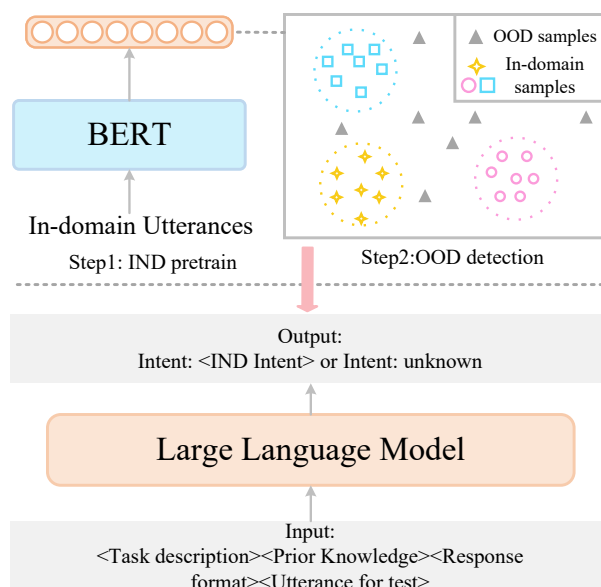


Figure 2: Comparison of the OOD detection method between previous method (Upper part) and LLM-based method (Lower part). Previous method trains a feature extractor using IND samples in the first stage, and estimates the confidence score of the sample using the designed scoring function and features; Our end-to-end OOD detection based on LLM adds task descriptions to prompts, and LLM directly outputs detection results.

- 1) injecting domain knowledge
- 2) strengthening knowledge transfer from IND to OOD, and
- 3) understanding long instructions.<sup>2</sup>

## 2. Related Work

### 2.1. LLM

LLM has become a popular paradigm for research and applications in natural language processing tasks. ChatGPT is a generative foundational model belonging to the GPT-3.5 series in the OpenAI GPT family, which includes its predecessors, GPT, GPT-2, and GPT-3. Recently, there has been an increasing interest in utilizing LLMs for various natural language processing (NLP) tasks. Several studies have been conducted to systematically investigate the performance of ChatGPT on different downstream tasks, including machine translation (Jiao et al., 2023), information extraction (Wei et al., 2023), summarization (Yang et al., 2023) and clustering (Song et al., 2023a). However, it is unclear about the performance of ChatGPT in OOD detection.

<sup>2</sup>Our codes can be found on [https://github.com/Yupeiwang/ood\\_llm\\_eval](https://github.com/Yupeiwang/ood_llm_eval)

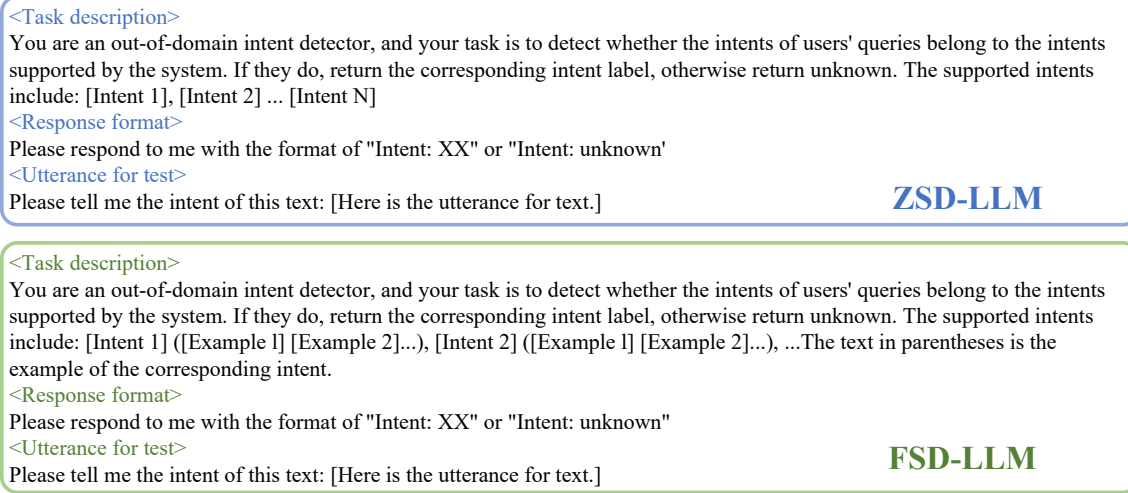


Figure 3: The demonstration of the two prompts we use to assist ChatGPT in performing OOD intent detection. FSD-LLM incorporates examples of intentions in the prompt as prior knowledge.

## 2.2. OOD Detection

Previous OOD detection methods can be divided into two categories: supervised OOD detection (Fei and Liu, 2016; Kim and Kim, 2018; Larson et al., 2019; Zheng et al., 2020) and unsupervised OOD detection (Shu et al., 2017; Lee et al., 2018; Ren et al., 2019; Lin and Xu, 2019; Xu et al., 2020; Zeng et al., 2021; Mou et al., 2022). The former indicates that there are some extensive labeled OOD samples in the training data. Classic supervised OOD algorithms consider the OOD detection problem as an  $N+1$  classification problem (Fei and Liu, 2016; Larson et al., 2019). Unsupervised methods generally perform in two stages: learning intent representation and estimating confidence scores (Mou et al., 2022). Most previous research has focused on fine-tuning small-scale pre-trained language models (PLMs), such as BERT, to learn intent features from the training data (Wang et al., 2023b). However, with the recent advancements in LLM, there is increasing interest in exploring their potential for OOD intent detection. Compared to small-scale PLMs, LLMs have greater capacity for learning and generalization from data, making them promising candidates for OOD intent detection.

## 3. Methodology

### 3.1. Problem Formulation

Given predefined set of intents, denoted as  $\mathcal{S} = \{l_1, l_2, \dots, l_N\}$ , it contains  $N$  intents supported by the system. The input is the user's natural language query  $q = \{t_1, t_2, \dots, t_n\}$ , where  $t_i$  represents  $i$ th token in the query. The output is an intent

label  $l_{pre}$  that belongs to the set  $\mathcal{S} \cup \{OOD\}$ .

### 3.2. Prompt Engineer

We evaluate the OOD intent detection capability of ChatGPT in an end-to-end manner. We heuristically propose two prompts based on different IND prior:

**Zero-shot Detection (ZSD-LLM):** This method only provides the IND intent set in the prompt as prior knowledge without supplying any IND samples. It can be utilized in scenarios where user privacy protection is required. The prompt template is: <Task description><Prior:  $\mathcal{S}$ ><Response format><Utterance for test>.

**Few-shot Detection (FSD-LLM):** This method provides several samples for each intent in the prompt, allowing ChatGPT to extract useful knowledge from these samples and apply it to distinguish between IND and OOD intents. The prompt template is: <Task description><Prior:  $D = \{(q_1, l_1), (q_2, l_2) \dots, (q_n, l_n)\}$ ><Response format><Utterance for test>.

We show this two methods in Figure 3. About the exploration of different prompts, we discuss it in Section 5.4.

## 4. Experiment

### 4.1. Setup

#### 4.1.1. Dataset & Metric

**Dataset** We conduct experiments on two widely used benchmark, CLINC (Larson et al., 2019) and Banking (Casanueva et al., 2020). CLINC consists of 150 intents distributed across 10 domains,

Model	Split = 25%						Split = 50%						Split = 75%					
	ALL		IND		OOD		ALL		IND		OOD		ALL		IND		OOD	
	ACC	F1	ACC	F1	Recall	F1	ACC	F1	ACC	F1	Recall	F1	ACC	F1	ACC	F1	Recall	F1
SCL	74.36	61.06	71.57	60.18	76.46	77.76	75.45	69.94	80.43	69.98	67.44	68.25	79.43	84.54	82.35	84.86	71.06	66.34
KNN-CL	88.21	77.65	78.16	77.94	91.77	92.36	<b>81.98</b>	<b>83.67</b>	<b>85.13</b>	<b>83.72</b>	<b>85.25</b>	<b>81.96</b>	81.69	70.21	<b>86.30</b>	<b>86.03</b>	85.13	71.88
UniNL	<b>89.41</b>	<b>80.04</b>	78.59	<b>79.36</b>	<b>92.96</b>	<b>93.02</b>	81.42	82.66	84.68	82.70	78.24	81.18	<b>82.78</b>	<b>86.36</b>	81.78	<b>86.59</b>	<b>85.66</b>	<b>73.34</b>
ChatGPT	47.5	42.68	73.16	42.02	39.09	55.17	46.46	54.91	71.32	55.47	22.24	33.77	50.58	56.97	62.85	57.58	15.62	22.03

Table 1: The performance comparison between ChatGPT and baselines of Banking. We select 25%, 50%, and 75% of all intents as IND intents. Three average values are taken for each experiment.

Model	Split = 25%						Split = 50%						Split = 75%					
	ALL		IND		OOD		ALL		IND		OOD		ALL		IND		OOD	
	ACC	F1	ACC	F1	Recall	F1	ACC	F1	ACC	F1	Recall	F1	ACC	F1	ACC	F1	Recall	F1
SCL	87.64	88.32	91.44	89.08	74.11	77.76	85.82	83.64	84.89	83.58	86.46	87.94	89.18	90.55	88.75	90.58	89.86	<b>87.13</b>
KNN-CL	<b>92.04</b>	83.31	84.86	82.99	<b>93.85</b>	<b>94.97</b>	90.33	88.52	88.53	88.47	<b>91.57</b>	<b>91.95</b>	89.18	92.03	88.49	92.10	<b>92.30</b>	77.66
UniNL	87.8	<b>89.79</b>	<b>97.27</b>	<b>89.89</b>	77.3	86.14	<b>90.95</b>	<b>93.15</b>	<b>95.79</b>	93.25	80.03	85.75	<b>91.77</b>	<b>94.01</b>	<b>93.84</b>	<b>94.09</b>	83.95	82.74
ChatGPT	63.86	58.86	81.26	58.51	58.15	71.82	59.84	69.9	82.4	70.14	37.29	51.43	64.24	70.18	74.79	70.44	33.16	41.71

Table 2: The performance comparison between ChatGPT and baselines of CLINC. We select 25%, 50%, and 75% of all intents as IND intents. Three average values are taken for each experiment.

Statistic	Banking	CLINC
Avg utterance length	9	12
Intent	150	77
Training set size	15000	9003
Training sample per class	100	-
Development set size	3000	1000
Development sample per class	20	-
Testing set size	5500	3080
Testing sample per class	30	-

Table 3: Statistics of datasets.

with each domain containing 15 intents. Banking contains intents from a single domain, totaling 77 intents. Consistent with previous research, we conduct OOD detection under three settings: 25%, 50%, and 75%. Here, 25% refers to selecting 25% of the intents as IND, with the remaining intents considered as OOD. We show the detailed statistics of the datasets in Table 3.

**Metric** We employ six commonly used OOD detection metrics to evaluate the performance, including IND metrics: accuracy and macro-F1, OOD metrics: recall and macro-F1, as well as overall accuracy and macro-F1.

#### 4.1.2. Baselines

We compare ChatGPT with the following three state-of-the-art discriminative two-stage methods:

**SCL** (Zeng et al., 2021) It proposes a supervised contrastive learning objective to minimize intra-class variance by pulling together in-domain intents belonging to the same class and maximize inter-class variance by pushing apart samples from different classes.

**KNN-CL** (Zhou et al., 2022) It proposes a KNN-based contrastive loss for IND pre-training. KNN-CL selects k-nearest neighbors from samples of the same class as positives and uses samples of the different classes as negatives.

**UniNL** (Mou et al., 2022) It proposes a unified Neighborhood Learning to align representation learning with the scoring function to improve OOD detection performance. KNCL objective is employed for IND pre-training and a KNN-based score function is used for OOD detection.

## 4.2. ZSD-LLM Results

Our results are shown in Table 1 and 2. The results show that ZSD-LLM performs worse than the best baselines on all metrics. we analyze the results from three aspects:

(1) **The performance of IND intent recognition.** There is a certain gap between ChatGPT and strong baselines (UniNL, KNN-CL). Taking Banking-50% as an example, in terms of IND indicators, ChatGPT’s performance is 13.36% (IND-ACC) and 27.23% (IND-F1) lower than UniNL. However, the gap between ChatGPT and SCL is slightly smaller, and it even surpasses SCL in some settings, such as Banking 25% (71.57 -> 73.16). This demonstrates ChatGPT’s strong zero-shot capability.

(2) **The performance of OOD sample detection.** Compared to IND classification, the performance gap between ChatGPT and baselines is larger on OOD metrics. Specifically, ChatGPT’s OOD-Recall is reduced by 56%, and OOD-F1 is reduced by 47.41% compared with UniNL for banking-50%. It’s generally observed across three IND intent splits in the two datasets. We speculate that the reason for the lower OOD metrics is that a large number of OOD samples are misclassified as IND intents by ChatGPT. Such results indicate that this kind of zero-shot prompting is not enough to provide ChatGPT with sufficient prior knowledge to complete the OOD detection task.

(3) **Comparison between datasets.** The performance comparison of ChatGPT between the two datasets shows the same trend as the baselines.

The detection ability of the multi-domain dataset (CLINC) is better than that of the single-domain dataset (Banking). On the simpler dataset CLINC, the gap between ChatGPT and UniNL is smaller than that on Banking. (ALL-ACC: 41.91 -> 23.84 for 25%, 34.96 -> 31.11 for 50%, 32.20 -> 27.53 for 75%). This indicates that the high granularity of intent division is the reason for the poor performance of ChatGPT.

(3) **Task fail.** In addition to the above results, we use the original OOD data from CLINC for OOD detection. This results in a total of 150 intents, which can be used to test ChatGPT’s ability to perform large-scale system OOD detection. Experimental results show that ChatGPT predicts new intents in approximately 8.49% of the test samples, neither returning an IND intent nor ‘unknown’, leading to task failure. This reflects the instability of LLM in performing OOD detection.

### 4.3. FSD-LLM Results

Due to the length limitation of ChatGPT’s conversations, we reduce the number of IND intents and randomly select N=5,10,20,30,40 intents as IND intents, with the number of OOD intents fixed at 20. Under each setting, we test four groups of experiments with K=0,1,3,5 (K is the number of samples provided for each intent). We show the detailed FSD-LLM results in Table 4 and the changing trend in Figure 4. We discover that:

(1) **FSD-LLM demonstrates strong competitiveness compared to the baseline in situations with a limited number of INDs**

When N = 5, K = 5, ChatGPT outperforms UniNL by 0.76 and 3.16 on ALL-ACC and ALL-IND respectively. When N = 10 and N = 20, ChatGPT is superior to UniNL in IND classification, but inferior in UniNL in OOD detection. When N = 30 and N = 40, UniNL widens the gap with ChatGPT.

(2) **The more the number of intents, the more demonstrations are needed for IND intent recognition.**

From Figure 4, We find that when N=5,10, FSD- OOD achieves better F1-OOD and F1-IND performance at K=1. Even ZS-LLM achieves best ACC-IND. However, as N increases to 30 or 40, both ACC-IND and F1-IND show an upward trend with the increase of K. This suggests that the more the number of intents, the more prior knowledge about intents is needed to help distinguish between different intents.

(3) **Too many demonstrations may introduce noise into OOD detection.**

OOD-Recall shows an overall trend of initially increasing and then decreasing. This demonstrates the model’s negative transfer from IND to OOD data which means . We speculate that this could be due to significant feature distribution differences

IND num	Few-shot	ALL		IND		OOD	
		ACC	F1	ACC	F1	ACC	F1
5	0	77.19	69.06	<b>96.36</b>	66.09	72.61	83.92
	1	88.80	80.11	90.00	77.60	88.50	92.67
	3	85.89	77.02	96.02	74.34	83.33	90.41
	5	<b>89.11</b>	<b>81.26</b>	90.01	<b>78.93</b>	88.89	<b>92.88</b>
	UniNL	88.35	78.1	78.35	75.21	<b>90.85</b>	92.58
10	0	63.44	64.85	86.87	64.64	51.71	66.94
	1	79.19	74.75	84.00	73.85	76.77	83.75
	3	78.81	76.67	<b>91.00</b>	76.10	72.77	82.35
	5	82.89	<b>88.17</b>	89.80	<b>79.57</b>	79.50	86.18
	UniNL	<b>84.17</b>	78.11	74.55	77.1	<b>89.05</b>	<b>88.25</b>
20	0	61.76	63.01	74.78	63.11	48.54	60.96
	1	70.89	72.29	79.29	72.36	62.44	70.89
	3	75.95	76.30	80.81	76.24	71.07	77.35
	5	77.84	<b>79.57</b>	<b>88.27</b>	<b>79.74</b>	65.82	76.33
	UniNL	<b>80.51</b>	78.87	73.21	78.71	<b>87.89</b>	<b>81.92</b>
30	0	56.91	58.13	74.31	58.64	30.56	42.72
	1	69.06	72.00	78.64	72.28	54.40	63.64
	3	72.18	72.65	84.31	73.00	52.63	62.11
	5	74.74	81.38	<b>91.95</b>	82.03	47.62	61.86
	UniNL	<b>80.76</b>	<b>83.07</b>	82.53	<b>83.29</b>	<b>78.09</b>	<b>76.61</b>
40	0	55.80	60.26	67.25	60.75	31.86	40.66
	1	66.32	70.66	76.53	71.08	45.60	53.99
	3	69.08	75.32	83.12	75.87	40.91	53.11
	5	69.35	75.84	86.84	76.55	32.80	47.15
	UniNL	<b>83.69</b>	<b>87.77</b>	<b>86.86</b>	<b>88.05</b>	<b>77.32</b>	<b>76.53</b>

Table 4: Performance of ChatGPT under different few-shot settings with varying five sets of IND numbers.

between the IND and OOD data, making it challenging for the model to learn useful features from the IND samples for OOD data.

## 5. Qualitative Analysis

### 5.1. Effect of IND intent number

In Section 4.3, we observe the varying performance of ChatGPT under different numbers of IND intents. In this section, we provide a detailed analysis of the changes in the effectiveness of the ZSD-LLM methods as N increases. Figure 5 shows the trend of the changes. The results reveal that as the number of IND intents increases:

(1) **ChatGPT is sensitive to the number of intents compared with UniNL.** As shown in Figure 5a and 5b, ChatGPT has the best OOD detection performance when N=5, but as the number of intents increases, both metrics consistently decrease. When it reaches 30 and 40, ACC and F1 decrease by 21.33 and 8.8, respectively. However, UniNL consistently demonstrates robust results across all numbers. UniNL still achieves an ACC of 83.69 and F1 of 87.77 when N=40.

(2) **The increase in intents leads to more severe confusion between labels.** Figure 5c shows a continuous decrease in IND-ACC, indicating that more IND samples are misclassified. We find that compared to IND samples being misclassified as OOD, the proportion of being misclassified as incorrect IND intents is increasing. This may be due to ChatGPT’s understanding of intent labels not aligning with human-defined labels, leading to con-

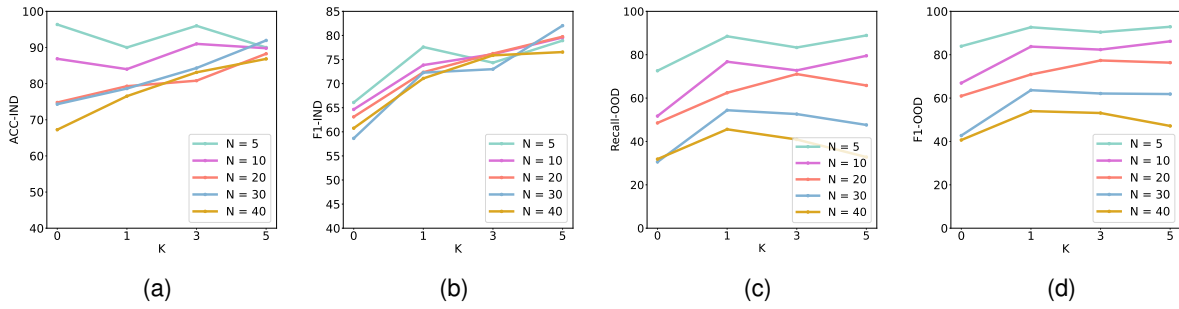


Figure 4: The effect of few-shot on different IND number. We show the changes of four metrics under different demonstration quantities. Due to the limitation of ChatGPT’s input length, we conduct three sets of experiments with 1-shot, 3-shot, and 5-shot settings.

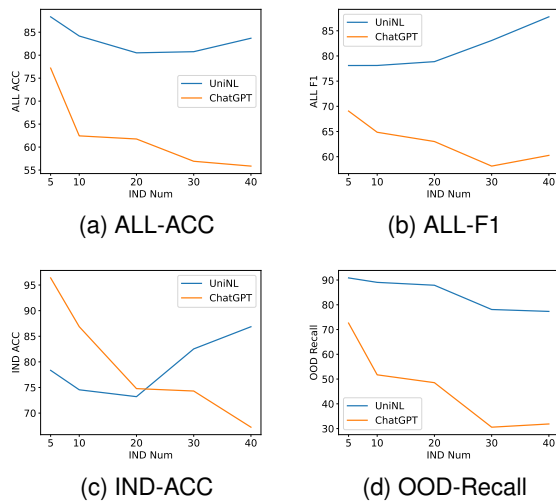


Figure 5: Changes in the ALL-ACC, ALL-F1, IND-ACC and OOD-Recall of ChatGPT and UniNL as the number of IND intent increases for banking-50%.

fusion between different label meanings. As the number of intents increases, this confusion intensifies.

(3) **The increase in intents causes a sharp drop in the OOD-recall rate.** Figure 5d shows that with the increase in the number of intents, OOD samples are more likely to be misclassified as IND. This is because the increase in IND intents number introduces more interference to ChatGPT’s OOD detection.

We believe that the advantage of ChatGPT over discriminative models lies in its OOD detection with fewer intents, where it can accurately make judgments based on its internal knowledge. However, as the number of labels increases, the confusion in label meanings becomes more prominent, resulting in a decline in both IND intent classification and OOD sample detection.

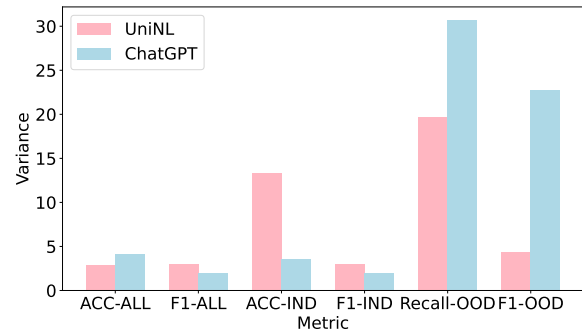


Figure 6: The variances of ChatGPT and UniNL on various metrics across five different data splits.

## 5.2. The robustness of ChatGPT OOD detection

The robustness of OOD detection can be reflected in its ability to maintain stable across different intent partitions. To verify it, we randomly select five different IND intent set split (using five different seeds when selecting IND intents) for the experiment. Figure 6 shows the variance of the results from the five sets of experiments. We observe that UniNL and ChatGPT perform similarly. However, in terms of the IND metrics IND-ACC and IND-F1, UniNL shows greater fluctuations. In terms of OOD metrics, ChatGPT is inferior to UniNL. This highlights the differences between the two models in performing OOD intent detection. ChatGPT excels in IND intent recognition, but its stability in OOD detection is relatively poor.

## 5.3. Comparison of Different LLMs

We do ZSD-LLM on other mainstream LLMs and compare them with ChatGPT.

- **Text-davinci-002, text-davinci-003**<sup>3</sup> belong to InstructGPT and text-davinci-003 is an improved

<sup>3</sup><https://platform.openai.com/docs/models>

Model	ALL		IND		OOD	
	ACC	F1	ACC	F1	Recall	F1
text-davinci-002	54.24	60.72	73.6	60.85	34.38	48.53
text-davinci-003	55.87	64.14	79.03	65.15	30.81	43.98
Claude	56.58	52.76	70.72	52.74	43.59	53.12
Llama2-70b-Chat	55.5	57.8	67.0	57.84	44.0	56.96
ChatGPT	61.76	63.01	74.78	63.11	48.54	60.96
GPT4	<b>68.5</b>	<b>73.55</b>	<b>87.5</b>	<b>74.07</b>	<b>49.5</b>	<b>63.26</b>

Table 5: OOD detection performance of six different LLMs.

version of text-davinci-002. Compared with GPT-3, the biggest difference of InstructGPT is that it is fine-tuned for human instructions.

- **Claude** is an artificial intelligence chatbot developed by Anthropic<sup>4</sup>.
- **Llama2-70B-Chat**<sup>5</sup> is developed and open-sourced by Meta AI. Llama-2-70b-Chat is the native open-source version with high-precision results.
- **GPT4**<sup>6</sup> is the latest and most advanced multi-modal large model from OpenAI. GPT-4 can generate more factual and accurate statements than GPT-3.5 and other language models.

Results are shown in Table 5. GPT4 leads in all six metrics compared to other models. Surprisingly, GPT3 performs better in IND intent recognition, with text-davinci-003 even surpassing ChatGPT and text-davinci-002 shows similar performance to ChatGPT in IND metrics. However, ChatGPT exhibits significantly better results than the GPT-3 series in OOD metrics. The differences in performance may be attributed to ChatGPT's inclusion of SFT (Supervised Fine-Tuning) during the optimization phase, which gives it an advantage in understanding human instructions. In contrast, the GPT-3 series is slightly inferior in understanding tasks, making it more inclined towards intent detection tasks. Llama2-70b-Chat and Claude also exhibit a similar phenomenon to ChatGPT, but overall, ChatGPT outperforms Llama2-70B-Chat and Claude.

#### 5.4. Effect of different prompts

Prompt engineering is a crucial strategy for LLM. To verify the impact of different prompts, we devise four additional variations. They are:

**prompt.detector**: Modify the role positioning of the LLM from intent detector to OOD detector.

**prompt.discovery**: Adopt a new task description for intent discovery. Directly make the LLM return the labels of OOD intent.

<sup>4</sup><https://www.anthropic.com/product>

<sup>5</sup><https://huggingface.co/meta-llama/Llama-2-70b-chat-hf>

<sup>6</sup><https://platform.openai.com/docs/models>

Model	ALL		IND		OOD	
	ACC	F1	ACC	F1	Recall	F1
prompt.original	53.10	51.32	56.82	51.18	49.47	56.62
prompt.detector	51.07	53.20	61.61	53.25	40.78	51.11
prompt.discovery	45.63	51.05	63.29	51.32	28.39	40.66
prompt.order	42.47	48.57	51.75	48.74	33.40	42.09
prompt.reason	48.67	47.65	55.47	47.56	42.03	50.92

Table 6: The performance of ChatGPT on various prompts.

**prompt.order**: Change the order of each part, using <Utterance for test> <Task description> <Utterance for test>.

**prompt.reason**: Output the reason before outputting the results.

As the version of ChatGPT changes over time, we choose to use the open-source Llama2-70B-Chat for the prompt experiment. Results are shown in the Table 6. We find that in both detection and discovery modes, LLM tends to perform IND detection rather than OOD detection. The prompt.order led to a decline in performance, which may be due to the overly long instructions causing LLM to forget the information from the beginning sentence. The Reason mode exacerbate the LLM's use of incorrect domain knowledge. About the exploration of a better prompt, we leave it for future research.

## 6. Challenge & Further Discussion

Based on the above experiments and analysis, we identify the challenging scenarios that LLMs encounter and offer guidance for future reference.

### 6.1. Conflict between Domain-Specific Knowledge and General Knowledge

The majority of errors are caused by the model's incorrect utilization of certain knowledge, which may be due to discrepancies between LLM and humans in both intent and task understanding. We refer to it as conflicts between generic knowledge within the model and domain-specific knowledge required for the task. Specifically, they result in three types of errors: **false association**, **focus deviation**, and **lack of domain knowledge**. We display the relevant cases in Figure 7(a).

Our FSD-LLM can inject domain-specific knowledge in Section 4.3. However, the effectiveness of demonstrations seems to be influenced by various factors such as the numbers of IND intents, the number of demonstrations. Besides, the knowledge conflict may be more pronounced in larger models than in smaller ones (Section 5.3), so how to inject domain-specific knowledge into LLM and eliminate noise interference from useless general knowledge will be a future research direction.

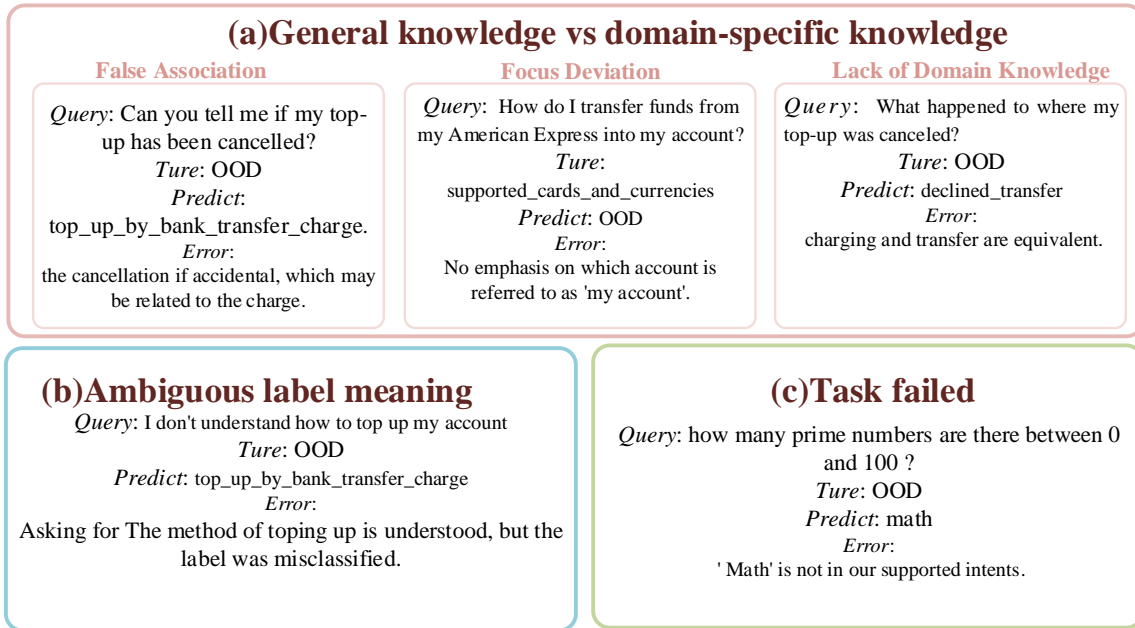


Figure 7: We identify three challenges that ChatGPT faces when performing intent OOD detection. We further subdivide the categories of *general knowledge vs domain-specific knowledge* into *False Association*, *Focus Deviation*, and *Lack of Domain Knowledge*. We provide specific cases to illustrate these errors, including the query, true label, predicted label, and the reasons for ChatGPT's misclassification.

## 6.2. Difficulty of Knowledge Transfer from IND to OOD

Experiments in Section 4.3 show that FSD-LLM achieve limited improvements in OOD sample detection. Ambiguous label meaning is quite common as shown in Figure 7(b). Too many demonstrations even lead to a decrease in results especially when the number of intents is low. One improvement method is to add OOD samples in prompt, but estimating the number of OOD intents and providing sufficiently comprehensive OOD data is challenging. Future research can focus on **how to enable models to learn transfer knowledge from IND's prior knowledge to OOD detection**. One optimization direction is to focus on how to select high-quality examples to inject diverse and noise-free prior knowledge.

## 6.3. Sensitivity to input length

When it comes to OOD intent detection with a large number of intents, ChatGPT exhibits errors in understanding instructions as mentioned in Section 4.2. When the prompt becomes longer, it is easy to exceed the processing capacity of LLMs, leading to erroneous outputs as shown in Figure 7(c). This limitation restricts the versatility of models like ChatGPT in handling diverse task scenarios and calls for future research efforts to address this issue.

## 6.4. Future Insights

We look forward to the emergence of LLM that can set new benchmarks and open up new areas of application. Unfortunately, such models have not yet appeared. On the contrary, we discovery some shortcomings in their ability to handle OOD tasks. At the same time, we also find the powerful zero-shot and few-shot capabilities of LLMs. Through the analysis, we believe that the further improvement directions for LLMs are **1) injecting domain knowledge, 2) strengthening knowledge transfer from In-Distribution (IND) to OOD, and 3) understanding long instructions**. We hope that our work can bring a deeper understanding of LLMs to the academic community, and we look forward to future work that can improve the application of LLMs in domain tasks.

## 7. Conclusion

In this paper, we conduct a comprehensive evaluation of ChatGPT for OOD intent detection. We first compare the performance of ChatGPT with traditional discriminative models and identify a significant performance gap. Additionally, we observe that ChatGPT excels in handling tasks with a small number of intents but struggles with tasks involving a large number of intents. While incorporating demonstration examples shows some improve-



ments, there is still considerable room for enhancement. We recommend future research to focus on improving large-scale models for OOD tasks by incorporating domain-specific knowledge into the models and how to learn transfer relationship from OOD detection.

## 8. Limitations

In this paper, we investigate the advantages, disadvantages and challenges of LLMs in open-domain intent OOD detection. Although we conduct extensive experiments, there are still several directions to be improved: (1) We propose FSD-LLM to do few-shot OOD detection for LLM, but the demonstration examples are randomly selected. In this paper, we do not consider the diversity and quality of the demonstration examples. (2) We use closed source LLMs in this paper like ChatGPT and GPT4. Although we ensure that all experiments are based on the same version (gpt-3.5-turbo-0301, gpt-4-0613), further updates of ChatGPT may lead to results in the future that will differ from those reported in this paper.

## 9. Acknowledgements

Thanks to the State Key Laboratory of Massive Personalized Customization System Technology for their support. This work was supported by the National Natural Science Foundation of China (NSFC No.62076031 and No.62076036).

## 10. Bibliographical References

- Abhijit Bendale and Terrance E. Boult. 2016. Towards open set deep networks. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1563–1572.
- Iñigo Casanueva, Tadas Temčinas, Daniela Gerz, Matthew Henderson, and Ivan Vulić. 2020. Efficient intent detection with dual sentence encoders. *arXiv preprint arXiv:2003.04807*.
- A. Castor and L. E. Pollux. 1992. The use of user modelling to guide inference and learning. *Applied Intelligence*, 2(1):37–53.
- Qian Chen, Zhu Zhuo, and Wen Wang. 2019. Bert for joint intent classification and slot filling. *arXiv preprint arXiv:1902.10909*.
- Hyunsoo Cho, Choonghyun Park, Junyeop Kim, Hyuhng Joon Kim, Kang Min Yoo, and Sang goo Lee. 2023. [Probing out-of-distribution robustness of language models with parameter-efficient transfer learning](#).
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. A survey for in-context learning. *arXiv preprint arXiv:2301.00234*.
- Geli Fei and Bing Liu. 2016. Breaking the closed world assumption in text classification. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 506–514.
- Dan Hendrycks and Kevin Gimpel. 2017. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *ArXiv*, abs/1610.02136.
- Wenxiang Jiao, Wenxuan Wang, Jen tse Huang, Xing Wang, and Zhaopeng Tu. 2023. [Is chatgpt a good translator? yes with gpt-4 as the engine](#).
- Joo-Kyung Kim and Young-Bum Kim. 2018. Joint learning of domain classification and out-of-domain detection with dynamic class weighting for satisficing false acceptance rates. *ArXiv*, abs/1807.00072.
- Stefan Larson, Anish Mahendran, Joseph Peper, Christopher Clarke, Andrew Lee, Parker Hill, Jonathan K. Kummerfeld, Kevin Leach, Michael Laurenzano, Lingjia Tang, and Jason Mars. 2019. An evaluation dataset for intent classification and out-of-scope prediction. In *EMNLP/IJCNLP*.
- Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. 2018. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *ArXiv*, abs/1807.03888.
- Ting-En Lin and Hua Xu. 2019. Deep unknown intent detection with margin loss. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5491–5496.
- Yutao Mou, Xiaoshuai Song, Keqing He, Chen Zeng, Pei Wang, Jingang Wang, Yunsen Xian, and Weiran Xu. 2023. [Decoupling pseudo label disambiguation and representation learning for generalized intent discovery](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9661–9675, Toronto, Canada. Association for Computational Linguistics.
- Yutao Mou, Pei Wang, Keqing He, Yanan Wu, Jingang Wang, Wei Wu, and Weiran Xu. 2022. Uninl: Aligning representation learning with scoring function for ood detection via unified neighborhood learning. *ArXiv*, abs/2210.10722.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong

- Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke E. Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Francis Christiano, Jan Leike, and Ryan J. Lowe. 2022. Training language models to follow instructions with human feedback. *ArXiv*, abs/2203.02155.
- Wenbo Pan, Qiguang Chen, Xiao Xu, Wanxiang Che, and Libo Qin. 2023. [A preliminary evaluation of chatgpt for zero-shot dialogue understanding](#).
- Jie Ren, Peter J. Liu, Emily Fertig, Jasper Snoek, Ryan Poplin, Mark A. DePristo, Joshua V. Dillon, and Balaji Lakshminarayanan. 2019. Likelihood ratios for out-of-distribution detection. *ArXiv*, abs/1906.02845.
- Lei Shu, Hu Xu, and Bing Liu. 2017. [Doc: Deep open classification of text documents](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2911–2916.
- Xiaoshuai Song, Keqing He, Pei Wang, Guanting Dong, Yutao Mou, Jingang Wang, Yunsen Xian, Xunliang Cai, and Weiran Xu. 2023a. [Large language models meet open-world intent discovery and recognition: An evaluation of ChatGPT](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10291–10304, Singapore. Association for Computational Linguistics.
- Xiaoshuai Song, Yutao Mou, Keqing He, Yueyan Qiu, Jinxu Zhao, Pei Wang, and Weiran Xu. 2023b. [Continual generalized intent discovery: Marching towards dynamic and open-world intent recognition](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4370–4382, Singapore. Association for Computational Linguistics.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. Llama: Open and efficient foundation language models. *ArXiv*, abs/2302.13971.
- Amrita S Tulshan and Sudhir Namdeorao Dhage. 2018. Survey on virtual assistant: Google assistant, siri, cortana, alexa. In *International symposium on signal processing and intelligent recognition systems*, pages 190–201.
- Jindong Wang, Xixu Hu, Wenxin Hou, Hao Chen, Runkai Zheng, Yidong Wang, Linyi Yang, Haojun Huang, Wei Ye, Xiubo Geng, Binxin Jiao, Yue Zhang, and Xing Xie. 2023a. [On the robustness of chatgpt: An adversarial and out-of-distribution perspective](#).
- Pei Wang, Keqing He, Yutao Mou, Xiaoshuai Song, Yanan Wu, Jingang Wang, Yunsen Xian, Xunliang Cai, and Weiran Xu. 2023b. [APP: Adaptive prototypical pseudo-labeling for few-shot OOD detection](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3926–3939, Singapore. Association for Computational Linguistics.
- Xiang Wei, Xingyu Cui, Ning Cheng, Xiaobin Wang, Xin Zhang, Shen Huang, Pengjun Xie, Jinan Xu, Yufeng Chen, Meishan Zhang, Yong Jiang, and Wenjuan Han. 2023. [Zero-shot information extraction via chatting with chatgpt](#).
- Yanan Wu, Keqing He, Yuanmeng Yan, QiXiang Gao, Zhiyuan Zeng, Fujia Zheng, Lulu Zhao, Huixing Jiang, Wei Wu, and Weiran Xu. 2022a. Revisit overconfidence for ood detection: Reassigned contrastive learning with adaptive class-dependent threshold. In *NAACL*.
- Yanan Wu, Zhiyuan Zeng, Keqing He, Yutao Mou, Pei Wang, and Weiran Xu. 2022b. Distribution calibration for out-of-domain detection with bayesian approximation. In *International Conference on Computational Linguistics*.
- Hong Xu, Keqing He, Yuanmeng Yan, Sihong Liu, Zijun Liu, and Weiran Xu. 2020. [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.
- Dao Xuan-Quy, Le Ngoc-Bich, Phan Xuan-Dung, Ngo Bac-Bien, and Vo The-Duy. 2023. [Evaluation of chatgpt and microsoft bing ai chat performances on physics exams of vietnamese national high school graduation examination](#).
- Jingkang Yang, Kaiyang Zhou, Yixuan Li, and Ziwei Liu. 2021. Generalized out-of-distribution detection: A survey. *arXiv preprint arXiv:2110.11334*.
- Xianjun Yang, Yan Li, Xinlu Zhang, Haifeng Chen, and Wei Cheng. 2023. [Exploring the limits of chatgpt for query or aspect-based text summarization](#).
- Weihaio Zeng, Keqing He, Zechen Wang, Dayuan Fu, Guanting Dong, Ruotong Geng, Pei Wang, Jingang Wang, Chaobo Sun, Wei Wu, and Weiran Xu. 2022. [Semi-supervised knowledge-grounded pre-training for task-oriented dialog](#)

systems. In *Proceedings of the Towards Semi-Supervised and Reinforced Task-Oriented Dialog Systems (SereTOD)*, pages 39–47, Abu Dhabi, Beijing (Hybrid). Association for Computational Linguistics.

Zhiyuan Zeng, Keqing He, Yuanmeng Yan, Zijun Liu, Yanan Wu, Hong Xu, Huixing Jiang, and Weiran Xu. 2021. [Modeling discriminative representations for out-of-domain detection with supervised contrastive learning](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 870–878, Online. 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.

Yunhua Zhou, Peiju Liu, and Xipeng Qiu. 2022. [KNN-contrastive learning for out-of-domain intent classification](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5129–5141, Dublin, Ireland. Association for Computational Linguistics.