

# Prompt Learning for Domain Adaptation in Task-Oriented Dialogue

**Makesh Narsimhan Sreedhar**  
University of Wisconsin - Madison \*  
msreedhar@wisc.edu

**Christopher Parisien**  
NVIDIA  
cparisien@nvidia.com

## Abstract

Conversation designers continue to face significant obstacles when creating production-quality task-oriented dialogue systems. The complexity and cost involved in schema development and data collection is often a major barrier for such designers, limiting their ability to create natural, user-friendly experiences. We frame the classification of user intent as the generation of a *canonical form*, a lightweight semantic representation using natural language. We show that canonical forms offer a promising alternative to traditional methods for intent classification. By tuning soft prompts for a frozen large language model, we show that canonical forms generalize very well to new, unseen domains in a zero- or few-shot setting. The method is also sample-efficient, reducing the complexity and effort of developing new task-oriented dialogue domains.

## 1 Introduction

Task-oriented dialogue systems in Conversational AI are challenging for developers to create. The current generation of dialogue frameworks requires developers to define actions (*intents*) and parameters (*slots*) that the natural language understanding (NLU) module accepts. This is then used to populate API service calls that operate in the backend to fulfill the user request. Casting natural language utterances from the user to a discrete set of intents and slots is often not very intuitive. This in turn leads to a situation where developers rely on hand-crafted rule-based grammars or a large annotated set of training samples for machine learning models to implement a given design. Any change to the design of the dialogue system would then require the developers to revisit and modify the implementation which is very often a time-consuming process. In this work, we aim to make dialogue system design easier and more intuitive.

The tremendous success of pre-trained language models such as BERT (Devlin et al., 2019) have made them the de facto standard for most intent classification and slot-filling tasks. However, these models are not immune to the challenge of adapting and extending existing models to new domains. One adaptation approach that has exploded in popularity in recent times is the usage of prompts with these language models. With a task description and few samples showing the input-output pairs, these language models become extremely effective at solving these tasks, especially at larger model sizes.

Manually specifying prompts suffers from sensitivity to phrasing; we get widely varying results based on how we frame the prompt. Prompt tuning (Lester et al., 2021) and p-tuning (Liu et al., 2021) have emerged as strong alternatives to manual prompt designing and they help optimize task-specific prompt tokens to get the best performance while keeping the language model itself frozen. In this work, we explore the task of intent classification using these large language models and p-tuning. Generative methods for classification tasks have not been widely adopted because generation is inherently difficult to control and utilize for further downstream tasks. Using our experiments on the Schema Guided Dialogue (Rastogi et al., 2019) dataset and the Virtual Assistant Benchmark (Liu et al., 2019), we show that with p-tuning we can achieve promising zero-shot and few-shot generalization capabilities to unseen domains.

In the task of intent classification, the intent labels provided as part of the dataset are usually terse and rigid. Generative models generalize better when intent labels are more descriptive but structured at the same. We borrow some aspects and terminology from semantic parsing to cast the intent labels to a more compositional format, known as *canonical forms*. In the traditional sense, canonical forms are paraphrases of the user utterances to

---

\* Work done as part of an internship with NVIDIA

convert them to a form that the semantic parser can operate on to output logical representations. In our use case, we loosely use the term, *canonical forms*, to refer to intent labels that are more descriptive than the discrete ones but are not too verbose, *e.g.*, "transfer\_money" → "transfer money to bank account". We manually frame these canonical forms and do not rely on any grammar, simplifying the approach.

We observe that using such *canonical forms* as labels for the intent classification task allows the model to generalize better to domains that are adjacent, but not seen at train time (*e.g.*, *Flight Reservations* → *Bus Bookings*). We also find that it is beneficial to do a two-stage P-tuning for domain adaption, *i.e.*, once we have a p-tuned large language model on a wide set of domains, we can continue p-tuning this model on a small set of labelled samples from the target domain to allow the model to generalize better. We find that this few-shot approach works very well and this has promising implications for developers for dialogue systems; with minimal effort it would be feasible to adapt an existing model pre-trained on multiple domains to a new domain. In summary, our contributions are:

- We cast the problem of intent classification into a generative approach and rewrite intent labels in a more descriptive format (*canonical forms*).
- When using such canonical forms, generative approaches with Large Language Models (LLMs) show promising results when compared with traditional methods for intent classification.
- Generative models generalize very well to unseen domains in zero-shot and few-shot settings when compared with BERT-style approaches.
- We demonstrate the sample efficiency of p-tuning LLMs where we can achieve close to full dataset performance with a fraction of the data.

## 2 Method

In this section, we describe the creation of canonical forms and the prompt tuning technique we adopt for intent classification in the task-oriented dialogue setting.

### 2.1 Canonical Forms

Canonical forms are usually paraphrases of the user utterance to a standardized form that can be utilized by downstream systems. These forms are traditionally obtained by using a set of grammar rules written by experts. The output of this process is a natural language sequence, but structured in a form that makes it better suited for a semantic parser. Different semantic parsers employ different canonical forms and thus transfer across datasets is quite challenging.

Utterance	Canonical Form
what is the newest published article?	article that has the largest publication date
who has published the most articles?	person that is author of the most number of article

Table 1: Examples of canonical forms corresponding to user utterances from the Overnight (Wang et al., 2015) semantic parsing dataset.

Our work uses canonical forms as a method of obtaining the intent of a user utterance. Traditionally, intent labels tend to be terse, which makes it difficult for models to generalize to unseen domains. The expressive and compositional nature of language models can be exploited if the intent labels are more verbose, allowing them to extrapolate the generated intents to capture even novel domains. At the same time, if the intent labels tend to be very long and riddled with descriptions, the language models become susceptible to hallucinations. Our work proposes the use of canonical forms as a way of establishing a balance between being terse and too verbose. We map intent labels to short descriptive phrases, *e.g.*, "check\_balance" → "check balance in bank account". Unlike traditional canonical forms, we do not use any formal grammar to perform this mapping and the phrases are manually specified. We believe that such an approach would reduce the burden on developers and designers of conversational systems.

### 2.2 P-tuning

Large Language Models (LLMs) have exhibited remarkable generalization capability when queried using *prompts* that contain examples of the task to be performed. However, the performance of LLMs varies widely depending on how such prompts are constructed. In order to overcome this issue of LLM sensitivity to the format of the prompt, multiple studies have come up with methods for automated prompt construction using discrete tokens (Lester et al., 2021) as well as soft tokens (Liu et al.,

2021).

In this work, we utilize the p-tuning approach that appends learnt *soft* tokens into the prompt that is fed to the LLM. The soft tokens traditionally do not have a mapping to words/subwords in the model vocabulary and are simply vectors optimized using gradient descent. Following the setup proposed by Liu et al. (2021), we use an LSTM model to learn and predict these soft tokens. The parameters of the LLM are frozen and only the parameters of this LSTM model are updated during p-tuning. We initialize the LSTM with random weights at the beginning of the p-tuning process and then update it during the training stage to output the optimal soft tokens. At the end of the training phase, we store these soft tokens and append them with the prompt to the LLM to get its prediction. The advantage of p-tuning is that we freeze the LM weights and update only the weights of the LSTM (14M parameters). This results in modifying only a very small fraction of the weights compared to traditional fine-tuning where all of the weights are updated.

The LM of choice in our experiments are the Megatron-GPT (Narayanan et al., 2021) models that are *decoder-only* transformers.

### 3 Experimental Setting

In this section, we describe the datasets used, the baselines we use for comparison and the evaluation metrics.

#### 3.1 Datasets

We consider two widely known datasets in the dialogue community, the Schema Guided Dialogue (SGD) dataset (Rastogi et al., 2019) and the Virtual Assistant dataset (Liu et al., 2019).

**Schema Guided Dialogue** - This dataset covers 16 domains and has over 16k annotated conversations. The domains span a variety of user actions, including setting calendars and alarms, travel booking (car rentals, flights, buses and trains), music, weather, movies, and more. The dataset also contains *multi-domain* dialogues where the utterances switch between domains. For the purpose of our experiments, we consider only the *single-domain* dialogues with 37 intents across all utterances.

**Virtual Assistant Dataset** - This dataset covers 21 domains with 64 intents across all

utterances. As the name suggests, the domains relate to user queries over a wide range of topics, including operating smart-home devices, media consumption, weather and travel. It has over 25k annotated user utterances that identify intents and slot values.

#### 3.2 Prompt Template

The prompts that we use for intent classification have the following format

$\langle v_1..v_n \rangle$  utterance intent : canonical

where  $\langle v_1, v_2, \dots, v_n \rangle$  indicate the virtual tokens.

During the training stage of p-tuning, the model is shown the entire sequence, but the loss is computed only on the *answer* which in this case is the predicted canonical form. During inference, the context to the model includes the sequence until the word "*intent:*" and the model completes the sequence with its prediction for the intent. We use 100 virtual tokens with our prompt-encoder being an LSTM model with 2 layers.

#### 3.3 Evaluation Method

**Intent Classification** Evaluating generative models for a classification task is not straightforward. This is further complicated by the fact that our model generates a canonical form identifying the intent of a given user utterance. We propose two methods to cast this generation problem to a classification setting. The difficulty arises from the fact that generated sequences very often differ from the exact gold truth sequence that the model sees as part of training. We utilize two approaches based on associating the generated canonical form to its closest label, *i.e.*, a nearest neighbor search. Once the canonical form label has been identified as the prediction, it becomes trivial to compute the classification accuracy. Since we already have a one-to-one mapping between canonical form labels and the discrete intent label, we can easily measure the performance of our model.

- Using Fasttext Embeddings (Bojanowski et al., 2016): We take the mean of all the embedding vectors of the generated canonical form and consider the vector obtained to be the representation of the whole sequence. We compute similar vectors for all the canonical form labels and consider the canonical form label that has the maximum cosine similarity

with the generated one as the model’s prediction.

- Using Sentence Transformers (Reimers and Gurevych, 2019): We use the *miniLM-QA* (Wang et al., 2020) transformer model that has been pretrained on multiple datasets on the text entailment/semantic search task, *i.e.*, given a query and a set of keys (documents/labels), it ranks the keys in order of relevance. We give as input to the model the generated canonical form (query) and the list of canonical form labels (keys). The model then returns the closest canonical form label to the generated canonical form which we consider as the prediction.

### 3.4 Baselines

We consider the following baselines for the intent classification task.

- **BERT-based finetuned model** (Intent Classification): We finetune BERT models on the datasets described in section 3.1. While some of the Megatron-GPT models we use are larger than the BERT model in terms of number of parameters, it should be noted that the LM parameters are frozen during the training stage of p-tuning and only the weights of the LSTM (14M parameters) are updated.

### 3.5 Evaluation Settings

We evaluate the performance of our model in two settings: in-domain and out-of-domain.

#### 3.5.1 In-Domain

This setting corresponds to the traditional dataset splits where the train and test sets come from similar distributions. We p-tune the Megatron-GPT models on the train set and evaluate them on the test set for intent classification.

#### 3.5.2 Out-of-Domain

In this setting, we aim to explore the generalization capability of LLMs. We hold out certain domains from the train set and use utterances from the held out domains as our test. This helps us understand how well these LLMs can generalize to unseen domains. The held out sets that we consider are:

- **Schema Guided Dialogue** (SGD): We hold out utterances corresponding to *bus bookings* and *hotel reservations* to form our test set.

The train set includes utterances from adjacent domains: flight booking and restaurant reservations. This should be a relatively easy setting for the language model to generalize to.

- **Virtual Assistant**: To make things more challenging, we hold out utterances corresponding to *operating IOT devices* and *media consumption commands* (*e.g.*, commands that are variants of "play" - play movie, play audiobook). The train set does not have utterances from similar domains and this setting is more challenging for the model.

We consider the generalization capability of the model in two modes:

- **Zero-shot**: P-tune the model on the train set and evaluate zero-shot on the unseen domain test set.
- **Few-shot**: After p-tuning on the train set, we do a second stage p-tuning on a set of  $k$  samples from the target domain. Unless otherwise noted,  $k$  here is 5, 10, 50 or 100 samples.

The few-shot paradigm may be very useful for dialogue system developers in a limited-resource setting. Developers can implement new domains using existing language models and a small set of curated examples, without the burden and expense of retraining or providing a large number of labelled samples.

## 4 Results

In this section, we review the quantitative performance of the models for intent classification.

### 4.1 Intent Classification

We compute and list the accuracy of the baselines and our p-tuned GPT model in identifying the intent given the user utterance.

#### 4.1.1 In-domain

We find that both the p-tuned GPT model as well as the BERT baseline perform very well on the standard in-domain split where both the train and test set come from the same distribution (Table 2). The classification accuracy of Megatron-GPT increases as we increase the model size. The trend of results remains consistent for both the SGD and Assistant datasets.

Model	SGD	Assistant
BERT-Large	0.88	0.91
Megatron-GPT - 345M	0.87	0.88
Megatron-GPT - 1.3B	0.91	0.92
Megatron-GPT - 5B	0.95	0.94

Table 2: Classification Accuracy on test sets of the SGD and Assistant datasets

#### 4.1.2 Out-of-Domain

The out-of-domain setting is where the advantage of using a LLM becomes apparent. It is not feasible to expect a finetuned BERT model to generalize to an unseen domain not present in the train set. Such models continue to predict that the intent belongs to one of the intent labels they see during training. The p-tuned Megatron-GPT models, on the other hand, show impressive zero-shot and few-shot generalization capabilities on the SGD dataset (Table 3). For instance, having seen intents such as "buy flight roundtrip tickets" when presented with utterances for *Flight Reservations* in training, we can expect the model to reasonably generalize to utterances from *Bus Reservations* with utterances like "Get me a return trip on the bus" with the model's prediction for the intent being "buy bus roundtrip tickets".

Mode	Bus Booking			Hotel Reservation		
	345M	1.3B	5B	345M	1.3B	5B
Zero Shot	0.755	0.762	0.787	0.379	0.448	0.467
FS - 10 samples	0.907	0.789	0.942	0.793	0.720	0.939
FS - 50 samples	0.953	0.965	0.975	0.957	0.968	0.970

Table 3: Zero-shot and Few Shot (FS) performance on the held out domains of the SGD dataset. The columns indicate the size of the Megatron-GPT model.

In the Assistant dataset, the p-tuned models face the same issue as the BERT models: they struggle to generalize to completely unseen domains and the performance is close to random (Table 4). Unlike in SGD, the held-out domains do not have sufficiently similar domains in training from which to generalize. However, the few-shot setting holds promise as the performance of the models improves with few samples. Since the held out domains have far more intents compared to the held out domains from the SGD dataset, we employ stratified sampling to ensure that the few-shot examples are representative of all intents in the domain.

Mode	IOT devices			Media Consumption		
	345M	1.3B	5B	345M	1.3B	5B
Zero Shot	0.096	0.011	0.022	0.037	0.008	0.012
FS - 10 samples	0.62	0.71	0.75	0.58	0.62	0.68
FS - 50 samples	0.69	0.83	0.87	0.67	0.86	0.89

Table 4: Zero-shot and Few Shot (FS) performance on the held out domains of the Assistant dataset. The columns indicate the size of the Megatron-GPT model.

## 5 Discussion

The results on zero-shot and few-shot settings for unseen domains demonstrate that p-tuning a LLM to have intents that are more verbose than discrete labels can be very helpful.

In this section, we analyze the impact of the structure of canonical forms, what helps the language model generalize, how sample efficient are these language models and what all this means for a developer of chatbots and dialogue systems.

### 5.1 How important is framing the right canonical form?

The phrasing of canonical forms has a significant impact on zero-shot cross domain generalization. In our initial experiments, we observed that the language models, especially the smaller ones, sometimes rely on spurious correlations to predict the intent. For instance, if the intent *Search-FlightOneWay* is mapped to the canonical form *search tickets for flight one way*, the model correlates the word *ticket* in both the user utterance and canonical form to identify the intent. When we use this model to predict the intent of user utterances related to *bus bookings* in a zero-shot manner, the model predicts that that the intent is related to a *flight booking* as most utterances in the *bus domain* contain the word *ticket*.

Mode	Accuracy		
	345M	1.3B	5B
ZS - Original	0.08	0.13	0.21
ZS - Modified	0.755	0.762	0.787

Table 5: Zero shot (ZS) performance on utterances from *Bus Bookings*. **Original** refers to having the canonical form for flight bookings as *search tickets for flight one way* which led to incorrect generalizations. **Modified** refers to having the improved canonical form for flight bookings as *search for flights one way*.

Rephrasing the canonical form for the intent

*SearchFlightOneWay* to *search for flights one way* helps the model to avoid making the spurious correlation and the performance in the zero-shot setting (Table 5) is significantly improved.

Mode	Accuracy	
	345M	1.3B
ZS - Original	0.08	0.13
FS 10 samples- Original	0.76	0.72
FS 20 samples - Original	0.84	0.87

Table 6: Zero shot (ZS) performance on utterances from *Bus Bookings*. **Original** refers to having the canonical form for flight bookings as *search tickets for flight one way* which led to incorrect generalizations. Adding a small number of examples resolves the error.

However, the few-shot setting (Table 6) alleviates this problem of sensitivity of the model to the canonical form structure. When we provide the model with a few samples from the the target domain, it learns to associate that the important words to distinguish between the domains are *flight* and *bus* and not *ticket*.

## 5.2 What do good canonical forms looks like?

Based on our experiments, a set of good canonical forms has the following properties:

- **Similarity in structure:** Use similar verbs for similar actions/domains, *e.g.*, **book** a flight, **book** bus tickets, **search** for hotels, **search** for restaurant reservations.
- **Compositional:** Using similar structures for canonical forms in similar domains naturally lends to compositionality. This makes it easier for the model to generalize in the zero-shot/few-shot setting while still allowing the developers to easily map the generations to a supported service on the backend.
- **Looks like natural language:** Since LLMs are pretrained on very large corpora of natural language, the benefit of pre-training is realized when the canonical forms resemble natural language rather than complex semantic forms. Making discrete intents look more like typical verb phrases brings out the expressive nature of language models.

Future work will explore and refine methods to automate the creation of canonical forms.

## 5.3 Do we need the entire training set for p-tuning?

We look for the fewest labelled samples for p-tuning needed to get an accuracy close to accessing the entire train set. We randomly sample  $k$  samples per intent ( $k \in 5, 10, 20, 30$ ) to form the train set the model is p-tuned on, and evaluate on the same test set as above. The train and test sets are from the in-domain setting for both SGD (Table 7) and Assistant (Table 8) datasets.

#Samples/Intent	Train Size	Accuracy		
		345M	1.3B	5B
10	370	0.77	0.81	0.827
20	740	0.82	0.83	0.844
30	1110	0.84	0.85	0.87

Table 7: Accuracy on the SGD test set when using only  $k$  samples per intent. The columns indicate the size of the Megatron-GPT model used.

#Samples/Intent	Train Size	Accuracy		
		345M	1.3B	5B
10	640	0.69	0.81	0.84
20	1280	0.74	0.84	0.91
30	1920	0.79	0.87	0.91

Table 8: Accuracy on the Assistant test set when using only  $k$  samples per intent. The columns indicate the size of the Megatron-GPT model used.

### 5.3.1 Comparison with BERT

We observe that Megatron-GPT is more sample efficient than BERT-type models, even when adjusting for the number of parameters. We use the 345M parameter version of the Megatron-GPT for a fair comparison. We finetune BERT-Large and p-tune the GPT model on the same training subset of the SGD dataset. Results are shown in Table 9.

With a small number of samples (10 per intent), both Megatron-GPT and BERT-Large have very similar performance. But with small increases in the number of labelled samples per intent in the train set, we observe that the performance of the GPT model improves faster than the BERT model.

## 5.4 What does this mean for dialogue system developers?

Task-oriented dialogue systems are challenging to create. Most common frameworks cast utterances

#Samples/Intent	Accuracy	
	345M	BERT
10	0.77	0.75
20	0.82	0.767
30	0.84	0.773

Table 9: Accuracy on the SGD test set when using only k samples per intent. MegatronGPT-345M is more sample efficient than BERT-Large.

into discrete intents and slots, but it is often not clear how to define these concepts for a given design. Such frameworks also employ NLU models that often require the creation of either rule-based grammars or a significantly large corpus of labelled samples. While ML-based approaches have come a long way, distributional shifts in the way utterances are structured can degrade performance. By leveraging LLMs, our approach reduces the effort involved in framing intents and training classifiers. Because of the flexibility in canonical form schemas and the sample efficiency of p-tuning, we argue that development of new task-oriented dialogues becomes simpler and faster. We envision a setting where a model publisher trains and releases a general-purpose p-tuned language model covering a broad set of cases. A conversation designer may then write a small set of example queries, submit a brief p-tuning job, and deploy a new application with minimal cost.

## 6 Conclusion

We explore the use of Large Language Models and p-tuning for intent classification in task-oriented dialogue systems. We show framing intent labels into more verbose forms allows LMs to exploit the underlying structure better and exhibit impressive zero-shot and few-shot generalization. We also analyze how important the phrasing of the verbose forms are and how many samples are needed to get good quantitative performance. We hope that this work on using sample efficient LLMs serves to motivate further research in making ToD systems simpler and quicker to develop.

## 7 Acknowledgements

The authors would like to thank Zhilin Wang, Virginia Adams, Sandeep Subramanian, Vlad Getselevich, Prasoon Varshney, and Jonathan Cohen for many useful discussions during the course of this

work.

## References

- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2016. Enriching word vectors with subword information. *arXiv preprint arXiv:1607.04606*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#).
- Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2021. [Gpt understands, too](#).
- Xingkun Liu, Arash Eshghi, Pawel Swietojanski, and Verena Rieser. 2019. [Benchmarking natural language understanding services for building conversational agents](#).
- Deepak Narayanan, Mohammad Shoeybi, Jared Casper, Patrick LeGresley, Mostofa Patwary, Vijay Anand Korthikanti, Dmitri Vainbrand, Prethvi Kashinkunti, Julie Bernauer, Bryan Catanzaro, Amar Phanishayee, and Matei Zaharia. 2021. [Efficient large-scale language model training on gpu clusters using megatron-lm](#).
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2019. [Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset](#).
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. [Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers](#).
- Yushi Wang, Jonathan Berant, and Percy Liang. 2015. [Building a semantic parser overnight](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1332–1342, Beijing, China. Association for Computational Linguistics.