



Institut für Maschinelle Sprachverarbeitung
Universität Stuttgart
Pfaffenwaldring 5b
70569 Stuttgart

Master Thesis

Prompt-based methods for Dialog State Tracking

Start Date: 01.08.2022

End Date: 01.02.2023

Mandava, Sai Pavan

M.Sc. Computational Linguistics

Mat.Nr.: 3461015

st169661@stud.uni-stuttgart.de

Supervisor

Prof. Dr. Thang Vu

thang.vu@ims.uni-stuttgart.de

Erklärung (Statement of Authorship)

Hiermit erkläre ich, dass ich die vorliegende Arbeit selbstständig verfasst habe und dabei keine andere als die angegebene Literatur verwendet habe. Alle Zitate und sinngemäßen Entlehnungen sind als solche unter genauer Angabe der Quelle gekennzeichnet. Die eingereichte Arbeit ist weder vollständig noch in wesentlichen Teilen Gegenstand eines anderen Prüfungsverfahrens gewesen. Sie ist weder vollständig noch in Teilen bereits veröffentlicht. Die beigelegte elektronische Version stimmt mit dem Druckexemplar überein.

(Mandava, Sai Pavan)

Non-binding Translation:

This thesis is the result of my own independent work, and any material from work of others which is used either verbatim or indirectly in the text is credited to the author including details about the exact source in the text. This work has not been part of any other previous examination, neither completely nor in parts. It has neither completely nor partially been published before. The submitted electronic version is identical to this print version.

(Mandava, Sai Pavan)

Abstract

Dialog State Tracking (DST) is an essential component of task-oriented dialogue systems, which helps the dialog system understand the user’s requirements for completing specific tasks. In this thesis, prompt-based methods for DST in task-oriented dialogue systems are explored by utilizing the MultiWOZ dataset. This approach does not rely on the pre-defined set of slots and their possible values. It can also be costly to label all the slots and values to train the DST models, especially for new domains. The prompt-based approach focuses on learning the DST model efficiently under low-resource few-shot settings. To examine the impact of prompt-based methods, a baseline pre-trained language model, SOLOIST¹, is fine-tuned to generate belief states as a word sequence. In the prompt-based approach, the SOLOIST model is fine-tuned on limited labeled training data to generate the slots directly from values. Further, multi-prompt methods² are applied to investigate the potential improvement in the slot generation performance. Experimental results show prompt-based methods significantly outperformed the baseline model under low-resource settings. Analysis of outputs shows prompt-based approach has some drawbacks due to the existing belief states annotation system in the MultiWOZ dataset. One limitation is that the prompt-based methods cannot generate multiple slots for repeated value candidates, as the slots are generated by passing values to the prompt function. The data, code and steps to reproduce results are publicly available [here](https://git.pavanmandava.com/pavan/master-thesis)³.

¹A Single Pre-trained Language Model (GPT-2) for task-oriented dialog systems

²Multi-prompt methods: *Prompt Decomposition, Prompt Ensembling, Prompt Augmentation*

³<https://git.pavanmandava.com/pavan/master-thesis>

Contents

1	Introduction	5
2	Background & Related Work	7
2.1	Dialog State Tracking (DST)	7
2.2	Pre-trained Language Models (PLMs)	7
2.3	SOLOIST Model	8
2.4	Prompt Learning	9
2.5	Prompt-based DST	10
2.6	MultiWOZ Dataset	11
3	Methods	11
3.1	Dataset	11
3.2	SOLOIST Baseline	13
3.3	Prompt-based few-shot DST	13
3.4	Multi-prompt methods for DST	17
3.5	Evaluation Metrics	19
3.6	Implementation Details	20
4	Results	20
4.1	SOLOIST Baseline	20
4.2	Prompt-based methods	21
4.3	Multi-prompt methods	22
4.3.1	Prompt Ensembling results	22
4.3.2	Prompt Augmentation results	23

5	Analysis	24
5.1	Error analysis of baseline model	24
5.2	Analysis of prompt-based methods	25
5.2.1	Value-based Prompt	25
5.2.2	Impact of Inverse Prompt	25
5.2.3	Repeated values in Belief States	26
5.2.4	Error Analysis of Value Extraction	27
6	Conclusion	29

1 Introduction

Dialog State Tracking (DST) is an essential module in dialog systems, which is responsible for tracking the user goals in the form of dialog states based on the entire dialog history. In dialog systems, “dialog states” - also known as “belief states” contains a set of $(slot, value)$ pairs for each turn of the dialog history. The $(slot, value)$ pairs hold specific pieces of information required for the dialog system to perform the task and help in generating the responses. The values of the slots can change when the user provides more information or accepts the system recommendations. Existing data-driven methods and neural models for individual dialog modules (NLU, DST, NLG) and end-to-end dialog systems show promising results, but they need large amounts of task-specific training data, which is rarely available for new tasks. These neural DST models do not generalize well on new domains with limited data (Li et al., 2021). For task-specific DST, collecting dialog state labels can be costly and time-consuming, requiring domain experts to indicate all possible $(slot, value)$ pairs for each turn of the dialog history. A typical task-oriented dialog system contains an ontology for each domain, with a pre-defined set of slots and all possible values for each domain. In real-world applications, defining all possible slots and values for DST is difficult due to the increasing number of new domains and the evolving needs of the users.

Prompt-based learning (*“pre-train, prompt, and predict”*) is a new paradigm in NLP that aims to predict the probability of text directly from the pre-trained LM. This framework is powerful as it allows the language model to be *pre-trained* on massive amounts of raw text, and by defining a new prompting function the model can perform *few-shot* or even *zero-shot* learning (Liu et al., 2021). The large pre-trained language models (PLMs) are supposed to be useful in few-shot scenarios where the task-related training data is limited, as they can be probed for task-related knowledge efficiently by using a prompt. One example of such large pre-trained language models is GPT-3 (Brown et al., 2020) - *“Language Models are Few-Shot Learners”*. Madotto et al. (2021) created an end-to-end chatbot (Few-

Shot Bot) using *prompt-based few-shot learning* learning and achieved comparable results to those of state-of-the-art. Prompting methods are particularly helpful in few-shot learning where domain-related data is limited. *Fixed-prompt LM tuning* is a fine-tuning strategy for downstream tasks, where the LM parameters are tuned with fixed prompts to help LM understand the task. This can be achieved by applying a discrete textual prompt template to the data used for fine-tuning the PLM.

Prompt-based learning for few-shot DST with limited labeled domains is still under-explored. Recently, [Yang et al. \(2022\)](#) proposed a new prompt learning framework for few-shot DST. This work designed a *value-based prompt* and an *inverse prompt* mechanism to efficiently train a DST model for domains with limited training data. This approach doesn't depend on the ontology of slots and the results show that it can generate slots by prompting the tuned PLM and outperforms the existing state-of-the-art methods under few-shot settings. In this thesis, the prompt-based few-shot methods for DST are explored by implementing the following three tasks:

1. Prompt-based few-shot DST - reproduce the results from [Yang et al. \(2022\)](#)
 - Implement prompt-based methods for DST task under few-shot settings
 - Implement a baseline model for comparing the prompt-based methods
2. Evaluation and analyses of belief state predictions
 - Evaluate the DST task using Joint Goal Accuracy (JGA) metric
 - Improvements observed from the prompt-based methods
 - Drawbacks of the prompt-based methods
3. Extend prompt-based methods to utilize various *multi-prompt* techniques
 - Can different multi-prompt techniques help the PLM better understand the DST task?
 - Evaluation of multi-prompt methods and what's the influence of various multi-prompt techniques?

2 Background & Related Work

2.1 Dialog State Tracking (DST)

Task-oriented dialog systems, both modular and end-to-end systems, are capable of handling a wide range of tasks (such as ticket booking, restaurant booking, etc.) across various domains. A task-oriented dialogue system has stricter requirements for responses because it needs to accurately understand and process the user’s message. Therefore, modular methods were suggested as a way to generate responses in a more controlled manner. A typical modular-based system uses a modular pipeline, which has four modules that execute sequentially - Natural Language Understanding (NLU), Dialog State Tracking (DST), Policy Learning (POL), and Natural Language Generation (NLG). The DST module is essential for enabling the system to comprehend the user’s requests by tracking them in the form of slots and values (belief states) at every turn. For instance, in a dialogue system that helps users book flights, the system might track slots such as destination, departure, travel date, and number of travelers. By keeping track of these slots and their values, the system can understand the user requirements and provides this information to the next module. For example, consider the user message - “*Plan a train trip to Berlin this Friday for two people*” - the DST module is supposed to extract (*slot, value*) pairs as follows: $\{(destination, Berlin), (day, Friday), (people, 2)\}$. In this thesis, the focus is on the DST module for extracting slots and values.

2.2 Pre-trained Language Models (PLMs)

Large pre-trained language models are trained on huge amounts of textual data and have achieved state-of-the-art performance in a variety of NLP tasks, such as machine translation, text classification, text generation, and summarization. These models are trained on large datasets and are able to learn the probability distribution of the words. Pre-trained language models based on transformer architectures (Vaswani et al., 2017), such as BERT (Devlin et al., 2019) and GPT (Radford et al., 2018), have also achieved state-of-the-art performance on many NLP tasks. GPT-2

(Radford et al., 2019) is a transformer-based auto-regressive language model trained on large amounts of open web text data. GPT-2 is trained with a simple objective: predict the next word, given all previous words within some text. The architecture and training objective of the PLMs plays an important role in determining their applicability to particular prompting tasks (Liu et al., 2021). For example, left-to-right auto-regressive LMs predict the next word by assigning a probability to the sequence of words. For tasks that require the PLM to generate text from *prefix* prompts (the entire prompt string followed by generated text), the left-to-right LMs tend to mesh well with the left-to-right nature of the language model.

The baseline model of this thesis, SOLOIST (Peng et al., 2021), uses a 12-layer GPT-2 for building the task-oriented dialog system. SOLOIST uses the publicly available 117M-parameter GPT-2 as initialization for task-grounded pre-training. The prompt-based methods in this thesis utilize the pre-trained SOLOIST and fine-tune it to the downstream DST task.

2.3 SOLOIST Model

SOLOIST (Peng et al., 2021) is a task-oriented dialog system that uses transfer learning and machine teaching to build task bots at scale. SOLOIST uses the *pre-train, fine-tune* paradigm for building end-to-end dialog systems using a transformer-based auto-regressive language model GPT-2 (Radford et al., 2019), which subsumes different dialog modules (i.e., NLU, DST, POL, NLG) into a single model. In a *pre-train, fine-tune* paradigm, a fixed *pre-trained* LM is adapted to different downstream tasks by introducing additional parameters and *fine-tuning* them using task-specific objective functions. In the pre-training stage, SOLOIST is initialized with the 12-layer GPT-2 (117M parameters) and further trained on large heterogeneous dialog corpora. The primary goal at this stage is to learn task completion skills such as belief state prediction (DST) and response generation. In the fine-tuning stage, the pre-trained SOLOIST model can be used to solve new tasks by just using a handful of task-specific dialogs.

In this thesis, the pre-trained SOLOIST is the baseline model for generating the belief states. For the baseline DST task, the pre-trained SOLOIST is fine-tuned on the belief predictions task for open-ended text generation. For prompt-based methods, the baseline SOLOIST is fine-tuned for generating belief state slots using prompts. The results and outputs from the baseline model are compared to the prompt-based model for detailed analyses.

2.4 Prompt Learning

Prompt-based learning (also dubbed as “*pre-train, prompt, and predict*”) is a new paradigm that aims to utilize PLMs more efficiently to solve downstream NLP tasks (Liu et al., 2021). In this paradigm, instead of adapting pre-trained LMs to downstream tasks by designing the task-specific training objectives, downstream tasks are reformulated to look more like those solved during the original LM training with the help of a textual *prompt*. To perform prediction tasks, the original input x is modified using a *template* into a textual *prompt* x' that has some unfilled slots, and then the PLM is used to probabilistically fill the unfilled information to obtain a final string z , from which the final output y can be derived. For text generation tasks, the generated answer z itself is the output y .

Name	Notation	Example
<i>Input</i>	x	I missed the bus today.
<i>Output</i>	y	sad
<i>Prompt Function</i>	$f_{prompt}(x)$	[X] I felt so [Z]
<i>Prompt</i>	x'	I missed the bus today. I felt so [Z]
<i>Answered Prompt</i>	$f_{fill}(x', z)$	I missed the bus today. I felt so sad
<i>Answer</i>	z	happy, sad, scared

Table 1: Terminology and notations of prompting methods

For example, to recognize the emotion in the text, where *input* x = “I missed the bus today.”, the *prompt function* (also called *template*) may take the form such as “[X] I felt so [Z]”. [X] takes the input text and [Z] is supposed to be generated by the LM. Then, the *prompt* x' would become “I missed the bus today. I felt so [Z]” and ask the PLM to fill the slot [Z] with an emotion-bearing word. For some text generation tasks, the answer mapping from z to y may not be required, as the generated text itself becomes output. There are two main varieties of prompts: *cloze prompts*, where the slot [Z] is to be filled in the middle of the text, and *prefix prompts*, where the input text comes entirely before [Z]. In general, for tasks that are being solved using a standard auto-regressive LM, prefix prompts tend to be more helpful, as they mesh well with the left-to-right nature of the model.

Prompt-based methods can be used without any explicit training of the LM for the downstream task, simply by taking a suitable pre-trained LM and applying the prompts defined for the task. This approach is traditionally called *zero-shot learning*. *Few-shot learning* is another approach where only a small number of data samples are used to train the language model. Prompting methods are particularly useful under few-shot settings, as there is generally not enough training data to fully specify the desired behavior. *Fixed-prompt LM tuning* is a training strategy that fine-tunes the parameters of the LM, as in the standard *pre-train fine-tune* paradigm, by using discrete prompts (*hard prompts*) to help PLM understand the downstream task. This approach can potentially lead to improvements, particularly in few-shot scenarios.

2.5 Prompt-based DST

Previous work by Lee et al. (2021) uses belief state slots in the prompts, along with the natural language descriptions of the schema for generating the corresponding values. This *slot-based* prompt DST approach uses encoder-decoder LM with a bi-directional encoder. This method relies on the known ontology of the slots and requires a lot of training data for fine-tuning PLM. In real-world applications, defining all possible slots is difficult due to the rising new domains and users’ continuous

needs. [Yang et al. \(2022\)](#) proposed a new prompt-learning framework for DST that uses values in prompts (*value-based*) and generates slots directly from the PLM. This *value-based* prompt approach does not rely on the ontology of the slots and their natural language descriptions. In task-oriented dialog systems, the prompt-based DST methods are still under-explored. In this thesis, the value-based prompt approach is applied for few-shot DST.

2.6 MultiWOZ Dataset

MultiWOZ ([Budzianowski et al., 2018](#)) is a multi-domain task-oriented dialogue dataset that contains over 10K dialogues across 8 domains. It is a fully-labeled collection of human-human written conversations and has been a widely used dataset for benchmarking DST methods. [Eric et al. \(2019\)](#) released MultiWOZ 2.1 after fixing the noisy dialog state annotations and utterances that negatively impact the performance of DST models. In this thesis, MultiWOZ 2.1 is used to benchmark both baseline and prompt-based methods.

3 Methods

This section describes the research methods and experimental setup of the work conducted in this thesis. This thesis work can be divided into the following tasks: SOLOIST baseline implementation for few-shot DST, prompt-based methods for few-shot DST, evaluation and analysis of belief state predictions, and multi-prompt methods for DST.

3.1 Dataset

The baseline and prompt-based methods are benchmarked on MultiWOZ 2.1 ([Eric et al., 2019](#)) dataset. The MultiWOZ dataset contains 8438/1000/1000 single-domain and multi-domain dialogues for training/validation/testing respectively. Each dialogue can have multiple turns and each turn can include multiple (*slot, value*) pairs.

Dialogues from only five domains (*Restaurant*, *Hotel*, *Attraction*, *Taxi*, *Train*) and one sub-domain (*Booking*) are used in the experiments, as the other two domains (*Hospital*, *Police*) only appear in the training set. To observe the performance under few-shot settings, dialogues are randomly sampled for each domain and six different data splits are created. Each data split contains all five domains and dialogues are evenly distributed. Only single-domain dialogues including booking sub-domain are picked for creating the data splits. Validation and test sets are not sampled after domain filtering. Table 2 provides data statistics and the summary of data splits used in few-shot experiments.

Data Splits	# Dialogues	# Total Turns	# (slot, value)
<i>5-dpd</i>	25	100	294
<i>10-dpd</i>	50	234	758
<i>50-dpd</i>	250	1114	3535
<i>100-dpd</i>	500	2292	7408
<i>125-dpd</i>	625	2831	9053
<i>250-dpd</i>	1125	5187	17214
<i>valid</i>	190	900	3106
<i>test</i>	193	894	3411

Table 2: Data statistics and data split summary for few-shot experiments. The term *dpd* means “*dialogues per domain*”. Each split contains dialogues for all five domains. In data split *250-dpd*, the domain “*Attraction*” contains only 125 dialogues.

In the MultiWOZ 2.1 dataset, 16 dialog slots are used to understand the user requirements. For the prompt-based experiments, these slots are converted to look like natural language words for fine-tuning the slot generation process. Table 3 lists the slots from all five domains and *booking* sub-domain.

Slots
<i>area, arrive, day, departure, destination, food, internet, leave, name, parking, people, price, stars, stay, time, type</i>

Table 3: Slots from MultiWOZ 2.1 dataset used in prompt-based experiments

3.2 SOLOIST Baseline

SOLOIST (Peng et al., 2021) is the baseline model for the prompt-based methods. SOLOIST is initialized with the 12-layer GPT-2 (Radford et al., 2019) and further trained on two task-oriented dialog corpora (Schema and Taskmaster). The task-grounded pre-training helps the SOLOIST model to solve two dialog-related tasks: *belief state prediction* and *response generation*. In the belief state predictions task, the model takes dialog history as input and generates the belief states as a sequence of words. In this thesis, for the baseline implementation, the pre-trained SOLOIST is fine-tuned on MultiWOZ 2.1 data splits to perform the belief predictions task. During inference time, the fine-tuned SOLOIST baseline doesn’t need the pre-defined set of slots and their possible values, and it uses top-K (Fan et al., 2018) and top-p or nucleus (Holtzman et al., 2020) sampling for generating the belief states. In the prompt-based DST task, the same pre-trained SOLOIST model is fine-tuned for prompt-based slot generation.

3.3 Prompt-based few-shot DST

This task aims to apply prompt-based methods proposed by (Yang et al., 2022) and reproduce the results. This task utilizes the *value-based prompt* and *inverse prompt* for fine-tuning the pre-trained SOLOIST, which can generate the belief state slots directly at inference time. The prompt-based methods are evaluated on the same data splits (Table 2) of the MultiWOZ 2.1 dataset.

Value-based prompt An intuitive idea for generating (*slot*, *value*) pairs is to use slots in prompts and generate the corresponding values (Lee et al., 2021). For

example, given the utterance - “Plan a trip to Berlin” and slot (*destination*), the prompt to the PLM could become “Plan a trip to Berlin. *destination* = [z]” and the PLM is expected to generate [z] as “Berlin”. However, this approach relies on the ontology of the slots, and the fixed set of slots can change in real-world applications. Yang et al. (2022) proposed *value-based prompt* that uses values in the prompts and generates corresponding slots. This method doesn’t require any pre-defined set of slots and can generate slots directly from the PLM. Consider this prompt template: “*belief states: value* = [v], *slot* = [s]”, the prompt function f can be of form $f(v) = [\text{dialog history}] \text{belief states: value} = [v], \text{slot} = [s]$, given the value candidate $v = \text{“Berlin”}$, the PLM can generate *slot* [s] = “*destination*”. The overall training objective of value-based prompt generation is minimizing the negative log-likelihood of slots in the training dataset D :

$$\mathcal{L} = - \sum_t^{|D|} \log P(s_t | c_t, f(v_t)) \quad (1)$$

where $P(s_t | c_t, f(v_t))$ is the probability of generating slot s_t given dialog history c_t and prompt-function f is filled with value v_t for each turn t . The loss \mathcal{L} from this step is combined with the loss from inverse prompt (next step) in order to compute the final loss. During training, the annotated values from the dataset are utilized to fill in the value-based prompts.

Inverse Prompt The *inverse prompt* mechanism (Yang et al., 2022) aims to generate the values by filling the prompts with generated slots. After generating slot s using the value-based prompt, this generated slot is presented to the inverse prompt function I . The inverse prompt aims to generate the value v' which is supposed to be close to the original value v . The prompt template for inverse prompt function can be of form $I = \text{“belief states: slot} = [s], \text{value} = [v]\text{”}$. [s] is filled with the generated slot from value-based prompt, [v] is expected to be generated by the PLM. This inverse prompt can be considered as an auxiliary task for the value-based prompt, which can improve performance by helping the PLM understand the task, especially

under low-resource scenarios. The loss function $\tilde{\mathcal{L}}$ for the inverse prompt mechanism:

$$\tilde{\mathcal{L}} = - \sum_t^{|D|} \log P(v'_t | c_t, I(s_t)) \quad (2)$$

where $P(v'_t | c_t, I(s_t))$ is the probability of generating the value v'_t by filling in the inverse prompt $I(s_t)$ with the generated slot s_t .

The final loss \mathcal{L}^* is computed by combining loss from value-based prompt \mathcal{L} and the inverse prompt loss $\tilde{\mathcal{L}}$:

$$\mathcal{L}^* = \mathcal{L} + w * \tilde{\mathcal{L}} \quad (3)$$

where w is a decimal value (0,1) which is used to adjust the influence of inverse prompt.

Training For training the prompt-based methods, the pre-trained Soloist (GPT-2 117M) is fine-tuned on value-based prompt and inverse prompt. All the MultiWOZ 2.1 data splits (Table 2) are used in the fine-tuning process in order to evaluate the performance under few-shot settings. The training strategy fixed-prompt LM tuning is adapted for tuning the prompt-based methods, where the fixed discrete prompts are used to fine-tune the parameters of the LM. Table 4 shows the prompt templates used in the fine-tuning process. The prompts are appended to the dialog history before providing them as input to the PLM and probabilistically generate the missing slots. The inverse prompt is only used during the training phase. Experiments are also performed to evaluate the influence of inverse prompt by omitting it during the training process.

Type	Prompt templates
value-based prompt	belief states: value = [v], slot = [s]
inverse prompt	belief states: slot = [s], value = [v]

Table 4: Prompt templates used during the training phase.

Testing (Slot Generation) During the testing phase, only value-based prompts are used to generate the slots. The filled prompt together with the dialog history is given as input to the PLM, and the next word with the highest probability is the generated slot. While testing, the value candidates are not known. A set of rules are applied to extract the candidate values directly from the user utterances. This sort of value extraction from utterances is previously explored by [Min et al. \(2020\)](#).

Value Extraction Value candidates are extracted directly from the dialog history and are provided to the value-based prompts for generating slots at inference time. Stanford CoreNLP Stanza ([Qi et al., 2020](#)) tool is used to first extract POS tags and named entities, a set of rules are then applied to extract the candidate values.

- Adjectives (JJ) and Adverbs (RB) are considered as possible values
 - E.g., *expensive, moderate, important*
- Consider the previous negator ‘not’
 - E.g., *not expensive, not important (= dont care)*
- Consider all named entities (name of place, time, date/day, numbers)
 - E.g., *cambridge, friday, 08:30*
- Custom set of Regex NER rules are applied for recognizing named entities
 - E.g., *restaurant names, attraction names*
- Stop words and repeated candidate values are filtered out

Prompt Decomposition For utterances where multiple (*slot, value*) pairs are expected to be predicted, directly using a single prompt for generating multiple slots is challenging. Prompt decomposition is a multi-prompt method that breaks down the prompt into sub-prompts and generates the slots separately for each sub-prompt. For each extracted value from the utterances, a value-based prompt is constructed

and the corresponding slot is generated. This sort of prompt decomposition has been explored by Cui et al. (2021) for the named entity recognition (NER) task. This approach is applied in both the training and testing phases.

Utterance:	Book a flight to Berlin on friday at 08:30.
Prompt 1:	belief states: value = <i>Stuttgart</i> , slot = [s]
Prompt 2:	belief states: value = <i>friday</i> , slot = [s]
Prompt 3:	belief states: value = <i>08:30</i> , slot = [s]

Table 5: Sub-prompts for an utterance with multiple values.

3.4 Multi-prompt methods for DST

The *value-based* prompt described in the previous section utilizes a *single* prompt for making predictions. However, a significant body of research has demonstrated that the use of multiple prompts can further improve the efficacy of prompting methods (Liu et al., 2021). There are different ways to extend the single prompt learning to use multiple prompts. This task explores two more multi-prompt learning methods: *Prompt Ensembling* and *Prompt Augmentation*. Experiments are performed on all the data splits of MultiWOZ 2.1 dataset. This task aims to answer the following questions - Q1: Can different *multi-prompt* techniques together help the PLM better understand the DST task? Q2: Can the use of multiple discrete prompts improve the performance of prompt-based model?

Prompt Ensembling This method uses multiple *value-based* prompts during the training and inference time. This idea can leverage the complementary advantages of different prompts and stabilize the performance on the downstream task. Yang et al. (2022) applied prompt ensembling to the value-based prompt by training a separate model for each prompt. Another way is to train a single model with multiple prompts as it is much faster and more memory efficient than having to train a separate model

for each prompt (Schick and Schütze, 2021). Prompt ensembling is applied only to value-based prompts, and the inverse prompt uses a single prompt. In this task, four hand-crafted prompt templates are chosen for value-based prompts and trained on a single model. The probability of generated slot s_t over multiple prompt functions is calculated by weighted averaging the probability of each prompt:

$$P(s_t | c_t) = \sum_k^{|K|} \alpha_k * P(s_t | c_t, f_k(v_t)) \quad (4)$$

where $|K|$ represents the number of prompt functions, f_k is the k -th prompt function, α_k is the weight of prompt k . During inference, a simple majority voting is used to pick the generated slot from multiple prompts. When there’s no simple majority in the generated slots, the slot with the highest probability is picked. Table 6 lists all the prompt templates used in prompt ensembling.

Prompt ensemble templates	
f_1	belief states: [v] = [s]
f_2	[v] is the value of [s]
f_3	[v] is of slot type [s]
f_4	belief states: value = [v], slot = [s]

Table 6: Prompt templates used for prompt ensemble.

Prompt Augmentation *Prompt Augmentation*, sometimes also called *demonstration learning* (Gao et al., 2021), provides a few additional *answered prompts* that can demonstrate to the PLM, how the actual value-based prompt can be answered. These demonstrations take advantage of the language models’ ability to learn repetitive patterns. The sample selection of answered prompts is hand-crafted from the training data. At inference time, the answered prompts are appended to the input before asking the PLM to generate the slot. Table 7 below provides an example of demonstration learning.

Demonstration learning	
Book a cheap flight to Frankfurt.	<i>Frankfurt</i> is of slot <i>destination</i>
Plan a train trip to Berlin.	<i>Berlin</i> is of slot <i>destination</i>
Book a taxi to the University.	<i>University</i> is of slot <i>destination</i>
Book a train to Stuttgart.	<i>Stuttgart</i> is of slot [s]

Table 7: Example prompt augmentation with answered prompts

3.5 Evaluation Metrics

The standard evaluation metric joint goal accuracy (JGA) is adopted to evaluate the belief state predictions of baseline and prompt-based methods. This metric compares all the predicted belief states to the ground-truth states at each turn. The prediction is correct only if all the predicted belief states match the ground-truth states. Both slots and values must exactly match for the belief state prediction to be correct. The rule-based methods used in value extraction can lead to many false positives in the value candidates. In order to exclude the influence of wrongly extracted values, [Yang et al. \(2022\)](#) proposed JGA*, the joint goal accuracy is computed only for the belief states where the values are extracted correctly. These evaluation metrics answer the following questions: Q1: How do the prompt-based methods perform overall compared to the Soloist baseline? Q2: Can the prompt-based methods perform better under low-resource settings? Q3: For prompt-based methods, does JGA* metric hold a better score than JGA? Q4: Can multi-prompt techniques together perform better than a single-prompt?

Analysis The belief state predictions from the SOLOIST baseline and prompt-based methods are analyzed to identify the potential improvements and drawbacks. A detailed qualitative analysis is performed on the wrong belief state predictions. Additionally, error analysis is also performed on the rule-based value extraction methods to identify the impact on the slot generation process.

3.6 Implementation Details

For the SOLOIST baseline, the existing implementation by Peng et al. (2021) is adapted to the few-shot experiments conducted in this thesis. There’s no publicly available implementation of the prompt-based methods for DST. Huggingface Transformers (Wolf et al., 2020) library is used to implement the prompt-based DST methods from scratch. Adam (Kingma and Ba, 2015) optimization algorithm is used during the fine-tuning process of both baseline and prompt-based methods. The rule-based value extraction methods are implemented using Stanford CoreNLP client stanza (Qi et al., 2020). The inverse prompt weight w in Eq. 3 is set to 0.1. The prompt weight α_k in Eq. 4 is set to the same value ($1/4$) for all the prompts used in prompt ensembling.

4 Results

This section presents the experimental results evaluated on all the methods described in the previous sections. Few-shot experiments are performed on all the data splits (see Table 2) for every method. The baseline Soloist model is evaluated only on the JGA metric. For the prompt-based methods, in addition to the JGA metric, JGA* is also computed.

4.1 SOLOIST Baseline

Table 8 shows the results of the baseline model under few-shot experiments. Experimental results show the baseline model performed poorly and struggled to generate belief states under low-resource settings (*5-dpd*, *10-dpd*, *50-dpd*). Under low-resource data splits, the limited size of data samples made it challenging for the baseline to generate unseen belief states. The results also show that more data may be necessary as the model achieves better results on the data splits with a higher number of data samples (*125-dpd*, *250-dpd*).

Data split (# dialogs)	JGA
<i>5-dpd</i> (25)	9.06
<i>10-dpd</i> (50)	14.20
<i>50-dpd</i> (250)	28.64
<i>100-dpd</i> (500)	33.11
<i>125-dpd</i> (625)	35.79
<i>250-dpd</i> (1125)	40.38

Table 8: Few-shot experimental results of the SOLOIST baseline model. The term “*dpd*” stands for “*dialogues per domain*”.

4.2 Prompt-based methods

Table 9 shows the results of the prompt-based model under few-shot experiments. Only a single value-based prompt is used in these experiments. Experimental results show the prompt-based model significantly outperforms the baseline model in all data splits. For low-resource data splits like *5-dpd*, *10-dpd*, and *50-dpd*, the prompt-based model shows a substantial improvement over the baseline, achieving an increase in the JGA metric by 21, 28, and 18 points respectively.

Data split (# dialogs)	JGA	JGA*
<i>5-dpd</i> (25)	30.66	71.04
<i>10-dpd</i> (50)	42.65	86.43
<i>50-dpd</i> (250)	47.06	91.63
<i>100-dpd</i> (500)	47.74	92.31
<i>125-dpd</i> (625)	46.49	91.86
<i>250-dpd</i> (1125)	47.06	92.08

Table 9: Few-shot experimental results from the prompt-based model. Only a single *value-based prompt* is used. The term “*dpd*” stands for “*dialogues per domain*”.

The prompt-based results also show that by increasing the number of data samples in experiments, the model only achieved minor performance improvements. For example, the prompt-based methods perform nearly identical on the data splits 50-dpd and 250-dpd. This suggests the prompt-based approach understands the DST task better under low-resource scenarios. The higher values of the JGA* metric across all data splits indicate the potential drawbacks of the rule-based value extraction methods.

4.3 Multi-prompt methods

4.3.1 Prompt Ensembling results

Table 10 shows the results of prompt ensembling under few-shot settings. The results from the prompt ensemble show a slight improvement over a single value-based prompt. Contrary to expectations, the prompt ensemble model did not show significant performance improvement on the JGA metric. The results also show that the performance of the prompt ensemble model is similar when trained on large data splits, i.e. 50-dpd, 100-dpd, 125-dpd, 250-dpd.

Data split (# dialogs)	JGA	JGA*
5-dpd (25)	30.09	69.23
10-dpd (50)	42.84	86.99
50-dpd (250)	47.62	91.74
100-dpd (500)	48.08	92.87
125-dpd (625)	46.96	92.08
250-dpd (1125)	48.08	92.87

Table 10: Few-shot experimental results from prompt ensembling (multi-prompt method). Four *value-based prompts* are used at training and inference time. The term “dpd” stands for “*dialogues per domain*”.

4.3.2 Prompt Augmentation results

Table 11 shows the results of prompt augmentation under few-shot settings.

Data split (# dialogs)	JGA	JGA*
5-dpd (25)	27.8	68.1
10-dpd (50)	38.91	74.43
50-dpd (250)	39.52	82.81
100-dpd (500)	42.42	83.71
125-dpd (625)	40.16	82.92
250-dpd (1125)	41.52	85.07

Table 11: Experimental results from demonstration learning (multi-prompt method). The term “dpd” stands for “*dialogues per domain*”.

Prompt augmentation (also called *demonstration learning*) provides additional context to the language models in the form of “*answered prompts*” at inference time. The hand-crafted answered prompts are supposed to help the language model understand the DST task better and generate accurate responses. Table 11 shows the experimental results from the prompt augmentation method. Results show that the demonstration learning struggled to generate the belief states accurately. The performance is inadequate across all data splits when compared to other prompt-based methods. Only a limited number of answered prompts can be provided to the GPT-2 LM due to the max input sequence length of 1024, which led to bias during the slot generation process.

Overall, the multi-prompt methods (prompt ensembling and prompt augmentation) struggled to improve the performance of the DST task. However, the prompt ensembling approach with multiple value-based prompts showed minor improvements over a single value-based prompt.

5 Analysis

5.1 Error analysis of baseline model

Wrong belief state predictions	
Dialog History	user: we need to find a guesthouse of moderate price.
	system: do you have any special area you would like to stay? or possibly a star request for the guesthouse?
	user: i would like it to have a 3 star rating.
True belief states	<i>(type, guesthouse) (pricerange, moderate) (stars, 3)</i>
Generated states	<i>(parking, yes) (stars, 3)</i>
Dialog History	user: i need an expensive place to eat in the west.
	system: is there a specific type of food you would like?
	user: yes, i would like eat indian food.
True belief states	<i>(area, west) (food, indian) (pricerange, expensive)</i>
Generated states	<i>(area, west) (food, indian) (pricerange, cheap) (area, east)</i>

Table 12: Examples of a wrongly generated belief states by the baseline model.

The belief predictions task of the SOLOIST baseline utilizes *top-k* and *top-p* sampling in order to generate the *(slot, value)* pairs. As the baseline model uses open-ended generation, it is susceptible to generating random slot-value pairs that are not relevant. The baseline performance was also affected by the repeated slot generations and in some cases incorrect values. Table 12 shows examples of some of the errors made by the baseline model. In the first example, the baseline system missed two true states and generated a totally incorrect belief state. For the second example, the slot *area* is repeated with a different value and the value for the slot *pricerange* is incorrectly generated.

5.2 Analysis of prompt-based methods

5.2.1 Value-based Prompt

Dialog History	user: I need to be picked up from pizza hut city centre after 04:30
True belief states	(<i>departure, pizza hut city centre</i>) (<i>leave, 04:30</i>)
Generated states	(<i>destination, pizza hut city centre</i>) (<i>arrive, 04:30</i>)
Dialog History	user: I need a taxi to arrive by 16:45 to take me to the parkside police station.
True belief states	(<i>destination, parkside police station</i>) (<i>leave, 16:45</i>)
Generated states	(<i>destination, parkside police station</i>) (<i>arrive, 16:45</i>)

Table 13: Incorrect belief states generated by value-based prompt.

The value-based prompt trained on low-resource data splits (i.e., *5-dpd*, *10-dpd*) struggled to distinguish between the slots like *departure* vs *destination* and *leave* vs *arrive*. In many instances, it wrongly generated the slot *destination* instead of *departure* and slot *arrive* instead of *leave*. Table 13 shows some example outputs where the slots are generated incorrectly. In both examples, the slot *arrive* is incorrectly generated. These incorrect slot generations are due to the limited training available for these examples. Overall, the prompt-based methods perform significantly better than the baseline even under low-resource settings, due to the constrained generation of slots using value-based prompts.

5.2.2 Impact of Inverse Prompt

The inverse prompt mechanism can be considered as an auxiliary task that complements the value-based prompt and helps generate the slots more accurately. Experiments are performed to analyze the impact of the inverse prompt by omitting it while training the value-based prompt.

Data split (# dialogs)	w/o inverse prompt		with inverse prompt	
	JGA	JGA*	JGA	JGA*
5-dpd (25)	26.81	64.25	30.66	71.04
10-dpd (50)	41.1	82.35	42.65	86.43
50-dpd (250)	45.7	90.7	47.06	91.63
100-dpd (500)	47.74	91.86	47.74	92.31
125-dpd (625)	45.02	90.61	46.49	91.86
250-dpd (1125)	46.15	91.4	47.06	92.08

Table 14: Experimental results showing value-based prompt performance with and without inverse prompt mechanism while training. The term “dpd” stands for “dialogues per domain”, “w/o” means “without”.

The results from table 14 show the inverse prompt mechanism helped improve the performance of the prompt-based model, especially under low-resource data splits (i.e., 5-dpd, 10-dpd). For the data split 5-dpd, where the training data is very limited, the inverse prompt brings noticeable improvements by achieving a 5% increase in performance on JGA and a 7% increase on JGA* metric. For the data splits with a higher number of data samples (i.e., 100-dpd, 125-dpd, 250-dpd), only minor improvements can be observed in the performance when the inverse prompt is included in the training. The experimental results conclude that the inverse prompt mechanism has a noticeable impact on the prompt-based model under extremely low-resource settings.

5.2.3 Repeated values in Belief States

In the prompt-based methods, the value-based prompt takes the candidate values and generates the corresponding slots. The belief states can have repeated values in the (slot, value) pairs. In other words, the user requirements may lead to having repeated values in the belief state (slot, value) pairs.

History	user: hi, can you help me find a 3 star place to stay?
	system: Is there a particular area or price range you prefer?
	user: how about a place in centre of town that is of type hotel.
	system: how long would you like to stay, and how many people?
True states	user: I'll arrive on saturday and stay for 3 nights with 3 people.
	(area, centre) (stars, 3) (type, hotel) (day, saturday) (stay, 3) (people, 3)

Table 15: An example instance with repeated values in the (slot, value) pairs

The data instance listed in table 15 contains multiple (slot, value) pairs. For the belief slots *stars*, *stay*, and *people*, the value is the same. The value-based prompt can only generate one slot with the repeated value 3. This is a main drawback of the value-based prompt under the existing belief state annotation system.

5.2.4 Error Analysis of Value Extraction

History	user: I want a place to stay that has free wifi and free parking.
	system: do you have a preference for area or price range?
True states	user: I don't have a preference. I want a hotel not guesthouse.
	(area, <u>dont care</u>) (internet, <u>yes</u>) (parking, <u>yes</u>) (price, <u>dont care</u>) (type, hotel)
Extracted values	free, hotel
History	user: I need a guesthouse with free wifi please.
	system: which area would you prefer?
True states	user: I also need free parking, and I prefer a 4 star place.
	(internet, <u>yes</u>) (parking, <u>yes</u>) (stars, 4) (type, guesthouse)
Extracted values	free, guesthouse, 4

Table 16: Example data instances where values cannot be extracted (underlined).

At inference time, the value-based prompt requires the belief state values in order to generate slots. The value extraction methods apply a set of rules on POS tags and named entities to extract value candidates directly from utterances. The rule-based extraction has an accuracy of 79% over all the values and a turn-level accuracy of 49% on the test split. Table 16 highlights instances where the values cannot be extracted using rule-based methods. In the first example, the value “*dont care*” does not appear in the utterances and cannot be extracted from POS tags. When the user requirement is *free* wifi or *free* parking, the existing annotation system for belief states considers it as the value “*yes*”. The rule-based methods adopted for value extraction can only extract the value “*free*” from the utterances. The values “*dont care*” and “*yes*” also occur twice in the examples shown in table 16, as described in the previous section (sec 5.2.4) the value-based prompt cannot handle repeated values for slot generation.

History	user: I kind of need some help finding a nice hotel in the north part of town.
True states	(<i>area</i> , <i>north</i>) (<i>price</i> , <i>expensive</i>) (<i>type</i> , <i>hotel</i>)
Extracted values	<u><i>kind</i></u> , <u><i>nice</i></u> , <i>hotel</i> , <i>north</i>

Table 17: Example instance where values are extracted incorrectly (underlined).

After extracting POS tags using the CoreNLP client, all the *adjectives* and *adverbs* from the utterances are considered as candidate values. This approach can lead to false positives in value candidates. Table 17 shows a data instance where some values are extracted incorrectly. The existing annotation system associates the user utterance “a nice hotel” with the value “*expensive*” for slot price, this cannot be achieved under the current rule-based methods. The value “*kind*” is also extracted incorrectly due to considering all the *adverbs* as possible values. The rule-based value extraction methods used in this thesis have limitations, which led to the performance degradation of prompt-based DST.

6 Conclusion

This work explored the use of prompt-based methods for dialog state tracking (DST) in task-oriented dialogue systems. The prompt-based methods, which include value-based prompt and inverse prompt, learned the DST task efficiently under low-resource few-shot settings without relying on the pre-defined set of slots and values. Experiments show that the prompt-based methods significantly outperformed the baseline Soloist model under low-resource settings. Analysis of generated belief states shows the prompt-based approach has some limitations. Additionally, multi-prompt methods such as prompt ensembling and prompt augmentation are applied to the DST task. Results show that the prompt ensemble model achieved minor improvements, and the performance of prompt augmentation is limited due to the bias in answered prompts. Error analysis of value extraction highlights the limitations of the rule-based methods. Further research is necessary to overcome the limitations of prompt-based methods and value extraction methods.

References

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf>.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. MultiWOZ - a large-scale multi-domain Wizard-of-Oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1547. URL <https://aclanthology.org/D18-1547>.
- Leyang Cui, Yu Wu, Jian Liu, Sen Yang, and Yue Zhang. Template-based named entity recognition using BART. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1835–1845, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.findings-acl.161. URL <https://aclanthology.org/2021.findings-acl.161>.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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, June 2019.

- Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL <https://aclanthology.org/N19-1423>.
- Mihail Eric, Rahul Goel, Shachi Paul, Abhishek Sethi, Sanchit Agarwal, Shuyang Gao, and Dilek Hakkani-Tür. Multiwoz 2.1: Multi-domain dialogue state corrections and state tracking baselines. *CoRR*, abs/1907.01669, 2019. URL <http://arxiv.org/abs/1907.01669>.
- Angela Fan, Mike Lewis, and Yann Dauphin. Hierarchical neural story generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-1082. URL <https://aclanthology.org/P18-1082>.
- Tianyu Gao, Adam Fisch, and Danqi Chen. Making pre-trained language models better few-shot learners. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3816–3830, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.295. URL <https://aclanthology.org/2021.acl-long.295>.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rygGQyrFvH>.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- Chia-Hsuan Lee, Hao Cheng, and Mari Ostendorf. Dialogue state tracking with a language model using schema-driven prompting. *CoRR*, abs/2109.07506, 2021. URL <https://arxiv.org/abs/2109.07506>.

- Shiyang Li, Semih Yavuz, Kazuma Hashimoto, Jia Li, Tong Niu, Nazneen Rajani, Xifeng Yan, Yingbo Zhou, and Caiming Xiong. Coco: Controllable counterfactuals for evaluating dialogue state trackers. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=eom0IUrf__F.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *CoRR*, abs/2107.13586, 2021. URL <https://arxiv.org/abs/2107.13586>.
- Andrea Madotto, Zhaojiang Lin, Genta Indra Winata, and Pascale Fung. Few-shot bot: Prompt-based learning for dialogue systems. *CoRR*, abs/2110.08118, 2021. URL <https://arxiv.org/abs/2110.08118>.
- Qingkai Min, Libo Qin, Zhiyang Teng, Xiao Liu, and Yue Zhang. Dialogue state induction using neural latent variable models. In Christian Bessiere, editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 3845–3852. International Joint Conferences on Artificial Intelligence Organization, 7 2020. doi: 10.24963/ijcai.2020/532. URL <https://doi.org/10.24963/ijcai.2020/532>. Main track.
- Baolin Peng, Chunyuan Li, Jinchao Li, Shahin Shayandeh, Lars Liden, and Jianfeng Gao. SOLOIST: building task bots at scale with transfer learning and machine teaching. *Transactions of the Association for Computational Linguistics*, 9:807–824, 2021. doi: 10.1162/tacl.a.00399. URL <https://aclanthology.org/2021.tacl-1.49>.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. Stanza: A python natural language processing toolkit for many human languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 101–108, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-demos.14. URL <https://aclanthology.org/2020.acl-demos.14>.

- Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Timo Schick and Hinrich Schütze. Few-shot text generation with natural language instructions. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 390–402, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.32. URL <https://aclanthology.org/2021.emnlp-main.32>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online, October 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-demos.6. URL <https://aclanthology.org/2020.emnlp-demos.6>.
- Yuting Yang, Wenqiang Lei, Juan Cao, Jintao Li, and Tat-Seng Chua. Prompt learning for few-shot dialogue state tracking. *CoRR*, abs/2201.05780, 2022. URL <https://arxiv.org/abs/2201.05780>.