

Contextual Dynamic Prompting for Response Generation in Task-oriented Dialog Systems

Sandesh Swamy
AWS AI Labs
sanswamy@amazon.com

Chacha Chen*
University of Chicago
chacha@uchicago.edu

Narges Tabari
AWS AI Labs
nargesam@amazon.com

Rashmi Gangadharaiah
AWS AI Labs
rgangad@amazon.com

Abstract

Response generation is one of the critical components in task-oriented dialog systems. Existing studies have shown that large pre-trained language models can be adapted to this task. The typical paradigm of adapting such extremely large language models would be by fine-tuning on the downstream tasks which is not only time-consuming but also involves significant resources and access to fine-tuning data. Prompting (Schick and Schütze, 2020) has been an alternative to fine-tuning in many NLP tasks. In our work, we explore the idea of using prompting for response generation in task-oriented dialog systems. Specifically, we propose an approach that performs *contextual dynamic prompting* where the prompts are learnt from dialog contexts. We aim to distill useful prompting signals from the dialog context. On experiments with MultiWOZ 2.2 dataset (Zang et al., 2020), we show that contextual dynamic prompts improve response generation in terms of *combined score* (Mehri et al., 2019a) by 3 absolute points, and a massive 20 points when dialog states are incorporated. Furthermore, human annotation on these conversations found that agents which incorporate context were preferred over agents with vanilla prefix-tuning.

1 Introduction

With the advent of large language models (LLMs), a vast majority of NLP tasks, including dialog systems, further fine-tune these LMs for their downstream tasks. Although these approaches provide substantial improvements over traditional task-specific models (Ham et al., 2020; Hosseini-Asl et al., 2020; He et al., 2022), it is a time consuming process that also involves significant use of energy/resources in the form of compute. These approaches also require tuning and storing parameters for each downstream task.

A more recent line of work, explores “prompting” LLMs to elicit the necessary knowledge required for the downstream tasks (Shin et al., 2020; Gao et al., 2020; Schick and Schütze, 2020; Petroni et al., 2019; Lee et al., 2021; Zhu et al., 2022). Prompts composed of tokens or short pieces of text (*discrete* prompts) inserted at the end of the input examples. These prompts are typically manually defined based on the specific downstream task. The main motivation behind these approaches stems from the idea that the large corpora that these language models are trained on contain relevant information which is pertinent to the task on hand.

Adapter-tuning was proposed as an alternate approach to fine-tuning. These methods only train task-specific layers that are inserted within pre-trained LMs. Such a lightweight approach that add about 4% task-specific parameters has shown to obtain comparable performances to their fine-tuning counterparts (Rebuffi et al., 2017; Houlsby et al., 2019; Lin et al., 2020a).

Drawing inspiration from prompting, *prefix-tuning* approaches (Li and Liang, 2021) were proposed as another alternative to fine-tuning. These approaches pre-pend a sequence of task-specific *continuous* vectors (aka prefix-) to the input. In contrast to prompting, the prefix consists of free parameters that do not correspond to actual real tokens. Such an approach is more prevalent since it only optimizes the prefix and does not tune parameters of the entire LM.

Most of the existing approaches use static prompts, i.e., the same set of tokens are used as “prompt tokens” regardless of input. However, we believe that taking context into consideration is critical especially in response generation since the current response has to fit not only the domain but also the information being requested in previous turns. For example: In the MultiWOZ dataset, if a customer asks about train bookings, the agent response has to restrict itself to that particular do-

* Work done during an internship at AWS AI Labs

main. To address this problem, we explore the idea of generating input-dependent or *contextual* prompts. We want the prompts to capture and encode different signals for different turns of dialogs depending on the context, hence, we call our approach *dynamic context prompting*. This way, we hope to distill useful signals into the prompts and provide the model with adequate signals to generate a desired system response. In this work, we explore the potential of using dialog context within a prefix tuning approach for the task of response generation in task-oriented dialog systems (TOD). The contributions of this paper are summarized as:

- we propose a context-dependent prefix-tuning method for dialog response generation in TOD systems.
- to illustrate the benefits of such an approach, we conduct experiments on the MultiWOZ dataset. We show that our model significantly outperforms the original task-dependent design of the prefix-tuning method.

2 Related Work

2.1 Dialog Generation

With the prevalence of LLMs, the quest for an answer to “how do we effectively adapt such models for dialog generation?” has been on the forefront of researchers’ minds in the dialog community. For task-oriented dialogs, fine-tuning large pre-trained models such as GPT-2 or T5 has made great progress on benchmarks recently (Ham et al., 2020; Hosseini-Asl et al., 2020). Built upon these advances, more recent line of work investigates the effectiveness of using multi-task learning (Su et al., 2021; Lin et al., 2020b; Yang et al., 2021), or pre-training the model on external dialog corpora (Peng et al., 2021; Liu et al., 2021). More recently, prompting has been used to address the sub-task of dialog state tracking (Lee et al., 2021; Zhu et al., 2022). Different from those works, we focus on the task of dialog response generation.

2.2 Prompt-based Learning

As an alternative to the fine-tuning paradigm, prompting involves a sequence of tokens appended to the input text, which can then induce the model to engage in a certain behavior suited to the task. Since the release of GPT-2 (Radford et al., 2018, 2019; Brown et al., 2020), many prompt-related papers have emerged. Most of the leading approaches

in prompting use **task-specific prompts**, ranging from discrete prompts (Shin et al., 2020; Gao et al., 2020; Schick and Schütze, 2020; Petroni et al., 2019) to continuous “soft prompts” (Li and Liang, 2021; Lester et al., 2021). These methods have a fixed prompt for each task. However, in dialog systems specifically, the context varies for every turn. In our work, we aim to design prompts which are **context-dependent**.

3 Problem Statement

Response generation is one of the tasks carried out in dialog systems usually in addition to dialog state tracking (DST). Given a dialog context (previous turns between the system and the user) $C = [u_1, s_1, \dots, u_{n-1}, s_{n-1}]$ and the current user utterance u_n , the goal of response generation is to generate system response s_n . Note that in the actual task, we generate *delexicalized* system responses, given all the groundtruth previous turns as input, following previous works (Hosseini-Asl et al., 2020; Wen et al., 2015).

Techniques mentioned in (Ham et al., 2020; Hosseini-Asl et al., 2020) rely on fully fine-tuning LLMs to carry out this task. In contrast, our approach builds on the prefix-tuning framework, but incorporates dialog context, C , as an additional signal for the prefix tokens. As a supplement to context C , we added dialog state information D (up to the current turn) to further help response generation.

4 Contextual Dynamic Prompting Framework

4.1 Prefix-tuning for Response Generation

Our work is built on top of prefix tuning for generation tasks (Li and Liang, 2021), which adds a fixed set of tunable prefix tokens/prompts to the original input x to obtain a new input, [PREFIX; x]. Following the denotation in (Li and Liang, 2021), we use $P_\theta[i, :]$ to denote the i_{th} prefix. $P_\theta[i, :]$ is generated by:

$$P_\theta[:, :] = MLP_\theta(P'), \quad (1)$$

where P' is a fixed smaller matrix as input to a feedforward neural network (MLP_θ). The training objective of prefix-tuning is same as fine-tuning, i.e., the following log-likelihood objective:

$$\max_{\theta} \log p_{\phi}(y|x),$$




Figure 1: The figures above indicate the differences between the vanilla prefix-tuning approach compared to our approach. In both these variants, only the prefix tokens are tuned.

where y is the decoder output and x is the input. θ represents the trainable parameters in the prefix tuning feedforward neural network and ϕ denotes all other parameters that include the frozen parameters of the large language model.

For our task of response generation, we concatenate the prefix with the dialog context and the current user utterance as input [PREFIX; $u_1, s_1, \dots, u_{n-1}, s_{n-1}, u_n$]. The target output is the system response s_n as seen in Figure 1 (a).

We adopt T5 (Raffel et al., 2020) as the pre-trained language model. T5 employs an encoder-decoder framework which is prevalent in seq2seq tasks (Sutskever et al., 2014; Cho et al., 2014).

4.2 Contextual Prefix-tuning

In vanilla prefix-tuning, the parameters of the prefix are fixed after training for any particular task to be reused. However, a dialog system involves having multiple turns of conversation between a system and the user. It is imperative in such systems to dynamically incorporate contextual information to carry out a meaningful conversation with the user. We explore how we can distill the dialog context information into the prefix with a prompt encoder.

Different from the original design, we want to encode additional signals into the prefix that differs for each input instances. In other words, we want to generate **contextual prefix** or **contextual dynamic prompts**.

Formally, we modify the equation (1) as follows:

$$P_\theta[:, :] = MLP_\theta(\text{encoder}(C)), \quad (2)$$

where $C = [u_1, s_1, \dots, u_{n-1}, s_{n-1}]$ represents the dialog context. We first obtain the representation of the dialog context by feeding C into a T5 encoder which is kept frozen as shown in Figure 1 (b). Subsequently, we use the prompt encoder, i.e., the feedforward neural network, to get the prefix. The generated prefix P_θ is then concatenated with only the current user utterance. Instead of concatenating the whole context as the input to the T5 decoder,

we first distill the signal into the prefix tokens. As a consequence of freezing the T5 encoder which generates the context representation, we still have the same number of tunable parameters as the original prefix-tuning framework.

4.3 Input-dependent Prefix-tuning with Dialog State

In most task-oriented dialog systems, we also have access to the dialog state at every turn in addition to dialog context. The dialog state has information such as requested slots and filled slots at every turn. We provide the dialog state D in addition to the context C to obtain contextual dynamic prompts. As a result, we will now modify equation (2) as:

$$P_\theta[:, :] = MLP_\theta(\text{encoder}(C; D_{n-1})), \quad (3)$$

we only provide the most recent dialog state D_{n-1} which is an amalgamation of all previous dialog states $D_{<n-1}$.

5 Experimental Settings

5.1 Dataset and Metrics

We evaluate our proposed framework and model on the MultiWOZ 2.2 dataset (Zang et al., 2020; Budzianowski et al., 2018) which is a large-scale, multi-domain, human-human task-oriented dialog dataset collected via the Wizard-of-Oz framework where one participant plays the role of the system. It consists of seven domains including hotel, restaurant, attraction, train, taxi, hospital, and police, and an additional domain general for acts such as greeting or goodbye. Due to its multi-domain setting, complex ontology, and flexible human expressions, developing dialog systems on MultiWOZ is extremely challenging. The training data contain 8437 dialogs, the dev and test set contain 1000 dialogs each.

We use four evaluation metrics: **BLEU** (Papineni et al., 2002), **Inform**, and **Success** rates, and **combined score**. *Inform* measures whether the

	MultiWOZ 2.2						
	BLEU	Inform	Success	Combined Score	Av. len.	#uniq. words	#uniq. 3-grams
Prefix-Tuning	19.19	54.7	48.0	70.54	13.83	245	1671
Prefix-Tuning (with DS)	19.36	51.8	47.0	68.76	13.08	231	1626
Contextual Dynamic Prompt	19.16	58.1	50.5	73.46	14.16	231	1532
Contextual Dynamic Prompt (with DS)	17.94	77.2	68.8	90.94	14.02	282	2390

Table 1: **Performance Comparison.** All model performance are based on features from all modalities. Contextual Dynamic Prompt (with DS) has the best performance in combined score.

system provides an appropriate entity and *Success* measures whether the system answers all the requested attributes. Specifically, the *Inform* rate relates to attributes that allow the user to constrain database searches, e.g., restaurant location or price range (the informational slots) and the *Success* rate focuses on request-able slots, that can be asked by the user, e.g., phone number. Both are calculated on the level of dialogs. The combined score is calculated following (Mehri et al., 2019b) as $BLEU + 0.5 * (Inform + Success)$. We followed a standard script¹ to report different measures.

5.2 Human Evaluation

We chose a 10% subset of the evaluation set (randomly shuffled) conversations with a total of 728 turns across them and provided annotators with the responses generated by each of the methods described in section 4. Annotators were asked to rate each agent on a turn-level and to also pick the agent which carried out the best conversation. If annotators felt more than one agent did well, they could choose multiple agents. The agent numbers, when provided to annotators, were shuffled to avoid bias. Each agent is described as:

- Agent 1: Incorporates only prefix-tuning
- Agent 2: Incorporates prefix-tuning with Dialog State
- Agent 3: Incorporates contextual dynamic prompts
- Agent 4: Incorporates contextual dynamic prompts with Dialog State

When annotating on turn level, from these 728 turns, we saw that the agents tied on 596 occasions, agent 1 had outright win on 12 occasions, agent 2 on 22, agent 3 on 33 occasions, and agent 4 on 65 occasions. This shows that our technique of using contextual dynamic prompts for generating responses is effective (Examples in Appendix B).

¹https://github.com/Tomiinek/MultiWOZ_Evaluation

Additionally, on the conversation level, we noticed that across 100 conversations, 37 were tied, and agents 3 and 4 were preferred in a total of 53 conversations confirming our hypothesis that incorporating context into prompts leads to better responses. We request readers to refer to Appendix A and B for more details about the annotation task.

6 Results

As shown in Table 1, contextual dynamic prompting with dialog states obtains a combined score of *90.94*, a 20 point jump from our baseline (prefix-tuning). In addition, even though we can’t explicitly explain the drop in BLEU, the massive jumps in both success and inform suggest more transparency and coherence for the responses generated by the input-dependent prefix-tuning as these metrics focus on quality of informational and request-able slots in each turn. When comparing our results with the human annotations, we also see that *Agent 4* - which uses contextual dynamic prompting - wins 38 conversations (out of 100). This is majority of wins compared to *Agent 1* winning only 3 conversations, and *Agent 2* winning 7. This further emphasized that adding contextual dynamic prompts leads to better quality of responses. We show some examples from the human annotation task where *Agent 4* was chosen as the winning agent in Tables 2 and 3 below.

7 Conclusion

In our work, we proposed an approach that performs *contextual dynamic prompting* where prompts are learnt from dialog contexts with the goal of distilling useful prompting signals. In our experiments, we showed that contextual dynamic prompts improve response generation in terms of *combined score* (Mehri et al., 2019a) by 3 points, and by 20 points when *dialog states* are incorporated compared to the baseline. Our technique does not expose the models to additional knowledge sources. Human annotation on these conversations

Turn num	User turn	Agent 1 response	Agent 2 response	Agent 3 response	Agent 4 response
1	i am looking for a place to dine. the restaurant should be expensive and should serve chinese food.	there are [choice] [pricerange] [food] restaurants in town. do you have a preference for area or price range?	i have [choice] [pricerange] [food] restaurants. do you have a preference for area?	there are [choice] [pricerange] [food] restaurants. do you have a price range in mind?	there are [choice] [food] restaurants in the [pricerange] price range. is there a particular area you are interested in?
2	not in any specific area please. i need to book a table for 3 at 14:00 on saturday.	booking was successful. the table will be reserved for 15 minutes. reference number is : [ref]. is there anything else i can help you with?	booking was successful. the table will be reserved for 15 minutes. reference number is : [ref].	i was able to book you at [name] for [bookday] at [booktime]. your reference number is [ref]. is there anything else i can help you with?	i have made a reservation for you at [name]. the table will be reserved for 15 minutes. the reference number is [ref]. is there anything else i can help you with?
3	how about 13:00 then?	booking was successful. the table will be reserved for 15 minutes. reference number is : [ref].	booking was successful. the table will be reserved for 15 minutes. reference number is : [ref].	booking was successful. the table will be reserved for 15 minutes. reference number is : [ref]. is there anything else i can help you with?	ok, i have you booked at [name]. your reference number is [ref]. is there anything else i can help you with?
4	what is the reference number for my booking?	your reference number is [ref]. is there anything else i can help you with?	the reference number is [ref]. is there anything else i can help you with?	your reference number is [ref]. is there anything else i can help you with?	the reference number is [ref].
5	no thank you, that is all i needed!	you're welcome. have a great day!	you're welcome. have a great day!	you're welcome. have a great day!	you're welcome. have a great day!

Table 2: An example where Agent 4 was picked as the best performing agent in a conversation.

Turn num	User turn	Agent 1 response	Agent 2 response	Agent 3 response	Agent 4 response
1	can you help me find a restaurant that serves african food with a moderate price range please	there are no [food] restaurants in the [pricerange] price range. would you like to try a different type of food?	there are no [food] restaurants in the [pricerange] price range. would you like to try a different type of food?	there are no [food] restaurants in the [pricerange] price range. would you like to try a different type of food?	[name] is a [food] restaurant in the [area]. would you like to try it?
2	yes how about asian food in the same price range?	there are [choice] [food] restaurants in the [area] and [choice] in the [area]. do you have a preference?	there are [choice] [food] restaurants in the [pricerange] price range. do you have a preference for area?	there are [choice] [pricerange] [food] restaurants in the [area]. do you have a price range in mind?	[name] is a [food] restaurant in the [pricerange] price range.
3	will you please book a table for 6 at 13:45 on friday?	booking was successful. the table will be reserved for 15 minutes. reference number is : [ref]. is there anything else i can help you with?	booking was successful. the table will be reserved for 15 minutes. reference number is : [ref].	i was able to book you at [name] for [bookday] at [booktime]. your reference number is [ref]. is there anything else i can help you with?	i have made a reservation for [bookpeople] at [booktime] on [bookday] at [booktime]. the table will be reserved for 15 minutes. the reference number is [ref]. is there anything else i can help you with?
4	that's all i needed. thanks a lot for the help!	you're welcome. have a great day!	you're welcome. have a great day!	you're welcome. have a great day!	you're welcome. have a great day!

Table 3: An example where Agent 4 was picked as the best performing agent in a conversation.

found that agents which incorporate context into prefix-tuning. prompts were preferred over agents with vanilla

Limitations

While our work explores a new technique of contextual dynamic prompts for response generation, we carried out our experiments on a dataset which is in the English language. A potential limitation of this work would be the transfer of our findings on an English dataset to a multi-lingual dataset or a mono-lingual dataset on a language other than English. We plan to address this in our future work and also request the help of the research community in doing so.

Acknowledgements

We would like to thank the anonymous reviewers for their feedback on this work. We also thank the AWS Console/UXP2 Science team at AWS AI Labs for discussions around this work. Lastly, we'd like to thank the human annotators for their efforts.

References

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Ultes Stefan, Ramadan Osman, and Milica Gašić. 2018. 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 (EMNLP)*.
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. [On the properties of neural machine translation: Encoder–decoder approaches](#). In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pages 103–111, Doha, Qatar. Association for Computational Linguistics.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2020. Making pre-trained language models better few-shot learners. *arXiv preprint arXiv:2012.15723*.
- Donghoon Ham, Jeong-Gwan Lee, Youngsoo Jang, and Kee-Eung Kim. 2020. End-to-end neural pipeline for goal-oriented dialogue systems using gpt-2. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 583–592.
- Wanwei He, Yinpei Dai, Yinhe Zheng, Yuchuan Wu, Zheng Cao, Dermot Liu, Peng Jiang, Min Yang, Fei Huang, Luo Si, Jian Sun, and Yongbin Li. 2022. [Galaxy: A generative pre-trained model for task-oriented dialog with semi-supervised learning and explicit policy injection](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(10):10749–10757.
- Ehsan Hosseini-Asl, Bryan McCann, Chien-Sheng Wu, Semih Yavuz, and Richard Socher. 2020. A simple language model for task-oriented dialogue. *Advances in Neural Information Processing Systems*, 33:20179–20191.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for NLP](#). *CoRR*, abs/1902.00751.
- Chia-Hsuan Lee, Hao Cheng, and Mari Ostendorf. 2021. [Dialogue state tracking with a language model using schema-driven prompting](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4937–4949, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*.
- Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*.
- Zhaojiang Lin, Andrea Madotto, and Pascale Fung. 2020a. [Exploring versatile generative language model via parameter-efficient transfer learning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 441–459, Online. Association for Computational Linguistics.
- Zhaojiang Lin, Andrea Madotto, Genta Indra Winata, and Pascale Fung. 2020b. Mintl: Minimalist transfer learning for task-oriented dialogue systems. *arXiv preprint arXiv:2009.12005*.
- Qi Liu, Lei Yu, Laura Rimell, and Phil Blunsom. 2021. Pretraining the noisy channel model for task-oriented dialogue. *Transactions of the Association for Computational Linguistics*, 9:657–674.
- Shikib Mehri, Tejas Srinivasan, and Maxine Eskenazi. 2019a. [Structured fusion networks for dialog](#). In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 165–177, Stockholm, Sweden. Association for Computational Linguistics.
- Shikib Mehri, Tejas Srinivasan, and Maxine Eskenazi. 2019b. Structured fusion networks for dialog. *arXiv preprint arXiv:1907.10016*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.

- Baolin Peng, Chunyuan Li, Jinchao Li, Shahin Shayan-deh, Lars Liden, and Jianfeng Gao. 2021. Soloist: Building task bots at scale with transfer learning and machine teaching. *Transactions of the Association for Computational Linguistics*, 9:807–824.
- Fabio Petroni, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. 2019. Language models as knowledge bases? *arXiv preprint arXiv:1909.01066*.
- Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training. <https://openai.com/blog/language-unsupervised/>.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. 2017. [Learning multiple visual domains with residual adapters](#). *CoRR*, abs/1705.08045.
- Timo Schick and Hinrich Schütze. 2020. Exploiting cloze questions for few shot text classification and natural language inference. *arXiv preprint arXiv:2001.07676*.
- Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. 2020. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. *arXiv preprint arXiv:2010.15980*.
- Yixuan Su, Lei Shu, Elman Mansimov, Arshit Gupta, Deng Cai, Yi-An Lai, and Yi Zhang. 2021. Multi-task pre-training for plug-and-play task-oriented dialogue system. *arXiv preprint arXiv:2109.14739*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2, NIPS’14*, page 3104–3112, Cambridge, MA, USA. MIT Press.
- Tsung-Hsien Wen, Milica Gasic, Nikola Mrksic, Pei-Hao Su, David Vandyke, and Steve Young. 2015. Semantically conditioned lstm-based natural language generation for spoken dialogue systems. *arXiv preprint arXiv:1508.01745*.
- Yunyi Yang, Yunhao Li, and Xiaojun Quan. 2021. Ubar: Towards fully end-to-end task-oriented dialog system with gpt-2. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 14230–14238.
- Xiaoxue Zang, Abhinav Rastogi, Srinivas Sunkara, Raghav Gupta, Jianguo Zhang, and Jindong Chen. 2020. Multiwoz 2.2: A dialogue dataset with additional annotation corrections and state tracking baselines. In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI, ACL 2020*, pages 109–117.
- Qi Zhu, Bing Li, Fei Mi, Xiaoyan Zhu, and Minlie Huang. 2022. Continual prompt tuning for dialog state tracking. *arXiv preprint arXiv:2203.06654*.

A Human Evaluation Task

We explored contextual dynamic prompting strategies for the response generation task using the MultiWOZ 2.2 (Budzianowski et al., 2018; Zang et al., 2020) dataset and noticed that the combined score that we obtained was significantly better than the baseline prefix-tuning method of response generation. To understand if the agents which incorporated contextual dynamic prompts did indeed provide a better conversational experience, we designed a small human evaluation task to test our hypothesis.

We picked a random subset of 10% of the conversations from the original MultiWOZ test data to perform this analysis. Once we obtained this random set, we ran our four model variants as described in Section 4 on the conversations to obtain system responses for each of them. We then presented the different agents’ responses to the annotator as shown in Table 4 below. In order to avoid potential biases, we shuffled the order of the agents between our annotators i.e., Agent 1 for annotator *a* would not be Agent 1 for annotator *b*. We kept track of which agents corresponded to which of our four methods prior to distribution of data amongst the annotators.

The annotators were given instructions to read every turn of conversation and provide a number between 1 and 4 for the agent which they thought performed the best for that turn. If the annotators found that there was a tie, they could pick more than one agent as [agent_a, agent_b]. In addition to this instruction, annotators were asked to read the entire conversation and pick the agent which performed the best - once again with an option to pick multiple. Table 5 below shows an example annotation style for a single conversation spanning 6 turns. There is an annotation at every turn and a single annotation at the end of the conversation.

We tallied results and re-mapped all agents back to their methods and found that agents 3 and 4 as mentioned in Section 5.2 were preferred at the conversation level in a total of 53 of the 100 conversations while agents 1 and 2 were only preferred 10 conversations in the entire set of 100.

B Examples

In the table below, we provide another examples from the human annotation task where annotators picked agent 4 as the best performing agent for the conversation (we performed the re-mapping to fit the agent numbers in Section 5.2). Table 2 shown in the main paper indicates a scenario where Agent 4 was rightly able to ask the customer about their area of preference for dining, and also considers a booking time change from the user. Table 6 shows an example where agent 4 is able to take into consideration all the different constraints mentioned by the user in terms of slots in the response that it generates. Table 3 shows an example where the agent 4 is able to correctly respond to the user request for number of people to reserve the table for - which none of the other agents captured. All these examples showed that agent 4, which uses contextual dynamic prompts with dialog states was indeed able to provide smoother responses to the customer input compared to the other agents.

C Compute Infrastructure and Experimental settings used

All our experiments were carried out on public AWS p3 GPU instances which are backed by Tesla V100 GPU cores. All four variants of our models required 18-24 hours to reach convergence on these machines. Learning rates were varied based on performance of models on the validation split of our dataset.

Turn num	User turn	Agent 1 response	Agent 2 response	Agent 3 response	Agent 4 response
1					
2					
3					
4					
5					
6					
7					

Table 4: The format which is presented to annotators while performing turn-level and conversation-level annotation. The agents are shuffled between the annotators to avoid biasing them.

Turn num	Turn level	Conversation level
1	2	
2	[3,4]	
3	2	
4	3	
5	4	
6	[3,4]	3

Table 5: We asked annotators to provide two levels of annotation for each conversation - turn level and conversation level. Annotators were allowed to pick more than agent in case of a tie.

Turn num	User turn	Agent 1 response	Agent 2 response	Agent 3 response	Agent 4 response
1	i am looking for a train that departs from peterborough and goes to cambridge.	there are [choice] trains leaving from [departure] to [destination]. what day would you like to travel?	what day would you like to travel?	there are [choice] trains leaving [departure] from [departure] to [destination]. what day would you like to travel?	i have [choice] trains departing from [departure] to [destination]. what day would you like to travel?
2	we will be leaving on sunday afternoon.	[trainid] leaves [departure] at [leaveat] and arrives in [destination] at [arriveby]. would you like me to book that for you?	i have [choice] trains available. what time would you like to leave?	[trainid] leaves at [leaveat] and arrives in [destination] at [arriveby]. would you like me to book it for you?	[trainid] leaves [departure] at [leaveat] and arrives at [destination] at [arriveby]. would that work for you?
3	yes, it should leave after 20:15 and leave on monday, not sunday. please give me a price, travel time, and arrival time for any train.	[trainid] leaves at [leaveat] and arrives in [destination] at [arriveby]. the price is [price].	[trainid] leaves at [leaveat] and arrives at [arriveby]. would you like me to book that for you?	[trainid] leaves at [leaveat] and arrives at [arriveby]. the price is [price]. would you like me to book it for you?	[trainid] leaves [departure] at [leaveat] and arrives in [destination] at [arriveby]. the cost is [price] and the travel time is [duration].
4	yes please. what is the travel time? i'll also need a hotel in the expensive price range. can you help me with that as well?	the travel time is [duration]. is there anything else i can help you with?	the travel time is [duration].	the travel time is [duration]. is there anything else i can help you with?	the travel time is [duration]. is there anything else i can help you with?
5	yes, i'd like a 2 star rating and i'll need it to be in the east side of town.	i have [choice] [type] in the [area] with [stars] stars. do you have a preference?	there are [choice] [type] in the [area] that are [stars] stars. do you have a preference?	there are [choice] [type] in the [area] that are [pricerange] and [pricerange]. do you have a preference?	[name] is a [stars] star hotel in the [area]. would you like me to book it for you?
6	yes, i would like to book it for 3 people starting that same day for 4 nights.	booking was successful. reference number is : [ref]. is there anything else i can help you with?	booking was successful. reference number is : [ref].	booking was successful, the total fee is [price] payable at the station. reference number is : [ref].	booking was successful. the total fee is [price] payable at the station. reference number is : [ref].
7	that is all thank you goodbye.	you're welcome. have a great day!	you're welcome. have a great day.	you're welcome. have a great day.	you're welcome. have a great day.

Table 6: Another example where Agent 4 was picked as the best performing agent in a conversation.