

Appendix - A Zero-Shot Approach for Multi-User Task-Oriented Dialog Generation

Shiv Surya

Yohan Jo

Arijit Biswas

Alexandros Potamianos

Amazon

{shisurya, jyoha, barijit, potamian}@amazon.com

A Appendix-Training and inference settings.

A.1 Appendix-Training and inference settings of dialog generators.

Our proposed dialog generators are initialized from GPT-2-XL (Radford et al., 2019) that uses a 48-layer decoder-only transformer and 1600-dimensional hidden states. We set the maximum sequence length to 512 tokens as it is large enough to contain most utterances and prompt inputs. To optimize the models, we use an Adam solver (Kingma and Ba, 2015) with an initial learning rate of $10e-5$ and minibatch of size 8. We train the models with teacher forcing until convergence. For the social generator, we evaluate perplexity on the validation set and use it as stopping criterion. For task generation, we evaluate semantic similarity on the validation set and use it as stopping criterion. The dialog context is restricted to 5 previous utterances for both generators during training.

Both dialog generators use special tokens in the prompt to demarcate dialog history, belief state and social act. This aids the generator to attend to the structured prompt better. At inference time, we generate dialogs with using top-p sampling scheme with $p=0.85$ (Holtzman et al., 2019), ngram repetition penalty with $n=4$ applied on generated dialog turns, a softmax temperature of 0.7 and 0.9 for the social and task generators, an utterance length constraint of 100 tokens. Our models are implemented in Pytorch (Paszke et al., 2017), use HuggingFace Transformers library (Wolf et al., 2020) and are trained on NVIDIA V100 GPU. Each dialog generator takes 200 GPU hours of compute for finetuning on the dialog dataset.

A.2 Appendix-Inference settings of fusion layer.

For fusing the generators, we choose a β randomly from a range of $\beta = [0.7, 1.0]$ for task oriented

turns and $[0.3, 0.7]$ for the social turns. This threshold of 0.7 was chosen to make sure that the task oriented turn does have the information-belief states from the single user utterance being rewritten and tuned by optimising for SS(semantic similarity) metric on the validation set. Figure 1 shows examples of how the generated dialog utterance varies as we vary β by biasing the fusion from the social dialog generator to the task-oriented dialog generator. In examples D.1 and D.2, we see that the utterance are off topic and not related to user belief when $\beta = 0.1$. For higher value of beta (biased towards task oriented dialog generator), the response is terse but reflects the belief state. Intermediate values of beta blends the generators and generates more natural responses.

A.3 Appendix-Training and inference settings of turn planner.

Turn planners in our approach are initialized from BERT (336M parameters) model (Devlin et al., 2019) that uses a 24-layer encoder-only transformer, 1024-dimensional hidden state and 16 attention heads. We set the maximum sequence length to 256 tokens, as it is large enough to contain dialog history. We use Lamb optimizer (You et al., 2019) with an initial learning rate of $5e-5$, 500 warmup steps and minibatch of size 8. For both the coarse and fine-grained turn planner, we use class-weighted accuracy on the validation set as stopping criterion. At inference time, we use a threshold that gives greater than 70% precision on predicting the task turns. Our models are implemented in Pytorch (Paszke et al., 2017), use HuggingFace Transformers library (Wolf et al., 2020) and are trained on NVIDIA V100 GPU. Each model based planner takes 15 GPU hours of compute for finetuning on a dialog dataset.

A.4 Appendix-Sample inference walkthrough

The architecture of our proposed multi-user task-oriented dialog generator is shown in Figure 2 of the main paper. We consider a single user utterance with some dialog history and belief state as input. To rewrite this user utterance, we start with a social turn and a random choice of social act in the first step. The social act and dialog history are given as input prompt to the social generator. The task generator is prompted with the belief state and dialog history. The next utterance is decoded by late fusion of both the dialog generators biased towards the social dialog generator. The generated turn is given as input to the turn planner which predicts the turn type and social act for the second turn in dialog. Depending on the turn type, the fusion is biased towards the social or task-oriented generator. This continues until the planner predicts task-oriented turn which is decoded by late fusion of both the dialog generators with a bias towards the task-oriented dialog generator.

B Human Evaluation

A snapshot of the detailed instructions describing the task and the human evaluation template with annotations are shown in Figure 3–4. The pool of workers for our human evaluations were located in the US, and had a 90% approval rate for at least 200 previous annotations. Based on approximate completion time for our task, we estimate a pay rate of \$12 per hour.

C Additional Qualitative Results

Additional qualitative examples are shown in Figure 2. Example D.1 shows a multi-user dialog where users chit-chat about visiting a park in the same location reflected in user belief and describing reasons they would like to visit it. One of the users then makes a request for a college in this location reflecting the belief state of the single-user utterance. Similarly D.6 captures an interesting social deliberation and the request to book a taxi. However, the destination referenced in the user belief is missed in this case. In example D.3, the user refers to a park entry fee of a dollar as expensive showing some of the limitations of language models in expressing commonsense. This limitation is also shown in D.5 where a generated utterance captures the intent to take a taxi and remarking how close it is walking and the good weather in the same utterance. Example D.6 shows a disconnect

in the topic between turns where a user talks about different location and asks for the admission fee despite not deciding on a location. This is also seen in example D.2.

References

- 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.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. The curious case of neural text degeneration. In *International Conference on Learning Representations*.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch. In *NIPS-W*.
- 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.
- 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. 2020. [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. Association for Computational Linguistics.
- Yang You, Jing Li, Sashank Reddi, Jonathan Hseu, Sanjiv Kumar, Srinadh Bhojanapalli, Xiaodan Song, James Demmel, Kurt Keutzer, and Cho-Jui Hsieh. 2019. Large batch optimization for deep learning: Training bert in 76 minutes. In *International Conference on Learning Representations*.

<p>D.1</p> <p>Previous Utterance from Agent:</p> <p>A:I have successfully booked your room at the Limehouse. Your reference number is 4VB8JH2B. As for the restaurant, do you have any preferences as far as price range goes?</p> <p>USER BELIEF STATE:</p> <p>restaurant area: north, restaurant food: India, restaurant price range: don't care</p> <p>NEXT UTTERANCE:</p> <p>U:No preferences for the price range, but I'd like it to be in the same area as the hotel</p> <p>GENERATED MULTI-USER DIALOG:</p> <p>Beta=0.1(biased towards social generator):I would like to stay in the downtown area. My name is John. Beta=0.5: I don't have a preference. I just want to eat Indian food. Beta=0.9(biased towards task oriented dialog generator): I don't have a preference.</p>	<p>D.2</p> <p>Previous Utterance from Agent:</p> <p>A: It is not. It serves British food. There are 2 modern European restaurants in the centre. Would you prefer one of those?</p> <p>USER BELIEF STATE:</p> <p>restaurant area: centre , restaurant book day: Tuesday, restaurant book people: 5</p> <p>NEXT UTTERANCE:</p> <p>U: Yes please pick a moderate one and book me a table for 5 people on Tuesday at 14:45.</p> <p>GENERATED MULTI-USER DIALOG:</p> <p>Beta=0.1(biased towards social generator): Yes, I'd like to have a British food. Can you recommend me one? Beta=0.5:I'd like to try the one in the centre. Can you book a table for 5 people at 14:45 on Tuesday? Beta=0.9(biased towards task oriented dialog generator): I would like to book a table for 5 people at 14:45 on Tuesday.</p>	<p>D.3</p> <p>Previous Utterance from Agent:</p> <p>A:What kind of food would you like today?</p> <p>USER BELIEF STATE:</p> <p>hotel stars: 4 ,hotel type: guesthouse</p> <p>NEXT UTTERANCE:</p> <p>U:I am not interested in food, thanks, but I'd like help finding a 4 star guesthouse.</p> <p>GENERATED MULTI-USER DIALOG:</p> <p>Beta=0.1(biased towards social generator):I'd like to have a steak, please. And I'd like to have a cup of coffee, please. Thank you. Beta=0.5:I'm not picky. I just want a place to stay. Beta=0.9(biased towards task oriented dialog generator):I don't care.</p>
--	---	--

Figure 1: Examples of utterances generated by fusion of our dialog generators by biasing the fusion from the social dialog generator to the task-oriented dialog generator.

<p>D.1</p> <p>DIALOG HISTORY:</p> <p>U: Are there any entertainment spots in the town centre? A: Nothing specifically marked entertainment, although we do offer quite a lot to do, from theatres to parks.</p> <p>USER BELIEF STATE:</p> <p>attraction area: centre , attraction type: college</p> <p>UTTERANCE BEING REWRITTEN:</p> <p>Can you recommend a college in town centre?</p> <p>GENERATED MULTI-USER DIALOG:</p> <p>U0: Well, I do like the idea of a park in the centre. It's something new and exciting. U1: I'm sure you do. The problem is that the city centre is quite full. It's hard to find a place to go to in a city centre that isn't full of office buildings and shopping malls. U0: I see. Perhaps a nice park would be good. There are many nice looking trees and plants around the park. It would be a nice change from the city centre. U1: I'm also looking for a college to visit in the centre of town.</p>	<p>D.2</p> <p>DIALOG HISTORY:</p> <p>U:I'm visiting Cambridge and looking to do the tourist thing! I'm interested in information about potential nightclubs in the centre of town! A:Sure, there are 5 different nightclubs in the centre area, were you wanting to check them all out or would you like a recommendation?</p> <p>USER BELIEF STATE:</p> <p>attraction area: centre, attraction type: nightclub ,attraction postcode: requested, attraction phone: requested</p> <p>UTTERANCE BEING REWRITTEN:</p> <p>Great can I get the phone number and postcode of one you'd recommend?</p> <p>GENERATED MULTI-USER DIALOG:</p> <p>U0: I want a recommendation for a nightclub that I think is good for me, but I don't want to spend a lot of time in one place. U1: There's a really good nightclub in the city centre, I think you'll enjoy it if you go there.It has a very good reputation for its live jazz and the people that work there are friendly and helpful.If you want to book a table for 5 at the club, I would be happy to help you out.If you just want to have a drink, I would also recommend that you book a table at the restaurant or in the park. U0: I think I'll go to the restaurant then.If you can get me a table for 7 at that time, that would be great.I would also like to book a taxi to take me back to the airport after I've enjoyed my stay in the city.My friend told me that it is very easy to get a taxi. U1:Well, I think I will go to the nightclub. Can you get me the phone number and postcode for it?</p>	<p>D.3</p> <p>DIALOG HISTORY:</p> <p>NA</p> <p>USER BELIEF STATE:</p> <p>attraction area: east, attraction type: park</p> <p>UTTERANCE BEING REWRITTEN:</p> <p>Can you help me find a park on the east side please?</p> <p>GENERATED MULTI-USER DIALOG:</p> <p>U0: I'm looking for a nice place to visit. U1: How about a park? U0: Park is good. It's very quiet and beautiful. U1: How much is the admission fee? U0: Usually it's free. But sometimes it's a little expensive. Do you know what the fee is for a park? U1: Usually free. But some people charge a dollar for it. It's a little expensive. U0: How about a park in the city center? U1: Sure, that would be fine.</p>	
<p>D.4</p> <p>DIALOG HISTORY:</p> <p>U: I have two places. Nusha and Tipin. Which would you prefer? A: I would like Nusha. Can you give me the phone number and address for Nusha? U: The address for Nusha is unit g6, Cambridge leisure park, cliffon road. Phone number is 01223902158. A: I also need an Irish restaurant that is expensive. U: What area would you like?</p> <p>USER BELIEF STATE:</p> <p>restaurant food:irish, restaurant price range:expensive</p> <p>UTTERANCE BEING REWRITTEN:</p> <p>The area doesn't matter.</p> <p>GENERATED MULTI-USER DIALOG:</p> <p>U0: South. East and West. It's not very particular.It'll be a big place. U1: It'd be very expensive. Can you recommend something? U1: Irish food is good. It's very good. If you want to go to a restaurant that is very expensive, I would recommend the restaurant. It'll be a big restaurant. It has a lot of people.It'd be very expensive. U0: I would like to book a table for three at 16:45 on Sunday.</p>	<p>D.5</p> <p>DIALOG HISTORY:</p> <p>U:I need help with attractions! Can you help me out? A:There are so many attractions to choose from! Are you interested in any particular area?</p> <p>USER BELIEF STATE:</p> <p>attraction area: centre, attraction: type: entertainment</p> <p>UTTERANCE BEING REWRITTEN:</p> <p>Yeah, I'm looking for an entertainment place in the centre of town.</p> <p>GENERATED MULTI-USER DIALOG:</p> <p>U0: I want to go to the Epcot Center. U1: That sounds great.How much is it to get there? U0: It's only about \$ 5.00 per person. U1: That's cheap for a lot of people. U0: Yeah, I think so. It's a good idea to do some shopping in the center.It'll be fun! U1: I'll take a taxi to get there. It's only about 20 minutes walking and it's a nice evening. U0: Sounds good.Let's go! U1: Wait, I'm looking for an entertainment place in the centre of town.</p>	<p>D.6</p> <p>DIALOG HISTORY:</p> <p>NA</p> <p>USER BELIEF STATE:</p> <p>taxi departure maharajah tandoori restaurant</p> <p>UTTERANCE BEING REWRITTEN:</p> <p>I'd like to book a taxi to depart from Maharajah Tandoori Restaurant.</p> <p>GENERATED MULTI-USER DIALOG:</p> <p>U0:I need to get a taxi. U1: It's unnecessary.You'll be fine. U0: But I'm not a child. I have a business to attend to.Why should I waste time looking for a taxi when I can get it from the airport? U1: Because you're a grown-up.That's your right. U0:I just can't help it.</p>	<p>D.7</p> <p>DIALOG HISTORY:</p> <p>U: I'm trying to find a concert hall in the east part of town. A: Unfortunately, I don't have any listings for concert halls on the east side. Can I check another area of town for you? USER BELIEF STATE:</p> <p>restaurant food:irish, restaurant price range:expensive</p> <p>UTTERANCE BEING REWRITTEN:</p> <p>How about an entertainment attraction instead?</p> <p>GENERATED MULTI-USER DIALOG:</p> <p>U0: How about a place in the west part of town? U1: I'm sorry, but I'm not sure.Do you have any suggestions for a place to go to for entertainment in the east? U0: Well, I'd like to go to a place in town that is not too far from the airport.How about a place near the airport? U1: Ok, I'll try that.How about an entertainment venue near the airport? U0: That would be great.How much is the admission fee? U1: I'm not sure. Do you have any information on entertainment venues?.</p>

Figure 2: Additional examples of multi-user task-oriented dialogs generated by our approach is shown here.

Task Example

Consider a dialog sequence between a user and an agent discussing a place of dining. In the third utterance, the user **U** specifies that they want a cheap restaurant. The belief of the user is reflected in the **User Belief State** shown below.

Dialog History:

U: I am looking for a place to dine. The restaurant should serve mediterranean food and should be in the south

A: I'm afraid there aren't any. Would you like to try another area or cuisine?

U: What restaurants are cheap?

User Belief State: restaurant area: south, restaurant price range: cheap

Utterance being rewritten: What restaurants are cheap?

In our task, the machine is asked to rewrite the last utterance as a dialog between two users. Below we show an example of the single utterance "What restaurants are cheap?" being **rewritten as a conversation between two users**.

Rewritten dialog sequence between two users U1 and U2

U1: Another area? I don't want to go to a place that isn't in the south. It's important to me that the restaurant be in the center of the city. If you don't have a place in the south, I'll go to a restaurant in the center. It'll be a new experience for me.

U2: Maybe you are right. Let me see what I can find. There are some places in the center that I would like to try. If you can't find anything in the center, then try the south. There are many good restaurants in the south.

U1: Yes, I would like a restaurant in The South. Cheap is fine.

In this page, you will be shown similar dialogs generated by two machines. You need to pick the better dialog based on a criteria defined below.

You need to carefully read and compare such dialog rewrites from two machines and make a judgment about the following criteria:

- Content Quality:** In the process of evaluating content quality, you should decide which of the two dialogs is better, based on **fluency**, **logical coherence** and proper **ordering**.
- Belief State Faithfulness:** In the process of evaluating Belief State Faithfulness, you should decide which of the two dialogs better reflect the provided user beliefs for the dialog turn being rewritten. **Please evaluate this criterion only based on faithfulness to user belief and do not consider fluency or coherence for this part.**

For **Belief State faithfulness**, you might want to look at the following details in the dialog:

- Does the dialog **mention** the user belief states provided?
- Is the **intent** expressed in the generated dialog matching the single turn being rewritten?
- If both dialogs are reflecting user beliefs, which dialog has **greater** belief state coverage.

For **Content Quality**, you might want to look at the following details in the dialog:

- Does the dialog show repetitive content or has abrupt changes in the topic of dialog?
- Is the flow of dialog **logically coherent** across turns?
- If both dialogs have minor logical inconsistencies, which dialog has **lesser** inconsistencies.

Figure 3: Snapshot of detailed instructions for our human evaluation template.

Dialog History: No History

User Belief State: attraction area: centre, attraction type: entertainment

Utterance being rewritten: Are there any entertainment spots in the town centre?

Dialog 0

(0):U1:I'm so bored all day.
(1):U2: Why don't you do something fun?
(2):U1: I don't know what to do. It's so boring
(3):U2: Well, maybe you should find a job and do something useful.
(4):U1: That's a good idea.
(5):U2: How about we try to find something entertaining in the centre of town.

Dialog 1

(0):U1:I'm looking for a good novel. Do you have anything on the shelf?
(1):U2: Sure. We have a variety of novels on the shelf. Some are detective mystery stories, some are science fiction, some are biography, some are fiction, some of them are biography, and some are just plain silly. We're also got a lot of fun stuff on the shelf.
(2):U1: What kind of stuff is that?
(3):U2: Some of the stuff is just for fun. We have everything from science fiction to biography. We're also got a few detective mystery stories and some science fiction. We'll have a lot more later. What do you want to do?
(4):U1: I'd like to go to an entertainment place in the centre of town.

Given the two dialogs, the answer to the questions are as follows (we additionally provide reasons for the answers):

Q1- Content Quality: Which of the two dialogs is more fluent, coherent and has proper ordering?

Dialog 0 Dialog 1

Reason: In the 2nd dialog "U1 and U2 are engaged in a casual conversation about a novel and then switch to the user intent to find an attraction" In the first dialog, "both U1 and U2 are engaged in deliberating the intent to find an attraction." The first dialog is more logically coherent and concise. Hence, the **first dialog** is better in content quality.

Q2- Belief State Faithfulness: Which of the two dialogs are better in terms of reflecting the provided user belief states?

Dialog 0 Dialog 1

Reason: Both dialogs reflect the intent of the user to request an attraction in the center of town but Dialog 0 has intent matching across utterances.

Figure 4: Snapshot of an example task for our human evaluation template with annotations