

Simulated Chats for Building Dialog Systems: Learning to Generate Conversations from Instructions

Biswesh Mohapatra^{1,2,*}, Gaurav Pandey¹, Danish Contractor¹, Sachindra Joshi¹

¹ IBM Research AI, New Delhi

² International Institute of Information Technology, Bangalore

biswesh.mohapatra@iiitb.org

{gpandey1, dcontrac, jsachind}@in.ibm.com

Abstract

Popular dialog data sets such as MultiWOZ (Budzianowski et al., 2018) are created by providing crowd workers an *instruction*, expressed in natural language, that describes the task to be accomplished. Crowd workers play the role of a *user* and an *agent* to generate dialogs to accomplish tasks involving booking restaurant tables, calling a taxi etc. In this paper, we present a data creation strategy that uses the pre-trained language model, GPT2 (Radford et al., 2018), to *simulate* the interaction between crowd workers by creating a *user* bot and an *agent* bot. We train the simulators using a smaller percentage of actual crowd-generated conversations and their corresponding instructions. We demonstrate that by using the simulated data, we achieve significant improvements in low-resource settings on two publicly available datasets - MultiWOZ dataset (Budzianowski et al., 2018) and the Persona chat dataset (Zhang et al., 2018a).

1 Introduction

Building dialog systems typically requires a large collection of conversation logs for training models. There exists a scarcity of the data required to train a dialog system for most tasks. Various methods have been proposed to tackle this issue including paraphrase techniques to generate artificial training data (Kumar et al., 2021; Du and Black, 2018), generating annotations including intent-slots and dialog acts (Yoo et al., 2019, 2020a,b) or even injecting noise to improve robustness in dialog act prediction for ASR data (Wang et al., 2020).

Crowd-sourcing is a popular method for generating new large datasets. For instance, to create datasets for task oriented dialogs, crowd workers may be provided with *instructions* that describes the task; workers then play the roles of a *user* and an *agent* to generate conversations (Budzianowski

* Work done during internship at IBM Research AI

Instruction : You are looking for a **train**. The train should **arrive by 11:45** and should leave on **Saturday**. The train should go to **Cambridge** and should depart from **Ely**. Once you find the train you want to make a booking for **8 people**. Make sure you get the **reference number**.

CHAT GENERATED BY CROWD-WORKERS:

User: I need to find a train for **Saturday**.

Agent: I have many trains that depart **saturday**. Where will you be departing from and where is your destination?

User: I will be departing from **Ely** and traveling to **Cambridge**.

Agent: Okay, and what time do you want to leave after or arrive by?

User: I need to arrive by **11:45**.

Agent: There are 3 trains here. Do you want me to book any tickets?

User: Yes please. I need it to be booked for **8 people**.

Agent: Booking was successful, the total fee is 28.16 GBP payable at the station. Reference number is : ZVERHBT3.

User: Thank you, that is all I need.

Agent: You're welcome, thank you for calling the Cambridge TownInfo centre. Have a great day.

CHAT GENERATED BY SIMULATORS:

User: I am looking for a train to **Cambridge**.

Agent: There are many trains available. Where will you be departing from?

User: I will be departing from **Ely**.

Agent: How about TR7994 that leaves at **05:35**?

User: That would work, can you book it for **8 people**?

Agent: I have booked **8** tickets on that train, the total fee is 35.2 gbp payable at the station. Reference number is ZXQSGIN7.

User: thank you. that is all i need.

Agent: You are quite welcome. I am glad I could help. Enjoy your trip !

Figure 1: An instruction from the MultiWOZ dialog corpus (Budzianowski et al., 2018) and the corresponding conversation generated by real crowd-sourced workers and crowd-worker simulators.

et al., 2018). The *user* worker begins the conversation by stating their requirement and the *agent* worker provides information to the user by querying a knowledge base (KB), if required. Together, both workers interact with each other via natural language to generate conversations. Similarly in Persona chat (Zhang et al., 2018a), the workers are provided different *personalities* to role play conversations. However, creating large crowd-sourced datasets can be time consuming and expensive.

Pre-trained transformer-based language models such as GPT-2 (Radford et al., 2018), that are trained on a large number of documents crawled from the web have achieved extensive generalization in natural language understanding and generation across a variety of diverse tasks (Budzianowski et al., 2018; Rajpurkar et al., 2016; Welleck et al.,

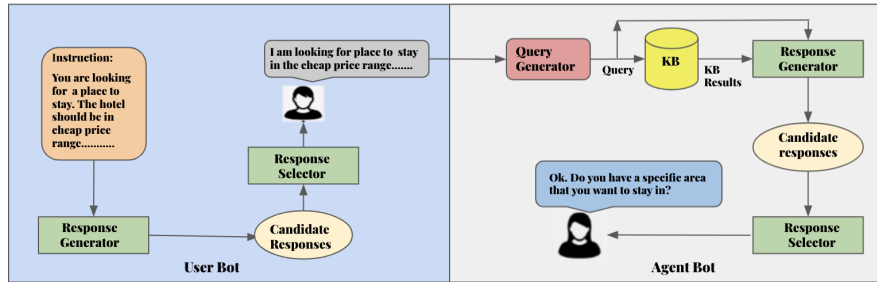


Figure 2: Generation of a conversation conditioned on the instructions and knowledge base (KB). Although not displayed in this diagram, each module (in green) also receives the dialog history as input.

2019). Recent works have exploited the prior knowledge in these models to train effective models for machine translation (Araabi and Monz, 2020), language understanding in low resource settings (Dou et al., 2019) and few-shot language models (Brown et al., 2020).

In this paper we demonstrate how such large pre-trained models can also be used to follow *instructions* and generate conversations. We create a *user* simulator and an *agent* simulator. The *user* simulator has access to the *instructions* while the *agent* simulator has access to a knowledge base (KB). The *agent* simulator maps the current dialog context to a *belief state* (query), that can be executed over a knowledge base (KB), to retrieve a set of results if required. Thus, the simulators are trained to interact with each other to generate conversations conditioned on the instructions and the KB. In our work we train these simulators using just 5-20% of crowd-sourced conversations by fine-tuning the pre-trained language models — GPT2 (Radford et al., 2018) and Longformer (Beltagy et al., 2020). We use the external knowledge present in these language models to help generate effective artificial data on low-resourced datasets. An example of a generated conversation is shown in Figure 1.

Our experiments further show that from a small number of existing conversations we are able to train meaningful *user* and *agent* bots that in-turn generate new conversations. This in principle, is somewhat similar to a noisy student-teacher model (Xie et al., 2020) where a weaker teacher model is used to generate labels(dialogs in our case) which is then used to train a new student model that significantly outperforms the teacher model in end task. Due to its simplicity and generality, our model could be used on a wide variety of dialog systems by taking different forms of instructions.

Contributions: (1) We present a novel technique that effectively uses weak generative models to create new artificial data which are used to train final end task models (2) We introduce a simple yet effective dialog-generation framework¹ that mimics the roles played by crowd workers to generate complete conversations. (3) We demonstrate the generality of our model by generating data for two different types of dialog tasks - task oriented conversations and persona-guided conversations. We show that pre-trained language models can be successfully used for generating artificial data in low resource dialog settings leading to a 7-13% improvement in combined score in MultiWOZ 2.0 and 2-10% improvement in Hits@1 metric in Persona Chat. (4) We present a human-study to assess the quality of our simulated dialogs. We find that the generated conversations are grammatically sound and meaningfully move the conversations forward.

2 Related Work

The method of interacting different models to converse with each other has seen some recent successes (Shah et al., 2018; Papangelis et al., 2019). (Hou et al., 2019) has previously used simulators to generate conversational artificial data. However the work uses hand-crafted templates for generating dialogs. Our approach is more general and less cumbersome as demonstrated through the effectiveness of our approach on two different tasks. (Lin et al., 2020a) tries to use a similar simulator approach but uses human in the loop in order to bring variations to the dialogs. On the other hand, our approach doesn't require any human involvement apart from providing diverse instructions which are easy to produce in large quantities.

Unlike existing data augmentation methods, like those based on paraphrase generation (Malandrakis

¹<http://ibm.biz/simulatedchats>

BOT	MODULE	INPUT	OUTPUT
User	Response Generator	[CLS]You are looking for a....[GOAL][St@rt][User] I need a train to Cambridge.[Agent] Sure, which day would you like to travel?[SEP]	[User] I would like to travel this [DAY] at [TIME].
	Response Selector	[CLS][St@rt][User] I need a....[Agent] Sure, which day would you like to travel?[CAN][User] I would like to travel this [DAY] at [TIME].[SEP]	0.92
Agent	Query Generator	[CLS][St@rt][User] I need a....[Agent] Sure, which day would you like to travel?[User] I would like to travel this Sunday at 9PM.[SEP]	[Q] Train Destination=Cambridge Day=Sunday Time=9PM [Q]
	Response Generator	[CLS][Q] Train Destination=Cambridge Day=Sunday Time=9PM [KB] Total = 2 [St@rt][User] I need a....[Agent] Sure,.....travel?[User] I would..... at 9PM.[SEP]	[Agent] I can see [value_count] trains available for [train_destination]. How many tickets do you need?
	Response Selector	[CLS][St@rt][User] I need a....[Agent] Sure....travel? [User] I would..... at 9PM.[CAN][Agent] I can....for [train_destination]. How many tickets do you need?[SEP]	0.83

Figure 3: The input and output formats for User and Agent Bot. [GOAL] indicates end of instruction, [CAN] indicates candidate for selector, [St@rt] indicates start of conversation, [Q] and [KB] indicate Queries and KBs.

et al., 2019; Gao et al., 2020; Anaby-Tavor et al., 2020), simulators create completely new conversations which create more diverse examples, helping train better end-task models (Section 4).

Teacher model has been used to train student models, based on the idea of knowledge distillation (Hinton et al., 2015), teach dynamic loss functions (Wu et al., 2018) or for adaptation under meta-learning settings (Qian et al., 2021). To the best of our knowledge, we are the first to adopt such a model for data generation in dialog systems.

3 Model

In order to generate the data, we train our GPT-2 based agent and user bots(teacher models) using a subset of original data to simulate low-resource environment(5% or 20%). In case of task oriented dialogs, we require a third model for generating belief states as well. Longformer based selector models are trained to chose from a list of responses generated by the teacher generator models as seen in Figure 2. Note that all the modules in figure (shown in green) also receive dialog history as input which has not been shown in the figure for ease of presentation. Finally the generated data from the process is mixed with the original low resourced data to create a new dataset. To test the effectiveness of new dataset, we compare the performance of newly trained student model on the new data to that of the teacher model and other baseline models on the respective end tasks of the datasets.

3.1 Overview

We assume that the dialog comprises of a sequence of utterances between a user and an agent i.e. $\mathcal{D} = (u_1, a_1, \dots, u_n, a_n)$ where u_i is a user utterance while a_i is an agent utterance. A turn is a pair of user and agent utterance. At any given turn m , the sequence of utterances prior to the turn, that

is, $c_m = (u_1, a_1, \dots, u_{m-1}, a_{m-1})$ is referred to as dialog context or dialog history. Apart from the dialog \mathcal{D} , we have access to a set of *instructions* \mathcal{I} and a knowledge base \mathcal{KB} . The aim is to learn a model that can generate the dialog \mathcal{D} conditioned on the instructions \mathcal{I} and the knowledge base \mathcal{KB} . That is, we wish to model $p(\mathcal{D}|\mathcal{I}, \mathcal{KB})$.

The dialog generation framework mimics the human-to-human data collection approach used in MultiWOZ (Budzianowski et al., 2018). The dialog is generated in a sequence of turns. The user bot has access to instructions \mathcal{I} while the agent bot can query the knowledge base \mathcal{KB} . Thus, the joint distribution of the dialog decomposes as follows:

$$p(\mathcal{D}|\mathcal{I}, \mathcal{KB}) = \prod_{i=1}^n p(u_i|c_i, \mathcal{I})p(a_i|c_i, u_i, \mathcal{KB}). \quad (1)$$

The dialog history for the first turn, c_1 , is an empty set. The first factor in the product on the left corresponds to user bot which conditions on the instructions, as well as, the dialog history to output the user utterance. The second factor models the distribution of the agent bot over the responses, conditioned on the dialog history and knowledge base. A pictorial representation of the interaction between the two bots is shown in Figure 2. We discuss the various modules of both the bots in further detail below. The input and output formats for the various networks of these modules are shown in Figure 3.

3.2 User Bot

The user bot generates utterances conditioned on the dialog history and the instructions, that is, it models $p(u_i|c_i, \mathcal{I})$. For the sake of readability, we will remove the turn index i from the distribution. As shown in Figure 2, this distribution is modeled in two steps. Firstly, the dialog history and the

instructions are fed to a *response generator* module which outputs a pool of candidate responses $P^u = (\bar{u}_1, \dots, \bar{u}_r)$. A *response selector* module then assigns a score s_k^u to each response \bar{u}_k in the pool. Based on these scores, we define the distribution $p(u|c, \mathcal{I})$ as follows:

$$p(u|c, \mathcal{I}) = \begin{cases} \frac{\exp(s_k^u)}{\sum_{j=1}^r \exp(s_j^u)}, & \text{if } u = \bar{u}_k \in P^u \\ 0, & \text{if } u \notin P^u \end{cases} \quad (2)$$

The candidate response with the highest probability is selected as the next user utterance and sent to the agent bot. Next, we discuss the various modules in the user bot and how they are trained.

3.2.1 Response Generator

The aim of response generator module is to output a pool of candidate user utterances for the given dialog history and the instructions. To achieve this, an autoregressive distribution over the tokens of the utterance u is defined. We finetune the pretrained GPT-2 network to model this distribution. Specifically, given the tokens in the instructions and the dialog history, the GPT-2 network is trained to output the tokens of the user utterance. The utterance generated are in delexicalised format which are lexicalised from the values present in instruction before being shown to the agent bot (once selected by the selector) as shown in Figure 3.

While it is possible to sample an utterance from the GPT-2 network via greedy sampling or beam search, this poses several issues. Firstly, autoregressive distributions tend to assign high probability to short utterances. Secondly, commonly occurring utterances in the corpus tend to have higher probability than the informative responses that are less frequent. We noticed that in lower data settings, the greedy response may not always be a relevant response. Nucleus sampling generates diverse responses which helps the response selector to pick more informative responses w.r.t the given context.

Hence, once the network has been trained, we sample multiple user responses from the network via nucleus sampling (Holtzman et al., 2019) to obtain a pool of candidate responses $P^u = (\bar{u}_1, \dots, \bar{u}_r)$. This pool of candidates is fed to the response selector module as shown in Figure 2.

3.2.2 Response Selector

The aim of the response selector module is to assign a score to each candidate response in the pool based on its relevance to the context. We achieve this by

feeding the tokens of the context and the candidate response (concatenating them with [CAN] token) to a Longformer network architecture (Beltagy et al., 2020). The network outputs a contextualized embedding for each token. We feed the embedding of the [CLS] token through a linear layer followed by a sigmoid unit. The output of the network corresponds to the score assigned to the response for the given context.

The network is trained to assign high scores to the positive (or ground-truth) responses while assigning low score to the negatively sampled responses. For each gold context-response pair, we provide a total of 10 negative response samples. These samples contain 5 random responses, 2 responses which are already part of the context (to stop the response selector from picking such responses) and 3 responses formed by concatenating 2 random responses to discourage the selector from picking longer candidate responses.

The network is trained via the triplet loss (Chechik et al., 2010; Hoffer and Ailon, 2015). Specifically, given the context c , the ground-truth response u_p and a negatively sampled response u_n , the triplet loss is defined as follows:

$$L(c, u_p, u_n) = \max(0, s(c, u_n) - s(c, u_p) + \alpha), \quad (3)$$

where $s(c, u)$ is the score assigned by the network to the response u for the given context c . We use $\alpha = .05$ in our experiments.

3.3 Agent Bot

The agent bot (distinct from user bot) models the distribution of the agent response a conditioned on the context c , the user utterance u and the knowledge base \mathcal{KB} , that is, $p(a|c, u, \mathcal{KB})$. This distribution is modeled in four steps as shown in Figure 2. Firstly, the agent bot feeds the context and the last user utterance to the *belief state generator* module which outputs a belief state of slot-value pairs (also referred to as query). Next, the query is executed over the knowledge base and a set of entities e , whose attributes match the values in the query, are returned. The total number of entities, the belief state, the dialog history and the previous user utterance are fed to the *response generator* which outputs a pool $P^a = (\bar{a}_1, \dots, \bar{a}_r)$ of candidate responses in delexicalised format as seen in Figure 3. Finally, the responses in the pool are scored by the *response selector*. Based on these scores, we define

the distribution of the agent response as follows:

$$p(a|c, u, \mathcal{KB}) = \begin{cases} \frac{\exp(s_k^a)}{\sum_{j=1}^m \exp(s_j^a)}, & \text{if } a = \bar{a}_k \in P^a \\ 0, & \text{if } a \notin P^a \end{cases} \quad (4)$$

where s_k^a is the score of the k^{th} candidate response. Note that in equation, we do not show agent utterance being dependant on the belief state since it is calculated internally using context c and previous user utterance u . The candidate response with the highest probability is selected and sent to the user bot to generate the next turn. This interaction between the user and agent bots is repeated until the user bot outputs the end-of-dialogue token.

Next, we discuss in detail about the modules in the agent bot and how these modules are trained. Note that these modules do not share weights with the corresponding modules of the user bot.

3.3.1 Belief State (query) Generator

The aim of the belief state generator is to generate a belief state for the given dialog history and last user utterance. Here, belief state contains the current domain followed by a sequence of key-value pairs of the form $\langle \text{attribute_name}=\text{attribute_value} \rangle$. To achieve this, we define a distribution over the belief states that can be executed over the knowledge base. The belief state generator treats the belief state as a sequence of tokens $q = (w_1^q, \dots, w_t^q)$. We train a GPT-2 network to model the distribution of the belief state tokens given the tokens of the dialog history and user utterance. Once the belief state generator has been trained, a belief state is sampled by greedy sampling and executed over the knowledge base.

3.3.2 Response Generator

This module mimics the response generator of the user bot with the exception that the input to the GPT-2 network comprises the context, the last user utterance, the belief state and the total number of KB entities satisfying the belief state. We provide only the number of entities instead of entire entities to the agent response. This is done as the response would be different for 0, 1 and more than 1 matched entities and further information about entities could be filled while lexicalising the response. The GPT-2 network is used to define an autoregressive distribution over the tokens of the agent response and is trained using maximum likelihood. Once the module is trained, a pool of candidate responses P^a is sampled via nucleus sampling. The

<p>Instruction : You are looking for a particular hotel . Its name is called bridge guest house . Make sure you get hotel type and phone number .</p> <p>User: hi, i am looking for information on the bridge guest house .</p>
<p>GREEDY RESPONSE: Agent Response: i have [value_count] guesthouses and [value_count] hotel -s that fit that criteria . do you have a preference for price range?</p>
<p>NUCLEUS DECODING SAMPLES(Top 5):</p> <p>Response 1: i have [value_count] guesthouses available, how many are in your area or price range? Score : 0.54</p> <p>Response 2: i have [value_count] options, [hotel_name] and [hotel_name], both of which offer free wifi and parking. do you have any other preferences? Score : 0.31</p> <p>Response 3: there are several guesthouses in the [value_pricerange] price range. do you have a preference? Score : 0.54</p> <p>Response 4: i have [value_count] results. what area would you like to stay in? Score : 0.52</p> <p>Response 5: i found the [hotel_name], which is located on the [value_area] side of town in the [value_pricerange] price range. would you like to book a room? Score : 0.89</p>

Figure 4: A goal along with context from the MultiWOZ dialog corpus where responses are generated using Greedy and Nucleus Sampling methods. Response highlighted in blue(highest score) was finally chosen by the model.

response is lexicalised using the values from the belief state before being shown to the user bot. Figure 4 illustrates the advantages of using nucleus sampling for our decoders followed by use of a response selector.

3.3.3 Response Selector

This module outputs the score of each agent response in the candidate pool. To achieve this, the context, the last user utterance and the agent response are fed to the Longformer network architecture. The training of this network as well as the selection of negative samples mimics the training of the response selector for the user bot. Once the model has been trained, it outputs a score s^a for each agent response in the candidate pool.

All the user and agent utterances, belief states and KB results created form the generated dialog.

4 Experiments

4.1 Datasets

To demonstrate the strength of our work we experiment on two different types of tasks - (i) Task oriented dialogs using the MultiWOZ 2.0 dataset

(Budzianowski et al., 2018) (ii) Persona-based conversation generation using the PersonaChat dataset (Zhang et al., 2018b).

4.1.1 Task-Oriented Dialog

MultiWOZ (Budzianowski et al., 2018) is a large scale multi-domain dialogue dataset consisting of 8438 training, 1000 validation and 1000 test conversations distributed across 7 domains. Each conversation is associated with *instructions* which were used by the crowd workers to generate the conversations. 30% of the dataset consists of conversations with a *single* goal while the rest are *multi-goal* dialogues, i.e, conversations accomplish more than one task – example, booking a train followed by making a restaurant reservation.

4.1.2 Persona-based Conversation

PersonaChat (Zhang et al., 2018a) is a large scale non task-oriented dataset containing a set of 1155 distinct characters, each consisting of at least 5 profile sentences. The dataset is collected via Amazon Mechanical Turk where each of the pair of speakers condition their dialogue on a given profile, which is provided. It contains a total of 10,907 dialogs out of which 1000 dialogs are used for validation while 968 dialogs are used for testing.

4.2 Data Generation using Simulators

MultiWoz: As mentioned previously, our simulator allows the generation of new conversations based on *instructions*. In our experiments, we operate our simulators using 5% (421/8438), 20% (1684/8438) and 100% of the original training data. For 5% and 20%, we use the instructions of the remaining training datasets (i.e. remaining 95% and 80% respectively) to generate simulated conversations. The simulated conversations are added to the original conversations, thereby ensuring that the size of the datasets remains unchanged. In case of 100% we train our simulators on the entire training data and then generate simulated conversations using the instructions of the same data. The simulated conversations are then appended to the original conversations. The resulting dataset has twice as many conversations as the original dataset.

Recall that each conversation requires KB and belief state by the agent. Our agent simulator generates queries for the simulated data using belief state generators as described earlier. While training the end-task dialog models using the simulated data, we use these generated values as the oracle

belief state for our simulated data. Similar to existing work on this dataset, we use delexicalised agent utterances using the format followed by MultiWOZ (Budzianowski et al., 2018) which are later updated with KB values based on the results of the query. Hyper-parameter settings are available in supplementary notes.

PersonaChat: In case of PersonaChat dataset, we train a single user bot to mimic both the users of a conversation. To generate the utterance for a specific user, the corresponding persona is fed to the bot along with the dialog context. Thus, a single bot is able to simulate a conversation between two distinct personas. We use 5%(447/8939), 20%(1788/8939) and 100% of the training data in our experiments just like in MultiWoz.

4.3 End-Task Models

MultiWoz: We experiment with two recent end-task models: Soloist (initialized with GPT2-small) (Peng et al., 2021) and MinTL-T5 (initialized with T5-small) (Lin et al., 2020b). Soloist is a transformer based auto-regressive model that incorporates dialog modules, including the query generator, into a single network. The original model was pre-trained on a variety of dialog tasks and then applied to MultiWoz in few-shot settings. However, we use an untrained instance of Soloist, initialized only using GPT2-small, as our goal is only to demonstrate that simulated-data based augmentation can help train useful end-task models. MinTL-T5 is another recent model that also uses pre-trained transformers along with an improved method for updating belief states.

PersonaChat: We use GPT2-small based end-task model to test effectiveness of simulated chat.

4.4 Baselines

MultiWoz: As baselines, we study the performance of our end-task models based on Soloist and MinTL-T5, when they are trained in the absence of data augmentation. We look at non-augmentation based recent baseline model DAMD (Zhang et al., 2020). Additionally, we compare the performance of our simulation based augmentation, against a recent approach - DAMD-MADA (Zhang et al., 2020) which uses dialog-states based augmentation and PARG-TSCP (Gao et al., 2020) which uses paraphrases to help improve response generation done by TSCP (Lei et al., 2018). Additionally, we experiment with a T5 (Raffel et al., 2020) based

Model	5%				20%				100%			
	B	I	S	C	B	I	S	C	B	I	S	C
DAMD (Zhang et al., 2020)	9.5	48.4	25.8	46.6	12.4	54.1	32.3	55.6	16.9	72.7	60.3	83.5
DAMD-MADA (Zhang et al., 2020)	9.5	50.5	33.9	51.8	13.1	60.1	42.9	64.6	16.6	76.3	60.4	85.0
PARG - TSCP (Gao et al., 2020)	13.1	53.5	39.2	59.4	13.0	63.6	48.9	69.2	15.4	80.1	63.1	87.0
Soloist (Peng et al., 2021)	12.5	52.5	32.9	55.2	14.0	60.9	50.0	69.5	16.8	80.5	63.2	88.6
Soloist (Paraphrase)	10.9	57.8	38.4	59.0	13.9	62.9	52.7	71.7	16.4	82.2	62.6	88.8
Soloist (Sim. Aug.)	14.0	55.3	41.4	62.3	14.8	70.5	56.4	78.3	17.6	76.5	60.9	86.3
MinTL-T5-Small (Lin et al., 2020b)	12.5	50.9	33.9	55.7	15.8	63.5	48.8	72.0	17.4	80.1	64.7	89.8
MinTL-T5-Small (Sim. Aug)	13.1	57.6	36.1	60.0	16.0	68.0	55.1	76.6	18.5	79.5	57.1	86.8

Table 1: Performance of models using varying sizes of MultiWOZ 2.0 dataset (B,I,S,C stand for BLEU, Inform, Success and Combined scores respectively). ‘Sim. Aug.’ refers to the use of our simulated data. Bold values indicate the highest scores.

Model	5%		10%		20%		100%	
	PPL	Hits@1	PPL	Hits@1	PPL	Hits@1	PPL	Hits@1
GPT2-small	40.4	11.1	41.2	13.3	33.6	13.8	35.8	14.6
GPT2-small (Sim. Aug.)	41.1	12.2	39.1	14.5	40.2	14.0	43.0	15.9
Lost In Conversation (Dinan et al., 2019)	-	-	-	-	-	-	-	17.3

Table 2: Performance of models using varying sizes of Persona dataset. PPL stands for Perplexity

paraphrase generation model² fine-tuned on the PAWS dataset (Zhang et al., 2019) – we use this model to generate paraphrases and augment training data and refer this model as Soloist(Paraphrase) in Table 1(Details in Supplementary).

PersonaChat: We compare the performance of an end-task model based on GPT2-small with and without augmented data. We report the performance of the ‘Lost in Conversation’ model (Dinan et al., 2019), the winner of the ConvAI2 challenge.³

4.5 Metrics

MultiWoz: We evaluate the usefulness of our generated data by using it to train a dialog model for the end-task. We therefore use BLEU (B), *Inform* (I) and *Success* (S) rates as defined by Budzianowski et al., along with Combined(C) score (Mehri et al., 2019) given by, $BLEU + 0.5 \times (Inform + Success)$. While BLEU evaluates the fluency of the generated response, *Inform* and *Success* measure the relevance of the agent utterances. Specifically, the *Inform* Rate measures the correctness of the entity provided by the agent, while the *Success* Rate measures how often the agent was able to provide correct attributes when requested by the user. We note that there are minor (but significant) differences in delexicalization used by different models and this makes the direct comparison using the metrics inaccurate. In our experiments, we use the delexicalisation scheme

used by (Budzianowski et al., 2018) and their⁴ task-evaluation scripts to report results. Hence we see slight drop in the scores in table 1 for Soloist and MinTL models compared to the scores cited in their respective papers(see supplementary for details).

PersonaChat: We use *Hits@1* and *Perplexity* as in (Zhang et al., 2018a) to evaluate the models. While *Perplexity*(PPL) measures the log likelihood of the correct sequence, *Hits@1* scores the responses in a next-utterance (response) prediction task – given an input context and persona, the goal is to predict the correct(ground-truth) response from a set consisting of other incorrect responses.

4.6 Results

We study the following research questions: (1) Would the new student model trained on simulated conversations along with crowd generated low(95% + 5%) and medium (20% + 80%) data perform better than teacher models trained only on only low(5%) and medium(20%) resourced data? (2) How does simulated-data based augmentation compare with recent work on augmentation? (3) How does the student model perform compared to models trained on 100% human generated data? (4) Can we use this technique to improve the models trained already on 100% human generated data? (5) What is the qualitative difference between simulated and crowd-sourced chats?

²https://huggingface.co/Vamsi/T5_Paraphrase_Paws

³<https://parl.ai/projects/convai2/>

⁴<https://github.com/budzianowski/multiwoz>

4.6.1 Use of simulated data in End-task

MultiWoz: Table 1 shows the use of simulated data helps improve performance in low data settings (5%) and medium data settings (20%). The use of simulated data helps improve performance of both Soloist and MinTL-T5 (gains of 7-8% in combined metric) in 5% data setting. We also see a higher improvement in medium data setting i.e. an increase of 7% in combined score for MinTL and 13% for Soloist model suggesting the effectiveness of our method in low and medium resource setting.

We further compare the performance of our augmented data w.r.t the original 100% dataset for both Soloist and MinTL. Adding additional 80% augmented data to our Soloist model trained on 20% dataset substantially increases the combined score from 69.5 to 78.3 although it lags behind the model when trained on 100% human generated dataset which gives a combined score of 88.6. Similarly for MinTL, the combined score improves from 72 to 76.6 but falls short of the performance on original 100% human generated dataset. Our model behaves according to the common knowledge that noisy student models do not perform as good as a teacher model trained on a similar sized data(100%). Through this experiment we show how the diverse knowledge contained in Longformer could be transferred to our augmented dataset by using it as a selector. The teacher generator model generates a list of diverse candidate responses(by top-p sampling) which is provided to the selector to pick the most relevant response helping the selector induce its knowledge through the process. Thus we see that the method improves the models trained on low-medium sized datasets and can be used effectively when larger datasets are not available.

Our data generation technique is not able to increase the performance of models trained on 100% human generated data(i.e.200% not performing better than 100%) the reasons for which are discussed under human study section. Additional qualitative results on the MultiWOZ dataset are available in the supplementary material.

Comparing with other baseline models, student models trained on simulated data(Sim. Aug.) on Soloist and MinTL outperforms existing end-task models such as dialog-state based augmentation (DAMD-MADA) and paraphrase-based augmentation(PARG and Soloist(paraphrase)) in low and medium data settings as seen in Table 1. Soloist

(Sim. Aug) gets combined score of 62.3 in 5% data compared to 59.4 obtained by the best performing augmentation based baseline model PARG-TSCP. Similarly Soloist (Sim. Aug) scores 78.3 in 20% data compared to 69.2 obtained by PARG-TSCP.

Persona Chat: Table 2 shows improvement in *Hits@1* when the GPT2-small based end-task model is trained on simulation-based augmented data. Gain in *Hits@1* (2-10%) demonstrates that the model is able to learn the context and persona of the given characters better which results in better conversations. The augmented data helps improve the performance of a simple GPT2-small model(fine-tuned on dataset) in *Hits@1* from 14.6 to 15.9 which is very close to 17.3 achieved by Lost in Conversation. The Perplexity (PPL) gives mixed results suggesting that the language *style* of simulated conversations differs from the language *style* of the original dataset. This is because GPT2-small incorporates its pre-trained knowledge in the simulated conversations. However the fact that *Hits@1* consistently increases across all dataset sizes suggests that the generated simulated conversations help the model capture the context and persona better despite changing its language *style*.

	Original Data	Simulated Data
Relevance	4.7	4.0
Grammar	4.6	4.5
User Bot Fluency	4.5	4.1
Agent Bot Fluency	4.6	4.1

Table 3: Human evaluation scores(scale of 1-5) on original and simulated data

4.6.2 Human Study

To assess the qualitative difference between simulated data and crowd-sourced data, we conducted a blind human-study involving six participants. Participants were presented the crowd-annotation instructions from MultiWOZ and were asked to assess the quality of a pair of dialogs corresponding to the same instruction – one generated by the crowd workers (from the original dataset) and the other generated by our simulators. The participants were blind to the source of the dialog (crowd or simulator). Each dialog was scored on the Likert scale(1-5) by answering the following questions: 1) ‘How relevant is the dialog w.r.t the dialog generation instruction?’ 2) ‘How good is the grammar of the utterances?’ 3) ‘Are the *user* utterances responding to *agent* utterances fluently and meaningfully?’ 4) ‘Are the *agent* utterances responding to

user utterances fluently and meaningfully?'. Each participant evaluated 25 pairs resulting in a total evaluation set of 150 pairs.

As seen in Table 3, the simulated data is of high quality with the bots scoring well on fluency as well as grammar. As expected, there is a slight deterioration in relevance to the instructions compared to crowd-sourced conversations. This happens because the simulated conversations may not use all the information present in the instructions. This also answers why the simulated data doesn't increase the performance of models trained on 100% dataset in Table 1. In lower data settings, the original dialogs of the remaining instructions (i.e. remaining 95% in case of 5% training data) were not part of dialog used for the end-task model. Hence, the simulated data provided new dialogs that were never seen by the model. In 100% data setting, since the model had already seen the original dialogs, the simulated dialogs did not improve the performance as they lacked some relevance w.r.t the instructions when compared with human generated data. The same issue causes the model trained on 100% original dataset to perform better than our augmented datasets i.e. (20+80)% and (5+95)%.

5 Conclusion

In this paper, we demonstrated a dialog generation framework that mimics the data creation process employed by crowd workers. We find that our method is able to generate meaningful conversations that aids the training of end-task dialog models in low resource data settings. The use of additional simulated data to train end-task dialog models result in a performance improvement of 7-13% in low resource settings of MultiWOZ 2.0 dataset and 2-10% increase in Hits@1 in case of PersonaChat. The simulation-framework does not make strict assumptions about the domain or dataset and can be applied to diverse dialog tasks such as task-oriented dialog and persona-based chat. In future, it would be interesting to compare the strengths of different augmentation methods and how they may be effectively combined.

References

Ateret Anaby-Tavor, Boaz Carmeli, Esther Goldbraich, Amir Kantor, George Kour, Segev Shlomov, Naama Tepper, and Naama Zwerdling. 2020. *Do not have enough data? deep learning to the rescue!* In *The Thirty-Fourth AAAI Conference on Artificial Intelli-*

gence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020, pages 7383–7390.

Ali Araabi and Christof Monz. 2020. *Optimizing transformer for low-resource neural machine translation*. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3429–3435, Barcelona, Spain (Online). International Committee on Computational Linguistics.

Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. *Longformer: The long-document transformer*. *arXiv preprint arXiv:2004.05150*.

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. 2020. *Language models are few-shot learners*. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.

Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, 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*, pages 5016–5026, Brussels, Belgium. Association for Computational Linguistics.

Gal Chechik, Varun Sharma, Uri Shalit, Samy Bengio, Soeren Ed Sonnenburg, Vojtech Ed Franc, Elad Ed Yom-Tov, and Michele Ed Sebag. 2010. *Large scale online learning of image similarity through ranking*. *Journal of Machine Learning Research*.

Emily Dinan, Varvara Logacheva, Valentin Malykh, Alexander H. Miller, Kurt Shuster, Jack Urbanek, Douwe Kiela, Arthur Szlam, Iulian Serban, Ryan Lowe, Shrimai Prabhumoye, Alan W. Black, Alexander I. Rudnicky, Jason Williams, Joelle Pineau, Mikhail S. Burtsev, and Jason Weston. 2019. *The second conversational intelligence challenge (con-vai2)*. *CoRR*, abs/1902.00098.

Zi-Yi Dou, Keyi Yu, and Antonios Anastasopoulos. 2019. *Investigating meta-learning algorithms for low-resource natural language understanding tasks*. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the*

- 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pages 1192–1197, Hong Kong, China. Association for Computational Linguistics.
- Wenchao Du and Alan Black. 2018. Data augmentation for neural online chats response selection. In *Proceedings of the 2018 EMNLP Workshop SCAI: The 2nd International Workshop on Search-Oriented Conversational AI*, pages 52–58, Brussels, Belgium. Association for Computational Linguistics.
- Silin Gao, Yichi Zhang, Zhijian Ou, and Zhou Yu. 2020. Paraphrase augmented task-oriented dialog generation. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Geoffrey E. Hinton, Oriol Vinyals, and J. Dean. 2015. Distilling the knowledge in a neural network. *ArXiv*, abs/1503.02531.
- Elad Hoffer and Nir Ailon. 2015. Deep metric learning using triplet network. In *International Workshop on Similarity-Based Pattern Recognition*, pages 84–92. Springer.
- 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*.
- Yutai Hou, Meng Fang, Wanxiang Che, and Ting Liu. 2019. A corpus-free state2seq user simulator for task-oriented dialogue. In *Chinese Computational Linguistics. CCL 2019. Lecture Notes in Computer Science, vol 11856. Springer, Cham*.
- Varun Kumar, Ashutosh Choudhary, and Eunah Cho. 2021. Data augmentation using pre-trained transformer models.
- Wenqiang Lei, Xisen Jin, Min-Yen Kan, Zhaochun Ren, Xiangnan He, and Dawei Yin. 2018. Sequicity: Simplifying task-oriented dialogue systems with single sequence-to-sequence architectures. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1437–1447, Melbourne, Australia. Association for Computational Linguistics.
- Chien-Wei Lin, Vincent Auvray, Daniel Elkind, Arijit Biswas, Maryam Fazel-Zarandi, Nehal Belgamwar, Shubhra Chandra, Matt Zhao, Angeliki Metallinou, Tagyoung Chung, Charlie Shucheng Zhu, Suranjit Adhikari, and Dilek Hakkani-Tür. 2020a. Dialog simulation with realistic variations for training goal-oriented conversational systems. *CoRR*, abs/2011.08243.
- Zhaojiang Lin, Andrea Madotto, Genta Indra Winata, and Pascale Fung. 2020b. MinTL: Minimalist transfer learning for task-oriented dialogue systems. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3391–3405, Online. Association for Computational Linguistics.
- Nikolaos Malandrakis, Minmin Shen, Anuj Goyal, Shuyang Gao, Abhishek Sethi, and Angeliki Metallinou. 2019. Controlled text generation for data augmentation in intelligent artificial agents. *Proceedings of the 3rd Workshop on Neural Generation and Translation*.
- Shikib Mehri, Tejas Srinivasan, and Maxine Eskenazi. 2019. 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.
- Alexandros Papangelis, Yi-Chia Wang, Piero Molino, and Gokhan Tur. 2019. Collaborative multi-agent dialogue model training via reinforcement learning. In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 92–102, Stockholm, Sweden. Association for Computational Linguistics.
- 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.
- Kun Qian, Wei Wei, and Zhou Yu. 2021. A student-teacher architecture for dialog domain adaptation under the meta-learning setting. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(15):13692–13700.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2018. Language models are unsupervised multitask learners.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Pararth Shah, Dilek Hakkani-Tür, Gokhan Tür, Abhinav Rastogi, Ankur Bapna, Neha Nayak, and Larry Heck. 2018. Building a conversational agent overnight with dialogue self-play. *arXiv preprint arXiv:1801.04871*.
- Longshaokan Wang, Maryam Fazel-Zarandi, Aditya Tiwari, Spyros Matsoukas, and Lazaros Polymenakos. 2020. Data augmentation for training dialog models robust to speech recognition errors. *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*.
- Sean Welleck, Jason Weston, Arthur Szlam, and Kyunghyun Cho. 2019. Dialogue natural language inference. pages 3731–3741.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R'emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. Huggingface's transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.

Lijun Wu, Fei Tian, Yingce Xia, Yang Fan, Tao Qin, Lai Jian-Huang, and Tie-Yan Liu. 2018. Learning to teach with dynamic loss functions. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.

Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V. Le. 2020. Self-training with noisy student improves imagenet classification. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695.

Kang Min Yoo, Hanbit Lee, Franck Dernoncourt, Trung Bui, Walter Chang, and Sang goo Lee. 2020a. Variational hierarchical dialog autoencoder for dialog state tracking data augmentation.

Kang Min Yoo, Hanbit Lee, Franck Dernoncourt, Trung Bui, Walter Chang, and Sang goo Lee. 2020b. Variational hierarchical dialog autoencoder for dialog state tracking data augmentation.

Kang Min Yoo, Youhyun Shin, and Sang-goo Lee. 2019. Data augmentation for spoken language understanding via joint variational generation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33:7402–7409.

Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018a. Personalizing dialogue agents: I have a dog, do you have pets too? *CoRR*, abs/1801.07243.

Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018b. Personalizing dialogue agents: I have a dog, do you have pets too? In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia. Association for Computational Linguistics.

Yichi Zhang, Zhijian Ou, and Zhou Yu. 2020. Task-oriented dialog systems that consider multiple appropriate responses under the same context. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):9604–9611.

Yuan Zhang, Jason Baldridge, and Luheng He. 2019. PAWS: Paraphrase Adversaries from Word Scrambling. In *Proc. of NAACL*.

A Appendix Overview

Section B provides information on the hyperparameter settings of models used in the experiments. Section C does the cost analysis and Section D discusses about the inconsistencies in evaluation of

various models in details. The paraphrase based augmentation method used to train an end-task Soloist (Peng et al., 2021) model has been described in Section E. We also show further experimental results in Section F and G while additional qualitative study is shown in Section H.

B Hyperparameter Settings

MultiWoz: We create separate user bots and agent bots along with their constituent modules consisting of query models (for tracking belief state), response generators and response selectors. We use GPT2-small (12 layered, 768 hidden size, 117M parameters) from the 'Transformers' library by Huggingface (Wolf et al., 2019) for the response generator. For response selectors, we use Longformers (12 layered, 1024 hidden size, 149M parameters) (Beltagy et al., 2020) for both user and agent models. We train on 5%, 10%, 30% and 100% of the training data with a learning rate of 1e-5. Adam optimizer with default settings is used for all the models.

PersonaChat: Similar to MultiWOZ, the response generators use GPT2-small while response selectors use Longformers. There is no belief state generator and only single user model is trained i.e. no separate agent model exists. Adam optimizer is used with a learning rate of 1e-5.

C Cost comparison

The response generator, belief state generator and response selector models (total 5) each take 1 day on a single V100 GPU to generate the dialogues (for 100% data). MultiWOZ data creation, on the other hand, required 1249 workers for the entire process. An Amazon EC2 P3 instance costs \$3.06 per hour in an On-Demand setup costing less than \$400 for the entire process. Generating 10.4K dialogs with 1249 workers (2 workers per conversation) means 15-16 dialogues per worker and assuming they take an hour to generate the conversation with a minimum wage of \$6 per hour payment, it leads upto \$7.5k. Our method is clearly both cost and time effective when compared with the crowdsource workers.

D Evaluation Inconsistency

We noticed the delexicalisation used in models such as PARG (Gao et al., 2020), DAMD (Zhang et al., 2020) and MinTL (Lin et al., 2020b) was different from the delexicalisation used in original

MultiWOZ code. Since, the delexicalised agent responses are used by the official evaluation script to score the dialogs there were discrepancies in the evaluation of different models. Some delexicalisation differences included use of tokens such as [restaurant_phone] or [hotel_address] in original format while some models used [value_phone] for such requestables. This lead to difference in evaluation script since the official script looks at the domain of requestable i.e. [train_address] should have been used if the utterance domain is 'Train' and not [hotel_address]. Replacing these tokens with a generic [value_address] as used in other models reduces the complexity of such dialogs leading to higher innform and success scores. The Soloist and MinTL models used in our experiments are trained on the delexicalisation used in the MultiWoz original code⁵. This leads to a small drop in the overall performance of the models.

E Soloist Paraphrase Model

In order to train a soloist (Peng et al., 2021) end task model using paraphrases generated from original training data, we use T5 (Raffel et al., 2020) based paraphrase generation model⁶ fine-tuned on the PAWS dataset (Zhang et al., 2019). We also add training paraphrase data from PARG (Gao et al., 2020) to create a mixture of corresponding paraphrases for each utterance. In low resource setting(5% and 20% of training dataset) we generate enough paraphrases to take the whole size of augmented data to 100% of original training data. i.e. We add paraphrases equaling the size of 95% of total data in case of 5% and 80% of total data in case of 20%. Adding paraphrases from T5 fine-tuned model and PARG provides the model with diversity inn training data for final end-task model.

F MultiWoz 2.1 Performance

In order to check our performance on the revised version of MultiWoz, we experimented with MultiWoz 2.1 as well. Since not many baseline models have results on this version of the dataset, we used MultiWOZ 2.0 as the dataset to compare our model with other baselines. Table 4 presents results using the Soloist end-task model for MultiWOZ 2.1. Additional data generated by our simulators results in a significant improvement on the Combined metric

⁵<https://github.com/budzianowski/multiwoz>

⁶https://huggingface.co/Vamsi/T5_Paraphrase_Paws

for both the *oracle* belief states(O) and the *generated* belief states(G) in the end-task model. The use of our simulated data results in a 18-25% improvement in combined metrics for low-medium resource setting. The performance of generated belief states shows that the query generator produces accurate queries most of the time.

ORIGINAL INSTRUCTIONS : You are looking for a **restaurant** . The restaurant should be in the **expensive** price range and should serve **italian** food . Once you find the **restaurant** you want to book a table for **5 people** at **11:30** on **sunday** . If the booking fails how about **10:30** . Make sure you get the **reference number** .

User: hi, i'm trying to find an **expensive italian** restaurant.
Agent: i have found several [value_pricerange] [value_food] restaurant -s . is there a specific area of town that you are looking for ?
User: i would like to go to the **south**, please.
Agent: there is an [value_pricerange] restaurant called [restaurant_name] in the [value_area] . would you like to make a reservation?
User: yes, i'd like to book a table for **5** at **11:30** on **sunday**.
Agent: i am sorry , but we did not get that reservation . would you like to try another day or time slot?
User: can you try for **10:30** instead?
Agent: booked at [restaurant_name] for [value_count] people . reference number is [restaurant_reference] .
User: thank you so much. that's all i needed. bye.
Agent: thank you for contacting towninfo centers . enjoy your dining experience.

PERTURBED INSTRUCTIONS : You are looking for a **restaurant** . The restaurant should be in the **cheap** price range and should serve **indian** food . **Restaurant should be in north**. Once you find the restaurant you want to book a table for **5 people** at **11:30** on **sunday**. If the booking fails how about 10:30. Make sure you get the **reference number**.

User: hi! can you help me find a **cheap indian** restaurant in the **north**?
Agent: sure , [restaurant_name] is in the [value_area] part of town . it is located at [restaurant_address].
User: i would like to book a table for **5 people** at **11:30** on **sunday**.
Agent: i am sorry booking was unsuccessful , do you need to find another restaurant ?
User: can you try for **10:30** ?
Agent: i have booked you a table for [value_count] people at [value_time] on [value_day] . your reference number is [restaurant_reference] .
User: thank you so much. that is all i needed today. bye.
Agent: enjoy your meal !

Figure 5: Examples of generated samples using original instructions from MultiWOZ dataset against perturbed instructions with minor changes. The generated dialogues show the robustness of our generator model which is able to generate an entirely new conversation with slight variations in the goal.

G Single Goal Dialogs

We also evaluate the performance of our model when trained on single goal dialogs of MultiWOZ 2.1 to test whether the model can learn generating multiple goal dialogs using just single goal data. We train separate models(generators and selectors) for each domains and the simulated single dialog chats were combined (pair of

Data Size	Belief State	Without Augmentation				With Augmentation			
		BLEU (B)	Inform (I)	Success (S)	Comb. (C)	BLEU (B)	Inform (I)	Success (S)	Comb. (C)
5%	O	12.5	47.3	33.8	53.1	12.2	67.8	40.4	66.3
10%	O	13.8	51.9	37.6	58.5	13.8	68.1	49.5	72.6
20%	O	15.2	61.3	48.7	70.2	13.7	81.8	64.1	86.7
100%	O	16.4	85.8	74.4	96.5	17.2	80.1	67.7	91.2
5%	G	12.5	39.6	26.8	45.6	12.0	54.8	32.3	55.6
10%	G	13.5	46.2	31.9	52.6	13.7	52.6	38.5	59.3
20%	G	15.1	55.4	42.8	64.2	13.5	71.2	54.7	76.5
100%	G	15.9	80.1	64.3	88.1	16.2	75.8	58.5	83.4

Table 4: Performance of the query generator and the end-task model Soloist (Peng et al., 2021) when trained on different dataset sizes with and without the use of additional simulated data for MultiWOZ 2.1. ‘O’ and ‘G’ in column two refer to the use of ‘oracle’ and ‘generated’ belief states respectively.

Dataset	Belief State	BLEU		Inform		Success		Combined	
		w.	w/o.	w.	w/o.	w.	w/o.	w.	w/o.
5%	Oracle	9.3	8.0	88.5	82.0	64.1	60.6	85.7	79.4
10%	Oracle	10.8	10.8	92.5	84.5	75.2	69	94.7	87.6
30%	Oracle	11.9	12.4	90.3	82.3	73.0	65.5	93.6	86.3
100%	Oracle	-	14.9	-	82.8	-	78.3	-	95.5

Table 5: Score against single goal conversations in test dataset with oracle belief state

Dataset	Belief State	BLEU		Inform		Success		Combined	
		w/o.	w.	w/o.	w.	w/o.	w.	w/o.	w.
5%	Oracle	7.1	9.2	63.2	73.2	34.4	42.6	55.9	67.1
10%	Oracle	9.6	10.8	63.8	78.2	38.9	52.9	61.0	76.4
30%	Oracle	9.8	12.4	66.6	77.0	34.9	52.3	60.6	77.1
100%	Oracle	15.9	-	72.8	-	63.7	-	84.2	-

Table 6: Score against entire test dataset containing both single and multiple goal conversations with oracle belief state

dialogs) from different domains using a basic script and trained the final end-task model. Our method achieves significant improvement over non-augmented(w/o.) dataset as seen in Table 5, 6, 7 and 8. Tables 5 and 7 show the performance of the model for the oracle and generated belief state on single domain goal conversations of the test dataset. We see an improvement of 7-8 % in combined score across all dataset sizes on applying our augmentation technique(w.). We are able to achieve a combined score of 94.7% with just 10% of the dataset which is very close to the combined score of 95.53% when trained on the entire dataset.

Tables 6 and 8 show the performance of the model when we use the oracle and generated belief state on the entire test dataset. We see a massive improvement in both the oracle and generated belief state setting. While the oracle belief state results improve the combined score by 20.02%, 25.2% and 17.98% for 5%, 10% and 30% of the dataset

respectively, we see an even bigger improvement of 146.11%, 21.82% and 29.12% when using generated belief states for 5%, 10% and 30% of dataset respectively. The augmentation helps in improving the combined score by a huge margin thus bringing them close to the score of entire dataset(100 percent). The results show that simulated data generated from single goal dialogs can also do a good job at generalising to multiple goal dialogs. This insight would be useful in combining various single goal dialogs from different datasets.

H Qualitative Study - Instruction Perturbation

We now present a qualitative study demonstrating how our simulator is able to accommodate changes to instructions and reflect them in a conversation. Figure 5 shows the generated dialogs from an original instruction in MultiWOZ and another from instructions created by perturbing the original instructions. The generated dialogs demonstrate the

Dataset	Belief State	BLEU		Inform		Success		Combined	
		w/o.	w.	w/o.	w.	w/o.	w.	w/o.	w.
5%	Generated	7.6	9.6	40.0	80.9	30.1	60.6	42.7	80.4
10%	Generated	10.7	10.7	77.4	84.5	61.5	66.3	80.1	86.2
30%	Generated	12.2	11.8	77.8	84.1	60.2	63.3	81.2	85.5
100%	Generated	14.8	-	81.4	-	76.1	-	93.5	-

Table 7: Score against single goal conversations in test dataset with generated belief state

Dataset	Belief State	BLEU		Inform		Success		Combined	
		w/o.	w.	w/o.	w.	w/o.	w.	w/o.	w.
5%	Generated	6.8	9.8	19.3	54.7	10.2	31.9	21.6	53.1
10%	Generated	9.5	10.7	52.3	61.2	29.9	40.6	50.6	61.6
30%	Generated	9.5	12.4	50.9	59.4	24.9	38.3	47.4	61.2
100%	Generated	15.9	-	66.2	-	55.4	-	76.7	-

Table 8: Score against entire test dataset containing both single and multiple goal conversations with generated belief state

robustness of our generator model which is able to produce new and meaningful conversations using new entities from the perturbed instructions. Further, the dialogues generated are very different from each other which shows the wide variety of conversations the simulators are capable of producing, when provided with similar goals.