

Zero-Shot Dialogue State Tracking via Cross-Task Transfer

Zhaojiang Lin^{1*}, Bing Liu², Andrea Madotto^{1*}, Seungwhan Moon², Paul Crook²,
Zhenpeng Zhou², Zhiguang Wang², Zhou Yu³, Eunjoon Cho², Rajen Subba²,
Pascale Fung¹

¹The Hong Kong University of Science and Technology

²Facebook, ³Columbia University

zlinao@ust.hk, bingl@fb.com

Abstract

Zero-shot transfer learning for dialogue state tracking (DST) enables us to handle a variety of task-oriented dialogue domains without the expense of collecting in-domain data. In this work, we propose to transfer the *cross-task* knowledge from general question answering (QA) corpora for the zero-shot DST task. Specifically, we propose TransferQA, a transferable generative QA model that seamlessly combines extractive QA and multi-choice QA via a text-to-text transformer framework, and tracks both categorical slots and non-categorical slots in DST. In addition, we introduce two effective ways to construct unanswerable questions, namely, *negative question sampling* and *context truncation*, which enable our model to handle “none” value slots in the zero-shot DST setting. The extensive experiments show that our approaches substantially improve the existing zero-shot and few-shot results on MultiWoz. Moreover, compared to the fully trained baseline on the Schema-Guided Dialogue dataset, our approach shows better generalization ability in unseen domains.

1 Introduction

Virtual assistants are designed to help users perform daily activities, such as travel planning, online shopping and restaurant booking. Dialogue state tracking (DST), as an essential component of these task-oriented dialogue systems, tracks users’ requirements throughout multi-turn conversations as dialogue states, which are typically in the form of a list of slot-value pairs. Training a DST model often requires extensive annotated dialogue data. These data are often collected via a Wizard-of-Oz (Woz) (Kelley, 1984) setting, where two workers converse with each other and annotate the dialogue states of each utterance (Wen et al., 2017; Budzianowski et al., 2018; Moon et al., 2020), or with a Machines Talking To Machines

(M2M) framework (Shah et al., 2018), where dialogues are synthesized via the system and user simulators (Campagna et al., 2020; Rastogi et al., 2020; Lin et al., 2021b). However, both of these approaches have inherent challenges when scaling to large datasets. For example, the data collection process in a Woz setting incurs expensive and time-consuming manual annotations, while M2M requires exhaustive hand-crafted rules for covering various dialogue scenarios.

In industrial applications, virtual assistants are required to add new services (domains) frequently based on user’s needs, but collecting extensive data for every new domain is costly and inefficient. Therefore, performing zero-shot prediction of dialogue states is becoming increasingly important since it does not require the expense of data acquisition. There are mainly two lines of work in the zero-shot transfer learning problem. The first is cross-domain transfer learning (Wu et al., 2019; Kumar et al., 2020; Rastogi et al., 2020; Lin et al., 2021a), where the models are first trained on several domains, then zero-shot to new domains. However, these methods rely on a considerable amount of DST data to cover a broad range of slot types, and it is still challenging for the models to handle new slot types in the unseen domain. The second line of work leverages machine reading question answering (QA) data to facilitate the low-resource DST (i.e., cross-task transfer) (Gao et al., 2020). However, the method of Gao et al. (2020) relies on two independent QA models, i.e., a span extraction model for non-categorical slots and a classification model for categorical slots, which hinders the knowledge sharing from the different types of QA datasets. Furthermore, unanswerable questions are not considered during their QA training phase. Therefore, in a zero-shot DST setting, the model proposed by Gao et al. (2020) is not able to handle “none” value slots (e.g., unmentioned slots) that present in the dialogue state.

* Work done during internship at Facebook

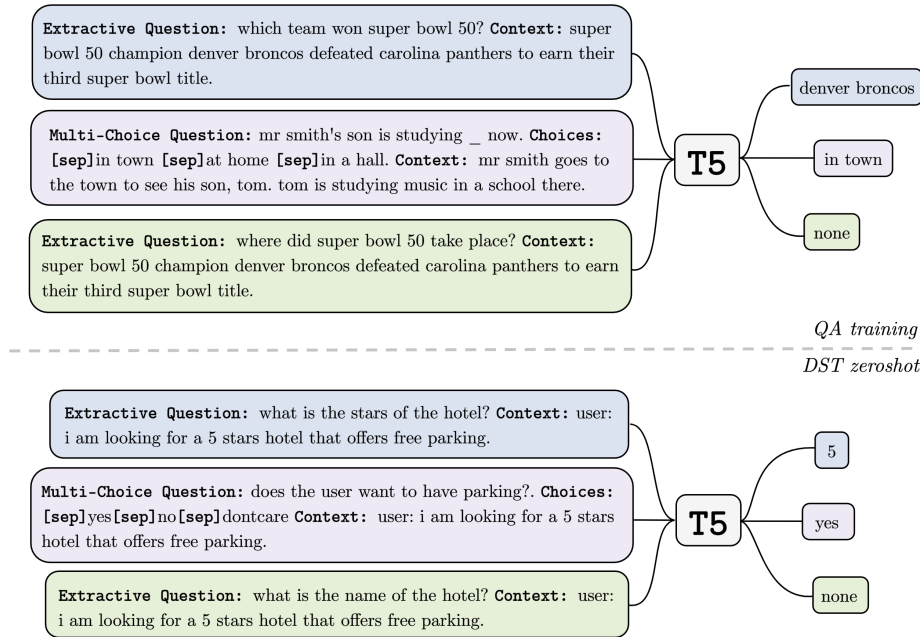


Figure 1: A high-level representation of the cross-task transfer for zero-shot DST (best viewed in color). During the QA training phase (top figure), the unified generative model (T5) is pre-trained on QA pairs of extractive questions (blue), multiple-choice questions (purple), and negative extractive questions (green). At inference time for zero-shot DST (bottom figure), the model predicts slot values as answers for synthetically formulated extractive questions (for non-categorical slots) and multiple-choice questions (for categorical slots). Note that the negative QA training allows for the model to effectively handle “none” values for unanswerable questions.

In this paper, to address the above challenges, we propose TransferQA, a unified generative QA model that seamlessly combines extractive QA and multi-choice QA via a text-to-text transformer framework (Raffel et al., 2020; Khashabi et al., 2020). Such design not only allows the model to leverage both extractive and multi-choice QA datasets, but also provides a simple unified text-to-text interface for tracking both categorical slots and non-categorical slots. To handle the “none” value slots in a zero-shot DST setting, we introduce two effective ways to construct unanswerable questions, namely, *negative question sampling* and *context truncation*, which simulate the out-of-domain slots and in-domain unmentioned slots in multi-domain DST. We evaluate our approach on two large multi-domain DST datasets: MultiWoz (Budzianowski et al., 2018; Eric et al., 2020) and Schema-Guided Dialogue (SGD) (Rastogi et al., 2020). The experimental results suggest that our proposed model, *without using any DST data*, achieves a significantly higher joint goal accuracy compared to previous zero-shot DST approaches. Our contributions are summarized as the following:

- We propose TransferQA, the first model that

performs domain-agnostic DST *without using any DST training data*.

- We introduce two effective ways to construct unanswerable questions, namely, negative question sampling and context truncation, which enable our model to handle “none” value slots in the zero-shot DST setting;
- We demonstrate the effectiveness of our approach in two large multi-domain DST datasets. Our model achieves 1) the state-of-the-art zero-shot and few-shot results on MultiWoz and 2) competitive performance compared to a fully trained baseline on the SGD dataset.

2 Methodology

2.1 Text-to-Text Transfer Learning for DST

In multi-choice QA, each sample consists of a context passage \mathcal{C} , a question q_i , multiple answer candidates $\mathcal{A} = \{a_1, a_2, \dots, a_n\}$, and the correct answer a_i . In extractive QA, answer candidates are not available, and \mathcal{A} become an empty set $\mathcal{A} = \emptyset$. Therefore, in QA training, the models learn to predict the answer a_i to a question q_i by reading the

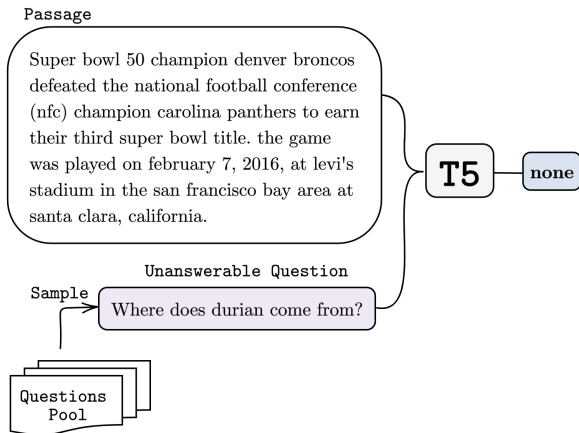


Figure 2: Negative sampling strategy for adding unanswerable questions to the training. Given a passage, we randomly sample a question from other passages and train the QA model (T5) to predict “none”.

context passage \mathcal{C} and answer candidates \mathcal{A} (if available), while in DST inference, the models predict the value a_i of a slot q_i by reading the dialogue history \mathcal{C} and value candidates \mathcal{A} (in categorical slots).

QA Training. As illustrated in Figure 1, we prepend special prefixes to each input source. For instance, in multi-choice QA, “Multi-Choice Question:” is added to the question sequence; and “Choices:” is added to the answer candidates sequence, where each candidate is separated by a special token “[sep]”. All the input sources are concatenated into a single sequence as input to a sequence-to-sequence (Seq2Seq) model. Then, the model generates the correct answer a_i token by token.

$$a_i = \text{Seq2Seq}([q_i, \mathcal{A}, \mathcal{C}]). \quad (1)$$

It is worth noting that some of the questions q_i are unanswerable given the context. In these cases, $a_i = \text{none}$.

The training objective of our QA model is minimizing the negative log-likelihood of a_i given q_i , \mathcal{A} and \mathcal{C} , that is

$$\mathcal{L} = -\log p(a_i | q_i, \mathcal{A}, \mathcal{C}). \quad (2)$$

We initialize the model parameters with T5 (Rafael et al., 2020), an encoder-decoder Transformer with relative position embeddings (Shaw et al., 2018). The model is pre-trained on 750GB of clean and natural English text with a masking language modeling objective (masking out 15% of input spans, then predicting the missing spans using the decoder).

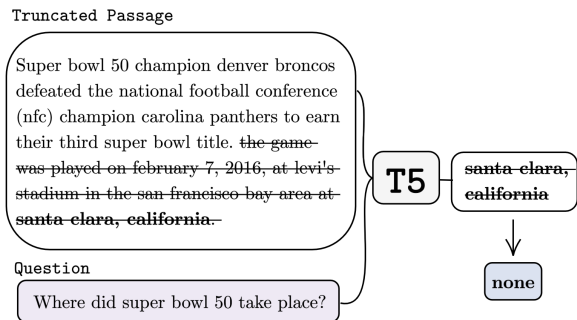


Figure 3: Context truncation strategy for generating none values. We truncate the passage to make sure the answer span is not present in the context and thus the QA model (T5) learns to predict “none”.

DST Zero-Shot. In DST, we consider tracking a slot value as finding the answer to a slot question from a dialogue history. Therefore, we first formulate all the slots as natural language questions, with each question roughly following the format “what is the $\langle \text{slot} \rangle$ of the $\langle \text{domain} \rangle$ that user wants?”. The context is a dialogue history which consists of an alternating set of utterances from two speakers, $\mathcal{C} = \{U_1, S_1, \dots, S_{t-1}, U_t\}$. “user:” and “system:” prefixes are added to the user and system utterance, respectively. Following the QA training phase, “Multi-Choice Question:” and “Extractive Question:” prefixes are added to the categorical and non-categorical slot question sequence. Then, the slot question q_i , value candidates \mathcal{A} , and dialogue context \mathcal{C} are concatenated into a single sequence as model input, and the model decodes the answer a_i with greedy decoding.

2.2 Unanswerable Questions

In DST, at any given turn of the conversation, the slots not mentioned by the user are marked with “none” in the dialogue state. Especially in multi-domain dialogues, there are typically two kinds of “none” value slots: *out-of-domain* and *in-domain unmentioned*. The *out-of-domain* slots are the slots in other domains that are irrelevant to the current conversation. For example, when the user asks for a hotel in the center, all the slots that are not in the hotel domain (e.g., *restaurant-price*) have the value “none”. The second, *in-domain unmentioned*, are those slots in the domain of interest but not yet mentioned by the user. For example, the user asks about a hotel in the center, and thus the slot *hotel-star* is “none” since the user does not specify this information. Therefore, we introduce two methods to simulate the *out-of-domain* and *in-domain*

unmentioned slots in the QA training phase.

Negative Question Sampling. The out-of-domain slots in DST are similar to out-of-context questions in QA, that is, the model must predict “none” when the question is irrelevant to the context. To construct this kind of unanswerable question, we adapt the negative sampling strategy (Mikolov et al., 2013). As illustrated in Figure 2, during QA training, we sample these negative questions from a pool of questions collected from other passages.

Context Truncation. The in-domain unmentioned slots often appear in the middle of conversations, where some of the in-domain slots have not yet mentioned by the user. We simulate such scenario by truncating the context passage from the first sentence that contains the answer span. As illustrated in Figure 3, given a question and a passage from a QA training set, we first truncate the passage according to the answer span annotation, then we pair the question and the truncated passage as an unanswerable sample.

3 Experiments

3.1 Datasets

QA datasets. For the QA training, we use six extractive QA datasets such as SQuAD2.0 (Rajpurkar et al., 2018)¹, NewsQA (Trischler et al., 2017), TriviaQA (Joshi et al., 2017), SearchQA (Dunn et al., 2017), HotpotQA (Yang et al., 2018), Natural Question (Kwiatkowski et al., 2019) from MRQA-2019 (Fisch et al., 2019), and two multiple-choice datasets such as RACE (Lai et al., 2017) and DREAM (Sun et al., 2019). The main train/dev statistics are reported in Table 1.

DST datasets. The evaluation is conducted on two multi-domain task-oriented dialogue benchmark, MultiWoz (Budzianowski et al., 2018; Eric et al., 2020) and Schema-Guided-Dialogue (SGD) (Rastogi et al., 2020). Both datasets provide turn-level annotations of dialogue states. In MultiWoz, we follow the pre-processing and evaluation setup from Wu et al. (2019), where restaurant, train, attraction, hotel, and taxi domains are used for training and testing. In SGD, the test set has 18 domains, and 5 of the domains are not presented in the training set.

¹Note that original MRQA-2019 dataset use SQuAD (Rajpurkar et al., 2016), here we also add the unanswerable questions from SQuAD2.0.

Dataset	Type	Train	Dev
SQuAD2.0	extractive	130,319	11,873
NewsQA	extractive	74,160	4,212
TriviaQA	extractive	61,688	7,785
SearchQA	extractive	117,384	16,980
HotpotQA	extractive	72,928	5,904
NaturalQA	extractive	104,071	12,836
RACE	multiple-choice	87,866	4,887
DREAM	multiple-choice	6,116	2,040

Table 1: Datasets used in the QA pre-training. Statistics of extractive datasets (except SQuAD2.0) are taken from MRQA-2019 (Fisch et al., 2019), and that of multiple-choice datasets are from RACE (Lai et al., 2017) and DREAM (Sun et al., 2019).

3.2 Evaluation

Joint Goal Accuracy (JGA) and Average Goal Accuracy (AGA) are used to evaluate our models and baselines. For JGA, the model outputs are only counted as correct when all of the predicted values exactly match the oracle values. AGA is the average accuracy of the active slots in each turn.

In order to make consistent comparisons to the previous works on cross-domain zero-shot/few-shot DST (Wu et al., 2019; Kumar et al., 2020; Zhou and Small, 2019) in MultiWoz, we compute JGA per domain as in Wu et al. (2019)². In SGD dataset, we use the official evaluation script³.

3.3 Implementation

We implement TransferQA based on T5-large (Raffel et al., 2020)⁴. All models are trained using the AdamW (Loshchilov and Hutter, 2018) optimizer with an initial learning rate of 0.00005. In the QA training stage, we set the ratio of generating an unanswerable question $\alpha = 0.3$, in which the ratio of negative sampled questions and truncated context is 0.95 : 0.05, and we train the models with batch size 1024 for 5 epochs.

In the DST zero-shot testing, we first treat all the slots as non-categorical and generate all the slot values. The slots that have no “none” values are considered as active slots. Then the model gener-

²<https://github.com/jasonwu0731/trade-dst>

³https://github.com/google-research/google-research/tree/master/schema_guided_dst

⁴Source code is available in <https://github.com/facebookresearch/Zero-Shot-DST>

Model	Joint Goal Accuracy					
	Attraction	Hotel	Restaurant	Taxi	Train	Average
TRADE [†] (Wu et al., 2019)	20.06	14.20	12.59	59.21	22.39	25.69
MA-DST [†] (Kumar et al., 2020)	22.46	16.28	13.56	59.27	22.76	26.87
SUMBT [‡] (Lee et al., 2019)	22.60	19.80	16.50	59.50	22.50	28.18
TransferQA (Ours)	31.25	22.72	26.28	61.87	36.72	35.77
<i>w/ Oracle Slot Gate</i>	<i>56.81</i>	<i>53.90</i>	<i>56.81</i>	<i>63.22</i>	<i>49.57</i>	<i>56.06</i>

Table 2: Zero-shot results on MultiWoz 2.1 (Eric et al., 2020). Results marked with [†] and [‡] are from Kumar et al. (2020) and Campagna et al. (2020). We also report the averaged zero shot joint goal accuracy among five domains. Note that this averaged per-domain accuracy is not comparable to the JGA in full shot setting.

	Domain	SGD-baseline		TransferQA	
		JGA	AGA	JGA	AGA
Unseen	Buses*	9.7	50.9	15.9	63.6
	Messaging*	10.2	20.0	13.3	37.9
	Payment*	11.5	34.8	24.7	60.7
	Trains*	13.6	63.5	17.4	64.9
	Alarm*	57.7	1.8	58.3	81.7
Seen	RentalCars	8.6	48.0	10.8	73.8
	Music	15.5	39.9	8.9	62.4
	RideSharing	17.0	50.2	31.2	61.7
	Media	18.0	30.8	30.2	67.5
	Homes	18.9	72.7	31.7	80.6
	Restaurants	22.8	55.8	16.3	68.9
	Events	23.5	57.9	15.6	56.8
	Flights	23.9	65.9	3.59	42.9
	Hotels	28.9	58.2	13.5	60.1
	Movies	37.8	68.6	24.0	56.2
	Services	40.9	72.1	37.2	75.6
	Travel	41.5	57.2	14.0	24.2
	Weather	62.0	76.4	40.3	59.4
All Domain	25.4	56.0	20.7	62.2	
<i>Oracle Slot Gate</i>	-	-	<i>48.0</i>	<i>76.6</i>	

Table 3: Zero-Shot results by domain in Schema Guided Dialogue (SGD) dataset (Rastogi et al., 2020). The SGD-baseline is trained with the whole training set, and the results are reported by Rastogi et al. (2020). Domains that appear in the test set but are not present in the training set are marked with *. For TransferQA, all the domains are unseen because the model is not trained with any DST data.

ates the value of active categorical slots by using a multi-choice QA formulation. In SGD, we follow the split of non-categorical and categorical slots in the dataset, while in MultiWoz, we follow the split of MultiWoz2.2 (Zang et al., 2020), except that all the number-type slots are considered as non-categorical slots. We also apply the canonicalization technique proposed by Gao et al. (2020) in

MultiWoz.

For the few-shot experiments, the QA pre-trained models are fine-tuned with 1%, 5% and 10% of the target domain data for 20 epochs. Other hyper-parameters are the same as in the QA training. We use 8 Tesla V100 GPUs for all of our experiments.

3.4 Baselines

TRADE. Transferable dialogue state generator (Wu et al., 2019), which utilizes a copy mechanism to facilitate domain knowledge transfer.

SUMBT. Slot-utterance matching belief tracker (Lee et al., 2019) based on the language model BERT (Devlin et al., 2018).

SGD-baseline. A schema-guided approach (Rastogi et al., 2020), which uses a single BERT (Devlin et al., 2018) model and schema descriptions to jointly predict the intent and dialogue state of unseen domains.

MA-DST. A multi-attention model (Kumar et al., 2020) which encodes the conversation history and slot semantics by using attention mechanisms at multiple granularities.

DSTQA. Dialogue state tracking via question answering over the ontology graph (Zhou and Small, 2019).

STARC. Applying two machine reading comprehension models based on RoBERTa-Large (Liu et al., 2019) for tracking categorical and non-categorical slots (Gao et al., 2020).

4 Results

4.1 Zero-Shot

In Table 2, three of the baselines, TRADE, MA-DST and SUMBT, are evaluated in the cross-

Model	Hotel			Resturant			Attraction			Train			Taxi		
	1%	5%	10%	1%	5%	10%	1%	5%	10%	1%	5%	10%	1%	5%	10%
TRADE	19.7	37.4	41.4	42.4	55.7	60.9	35.8	57.5	63.1	59.8	69.2	71.1	63.8	66.5	70.1
DSTQA	N/A	50.1	53.6	N/A	58.9	64.5	N/A	70.4	71.6	N/A	70.3	74.5	N/A	70.9	74.1
STARC	45.9	52.5	57.3	51.6	60.4	64.6	40.3	65.3	66.2	65.6	74.1	75.0	72.5	75.3	79.6
TransferQA	43.4	52.1	55.7	51.7	60.7	62.9	52.3	63.5	68.2	70.1	75.6	79.0	75.4	79.2	80.3

Table 4: Few-shot performance on MultiWoz 2.0 in terms of Joint Goal Accuracy (JGA). N/A for results not presented in the original paper. All models are evaluated with 1%, 5%, and 10% in-domain data.

domain setting, where the models are trained on the four domains in MultiWoz then zero-shot on the held-out domain. Our TransferQA, *without any DST training data*, achieves significantly higher JGA (7.59% on average) compared to the previous zero-shot results. Table 3 summarizes the results on SGD dataset, where the SGD-baseline (Rastogi et al., 2020) is trained with the whole SGD training set. TransferQA zero-shot performance is consistently higher in terms of JGA and AGA in the unseen domains, and competitive in seen domains. The results on both datasets shows the effectiveness of cross-task zero-shot transferring. In the cross-domain transfer scenario, despite the large amount of dialogue data, only a limited number of the slots appear in the source domain. For example, MultiWoz has 8,438 dialogues with 113,556 annotated turns, but only 30 different slots in 5 domains. Thus, cross-domain transferring requires the models generalize to new slots after being trained with fewer than 30 slots. By contrast, in cross-task transferring, each question in QA datasets can be considered as a slot. Therefore, a model which trained with diverse questions (around 500,000) on QA datasets is more likely to achieve better generalization.

4.2 Few-Shot

Table 4 shows the few-shot results on MultiWoz 2.0⁵, where TRADE (Wu et al., 2019) and DSTQA (Zhou and Small, 2019) are trained on four source domain on MultiWoz then finetuned with the target domain data, while STARC (Gao et al., 2020) and our model TransferQA are first trained on the same QA datasets then finetuned with the target domain data. We experiment with 1%, 5% and 10% of the target domain data. The results show that both cross-task transferring approaches (i.e., STARC and TransferQA) outperform cross-domain transferring approaches (i.e., TRADE and

DSTQA) in 4 out of 5 domains. Compared to STARC, TransferQA achieves around 1% lower JGA in the hotel domain, but consistently higher JGA on other domains under different data ratio settings. Especially when only 1% of in-domain data are available, our model outperforms STARC in most domains (except hotel) by a large margin (e.g., 11.95% in the attraction and 4.49% in the train domain). This significant improvement can be attributed to the generated unanswerable samples, which bridge the gap between the source data distribution and the target data distribution.

5 Analysis

5.1 Impact of Unanswerable Questions

In Table 5, we study the effect of the two unanswerable question generation strategies Context Truncation (CT) and Negative Question Sampling (NQS) described in Section 2.2. Applying both CT and NQS gives the best result in terms of average JGA for both TransferQA-large and TransferQA-base. While removing the CT strategy during the QA training only affects the performance in the train domain, removing both NQS and CT decreases the JGA dramatically in all the domains. This is due to the ratio of unanswerable (none) slots in MultiWoz is high (55.25%), and removing the simulated unanswerable questions during QA training affects the Slot Gate Accuracy (SGA) in DST inference. Indeed, by adding NQS and CT, we observed large JGA improvement (around 30%) in the taxi domain which has highest unanswerable slots ratio (71.85%), and relatively small JGA improvement (around 10%) in attraction and train domains where the ratios of unanswerable slots are 47.70% and 49.58%. Overall, these results demonstrate the importance of generating unanswerable questions.

In Figure 4, we show the effect of using different ratios α for generating unanswerable questions, while when it is too low, the model is not able to capture the unmentioned slots; when the ratio

⁵Few shot experiments are conducted on MultiWoz 2.0 for comparing with previous works.

Model	Joint Goal Accuracy						Slot Gate Accuracy					
	Attraction	Hotel	Restaurant	Taxi	Train	Average	Attraction	Hotel	Restaurant	Taxi	Train	ALL
TransferQA-base	28.48	22.75	20.92	61.16	31.15	32.89	60.10	78.41	76.36	86.06	85.15	78.22
- CT	29.51	21.66	23.37	58.90	24.13	31.51	64.81	78.63	78.83	83.98	80.11	78.02
- CT - NQS	23.98	15.54	18.16	27.03	12.48	19.44	65.84	75.27	75.86	71.70	74.00	73.94
TransferQA-large	31.25	22.72	26.28	61.87	36.72	35.77	60.62	77.84	81.73	86.48	87.21	79.95
- CT	32.47	22.69	27.71	62.96	32.17	35.60	66.99	79.56	82.72	88.88	86.79	81.48
- CT - NQS	24.69	16.22	23.01	31.54	23.05	23.70	69.34	74.82	80.04	78.87	83.45	77.95

Table 5: Ablation study on the effectiveness of two unanswerable question generation strategies: Context Truncation (CT) and Negative Question Sampling (NQS). The experiments are conducted on MultiWoz 2.1 with different model size. Slot Gate Accuracy measures how well the model can classify unanswerable slots.

of unanswerable questions is too high, the model tends to over-predict “none”. In general, we find that $\alpha = 0.3$ and $\alpha = 0.6$ gives the highest JGA.

5.2 Error Analysis

To understand the current limitation of cross-task transfer learning, we conducted an error analysis on the results of MultiWoz 2.1 zero-shot. We found that 79.79% of the errors come from the slot gate prediction (i.e., whether the slot is unanswerable or answerable), of which 37.54% are false positive errors (i.e., the slot is unanswerable and the model predict answerable), 42.25% are false negative errors (i.e., the slot is answerable and the model predicts unanswerable), and only 20.21% of errors come from wrong value predictions of answerable slots. In Table 6, we show two typical errors that we found in the zero-shot DST setting. The first, as shown in the example of dialogue *MUL2321*, is the model predicting slot values that have not been confirmed by the user yet (e.g., pricerange="expensive" etc.). The second error, as shown in dialogue *PMUL0089*, is the model not capturing slot values when the user does not explicitly mention the domain (e.g., a place to stay refers to the hotel domain). These errors occurred because of question-context mismatching, and they might be addressed with well designed or leaned slot questions (Li and Liang, 2021; Wallace et al., 2019). We leave this exploration to the future work.

5.3 Oracle Study

We further conducted an oracle study on our model by providing the gold slot gate information. The results are shown in the last row of Table 3 and Table 2. We found that this oracle information dramatically increases the JGA (20.7% \rightarrow 48.0% in SGD, 35.77% \rightarrow 56.06% in MultiWoz). Therefore, by improving the accuracy of predicting “none” value slots, we have the potential to increase the overall




Figure 4: Joint Goal Accuracy (JGA) w.r.t. the probability of generating unanswerable questions α . Highlighted a region of ratio where the model achieves the highest JGA.

zero-shot DST performance by a large margin.

6 Related Work

Machine Reading for Question Answering (MRQA) is an important task for evaluating how well computer systems understand human language (Fisch et al., 2019). In MRQA, a model must answer a question by reading one or more context documents. There are mainly two types of MRQA tasks. The first is extractive QA (Rajpurkar et al., 2016, 2018; Trischler et al., 2017; Joshi et al., 2017; Dunn et al., 2017; Yang et al., 2018; Kwiatkowski et al., 2019), where the answer to each answerable question appears as a span of tokens in the passage. A popular approach for this task is to predict the start token and end token of the answer span (Devlin et al., 2019). The second is multi-choice QA (Lai et al., 2017; Sun et al., 2019), where the answer candidates are provided. In this task, classification-base models are usually applied to predict the correct candidate.

Dialogue History (MUL2321)			
S: yes I can. what restaurant are you looking for?			
U: It is called maharajah tandoori restaurant .			
S: I've located the maharajah tandoori restaurant for you. It serves indian food, it's in the west area and is in the expensive price range.			
U: Can you book a table for 7 people at 12:30 on tuesday ?			
Slots	Questions	Gold Values	Predicted Values
restaurant-book day	what is the day for the restaurant reservation?	tuesday	tuesday
restaurant-book people	how many people for the restaurant reservation?	7	7
restaurant-book time	what is the book time of the restaurant that the user is interested in?	12:30	12:30
restaurant-name	what is the name of the restaurant that the user is interested in?	maharajah tandoori	maharajah tandoori
restaurant-pricerange	what is the price range of the restaurant that the user is interested in?	none	expensive
restaurant-area	what is the area of the restaurant that the user is interested in?	none	west
restaurant-food	what kind of food does user want to eat in restaurant?	none	indian
Dialogue History (PMUL0089)			
U: Can you help me find a cheap place to stay in the east part of town?			
Slots	Questions	Gold Values	Predicted Values
hotel-area	what is the area of the hotel that the user wants?	east	none
hotel-pricerange	what is the price range of the hotel or guesthouse that the user wants?	cheap	none

Table 6: Two typical errors of TransferQA zeroshot in MultiWoz 2.1. The first (top example) is predicting the values that not confirmed by the user yet, and the second (bottom example) is missing the values of implicit mentioned domain.

Dialogue State Tracking is an essential yet challenging task in conversational AI research (Williams and Young, 2007; Williams et al., 2014). Recent state-of-the-art models (Lei et al., 2018; Zhang et al., 2020; Wu et al., 2020; Peng et al., 2020; Zhang et al., 2019; Kim et al., 2019; Lin et al., 2020; Chen et al., 2020; Heck et al., 2020; Mehri et al., 2020; Hosseini-Asl et al., 2020; Yu et al., 2020; Li et al., 2020; Madotto et al., 2020) trained with extensive annotated dialogue data have shown promising performance in complex multi-domain conversations (Budzianowski et al., 2018; Eric et al., 2020). However, collecting large amounts of data for every dialogue domain is often costly and inefficient. To reduce the expense of data acquisition, zero-shot (few-shot) transfer learning has been proposed as an effective solution. Wu et al. (2019) adapt a copy mechanism for transferring prior knowledge of existing domains to new ones, while Zhou and Small (2019) use the ontology graph to facilitate domain knowledge transfer. Campagna et al. (2020) leverage the ontology and in-domain templates to generate a large amount of synthesized data for domain adaptation, and Rastogi et al. (2020) apply schema descriptions for tracking out-of-domain slots. Despite the effectiveness of these approaches, a considerable amount of DST data are still required to cover a broad range of slot categories (Gao et al., 2020).

On the other hand, Gao et al. (2020) propose

to utilize abundant QA data to overcome the data scarcity issue in DST tasks. The authors first train a classification model and a span-extraction model on multi-choice QA and extractive QA datasets independently. Then, they use the two QA models to track categorical and extractive slots. Compared to this approach, our method is fundamentally different in two aspects: 1) our model can effectively handle “none” value slots (e.g., unmentioned and out-of-domain slots) in the zero-shot setting, which is important to DST performance as there are many “none” slots in multi-domain dialogues; 2) our method provides a simple text-to-text input-output interface for tracking both categorical and extractive slots with a single generative model.

7 Conclusion

In this paper, we present TransferQA, a unified generative model that performs DST *without using any DST training data*. TransferQA uses the text-to-text transfer learning framework that seamlessly combines extractive QA and multi-choice QA for tracking both categorical slots and non-categorical slots. To enable our model to zero-shot “none” value slots, we introduce two effective ways to construct unanswerable questions, i.e., negative question sampling and context truncation. The experimental results on the MultiWoz and SGD datasets demonstrate the effectiveness of our approach in both zero-shot and few-shot settings. We also show that improving the “none” value slot accuracy has

the potential to increase the overall zero-shot DST performance by a large margin, which can be explored in future work.

References

- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gasic. 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.
- Giovanni Campagna, Agata Foryciarz, Mehrad Moradshahi, and Monica S Lam. 2020. Zero-shot transfer learning with synthesized data for multi-domain dialogue state tracking. *arXiv preprint arXiv:2005.00891*.
- Lu Chen, Boer Lv, Chunxin Wang, Su Zhu, Bowen Tan, and Kai Yu. 2020. Schema-guided multi-domain dialogue state tracking with graph attention neural networks. In *AAAI 2020*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT (1)*.
- Matthew Dunn, Levent Sagun, Mike Higgins, V Ugur Guney, Volkan Cirik, and Kyunghyun Cho. 2017. Searchqa: A new q&a dataset augmented with context from a search engine. *arXiv preprint arXiv:1704.05179*.
- Mihail Eric, Rahul Goel, Shachi Paul, Abhishek Sethi, Sanchit Agarwal, Shuyang Gao, Adarsh Kumar, Anuj Goyal, Peter Ku, and Dilek Hakkani-Tur. 2020. Multiwoz 2.1: A consolidated multi-domain dialogue dataset with state corrections and state tracking baselines. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 422–428.
- Adam Fisch, Alon Talmor, Robin Jia, Minjoon Seo, Eunsol Choi, and Danqi Chen. 2019. Mrqa 2019 shared task: Evaluating generalization in reading comprehension. In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, pages 1–13.
- Shuyang Gao, Sanchit Agarwal, Tagyoung Chung, Di Jin, and Dilek Hakkani-Tur. 2020. From machine reading comprehension to dialogue state tracking: Bridging the gap. *arXiv preprint arXiv:2004.05827*.
- Michael Heck, Carel van Niekerk, Nurul Lubis, Christian Geishausser, Hsien-Chin Lin, Marco Moresi, and Milica Gašić. 2020. Trippy: A triple copy strategy for value independent neural dialog state tracking. *arXiv preprint arXiv:2005.02877*.
- Ehsan Hosseini-Asl, Bryan McCann, Chien-Sheng Wu, Semih Yavuz, and Richard Socher. 2020. A simple language model for task-oriented dialogue. *arXiv preprint arXiv:2005.00796*.
- Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611.
- John F Kelley. 1984. An iterative design methodology for user-friendly natural language office information applications. *ACM Transactions on Information Systems (TOIS)*, 2(1):26–41.
- Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hananeh Hajishirzi. 2020. Unifiedqa: Crossing format boundaries with a single qa system. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 1896–1907.
- Sungdong Kim, Sohee Yang, Gyuwan Kim, and Sang-woo Lee. 2019. Efficient dialogue state tracking by selectively overwriting memory. *arXiv preprint arXiv:1911.03906*.
- Adarsh Kumar, Peter Ku, Anuj Kumar Goyal, Angeliki Metallinou, and Dilek Hakkani-Tur. 2020. Ma-dst: Multi-attention based scalable dialog state tracking. *arXiv preprint arXiv:2002.08898*.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. Race: Large-scale reading comprehension dataset from examinations. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 785–794.
- Hwaran Lee, Jinsik Lee, and Tae-Yoon Kim. 2019. Sumbt: Slot-utterance matching for universal and scalable belief tracking. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5478–5483.
- 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.

- Shiyang Li, Semih Yavuz, Kazuma Hashimoto, Jia Li, Tong Niu, Nazneen Rajani, Xifeng Yan, Yingbo Zhou, and Caiming Xiong. 2020. Coco: Controllable counterfactuals for evaluating dialogue state trackers. *arXiv preprint arXiv:2010.12850*.
- Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*.
- Zhaojiang Lin, Bing Liu, Seungwhan Moon, Paul A Crook, Zhenpeng Zhou, Zhiguang Wang, Zhou Yu, Andrea Madotto, Eunjoon Cho, and Rajen Subba. 2021a. Leveraging slot descriptions for zero-shot cross-domain dialogue statetracking. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5640–5648.
- Zhaojiang Lin, Andrea Madotto, Genta Indra Winata, and Pascale Fung. 2020. 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.
- Zhaojiang Lin, Andrea Madotto, Genta Indra Winata, Peng Xu, Feijun Jiang, Yuxiang Hu, Chen Shi, and Pascale Fung. 2021b. Bitod: A bilingual multi-domain dataset for task-oriented dialogue modeling. *arXiv preprint arXiv:2106.02787*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Ilya Loshchilov and Frank Hutter. 2018. Decoupled weight decay regularization. In *International Conference on Learning Representations*.
- Andrea Madotto, Zhaojiang Lin, Zhenpeng Zhou, Seungwhan Moon, Paul Crook, Bing Liu, Zhou Yu, Eunjoon Cho, and Zhiguang Wang. 2020. Continual learning in task-oriented dialogue systems. *arXiv preprint arXiv:2012.15504*.
- Shikib Mehri, Mihail Eric, and Dilek Hakkani-Tur. 2020. Dialoglue: A natural language understanding benchmark for task-oriented dialogue. *arXiv preprint arXiv:2009.13570*.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Seungwhan Moon, Satwik Kottur, Paul A Crook, Ankita De, Shivani Poddar, Theodore Levin, David Whitney, Daniel Difranco, Ahmad Beirami, Eunjoon Cho, et al. 2020. Situated and interactive multimodal conversations. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1103–1121.
- Baolin Peng, Chunyuan Li, Jinchao Li, Shahin Shayan-deh, Lars Liden, and Jianfeng Gao. 2020. Soloist: Few-shot task-oriented dialog with a single pre-trained auto-regressive model. *arXiv preprint arXiv:2005.05298*.
- 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, Robin Jia, and Percy Liang. 2018. Know what you don’t know: Unanswerable questions for squad. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789.
- 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.
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2020. Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8689–8696.
- 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*.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. Self-attention with relative position representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 464–468.
- Kai Sun, Dian Yu, Jianshu Chen, Dong Yu, Yejin Choi, and Claire Cardie. 2019. Dream: A challenge data set and models for dialogue-based reading comprehension. *Transactions of the Association for Computational Linguistics*, 7:217–231.
- Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordani, Philip Bachman, and Kaheer Suleman. 2017. Newsqa: A machine comprehension dataset. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 191–200.
- Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. 2019. Universal adversarial triggers for attacking and analyzing nlp. 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 2153–2162.

- Tsung-Hsien Wen, David Vandyke, Nikola Mrkšić, Milica Gasic, Lina M Rojas Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2017. A network-based end-to-end trainable task-oriented dialogue system. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 438–449.
- Jason D Williams, Matthew Henderson, Antoine Raux, Blaise Thomson, Alan Black, and Deepak Ramachandran. 2014. The dialog state tracking challenge series. *AI Magazine*, 35(4):121–124.
- Jason D Williams and Steve Young. 2007. Partially observable markov decision processes for spoken dialog systems. *Computer Speech & Language*, 21(2):393–422.
- Chien-Sheng Wu, Steven CH Hoi, Richard Socher, and Caiming Xiong. 2020. Tod-bert: Pre-trained natural language understanding for task-oriented dialogue. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 917–929.
- Chien-Sheng Wu, Andrea Madotto, Ehsan Hosseini-Asl, Caiming Xiong, Richard Socher, and Pascale Fung. 2019. Transferable multi-domain state generator for task-oriented dialogue systems. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 808–819.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380.
- Tao Yu, Rui Zhang, Alex Polozov, Christopher Meek, and Ahmed Hassan Awadallah. 2020. Score: Pre-training for context representation in conversational semantic parsing. In *International Conference on Learning Representations*.
- Xiaoxue Zang, Abhinav Rastogi, 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*, pages 109–117.
- Jian-Guo Zhang, Kazuma Hashimoto, Chien-Sheng Wu, Yao Wan, Philip S Yu, Richard Socher, and Caiming Xiong. 2019. Find or classify? dual strategy for slot-value predictions on multi-domain dialog state tracking. *arXiv preprint arXiv:1910.03544*.
- Yichi Zhang, Zhijian Ou, and Zhou Yu. 2020. Task-oriented dialog systems that consider multiple appropriate responses under the same context. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9604–9611.
- Li Zhou and Kevin Small. 2019. Multi-domain dialogue state tracking as dynamic knowledge graph enhanced question answering. *arXiv preprint arXiv:1911.06192*.