

Joint Reasoning on Hybrid-knowledge sources for Task-Oriented Dialog

Mayank Mishra

IBM Research

mayank.mishra1@ibm.com

Danish Contractor

IBM Research

danish.contractor@ibm.com

Dinesh Raghu

IBM Research

diraghu1@in.ibm.com

Abstract

Traditional systems designed for task oriented dialog utilize knowledge present only in structured knowledge sources to generate responses. However, relevant information required to generate responses may also reside in unstructured sources, such as documents. Recent state of the art models such as HyKnow (Gao et al., 2021b) and SEKNOW (Gao et al., 2021a) aimed at overcoming these challenges make limiting assumptions about the knowledge sources. For instance, these systems assume that certain types of information, such as a phone number, is *always* present in a structured knowledge base (KB) while information about aspects such as entrance ticket prices, would always be available in documents.

In this paper, we create a modified version of the MutliWOZ-based dataset prepared by (Gao et al., 2021a) to demonstrate how current methods have significant degradation in performance when strict assumptions about the source of information are removed. Then, in line with recent work exploiting pre-trained language models, we fine-tune a BART (Lewis et al., 2020) based model using prompts (Brown et al., 2020; Sun et al., 2021) for the tasks of querying knowledge sources, as well as, for response generation, without making assumptions about the information present in each knowledge source. Through a series of experiments, we demonstrate that our model is robust to perturbations to knowledge modality (source of information), and that it can fuse information from structured as well as unstructured knowledge to generate responses.

1 Introduction

Most existing work on task-oriented dialog systems assumes that the knowledge required for completing a task (eg: booking a restaurant reservation), resides in structured knowledge sources. Thus, typical task-oriented dialog systems require generating a *belief state*, that can be used to query a knowledge

base to fetch entity results; these results are then used to generate responses. Recognizing that information is not always present in structured resources, recently methods that can additionally use unstructured knowledge (eg: document collections), have also been developed (Kim et al., 2020; Gao et al., 2021a). However, current state-of-the-art models designed for such tasks make limiting assumptions about the nature of knowledge sources, that make them unsuitable for use in real-world settings.

Limitations of existing methods: First, current task-oriented dialog systems designed to reason over hybrid knowledge sources assume that a knowledge base and the unstructured knowledge source encode separate pieces of information about entities (eg: the zip-code is always in structured knowledge, ticket prices are always available in unstructured text) (Kim et al., 2020; Zhang et al., 2021). This is not reflective of real-world knowledge, where independent information systems are often fused to enable applications.

Second, existing systems are trained to learn the source of different pieces of information, thus, making them unsuitable for situations where any field that was previously in a structured knowledge source is now available in an unstructured knowledge source (and vice versa). In effect, a simple change in the modality of information can result in a failure of the model to utilize the information present in knowledge, as existing models memorize the source of every piece of information.

Third, such systems assume that each knowledge grounded response can contain information from only *one* source type (Kim et al., 2020; Gao et al., 2021a; Zhang et al., 2021) – either structured or unstructured knowledge. This is an artificial constraint imposed to make modelling easier, but real-world conversations can routinely require systems to fuse information from more than one knowledge type (eg: See Dialog turn 4 in Figure 1).

Contributions: In this paper, we present our work

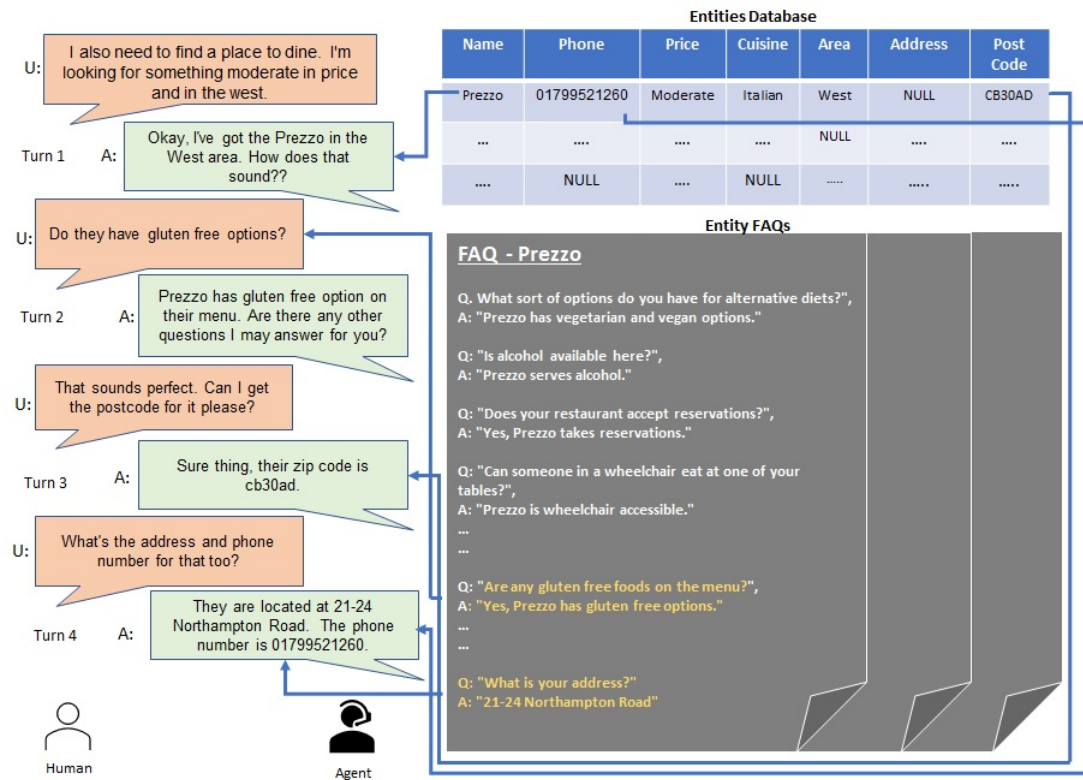


Figure 1: Example of a dialog requiring the use of data from two different sources. Agent Turn 4, requires incorporating information from both structured (DB Table) and unstructured data (a document consisting of FAQs for the entity).

aimed at removing each of these strict assumptions from task-oriented dialog systems. Current methods for joint-reasoning in task oriented dialogs have been developed using an augmented version of MultiWOZ 2.1 which contains additional dialog turns based on new unstructured information (Gao et al., 2021a). Unfortunately, no attempt has been made to distribute information across knowledge sources. We therefore create a modified version of this dataset (called HYBRIDTOD) that optimally redistributes information across structured and unstructured knowledge so that most dialogs in the train dataset are affected by this change.

A trivial method of redistributing information across structured and unstructured knowledge sources would be to arbitrarily move structured fields for some entities to the unstructured knowledge source. However, since the universe of entities in the dataset is very large and not all entities are directly referred to in the dialogs, such a method of redistributing information may not be as effective if the dialogs do not use the slot-values that have been redistributed. We therefore, develop an automated graph based approach which uses the max-cut of the graph to optimally redistribute infor-

mation from structured to unstructured knowledge sources.

Lastly, in line with recent work exploiting pre-trained language models, we fine-tune BART (Lewis et al., 2020) using prompts for the tasks of querying knowledge as well as response generation without making assumptions about the information present in each knowledge source. Specifically, we do Prompt+LM finetuning (Liu et al., 2021a) in which both the prompt and model parameters are trainable (Ben-David et al., 2021; Liu et al., 2021b; Han et al., 2021). Through a series of experiments, we demonstrate that our model is robust to perturbations to knowledge modality (source of information), and it can fuse information from structured as well as unstructured knowledge to generate responses.

In summary we make the following contributions¹: (1) We prepare a new version of the MultiWOZ-DSTC9 combined dataset (Kim et al., 2020; Gao et al., 2021a) called HYBRIDTOD to study the reasoning on hybrid knowledge sources for task oriented dialog systems. (2) We demon-

¹The constructed dataset and code used is available at <https://github.com/mayank31398/HybridToD>

Slot Type	Slot Values	Question Template	Answer Template
price	cheap	What is the price range?	It has \${price} pricing.
	expensive	How costly is \${restaurant name}?	\${restaurant name} is \${price}
cuisine	Italian	What is the cuisine?	\${restaurant name} caters for \${cuisine} cuisine.
	Thai	What type of food is served here?	You can find \${cuisine} food here

Table 1: Examples of templates used for moving slot values from the structured to the unstructured knowledge source.

strate that our model (referred to as JOINTLM) is also able to fuse information from both knowledge modalities and beats existing state-of-the-art systems on standardized metrics. (3) We present detailed ablation studies demonstrating the value of our modelling choices.

2 Related Work

Modeling Task Oriented Dialogs: Multiple flavours of this problem have been defined to address different aspects of modeling - eg: belief state tracking to assess whether a model is able to correctly decode the query needed given a current conversational context (Dey and Desarkar, 2021; Li et al., 2021; Yang et al., 2021), generating responses given belief states to assess whether a model is able to correctly predict the knowledge attributes to be used in a response (Yang et al., 2021; Chen et al., 2019; Gao et al., 2020; Mohapatra et al., 2021), end-to-end modeling of dialog systems where models are assessed on the correctness of the response generated including the values used from the knowledge base (Bordes et al., 2017; Raghu et al., 2021b), etc. Recent work that assumes that belief state annotations are latent and not available for training have also been developed (Raghu et al., 2021a).

Knowledge Grounded Dialog: Dialog systems that generate responses on information grounded in external knowledge have also been developed. Unlike, work on task oriented dialogs, which primarily focuses on using structured knowledge to complete a ‘goal’ or accomplish a ‘task’ (eg: POI recommendation for in car navigation (Eric et al., 2017), restaurant, hotel or flight booking (El Asri et al., 2017), etc), most existing knowledge grounded systems are designed to address informational needs of users (eg: answering queries based on collections of documents, making response recommendations to contact center agents). Finally, contemporaneous to our work, knowledge grounded response generation tasks that combine information from hybrid knowledge sources have also been proposed

(Nakamura et al., 2022). Here, unlike task oriented dialog systems, which require the retrieval of an entity to make recommendations or accomplish a task, in such tasks, the goal is to answer an informational seeking query in a chat-conversation. Models are required to use the dialog context to fetch related tables (often flattened and encoded as independent table cells), along with documents to generate a response.

3 The HYBRIDTOD Dataset

The dataset prepared by (Gao et al., 2021b) (referred to as the SEKNOW-MULTIWOZ dataset in this paper) is the only publicly available task-oriented dialog dataset in which the dialogs are grounded on two types of knowledge sources: structured and unstructured (FAQs). However, SEKNOW-MULTIWOZ is not indicative of a real-world setting due to two major limitations: (1) It has a strict, slot-type to knowledge-source type mapping. For example, the slot-type ‘cuisine’ is always in the structured source while ‘timings’ of operation would always be mentioned in unstructured documents, and (2) an agent response contains information from only one source (i.e., either from structured or unstructured). To alleviate these limitations, we systematically modify the knowledge sources in SEKNOW-MULTIWOZ to construct a new dataset that we refer to as HYBRIDTOD.

Dataset Construction: We first create an undirected graph $G = (V, E)$ where each vertex $v \in V$ is a unique slot-value and an edge $e \in E$ exists between two vertices, if the slot values represented by these vertices occur together in a training dialog utterance. For instance, in Figure 1 nodes associated with slot-values “21-24 Northampton Road” and phone number “01799521660” would have an edge between them due to Turn 4. Similarly, vertices corresponding to the values for slot-type ‘cuisine’ *Italian* and the slot-type ‘address’ *21-24 Northampton Road* would have had an edge between them if the utterance at Turn 4 was instead, “It is an *Italian restaurant located at 21-24 Northampton*

Domain	Context-Response pairs			Number of entities
	train	validation	test	train/validation/test
hotel	19370	2316	2295	33
restaurant	19716	2162	2188	110
attraction	8192	1226	1246	79
total	47278	5704	5729	222

Table 2: Number of context-response pairs in the dataset

Domain	SEKNOW-MULTIWOZ	HYBRIDTOD
hotel	10.97	6.79
restaurant	8.12	5.25
attraction	9	6.38

Table 3: Average number of slot values by domain in the structured knowledge source for each dataset.

Road". Our goal is to move some of the slot-values (vertices) in G that are originally in the structured knowledge source to the unstructured knowledge source so as to alleviate some of the limitations of the original dataset.

In order to identify which vertices to move, we create a maxcut of the graph G using the Max-CutBM algorithm (Boumal et al., 2016). A maxcut results in a graph in which the most number of edges from the original graph are ‘cut’. After the application of MaxCut, all slot values in one graph partition are retained in the structured knowledge source, while the others are converted to text QA pairs using templates and included as part of the unstructured document associated with that entity. The templates for the restaurant slot types ‘price’ and ‘cuisine’ are shown in Table 1 for illustration. Since the edges of the graph are based on slot-value mentions in dialog utterances, applying a maxcut modifies the knowledge source in a way that it affects most dialog turns in the dataset; in other words, the max-cut ensures the maximum possible utterances in the dataset have information fused from both knowledge sources.

Since we move slot values from one partition of the graph to the unstructured knowledge source, a slot type can now have some values in structured knowledge and some in unstructured knowledge (as an FAQ). We find that our approach ends up modifying each entity referred to in the dataset, and that slot-values of the same type are now distributed across different types of knowledge.

Domain	SEKNOW-MULTIWOZ	HYBRIDTOD	UNSTRUCTURED TOD
hotel	36.52	40.58	46.48
restaurant	14.96	17.83	22.7
attraction	0	2.62	8

Table 4: Average number of FAQs for each domain in the unstructured knowledge source.

For experimentation, we also create a version of the dataset with all slot values² moved from the structured to the unstructured knowledge source. We refer to this dataset as UNSTRUCTURED TOD. To construct HYBRIDTOD and UNSTRUCTURED TOD dataset, we only consider dialogs from 3 domains: hotel, restaurant and attraction. We omit dialogs from other domains as they do not have associated knowledge. For example, the taxi domain only contains the information that the slot-type *phone* should match the regular expression $[\"^[0-9]{10}\$\"]$, but does not contain any instance of phone numbers present in the train dialogs.

Dataset Statistics: The number of context-response pairs (spread across the 3 domains: hotel, restaurant and attractions) for HYBRIDTOD are shown in Table 2. We also show the entity distribution by domain-type. The restaurant domain dominates the knowledge sources, occupying almost half of the total entities and the other half is constituted by hotel and attraction domains. Tables 3 and 4 show the distribution of entity slot-values in structured knowledge sources and FAQs in the unstructured knowledge source for each domain in the datasets. As can be seen, the average number of slot-values presented in structured knowledge are lesser in HYBRIDTOD as compared to SEKNOW-MULTIWOZ and correspondingly the number of FAQs in HYBRIDTOD are higher as compared to SEKNOW-MULTIWOZ. We present the detailed slot-type distribution of SEKNOW-MULTIWOZ and HYBRIDTOD in the appendix. We find that approximately 50% slot-values are moved to unstructured knowledge from the structured sources for each slot-type.

Limitations of the Dataset: Information about entities is only redistributed from the structured knowledge source to the unstructured knowledge source. In effect, information that was previously in unstructured knowledge sources continues to remain there. Redistributing information from unstructured documents to structured documents would require annotations to be able to extract facets to be converted to slot-types.

We describe our model, JOINTLM in the next section.

²The entity name is also a slot type but we always retain it in both knowledge sources.

4 JOINTLM

The problem of utilizing information and responding to users in task-oriented dialogs can be broken down into parts: (i) Querying Knowledge Source (structured and/or unstructured) to return entities (ii) Generating Responses (eg:sharing information about entities, requesting for more details from the user, etc).

We represent the dialog context as $c = (u_1, r_1, \dots, u_n)$, where (u_i, r_i) represent the user and the system response utterance at i^{th} turn respectively. We represent the entity e required for generating the response as the concatenation of its slot-values (from structured KB), represented as e^s , and FAQs from the unstructured knowledge source, represented as e^{us} :

$$\begin{aligned} [e^s] &= \langle struct \rangle \langle slot \rangle slot_1 \langle val \rangle value_1 \\ &\quad \langle slot \rangle slot_2 \langle val \rangle value_2 \dots \\ [e^{us}] &= \langle unstruct \rangle \langle doc \rangle document_1 \\ &\quad \langle doc \rangle document_2 \dots \\ [e] &= [e^s] [e^{us}] \end{aligned}$$

where $\langle struct \rangle$, $\langle unstruct \rangle$ are special tokens to demarcate the start of structured knowledge and unstructured knowledge of an entity respectively. $\langle slot \rangle$, $\langle val \rangle$ demarcate the slot-type and its value and $\langle doc \rangle$ denotes the start of a document from unstructured knowledge. We train JOINTLM to jointly model two tasks: entity retrieval and response generation. We use a hyperparameter α to weigh the two tasks during training, where α denotes the number of training samples used for entity retrieval task. Note that $\alpha = 0.5$ denotes equal number of examples for both the tasks.

4.1 Entity Retrieval

As discussed, prior to generating a response, we need to retrieve the relevant entity required to generate the response. We represent the inputs to the language model (LM) for this task as:

$$\begin{aligned} &\langle entity_retrieval_task \rangle \langle u \rangle u_1 \langle r \rangle r_1 \dots \\ &\langle u \rangle u_n \langle entity \rangle [e_j] \end{aligned}$$

where, $e_j \in \mathcal{E}$, the set of all entities, $\langle entity_retrieval_task \rangle$ and $\langle entity \rangle$ are special tokens for task prompting and demarcating the start of an entity. We train the model to generate the special tokens $z_j = \langle relevant \rangle$ or $z_j = \langle irrelevant \rangle$

for each entity e_j given the context c . We choose the best entity e as:

$$e = \underset{e_j}{\operatorname{argmax}} p(z_j = \langle relevant \rangle | c, e_j) \quad (1)$$

During training we use a subset of the entities in \mathcal{E} for creating the positive and negative set of entities. However, at inference time, we evaluate on all the entities in \mathcal{E} .

4.2 Response Generation

After scoring all entities, we use the context and the best entity e (the entity with the highest score for the $\langle relevant \rangle$ token) and generate response using the same LM. We represent the inputs for this task as:

$$\begin{aligned} &\langle response_task \rangle \langle u \rangle u_1 \langle r \rangle r_1 \dots \\ &\langle u \rangle u_n \langle entity \rangle [e] \end{aligned}$$

where $\langle response_task \rangle$ is a special token to prompt this task. We train the model to generate the response token-by-token.

4.3 Training details

We train our model to minimize $\sum_{(c,r)} \mathcal{L}(\theta, c, r)$, where

$$\begin{aligned} \mathcal{L}(\theta, c, r) &= -\alpha \log p_\theta(z_j | c, e_j) \\ &\quad - (1 - \alpha) \log p_\theta(r | c, e_j) \end{aligned}$$

The first term in the above objective represents the log-likelihood of retrieving the relevant entity and the second term is the log-likelihood of generating the response. Note that the term α (percentage of samples for each task) can be adjusted by changing the number of examples for the two tasks in a given batch of fixed size.

To train our model, we use early stopping with $patience = 5$ for the above objective on the validation set to prevent overfitting of our model. The loss was optimized using AdamW optimizer (Loshchilov and Hutter, 2017). We use a batch-size of 8 examples, with 4 examples for entity retrieval and 4 for response generation per batch. For the 4 examples for entity retrieval, 2 are positive and 2 are negative examples (effectively our batch is $2 + 2 + 4$). We use equation 1 during inference to pick the highest scored relevant entity.

Train Dataset	Test Dataset	Model	Bleu-1	Bleu-4	slot-values		
					prec.	recall	F1
HYBRIDTOD	SEKNOW-MULTIWoz	JOINTLM	30.63	8.60	50.48	45.37	47.79
		SEKNOW	29.20	7.83	43.16	28.65	33.14
HYBRIDTOD	HYBRIDTOD	JOINTLM	30.59	8.67	50.56	45.83	48.08
		SEKNOW	29.05	7.70	44.29	29.12	35.14
HYBRIDTOD	UNSTRUCTUREDToD	JOINTLM	30.30	8.44	51.05	45.37	48.04
		SEKNOW	27.43	6.68	42.96	19.62	27.11

Table 5: All models trained on HYBRIDTOD and evaluated on the rest of the datasets

Train Dataset	Test Dataset	Model	Bleu-1	Bleu-4	slot-values		
					prec.	recall	F1
SEKNOW-MULTIWoz	SEKNOW-MULTIWoz	JOINTLM	29.07	8.06	49.74	41.31	45.13
		SEKNOW	31.00	9.14	52.17	44.98	48.31
SEKNOW-MULTIWoz	HYBRIDTOD	JOINTLM	27.77	7.54	44.48	36.39	40.03
		SEKNOW	26.61	7.32	42.19	26.70	33.31
SEKNOW-MULTIWoz	UNSTRUCTUREDToD	JOINTLM	27.03	7.17	46.29	34.93	39.82
		SEKNOW	26.19	6.42	41.96	19.48	26.53

Table 6: All models trained on HYBRIDTOD and evaluated on the rest of the datasets

5 Experiments

Our experiments are aimed at answering the following questions: (1) How does JOINTLM perform compared to the baseline when trained and tested on HYBRIDTOD? (2) How does the change in slot-value distribution across structured and unstructured sources affect the performance of the models? (3) Is joint training of PromptLM for the two tasks of entity retrieval and response generation helpful? (4) How does JOINTLM compare with natural baselines for entity retrieval?

Experimental Setup: Task oriented dialog systems have to identify relevant entities (e.g. restaurants) from associated knowledge sources needed to generate a response. In order to identify these relevant entities, existing datasets provide the belief state annotations during training. Additionally, in our work for each dialog context, we associate a set of (positive) entities that exactly match the requirements present in the dialog context and a set of (negative) entities that do not match by an automated method. Note that the text snippets in the unstructured corpus do not have any annotations.

For all of our experiments, we use BART (Lewis et al., 2020) encoder-decoder based language model and finetune the pretrained model on the three datasets i.e, SEKNOW-MULTIWoz (Gao et al., 2021a), HYBRIDTOD and UNSTRUCTUREDToD datasets.

Baseline: We use the current state-of-the-art model for joint reasoning, SEKNOW (Gao et al., 2021a) model as our baseline. SEKNOW is designed to use belief state annotations – specifically, SEKNOW is trained to generate the belief state given the dialog context. These belief states are then used to query

the knowledge sources and generate a delexicalised response using the context and the generated belief state. The slot-values in the delexicalised response are then populated using an unordered set of entities returned by the belief state query on the structured knowledge source.

5.1 Evaluation Metrics

We report BLEU scores for assessing response generation performance and slot-value precision, recall and $F1$ for comparing the slot-value filling performance against the baseline. As described previously, since no new slot types were created from unstructured documents, the slot-value metrics are computed only using the slot-types that were originally present in the structured knowledge source.

We also report success@k for entity retrieval baselines to assess the performance of systems on the entity selection task. We define success@k as 1 if the top-k scored entities contain a relevant entity for response generation and 0 otherwise. However, note that it is not possible to measure success@k on SEKNOW since it generates the response using an unordered set of entities returned by the belief state query. We thus compare the two models only based on their performance on response generation.

5.2 Results

Knowledge-Source Memorization: We train and test both JOINTLM and the baseline model, SEKNOW on HYBRIDTOD and observe that JOINTLM outperforms SEKNOW by 13 points on slot-value F1 score (Row 1, Table 5). Also, the performance of SEKNOW drops from 48.31 (Row

Train Dataset	Test Dataset	Model	success@1	success@5	Bleu-1	Bleu-4	prec.	recall	F1
HYBRIDTOD	HYBRIDTOD	JOINTLM	84.50	86.57	30.59	8.67	50.56	45.83	48.08
		SEPLM	79.79	85.64	29.96	8.66	47.08	42.53	44.69
		TF-IDF	28.31	34.49	-	-	-	-	-

Table 7: Performance of models on the entity retrieval task.

1, Table 6) when trained/tested on the SEKNOW-MULTIWOZ dataset to 35.14 (Row 1, Table 5) when trained/tested on HYBRIDTOD dataset. This severe drop in performance is indicative of the fact that SEKNOW learns the source of slot-values and is unable to use information when the source of the particular slot-value can be varying (structured/unstructured) across entities.

Generalization of JOINTLM: To assess the generalization performance of the models, we train all the models on HYBRIDTOD and test on other datasets which have different slot-value distributions. As can be seen from Table 5, when trained on HYBRIDTOD, JOINTLM outperforms SEKNOW on all three dataset settings, SEKNOW-MULTIWOZ, HYBRIDTOD and UNSTRUCTURED-TOD across all response generation metrics. We also notice that JOINTLM trained on HYBRIDTOD is robust to change in the knowledge modality during inference (slot-value F1 stays at approx. 48). This is not the case for SEKNOW which exhibits large drop (31% from SEKNOW-MULTIWOZ to HYBRIDTOD and 45% from SEKNOW-MULTIWOZ to UNSTRUCTURED-TOD) in slot-value F1, as the distribution of slot-types changes in different datasets (Table 5).

We also train the models on SEKNOW-MULTIWOZ, and test on the other datasets and notice that JOINTLM outperforms SEKNOW on both HYBRIDTOD and UNSTRUCTURED-TOD (Table 6). However, SEKNOW has better slot-value F1 than JOINTLM on HYBRIDTOD. We hypothesize that this is because the belief state labels are more informative and provide a very strong signal for SEKNOW on SEKNOW-MULTIWOZ and this has the effect of SEKNOW learning the knowledge modality which is not the case for JOINTLM. This suggests that JOINTLM has better generalization performance.

5.3 Model Ablation Study

To study the importance of joint-training of our model, we also train a model without prompts using entity annotations, where two different BART (Lewis et al., 2020) models are trained for retrieval and generation. We call this model SEPLM. This

Dialog context

User: Suggest me some *Turkish* restaurants in Cambridge.

Agent: Yes there is a *Turkish* restaurant in Cambridge with two different locations, would you like the addresses?

User: I would like if its located in the *center* of the city price one located in the center of city and I don't care about the pricing.

Response

Ground truth response: I have two *Turkish* restaurants, both in the *center* and both *expensive*. May I recommend *Meze Bar*?

	SEKNOW	JOINTLM
SeKnow-MultiWOZ	I found a restaurant at <i>196 Mill Road City Centre</i> . Their phone number is <i>01223362372</i> .	<i>Meze Bar</i> is a <i>Turkish</i> restaurant in the <i>expensive</i>
HybridToD	I found a restaurant at <i>196 Mill Road City Centre</i> . Their phone number is {null}.	price range. Their address is <i>196 Mill Road City</i>
UnstructuredToD	I found a restaurant at {null}. Their phone number is {null}.	<i>Centre</i> .

Table 8: Response generated by SEKNOW and JOINTLM on the three datasets.

model is trained on HYBRIDTOD and is compared against JOINTLM on both entity retrieval and response generation (Table 7). We observe that JOINTLM outperforms SEPLM in both the tasks with a 5 points difference in success@1 and a 3 points difference in slot-value F1. This confirms that the joint modeling of the 2 tasks using prompting yields a better model than learning a separate model for the 2 tasks at hand. For a detailed evaluation on all other dataset combinations, please refer to the Appendix. For comparison with a non-neural entity retrieval baseline, we also report the success scores BM25 based TF-IDF retriever which are significantly worse than the neural retrievers used for JOINTLM and SEPLM. These experiments highlight the benefit of joint modeling of the two tasks.

5.4 Qualitative Study

In Table 8, we show the responses generated for a sample dialog by JOINTLM and SEKNOW on the three datasets used for our experiments. It should be noted that JOINTLM generates the same response for all the three datasets. However, SEKNOW is not able to populate the required slot-values for this entity (*Meze Bar*) in the response in HYBRIDTOD and UNSTRUCTURED-TOD when those slot-values are no longer available in the structured source.

6 Conclusion

In this paper we presented a new dataset, HYBRID-TOD that requires reasoning over both structured and unstructured knowledge sources to generate responses to dialogs. Unlike existing task-oriented dialog datasets, it does not restrict slot-types to specific knowledge sources. Through our experiments we demonstrated how existing methods do not adapt well to changing distributions of slot-type sources and that our model JOINTLM (trained using entity annotations rather than belief state), not only generates better responses by reasoning over both knowledge sources, it also learns a better retriever for entities. In future work, we also plan to train our models without using any annotations i.e without any supervision on entity label information.

7 Limitations

Our dataset and model are not intended to be directly used in a real-world system as they have some inherent limitations. As mentioned in Section 3, we only redistribute slot types from structured knowledge sources to unstructured knowledge sources. Due to a lack of resources we are unable to annotate unstructured documents – our dataset has a bias that certain information will always appear in unstructured information. In addition, we rely on a pre-trained language model, BART, to generate responses. We have not assessed to what extent the generated responses could exhibit any form of social bias or toxic language (when prompted). We do not recommend that our system be used in a real-world deployed chatbot without further study. Lastly, this work has been assessed only on English language data using a pretrained language model developed for English.

8 Acknowledgement

The authors would like to thank Raunak Sinha for his contributions on the scripts for the development of the dataset.

References

- Eyal Ben-David, Nadav Oved, and Roi Reichart. 2021. [PADA: A prompt-based autoregressive approach for adaptation to unseen domains](#). *CoRR*, abs/2102.12206.
- Antoine Bordes, Y-Lan Boureau, and Jason Weston. 2017. [Learning end-to-end goal-oriented dialog](#). In

International Conference on Learning Representations.

- Nicolas Boumal, Vladislav Voroninski, and Afonso S. Bandeira. 2016. The non-convex burer–monteiro approach works on smooth semidefinite programs. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS’16*, page 2765–2773, Red Hook, NY, USA. Curran Associates Inc.
- 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.
- Xiuyi Chen, Jiaming Xu, and Bo Xu. 2019. A working memory model for task-oriented dialog response generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2687–2693.
- Suvodip Dey and Maunendra Sankar Desarkar. 2021. [Hi-dst: A hierarchical approach for scalable and extensible dialogue state tracking](#). In *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 218–227.
- Layla El Asri, Hannes Schulz, Shikhar Sharma, Jeremie Zumer, Justin Harris, Emery Fine, Rahul Mehrotra, and Kaheer Suleman. 2017. [Frames: a corpus for adding memory to goal-oriented dialogue systems](#). In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 207–219, Saarbrücken, Germany. Association for Computational Linguistics.
- Mihail Eric, Lakshmi Krishnan, Francois Charette, and Christopher D. Manning. 2017. [Key-value retrieval networks for task-oriented dialogue](#). In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 37–49, Saarbrücken, Germany. Association for Computational Linguistics.
- Silin Gao, Ryuichi Takanobu, and Minlie Huang. 2021a. [End-to-end task-oriented dialog modeling with semi-structured knowledge management](#). *CoRR*, abs/2106.11796.
- Silin Gao, Ryuichi Takanobu, Wei Peng, Qun Liu, and Minlie Huang. 2021b. [HyKnow: End-to-end task-oriented dialog modeling with hybrid knowledge management](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1591–1602, Online. Association for Computational Linguistics.

- Silin Gao, Yichi Zhang, Zhijian Ou, and Zhou Yu. 2020. [Paraphrase augmented task-oriented dialog generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 639–649, Online. Association for Computational Linguistics.
- Xu Han, Weilin Zhao, Ning Ding, Zhiyuan Liu, and Maosong Sun. 2021. [PTR: prompt tuning with rules for text classification](#). *CoRR*, abs/2105.11259.
- Seokhwan Kim, Mihail Eric, Karthik Gopalakrishnan, Behnam Hedayatnia, Yang Liu, and Dilek Hakkani-Tur. 2020. [Beyond domain APIs: Task-oriented conversational modeling with unstructured knowledge access](#). In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 278–289, 1st virtual meeting. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Shuyang Li, Jin Cao, Mukund Sridhar, Henghui Zhu, Shang-Wen Li, Wael Hamza, and Julian McAuley. 2021. [Zero-shot generalization in dialog state tracking through generative question answering](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1063–1074.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021a. [Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing](#). *CoRR*, abs/2107.13586.
- Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2021b. [GPT understands, too](#). *CoRR*, abs/2103.10385.
- Ilya Loshchilov and Frank Hutter. 2017. [Decoupled weight decay regularization](#). *arXiv preprint arXiv:1711.05101*.
- Biswesh Mohapatra, Gaurav Pandey, Danish Contractor, and Sachindra Joshi. 2021. [Simulated chats for building dialog systems: Learning to generate conversations from instructions](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1190–1203, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Kai Nakamura, Sharon Levy, Yi-Lin Tuan, Wenhua Chen, and William Yang Wang. 2022. [HybridDialogue: An information-seeking dialogue dataset grounded on tabular and textual data](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 481–492, Dublin, Ireland. Association for Computational Linguistics.
- Dinesh Raghu, Nikhil Gupta, and Mausam. 2021a. [Un-supervised Learning of KB Queries in Task-Oriented Dialogs](#). *Transactions of the Association for Computational Linguistics*, 9:374–390.
- Dinesh Raghu, Atishya Jain, Mausam, and Sachindra Joshi. 2021b. [Constraint based knowledge base distillation in end-to-end task oriented dialogs](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 5051–5061, Online. Association for Computational Linguistics.
- Yu Sun, Shuohuan Wang, Shikun Feng, Siyu Ding, Chao Pang, Junyuan Shang, Jiaxiang Liu, Xuyi Chen, Yanbin Zhao, Yuxiang Lu, Weixin Liu, Zhihua Wu, Weibao Gong, Jianzhong Liang, Zhizhou Shang, Peng Sun, Wei Liu, Xuan Ouyang, Dianhai Yu, Hao Tian, Hua Wu, and Haifeng Wang. 2021. [ERNIE 3.0: Large-scale knowledge enhanced pre-training for language understanding and generation](#). *CoRR*, abs/2107.02137.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Yunyi Yang, Yunhao Li, and Xiaojun Quan. 2021. [Ubar: Towards fully end-to-end task-oriented dialog system with gpt-2](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 14230–14238.
- Weijie Zhang, Jiaoxuan Chen, Haipang Wu, Sanhui Wan, and Gongfeng Li. 2021. [A knowledge-grounded dialog system based on pre-trained language models](#). *CoRR*, abs/2106.14444.

A Appendix

A.1 Additional Results

We present additional results for the comparison of JOINTLM, SEPLM and SEKNOW (Gao et al., 2021a) when trained on HYBRIDTOD and tested on the other datasets (Table 9). We see that JOINTLM outperforms SEPLM and SEKNOW on all the datasets demonstrating the importance of joint modeling.




Figure 2: Figures 2a and 2b show the slot-value distribution by slot-types in the hotel and restaurant domains the three datasets.

Train Dataset	Test Dataset	Model	Bleu-1	Bleu-4	slot-values		
					prec.	recall	F1
HYBRIDTOD	SEKNOW-MULTIWOZ	JOINTLM	30.63	8.60	50.48	45.37	47.79
		SEPLM	30.03	8.63	47.26	42.76	44.89
		SEKNOW	29.20	7.83	43.16	28.65	33.14
HYBRIDTOD	HYBRIDTOD	JOINTLM	30.59	8.67	50.56	45.83	48.08
		SEPLM	29.96	8.66	47.08	42.53	44.69
		SEKNOW	29.05	7.70	44.29	29.12	35.14
HYBRIDTOD	UNSTRUCTURED TOD	JOINTLM	30.30	8.44	51.05	45.37	48.04
		SEPLM	29.78	8.41	47.08	41.63	44.19
		SEKNOW	27.43	6.68	42.96	19.62	27.11

Table 9: All models trained on HYBRIDTOD and evaluated on the rest of the datasets

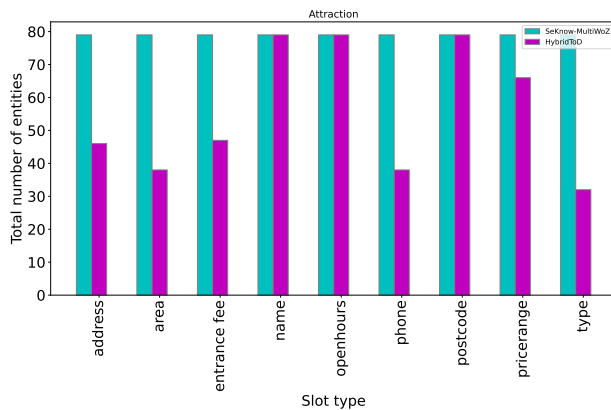


Figure 3: This figure shows the slot-value distribution by slot-types in the attraction domains in the three datasets.

A.2 Additional Dataset Statistics

We present the detailed slot-type distribution of SEKNOW-MULTIWOZ and HYBRIDTOD in Figure 2 and 3. We find that approximately 50% slot-values are moved to unstructured knowledge from the structured sources for each slot-type. The bar-graphs show the number of entities with a particular slot-type.

A.3 Hyperparameters and Training Details

For all our experiments, we use BART (Lewis et al., 2020) model from the HuggingFace Transformers library (Wolf et al., 2020). To train the BART model, we use early stopping with $patience = 5$ on the validation set to prevent overfitting of both the entity retriever and the response generator. We use learning rate = 10^{-5} with AdamW optimizer (Loshchilov and Hutter, 2017). We use a batch-size of 8 examples, with 4 examples for entity retrieval and 4 for response generation per batch. For the 4 examples for entity retrieval, 2 are positive and 2 are negative examples (effectively our batch is $2 + 2 + 4$). All the experiments are conducted on a single A100 80GB GPU.