

# DialoKG: Knowledge-Structure Aware Task-Oriented Dialogue Generation

Md Rashad Al Hasan Rony<sup>1,3</sup>, Ricardo Usbeck<sup>2</sup>, Jens Lehmann<sup>1,3</sup>

<sup>1</sup>University of Bonn, <sup>2</sup>University of Hamburg, <sup>3</sup>Fraunhofer IAIS Dresden

{rashad.rony, jens.lehmann}@iaais.fraunhofer.de

{lehmann}@uni-bonn.de

ricardo.usbeck@uni-hamburg.de

## Abstract

Task-oriented dialogue generation is challenging since the underlying knowledge is often dynamic and effectively incorporating knowledge into the learning process is hard. It is particularly challenging to generate both human-like and informative responses in this setting. Recent research primarily focused on various knowledge distillation methods where the underlying relationship between the facts in a knowledge base is not effectively captured. In this paper, we go one step further and demonstrate how the structural information of a knowledge graph can improve the system’s inference capabilities. Specifically, we propose DialoKG, a novel task-oriented dialogue system that effectively incorporates knowledge into a language model. Our proposed system views relational knowledge as a knowledge graph and introduces (1) a structure-aware knowledge embedding technique, and (2) a knowledge graph-weighted attention masking strategy to facilitate the system selecting relevant information during the dialogue generation. An empirical evaluation demonstrates the effectiveness of DialoKG over state-of-the-art methods on several standard benchmark datasets.

## 1 Introduction

Traditional task-oriented dialogue systems are designed to achieve specific goals such as restaurant reservation, hotel booking and car navigation. These systems are often empowered by external domain- or task-specific knowledge that enables them to generate informative dialogues (Eric et al., 2017; Wu et al., 2019; Qin et al., 2019; He et al., 2020b). The external knowledge in these systems is usually incorporated in the form of structured knowledge triples (Zhou et al., 2018; Liu et al., 2018a) or unstructured documents (Ye et al., 2020; Ghazvininejad et al., 2018). Figure 1 depicts a knowledge-grounded dialogue about reserving a hotel.

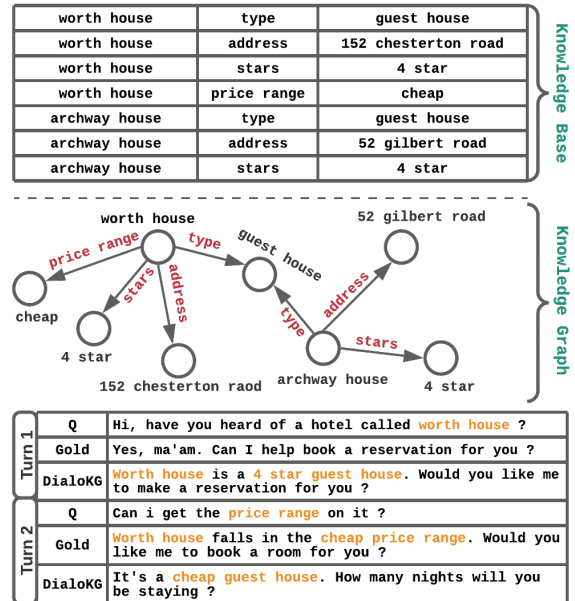


Figure 1: An illustration of knowledge-based multi-turn dialogue where DialoKG models the knowledge base as a Knowledge Graph. The user utterance is denoted by **Q**, the ground-truth response by **Gold**, and the words in orange are knowledge graph entries.

Recent research primarily concentrated on various knowledge filtering methods for selecting relevant knowledge (Wen et al., 2018; Kim et al., 2020; Wu et al., 2020b). These approaches treat the knowledge triples independently and leverage Pointer Networks and copy mechanisms to generate knowledge-grounded dialogues (Vinyals et al., 2015; Gu et al., 2016; Wu et al., 2019; Sukhbaatar et al., 2015; Raghu et al., 2021; Chaudhuri et al., 2021). Typically, these systems generate a template or sketch-response during training and learn to fill in the slots with knowledge graph entries. Such systems face two issues when they try to generate dialogues in a multi-domain setting. **Firstly**, they are unable to capture the underlying semantics of a knowledge graph, such as the relationship between entity and relation. This leads frequently to in-

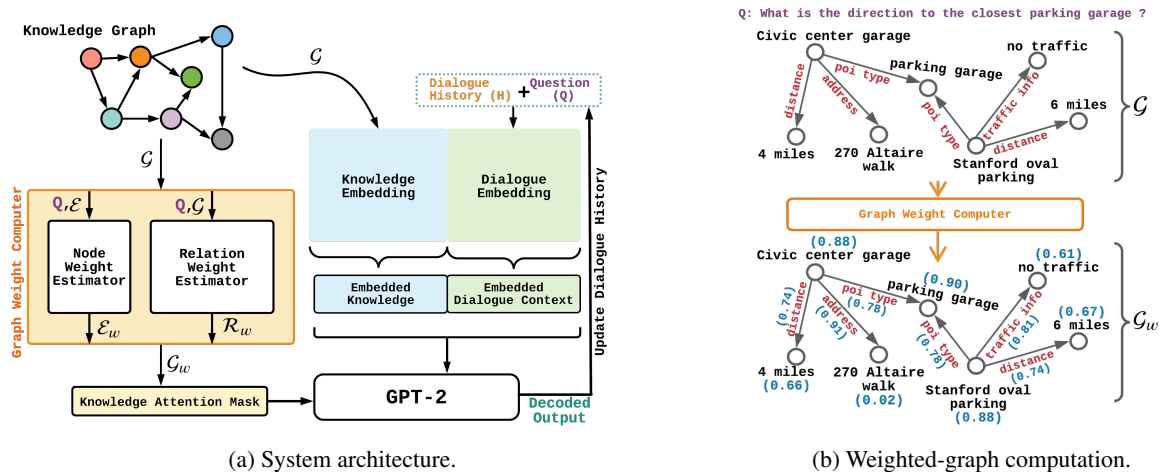


Figure 2: A high-level overview of DialoKG is shown in Figure (a). Figure (b) depicts the input and output of the *Graph Weight Computer* module of DialoKG.

correct and inappropriate dialogue generation (Lin et al., 2020). **Secondly**, they lack the ability to encode dynamic knowledge in a multi-domain setting, resulting in noisy dialogues (Madotto et al., 2020). Generally, integrating a knowledge base into the learning process and generating correct and coherent dialogues at the same time is a challenging task.

In this paper, we propose a novel task-oriented dialogue system, named DialoKG that employs structural information of the knowledge graph into a language model (LM) for generating informative dialogues (see Figure 2a). For this purpose, we exploit GPT-2 (Radford et al., 2019) - a language model developed based on a stack of Transformer decoders (Vaswani et al., 2017). Specifically, we introduce a novel structure-aware multiple embedding layer-based knowledge embedding technique that constructively embeds the underlying relationship between the knowledge triples. DialoKG interprets the knowledge as a knowledge graph; therefore, separate embedding layers for word token, entity, triple and token type enable the system to capture the graph features (e.g., subject, relation and object). This enables the system to generate correct and human-like dialogues and prevents generating erroneous responses such as "4 miles is located at 792 Bedoin Street Starbucks away". Furthermore, the ability to correctly capture the relationship in the knowledge graph eliminates the need for template-based or sketch-based response generation.

In order to guide the decoder on relevant parts of the knowledge graph, we propose a new knowledge attention masking method. For constructing the

knowledge attention mask, in each dialogue turn, a weighted graph is computed in two steps: 1) Entity weights are computed using a pre-trained language model that estimates the importance of an entity for the given utterance, and 2) relation weights are computed based on the concept of graph convolution networks (GCN) (Kipf and Welling, 2017). Both steps take the user utterance into consideration, i.e., the obtained weighted graph is question specific. A set of triples is then selected based on the most relevant entities and relations of the weighted graph to construct a knowledge attention mask for the language model. This allows the masked language model to focus on relevant graph triples. We hypothesise that this leads to the generation of more accurate responses and enhance the model’s capabilities of understanding the domain and task.

To assess the performance of DialoKG, we conduct experiments on three public benchmarks: SMD (Eric et al., 2017), CamRest (Wen et al., 2017) and Multi-WOZ 2.1 (Budzianowski et al., 2018). We evaluate the system generated responses using both human and automatic metrics. Furthermore, we analyse impact of the individual components on the overall performance to verify the effectiveness. Our experimental results show that DialoKG outperforms state-of-the-art models in knowledge-grounded dialogue generation and can generate human-like responses. We made our code publicly available <sup>1</sup>.

<sup>1</sup><https://github.com/rashad101/DialoKG>




Figure 3: An illustration of knowledge and dialogue embedding techniques.

## 2 Approach

### 2.1 Problem Definition

DialoKG aims to generate informative responses given a dialogue history, a question and a knowledge base. We define the dialogue history  $\mathcal{H}$  as a set of turns between two speakers, such that  $\mathcal{H} = \{U_1, S_1, \dots, U_t, S_t\}$ , where  $U_i$  and  $S_i$  are the sequences of words in turn  $i$ . We assume that the knowledge is stored in a multi-relational knowledge graph  $\mathcal{G}$ . Here,  $\mathcal{G}$  is a set of triples  $\mathcal{T}$  such that  $\mathcal{T} \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E}$ , where  $\mathcal{E}$  is the set of entities and  $\mathcal{R}$  the set of relations. A triple  $\mathcal{T} \in \mathcal{G}$  is denoted as  $(s, r, o)$  in which  $s \in \mathcal{E}$  and  $o \in \mathcal{E}$  denote the subject and object entities, respectively, and  $r \in \mathcal{R}$  is the relation between them. We use the terms "Knowledge Graph" and "Graph" interchangeably throughout this paper. Furthermore, we denote the user utterance of the current dialogue turn as  $Q$ . A GPT-2 (Radford et al., 2019) language model is used in this paper to generate responses. However, any Transformer decoder-based LM can be used. Formally, the probability distribution of generating a response by the language model is defined as:

$$p(S_t | \mathcal{H}, Q, \mathcal{G}) = \prod_{i=1}^n p(s_i | s_1, \dots, s_{i-1}, \mathcal{H}, Q, \mathcal{G}) \quad (1)$$

Here,  $S_t$  is the generated response in turn  $t$  and  $n$  is the maximum length of the generated response.

### 2.2 Knowledge and Dialogue Embedding

DialoKG takes a knowledge graph  $\mathcal{G}$ , dialogue history  $\mathcal{H}$ , and the current user utterance  $Q$  together as input and constructs a single input sequence as depicted in Figure 3. The first part of the sequence contains graph related information (i.e., subject, relation, and object) and the latter part dialogue specific information such as dialogue history ( $\mathcal{H}$ ) and the current user utterance ( $Q$ ). **Knowledge Specific Embedding.** To infuse structural information, DialoKG employs entity embedding, triple

embedding and type embedding, besides the usual word token and positional embedding. Such an embedding technique allows the system to encode the knowledge graph structure. To do this, knowledge graph triples are linearized into a sequence as input, as depicted in Figure 3. To facilitate order invariance of the knowledge embedding, we shuffle the order of the graph triples in the input sequence during training. In the token embedding layer [S], [R] and [O] are special tokens to separate subject, relation and object of a triple from each other in the sequence. Entity and triple embedding layers embed entity and triple-level information of the word token. For instance, ENT1 in the entity embedding layer indicates that the corresponding words in the token embedding layer are related to the first subject, which is *starbucks* in this case. Likewise, T1 and T2 in the triple embedding layer indicate that the corresponding words in the token embedding layer are related to the first and second triple, respectively. Finally, the type embedding indicates that the corresponding tokens are from the knowledge graph as opposed to the dialogue history.

**Dialogue Specific Embedding.** The dialogue specific part of the input sequence is separated from the knowledge specific part by a [SEP] token in the token embedding layer. Furthermore, the user utterance/question ( $Q$ ) of the current turn is separated by a [Q] token from the dialogue history. The type embedding layer stores information about whether the corresponding utterance is from the user or system. This way, the decoder can use information about typical dialogue turn patterns.

The positional embedding in both knowledge and dialogue embeddings encodes the position of each word token in the sequence. Finally, embeddings from all five layers are summed up as depicted in Figure 3. *Layer Normalization* (Ba et al., 2016) is then applied to obtain the final embedding representation of the complete input sequence. It

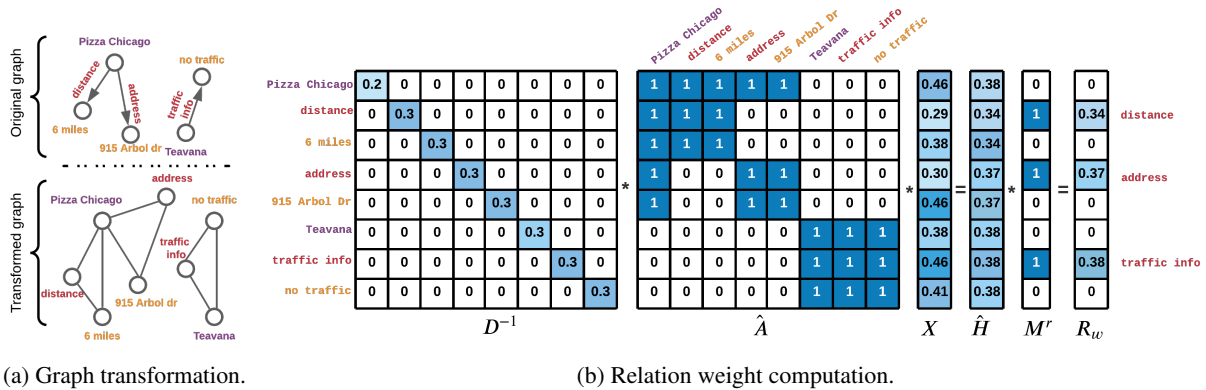


Figure 4: For the graph in Figure (a) and the question "Find me the quickest route to the restaurant?" the computation of the relation weight is shown in Figure (b), where  $\hat{A} = A + I$ .

normalizes the embedding representation of layers, which restricts the weights of the learning network from exploding.

We argue that the proposed design pattern of forming a single sequence and specifying each item in the input sequence further with additional embedding layers can improve the system’s understanding of the task and domain.

### 2.3 Knowledge Attention Mask Construction

To notify the decoder about the relevant KG triples for answering the current user question, a knowledge graph weighted-attention mask is constructed. Prior to the construction of the knowledge attention mask, a weighted-knowledge graph,  $\mathcal{G}_w$  is first computed by a *Graph Weight Computer* module, where the entity and relation weights are computed independently. We discuss the components of the *Graph Weight Computer* module below.

**Entity Weight Estimator.** A pre-trained language model RoBERTa (Liu et al., 2019), is used to compute the entity weights, similar to (Yasunaga et al., 2021). Each entity  $E_i \in \mathcal{E}$  of graph  $\mathcal{G}$  is concatenated with the user utterance  $\mathcal{Q}$  to obtain the probability score from the language model.

$$E_{iw} = LM_{head}(LM_{enc}([\mathcal{Q}; E_i])) \quad (2)$$

In Equation 2,  $LM_{head} \circ LM_{enc}$  represents the probability of the entity  $E_i$  computed by the language model. We consider  $E_{iw}$  as the weight of the entity  $E_i$ , which represents the relevance of the entity for the given user utterance  $\mathcal{Q}$ .

**Relation Weight Estimator.** We follow (Kipf and Welling, 2017; Vashishth et al., 2019) and leverage the concept of GCN to obtain the relation weight. In contrast to the previous works, our

proposed relation weight estimator transforms the input graph into an undirected graph  $\hat{\mathcal{G}}$ , where the relations are considered as nodes of a graph. This transformation technique allows the relation estimator to obtain a score for each relation. The graph transformation is demonstrated in Figure 4a. The relation weight is computed as follows:

$$R_w = \hat{H}M^r, \quad (3)$$

$$\hat{H} = D^{-1}(A + I)X$$

Here,  $D^{-1}(A + I)$  computes the row-normalized adjacency matrix, where  $D$  and  $A$  are respectively the degree matrix and adjacency matrix of the graph  $\hat{\mathcal{G}}$  as depicted in Figure 4b and  $I$  is the identity matrix. Let  $d_g = |\mathcal{E}| + |\mathcal{R}|$  be the total number of entities and relations in the graph  $\hat{\mathcal{G}}$ , then  $D, A, I \in \mathbb{R}^{d_g \times d_g}$ . A feature vector  $X \in \mathbb{R}^{d_g \times 1}$  is obtained by computing the cosine similarity between the embedding of knowledge graph entries (entities and relations) and the embedding of question. Furthermore, a relation mask  $M^r \in \mathbb{R}^{d_g \times 1}$  is constructed by setting a value of 1 and 0 to the positions that correspond to relations and entities, respectively, to attend to the values that correspond to the relations only. The feature vector  $X$  eliminates the need for additional learnable parameters to capture the relationships, unlike GCN. Finally, values that correspond to the entities in  $\hat{H}$  are masked out by multiplying with  $M^r$  to obtain final relation weights  $R_w \in \mathbb{R}^{d_g \times d_g}$ .

The computed weighted-graph assists the model to focus on the task by constructing a knowledge attention mask. We use the normalized score of  $R_w$  and  $E_w$  for constructing the knowledge attention mask. To filter-out irrelevant knowledge triples,

Dataset	#Dialogues	#Utterances	Avg. Length of Utt.	#Utt. with Entities	Avg. #Entities per Utt.
SMD (Eric et al., 2017)	3,031	15,928	9.22	4430	2.96
CamRest (Wen et al., 2017)	676	2,744	11.72	2366	2.43
MWOZ (Budzianowski et al., 2018)	2,877	19,870	16.68	6241	2.06

Table 1: Dataset statistics.

we obtain the top- $k$  entities and relations from the weighted graph as denoted as  $\hat{\mathcal{E}}$  and  $\hat{\mathcal{R}}$  respectively. Here,  $k$  is a hyper-parameter which we chose from a range of  $[0, \max(|\mathcal{E}|, |\mathcal{R}|)]$ , based on the validation score. Finally, based on the selected  $\hat{\mathcal{E}}$  and  $\hat{\mathcal{R}}$ , the knowledge attention mask is constructed as follows:

$$M_{i,j}^{kg} = \begin{cases} 0, & \text{if } ((s_i \vee o_i) \in \hat{\mathcal{E}}) \wedge (r_i \in \hat{\mathcal{R}}) \\ -\infty, & \text{otherwise} \end{cases}$$

Here  $r_i$ ,  $s_i$  and  $o_i$  corresponds to the relation, subject and object entity of triple  $\mathcal{T}_i$ . Any position that corresponds to the value of  $-\infty$  results in 0 after computing the *softmax* during the attention computation (discussed in the next sub-section). The final mask  $M \in \mathbb{R}^{n \times n}$  is obtained by appending the mask for the dialogue related sequence with the knowledge attention mask where  $n$  is sequence length. Padding is added to adjust the dimension of the metrics.

## 2.4 Decoder

A Transformer (Vaswani et al., 2017) based GPT-2 (Radford et al., 2019) model is used for generating the response. The attention, computed in each of GPT-2’s heads is formalized as follows:

$$\begin{aligned} \text{Attn}(Q, K, V) &= \text{softmax}\left(\frac{1}{\sqrt{d_k}}(QK^T) + M\right)V, \\ H_i &= \text{Attn}(QW_i^Q, KW_i^K, VW_i^V) \end{aligned} \quad (4)$$

where,  $\text{Attn}(\cdot)$  computes the masked attention,  $H_i$  is the  $i$ -th head,  $d_k = d_m/h$ . Here,  $d_m$  is the dimension of the model where  $h$  the number of heads.  $Q$ ,  $K$  and  $V$  are query, key and value where  $W_i^Q$ ,  $W_i^K$ ,  $W_i^V$  are trainable parameters. The objective of the model is to minimize the negative log-likelihood  $\mathcal{L}$  for next-token prediction. For a dialogue dataset  $D = \{D_1, D_2, \dots, D_j\}$ , we formally define  $\mathcal{L}$  as follows:

$$\mathcal{L}(D) = - \sum_j^{|D|} \sum_i^n \log p(s_i^j | s_1^j, \dots, s_{i-1}^j, \mathcal{H}^j, \mathcal{Q}^j, \mathcal{G}^j), \quad (5)$$

where  $n$  is the maximum response length and  $\mathcal{H}^j, \mathcal{Q}^j, \mathcal{G}^j \in D_j$ . Top- $k$  sampling (Fan et al., 2018) decoding is used to generate the next word token at each time step, during the inference.

## 3 Experimental Setup

### 3.1 Data

We evaluate DialoKG on three publicly available knowledge-grounded and task-oriented dialogue datasets: Stanford Multi-Domain dataset (SMD) (Eric et al., 2017), CamRest (Wen et al., 2017) and Multi-WOZ 2.1 (MWOZ) (Budzianowski et al., 2018). SMD consists of three domains: weather, navigation, and calendar. MWOZ contains five domains: train, hotel, restaurant, taxi and attraction. We use the splits provided with the datasets for train, validation, and test. Each dialogue is provided with a knowledge base. Table 1 shows the statistics of the benchmark datasets.

### 3.2 Hyper-parameter Settings

Throughout this paper, we use the GPT-2 (Radford et al., 2019) model with 117M parameters. AdamW (Loshchilov and Hutter, 2019) with  $\epsilon = 1e-8$  and learning rate of  $6.25e-5$  is employed as optimizer. GELU (Hendrycks and Gimpel, 2016) is used as activation function. The best hyper-parameters for each dataset were found using grid search and based on the results on the validation set. We run all experiments on a distributed training setting with 10 GPUs, each with 12 GB of memory. More implementation details can be found in Appendix A.

### 3.3 Evaluation Metrics

**Automatic Metrics.** Following the baseline models, we use BLEU (Papineni et al., 2002) and Entity F1 score (Eric et al., 2017) as automatic evaluation metrics. The Entity F1 score represents the model’s capability of generating knowledge grounded responses. It computes the F1 score between the set of entities present in the ground truth and system-generated responses. Several studies (Novikova et al., 2017; Liu et al., 2016) on evaluation metrics suggest that word-overlap based metrics such as BLEU are insufficient for evaluating natural language generation (NLG) systems. Hence, we use MoverScore (Zhao et al., 2019) as

Model	SMD			CamRest			MWOZ		
	BLEU	MoverScore	Ent. F1	BLEU	MoverScore	Ent. F1	BLEU	MoverScore	Ent. F1
GLMP (Wu et al., 2019)	13.9	54.2	59.6	15.1	57.2	58.9	6.9	51.2	32.4
MLM (Gangi Reddy et al., 2019)	17.0	64.0	54.6	15.5	57.0	62.1	-	-	-
Ent. Const. (Qin et al., 2019)	13.9	53.8	53.7	18.5	65.9	58.6	-	-	-
GPT2+KE (Madotto et al., 2020)	17.4	<u>66.4</u>	59.8	18.0	65.8	54.9	<b>15.0</b>	<u>60.9</u>	<u>39.6</u>
TTOS (He et al., 2020a)	17.4	59.8	55.4	20.5	67.0	61.5	-	-	-
DF-Net (Qin et al., 2020)	14.4	56.3	62.7	-	-	-	9.4	54.2	35.1
EER (He et al., 2020c)	17.2	60.9	59.0	19.2	66.1	65.7	13.6	57.2	35.6
FG2Seq (He et al., 2020b)	16.8	60.2	61.1	20.2	66.6	66.4	<u>14.6</u>	58.4	36.5
CDNet (Raghu et al., 2021)	<u>17.8</u>	61.1	<u>62.9</u>	<u>21.8</u>	<u>67.8</u>	<u>68.6</u>	11.9	55.8	38.7
<b>DialoKG</b>	<b>20.0</b>	<b>70.6</b>	<b>65.9</b>	<b>23.4</b>	<b>70.4</b>	<b>75.6</b>	12.6	<b>62.6</b>	<b>43.5</b>

Table 2: Performance of DialoKG and baseline models on three benchmark datasets. Best scores in **bold** and second-best underlined.

addition metric to evaluate the semantic similarity between the system generated response and the ground truth. We compute both MoverScore and BLEU scores on the sentence level.

**Human Evaluation.** To assess the quality of the system-generated responses, we conduct a human evaluation based on the following criteria: 1) Naturalness: how human-like and fluent the generated responses are, and 2) Correctness: how correct the knowledge-grounded responses are. We asked three annotators (two from Computer Science (CS) and one from a non-CS background) who are not part of this research work to evaluate the quality of the system-generated responses. We randomly sampled 90 dialogues in total from the benchmark datasets and asked annotators to evaluate the system-generated responses given the ground truth response and the knowledge graph triples on a scale of [1,5] (higher is better). The inter-annotator agreement score (Cohen’s kappa  $\kappa$ ) of the annotated data is 0.82. The human evaluation process is explained in detail in Appendix D.

### 3.4 Baselines

We compare DialoKG with the following state-of-the-art methods: **GLMP** (Wu et al., 2019), **MLM** (Gangi Reddy et al., 2019), **Ent. Const.** (Qin et al., 2019), **DF-Net** (Qin et al., 2020), **CDNet** (Raghu et al., 2021), **GPT2+KE** (Madotto et al., 2020), **TTOS** (He et al., 2020a) and **EER** (He et al., 2020c). Most of these approaches adopt memory networks to generate knowledge grounded dialogues, whereas **GPT2+KE** (Madotto et al., 2020) directly embeds the knowledge base into the model’s parameters and **TTOS** (He et al., 2020a) proposed a reinforcement learning-based framework.

## 4 Results and Analysis

### 4.1 Quantitative Results

We conduct both quantitative and qualitative analyses to assess system-generated responses. Table 2 summarizes the performance of DialoKG with respect to the baseline models. It is evident that DialoKG outperforms the baseline models significantly in Entity F1 score on CamRest, which contains mostly knowledge-grounded dialogues about restaurant reservations. A high Entity F1 score of 75.6 on CamRest shows DialoKG’s ability to generate knowledge-grounded with high accuracy. Although DialoKG achieves an improved Entity F1 score on the MWOZ dataset, it has a lower BLEU score since MWOZ often contains lengthy responses. However, the high MoverScore across all datasets demonstrates that DialoKG can generate highly semantically similar responses. Domain-wise results are reported in Appendix B due to space limitation.




Figure 5: Distribution of human evaluation scores.

### 4.2 Qualitative Results

We obtain human evaluation scores (naturalness and correctness) for the closest three models. Re-

Model	Naturalness	Correctness
EER (He et al., 2020c)	3.27	3.61
FG2Seq (He et al., 2020b)	3.33	3.87
CDNet (Raghu et al., 2021)	3.53	3.94
<b>DialoKG</b>	<b>4.33</b>	<b>4.01</b>

Table 3: Human evaluation results.

Approach	BLEU	$\Delta$	Ent. F1	$\Delta$
<b>DialoKG</b> (seq2seq)	14.5	-	59.4	-
+ Entity embedding	17.7	3.2 $\uparrow$	63.0	3.6 $\uparrow$
+ Triple embedding	19.2	1.5 $\uparrow$	67.8	4.8 $\uparrow$
+ Type embedding	20.1	0.9 $\uparrow$	68.4	0.6 $\uparrow$
+ Knowledge attention mask	23.4	3.3 $\uparrow$	75.6	7.2 $\uparrow$

Table 4: Ablation study.

sults in Table 3 show that our proposed dialogue system can generate more human-like responses. An improved score is also achieved in terms of correctness, reflecting DialoKG’s ability to generate highly accurate dialogues. Furthermore, Figure 5 shows the distribution of human evaluation scores. The figure allows a better direct comparison of the individual score levels. Details about the annotation process are reported in Appendix D.

### 4.3 Ablation Study

We conducted an ablation study to investigate the contribution of major components of DialoKG. The results on CamRest in Table 4 demonstrates that the *seq2seq* approach achieves the lowest scores, which represents the DialoKG model without the embedding layers: entity embedding, triple embedding, and type embedding. Inclusion of the entity and triple embedding layers significantly improved model’s performance in both BLEU and Entity F1 scores. The type embedding further improved DialoKG’s performance. The significant difference in results shows the effectiveness of the proposed embedding technique. Finally, we observed a remarkable improvement in DialoKG’s overall performance after the inclusion of knowledge attention mask. Question-aware weighted-graph computation used to construct knowledge attention mask, helped the model focus on the task at the inference time.

### 4.4 Effectiveness of Knowledge Embedding

The proposed graph embedding technique works best in combination with the knowledge attention mask. The graph embedding design allows DialoKG to handle disconnected graphs and

Top- <i>k</i> (entity)	Top- <i>k</i> (relation)	BLEU	MoverScore	Entity F1
3	5	10.8	65.3	48.2
3	7	11.0	65.4	48.9
5	5	16.9	68.0	62.1
5	7	17.4	68.1	62.5
7	5	19.3	70.4	64.4
7	7	<b>20.0</b>	<b>70.6</b>	<b>65.9</b>
All	All	15.9	67.2	59.0

Table 5: Effect of triple selection on the performance.

triples. This makes DialoKG suitable for large-scale graphs, where a cosine-similarity based triple selection may be used to fit the graph triples inside the model’s input capacity. The entity and triple embedding layers allow the model to preserve the structural information of a particular triple even though triples from different parts of the input sequence are selected based on the top-*k* entities and relations to construct the knowledge attention mask. Overall, the graph embedding technique improves the Entity F1 score by 5.4, 9.0, and 3.7 points on SMD, CamRest, and MWOZ, respectively. This indicates the effectiveness of the proposed embedding techniques for capturing graph triples.

### 4.5 Impact of Knowledge Attention Mask

To understand the effect of the knowledge-graph weighted attention mask, we experiment with the triple selection process described in DialoKG’s approach. Table 5 shows the performance of DialoKG with selected top-*k* entities and relations on the SMD dataset. We observe that DialoKG achieves the best performance on SMD when the top 7 entities and relations are chosen to construct the knowledge mask. Consider the question "Do you have any local coffee shops?" the ground truth is "There is Coupa, it is just 6 miles away but there is heavy traffic on our way". The ground truth contains traffic information in addition to the distance and name of the coffee shop. Selecting a high number of entities and relations increases the chance of generating such additional information related to the subject of the question. However, choosing too many entities harms the model since it is more likely to add irrelevant noise (see Table 5). Particularly, the entities that are irrelevant for answering the current user utterance act as a noise. For MWOZ, six entities and seven relations, and for CamRest, seven entities and five relations result in the best performance.

Knowledge Triples			Case 1			
Subject	Relation	Object	User		BLEU	MoverScore
Charlie Chan	food	chinese	Truth	I'd like to find a restaurant that serves cheap chinese food	-	-
Charlie Chan	address	Regent street city centre	CDNet	Charlie Chan is located in the centre and it is cheap! Would you like me to book it for you?	0.0	56.5
Charlie Chan	pricerange	cheap	FG2Seq	Charlie Chan is a chinese restaurant in the center of town.	0.0	58.3
Charlie Chan	area	center	DialoKG	Charlie Chan is a cheap restaurant.	14.5	67.3
Yu Garden	address	529 Newmarket road fen ditton	Case 2			
Yu Garden	pricerange	expensive	User	Thank you.		
Yu Garden	food	chinese	Truth	You're welcome . Have a great day!	-	-
Rice House	address	88 mill road city centre	CDNet	You're welcome.	34.6	69.6
Rice House	food	chinese	FG2Seq	Have a great day!	50.8	70.5
Rice House	pricerange	cheap	DialoKG	Thank you for using our service & have a good day!	0.0	66.4

Figure 6: Case study: comparison between ground truth and system-generated responses.

## 5 Case Study

Figure 6 shows two cases from the MWOZ dataset given a subset of the knowledge graph. In Case 1, we observe that in answering the user question, DialoKG correctly picked *Rice House* that serves *cheap* and *Chinese* food. However, in this case, multiple correct answers exist, e.g. *Charlie Chan* also falls into the same category of restaurant. Despite generating the correct answer based on the given knowledge and the user question, DialoKG receives a low Entity F1 score since the generated response entity does not match the ground truth. In Case 2, where the baseline systems focus on imitating the ground truth, DialoKG generates a fluent and engaging response. Despite generating a meaningful and semantically similar sentence, it obtained a BLEU score of 0.0 because of the low overlap with the ground truth response. However, a high MoverScore in both cases indicates DialoKG’s ability to generate a semantically similar response. Overall, we observe that DialoKG can generate human-like, engaging, and informative responses in a multi-turn dialogue setting.

### 5.1 Influence of Dialogue History

Dialogue history is particularly crucial since it gives the model the context for generating the response. In some cases where the entity information is missing in the current user utterance, the dialogue context provides the model with enough information to perform the inference and generate the correct response. For instance, for the question, *What is the food type they serve?*, the name of the restaurant is not given in the question, but the system can infer it from the dialogue history. However, from the experiments, we found that too much dialogue context may inject noisy and irrelevant information to answer the current question, in particular for knowledge-grounded responses

in MWOZ. To quantify this, we selected different numbers of dialogue turns as history for the model’s input depending on the characteristics of the dataset and visualised the result in Figure 7.




Figure 7: DialoKG’s performance on benchmark datasets for different number of dialogue contexts.

## 6 Related Work

**Task-Oriented Dialogue Systems.** Recent dialogue systems mainly leverage Memory Pointer Networks (Sukhbaatar et al., 2015; Madotto et al., 2018; Wu et al., 2019), Copy mechanisms (Gu et al., 2016; Lin et al., 2020; Chaudhuri et al., 2019) and similarity-based knowledge distillation techniques (Wen et al., 2018; Raghu et al., 2021) for the knowledge selection and dialogue generation task. In this research direction, learning to generate template responses and fill in the slot is a common practice (Wu et al., 2019). Dialogue history and knowledge entities are stored in shared memory, facilitating these systems to apply copy mechanisms over the memory space. A multi-level memory architecture is proposed by (Gangi Reddy et al., 2019) that handles the dialogue history and knowledge entries separately.



**Knowledge-Structure Aware Dialogue Generation.** Recently, several knowledge-grounded dialogue systems have attempted to capture structural knowledge to improve the performance of dialogue generation. A sequence-to-sequence model is proposed by (Liu et al., 2018b) that employs global and local attention to understand structural information. Several studies found GCN to be effective for modeling graph-based data. Hence, they construct a graph from a document (Moghe et al., 2020) or the interaction between two speakers (Wu et al., 2021) for generating informative dialogues. In a different approach, an enhanced entity representation is proposed by (He et al., 2020c) by considering the entity information and the structural and relational information of the knowledge entries. In contrast to the previous works, our proposed approach represents the graph’s structural information such as entities and triples through multiple embedding layers.

**Language Model Based Dialogue Generation.** Pre-training a language model with dialogue datasets (Zhang et al., 2020; Bao et al., 2020; Gu et al., 2021) and fine-tuning an already pre-trained model for various dialogue-related sub-tasks such as dialogue state tracking, action decision and response generation (Hosseini-Asl et al., 2020; Wu et al., 2020a; Galetzka et al., 2021) has received much attention in recent years. Recently, (Madotto et al., 2020) proposed a new method to embed the knowledge into the language model parameters. However, the authors noticed that the generated dialogues are sometimes noisy and requires high fine-tuning costs. Despite the success of language model-based approaches, integrating structured knowledge into the dialogue generation process remains a challenging task. Unlike the previous approaches, we designed a structure-aware embedding method and exploit GPT-2 to generate dialogues.

## 7 Conclusion

We have presented DialoKG, a novel knowledge-grounded task-oriented dialogue system improving the state-of-the-art across multiple benchmark datasets. DialoKG focuses on capturing the underlying semantics of the knowledge graph and pays attention to the relevant graph triples to understand the task and generate correct and human-like responses. The key contributions of DialoKG include 1) **Knowledge embedding technique**, that

embeds the structural information of a knowledge graph effectively, and 2) **Knowledge graph-weighted attention masking**, which guides the masked language model to attend to the relevant knowledge entries for generating correct and informative responses. Finally, we showed DialoKG’s ability to generate accurate, diverse, and human-like dialogues through quantitative and qualitative analysis. We performed an ablation study and studied the effect of dialogue history, knowledge embedding and knowledge attention masking.

## Acknowledgements

We acknowledge the support of the following projects: SPEAKER (BMW FKZ 01MK20011A), JOSEPH (Fraunhofer Zukunftsstiftung), OpenGPT-X (BMW FKZ 68GX21007A), the excellence clusters ML2R (BmBF FKZ 01 15 18038 A/B/C), ScaDS.AI (IS18026A-F) and TAILOR (EU GA 952215). The authors also acknowledge the financial support by the Federal Ministry for Economic Affairs and Energy of Germany in the project CoyPu (project number 01MK21007G).

## References

- Lei Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. [Layer normalization](#). *CoRR*, abs/1607.06450.
- Siqi Bao, Huang He, Fan Wang, Hua Wu, and Haifeng Wang. 2020. [PLATO: Pre-trained dialogue generation model with discrete latent variable](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 85–96, Online. Association for Computational Linguistics.
- 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.
- Debanjan Chaudhuri, Md Rashad Al Hasan Rony, Simon Jordan, and Jens Lehmann. 2019. Using a kg-copy network for non-goal oriented dialogues. In *International Semantic Web Conference*, pages 93–109. Springer.
- Debanjan Chaudhuri, Md Rashad Al Hasan Rony, and Jens Lehmann. 2021. Grounding dialogue systems via knowledge graph aware decoding with pre-trained transformers. In *European Semantic Web Conference*, pages 323–339. Springer.

- 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.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018. Hierarchical neural story generation. *arXiv preprint arXiv:1805.04833*.
- Fabian Galetzka, Jewgeni Rose, David Schlangen, and Jens Lehmann. 2021. [Space efficient context encoding for non-task-oriented dialogue generation with graph attention transformer](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7028–7041, Online. Association for Computational Linguistics.
- Revanth Gangi Reddy, Danish Contractor, Dinesh Raghun, and Sachindra Joshi. 2019. [Multi-level memory for task oriented dialogs](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3744–3754, Minneapolis, Minnesota. Association for Computational Linguistics.
- Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen-tau Yih, and Michel Galley. 2018. [A knowledge-grounded neural conversation model](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O.K. Li. 2016. [Incorporating copying mechanism in sequence-to-sequence learning](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1631–1640, Berlin, Germany. Association for Computational Linguistics.
- Jing Gu, Qingyang Wu, Chongruo Wu, Weiyan Shi, and Zhou Yu. 2021. [PRAL: A tailored pre-training model for task-oriented dialog generation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 305–313, Online. Association for Computational Linguistics.
- Wanwei He, Min Yang, Rui Yan, Chengming Li, Ying Shen, and Ruifeng Xu. 2020a. [Amalgamating knowledge from two teachers for task-oriented dialogue system with adversarial training](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3498–3507, Online. Association for Computational Linguistics.
- Zhenhao He, Yuhong He, Qingyao Wu, and Jian Chen. 2020b. [Fg2seq: Effectively encoding knowledge for end-to-end task-oriented dialog](#). In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8029–8033.
- Zhenhao He, Jiachun Wang, and Jian Chen. 2020c. Task-oriented dialog generation with enhanced entity representation. In *INTERSPEECH*, pages 3905–3909.
- Dan Hendrycks and Kevin Gimpel. 2016. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*.
- 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*.
- Byeongchang Kim, Jaewoo Ahn, and Gunhee Kim. 2020. Sequential Latent Knowledge Selection for Knowledge-Grounded Dialogue. In *ICLR*.
- Thomas N. Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*.
- Xiexiong Lin, Weiyu Jian, Jianshan He, Taifeng Wang, and Wei Chu. 2020. [Generating informative conversational response using recurrent knowledge-interaction and knowledge-copy](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 41–52, Online. Association for Computational Linguistics.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Mike Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. [How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2122–2132, Austin, Texas. Association for Computational Linguistics.
- Shuman Liu, Hongshen Chen, Zhaochun Ren, Yang Feng, Qun Liu, and Dawei Yin. 2018a. Knowledge diffusion for neural dialogue generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1489–1498.
- Tianyu Liu, Kexiang Wang, Lei Sha, Baobao Chang, and Zhifang Sui. 2018b. Table-to-text generation by structure-aware seq2seq learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- 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. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.

- Andrea Madotto, Samuel Cahyawijaya, Genta Indra Winata, Yan Xu, Zihan Liu, Zhaojiang Lin, and Pascale Fung. 2020. [Learning knowledge bases with parameters for task-oriented dialogue systems](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2372–2394, Online. Association for Computational Linguistics.
- Andrea Madotto, Chien-Sheng Wu, and Pascale Fung. 2018. [Mem2seq: Effectively incorporating knowledge bases into end-to-end task-oriented dialog systems](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1468–1478. Association for Computational Linguistics.
- Nikita Moghe, Priyesh Vijayan, Balaraman Ravindran, and Mitesh M. Khapra. 2020. [On incorporating structural information to improve dialogue response generation](#). In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pages 11–24, Online. Association for Computational Linguistics.
- Jekaterina Novikova, Ondřej Dušek, Amanda Cercas Curry, and Verena Rieser. 2017. [Why we need new evaluation metrics for NLG](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2241–2252, Copenhagen, Denmark. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Libo Qin, Yijia Liu, Wanxiang Che, Haoyang Wen, Yangming Li, and Ting Liu. 2019. [Entity-consistent end-to-end task-oriented dialogue system with KB retriever](#). 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 133–142, Hong Kong, China. Association for Computational Linguistics.
- Libo Qin, Xiao Xu, Wanxiang Che, Yue Zhang, and Ting Liu. 2020. [Dynamic fusion network for multi-domain end-to-end task-oriented dialog](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6344–6354, Online. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Dinesh Raghu, Atishya Jain, Mausam, and Sachindra Joshi. 2021. [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.
- Sainbayar Sukhbaatar, Arthur Szlam, Jason Weston, and Rob Fergus. 2015. [End-to-end memory networks](#). In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2, NIPS’15*, page 2440–2448, Cambridge, MA, USA. MIT Press.
- Shikhar Vashishth, Prateek Yadav, Manik Bhandari, and Partha Talukdar. 2019. [Confidence-based graph convolutional networks for semi-supervised learning](#). In *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pages 1792–1801. PMLR.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. [Pointer networks](#). In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.
- Haoyang Wen, Yijia Liu, Wanxiang Che, Libo Qin, and Ting Liu. 2018. [Sequence-to-sequence learning for task-oriented dialogue with dialogue state representation](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3781–3792, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrkšić, Milica Gašić, 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, Valencia, Spain. Association for Computational Linguistics.
- Chien-Sheng Wu, Steven C.H. Hoi, Richard Socher, and Caiming Xiong. 2020a. [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, Online. Association for Computational Linguistics.
- Chien-Sheng Wu, Richard Socher, and Caiming Xiong. 2019. [Global-to-local memory pointer networks for task-oriented dialogue](#). In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Han Wu, Kun Xu, and Linqi Song. 2021. [Csagn: Conversational structure aware graph network for conversational semantic role labeling](#). *arXiv preprint arXiv:2109.11541*.
- Sixing Wu, Ying Li, Dawei Zhang, and Zhonghai Wu. 2020b. [Improving knowledge-aware dialogue response generation by using human-written prototype](#)

dialogues. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1402–1411, Online. Association for Computational Linguistics.

Michihiro Yasunaga, Hongyu Ren, Antoine Bosselut, Percy Liang, and Jure Leskovec. 2021. [QA-GNN: Reasoning with language models and knowledge graphs for question answering](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 535–546, Online. Association for Computational Linguistics.

Hao-Tong Ye, Kai-Lin Lo, Shang-Yu Su, and Yun-Nung Chen. 2020. Knowledge-grounded response generation with deep attentional latent-variable model. *Computer Speech & Language*, 63:101069.

Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. Dialogpt: Large-scale generative pre-training for conversational response generation. In *ACL, system demonstration*.

Wei Zhao, Maxime Peyrard, Fei Liu, Yang Gao, Christian M. Meyer, and Steffen Eger. 2019. [MoverScore: Text generation evaluating with contextualized embeddings and earth mover distance](#). 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 563–578, Hong Kong, China. Association for Computational Linguistics.

Hao Zhou, Tom Young, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018. Commonsense knowledge aware conversation generation with graph attention. In *IJCAI*, pages 4623–4629.

## A Hyper-parameter Settings

We report the hyper-parameters used to train DialoKG in Table 6 for SMD, CamRest, and MWOZ. GPT-2 specific hyper-parameters are also reported in Table 6. All the hyper-parameters are found after a grid search and evaluation on the validation set. We sample learning rate from {6.25e-01, 6.25e-04, 6.25e-05} and maximum history token and knowledge token from {128, 256, 384, 512}.

	SMD	CamRest	MWOZ
Learning rate	6.25e-05	6.25e-04	6.25e-05
Adam epsilon	1e-08	1e-08	1e-08
Batch size	4	4	4
Gradient accumulation steps	4	4	4
Max history turn	4	4	1
Maximum history token	128	256	128
Maximum knowledge token	384	256	384
Top relations	7	7	6
Top entities	7	5	7
Epochs	40	25	30

Table 6: Training parameters.

For both training and evaluation, we use a batch size of 4. Hyper-parameters used during the inference are reported in Table 7. We used 12 NVIDIA TitanX GPUs, each with 12GB of memory to train models. It took 30, 18 and 45 minutes to train on SMD, CamRest and MWOZ data.

	SMD	CamRest	MWOZ
Temperature	0.68	0.85	0.18
Top-k	6	8	10
Top-p	0.9	0.9	0.9
Maximum response length	100	80	120
Top entities	7	7	6
Top relations	7	5	7

Table 7: Decoding parameters.

## B Results

We report the domain-wise results for SMD and MWOZ in Table 8 and Table 9 respectively. Baseline model’s results are reported from (Raghu et al., 2021) and (Madotto et al., 2020). The MWOZ dialogue dataset contains conversations on the following domains as reported in the baseline works: attraction, restaurant, and hotel. The domain-wise results demonstrate that DialoKG achieves improved performance in almost all domains in a multi-domain setup. This demonstrates DialoKG’s capacity to handle a dynamic knowledge base.

Models	BLEU	MoverScore	Entity F1	Schedule	Navigate	Weather
GLMP (Wu et al., 2019)	13.9	54.2	59.6	72.5	54.6	56.5
MLM (Gangi Reddy et al., 2019)	17.0	64.0	54.6	66.7	46.9	56.0
Ent. Const. (Qin et al., 2019)	13.9	53.8	53.7	55.6	54.5	52.2
GPT2+KE (Madotto et al., 2020)	17.4	66.4	59.8	72.6	53.5	57.7
TTOS (He et al., 2020a)	17.4	59.8	55.4	63.5	45.9	64.1
DF-Net (Qin et al., 2020)	14.4	56.3	62.7	73.1	57.9	57.6
EER (He et al., 2020c)	17.2	60.9	59.0	71.8	52.5	57.8
FG2Seq (He et al., 2020b)	16.8	60.2	61.1	73.3	56.1	57.4
CDNet (Raghu et al., 2021)	17.8	61.1	62.9	75.4	56.7	61.3
<b>DialoKG (Ours)</b>	<b>20.0</b>	<b>70.6</b>	<b>65.9</b>	<b>77.9</b>	<b>58.4</b>	<b>72.7</b>

Table 8: Domain-wise results on SMD dataset.

Models	BLEU	MoverScore	Entity F1	Attraction	Restaurant	Hotel
GLMP (Wu et al., 2019)	6.9	51.2	32.4	24.4	38.4	28.1
MLM (Gangi Reddy et al., 2019)	-	-	-	-	-	-
Ent. Const. (Qin et al., 2019)	-	-	-	-	-	-
GPT2+KE (Madotto et al., 2020)	<b>15.0</b>	60.9	39.6	<b>43.3</b>	37.1	33.4
TTOS (He et al., 2020a)	-	-	-	-	-	-
DF-Net (Qin et al., 2020)	9.4	54.2	35.1	28.1	40.9	30.6
EER (He et al., 2020c)	13.6	57.2	35.6	43.0	34.3	35.7
FG2Seq (He et al., 2020b)	14.6	58.4	36.5	37.2	38.9	34.4
CDNet (Raghu et al., 2021)	11.9	55.8	38.7	38.9	41.7	36.3
<b>DialoKG (Ours)</b>	12.6	<b>62.6</b>	<b>43.5</b>	39.8	<b>46.7</b>	<b>37.9</b>

Table 9: Domain-wise results on MWOZ dataset.

## C Knowledge Triples to Sequence Transformation

Figure 8 depicts how we linearize a graph into a sequence. The sequence begins with a [BOS] token, followed by the token [S] and a subject (*worth house*). The token [S] indicates that the following word in the sequence is a subject (in this case *worth house*). Then we append all the triples that are connected to the subject *worth house* where the relation and object is separated by the token [R] and [O], respectively. Similarly, the second subject is appended to the sequence separated by a preceding [S] token.

## D Human Evaluation

Figure 9 shows the interface of the annotation tool used to obtain human annotation scores. The interface displays a set of knowledge triples, a user utterance, the ground truth response, and a system-generated response for each point. Given the information displayed on the annotation tool, we asked the annotators to rate the system-generated responses against the ground-truth on a scale of [1,5] (higher is better). We explained the participants about the purpose of this research. The first two participants are male (over 30 years old), and

the third participant is female (more than 35 years old), both with several years of experience in the domain.

## E Example System Outputs

We show conversations performed by DialoKG on SMD and MWOZ 2.1 dataset in Figure 10 and Figure 11, respectively. The example conversations demonstrate that DialoKG can perform accurate and engaging conversations.




Figure 8: Illustration of graph to sequence transformation.

**Knowledge Base**

Subject	Relation	Object
chevron	distance	5 miles
chevron	traffic info	moderate traffic
chevron	poi type	gas station
chevron	address	783 arcadia pl
town and country	distance	5 miles
town and country	address	383 university ave
jacks house	poi type	friends house
jacks house	address	864 almanor ln
the clement hotel	traffic info	no traffic

**User utterance:** What is the address of chevron ?

**Ground truth:** 783 arcadia pl is the address for chevron gas station .

**System generated response:** chevron is located at 783 arcadia pl .

Naturalness  1  2  3  4  5

Correctness  1  2  3  4  5

Save Previous Next

Figure 9: The interface of the annotation tool to obtain the human annotation scores.

the westin	poi type	rest stop
the westin	address	329 el camino real
the westin	distance	4 miles
the westin	traffic info	no traffic
sigona farmers market	poi type	grocery store
sigona farmers market	address	638 amherst st
sigona farmers market	distance	4 miles
sigona farmers market	traffic info	no traffic
whole foods	distance	2 miles
whole foods	poi type	grocery store
whole foods	traffic info	heavy traffic
stanford shopping center	distance	2 miles
stanford shopping center	poi type	shopping center
stanford shopping center	traffic info	moderate traffic

Turn 1	Q	Give me the directions to the closest <a href="#">grocery store</a> ?
	Gold	There are <a href="#">whole foods</a> 2 miles away and <a href="#">sigona farmers market</a> 4 miles away .
	DialoKG	The nearest <a href="#">grocery store</a> is <a href="#">whole foods</a> located at 2 miles away .
Turn 2	Q	I need the one that avoids all <a href="#">heavy traffic</a> .
	Gold	<a href="#">sigona farmers market</a> is located 4 miles away with no traffic is located at 638 amherst st .
	DialoKG	There is <a href="#">heavy traffic</a> on our way to the nearest <a href="#">grocery store</a> , but I could find another one nearby.
Turn 3	Q	Thank you .
	Gold	Great, glad I could be help. Have a great day we will be there soon .
	DialoKG	You are welcome .

Figure 10: An example dialogue performed by DialoKG on SMD dataset.

acron guest house	type	guest house
acron guest house	address	154 chesteron road
acron guest house	price range	moderate
acron guest house	area	north
acron guest house	stars	4 star
warkworth house	type	guest house
warkworth house	address	warkworth terrace
warkworth house	price range	moderate
warkworth house	area	east
warkworth house	stars	4 stars
kirkwood house	type	guest house
kirkwood house	address	172 chesteron road
kirkwood house	price range	moderate
kirkwood house	stars	4 star

Turn 1	Q	Can I get some help locating a room for the night ?
	Gold	Absolutely. What type of lodging did you have in mind ?
	DialoKG	I have a great number of places to stay. Do you have a specific <a href="#">area</a> you would like to stay in ?
Turn 2	Q	I am looking for a <a href="#">guesthouse</a> with free wifi and parking, please .
	Gold	You're in luck, there are a great number of <a href="#">guesthouse</a> that might work. Did you have an <a href="#">area</a> in mind ?
	DialoKG	I have a great number of guest houses that offer both free wiki and parking. Do you have a specific <a href="#">area</a> you would like to stay in ?
Turn 3	Q	I don't care about the <a href="#">area</a> . Which one do you suggest ?
	Gold	How about <a href="#">acron guest house</a> ?
	DialoKG	Alright, how about <a href="#">acron guest house</a> ?

Figure 11: An example dialogue performed by DialoKG on MWOZ 2.1 dataset.