

OpenDialKG: Explainable Conversational Reasoning with Attention-based Walks over Knowledge Graphs

Seungwhan Moon, Pararth Shah, Anuj Kumar, Rajen Subba
Facebook Conversational AI

{shanemoon, pararths, anujk, rasubba@}fb.com

Abstract

We study a conversational reasoning model that strategically traverses through a large-scale common fact knowledge graph (KG) to introduce engaging and contextually diverse entities and attributes. For this study, we collect a new **Open-ended Dialog** \leftrightarrow **KG** parallel corpus called *OpenDialKG*, where each utterance from 15K human-to-human role-playing dialogs is manually annotated with ground-truth reference to corresponding entities and paths from a large-scale KG with 1M+ facts. We then propose the *DialKG Walker* model that learns the symbolic transitions of dialog contexts as structured traversals over KG, and predicts natural entities to introduce given previous dialog contexts via a novel domain-agnostic, attention-based graph path decoder. Automatic and human evaluations show that our model can retrieve more natural and human-like responses than the state-of-the-art baselines or rule-based models, in both in-domain and cross-domain tasks. The proposed model also generates a KG walk path for each entity retrieved, providing a natural way to explain conversational reasoning.

1 Introduction

The key element of an open-ended dialog system is its ability to understand conversational contexts and to respond naturally by introducing relevant entities and attributes, which often leads to increased engagement and coherent interactions (Chen et al., 2018). While a large-scale knowledge graph (KG) includes vast knowledge of all the related entities connected via one or more factual connections from conversational contexts, the core challenge is in the domain-agnostic and scalable prediction of a small subset from those reachable entities that follows natural conceptual threads that can keep conversations engaging and meaningful. Hence, we study a data-driven reasoning model

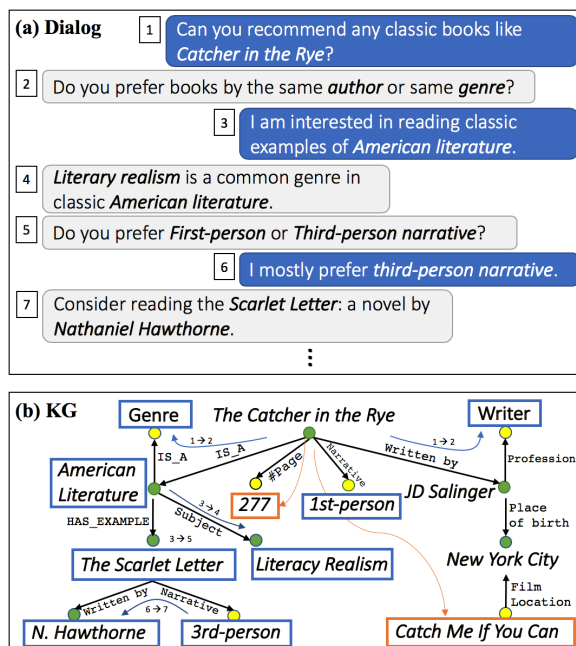


Figure 1: **Conversational reasoning** with a parallel (a) dialog and (b) knowledge graph (KG) corpus. Diverse topical jumps across open-ended multi-turn dialogs are annotated and grounded with a large-scale common-fact KG. To generate a KG entity response at each dialog turn, the model learns *walkable* paths within KG that lead to engaging and natural topics or entities given dialog context, while pruning *non-ideal* (albeit factually correct) KG paths among 1M+ candidate facts.

that map dialog transitions with KG paths, aimed at identifying a subset of ideal entities to mention as a response to previous dialog contexts.

Figure 1 illustrates a motivating dialog example between two conversation participants, which spans multiple related KG entities from a starting seed entity *The Catcher in the Rye*. Specifically, we observe that there exists a small subset of *walkable* patterns within a KG or a preferred sequence of graph traversal steps which often leads to more engaging entities or attributes than others (e.g. *Literacy Realism*, *Nathaniel Hawthorne*, etc. vs. *Catch Me If You Can*, *277*, etc. -

all connected via one- or multi-hop factual connections). Note also that the *walkable* degree of each entity varies by dialog contexts and domains, thus making conventional rule-based or entity-to-entity learning approaches intractable or not scalable for open-ended dialogs with 1M+ candidate facts. Therefore, pruning the search space for entities based on dialog contexts and their relation-based walk paths is a crucial step in operating knowledge-augmented dialog systems at scale.

To this end, we propose a new model called *DialKG Walker* that can learn natural *knowledge paths* among entities mentioned over dialog contexts, and *reason* grounded on a large common-sense KG. Specifically, we propose a novel *graph decoder* that attends on viable KG paths to predict the most relevant entities from a KG, by associating these paths with the given input contexts: dialog, sentence, and a set of starting KG entities mentioned in the previous turn. We then build a parallel zeroshot learning model that predicts entities in the KG embeddings space, and ranks candidate entities based on decoded graph path output.

To train the DialKG Walker model with ground-truth reference to KG entities, we collect a new human-to-human multi-turn dialogs dataset (91K utterances across 15K dialog sessions) using ParIAI (Miller et al., 2017), where conversation participants play a role either as a user or as an assistant, while annotating their mention of an entity in a large-scale common fact KG. This new dataset provides a new way for researchers to study how conversational topics could jump across many different entities within multi-turn dialogs, grounded on KG paths that thread all of them. To the best of our knowledge, our *OpenDialKG* is the first parallel Dialog \leftrightarrow KG corpus where each mention of a KG entity and its factual connection in an open-ended dialog is fully annotated, allowing for in-depth study of symbolic reasoning and natural language conversations.

Note that our approaches are distinct from the previous work on dialog systems in that we completely ground dialogs in a large-scale common-fact KG, allowing for domain-agnostic conversational reasoning in open-ended conversations across various domains and tasks (*e.g.* chit-chat, recommendations, etc.) We therefore perform extensive cross-domain and transfer learning evaluations to demonstrate its flexibility. See Section 5 for the detailed literature review.

Our contributions are as follows: we propose (1) a novel attention-based *graph decoder* that walks an optimal path within a large common-sense KG (100K entities, 1.1M facts) to effectively prune unlikely candidate entities, and (2) a zeroshot learning model that leverages previous sentence, dialog, and KG contexts to re-rank candidates from pruned decoder graph output based on their relevance and path scores, which allows for generalizable and robust classification with a large number of candidate classes. We present (3) a new parallel open-ended dialog \leftrightarrow KG corpus called *OpenDialKG* where each mention of an entity in dialog is manually linked with its corresponding ground-truth KG path. We show that the proposed approaches outperform baselines in both in-domain and cross-domain evaluation, demonstrating that the model learns domain-agnostic walking patterns that are generalizable for unseen domains.

2 Method

Figure 2 illustrates the overall architecture of the DialKG Walker model which retrieves a set of entities from a provided KG given multiple modalities of dialog contexts. Specifically, for each turn the model takes as input a set of KG entities mentioned at its current turn, a full sentence at the current turn, and all sentences from previous turns of dialog, which are encoded using Bi-LSTMs with self-attention modules (Section 2.2). The autoregressive graph decoder takes attention-based encoder output at each decoding step to generate a walk path for each starting KG entity, which is combined with zeroshot KG embeddings prediction results to rank candidate entities (Section 2.3).

2.1 Notations

We define the knowledge graph $\mathbf{G}_{\text{KG}} = \mathbf{V}_{\text{KG}} \times \mathbf{R}_{\text{KG}}$ which is composed of all common-sense entity nodes \mathbf{V}_{KG} and the relation set \mathbf{R}_{KG} that connects each pair of two nodes. Let us also denote $\mathbf{V}_r(v)$ to be a set of nodes directly connected to a node $v \in \mathbf{V}_{\text{KG}}$ by a relation $r \in \mathbf{R}_{\text{KG}}$. Similarly, we denote $\mathbf{V}_{R,n}(v)$ to be a set of nodes connected to v via n -hops with a set of relations R .

Each input is composed of three modalities: $\mathbf{x} = \{\mathbf{x}_e; \mathbf{x}_s; \mathbf{x}_d\}$, where $\mathbf{x}_e = \{\mathbf{x}_e^{(i)}\}$ is a set of entities mentioned in the current turn, \mathbf{x}_s is its surrounding sentence context in the same turn, and \mathbf{x}_d is its dialog context up to the previous turn.

Each output is a KG path sequence that con-

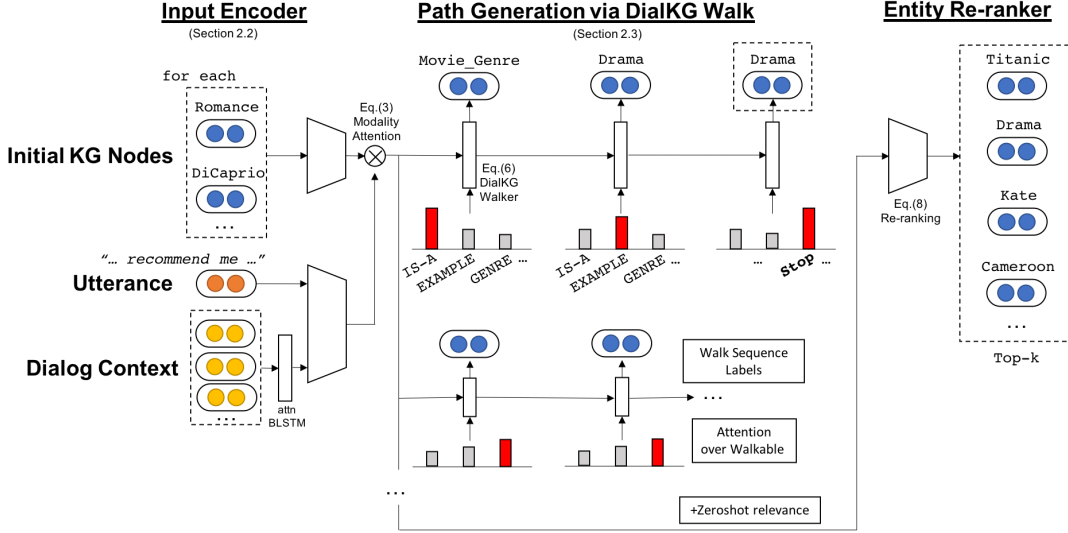


Figure 2: **Overall architecture.** $\mathbf{x} = \{\mathbf{x}_e; \mathbf{x}_s; \mathbf{x}_d\}$ is encoded with the input encoder (left), aggregated via multiple attention mechanism. The decoder (right) predicts both the optimal paths and the final entities $\mathbf{y} = \{\mathbf{y}_e; \mathbf{y}_r\}$ based on their zeroshot relevance scores as well as soft-attention based walk paths, which prunes unlikely entities.

nects \mathbf{x} with entities mentioned in the next turn, which is represented in two modalities: $\mathbf{y} = \{\mathbf{y}_e; \mathbf{y}_r\}$, where $\mathbf{y}_e = \{\mathbf{y}_e^{(i)}\}$ is a set of entity paths, where each entity path with length T is defined as $\mathbf{y}_e^{(i)} = \{\mathbf{y}_{e,t}^{(i)}\}_{t=1}^T$. Similarly, $\mathbf{y}_r = \{\mathbf{y}_r^{(i)}\}$ is a set of relation paths, each with length T that connects $\mathbf{x}_e^{(i)}$ and $\mathbf{y}_e^{(i)}$ via relations in \mathbf{R}_{KG} .

We formulate the future entity retrieval task as:

$$\mathbf{y} = \underset{\mathbf{y}'_e \subset \mathbf{V}(\mathbf{x}_e)}{\operatorname{argmax}} \operatorname{score}(\mathbf{f}_{\mathbf{x} \rightarrow \mathbf{y}}(\mathbf{x}), \mathbf{y}')$$

where $\mathbf{f}_{\mathbf{x} \rightarrow \mathbf{y}}$ is a function with learnable parameters that projects input samples at the current turn (\mathbf{x}) into the same space as the output representations (\mathbf{y}), *i.e.* entities to be mentioned in the next turn and their optimal paths. $\mathbf{V}(\mathbf{x}_e) \subset \mathbf{V}_{\text{KG}}$ denotes a set of KG entity nodes reachable from \mathbf{x}_e , defined accordingly to each decoding method.

2.2 Input Encoding

Entity representation: We construct KG embeddings to encode each entity mention (Bordes et al., 2013), in which semantically similar entities are distributed closer in the embeddings space. In brief formulation, the model for obtaining embeddings from a KG (composed of subject-relation-object (s, r, o) triples) is as follows:

$$P(\mathbb{I}_r(s, o) = 1 | \theta) = \operatorname{score}(\mathbf{e}(s), \mathbf{e}_r(r), \mathbf{e}(o)) \quad (1)$$

where \mathbb{I}_r is an indicator function of a known relation r for two entities (s, o) (1: valid relation, 0: unknown relation), \mathbf{e} is a function that extracts

embeddings for entities, \mathbf{e}_r extracts embeddings for relations, and $\operatorname{score}(\cdot)$ is a deep neural network that produces a likelihood of a valid triple.

Sentence representation: We represent textual context of surrounding words of a mention with a state-of-the-art attention-based Bi-LSTM language model (Conneau et al., 2017) with GloVe (Pennington et al., 2014) distributed word embeddings trained on the Wikipedia and the Gigaword corpus with a total of 6B tokens.

Dialog representation: To encode previous dialog history, we use a hierarchical Bi-LSTM (Yang et al., 2016) over a sequence of previous sentences with a fixed window size. We apply self-attention over sentences to attenuate and amplify sentence contexts based on their relevance to the task, allowing for more robust and explainable prediction.

Input aggregation: We aggregate input contexts $\bar{\mathbf{x}}$ from entities, sentences and dialogs, by applying the modality attention (Moon et al., 2018a,b), which selectively attenuates or amplifies each modality based on their importance on the task:

$$[\mathbf{a}_e; \mathbf{a}_s; \mathbf{a}_d] = \sigma(\mathbf{W}_m \cdot [\mathbf{x}_e; \mathbf{x}_s; \mathbf{x}_d] + \mathbf{b}_m) \quad (2)$$

$$\alpha_m = \frac{\exp(\mathbf{a}_m)}{\sum_{m' \in \{e, s, d\}} \exp(\mathbf{a}_{m'})} \quad \forall m \in \{e, s, d\}$$

$$\bar{\mathbf{x}} = \sum_{m \in \{e, s, d\}} \alpha_m \mathbf{x}_m \quad (3)$$

where $\alpha = [\alpha_e; \alpha_s; \alpha_d] \in \mathbb{R}^3$ is an attention vector, and $\bar{\mathbf{x}}$ is a final context vector that maximizes information gain.

2.3 Graph Decoder

Using the contextual information extracted from an entity and its surrounding text (Section 2.2), we build a network which predicts a corresponding KG entity based on its knowledge graph embeddings with the following objective:

$$\min_{\mathbf{W}} \mathcal{L}_f(\mathbf{x}, \mathbf{y}_e; \mathbf{W}_f, \mathbf{W}_p) + \mathcal{L}_{\text{walk}}(\mathbf{x}, \mathbf{y}_p; \mathbf{W}_p) \\ \mathcal{R}(\mathbf{W}): \text{regularization} \quad (4)$$

where $\mathcal{L}_f(\cdot)$ is a supervised loss for generating the correct entity at the next turn, and $\mathcal{L}_{\text{walk}}(\cdot)$ is a loss defined for taking the optimal path within a knowledge graph. $\mathbf{W} = \{\mathbf{W}_f, \mathbf{W}_p, \mathbf{W}_{\text{input}}\}$ are the learnable parameters for the final entity classifier (\mathbf{W}_f), the path walker model (\mathbf{W}_p), and the input encoder, respectively. $\mathcal{R}(\mathbf{W})$ denotes the weight decay regularization term.

2.3.1 Zeroshot Relevance Score

We compute zeroshot relevance score in the KG embeddings space, thus allowing for robust prediction for KG entities and domains unseen during training as well. Specifically, we use the supervised hinge rank loss for KG embeddings prediction as a choice of \mathcal{L}_f , defined for each sample (Moon and Carbonell, 2017).

$$\sum_i \sum_{\tilde{\mathbf{y}} \neq \mathbf{y}_e^{(i)}} \max[0, \tilde{\mathbf{y}} \cdot \mathbf{y}_e^{(i)} - \mathbf{f}(\bar{\mathbf{x}}^{(i)}) \cdot (\mathbf{y}_e^{(i)} - \tilde{\mathbf{y}})^\top] \quad (5)$$

where $\mathbf{f}(\cdot)$ is a transformation function that walks through the knowledge graph and projects a predicted future entity in the KG embeddings space, and $\tilde{\mathbf{y}}$ refers to the embeddings of negative samples randomly sampled from KG entities except the ground truth label of the instance. Intuitively, the model is trained to produce a higher dot product similarity between the projected embeddings of a sample with its correct label ($\mathbf{f}(\bar{\mathbf{x}}^{(i)}) \cdot \mathbf{y}_e^{(i)}$) than with an incorrect negative label in the KG label embeddings space ($\mathbf{f}(\bar{\mathbf{x}}^{(i)}) \cdot \tilde{\mathbf{y}}$), where the margin is defined as the similarity between a ground truth sample and a negative sample ($\tilde{\mathbf{y}} \cdot \mathbf{y}_e^{(i)}$).

2.3.2 KG Path Walker

Generating candidate KG entities solely based on their relevance score (Eq.5) is challenging due to the exponentially large search space. To this end, we define the attention-based DialKG graph decoder model which prunes unattended paths, which effectively reduce the search space. Decoding steps are formulated as follows (bias terms for

gates are omitted for simplicity of notation):

$$\mathbf{i}_t = \sigma(\mathbf{W}_{hi}\mathbf{h}_{t-1} + \mathbf{W}_{ci}\mathbf{c}_{t-1}) \\ \mathbf{c}_t = (1 - \mathbf{i}_t) \odot \mathbf{c}_{t-1} \\ \quad + \mathbf{i}_t \odot \tanh(\mathbf{W}_{zc}\mathbf{z}_t + \mathbf{W}_{hc}\mathbf{h}_{t-1}) \\ \mathbf{o}_t = \sigma(\mathbf{W}_{zo}\mathbf{z}_t + \mathbf{W}_{ho}\mathbf{h}_{t-1} + \mathbf{W}_{co}\mathbf{c}_t) \\ \mathbf{h}_t = \text{WALK}(\bar{\mathbf{x}}, \mathbf{z}_t) = \mathbf{o}_t \odot \tanh(\mathbf{c}_t) \quad (6)$$

where \mathbf{z}_t is a context vector at decoding step t , produced from the attention over walkable path which is defined as follows:

$$\alpha_t = \sigma(\mathbf{W}_{h\alpha}\mathbf{h}_{t-1} + \mathbf{W}_{x\alpha}\bar{\mathbf{x}}_t) \\ \mathbf{z}_t = \mathbf{h}_{t-1} + \sum_{\mathbf{r}_k \in \mathbf{R}_{\text{KG}}} \alpha_{t,k} \mathbf{r}_k \quad (7)$$

where $\alpha_t \in \mathbb{R}^{|\mathbf{R}_{\text{KG}}|}$ is an attention vector over the relations space, \mathbf{r}_k is relation embeddings, and \mathbf{z}_t is a resulting entity context vector after walking from its previous entity on an attended path.

We guide the graph decoder with the ground-truth walk paths by computing the following loss $\mathcal{L}_{\text{walk}}(\mathbf{x}, \mathbf{y}) = \sum_{i,t} \mathcal{L}_{\text{ent}} + \mathcal{L}_{\text{rel}}$ between predicted paths and each of $\{\mathbf{y}_e, \mathbf{y}_r\}$, respectively (\mathcal{L}_{ent} : loss for entity paths, and \mathcal{L}_{rel} for relation paths):

$$\sum_{\tilde{\mathbf{y}}_e \neq \mathbf{y}_{e,t}^{(i)}} \max[0, \tilde{\mathbf{y}}_e \cdot \mathbf{y}_{e,t}^{(i)} - \mathbf{h}_t^{(i)} \cdot (\mathbf{y}_{e,t}^{(i)} - \tilde{\mathbf{y}}_e)^\top] \\ + \sum_{\tilde{\mathbf{y}}_r \neq \mathbf{y}_{r,t}^{(i)}} \max[0, \tilde{\mathbf{y}}_r \cdot \mathbf{y}_{r,t}^{(i)} - \alpha_t \mathbf{r} \cdot (\mathbf{y}_{r,t}^{(i)} - \tilde{\mathbf{y}}_r)^\top]$$

Once the model is trained, at each decoding step, we can rank the potential paths based on the sum of their zeroshot relevance (left) and soft-attention-based output path (right) scores:

$$\mathbf{y}_{e,t}^{(i)} = \operatorname{argmax}_{\mathbf{y}_e^{(i)} \in \mathbf{V}_{R,1}(\mathbf{y}_{e,t-1}^{(i)})} \mathbf{h}_t \cdot \mathbf{y}_e^{(i)\top} + \sum \alpha_{t,k} \mathbf{r}_k \cdot \mathbf{y}_r^{(i)\top} \quad (8)$$

Adversarial Transfer Learning: if domain labels (\mathbf{y}_d) are available (e.g. movie, book, sports, etc.), we can utilize these labels to further aid training by extracting transferrable features and learning optimal paths conditioned on domain embeddings (Ganin et al., 2016). We implement adversarial transfer learning for DialKG Walker as follows and study this specific setting in one of our experiments to demonstrate that the model can better generalize over multiple domains:

$$\mathcal{L} = \mathcal{L}_f + \mathcal{L}_{\text{walk}} + \text{Entropy}(\sigma(\mathbf{W}_d \mathbf{x}), \mathbf{y}_d) \\ \mathbf{h}_t = \text{WALK}([\bar{\mathbf{x}}; (\mathbf{W}_d \mathbf{x})], \mathbf{z}_t) \quad (9)$$

3 Dataset: OpenDialKG

To empirically evaluate the proposed approach, we collected a new dataset, **OpenDialKG**, of chat conversations between two agents engaging in a dialog about a given topic (91K turns across 15K dialog sessions). Each dialog is paired with its corresponding “KG paths” that weave together the KG entities and relations that are mentioned in the dialog. This parallel corpus of textual dialogs and corresponding KG walks enables learning models that ground the implicit reasoning in human conversations to discrete KG operations.

Wizard-of-Oz setup The dialogs were generated in a Wizard-of-Oz setting (Shah et al., 2018) by connecting two crowd-workers to engage in a chat session, with the joint goal of creating natural and engaging dialogs. The first agent is given a seed entity and asked to initiate a conversation about that entity. The second agent is provided with a list of facts relevant to that entity, and asked to choose the most natural and relevant facts and use them to frame a free-form conversational response. Each fact is a 1-hop or 2-hop path initiating from the conversation topic. After the second agent sends their response, various new multi-hop facts from KG are surfaced to include paths initiating from new entities introduced in the latest message. This process allows the conversation participants to annotate any new fact or entity they want to introduce at each turn, along with the ground-truth KG walk path that connect the two KG entities. At this point the first agent is instructed to continue the conversation by choosing among the updated set of facts and framing a new message. This cycle continues for 6 messages per session on average spanning multiple KG paths, until one of the agents decides to end the conversation (*e.g.* the task goal is met).

We did two separate collections: a **recommendation** task where the second agent acts as an assistant who is providing useful recommendations to the user, and a **chit-chat** task where both agents act as users engaging in open-ended chat about a particular topic. To ensure sufficient separation of the dialog content, we used entities related to **movies** (titles, actors, directors) and **books** (titles, authors) for the recommendation task, and entities related to **sports** (athletes, teams) and **music** (singers) for the chit-chat task (Table 1). Seed entities for each domain are crawled from various public resources (*e.g.* IMDB top movies list, top

Task:	Recommendation		Chit-chat		(All)
Domain:	Movies	Books	Sports	Music	
# of dialogs	6,429	5,891	2,495	858	15,673
# of turns	37,838	34,035	14,344	4,992	91,209

Table 1: Task / domain distribution of *OpenDialKG*.

athletes list, etc.) and linked with the corresponding KG entities.

KG sources: We use the Freebase (Bast et al., 2014) KG which is a publicly available and comprehensive source of general-knowledge facts. To reduce noise, we filter tail-end entities based on their prominence scores, the resulting KG of which consists of total 1,190,658 fact triples over top 100,813 entities and 1,358 relations.

We randomly split the dialog sessions into train (70%), validation (15%), and test sets (15%).

4 Empirical Evaluation

Task: Given a set of KG entity mentions from current turn, and dialog history of all current and previous sentences, the goal is to build a robust model that can retrieve a set of natural entities to mention from a large-scale KG that resemble human responses. Note that end-to-end generation of sentences (*e.g.* based on the retrieved entities) is not part of this study - instead, we focus on the important challenge of scaling the conversational reasoning and knowledge retrieval task to open-domain dialogs, requiring an aggressive subset selection (from 1M+ facts subset of Freebase).

4.1 Baselines

We choose as baselines the following state-of-the-art approaches that augment external knowledge to dialog systems for various tasks (see Section 5 for details), and modify accordingly to fit to our entity retrieval task (*e.g.* we use the same 1M-facts FreeBase KG for all of the baselines):

- **seq2seq** (Sutskever et al., 2014) with dialog contexts + zeroshot: we apply the seq2seq approach for entity path generation, given all of the dialog contexts. To make this baseline stronger, we add a zeroshot learning layer in the KG embeddings space (replacing typical softmax layers to improve generality) for entity token decoding.
- **Tri-LSTM** (Young et al., 2018): encodes each utterance and all of its related facts

Input	Model	All Domains → All					Movie → Movie				
		r@1	3	5	10	25	r@1	3	5	10	25
E+S+D	seq2seq (Sutskever et al., 2014)	3.1	18.3	29.7	44.1	60.2	3.0	13.4	23.4	38.5	55.5
E+S	Tri-LSTM (Young et al., 2018)	3.2	14.2	22.6	36.3	56.2	1.5	10.3	17.4	30.7	51.1
E+S	Ext-ED (Parthasarathi and Pineau, 2018)	1.9	5.8	9.0	13.3	19.0	1.3	5.4	7.8	11.8	15.8
E	DialKG Walker (ablation)	10.7	22.9	32.0	44.9	57.4	5.3	13.5	18.5	25.2	39.1
E+S	DialKG Walker (ablation)	11.3	23.3	31.0	44.0	60.5	7.2	19.2	27.9	40.7	58.7
E+S+D	DialKG Walker (proposed)	13.2	26.1	35.3	47.9	62.2	7.8	20.0	27.9	40.4	58.6

Table 2: In-domain (train/test on the same domain) response generation performance on the *OpenDialKG* dataset (metric: recall@ k). Our proposed model is compared against state-of-the-art models as well as several ablation variations of the proposed model. All of the 100K+ KG entities are considered initial candidates for generation (before masking). E: entities, S: sentence, D: dialog contexts.

Input	Model	Movie → Book					Movie → Music				
		r@1	3	5	10	25	r@1	3	5	10	25
E+S+D	seq2seq (Sutskever et al., 2014)	2.9	21.3	35.1	50.6	64.2	1.5	12.1	19.7	34.9	49.4
E+S	Tri-LSTM (Young et al., 2018)	2.3	17.9	29.7	44.9	61.0	1.9	8.7	12.9	25.8	44.4
E+S	Ext-ED (Parthasarathi and Pineau, 2018)	2.0	7.9	11.2	16.4	22.4	1.3	2.6	3.8	4.1	8.3
E	DialKG Walker (ablation)	8.2	15.7	22.8	31.8	48.9	4.5	16.7	21.6	25.8	33.0
E+S	DialKG Walker (ablation)	12.6	28.6	38.6	54.1	65.6	6.0	15.9	22.8	33.0	47.5
E+S+D	DialKG Walker (proposed)	13.5	28.8	39.5	52.6	64.8	5.3	13.3	19.7	28.8	38.0

Table 3: Cross-domain (train/test on the different domain) response generation performance on the *OpenDialKG* dataset (metric: recall@ k). E: entities, S: sentence, D: dialog contexts.

within 1-hop from a KG to retrieve a response from a small (N=10) pre-defined sentence bank. We modify the retrieval bank to be the facts from the KG instead.

- **Extended Enc-Dec (Parthasarathi and Pineau, 2018)**: conditions response generation with external knowledge vector input. A response entity token is generated at its final softmax layer, hence not utilizing structural information from KG.

We also consider several configurations of our proposed approach to examine contributions of each component (input modalities (E): entities, (S): sentence, (D): dialog contexts).

- **(Proposed; E+S+D)**: is the proposed approach as described in Figure 2
- **(E+S)**: relies only on its previous sentence and excludes dialog history from input.
- **(E)**: only uses starting KG entities as input contexts, and excludes any textual context.

4.2 Results

Parameters: We tune the parameters of each model with the following search space (bold in-

dicating the choice for our final model): KG embeddings size: {64, **128**, 256, 512}, LSTM hidden states: {64, **128**, 256, 512}, word embeddings size: {100, 200, **300**}, max dialog window size: {2, **3**, 4, 5}. We optimize the parameters with Adagrad (Duchi et al., 2011) with batch size 10, learning rate 0.01, epsilon 10^{-8} , and decay 0.1.

In-domain evaluation: Table 2 shows the generation results of the top- k predictions of the model for in-domain train and test pairs (train & test on: all domains / train & test on: movie domain split). It can be seen that the proposed DialKG Walker model outperforms other state-of-the-art baselines, especially for recalls at small k s. Specifically, when textual contexts are added as input (E+S and E+S+D), the model learns to condition its walk path output on textual contexts, thus outperforming the non-textual ablation model (E). seq2seq and Tri-LSTM models consider the nodes connected via all possible relations as candidates in the final layer (without pruning), resulting in extensive search space and consequently poor recall performance. In addition, Tri-LSTM only considers the facts connected via 1-hop relations as input contexts, which limits its prediction for multi-hop facts. Ext-ED relies its prediction in

Input Dialog (start entity)	Response		
	Model	Walk Path	Predicted Entity
A: <i>Yes, I believe he [Muller] has played in Munich.</i>	GT	award won by → position	Forward
B: <i>He also won a <u>Bravo Award</u>. I think that's awesome!</i>	KG-Walker	award won by	Lionel Messi
A: [response]	Ext-ED	award won by	Muller
A: <i>Could you recommend a book by <u>Mark Overstall</u>?</i>	GT	wrote → has genre	Romance
B: [response]	KG-Walker	wrote → has genre	Romance
	Ext-ED	language	English
A: <i>Do you like <u>Lauren Oliver</u>. I think her books are great!</i>	GT	written by → wrote	Requiem
B: <i>I do, <u>Vanishing Girls</u> is one of my favorite books.</i>	KG-Walker	written by → wrote	Annabel
A: [response]	Tri-LSTM	released year	2015
A: <i>What about the <u>Oakland Raiders</u>?</i>	GT	Champion	Packers
B: <i>Oh yes, I do like them. I've been a fan since they were runner-up in Super Bowl II. What about you? // A: [response]</i>	KG-Walker	Champion	Packers
	seq2seq	Runner-up → Is_A	NFL Team
A: <i>Do you like <u>David Guetta</u>? I enjoy his music.</i>	GT	composer → composed	Club Can't Handle Me
B: <i>Oh, I love his lyrics to <u>Love is Gone</u> and the song <u>Wild Ones</u>. What are your favorites? // A: [response]</i>	KG-Walker	composer → composed	I Love It
	Tri-LSTM	composer	David Guetta

Table 4: **Error analysis:** DialKG Walker with attention (ours) vs. baselines. Ground-truth response (GT) and model predictions of walk paths and future entities for the underlined entity mentions are shown. Dialogs are only partially shown due to space constraints.

Model	% in top- k		
	$k=1$	$k=2$	$k=3$
(Parthasarathi and Pineau, 2018)	17.5	33.6	47.2
(Young et al., 2018)	30.8	50.1	70.3
(Sutskever et al., 2014)	31.5	57.7	73.1
KG Walker (proposed)	38.6	61.8	76.3

Table 5: **Human evaluation:** “Which response is the most natural for given dialog context?” (metric: % of cases chosen as top- k response by the raters)

the final softmax layer, which typically performs poorly for a large number of output class, compared to zeroshot learning approaches.

Cross-domain evaluation: Table 3 demonstrates that the DialKG Walker model can generalize to multiple domains better than the baseline approaches (train: movie & test: book / train: movie & test: music). This result indicates that our method also allows for zeroshot pruning by relations based on their proximity in the KG embeddings space, thus effective in cross-domain cases as well. For example, relations ‘scenario by’ and ‘author’ are close neighbors in the KG embeddings space, thus allowing for zeroshot prediction in cross-domain tests, although their training examples usually appear in two separate domains: movie and book.

Human evaluation: To compare the subjective quality of the models, *i.e.* the relative naturalness and relevance of the generated KG paths, we performed a human evaluation where paid raters were




Figure 3: **Transfer learning** results ($r@5$) of DialKG Walker at varying availability of target data with (a) Book and (b) Sports domains as a Target (Source: Movie). (TL:Adv): data transfer with adversarial discriminator for source and target domains, (TL:FT): model transfer with fine-tuning, (No-TL): target only.

shown partial dialogs taken from the test dataset, along with the top 2 paths output from each model. The rater was asked to choose the 3 most appropriate paths for continuing the dialog. We evaluated 250 dialogs, showing each dialog to 3 raters, for a total 750 tasks. We report the % of cases when a top- k chosen fact was generated by each of the models (Table 5). The numbers add up to more than 100% as models can generate identical paths. If such a path is chosen by the rater, it is counted towards each of the models that generated the path.

We show that the generated responses by our proposed methods achieve the highest scores in all top- k evaluation, validating that the model can output more natural human-like responses.

Transfer learning: In Figure 3, we show that cross-domain performance can greatly improve with a relatively small addition of in-domain target data, via the transfer learning approaches. Specifically, it can be seen that (TL:Adv), which simultaneously trains for both source and target data (effectively doubling the training size) with additional adversarial discriminator for source and target domains, achieves the best performance especially for domains that are semantically close (*e.g.* movie and book). (TL:FT) transfers knowledge from a pre-trained source model via fine-tuning (hence requiring significantly less training resources), and effectively avoids “cold start” training (Moon et al., 2015). This result shows that the DialKG model can quickly adapt to other new low-resource domains and improve upon the zero-shot cross-domain performance, demonstrating its potential capability to reason on open-ended conversations.

Error analysis: Table 4 shows some of the example output from each model (as well as ground-truth responses), given dialog contexts. In general, the DialKG Walker tends to explore more multi-hop relations than other baselines in order to generate natural and engaging entities, which consequently improves the diversity of answers. Note that if the graph decoder arrives at a sufficiently good entity to generate, it stops its traversal operation and outputs the most viable entity based on the relevance score. Some of the models do not take into account the dialog history, hence generating redundant topics from previous turns. There are some cases where the final entity prediction is different from the ground-truth, whereas its relation path is correctly predicted. The generated entities are often still considered valid and natural, because the proposed model uses zero-shot relevance score to best predict the candidates.

5 Discussion and Related Work

Knowledge augmented dialog systems: Young et al. (2018) propose to explicitly augment input text with concepts expanded via 1-hop relations (where KG triples are represented in the sentence embeddings space), and He et al. (2017) propose a system which iteratively updates KG embeddings and attends over connected entities for response generation. However, several challenges remain to scale the simulated knowledge graph used in the study to our open-ended and large-

scale KG with 1M+ facts. Other line of work (Parthasarathi and Pineau, 2018; Ghazvininejad et al., 2018; Long et al., 2017) uses embedding vectors obtained from external knowledge sources (*e.g.* NELL (Carlson et al., 2010), Wikipedia, Freebase (Bast et al., 2014), free-form text, etc.) as an auxiliary input to the model in dialog generation. Our model extends the previous work by (1) explicitly modeling output reasoning paths in a structured KG, (2) by introducing an attention-based multi-hop concept decoder to improve both recall and precision.

End-to-end dialog systems: Several models and corresponding datasets have recently been published. Most work focuses on task or goal oriented dialog systems such as conversational recommendations (Salem et al., 2014; Bordes et al., 2017; Sun and Zhang, 2018; Dalton, 2018), information querying (Williams et al., 2017; de Vries et al., 2018; Reddy et al., 2018), etc., with datasets collected mostly through bootstrapped simulations (Bordes et al., 2017), Wizard-of-Oz setup (Zhang et al., 2018; Wei et al., 2018), or online corpus (Li et al., 2016). Our OpenDialKG corpus is unique in that it includes open-ended natural human conversations over multiple scenarios (*e.g.* chit-chat and recommendation on various domains), where reasoning paths from each dialog are annotated with their corresponding discrete KG operations. Our work can also be viewed as extending the conventional state-tracking approaches (Henderson et al., 2014) to more flexible KG path as states.

KG embeddings and inference: Several methods have been proposed for KG inference tasks (*e.g.* edge prediction), which include neural models trained to discern positive and negative triples (Bordes et al., 2013; Wang et al., 2014; Nickel et al., 2016; Dettmers et al., 2018), or algorithms with discrete KG operations on structured data (Lao et al., 2011; Chen et al., 2015). KG embeddings have been shown effective in other NLP tasks when they are used as target labels for classification tasks, which also allows for effective transfer learning (Moon and Carbonell, 2017). For effective application of KG embeddings in NLP tasks, recent studies (Kartsaklis et al., 2018) proposed to map word embeddings and KG embeddings via end-to-end tasks. In contrast to the line of work on KG edge prediction, we aim to learn an optimal path within existing paths that resemble human reasoning in conversations.

6 Conclusions

We study conversational reasoning grounded on knowledge graphs, and formulate an approach in which the model learns to navigate a large-scale, open-ended KG given conversational contexts. For this study, we collect a newly annotated Dialog \leftrightarrow KG parallel corpus of 15K human-to-human dialogs which includes ground-truth annotation of each dialog turn to its reasoning reference in a large-scale common fact KG. Our proposed DialKG Walker model improves upon the state-of-the-art knowledge-augmented conversation models by 1) a novel attention-based *graph decoder* that penalizes decoding of unnatural paths which effectively prunes candidate entities and paths from a large search space (1.1M facts), 2) a zeroshot learning model that predicts its relevance score in the KG embeddings space, combined score of which is used for candidate ranking. The empirical results from in-domain, cross-domain, and transfer learning evaluation demonstrate the efficacy of the proposed model in domain-agnostic conversational reasoning.

References

- Hannah Bast, Florian Baurle, Bjorn Buchhold, and Elmar Haussmann. 2014. Easy access to the freebase dataset. In *WWW*.
- Antoine Bordes, Y-Lan Boureau, and Jason Weston. 2017. Learning end-to-end goal-oriented dialog. *ICLR*.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *NIPS*.
- Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam R Hruschka, and Tom M Mitchell. 2010. Toward an architecture for never-ending language learning. In *AAAI*.
- CY Chen, D Yu, W Wen, YM Yang, J Zhang, M Zhou, K Jesse, A Chau, A Bhowmick, S Iyer, G Sreenivasulu, R Cheng, A Bhandare, and Z Yu. 2018. Gunrock: Building a human-like social bot by leveraging large scale real user data. In *2nd Alexa Prize*.
- Yun-Nung Chen, William Yang Wang, and Alexander Rudnicky. 2015. Jointly modeling inter-slot relations by random walk on knowledge graphs for unsupervised spoken language understanding. In *NAACL*.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. Supervised learning of universal sentence representations from natural language inference data. In *EMNLP*.
- Jeff Dalton. 2018. Vote goat: Conversational movie recommendation. *SIGIR*.
- Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. 2018. Convolutional 2d knowledge graph embeddings. In *AAAI*.
- John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *JMLR*.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *JMLR*.
- Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Scott Wen-tau Yih, and Michel Galley. 2018. A knowledge-grounded neural conversation model. In *AAAI*.
- He He, Anusha Balakrishnan, Mihail Eric, and Percy Liang. 2017. Learning symmetric collaborative dialogue agents with dynamic knowledge graph embeddings. *ACL*.
- Matthew Henderson, Blaise Thomson, and Jason D Williams. 2014. The second dialog state tracking challenge. In *SIGDIAL*.
- Dimitri Kartsaklis, Mohammad Taher Pilehvar, and Nigel Collier. 2018. Mapping text to knowledge graph entities using multi-sense lstms. *EMNLP*.
- Ni Lao, Tom Mitchell, and William W. Cohen. 2011. Random walk inference and learning in a large scale knowledge base. In *EMNLP*.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios P Spithourakis, Jianfeng Gao, and Bill Dolan. 2016. A persona-based neural conversation model. *ACL*.
- Yinong Long, Jianan Wang, Zhen Xu, Zongsheng Wang, Baoxun Wang, and Zhuoran Wang. 2017. A knowledge enhanced generative conversational service agent. In *NIPS DSTC6 Workshop*.
- A. H. Miller, W. Feng, A. Fisch, J. Lu, D. Batra, A. Bordes, D. Parikh, and J. Weston. 2017. Parlai: A dialog research software platform. *EMNLP*.
- Seungwhan Moon and Jaime Carbonell. 2017. Completely heterogeneous transfer learning with attention: What and what not to transfer. *IJCAI*.
- Seungwhan Moon, Suyoun Kim, and Haohan Wang. 2015. Multimodal transfer deep learning with applications in audio-visual recognition. In *NIPS MMML Workshop*.
- Seungwhan Moon, Leonard Neves, and Vitor Carvalho. 2018a. Multimodal named entity recognition for short social media posts. *NAACL*.

- Seungwhan Moon, Leonard Neves, and Vitor Carvalho. 2018b. Zeroshot multimodal named entity disambiguation for noisy social media posts. *ACL*.
- Maximilian Nickel, Lorenzo Rosasco, and Tomaso Poggio. 2016. Holographic embeddings of knowledge graphs. *AAAI*.
- Prasanna Parthasarathi and Joelle Pineau. 2018. Extending neural generative conversational model using external knowledge sources. *EMNLP*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *EMNLP*.
- Siva Reddy, Danqi Chen, and Christopher D Manning. 2018. Coqa: A conversational question answering challenge. *arXiv preprint arXiv:1808.07042*.
- Yasser Salem, Jun Hong, and Weiru Liu. 2014. History-guided conversational recommendation. In *WWW*.
- Pararth Shah, Dilek Hakkani-Tur, Bing Liu, and Gokhan Tur. 2018. Bootstrapping a neural conversational agent with dialogue self-play, crowdsourcing and on-line reinforcement learning. In *NAACL*.
- Yueming Sun and Yi Zhang. 2018. Conversational recommender system. *SIGIR*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *NIPS*.
- Harm de Vries, Kurt Shuster, Dhruv Batra, Devi Parikh, Jason Weston, and Douwe Kiela. 2018. Talk the walk: Navigating new york city through grounded dialogue. *ECCV*.
- Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. Knowledge graph embedding by translating on hyperplanes. In *AAAI*.
- Wei Wei, Quoc Le, Andrew Dai, and Jia Li. 2018. Airdialogue: An environment for goal-oriented dialogue research. In *EMNLP*.
- Jason D Williams, Kavosh Asadi, and Geoffrey Zweig. 2017. Hybrid code networks: practical and efficient end-to-end dialog control with supervised and reinforcement learning. *ACL*.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical attention networks for document classification. In *NAACL*.
- Tom Young, Erik Cambria, Iti Chaturvedi, Minlie Huang, Hao Zhou, and Subham Biswas. 2018. Augmenting end-to-end dialog systems with common-sense knowledge. *AAAI*.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. Personalizing dialogue agents: I have a dog, do you have pets too? *ACL*.