

CorefDiffs: Co-referential and Differential Knowledge Flow in Document Grounded Conversations

Lin Xu¹, Qixian Zhou², Jinlan Fu¹, Min-Yen Kan¹, See-Kiong Ng¹

¹National University of Singapore, ²ByteDance
{cathyxl2016, qixianzhou.mail, jinlanjonna}@gmail.com
{knmny, seekiong}@nus.edu.sg

Abstract

Knowledge-grounded dialog systems need to incorporate smooth transitions among knowledge selected for generating responses, to ensure that dialog flows naturally. For document-grounded dialog systems, the inter- and intra-document knowledge relations can be used to model such conversational flows. We develop a novel Multi-Document Co-Referential Graph (Coref-MDG) to effectively capture the inter-document relationships based on commonsense and similarity and the intra-document co-referential structures of knowledge segments within the grounding documents. We propose CorefDiffs, a Co-referential and Differential flow management method, to linearize the static Coref-MDG into conversational sequence logic. CorefDiffs performs knowledge selection by accounting for contextual graph structures and the knowledge difference sequences. CorefDiffs significantly outperforms the state-of-the-art by 9.5%, 7.4% and 8.2% on three public benchmarks. This demonstrates that the effective modeling of co-reference and knowledge difference for dialog flows are critical for transitions in document-grounded conversation¹.

1 Introduction

Document-grounded conversations (Moghe et al., 2018; Dinan et al., 2018; Feng et al., 2021b) is a core class of knowledge-grounded dialogs that leverage text-based knowledge segments from documents to generate informative dialog responses. This task is typically divided into two sub-tasks, given the dialog history (Dinan et al., 2018): namely, knowledge selection and response generation. Knowledge selection, which determines the content of the generated responses (Moghe et al., 2018; Dinan et al., 2018), is the crucial sub-task for dialog flow management as it leads to the manifestation of knowledge transition (Meng et al., 2020), essential for naturalistic engaging conversations.

¹The source code has been released at <https://github.com/cathyxl/coref-diffs>

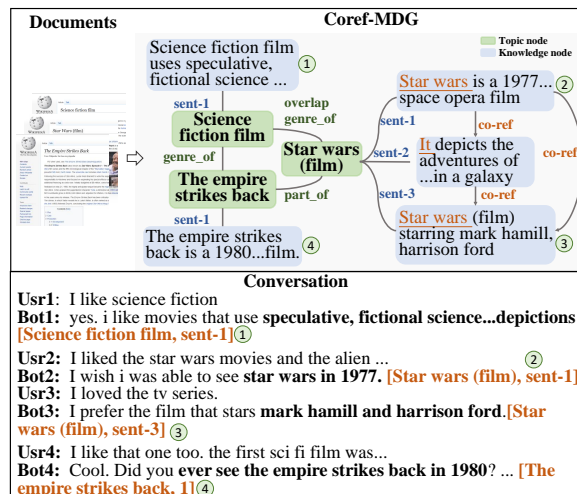


Figure 1: Co-Referential Multi-Document Graph (Coref-MDG). Topic vertices correspond to documents and are connected by commonsense/word overlap relations. Knowledge vertices are connected with its topic vertex by its document sentence index, e.g. *sent-1*, and connected to each other by co-reference (co-ref) relations. The Bot’s utterances are followed by its topic and knowledge segment, e.g. [Science fiction film, sent-1].

Most existing studies on document-grounded conversations (Lian et al., 2019; Zheng et al., 2020; Zhao et al., 2020) treat knowledge selection as a matching problem between the dialog context and individual knowledge segments, independently. However, for document-grounded conversations, we posit that there is an implicit alignment between the background knowledge and conversation logic which can be learned from the underlying structural relationships of the knowledge segments within and between the grounding documents. For example, the conversation in Figure 1 exhibits document-level topic flow, from *science fiction* -> *star wars* -> *the empire strikes back*, and deep dives into the specifics of the *star wars* document (Turns 2 to 3).

To effectively exploit the relationships of the knowledge segments to guide dialog flows would require a thorough comprehension of the intra-




Figure 2: Exploiting knowledge segment relationships. k1-7 represent knowledge segments and doc1-3 are the grounding documents they belong to.

document discourse structures and inter-document relationships for the knowledge selection process. Existing works either ignore such relations (as illustrated in Figure 2 (a)), or exploit limited local correlations (as depicted in Figure 2 (b)), for example by encoding knowledge segments within passage context Wu et al. (2021). In this work, we propose to capture both intra- and inter-document relationships of the knowledge segments (Figure 2 (c)) in the grounding documents to guide the smooth and natural knowledge selection and transitions for document-grounded conversations. However, how to apply such a static knowledge graph to dialog flow management has always been a problem. Many previous studies (Moon et al., 2019; Xu et al., 2021a,b) have used graph structures to constrain search (e.g. confining the next topic to neighboring areas), but have also ignored deeper integration of dialog contexts and knowledge graphs, such as optimal knowledge representation to capture dialogue flow information.

Based on the considerations above, we propose to first capture the inter- and intra-document knowledge relationships as a heterogeneous document graph, and then exploit the graph effectively for dialog flow management through fine-to-coarse contextualization — from the local word-level knowledge attentions, to knowledge interactions in document graphs, and finally to the knowledge transition flow along dialogue turns. Specifically, we design a two-level document graph consisting of topic (i.e. document) and knowledge vertices connected by inter- and intra-document relations (Figure 1). The topic vertices correspond one-to-one to the grounding documents, while the knowledge vertices refer to the knowledge segments from each document. The knowledge vertices are connected to the corresponding topic vertices they belong to. Meanwhile, the graph connects the knowledge segments within the same document by their co-referential mentions, and the documents are connected based on similarity or commonsense

relationships. Hence we call the graph Multi-Document Co-referential Graph (Coref-MDG).

We then propose our CorefDiffs method which leverages Coref-MDG’s graph structure and integrates dialog flow for knowledge contextualization and selection. CorefDiffs focuses on the inter-turn knowledge difference flow in the dialog histories by means of a novel differential linearization module.

Our contributions in this paper can be summarized as follows. 1) We develop Coref-MDG, a novel multi-document graph structure incorporating co-referential mentions. When leveraged in guiding document-grounded conversations in our CorefDiffs methodology, it empirically outperformed alternative graph structures; 2) We systematically study the different kinds of inter- and intra-document relations and show that document-level semantics, such as co-reference and sentence order, are significant factors for knowledge selection (Sec. 4.4); 3) Our CorefDiffs achieves state-of-the-art on WoW, Holl-E, multidoc2dial and CMU-DOG datasets, for both knowledge selection and response generation tasks.

2 Related Work

Document-grounded dialog Systems. Early works on document-grounded dialog systems (Ghazvininejad et al., 2018) focused on generating responses directly by copying words from the external documents. The subsequent availability of datasets with knowledge annotations (Dinan et al., 2018; Moghe et al., 2018) has led to the separation of the tasks of knowledge selection and response generation. For knowledge selection, most works (Dinan et al., 2018; Lian et al., 2019; Zheng et al., 2020; Zhao et al., 2020; Meng et al., 2021) in document-grounded conversations directly modeled correlations between dialog contexts and knowledge through independent matching and optimized the correlations by modeling knowledge sequence (Kim et al., 2019), increasing knowledge informativeness (Zheng et al., 2020) or distinguishing initiative roles (Meng et al., 2021). A recent work (Wu et al., 2021) boosted knowledge selection by encoding knowledge within the passage context, which demonstrates the importance of exploiting knowledge relations. Our work further explores more effective document structures and connections for this task. There is also an unpublished paper that used document

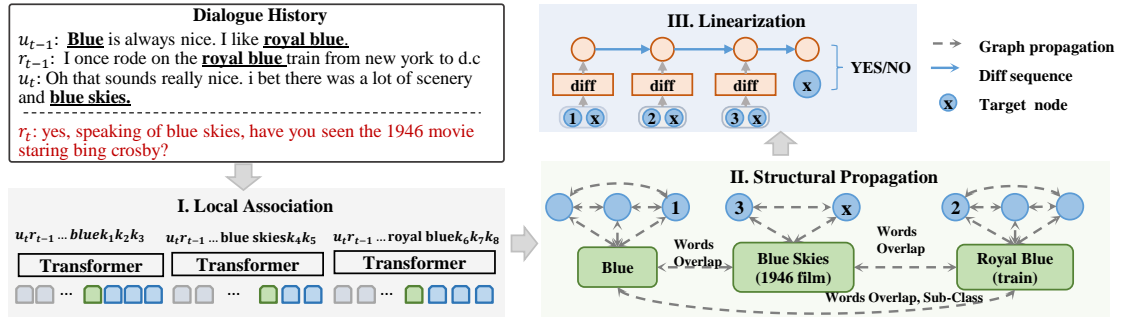


Figure 3: CorefDiffs Architecture. Green round rectangles and blue circles are topic and knowledge vertices, respectively. Steps I to III contextualize knowledge in an increasingly fine-grained manner: first, I) by vertex embedding by BERT; then, II) by propagating Coref-MDG information, and finally III) by linearizing the knowledge representations, according to the dialog’s historical knowledge sequence.

semantic graphs (Li et al., 2022), while our work considers end-to-end integration of document graph and dialog flow which gives better result compared to theirs.

Knowledge Graph for Conversations. Knowledge graphs were also often used in dialog management, such as dialog transition graphs (Xu et al., 2019, 2020a) constructed from common transitions present in a dialog corpus and off-the-shelf commonsense graphs (Zhou et al., 2018a). There were also some works (Liu et al., 2019; Xu et al., 2021a) transforming unstructured text into structures or combining triplets and texts into graphs. For example, (Xu et al., 2021a) constructed key phrases into graphs according to their order in stories. Interestingly, to the best of our knowledge, co-reference mentions have not been considered in such document graph construction although it has been proved critical in learning language models for reasoning intensive NLP tasks (Dasigi et al., 2019; Ye et al., 2020). To apply knowledge graphs for dialog, many existing works (Xu et al., 2020b,a) used the graph structures to confine the search space and optimized selection through hand-crafted rewards. In contrast, we incorporate the knowledge graph into dialog management by learning knowledge representations from the graph structure.

Sequence Learning in dialog. Sequence learning is essential for conversations. Several studies (Kim et al., 2019; Zhan et al., 2021b) explored the historical knowledge sequence to select knowledge for document-grounded dialog. For example, (Kim et al., 2019) captured knowledge sequence by a latent variable, while (Zhan et al., 2021b) further proposed to learn abstract topic sequence to mitigate the issues of knowledge sparsity and knowl-

edge transition noise. Inspired by the importance of exploiting knowledge difference (Zheng et al., 2020) for informative dialog, we extend the use of dialog knowledge differences into sequences, thus capturing the knowledge shift patterns from turns with longer distances as well as the sequential patterns of knowledge transitions in a dialogue.

3 Approach

Figure 3 shows the overall architecture of our approach. As shown in the Dialog History part, in each data sample, given a dialog history $U = \{u_{t-l}, r_{t-l}, \dots, r_{t-1}, u_t\}$ of l turns and a set of grounding documents $\mathcal{D} = \{d_1, \dots, d_i, \dots, d_{|\mathcal{D}|}\}$, where u_* and r_* are utterances from the user and chatbot, respectively. $d_i = \{k_1^i, k_2^i, \dots, k_{|d_i|}^i\}$ is a document containing a bunch of knowledge segments, our task is to select the most appropriate knowledge segment from the grounding documents \mathcal{D} (i.e. the knowledge selection subtask) and generate the chatbot’s next response r_t based on the selected knowledge (i.e. the response generation subtask). Each grounding document d_i has a phrase t_i as its topic. For example, the document of wikipedia blue has the topic phrase *blue*.

3.1 Coref-MDG Construction

We devise a Multi-Document Co-referential Graph (Coref-MDG) to capture the inter-document and the intra-document relations. Each data sample gets a specific Coref-MDG, denoted as $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, where \mathcal{V}, \mathcal{E} are vertices and edges respectively.

Vertices \mathcal{V} . Our Coref-MDG consists of two types of Vertices: topic and knowledge vertices, as shown in Figure 1. Each topic vertex r_t represents one of document d_i from \mathcal{D} while each knowledge vertex

represents a knowledge segment k_j^i from a document d_i , hence in total $M = |\mathcal{D}|$ topic vertices and $N = |d_1| + \dots + |d_{|\mathcal{D}|}|$ knowledge vertices.

Edges \mathcal{E} . There are also multiple types of edges in Coref-MDG. We can generally divide these edges into three categories according to their vertices: **1) edges between topics and knowledge vertices;** **2) topic edges**—these are inter-document or inter-topic edges between topic vertices; **3) knowledge edges**—intra-topic edges amongst knowledge vertices. For the first category, we simply use the order index of segment k_j^i appearing in its corresponding document d_i as the edge type, denoted as `sent_j` edge, thus knowledge vertices under different topics are not connected in Coref-MDG. The remaining two types of edges are constructed as follows.

3.1.1 Topic Edges

We posit that topic transitions in human-to-human conversations are likely to be based on the similarity or commonsense relations between two topics, such as from *sci-fi movie* to *sci-fi novel* (similarity), or from *UK* to *London* (commonsense). We introduce two corresponding types of topic edges for such topic transitions.

Word Overlap. We use the word overlaps between two topics (or documents) to measure their similarity. Specifically, we obtain the lemmas of topic phrases by spaCy² and judge whether the two topics have at least one identical lemma so as to determine whether these two topics vertices have a `word_overlap` edge.

Commonsense. Since the knowledge backend of the WoW (Dinan et al., 2019) came from the Wikipedia corpus, we use the WikiData³ to obtain commonsense relations between topics. We only collected relations for the topics in the training set and for simplicity, we kept the high-frequency relation types, for example `city_of`, while uniformly treating the low-frequency relation types as `others`.

3.1.2 Knowledge Edges

For the intra-document knowledge relations, we introduce the `coreference_link` edge. For each topic (i.e. document), the co-reference links (referring paths) within the corresponding document d_i can be extracted by a co-reference res-

olution model.⁴ For each co-reference link, every knowledge segment on this link is connected to its mentions by a `coreference_link` edge. Aside from our proposed co-reference edges, we also model two other knowledge edge type for comparison, `common_entity` and `partial_order`. The former connects knowledge segments that share entities, while the latter captures knowledge segment’s partial order. We will show later that co-reference performs best for dialog flow management.

3.2 Structural Propagation and Linearization

Next, we introduce how we contextualize each vertex in a dialog’s Coref-MDG with both the graph structure and dialog flow.

3.2.1 Node Initialization

Following (Karpukhin et al., 2020; Cheng et al., 2020; Wu et al., 2021), we adopt BERT (Devlin et al., 2019) to obtain the text representations to initialize topic and knowledge vertices, as shown by the Step I in Figure 3. Specifically, we concatenate the dialog context U with each grounding document’s topic phrase and knowledge segments, and feed them into the BERT encoder to get their associated representations. The concatenated input for a document d_i is thus:

$$[\text{cls}]\hat{U}_t[\text{sep}]t_i[\text{cls}]k_1^i\dots[\text{cls}]k_{|d_i|}^i[\text{sep}] \quad (1)$$

where $\hat{U} = [\text{usr}]u_t[\text{agt}]r_{t-1}\dots[\text{usr}]u_{t-l}$ is the spliced dialog context, and the role symbols `[usr]` and `[agt]` indicate utterances from the user or agent turn. We use the hidden state of the first `[cls]` token \mathbf{t}_i (note that we use **bold** here to refer to the representation of, in this case, topic phrase t_i) as the initialized representation for the topic vertex of d_i . Similarly, the outputs of the subsequent `[cls]` tokens, denoted as $\{\mathbf{k}_1^i, \mathbf{k}_2^i, \dots, \mathbf{k}_{|d_i|}^i\}$, are gathered and used to initialize the corresponding knowledge vertices of d_i . The process is formulated as:

$$\mathbf{t}_i, \mathbf{K}^i = \text{BERT}(U_t, d_i), i \in [1, |\mathcal{D}|] \quad (2)$$

where $\mathbf{t}_i, \mathbf{k}_j^i \in \mathbb{R}^{d_{init}}, \mathbf{K}^i = \{\mathbf{k}_j^i\}_{j=1}^{|d_i|}$. In this way, we obtain the initialized vertex embedding for a Coref-MDG as $\mathbf{H}^0 = \{\mathbf{t}_i, \mathbf{K}^i\}_{i=1}^{|\mathcal{D}|} \in \mathbb{R}^{(M+N) \times d_{init}}$.

²<https://spacy.io/>, MIT License

³https://www.wikidata.org/wiki/Wikidata:Main_Page

⁴<https://github.com/huggingface/neuralcoref>, MIT License

3.2.2 Residual Graph Propagation

Knowledge transitions in document-grounded dialogs can be divided into two types namely, transition across different documents (**out-topic**) and within the same document (**intra-topic**). Transitions across different documents occur between topic vertices in our Coref-MDG and usually requires multi-hop reasoning. We use the residual graph propagation (Step II in Figure 3) to model such transitions in Coref-MDG. Specifically, we devise a variant of Relational Graph Attention Layer(RGAT) (Busbridge et al., 2019) layer with concatenated residual connection (He et al., 2016), named RES-RGAT. This layer facilitates the deeper multi-hop information propagation by avoiding information loss and the over-smooth problem (Oono and Suzuki, 2019). The output $\mathbf{H}^{\text{out}} \in \mathbb{R}^{|\mathcal{G}| \times d_{\text{out}}}$ of one RES-RGAT is the concatenation of the propagated results and the input $\mathbf{H}^{\text{in}} \in \mathbb{R}^{|\mathcal{G}| \times d_{\text{in}}}$, which is formulated as:

$$\mathbf{H}^{\text{out}} = W[\mathbf{H}^{\text{in}}, \text{RGAT}(\mathbf{H}^{\text{in}}, \mathbf{R}, \mathcal{E}, \mathcal{G})] \quad (3)$$

where $\mathbf{R} \in \mathbb{R}^{E \times d_e}$ is the embedding look-up table for all the edge types in \mathcal{E} , E is the number of edge types, and $W \in \mathbb{R}^{d_{\text{out}} \times 2d_{\text{in}}}$ is used for dimension transform. We stack n layers of RES-RGAT to do enough propagation based on empirically determined n . With \mathbf{H}^0 as input, we obtain the propagated outputs for all vertices as $\mathbf{H}^G \in \mathbb{R}^{(M+N) \times d_G}$.

3.2.3 Differential Linearization

To integrate the dialog flow information into the knowledge representations after graph propagation, we propose a novel Differential Linearization (Step III in Figure 3) method. While knowledge sequence has been used for knowledge selection in dialog (Kim et al., 2019), knowledge shift sequence (or shifting sequence), defined as the sequence of knowledge differences within each consecutive turns, is a relatively novel notion for this task. We argue that the shifting sequence is a more useful feature for learning and predicting knowledge transitions since it focuses on the difference and interaction between knowledge, leading to sharper features. It also captures the transition patterns from turns using varying distances to the current turn to further aid in the selection.

To construct the shifting sequence, we first obtain the knowledge/topic vertices that appeared in the previous chatbot turns (since we note that the

labels of the user turns are inaccessible in practice). By collecting these knowledge/topic vertices' representations from \mathbf{H}^G , we can get the sequence $S = \{\mathbf{h}_{t-\tau}^G, \dots, \mathbf{h}_{t-1}^G\}$ for knowledge and topic vertices, respectively. Here τ is the length of turns. Since we treat topic and knowledge vertex sequence identically, we will refer to them as simply vertices in the following discussion. We compare the vertex i with these historical vertices in S with a comparison function \mathcal{F} to get the differential sequence for vertex i . By doing this for all vertices, we get $M + N$ such sequences:

$$\{\mathcal{F}(\mathbf{h}_{t-\tau}^G, \mathbf{h}_i^G), \dots, \mathcal{F}(\mathbf{h}_{t-1}^G, \mathbf{h}_i^G)\}_{i=1}^{M+N} \quad (4)$$

\mathcal{F} computes the interaction between two vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^d$ by element-wise difference and product, defined as $\mathcal{F}(\mathbf{a}, \mathbf{b}) = [\mathbf{a} - \mathbf{b}; \mathbf{a} \odot \mathbf{b}]$ (Chen et al., 2017).

With the sequence for vertex with index i , we finalize its representations in sequential transition dependency. Specifically, we feed each sequence into a stacked GRU (Cho et al., 2014) cells and use the last hidden state as the final linearized vertex representation. We concatenate the graph representation of vertex h_i^G and the linearized output:

$$\mathbf{h}_i^D = [\text{GRU}(\dots, \mathcal{F}(\mathbf{h}_{t-1}^G, \mathbf{h}_i^G)); \mathbf{h}_i^G] \in \mathbb{R}^{2d_G} \quad (5)$$

The vertex representation in graph after linearization is $\mathbf{H}^D \in \mathbb{R}^{(M+N) \times 2d_G}$, which will then be used to predict the next topic and knowledge segment.

3.3 Training

Note that topic selection is an auxiliary task in our framework, apart from the knowledge selection. As such, we split the representations for topic vertices and knowledge vertices from \mathbf{H}^D and obtain $\mathbf{H}_{tpc}^D \in \mathbb{R}^{M \times 2d_G}$ and $\mathbf{H}_{knl}^D \in \mathbb{R}^{N \times 2d_G}$, respectively. \mathbf{H}_{tpc}^D is fed into a linear layer to obtain the topic selection scores. For the knowledge vertices, we further include their connected topic vertex representations and the in-between edge embedding to calculate the knowledge selection scores similarly with a linear layer.

Following Wu et al. (2021), we implement the history loss as an auxiliary objective function in our framework to further utilize the dialog history information. Finally, the overall objective function we adopt is formulated as follows:

Method	dialogs	Avg Turns	Domain	Document
WoW	22311	9	Open Domain	Multiple
Holl-E	9071	10	Movie	Single
CMU-DoG	4112	31	Movie	Single
MultiDoc2Dial	4796	14	Info Seek	Multiple

Table 1: Dataset statistics.

$$\mathcal{L} = \mathcal{L}_{knl} + \mathcal{L}_{tpc} + \mathcal{L}_{hist}$$

$$\mathcal{L}_{hist} = \frac{1}{2l} \sum_{hi=1}^l (\mathcal{L}_{knl}^{hi} + \mathcal{L}_{tpc}^{hi}) \quad (6)$$

where l is a hyperparameter representing the history length, \mathcal{L}_{knl} and \mathcal{L}_{tpc} are knowledge and topic losses, respectively. All of the classification objective functions in \mathcal{L} are standard cross-entropy.

4 Experiments

Datasets. We validate our method on four public benchmarks for document-grounded conversation, WoW (Dinan et al., 2018), Holl-E (Moghe et al., 2018), CMU-DoG (Zhou et al., 2018b) and MultiDoc2Dial (Feng et al., 2021a). The dataset statistics are summarized in Table 1. We first conduct knowledge selection with our Coref-Diffs method and then feed the selections and dialogue history into text generation models to compare the final responses.

Evaluation metrics. We focused on evaluating the knowledge selection sub-task for the document-grounded dialog system, based on the knowledge and topic selection accuracies, denoted as **KL** and **TP**, respectively. We also explore the knowledge selection accuracy of all intra-topic data samples, whose knowledge transitions are within the same topic, denoted as **In-TP**. As for evaluating the sub-task of response generation given the dialog context and selected knowledge, we calculate the overlap of the generated response and the ground-truth with the unigram-F1 (**uF1**) and bigram-F1 (**bF1**).

Baselines. For the two commonly used datasets, WoW and Holl-E, we split the baselines into three categories based on their text encoder types. (i) **Non-Pretrained encoder:** Transformer+MemNet (Dinan et al., 2018) is the baseline released with the dataset WoW. Diffs(RNN) (Zheng et al., 2020) incorporates the knowledge difference feature in knowledge selection. (ii) **BERT encoder:** BERT+PoKS, a variant of PoKS with BERT (Devlin et al., 2019) encoder, learns knowledge selection by posterior knowledge distribution. SLKS (Kim et al., 2019) captures

historical knowledge sequence with a latent variable. PIPM (Chen et al., 2020) improves SLKS by addressing the problem of missing posterior distribution in test phase. CoLV (Zhan et al., 2021a) includes two collaborative variables for knowledge selection and response generation. KnowledgeGPT (Zhao et al., 2020) optimizes knowledge grounded dialog task by the pre-trained BERT encode and GPT-2 (Radford et al., 2019). (iii) **Passage-level BERT encoder:** DIALKI (Wu et al., 2021) encodes knowledge at passage level to capture knowledge segment relations as we do in CorefDiffs. For response generation, given that the above-mentioned methods adopted different generators, we uniformly replaced their generators with a prompt-based generator PrefixTuning (Li and Liang, 2021) for a fair comparison, thus forming the baselines with "*" in Table 3. For MultiDoc2Dial and CMU-DoG, we compare our method with the current state-of-the-art DPR+RAG (Lewis et al., 2020b) and DoHA (Prabhumoye et al., 2021), respectively. The generators used are fine-tuned BART-large.

4.1 Implementation Details.

The BERT-base models in all our experiments used the Huggingface Transformers⁵ (Wolf et al., 2020). We trained the model with Adam (Kingma and Ba, 2015) optimizer with initial learning rate $1e-5$. A linear scheduler with a warm-up strategy in 5k steps was used. The maximum history length l was empirically set to 4 for WoW, 2 for Holl-E, 3 for CMU-DoG and 4 for MultiDoc2Dial to achieve the best performance. The number of the stacked Res-RGAT was set to 2. It took around 5 and 10 epochs to achieve the reported performance by 4 nvidia V100 GPUs. We will release all the codes and the hyper-parameters settings for reproduction.

4.2 Automatic Evaluations

Knowledge Selection. The knowledge selection results on the four datasets are presented in Table 2 and 4. CorefDiffs significantly outperforms all other methods regardless of the encoders they used. Compared to the best performance achieved by DIALKI, CorefDiffs improves by 9.5% and 5.9% and is the first to achieve knowledge accuracy over 40% on both the WoW Test Seen and Unseen sets. For Holl-E, CorefDiffs also performs the best, with gains of at least 7.4% in knowledge selection. For

⁵<https://github.com/huggingface/transformers>

Dialog Context	<p><i>User</i>: Hair loss, I'm going bald!</p> <p><i>Bot</i>: There are surgeries and medications that can help [management of hair loss, 1]</p> <p><i>User</i>: What is the top option! I need help!</p> <p><i>Bot</i>: Typically, the reasons for hair loss in men has to do with genetics and male hormones.... [hair loss, 6]</p> <p><i>User</i>: Thanks for all your help/info today .</p>	<p>Three medications have evidence to support their use in male pattern hair loss: finasteride, dutasteride and minoxidil.</p> <p>Management of hair loss</p> <p>Hair loss</p> <p>The cause of male-pattern hair loss is a combination of genetics and male hormones..</p> <p>The management of hair loss, include medications and surgery</p> <p>Words overlap of</p> <p>sent 2</p> <p>sent 1</p> <p>sent 4</p> <p>sent 6</p>
Responses	<p><i>GD-SKT</i>: You are welcome. I would recommend taking three medications to help with male pattern hair loss: finasteride, dutasteride and minoxidil.</p> <p><i>DialKI</i>: well, I hope you have a great day! I know that hair loss can be caused by both genetic and environmental factors.</p> <p><i>Gold</i>: The three types of medications that seem to help the most are finasteride, dutasteride, and minoxidil. I hope this information will help you! [management of hair loss, 2]</p>	
Dialog Context	<p><i>User</i>: Seattle, I have never been to Seattle but I would love to visit.</p> <p><i>Bot</i>: I would too. I read it's the fastest growing city in 2016, with a 3.1 annual growth rate. [seattle, 4]</p> <p><i>User</i>: wow, I bet it is really busy and crowded.</p> <p><i>Bot</i>: I believe it said there were over 700,000 residents just in Seattle and it is the largest city in Washington also. [seattle, 2]</p> <p><i>User</i>: I would hate to drive there. I'm assuming people use a lot of public transportation.</p>	<p>In July 2016, seattle was again the fastest-growing major u.s. city, with a 3.1% annual growth rate</p> <p>Seattle</p> <p>With an estimated 713,700 residents , seattle is the largest city in both the state of Washington...</p> <p>A major gateway for trade with asia, seattle is the fourth-largest port in north america in terms of container</p> <p>co-ref</p> <p>sent 4</p> <p>sent 2</p> <p>sent 6</p> <p>co-ref</p> <p>co-ref</p>
Responses	<p><i>GD-SKT</i>: I'm not sure, but it is a major gateway for trade with Asia and the fourth largest port in north America.</p> <p><i>DialKI</i>: I'm not sure but I do know that it is the most populous city in the United States.</p> <p><i>Gold</i>: me too. I hate waiting in traffic. it's a major trade route with Asia. It has the fourth largest port in north America in terms of container shipping. [seattle, 6]</p>	

Figure 4: Two generation examples from WoW. The bold words in "[]" indicate the knowledge. For example, [hair loss, 6] represents the 6-th knowledge sentence in the document with topic **hair loss**. Our method chose the right knowledge for both examples compare to DIALKI owing to the well-designed graph structure.

Method	WoW (Seen)	WoW (Unseen)	Holl-E
TMN	22.5	12.2	22.7
DiffKS(RNN)	25.6	18.6	33.5
BERT+PoKS	25.5	14.1	27.6
SLKS	26.8	18.3	29.2
PIPM	27.8	19.7	30.7
CoLV	30.1	18.9	32.7
DukeNet	26.4	19.6	30.0
KnowledGPT	28.0	25.4	-
DIALKI	32.9	35.5	-
CorefDiffs	42.4	41.4	40.9
w/o Diff-Seq	40.8	39.5	39.7
w/o Diff	40.9	40.1	40.1
w/o Res-RGAT	35.5	36.5	39.5

Table 2: The knowledge selection results measured by accuracy on WoW and Holl-E.

MultiDoc2Dial our method outperforms state-of-the-art by 8.2% in knowledge selection accuracy. CMU-DoG has no ground-truth knowledge, so we only report generation results. The substantial enhancements across all datasets strongly suggest that CorefDiffs has benefited from modeling document structures and knowledge relations in the grounding documents with differential dialog flow learning.

Response Generation. Tables 3 and 4 show the results of response generation on all the four datasets. We applied PrefixTuning (Li and Liang, 2021) to generate responses with the corresponding dialog context and selected knowledge as the input for WoW and Holl-E, while for MultiDoc2Dial and CMU-DoG, we followed previous works using

Method	WoW (Seen)		WoW (Unseen)		Holl-E	
	uF1	bF1	uF1	bF1	uF1	bF1
SLKS(TM+Copy)	19.3	6.8	16.1	4.2	29.2	22.3
DukeNet(TM+Copy)	19.3	6.3	17.1	4.7	30.6	23.1
SLKS*	20.2	7.3	17.5	5.3	-	-
DiffKS*	21.5	7.6	20.0	6.3	30.7	23.9
KnowledGPT*	22.0	8.2	20.8	7.4	-	-
DIALKI*	22.0	8.0	22.2	8.1	-	-
CorefDiffs	25.2	10.7	25.8	10.8	38.4	31.8

Table 3: Response generation results on WoW and Holl-E. Methods with '(TM+Copy)' and '*' used generator Transformer + Copy mechanism and PrefixTuning. '-' indicates the method didn't do experiment on the dataset.

BART-Large. The PrefixTuning obtained a comparable performance with fewer learnable parameters and extrapolated better to unseen topics than fine-tuning method. Again, CorefDiffs obtains best performance in all generation metrics on four datasets, which we attribute to the large margins on knowledge selection.

Ablation Study. To study the impact of each of the modules in CorefDiffs, we conduct three experiments, as shown in the lower part of Table 2. For *w/o Diff-Seq*, we remove the Differential Linearization. *w/o Diff* uses the normal knowledge sequence instead of the shifting sequence. *w/o Res-RGAT* removes the Residual Graph Propagation. After removing Res-RGAT, we observe a steep drop in knowledge selection accuracy, which proves that knowledge representations updated by graph prop-

Method	MultiDoc2Dial		CMU-DOG
	uF1	KL	uF1
DoHA	-	-	22.8
DPR+RAG	33.7	24.9	-
CorefDiffs	39.3	33.1	23.9

Table 4: Experimental Results on MultiDoc2Dial and CMU-DOG.

agation on Coref-MDG is well aligned with knowledge distribution in next turn. This also shows that purely relying on local correlations within passage context (Step I in Figure 3) is not as good as using higher-level document structures. In addition, *w/o diff-seq* also presents lower knowledge selection accuracy, showing its importance to enhance dialog flow modeling upon graph propagation. More importantly, by comparing *w/o diff-seq* and *w/o Diff*, we notice that using shifting sequences outperforms normal ones, thus validating our earlier argument that shifting features are sharper and more effective for dialog knowledge flow.

4.3 Case Study

Why does Coref-MDG work on knowledge selection? To answer this question, we visualize two typical examples, shown in Figure 4. The *dialog Context* rows are dialog histories, and the generated responses of different methods are listed in the *Response* row. We compare our CorefDiffs with DIALKI and the Gold (ground-truth) response. The first example performed topic change from “hair loss” to “management of hair loss”. CorefDiffs chose the right knowledge topic, “management of hair loss”, while DIALKI repeated the knowledge mentioned in the earlier conversation turn. The reason is that CorefDiffs was able to do so is because it had referred to the `word_overlap` connection between “hair loss” and “management of hair loss”, whereas DIALKI did not consider such inter-topic relations. For the second example, the knowledge transition is intra-topic (knowledge in consecutive turns belonging to the same topic/document). Our method successfully predicts the right knowledge due to the co-reference relation between these knowledge sentences within the “seattle” document, whereas the response generated by DIALKI — even with passage-level knowledge correlations encoded — missed the long dependency from the second to the sixth sentence.

4.4 Graph Analysis

To study effects of the different type of relations in Coref-MDG on topic/knowledge selection accuracies. We did more experiments on WoW. We craft 3 Coref-MDG variants lie in three categories for relations between topics. (1) *w/o TP*: removing all topic edges; (2) *w/o TP overlap*: removing the word overlap edges; (3) *w/o TP wikigraph*: removing the commonsense edges. Another three variants for exploring the relations between knowledge vertices are as follows: (4) *w/o KG*: removing all knowledge edges (that is co-reference link); (5) *+KG common entity*: applying entity edges between knowledge instead; (6) *+KG partial order*: employing partial order edges between knowledge. We also remove the sentence order edges between topic and knowledge vertices and formed a variant (7) *w/o TP-KG*. The results of the above experiments are listed in Table 5, from which we get the following **conclusions**:

- (i) **Coref-MDG performs the best in knowledge selection compared to other graph structures.** Removing or replacing edge types in Coref-MDG, such as the edges between topic vertices (Exp. 1-3), knowledge vertices (Exp. 4-6), or topic and knowledge vertices (Exp. 7), can cause a drop in topic or knowledge selection both on Seen or Unseen settings. Moreover, sometimes using other kind of edge leads to worse results than their absence. For example, in Exp. 4 and 5, *w/o KG* performs better than *+ KG common entity* in Unseen.
- (ii) **Topic and knowledge selection accuracies are affected by their relevant relations in Coref-MDG.** In Exp. 1 and 7, without topic edges or topic-knowledge edges, the model achieves lowest TP. In Exp. 4-7, model achieves lower KL without suitable knowledge relations.
- (iii) **Topic and Knowledge relations also facilitate each other.** In Exp. 1 and 4 even removing topic relation or knowledge relations, the model still achieves better TP and KL compared to DIALKI (no graph relation used).
- (iv) **Knowledge relations improve intra-topic knowledge selection.** As shown by results in 4th and 7th columns, by comparing In-TP results in Exp. 1-3 and Exp. 4-6, after removing knowledge edges, the In-TP drops a lot, thus we conclude that relations between knowledge enhance the intra-topic knowledge selection.

Method	WoW Seen			WoW Unseen		
	KL	TP	In-TP	KL	TP	In-TP
DIALKI	32.9	70.0	42.3	35.5	71.6	43.5
CorefDiffs	42.4	76.1	51.1	41.4	77.7	49.2
1. w/o TP	42.1	74.0	50.6	39.8	75.2	47.2
2. w/o TP overlap	42.4	75.9	51.2	40.9	77.7	48.1
3. w/o TP wikigraph	42.3	75.9	50.9	41.1	77.5	48.8
4. w/o KG	35.4	75.7	44.6	37.1	77.2	46.1
5. + KG common entity	35.4	74.6	44.4	36.4	75.9	43.8
6. + KG partial order	36.6	75.9	45.7	37.1	76.8	45.5
7. w/o TP-KG	40.5	73.5	49.5	38.3	74.5	45.7

Table 5: Graph Comparisons in selection accuracy.

5 Conclusion

We show the significance of utilizing the document’s semantic structures and relations for managing dialog flow. We embody these relations in our novel multi-document graph Coref-MDG which models co-referential knowledge mention links and inter-document relations. Our analysis of Coref-MDG yields insights of how the difference among intra- or inter-document relations affect the final topic and knowledge selection accuracy. For example, we find that coreference links and topic-knowledge sentence order relations are critical relations. We then build dynamically-sensitive dialog flows via our CorefDiffs method, which integrates the modeling of dialog difference flow with the prior knowledge represented in Coref-MDG. CorefDiffs demonstrates that it is possible to seamlessly integrate static graph structures with dynamic dialog-specific flows, improving document-grounded conversations.

Ethical Impact

Document-grounded dialog technology has broad application prospects in open-domain dialog, emotional escort robots, intelligent assistants, etc. This work focuses on knowledge selection which plays a significant role in dialog management of multi-turn dialog for document-grounded conversations. All datasets we used in this work were privacy filtered and content moderated by the dataset authors (Dinan et al., 2019; Moghe et al., 2018). However, advanced dialog knowledge selection techniques may also enable bots to select harmful content on the Internet and generate inappropriate or biased responses to users. Future work should take this into consideration.

Acknowledgements

We thank all reviewers for their valuable comments. This work is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG-GC-2019-001-2A) and Industry Alignment Fund – Pre-positioning (IAF-PP) Funding Initiative. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of the National Research Foundation, Singapore.

References

- Dan Busbridge, Dane Sherburn, Pietro Cavallo, and Nils Y Hammerla. 2019. Relational graph attention networks. *arXiv preprint arXiv:1904.05811*.
- Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, Hui Jiang, and Diana Inkpen. 2017. [Enhanced LSTM for natural language inference](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1657–1668, Vancouver, Canada. Association for Computational Linguistics.
- Xiuyi Chen, Fandong Meng, Peng Li, Feilong Chen, Shuang Xu, Bo Xu, and Jie Zhou. 2020. [Bridging the gap between prior and posterior knowledge selection for knowledge-grounded dialogue generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3426–3437, Online. Association for Computational Linguistics.
- Hao Cheng, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2020. [Probabilistic assumptions matter: Improved models for distantly-supervised document-level question answering](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5657–5667, Online. Association for Computational Linguistics.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- Pradeep Dasigi, Nelson F. Liu, Ana Marasović, Noah A. Smith, and Matt Gardner. 2019. [Quoref: A reading comprehension dataset with questions requiring coreferential reasoning](#). 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 5925–5932, Hong Kong, China. Association for Computational Linguistics.

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2018. Wizard of wikipedia: Knowledge-powered conversational agents. In *International Conference on Learning Representations*.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. Wizard of wikipedia: Knowledge-powered conversational agents. In *International Conference on Learning Representations*.
- Song Feng, Siva Sankalp Patel, Hui Wan, and Sachindra Joshi. 2021a. [MultiDoc2Dial: Modeling dialogues grounded in multiple documents](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6162–6176, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Song Feng, Siva Reddy, Malihe Alikhani, He He, Yangfeng Ji, Mohit Iyyer, and Zhou Yu, editors. 2021b. *Proceedings of the 1st Workshop on Document-grounded Dialogue and Conversational Question Answering (DialDoc 2021)*. Association for Computational Linguistics, Online.
- Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen-tau Yih, and Michel Galley. 2018. A knowledge-grounded neural conversation model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Byeongchang Kim, Jaewoo Ahn, and Gunhee Kim. 2019. Sequential latent knowledge selection for knowledge-grounded dialogue. In *International Conference on Learning Representations*.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020a. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020b. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Sha Li, Di Jin Madhi Namazifar, Mohit Bansal, Heng Ji, Yang Liu, and Dilek Hakkani-Tur. 2022. Enhanced knowledge selection for grounded dialogues via document semantic graphs.
- Xiang Lisa Li and Percy Liang. 2021. [Prefix-tuning: Optimizing continuous prompts for 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 1: Long Papers)*, pages 4582–4597, Online. Association for Computational Linguistics.
- Rongzhong Lian, Min Xie, Fan Wang, Jinhua Peng, and Hua Wu. 2019. Learning to select knowledge for response generation in dialog systems. In *IJCAI International Joint Conference on Artificial Intelligence*, page 5081.
- Zhibin Liu, Zheng-Yu Niu, Hua Wu, and Haifeng Wang. 2019. [Knowledge aware conversation generation with explainable reasoning over augmented graphs](#). 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 1782–1792, Hong Kong, China. Association for Computational Linguistics.
- Chuan Meng, Pengjie Ren, Zhumin Chen, Zhaochun Ren, Tengxiao Xi, and Maarten de Rijke. 2021. Initiative-aware self-supervised learning for knowledge-grounded conversations. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 522–532.
- Chuan Meng, Pengjie Ren, Zhumin Chen, Weiwei Sun, Zhaochun Ren, Zhaopeng Tu, and Maarten de Rijke. 2020. Dukenet: A dual knowledge interaction network for knowledge-grounded conversation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1151–1160.
- Nikita Moghe, Siddhartha Arora, Suman Banerjee, and Mitesh M. Khapra. 2018. [Towards exploiting background knowledge for building conversation systems](#).

- In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2322–2332, Brussels, Belgium. Association for Computational Linguistics.
- Seungwhan Moon, Pararth Shah, Anuj Kumar, and Rajen Subba. 2019. [OpenDialKG: Explainable conversational reasoning with attention-based walks over knowledge graphs](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 845–854, Florence, Italy. Association for Computational Linguistics.
- Kenta Oono and Taiji Suzuki. 2019. Graph neural networks exponentially lose expressive power for node classification. In *International Conference on Learning Representations*.
- Shrimai Prabhumoye, Kazuma Hashimoto, Yingbo Zhou, Alan W Black, and Ruslan Salakhutdinov. 2021. [Focused attention improves document-grounded generation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4274–4287, Online. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Zequ Wu, Bo-Ru Lu, Hannaneh Hajishirzi, and Mari Ostendorf. 2021. [DIALKI: Knowledge identification in conversational systems through dialogue-document contextualization](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1852–1863, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jun Xu, Zeyang Lei, Haifeng Wang, Zheng-Yu Niu, Hua Wu, and Wanxiang Che. 2021a. Enhancing dialog coherence with event graph grounded content planning. In *Proceedings of the Twenty-Ninth International Conference on Artificial Intelligence*, pages 3941–3947.
- Jun Xu, Zeyang Lei, Haifeng Wang, Zheng-Yu Niu, Hua Wu, Wanxiang Che, Jizhou Huang, and Ting Liu. 2021b. Coherent dialog generation with query graph. *Transactions on Asian and Low-Resource Language Information Processing*, 20(6):1–23.
- Jun Xu, Haifeng Wang, Zheng-Yu Niu, Hua Wu, Wanxiang Che, and Ting Liu. 2020a. [Conversational graph grounded policy learning for open-domain conversation generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1835–1845, Online. Association for Computational Linguistics.
- Jun Xu, Haifeng Wang, Zhengyu Niu, Hua Wu, and Wanxiang Che. 2020b. Knowledge graph grounded goal planning for open-domain conversation generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9338–9345.
- Lin Xu, Qixian Zhou, Ke Gong, Xiaodan Liang, Jianheng Tang, and Liang Lin. 2019. End-to-end knowledge-routed relational dialogue system for automatic diagnosis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7346–7353.
- Deming Ye, Yankai Lin, Jiaju Du, Zhenghao Liu, Peng Li, Maosong Sun, and Zhiyuan Liu. 2020. [Coreferential Reasoning Learning for Language Representation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7170–7186, Online. Association for Computational Linguistics.
- Haolan Zhan, Lei Shen, Hongshen Chen, and Hainan Zhang. 2021a. [CoLV: A collaborative latent variable model for knowledge-grounded dialogue generation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2250–2261, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Haolan Zhan, Hainan Zhang, Hongshen Chen, Zhuoye Ding, Yongjun Bao, and Yanyan Lan. 2021b. [Augmenting knowledge-grounded conversations with sequential knowledge transition](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5621–5630, Online. Association for Computational Linguistics.
- Xueliang Zhao, Wei Wu, Can Xu, Chongyang Tao, Dongyan Zhao, and Rui Yan. 2020. [Knowledge-grounded dialogue generation with pre-trained language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3377–3390, Online. Association for Computational Linguistics.
- Chujie Zheng, Yunbo Cao, Daxin Jiang, and Minlie Huang. 2020. [Difference-aware knowledge selection for knowledge-grounded conversation generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 115–125, Online. Association for Computational Linguistics.
- Hao Zhou, Tom Young, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018a. Commonsense knowledge aware conversation generation with graph attention. In *IJCAI*, pages 4623–4629.

Kangyan Zhou, Shrimai Prabhume, and Alan W Black. 2018b. [A dataset for document grounded conversations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 708–713, Brussels, Belgium. Association for Computational Linguistics.

A Implementation Details

We set the maximum lengths of model input to 512, which is also the longest input limit for the BERT model, in order to fit the longer passage text as much as possible on both datasets. We employ a Linear layer to transform the output features of BERT from 768 to 320 to reduce memory usage. The edge embedding size is set to 64. The hidden size and headers of Res-RGAT are 1024 and 8 respectively while the alpha value of Graph Attention Network is 0.2. We utilize a unidirectional stacked GRU model for Differential Sequential Learning, the number of GRU layers is 2.

For response generation, we apply PrefixTuning (Li and Liang, 2021) on BART (Lewis et al., 2020a) large model to learn the responses generation model based on the knowledge selection results from the previous stage. We use the prefix length 200 and the hidden dimension of 800 for all the methods using PrefixTuning generator. The PrefixTuning generator takes about 4 hours and 30 epoch to become converged during training on 4 V100 32G GPUs, which is much faster and more resource saving than fine-tuning BART large.

B Dataset Processing Details

WoW. There are more than 130k different documents from Wikipedia in WoW training set. We keep 350 high-frequency relations from the Wiki knowledge graph, covering these 130k documents. The top-10 wiki relations with corresponding frequency are shown as follows:

1. ('subclass of', 27015)
2. ('facet of', 11381)
3. ('sport', 10646)
4. ('performer', 9482)
5. ('part of', 6892)
6. ('manufacturer', 5742)
7. ('instance of', 5551)
8. ('history of topic', 5517)
9. ('has part', 5445)
10. ('follows', 5077)

As shown in Table 6, for topic relations, we found the `word_overlap` edges is denser than the commonsense edges from wikiData, giving the average edge number of 8.11 and 2.89, respectively. While for knowledge relations, the `coreference_link` has much less average number of relations in one sample than other two types relations, which again proves that `coreference_link` with more accurate knowledge relations lead to better knowledge selection results without introducing wrong structures information to CorefDiffs framework.

	Topic Relations		Knowledge relations		
	WordOverlap	WikiGraph	Partial	EntityLink	Coreference
Freq	8.11	2.89	61.18	87.52	15.90

Table 6: Average number of different kinds of relations in one sample on the WoW training set.

Holl-E. Different from WoW, each sample of Holl-E has only one topic, which is the movie in this session of conversation. There are four types of information for each movie in Holl-E, which are plots, comments, reviews, and table information. So we simply divide all the knowledge sentences of each movie into four topics. As the absence of common sense relations of such topics in Holl-E, we count the co-occurrence relationship of all topics in the training set as the relations between topics in Holl-E. The relations between knowledge are as same as the WoW, using coreference relations in passage text. The relations between knowledge and topics are sentence order of knowledge sentence in the original text, which is also used in WoW. **CMU-DoG.** CMU-DoG is a document-grounded conversation dataset about movie, which is the same as Holl-E. The difference is that CMU-DoG includes only one grounding document(one topic) at each dialog turn. The relations of topics is absent as there is only one topic in grounding document. The relations between knowledge are as same as the WoW and Holl-E with coreference relations. The relations between knowledge and topics are sentence order of knowledge sentence in the passage, which is also consistent with WoW and Holl-E. On the other hand, CMU-DoG didn't contain the gold knowledge of knowledge selection task. We adopt unigram F1 score as similarity function, selecting the knowledge closest to the ground-truth response as gold knowledge to train the selector model.

MultiDoc2Dial. MultiDoc2Dial includes multiple grounding documents at each dialog turn. We con-

struct the graph following the steps of Holl-E. However, MultiDoc2Dial introduces a span prediction task to locate knowledge set in the original document instead of knowledge selection. But that’s ok, it easy for our framework to transfer downstream task by using two independent classifier to predict both start knowledge segment and end knowledge segment instead of one classifier for knowledge selection. Simultaneously, we replace the metric from knowledge accuracy to knowledge EM, which is used in MultiDoc2Dial. For convenience, we still use **KL** in Table 4 to denote the EM metric in MultiDoc2Dial.

C Analysis on Partial Order Edge

For partial order relations, we explored the effects of different hops. Hop- k partial order relation means each knowledge vertex is connected with k knowledge vertices behind according to the sentence order. As shown in Fig 5, hop-2 partial relation performed the best. A hop that was too large or too small could cause information loss or introduce many erroneous connections.

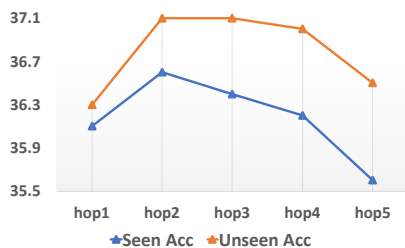


Figure 5: Knowledge accuracy for partial order with different hops.