

Graph-Guided Unsupervised Knowledge Identification for Dialogue Agents

Shrinivas Khiste^{1*} Tathagata Raha^{2*} Milan Aggarwal³ Sumit Bhatia³ Simra Shahid³
¹ IIT Kharagpur ² IIIT Hyderabad
³ Media and Data Science Research Lab, Adobe, India

Abstract

Dialogue systems that can effectively respond to user queries in a conversational style have become ubiquitous. Large Language Models (LLMs) have been extensively used as a key component in such systems owing to their linguistic capabilities and implicit knowledge. However, such models are prone to hallucinate while generating a response that can be detrimental, particularly in applications where accuracy is critical. While many works have attempted to address the hallucination concern by supplementing external knowledge in the input to the LLMs, most of them rely on supervised labels to train the knowledge identification module. Such labels might often not be available or difficult to obtain at scale. To address this, we propose our method RANKING, which leverages the structure of the external document to obtain a ranked subset of relevant sentences in an unsupervised manner that can be used for response generation. We model the dependencies in the form of a graph between the sentences present in the external document and the utterances till the given point in the dialogue. We demonstrate the efficacy of RANKING on a commonly used document-grounded conversation dataset (Doc2Dial) where it is observed that RANKING enables generating better responses than using the entire document.

1 Introduction

Dialogue systems that possess the ability to respond to user queries have become popular. The emergence of Chat-GPT (Ouyang et al., 2022) has further revolutionized the domain by facilitating human-like question-answering. Users now prefer to search for desired information through interactive dialogues, as opposed to retrieving and comprehending lengthy documents on their own.

Despite the remarkable capabilities of existing dialogue systems such as Chat-GPT, the cost, legal and privacy concerns surrounding such models

make them unsuitable for general business applications. Additionally, queries posed by the users of a business application may require domain-specific knowledge, which a Large Language Model (LLM) may not be privy to. Further, such models are susceptible to hallucinating their responses (Roller et al., 2021; Marcus, 2020) which can be detrimental in scenarios where accurate responses are critical. To mitigate this, there is a growing interest in incorporating external knowledge while using such LLMs (Kaur et al., 2022).

Numerous works have attempted to leverage the knowledge present in external documents to generate automated dialogue responses. Most approaches rely on supervised labels to train knowledge identification module to retrieve text from an external document that is relevant to the dialogue (Lian et al., 2019; Kim et al., 2020; Zheng et al., 2020; Wu et al., 2021). Obtaining such labels requires manually annotating sentences that should be selected by the model at each step in the dialogue which can be expensive and time-consuming. This creates a need to develop approaches that can identify relevant external knowledge in an unsupervised manner (Huang et al., 2021). Typical unsupervised knowledge selection methods employ semantic similarity between the embedded representations of the dialogue context and the text present in external knowledge (Huang et al., 2021; Dinan et al., 2019). However, this ignores the structure in which information is organized in an external document which has been shown to be useful (Djeddal et al., 2021). Such methods also fail to properly capture the relationship between the information flow in the dialogue utterances and structured information in the external document¹.

To address this, we propose our method RANKING (Graph-Guided Unsupervised Knowledge Identification for Dialogue Agents) that captures the dependencies between the knowledge sen-

*Work done as an intern at MDSR, Adobe, India

¹Please refer to Appendix A for detailed related work

tences, the user and the agent utterances from the dialogue through a graph by representing them as nodes. The graph also captures information flow in the document based on its structure (such as paragraphs). RANKING learns contextualized representations for each node and provides a ranked list of relevant sentences from external knowledge for response generation. We study different design choices for the underlying graph and observe that RANKING performs better than simple TF-IDF or BERT-based embedding similarity for selecting relevant sentences. Further, it is observed that RANKING helps in precisely selecting relevant information by filtering out irrelevant noisy sentences and performs better than providing the entire document as input for response generation.

2 Our Methodology

Problem Definition: Given a dialogue $\mathcal{D} = \{u_1, a_1, u_2, a_2, \dots, u_i\}$ (where u_j and a_j represent the user and agent utterances respectively); and a supporting grounding document ($\mathcal{K} = \{p_1, p_2, \dots, p_P\}$) comprising of paragraphs, the task of the modeled agent is to generate the response a_i . Here $p_i = \{s_1^i, s_2^i, \dots, s_{|p_i|}^i\}$; s_k^q represents the k^{th} sentence in the q^{th} paragraph; $|p_i|$ represents the total number of sentences in paragraph p_i and P represents the number of paragraphs in \mathcal{K} . The model has to select sentences from \mathcal{K} relevant to \mathcal{D} while generating the next agent response a_i .

Overview of RANKING: Figure 1 depicts the architecture of our method. We construct a Document-Dialogue graph \mathcal{G} by representing each utterance in \mathcal{D} as well as paragraphs and sentences in \mathcal{K} as nodes. The nodes are connected with edges to capture the semantic dependencies between them. Higher-order relationships between different types of nodes are encapsulated using the Node2Vec algorithm that updates the embeddings of each node by utilizing the underlying graph structure. Subsequently, the similarity between the node representing the most recent user utterance and each knowledge sentence is estimated to rank the sentences in \mathcal{K} . The top-ranked sentences are provided as input in addition to the dialogue history to a response generation module that consists of a pre-trained generative LLM.

2.1 Document-Dialogue Graph

We refer to our proposed graph structure as Document Dialogue Graph (Figure 2 in appendix B). It

consists of 4 types of nodes: sentence nodes (s_k^q), paragraph nodes (p_i), user utterance nodes (u_i), and agent utterance nodes (a_i). We first explain the graph structure followed by the intuition behind it. A sentence node in a paragraph is connected with the node representing the next sentence in the paragraph ($s_k^i \rightarrow s_{k+1}^i$). A paragraph node is connected to the next paragraph ($p_i \rightarrow p_{i+1}$) and also to the sentence nodes in that paragraph ($s_k^i \leftrightarrow p_i$). The user and agent utterance nodes are connected to the next utterance in the dialogue ($u_i \rightarrow a_i$, $a_i \rightarrow u_{i+1}$). Utterance nodes are also connected to the paragraph nodes as ($u_j \rightarrow p_i \rightarrow a_j$).

The graph nodes and the edges between them model various interactions needed to perform knowledge selection. Paragraph nodes and directed edges between them are important to encapsulate the hierarchical structure and information flow in the document. The flow of information between the dialogue and the document is facilitated through edges from user utterance nodes to paragraph nodes and back to agent utterance nodes. This allows it to ground the user query in knowledge sentences and match them with the agent response. The undirected edges between sentences and paragraphs facilitate the interactions between the local information present within a paragraph and the content in other paragraphs and dialogue utterances. We initialize the representation of each node by processing the corresponding text through a sparse featurizer - TF-IDF. The weight for an edge is computed as cosine similarity between the corresponding node representations. We ablate over different graph structures in the experiments section.

2.2 Learning Node Representations

To obtain the contextual representation of each node based on the graph structure, we use Node2Vec (Grover and Leskovec, 2016), an algorithm used to learn continuous feature representations of graph nodes. Specifically, it treats each node in the random walk sequence as a target node and the surrounding nodes as context nodes. It then uses stochastic gradient descent to maximize the likelihood of predicting the context nodes given the target node. We obtain representation $V_i = \text{Node2Vec}(\mathcal{G})[i]$ for i^{th} node in \mathcal{G} .

2.3 Sentence Selection

We select the most relevant knowledge sentences (s_k^q) based on the most recent user query (u_i) in the dialogue. This is achieved by computing the

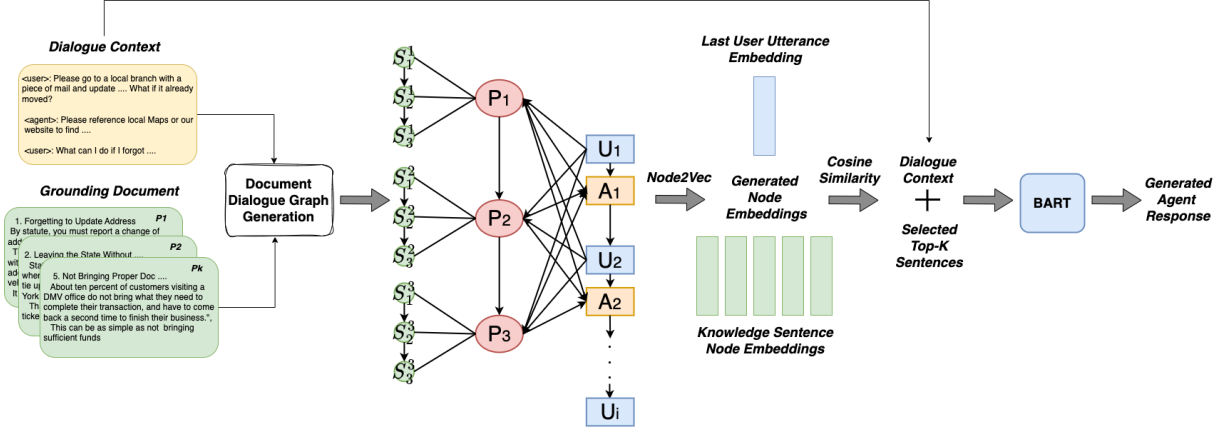


Figure 1: Model architecture of RANKING comprising of creating a graph derived using the dialogue history and text present in the grounding document. To capture the higher-order semantic dependencies between the nodes, Node2Vec is used to update the representation of each node by taking into consideration the graph structure. Finally, the similarity between the node representing the most recent user utterance and knowledge sentences is used to rank the latter and provided as input along with the dialogue history to the response generation model.

pairwise cosine similarity between each s_k^g with the representations of u_i . The top-k sentences ($\mathcal{K}_{sel} = \{s_1, s_2, \dots, s_k\}$) are selected based on this score and provided as additional context to a response generation module. The ordering of the sentences is maintained as seen in the document to avoid disturbing information flow.

2.4 Response Generation

A pre-trained generative LLM is fine-tuned to generate the agent response (a_i) conditioned on the dialogue context and the selected knowledge sentences. In our case, we use BART as the pre-trained LLM to perform response generation. The input to the generative LLM is provided as follows:

$$X = [CLS][\text{"query"}]u_i[\text{"usr"}]u_1[\text{"agt"}]a_1\dots[\text{"usr"}]u_i[\text{"grounding"}]s_1[SEP]\dots s_k[SEP]$$

Special tags such as $[\text{"query"}]$, $[\text{"usr"}]$, $[\text{"agt"}]$, $[\text{"grounding"}]$ etc. are used to instruct the model about the corresponding segments of the input.

3 Experiments

We now discuss the experimental setup to study the efficacy of RANKING and various design choices.

3.1 Dataset

We evaluate our proposed model using a goal-oriented document-grounded dialogue dataset Doc2Dial (Feng, 2021). It contains 3,474 dialogues with 44,149 turns for training and 661 dialogues

with 8539 turns for evaluation. It contains multi-turn conversations grounded in relevant documents from four domains for social welfare.

3.2 Implementation Details

For knowledge identification, we use the standard Node2Vec implementation² with $p = 1$, $q = 0.1$, number of walks = 200, output vector size of 128, walk length of 10 and window size 5. These are decided after hyper-parameter tuning using Grid Search. We set the number of knowledge sentences to be selected as $k = 15$. Training is performed for 10 epochs with an initial learning rate of $3e - 5$ using Adam optimizer (epsilon = $1e - 6$) (Kingma and Ba, 2014). The BART model is fine-tuned on a single A100 GPU with 40GB memory, 24 vCPU threads, and 170GB RAM.

3.3 Evaluation Metrics

We use Average Rank as the metric to evaluate the knowledge identification task. Given the actual set of grounding sentences - $s_1^g, s_2^g, \dots, s_m^g$ for a given turn in a dialogue; and the corresponding rank of these sentences as predicted by a model - r_1, r_2, \dots, r_m , the corresponding Rank (R) is estimated as $R = \frac{\sum_i r_i}{m}$.

Based on the above definition of rank, Average Rank is computed accordingly by averaging over the evaluation samples. BLEU (Papineni et al., 2002; Post, 2018) metric is used to evaluate the response generation module.

²<https://github.com/eliorc/node2vec>

Model	Average Rank
Sentence BERT	19.62
TF-IDF	18.76
RANKING (our method)	11.09

Table 1: Evaluation of RANKING on the validation set of Doc2Dial dataset. It can be seen that our method improves the ranking (lower is better) of relevant sentences significantly over just using TF-IDF or BERT embedding-based similarity.

Model	BLEU
Entire document as input	17.67
RANKING (unsupervised)	18.09
DIALKI (supervised)	19.01

Table 2: Response generation results on Doc2Dial dataset. RANKING improves the response quality compared to providing the entire document as input. RANKING also bridges the performance gap with supervised baseline - DIALKI which uses supervised labels for training knowledge identification module.

3.4 Comparison with Baselines

We compare RANKING with the method where the entire document \mathcal{K} is provided as input with the dialogue context \mathcal{D} as input to the response generation model. We also compare with other simple unsupervised approaches like BERT embedding-based and TF-IDF-based ranking which uses cosine similarity between corresponding feature embeddings. This also evaluates the efficacy of utilizing the graph structure toward computing better embeddings for ranking. We also compare with DIALKI (Wu et al., 2021), a method that uses supervised labels for training the knowledge identification module.

It can be seen in Table 1 that our model improves knowledge identification significantly over using the TF-IDF or BERT embedding-based ranking outlining the importance of graph connections in obtaining enriched and contextualized embeddings. Further, Table 2 shows improvement in response generation BLEU score compared to providing the whole document as input. This highlights the fact that response generation improves if the selected knowledge is more precise, less noisy and the proportion of irrelevant sentences is lesser. Further, it can be noticed that RANKING helps in bridging the performance gap between unsupervised methods and the supervised baseline - DIALKI. Please refer to appendix C for qualitative analysis.

Description	Average Rank
Information Flow between dialogue and document	
Agent ->Para ->User	21.231
Para<->User ->Agent	12.684
User ->Para ->Agent	11.697
Information Flow in the Document	
Undirected Sentence, Para Connection	11.697
Sentence ->Para	17.040
Undirected Adjacent Para Connection	11.341
Directed Adjacent Para Connection	11.112

Table 3: Ablation experiments for the design of the graph structure to highlight the importance of different edges and their directionality.

3.5 Ablation Study

We discuss the experiments performed to make design choices related to the graph structure. Results are shown in Table 3. We ablate over different ways of connecting various types of nodes as well as the directionality of edges. We observe that having a directed edge from the user utterance nodes to the paragraph nodes followed by an edge with the agent utterance node performs the best. For capturing the information flow within the document, we compare having directed vs. undirected edges between the sentence and paragraph nodes where the latter is found to perform better. Finally, directed edges between the adjacent paragraph nodes based on the order in which they appear in the document performs better than other options.

4 Conclusion and Future Work

We introduce a novel graph-based knowledge selection method RANKING that identifies relevant external knowledge in an unsupervised manner which can be used to generate better dialogue responses. RANKING captures the interdependencies between the dialogue utterances and text present in an external document. Most prior works rely on supervised labels to train the knowledge identification module which might not always be available. We show that our method improves knowledge selection over other unsupervised baselines and bridges the gap in performance with supervised methods. We also perform ablations to study the importance of graph structure. As future work, it can be explored to jointly train the knowledge identification and response generation modules to enable them to provide feedback to each other.

References

- Nico Daheim, David Thulke, Christian Dugast, and Hermann Ney. 2021. [Cascaded span extraction and response generation for document-grounded dialog](#). In *Proceedings of the 1st Workshop on Document-grounded Dialogue and Conversational Question Answering (DialDoc 2021)*, pages 57–62, Online. Association for Computational Linguistics.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. [Wizard of wikipedia: Knowledge-powered conversational agents](#).
- Hanane Djeddal, Thomas Gerald, Laure Soulier, Karen Pinel-Sauvagnat, and Lynda Tamine. 2021. [Does structure matter? leveraging data-to-text generation for answering complex information needs](#).
- Song Feng. 2021. [DialDoc 2021 shared task: Goal-oriented document-grounded dialogue modeling](#). In *Proceedings of the 1st Workshop on Document-grounded Dialogue and Conversational Question Answering (DialDoc 2021)*, pages 1–7, Online. Association for Computational Linguistics.
- Chang Gao, Wenxuan Zhang, and Wai Lam. 2022. [Unigdd: A unified generative framework for goal-oriented document-grounded dialogue](#).
- Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen tau Yih, and Michel Galley. 2018. [A knowledge-grounded neural conversation model](#).
- Aditya Grover and Jure Leskovec. 2016. [node2vec: Scalable feature learning for networks](#).
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. [REALM: retrieval-augmented language model pre-training](#). *CoRR*, abs/2002.08909.
- Xinxian Huang, Huang He, Siqi Bao, Fan Wang, Hua Wu, and Haifeng Wang. 2021. [PLATO-KAG: Unsupervised knowledge-grounded conversation via joint modeling](#). In *Proceedings of the 3rd Workshop on Natural Language Processing for Conversational AI*, pages 143–154, Online. Association for Computational Linguistics.
- Jivat Kaur, Sumit Bhatia, Milan Aggarwal, Rachit Bansal, and Balaji Krishnamurthy. 2022. [LM-CORE: Language models with contextually relevant external knowledge](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 750–769, Seattle, United States. Association for Computational Linguistics.
- Boeun Kim, Dohaeng Lee, Damrin Kim, Hongjin Kim, Sihyung Kim, Ohwoog Kwon, and Harksoo Kim. 2022. [Generative model using knowledge graph for document-grounded conversations](#). *Applied Sciences*, 12(7).
- Byeongchang Kim, Jaewoo Ahn, and Gunhee Kim. 2020. [Sequential latent knowledge selection for knowledge-grounded dialogue](#). In *International Conference on Learning Representations*.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2021. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#).
- Jiapeng Li, Mingda Li, Longxuan Ma, Wei-Nan Zhang, and Ting Liu. 2021. [Technical report on shared task in DialDoc21](#). In *Proceedings of the 1st Workshop on Document-grounded Dialogue and Conversational Question Answering (DialDoc 2021)*, pages 52–56, Online. Association for Computational Linguistics.
- Sha Li, Mahdi Namazifar, Di Jin, Mohit Bansal, Heng Ji, Yang Liu, and Dilek Hakkani-Tur. 2022. [Enhanced knowledge selection for grounded dialogues via document semantic graphs](#).
- Zekang Li, Cheng Niu, Fandong Meng, Yang Feng, Qian Li, and Jie Zhou. 2019. [Incremental transformer with deliberation decoder for document grounded conversations](#).
- Rongzhong Lian, Min Xie, Fan Wang, Jinhua Peng, and Hua Wu. 2019. [Learning to select knowledge for response generation in dialog systems](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 5081–5087. International Joint Conferences on Artificial Intelligence Organization.
- Xiexiong Lin, Weiyu Jian, Jianshan He, Taifeng Wang, and Wei Chu. 2020. [Generating informative conversational response using recurrent knowledge-interaction and knowledge-copy](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 41–52, Online. Association for Computational Linguistics.
- Gary Marcus. 2020. [The next decade in ai: Four steps towards robust artificial intelligence](#).
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.

- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Eric Michael Smith, Y-Lan Boureau, and Jason Weston. 2021. [Recipes for building an open-domain chatbot](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 300–325, Online. Association for Computational Linguistics.
- Kexin Wang, Zhixu Li, Jiaan Wang, Jianfeng Qu, Ying He, An Liu, and Lei Zhao. 2022. [Rt-kgd: Relation transition aware knowledge-grounded dialogue generation](#).
- Zeqiu 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.
- Yan Xu, Etsuko Ishii, Genta Indra Winata, Zhaojiang Lin, Andrea Madotto, Zihan Liu, Peng Xu, and Pascale Fung. 2021. [CAiRE in DialDoc21: Data augmentation for information seeking dialogue system](#). In *Proceedings of the 1st Workshop on Document-grounded Dialogue and Conversational Question Answering (DialDoc 2021)*, pages 46–51, Online. Association for Computational Linguistics.
- Semih Yavuz, Abhinav Rastogi, Guan-Lin Chao, and Dilek Hakkani-Tur. 2019. [Deepcopy: Grounded response generation with hierarchical pointer networks](#).
- 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.

A Related Work

Existing approaches mainly apply a pipelined approach to perform Knowledge Identification followed by Response Generation by utilizing separate models for each task (Daheim et al., 2021; Xu et al., 2021; Li et al., 2021; Zhao et al., 2020; Lin et al., 2020). Some works have also attempted to jointly model both the steps together as a single generative objective (Guu et al., 2020; Lewis et al., 2021; (Gao et al., 2022)).

Very limited works have addressed the issue of identifying the relevant knowledge in an unsupervised manner. For instance, Li et al., 2019, Yavuz et al., 2019, and Lin et al., 2020 leverage implicit soft fusion techniques to combine externally available knowledge resources, without an explicit supervised training step. Although some attempts have been made to perform an unsupervised selection of external knowledge based on semantic similarity (Ghazvininejad et al., 2018; Huang et al., 2021; Dinan et al., 2019), the top-1 knowledge selection criteria employed by such methods make it difficult to identify the knowledge that should be present in the target response to be generated (Huang et al., 2021). Consequently, we employ top-k selection to improve the robustness of knowledge selection.

Some works have attempted to address the knowledge identification problem using graphs (Kim et al., 2022; Li et al., 2022; Wang et al., 2022). They model transitions between user and agent utterances in a dialogue to better identify knowledge over a knowledge graph. Most of these approaches are supervised and do not take full advantage of the interactions between the dialogue and the document, unlike our work. Moreover, these approaches rely heavily on standard knowledge graphs that are either provided or built on top of documents. In contrast, we focus on designing an unsupervised pipeline based on a custom graph that incorporates specific structures to ensure optimal performance for knowledge-grounded dialogue generation.

B Graph Structure

Our proposed graph structure is called a Document Dialogue Graph (Figure 2). It consists of 4 types of nodes: Sentence Nodes (s_k^q), Paragraph Nodes (p_i), User utterance nodes (u_i), and Agent utterance nodes (a_i). We explain the graph structure first followed by the intuition behind it. A sentence node in a paragraph is connected with the node




Figure 2: Document Dialogue Graph Structure. The document dialogue graph consists of nodes representing spans/sentences (green nodes), paragraphs (red nodes), user utterances (blue nodes) and agent utterances (orange nodes) and has connections to model interactions between each of these.

representing the next sentence in the paragraph ($s_k^i \rightarrow s_{k+1}^i$). A paragraph node is connected to the next paragraph ($p_i \rightarrow p_{i+1}$) and also to the sentence nodes in that paragraph ($s_k^i \leftrightarrow p_i$). The user and agent utterance nodes are connected to the next utterance in the dialogue history ($u_i \rightarrow a_i$, $a_i \rightarrow u_{i+1}$). Utterance nodes are also connected to the paragraph nodes ($u_j \rightarrow p_i \rightarrow a_j$).

C Qualitative Analysis

Let us have a look at how the ground knowledge can be identified using examples from the dataset. This will also help us to reason about the importance of using a graph-based approach to identify relevant knowledge.

C.1 Example 1

Past dialogues reveal which topic is being discussed and can be used to locate the section which will have our grounding sentence. Thus noting where previous sentences have been grounded is important. This is done using edges between documents and dialogues. Refer to Figure 3 for an example.


<pre>"<user>: Hi. I need more information about your medical benefits package" "<agent>: In our package we cover preventive care services, such as:\n\u00a0 Health exams, including gender-specific exams\n\u00a0 Health education, including nutrition education", "<user>: What finger to do to use these services?", "<agent>: In order to use these services, you must be enrolled in VA health care, and you must have received care from us within the past 24 months. Be sure to tell the urgent care provider that you are using the VA urgent care benefit when you arrive.", "<user>: What should I do if the emergency care they give me is not covered by VA?", "<agent>: You must meet several requirements. Learn more about non-VA emergency medical care",</pre>	<p>All Grounded Here</p> 	<pre>"More about your medical benefits package", "We cover preventive care services, like : Health exams including gender - specific exams Health education including nutrition education", "..... "..... To use these services, you ll need to be enrolled in VA health care , and you'll need to have received care from us within the past 24 months.", "..... "For us to consider covering non - VA emergency care for a non - service - connected condition , you ll need to meet several requirements.", "Learn more about non - VA emergency medical care", "....."</pre>
---	--	--

Figure 3: The first dialogue mentions Medical Care Section and the rest of the user queries are all grounded in the Medical Care Section.

C.2 Example 2

Conversation flow sometimes follows the information flow in the section too. Thus edges between adjacent paragraphs and utterances are important. Refer to figure 4 for an example.

C.3 Example 3

There are two options for the agent to look for grounding knowledge: either follow the same topic on which the previous dialogue is based OR jump to a new topic which can be determined based on the current dialogue itself (not previous reference needed). Thus the two types of connections: between utterances and utterances and documents are important. Refer to Figure 5 for an example. An important thing to note is that the conversation does not jump back to a topic discussed in an earlier part of the conversation without the current dialogue being enough to decide that we have to jump to that topic. Thus it is unnecessary to model jumps to topics discussed earlier in the conversation and only recent topics have an effect on the current grounding which is done by our model with the help of edges between adjacent utterances in the conversation.

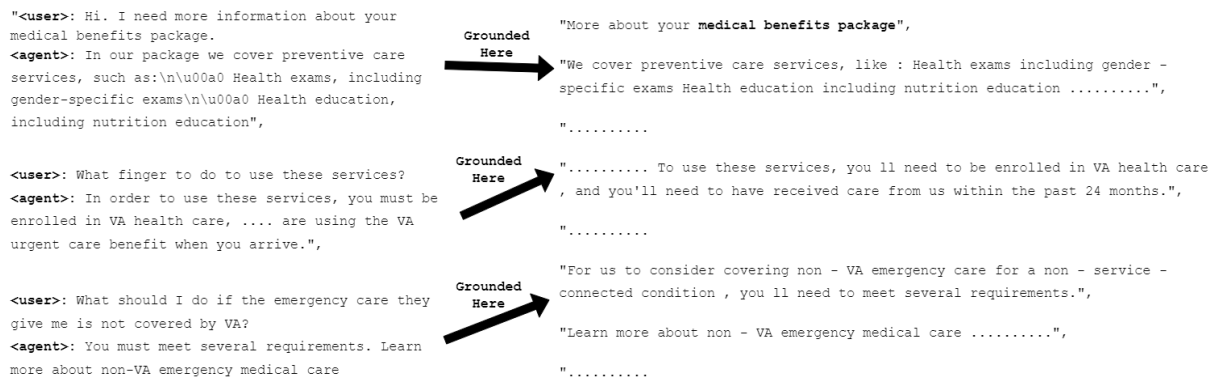


Figure 4: The information flow in the dialogue follows the information flow in the document.

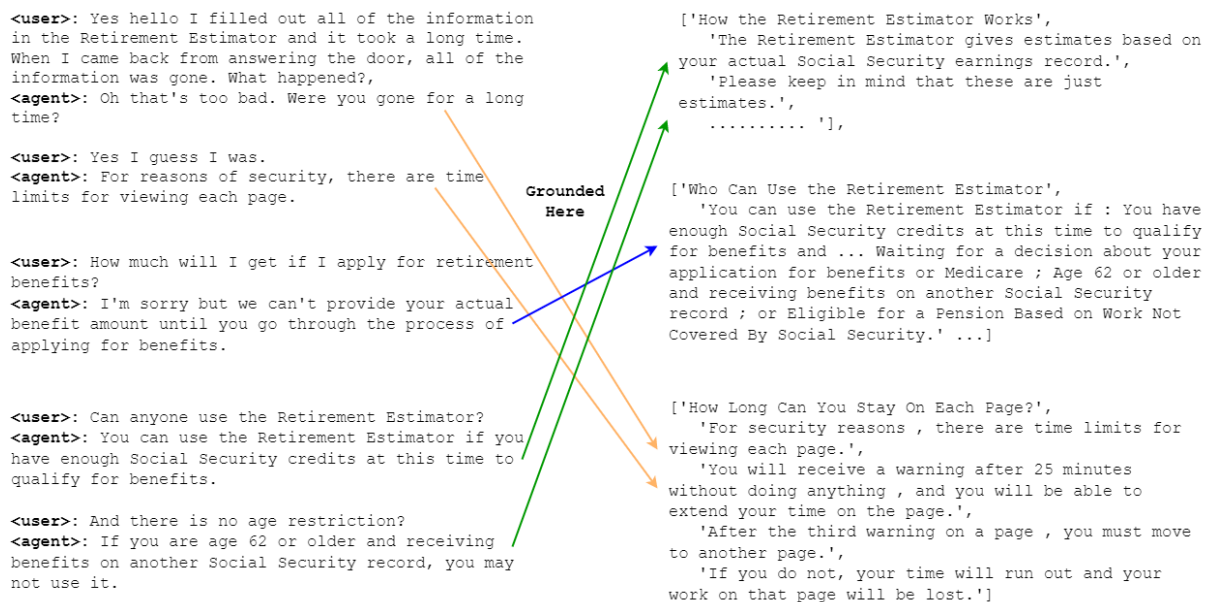


Figure 5: Grounding Knowledge can be based on the same topic as the previous utterance or a new topic. The same coloured lines refer to the same topic.