Knowledge-Based Systems 301 (2024) 112294

Contents lists available at ScienceDirect

Knowledge-Based Systems

journal homepage: www.elsevier.com/locate/knosys

DANTE: Dialog graph enhanced prompt learning for conversational question answering over KGs

Jingyang Li, Shengli Song *, Sitong Yan, Guangneng Hu, Chengen Lai, Yulong Zhou

School of Computer Science and Technology, Xidian University, Xi'an, Shaanxi, 710071, China

ARTICLE INFO

Keywords:
Conversational question answering
Knowledge graphs
Pre-train prompt and predict
Graph prompt

ABSTRACT

In this study, we focus on the task of selecting high-quality answers selection in knowledge-graph-based (KG-based) conversational question answering (ConvQA) system. Effectively exploring a user's intention and modeling historical interaction records are challenging. To address this challenge, we propose the **D**ialog gr**A**ph e**N**hanced promp**T** l**E**arning (**DANTE**) model, which simultaneously integrates sequential and structural information from questions and interactive logic. While the structural information was exploited in previous studies by simply converting it into linear strings in a "pre-train, predict" paradigm, DANTE comprises the use of a novel graph representation for jointly modeling the QA pairs, relevant KG paths, and dialog contexts. The dialog graph constructs in both the turn-level and dialog-level, where DANTE fuses the structural and sequential information deeply in a "pre-train, prompt, and predict" manner. The experimental results showed that DANTE improves the absolute points by 7.1% and 8.2% in terms of the P@1 and mean reciprocal rank metrics, respectively, on the ConvQuestions/ConvRef benchmark compared with state-of-the-art baselines.

1. Introduction

Knowledge-graphs-based (KGs-based) conversational question answering (ConvQA) with multi-turn dialogs is an essential task in information retrieval and human–machine interactions [1–3]. It involves mapping a user's query or utterance in the context of multi-turn dialog historical interaction records to formal queries for correct-answer retrieval or relevant information. With the growing popularity of Artificial General Intelligence (AGI), e.g., ChatGPT [4], LaMDA [5], as researchers, we not only perceive the significant potential of Large Language Models (LLMs), but also the shortcomings of their lack of semantic knowledge and explicit reasoning ability, which is the focus of the KG-based ConvQA task.

Understanding KGs-based ConvQA is challenging owing to the dual question–answer problem. On the question side, the request is mainly intent-implicit. In a multi-turn interaction setting, user goals and information needs can be ambiguous or may evolve throughout the conversation. With implicit intent, the user acquires the necessary knowledge by interacting with the system through multi-turn question–answer pairs, engaging in an exploratory process that often involves an anaphora ellipsis in query sentences [6]. On the answer side, knowledge should be context-aware. In a knowledge-oriented dialog, the user's queries can inherently exhibit complex dependencies on the knowledge mentioned in the dialog context, where questions requiring reasoning or inference based on the context information gathered thus far are possibly raised from the user side. For instance, in Fig. 1, the user could ask "Where is he from originally? " after asking about the author of the book "Moby-Dick". To answer such queries, the system should sufficiently consider the dependency relationships between the current query and preceding dialog context [7,8]. Overall, the selection of the high-quality and most relevant answers is a crucial goal of KGs-based ConvQA.

Since the introduction of ConvQuestion [1], researchers have explored various approaches to KGs-based ConvQA tasks, including graph-based reasoning [9,10], logical form generation models [11], reinforcement learning [12], and hybrid approaches that combine textual and KG-based information [13,14]. Pre-trained language models (PLMs) outperformed other SOTA approaches, including query and context encoding, embedding matching, and transfer learning, which have significantly advanced the KGs-based ConvQA task [15,16]. However, the previous research has exploited the structural information with PLMs by simply converting it into linear strings in a "pre-train, predict" paradigm. However, the challenges remain when dealing with KG-specific reasoning, query formalization, and context integration. Moreover, the capture of sufficient depth and synchronization of modeled sequences and structural information to obtain further semantics remains challenging. Specifically, it refers to the process of identifying and capturing synchronized representational patterns using shared cues and entities across the three elements involved within each dialog turn.

* Corresponding author.
E-mail address: shlsong@xidian.edu.cn (S. Song).

https://doi.org/10.1016/j.knosys.2024.112294
Received 31 August 2023; Received in revised form 1 July 2024; Accepted 28 July 2024
Available online 31 July 2024

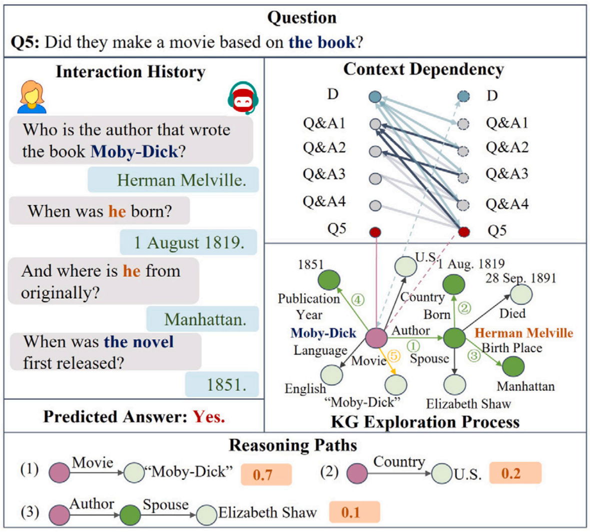

Fig. 1. Motivating example illustrating a sample conversation. With conversational interaction history (a), our proposed approach DANTE constructs a dialog graph that captures context dependency and (b) the KG exploration process (c) simultaneously to reason over KG paths for KGs-based ConvQA tasks. The predicted answer for Q5 is *"Yes"* based on the reasoning path *"<head: Moby-Dick, predict: Movie, tail: "Moby-Dick">"*.

Prompt-based PLMs methods have demonstrated remarkable autonomous capabilities across numerous NLP tasks [17,18]. We focus on the problem of integrating sequential and structural information deeply and simultaneously in a "pre-train, prompt, and predictive " manner. In this paper, we present the **D**ialog gr**A**ph e**N**hanced promp**T** l**E**arning (**DANTE**) model, which integrates sequential and structural information simultaneously from questions and interactive logic, respectively, to obtain the semantic benefits of both types of information to improve the high-quality answer-selection performance of the KG-based ConvQA system. Specifically, (1) capturing the question–answer (QA) and relevant KG paths dialog context simultaneously by constructing a dialog graph representation; (2) boosting the PLM with both sequential and graph structure information in the prompt learning mode; and (3) for high-quality answer selection, training DANTE in a dialog graph enhanced the prompt multi-task learning paradigm.

The main contributions of this study are as follows:

- We present DANTE, which designs a graph-enhanced prompt learning paradigm that can jointly train the model in a "pre-train, prompt, predict" manner with a multi-task objective. To the best of our knowledge, DANTE is the first approach that incorporates prompt learning in a KGs-based ConvQA system.
- DANTE constructs a novel graph representation for jointly modeling the sequential (fluent question–answer pairs) and structural (entire dialog history and context-aware knowledge in KGs) information at both the turn-level and dialog-level to obtain profound semantic fusion and explicit representation for the KGs-based ConvQA task.
- DANTE realizes high-quality and high-relevance answer-selection performance in the KGs-based ConvQA system, which improves the absolute points by 7.1% and 8.2% in terms of the P@1 and mean reciprocal rank metrics on the ConvQuestions benchmark compared with SOTA baselines.

2. Related work

In this section, we investigate the extant research on KGs-based ConvQA and focus on the studies most relevant to our proposed approach, including graph approaches on conversational tasks and prompt learning approaches.

2.1. KGs-based ConvQA

ConvQA has attracted significant attention and demonstrated significant potential [19]. However, many challenges have been identified, including complex question analysis and large-scale knowledge queries [20–23], thus giving rise to the KGs-based ConvQA to answer questions based on the provided graph-based structured knowledge, which in turn has the potential to revolutionize how humans interact with machines. The majority of recent studies on KGs-based ConvQA have employed the semantic parsing approach [11,24–26] and multi-task learning paradigm [3,27,28] to answer conversational questions.

For KG path ranking, Christmann et al. [1] proposed a graph exploration approach CONVEX, which involves addressing conversational inquiries within a KG by retaining the contextual flow of conversation through tracked entities and predicates, thereby automatically deducing absent or uncertain elements for subsequent queries. CONVEX tactfully extends the frontier to discover and prioritize potential answers to the provided questions. Bi et al. [29] proposed UMRNet based on an attention redistribution mechanism that was capable of handling the mapping problems between the question and KG relations. Kacupaj et al. [14] presented a contrastive representation learning-based approach to rank KG paths effectively. Kaiser et al. [12] introduced a reinforcement learning (RL) model that conceptualizes the process of answering as multiple agents traversing the KG in parallel. Bi et al. [30] proposed an effective model for handling multi-hop KG path reasoning tasks under weak supervision settings based on reward integration and policy evaluation. These agents learn from a continuous flow of conversational questions and their revised forms.

2.2. Graph approaches on conversations

Graph data have functioned as organized knowledge repositories in numerous systems by explicitly modeling the interactions among different entities, therefore, they are frequently employed in conversational tasks to infer contextual and commonsense knowledge.

Graph on Visual Dialog. To address the visual dialog task, Zheng et al. [31] proposed a graph-based approach that explicitly formalizes the task as an inference problem within a graphical model featuring partially observed nodes and unknown graph structures. Schwartz et al. [32] developed a factor-graph-based attention mechanism that operates on data utilities to resolve the details and nuances on visual dialog.Guo et al. [33] presented a Context-AwareGraph neural network that can iteratively update the graph structure using an adaptive top-K message passing mechanism to model the underlying context-aware relation inference.

Graph on Task-oriented Dialog. Graph-based approaches on task-oriented dialog focus on efficiently incorporating knowledge into end-to-end task-oriented dialog systems. In recent studies, Zhao et al. [34] explore a dialog state GAT consisting of a dialog context subgraph and an ontology schema subgraph to alleviate the cross-domain slot sharing issue. Wu et al. [35] proposed a Graph Memory Network-based (GMN-based) Seq2Seq model, GraphMemDialog, to acquire the latent structural insights concealed within the dialog history and to depict the dynamic interaction between the dialog history and knowledge bases. Yang et al. [36] studied target-oriented dialog using a commonsense KG and designed a global reinforcement learning framework that incorporates planned paths. This approach facilitates adaptable adjustments to the local response generation model to align with a global target.

Graph on KGs-based ConvQA. For graph-based inference on the KGs-based ConvQA task, Chen et al. [37] presented a graph-learning technique that can effectively capture conversational flow in dialogs, constructing a history-aware context graph of each conversation turn. Yasunaga et al. [38] presented a QA-GNN model that enables LMs to identify pertinent information from extensive KGs and execute unified reasoning across both the question-answering context and KGs. For ConvQA over heterogeneous sources, Christmann et al. [39] constructed a heterogeneous graph from multiple knowledge sources that was then iteratively reduced using GNN-incorporated question-level attention.

2.3. Prompt learning approaches

Prompt Learning on NLP Tasks. In recent research, the new training paradigm named "pre-trained, prompt, and predict" has been proposed and described as the key to LLMs in real-world tasks [18], and it allows the language models to be pre-trained on extensive volumes of unprocessed text. The model can realize few-shot or even zero-shot learning by formulating a novel prompting function and adapting to new scenarios with limited labeled data. The prompt learning of NLP tasks comprises prompt template engineering [17,40], prompt answer engineering [41,42], multi-prompt learning [43,44], and prompt-based training strategies [45]. In recent studies, prompt learning has been applied to perform cross-lingual relation extraction [46], knowledge transfer [47,48], few-shot dialog state tracking [49,50], and task-oriented dialog domain adaptation [51].

Graph Prompting Methods. In contrast to conventional methods, graph-prompting functions can induce task-specific contexts and apply templates enriched in structured knowledge [52]. Liu et al. [53] introduced a pre-training and prompting framework called GraphPrompt, which integrates pre-training and downstream tasks within a unified task template. Liu et al. [54] found that the incorporation of external knowledge benefits commonsense reasoning and developed generated knowledge prompting from a language model to provide additional knowledge when answering a question.

Inspired by the excellent performance of graph-prompting methods, we propose a dialog graph-enhanced prompt-learning approach for the KGs-based ConvQA. In a "pre-trained, prompt, and predictive " manner, the proposed DANTE model simultaneously captures the semantic benefits of sequential and structural information to improve the selection performance of high-quality answers.

3. Approach

3.1. Notation

Before introducing our approach, we first list DANTE's notation and meaning in terms of four aspects: conversation, knowledge graph, conversation graph, and model, as shown in Table 1. For the KGs-based ConvQA task, given a KG denoted by K, a user request q^t, conversation context C^t, and context entity set E_c^t, all feasible paths $\mathcal{P}_c^t$ within the KG are extracted. We model the problem as an information retrieval task that requires a score and rank $\mathcal{P}_c^t$ to select the most relevant context paths $p_c^t \in \mathcal{P}_c^t$ that result in entities or literals that correspond to the correct answer a^t as the answer to the question q^t.

3.2. Model overview

In a KGs-based ConvQA task, the input data include question q^t and answers a^t that are extracted from the KG. The proposed DANTE generates answers according to the following four steps: (1) The entire conversation history sequences s^t with prefix prompts and the dynamic dialog graph $\mathcal{G}^t$ are taken as model inputs. (2) The dialog representation $\mathcal{G}^{t\prime}$ is learned using the topic classification task. (3) The dialog topic information is used, and s^t is encoded, $\mathcal{P}$ is encoded as $s^{t\prime}$, $\mathcal{P}_c^{t}$ is encoded to calculate each candidate KG path via the cosine similarity and ranking it to select the answer from the highest-scoring path. (4) A fluent answer is generated via the PLM decoder with sequential information $s^{t\prime}$ and learned structure information $\mathcal{G}^{t\prime}$ simultaneously.

Table 1
Notation used herein and their meanings.

	Notation	Meaning
Conversation	C	Conversation context
	t	Dialog turn
	q^t	Question at turn t
	a^t	Answer at turn t
	C^t	Interaction history of C at turn t
Knowledge Graph	$\mathcal{K}$	Knowledge graph
	$\mathcal{E}$	Entities
	$\mathcal{T}^+$	Triples
	$\mathcal{P}_c$	Context KG paths
	$\mathcal{D}^{t+}$	Set of positive context paths for q^t
	$\mathcal{D}^{t-}$	Set of negative context paths for q^t
Dialog Graph	$\mathcal{G}^t$	Dialog graph at turn t
	$\mathcal{V}_C^t$	Turn-level vertex at turn t
	$\mathcal{V}_D$	Dialog-level vertex representing entire dialog history
	$\mathcal{V}_\mathcal{K}^t$	Vertex in $\mathcal{K}$ at turn t
	$\mathcal{V}_\mathcal{F}^t$	Vertex representing current focal entity
Model	s^t	Input sequence (contains C^t and q^t)
	θ^p	Prompt learning parameter
	$s^{t\prime}$	Embedding of input sequence s^t
	$\mathcal{P}_c^{\prime}$	Embedding of context KG paths $\mathcal{P}_c$
	$\mathcal{G}^{t\prime}$	Embedding of dialog graph $\mathcal{G}^t$
	$\phi^{S^{t\prime}}$	Joint embedding for s^t and $\mathcal{G}^t$
	$\phi^{\mathcal{P}_c{}^{\prime}}$	Joint embedding for $\mathcal{P}_c$

Three tasks are jointly trained in an end-to-end manner, as shown in Fig. 2, which includes a fluent answer-generation task, a contrastive KG-path ranking task, and a topic classification task. Furthermore, the implementation choices of the PLM (e.g., GPT2 [55]) or topic classification GNN (e.g., GCN [56]) are used for training and empirical effectiveness using appropriate techniques proposed to solve the subtasks in DANTE, such that the approach can be implemented with other choices.

3.3. Dialog graph construction

Graph Definition. Dialog graphs contain explicit relationships amo ng QA pairs, relevant KG paths, and dialog histories. From a technical perspective, the essence of a dialog graph is multi-hop reasoning and question-context-KG co-reference. Inspired by Guo et al. [33], we address this problem by adaptively capturing the related question-context and question-KG cues in the dynamic co-reference mode. Four types of nodes are defined: (1) the KG node $\mathcal{V}_\mathcal{K}^t$, which represents entities $\mathcal{E}$ in the KG $\mathcal{K}$; (2) the turn-level node $\mathcal{V}_C^t$ that represents the conversation interaction history with the question-and-answer information at each dialog turn t; (3) dialog-level node $\mathcal{V}_D^t$, representing all the dialog histories since turn 1 to turn t; (4) focal entity node $\mathcal{V}_\mathcal{F}^t$, representing the current focal entity, which should also be updated to follow the rule of having only one focal entity node in each graph.

To connect these nodes dynamically, four types of edges are defined. (1) Focus edges are undirected edges between a focal entity and its one-hop neighbors to establish the connection between the dialog context nodes $\mathcal{V}_C^t$ and $\mathcal{V}_\mathcal{K}^t$. (2) Context dependency edges are directed edges between turn-level nodes, from a newer node $\mathcal{V}_C^t$ point to the referred dependency turn-level node, e.g., $\mathcal{V}_C^{t-1}$ for adding an ellipsis to q^t. (3) Dialog edges are dual-directed edges between turn-level nodes $\mathcal{V}_C^t$ and $\mathcal{V}_D^t$, $\mathcal{V}_\mathcal{F}^t$, and $\mathcal{V}_D^t$ for representing and tracking the dialog topic. (4) Argument edges are undirected edges between KG nodes $\mathcal{V}_\mathcal{K}^t$ for exploiting deep KG paths. An example of a vivid dialog graph corresponding to Fig. 1 is presented in Fig. 3.

Feature Representation. The four types of nodes in the dialog graph originate from different concepts that require varying rules for embedding initialization.

For turn-level nodes $\mathcal{V}_C^t$, we concatenate the question and answer sequences at the current turn t with special tokens to generate the

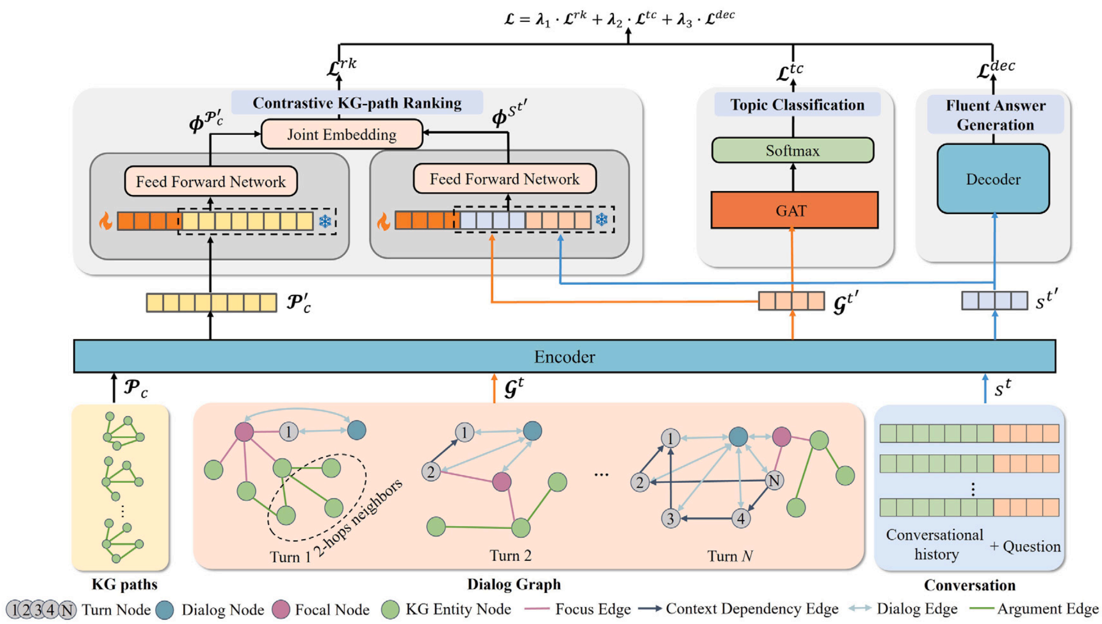

Fig. 2. DANTE architecture. Three types of embeddings are sent into the encoder as the input: (1) KG paths, (2) a dialog graph, and (3) the conversational sequence. DANTE follows a multi-task learning paradigm jointly trained from three sub-tasks, including contrastive KG-path ranking, topic classification, and fluent answer generation. The optimal objective $\mathcal{L}$ is derived through a careful balancing of the optimization weights among the three components, thus enhancing and maximizing the final answer output performance.

QApair sequence s_{qa}, which is then fed into the PLM to encode turn-level information. We only use embedding at the $[cls]$ index, while considering its consistent dimensions with other graph nodes and its comprehensive context representation at each turn.

$$s^t_{qa} = [[cls], q^t, [sep], a^t, [eos]] \tag{1}$$

For dialog-level nodes $\mathcal{V}^t_D$, an initial embedding is randomly assigned with an invariant random seed to enable adequate learning of additional information from other nodes during the message-passing process. The dialog-level node is connected to the majority of the nodes in the dialog graph, therefore, it can develop a global representation of the graph. Here, we introduce a topic classification subtask to efficiently guide the GNN learning process. Topic classification incorporates dialog-level node embedding after message passing, which effectively enhances the primary task with an appropriate weight. Further experiments demonstrate that this sub-task can be considered a "prompt tuning" for the ranking and generation object.

For KG nodes $\mathcal{V}^t_{\mathcal{K}}$ and focal nodes $\mathcal{V}^t_{\mathcal{F}}$, we use the prefix template $\mathcal{V}^t_{\mathcal{K}}\prime$ and $\mathcal{V}^t_{\mathcal{F}}\prime$ to prompt the PLMs to encode their embedding, and for a dialog-level node $\mathcal{V}^t_D$, we use a topic word token to initialize it. The corresponding details are presented in the following section. The feature representations can then be formulated as follows:

$$\mathcal{G}^{t'} = \textbf{PLMEncoder}(s^t_{qa}, \mathcal{V}^t_{\mathcal{K}}\prime, \mathcal{V}^t_{\mathcal{F}}\prime, \mathcal{V}^t_D) \tag{2}$$

Graph Representation. With graph definition and feature representation, dialog graph representation was learned in two steps: (1) dialog graph construction and (2) message passing.

Dialog graph construction is considered to be a dynamic evolving process in which links between specific nodes $\mathcal{V}$ can be obtained by updating the corresponding positions in the adjacency matrix $\mathcal{A}dj \in \mathbb{R}^{N*N}$. N denotes the nodes in the Dialog Graph. Algorithm 1 outlines the high-level pseudo-code for the step-by-step building process of the dialog graph $\mathcal{G}^t$.

For messages passing through the dialog graph, we employed the state-of-the-art GNN model GAT. Each node learns features from its neighbors with different attention levels by traversing the dialog inter action–focus transitions turn-by-turn. The message-passing procedure can be calculated as follows:

$$\mathcal{G}^{t'} = \textbf{GAT}(\mathcal{G}^{t'}, \ \mathcal{A}dj^t) \tag{3}$$

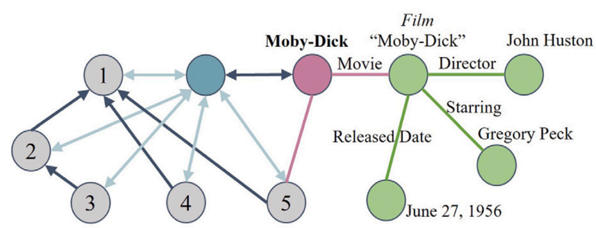

Fig. 3. Dialog graph representation of Q5 in Fig. 1. Context dependency is represented with directed edges between turn-level nodes, and KG entities are connected by undirected edges. The dialog node indicated in cyan connects turn nodes and the focal entity *Moby-Dick* node with dual directed edges.

where $\mathcal{A}dj^t$ is the adjacency matrix of the dialog graph at turn t.

3.4. Prompt engineering

In this study, we designed hard and soft prompts for contrastive ranking and fluent answer generation tasks.

Hard Prompt for Language Model Inputs. As shown in Fig. 2, the language model input comprises three parts: KG paths, conversational history sequences, and dialog graphs.

(1) KG paths are input to obtain a representation for every $\mathcal{P}_c \in \{\mathcal{D}^{t+} \cup \mathcal{D}^{t-}\}$. In this context, we identify the entities $\mathcal{E}_c$ and extract possible candidates for KG paths, which is similar to the approach described in Kaiser et al. [12]. To preserve the original KG path information, the triples $\mathcal{T}^+$ are concatenated with special tokens and fed into the PLM to encode the sentence embedding $\mathcal{P}'_c$:

$$\mathcal{P}'_c = \textbf{PLMEncoder}([[cls], \mathcal{T}^+_1, [sep], \mathcal{T}^+_2, [sep], \ldots, \mathcal{T}^+_n]) \tag{4}$$

(2) Conversational history sequences serve as the input for the PLM to produce fluent answers in an end-to-end fashion. Specifically, the conversation sequence s^t comprises the entirety of the historical QA pairs along with the current question q^t. This comprehensive structure enables the efficient organization of the question–answer prompt template, as outlined in Eq. (5). Subsequently, this prompt template is input into the PLM, which encodes it into a sentence embedding $s^{t\prime}$ and generates a fluent answer sequence $a^{t\prime}$.

$$s^{t\prime} = \textbf{PLMEncoder}([[cls], q^1, [sep], a^1, [sep], \dots, q^t, [mask]]) \tag{5}$$

$$a^{t\prime} = \textbf{PLMDecoder}(s^{t\prime}) \tag{6}$$

(3) Dialog graph $\mathcal{G}^t$. In the previous section, three types of nodes were defined using prompt templates to obtain improved initial node feature embeddings. For turn-level nodes $\mathcal{V}_C^t$, we organized them in a turn-level question–answer prompt template and fed them into the PLM to obtain a node embedded at the $[mask]$ token position (only the current turn question). For the focal entity nodes and KG entities, the triple $\tau \in \mathcal{D}^{t+}$ in the current dialog graph is used to formulate a head-predict-object prompt template. $\mathcal{T}_{head}^{\mathcal{V}_\mathcal{F}^t}$ and $\mathcal{T}_{predict}^{\mathcal{V}_\mathcal{F}^t}$ denote the head entity and predicate in the triples where the object is $\mathcal{V}_\mathcal{F}^t$. Meanwhile, $\mathcal{T}_{head}^{\mathcal{V}_\mathcal{K}^t}$ and $\mathcal{T}_{predict}^{\mathcal{V}_\mathcal{K}^t}$ represent the head entity and predicate in the triples where the object is $\mathcal{V}_\mathcal{K}^t$. These entities are all from the context KG paths $\mathcal{P}_c$ corresponding to the current turn of the answer $a^{t\prime}$. The prompt template is then fed into the PLM to obtain a node feature embedded at the $[mask]$ token position. For dialog-level nodes, $\mathcal{V}_D^t$ is a randomly initialized encoding, as mentioned previously.

$$\mathcal{V}_C^{t\prime} = \textbf{PLMEncoder}([[cls], q^t, [sep], a^t, [mask]]) \tag{7}$$

$$\mathcal{V}_\mathcal{F}^{t\prime} = \textbf{PLMEncoder}([[cls], \mathcal{T}_{head}^{\mathcal{V}_\mathcal{F}^t}, \mathcal{T}_{predict}^{\mathcal{V}_\mathcal{F}^t}, [mask]]) \tag{8}$$

$$\mathcal{V}_\mathcal{K}^{t\prime} = \textbf{PLMEncoder}([[cls], \mathcal{T}_{head}^{\mathcal{V}_\mathcal{K}^t}, \mathcal{T}_{predict}^{\mathcal{V}_\mathcal{K}^t}, [mask]]) \tag{9}$$

Algorithm 1: Dialog Graph Construction

input : Question q^t, answer a^t, KG triples $\mathcal{T}^+$, positive context paths $\mathcal{D}^+$, negative context paths $\mathcal{D}^-$, and entity dictionary ED
output: Dialog graph $\mathcal{G}^t$

```
initialize Dialog graph G^t with turn-level context nodes V_C, dialog-level nodes V_D, and entities E in QA pairs;
repeat
    //create connection between V_C^t and V_D
    add new V_C^t node and edge (V_C^t, V_D) to G^t at turn t;
    //create connection between V_C nodes
    update entity dictionary ED for each V_C node and all entities detected in (q^t ∪ q^{t-1} ∪ a^{t-1});
    for V_C^k, entity E^k ← ED do
        if entities of V_C^t in E^k then
            add edge (V_C^t, V_C^k) to G^t;
            set the direction of the edge from V_C^t to V_C^k;
        end
    end
    //update focal entity
    if new entity E^t occurs in (D^+ ∩ V_C^t) then
        set E^t as the focal entity E_K^t and update the focal entity representation;
        add edges (E_K^t, V_C^t) and (E_K^t, V_D)
    else
        set the focal entity E_K^{t-1} as the focal entity E_K^t;
        add edge (E_K^t, V_C^t);
    end
    //expand candidate KG paths
    add two-hop neighbor nodes and edges of E_K^t in the KG paths ⊆ {D^{t+}, D^{t-}} to G^t ;
until there is no V to add or G^t is sufficiently large;
```

Soft Prompt for Contrastive Ranking Embeddings. To achieve the objective of high-quality answer selection, we use a contrastive ranking task such as PRALINE [14] and improve its effectiveness with soft prompt tuning params. As illustrated in Fig. 2, for the contrastive learning process, the features produced by the PLM are all frozen, and a learnable prompt vector θ^p and readout operation are introduced to generate a subgraph representation for both contrastive sizes [53]. This intuitively enhanced the ability to learn a better representation and rich semantic information, which proved to be effective in our ablation study experiment.

$$\mathcal{P}_C^{t\prime} = \textbf{ReadOut}(\mathcal{W}[\theta_C^p \oplus \mathcal{P}_C^{t\prime}]) \tag{10}$$

$$\mathcal{S}^{t\prime} = \textbf{ReadOut}(\mathcal{W} \cdot [\theta_S^p \oplus (s^{t\prime}; \mathcal{G}^{t\prime})]) \tag{11}$$

The $\oplus$ is a plus operation. The ReadOut operation used is a sigmoid function that can be replaced by other activation functions.

3.5. Dialog-graph-enhanced multi-task prompt learning

As discussed in Section 4.1, DANTE consists of three modules to which a joint objective loss function can be applied.

Topic Classification Task. For the topic classification task, we used a GAT network for message passing and trained it in a dialog graph node classification manner. At time step t, the GAT model input was a node feature set $h = \{\vec{h}_1, \vec{h}_2, \dots, \vec{h}_N\}, \vec{h}_i \in \mathbf{R}^F$, which was obtained from $\mathcal{G}^t$. N denotes the number of KG nodes, the maximum of which is predefined based on the training data. F denotes the number of features in each node. In particular, we observed that the two-graph attention layer GAT exhibited the best performance, mainly because the two-hop neighbor had the greatest impact on the topic classification task. After the message passing with $\mathcal{A}dj^t$ in GAT, the node feature h' is aggregated as follows:

$$\vec{h}_i' = \sigma\left(\frac{1}{K}\sum_{k=1}^{K}\sum_{j\in N_i}\alpha_{ij}^k \mathbf{W}^k \vec{h}_i\right) \tag{12}$$

The dialog node feature is then fed into a Softmax layer to learn the topic classification. In ConvQuestions and ConvRef, we take the five domains (music, movie, tv_seris, soccer, and books) as the topic vocabulary $V^{Topic} = \{T_1, ..T_m\}$ for each conversation. The loss object of the topic classification task is formulated as follows:

$$\mathcal{L}^{tc} = -\sum_{j=1}^{m} logp(y_j^{Topic} | \vec{h}_i') \tag{13}$$

where $y_j^{Topic} \in V^{Topic} = \{T_1, ..T_m\}$ denote the gold labels.

Contrastive KG-path Ranking Task. The contrastive ranking module proposed by Kacupaj et al. [27] is used by employing two identical sequential networks to produce combined embeddings for $\mathcal{S}^{t\prime}$ and $\mathcal{P}^{t\prime}$ at turn t, where both sides contain a two linear layers feedforward network with a $Relu$ activation and are appended with a $tahn$ non-linear layer. During the training procedure, the module calculates the cosine similarity among all the potential candidates within the current batch. The two feedforward networks are trained collaboratively to enhance the similarity for the correct pairs and diminish the similarity for incorrect pairs. The loss object of the contrastive ranking task is formulated as follows:

$$\mathcal{L}^{rk} = \begin{cases} 1 - cos(\phi^{\mathcal{S}^{t\prime}}, \phi^{\mathcal{P}^{t\prime}}), & if\ y^{(rk)} = 1, \\ max(0, cos(\phi^{\mathcal{S}^{t\prime}}, \phi^{\mathcal{P}^{t\prime}}) - \alpha) & if\ y^{(rk)} = -1 \end{cases} \tag{14}$$

where $y^{(rk)} \in \{1, -1\}$ is the ground truth label for the ranking module. $cos(\doteq)$ refers to the normalized cosine similarity, and α is the margin. $\phi^{S^{t\prime}}$ is the joint embedding for s^t and $\mathcal{G}^t$, and $\phi^{\mathcal{P}_c{}'}$ is the joint embedding for $\mathcal{P}_c$.

Fluent Answer Generation Task. Fluent answers are typically generated based on the PLM sequence generation process, which is regarded as a fundamental training task of the PLM. The loss object can be formulated as

$$\mathcal{L}^{dec} = -\sum_{l=1}^{n} logp(y_l^{dec} |\ s, \mathcal{G}) \tag{15}$$

Table 2
Statistics of ConvQuestions and ConvRef datasets. ConvRef is an extension of ConvQuestions.

statistic	ConvQuestions	ConvRef
Number of domains	5	5
Number of dialogs	11 200	11 200
Number of reformulations	–	205 000
Training-dev-test split	6720/2240/2240	6720/2240/2240
Number of turns per dialog	5	5
Total number of turns	56 k	262 k

where $y_l^{dec} \in V^{(dec)}$ are the gold labels used to generate a for the decoder

A joint learning object was applied to the three tasks. To train all the modules and tasks simultaneously, we used a weighted average of individual losses formulated as follows:

$$\mathcal{L} = \lambda_1 \cdot \mathcal{L}^{rk} + \lambda_2 \cdot \mathcal{L}^{tc} + \lambda_3 \cdot \mathcal{L}^{dec} \tag{16}$$

$\lambda_1, \lambda_2, and \lambda_3$ are the hyperparameters for adjusting the bias of the modules in DANTE during the training process.

4. Experiments

4.1. Datasets

In this study, we focused on a KGs-based ConvQA task and conducted experiments on two large-scale multi-domain datasets: ConvQuestions and ConvRef. ConvQuestions, a ConvKBQA dataset created on Wikidata by crowdworkers on Amazon Mechanical Turk, comprises 11,200 dialogs over five domains: “Movies”, “TV Series”, “Music”, “Books”, and “Soccer”. Each conversation has a five-turn conversation with its ground-truth answers. ConvRef builds on the aforementioned dataset ConvQuestions by incorporating reformulations. It consists of 11,200 conversations with approximately 205,000 reformulations. The average reformulation length was approximately 7.6 words, whereas the initial questions per session had an average length of approximately 6.7 words. Both the datasets were divided into 6720, 2240, and 2240 data for the training, validation, and testing, respectively.(see Table 2).

4.2. Implementation & training details

We used the BART-based model as our PLM encoder–decoder model with a dropout of 0.1 and set the word embedding dimension (d = 768), and the number of layers ($L = 2$) of our GAT module. With the dropout rate set as 0.2 for each layer, the node feature dimensions were set to be the same as the word-embedding size (d = 768). During the training process, we restricted DANTE’s input sequence size to 150 tokens, which was sufficient for our task. For topic classification and generation sub-tasks, we select the relative weights λ_2 and λ_3 from {0.25, 0.1, 0.05}. For the cosine embedding loss in the path ranking, we employ a margin of α equal to 0.1 and a relative weight λ_1 of 1.0. We set the batch size as 64 and the learning rate from the {1e-3, 1e-4, 2e-5} optimizer using a Single GPU (NVIDIA V100 GPU), which required ∼48 h after ∼150 epochs for ConvQuestions and ∼5 days after ∼300 epochs for ConverRef to achieve the best performance.

4.3. Baselines

The proposed DANTE model was compared with the following state-of-the-art methods that were most relevant to our approach.

- **CONVEX** [1]: It infers the missing parts of the incomplete questions with the conversation history and then uses a graph exploration algorithm to search for candidate answers.
- **FOCAL ENTITY** [9]: It is a graph-based model focused on the transitions of implied entities from the conversation history.

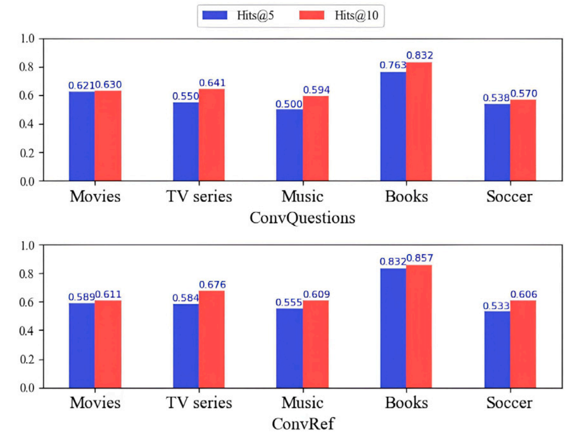

Fig. 4. DANTE’s ranking performance on all the five domains in terms of H@5 and H@10. The model achieves the best results in the *Books* domain on both the *ConvQuestion* and *ConvRef* datasets, and the blue and red columns represent H@5 and H@10, respectively.

- **OAT** [11]: This method defines a new logical form (LF) grammar and incorporates conversational contexts and KGs.
- **CONQUER** [12]: It relies on reinforcement learning and uses implicit negative feedback that arises when users rephrase questions that were unsuccessful in previous attempts.
- **PRALINE** [14]: It formulates KGs-based ConvQA as a KG path ranking problem and models the conversational context and KG paths to jointly learn embedding representations.
- **KRR** [57] : This method first rewrites the question based on the historical conversation with the supervision of transferring the knowledge base and then runs over the knowledge base to obtain the answer.
- **EXPLAIGNN** [39]: With heterogeneous graph neural networks that incorporate question-level attention, EXPLAIGNN obtains competitive performance. Moreover, integrating heterogeneous sources (KB, text, tables, and infoboxes) into the heterogeneous graph can substantially improve the answer performance.

4.4. Metrics

To evaluate the model performance on a KGs-based ConvQA task, we employed the same metrics as in previous studies. Firstly, we used Precision at 1 (P@1) to report the proportion of precise top-ranked answers. We also used the mean reciprocal rank (MRR) and Hit at 5 (H@5), which is the proportion of correct answers in the top five positions. The precision, recall, and F1 scores were used to conduct the domain identification task, and BLEU-4 and METEOR were used for answer generation.

4.5. Results

Overall Performance on ConvQuestions and ConvRef datasets. Table 3 provides a summary of the results, offering a comparison between DANTE and previous baseline methods. DANTE outperformed the baselines in almost all the metrics for both datasets. Here, we introduce the EXPLAIGNN, the latest but not officially published approach. Specifically, for P@1, DANTE achieved the third-best performance. For H@5 and MRR, the margins are prominent, with 4.9% and 3.6% points better than the latest EXPLAIGNN on ConvQuestion. On the ConvRef dataset, DANTE outperforms CONVEX, CONQUER, and PRALINE across all the metrics, where the margin for all the metrics is more than 2% absolute points. Moreover, it surpasses PRALINE on P@1 and MRR by 7.1% and 8.2% points, respectively.

Table 3
Overall results on the two datasets. The incorporation of the dialog graph and prompt learning within the DANTE framework yielded notable improvements in empirical outcomes. These enhancements resulted in superior results compared to the majority of the baseline methods. The EXPLAINGNN (heterogeneous sources) uses additional information to realize a better performance. The best values are indicated in bold.

Dataset	ConvQuestions			ConvRef		
Model	P@1	H@5	MRR	P@1	H@5	MRR
CONVEX [1]	0.184	0.219	0.200	0.225	0.257	0.241
FOCAL ENTITY [9]	0.248	0.248	0.248	–	–	–
OAT [11]	0.166	–	0.175	–	–	–
OAT [11] (gold seed entity)	0.250	–	0.260	–	–	–
CONQUER [12]	0.240	0.329	0.279	0.353	0.429	0.387
PRALINE [14]	0.294	0.464	0.373	0.335	0.599	0.441
KRR [57] (gold seed entity)	**0.397**	0.397	0.397	–	–	–
EXPLAIGNN [39] (KB-only)	0.330	0.480	0.399	–	–	–
EXPLAIGNN [39] (heterogeneous sources)	0.363	0.546	0.447	–	–	–
DANTE	0.352	**0.595**	**0.483**	**0.406**	**0.619**	**0.523**

Table 4
Detailed outcomes across various domains within both benchmarks, assessed using ranking metrics. DANTE achieves better results in eight of ten scenarios. The best values are indicated in bold.

Dataset	ConvQuestions									
Domain	Movies		TV Series		Music		Books		Soccer	
Models	H@5	MRR	H@5	MRR	H@5	MRR	H@5	MRR	H@5	MRR
CONVEX [1]	0.355	0.305	0.269	0.218	0.293	0.237	0.303	0.246	0.284	0.234
CONQUER [12]	0.357	0.316	0.382	0.325	0.320	0.263	0.464	0.417	0.310	0.268
PRALINE [14]	0.561	0.426	0.457	0.378	0.405	0.279	0.739	0.599	0.492	0.344
DANTE	**0.621**	**0.530**	**0.550**	**0.484**	**0.500**	**0.365**	**0.763**	**0.626**	**0.538**	**0.408**
Dataset	ConvRef									
Domain	Movies		TV Series		Music		Books		Soccer	
Models	H@5	MRR	H@5	MRR	H@5	MRR	H@5	MRR	H@5	MRR
CONQUER [12]	0.436	0.405	0.442	0.392	0.398	0.357	0.554	0.502	0.360	0.316
PRALINE [14]	0.567	0.429	0.545	0.466	0.495	0.329	**0.835**	0.659	**0.564**	0.378
DANTE	**0.589**	**0.531**	**0.584**	**0.525**	**0.555**	**0.410**	0.832	**0.720**	0.533	**0.427**

Ranking Performance Across Domains. In addition, we explored the ranking performance of DANTE across various domains within both benchmarks. Table 4 lists the comprehensive ranking outcomes for the metrics H@5 and MRR, both of which are ranking evaluation measures. On the ConvQuestions benchmark, DANTE outperforms on every domain by at least 2.4% points on H@5 and 2.7% points, which proves that DANTE facilitates the retrieval of more relevant and high-quality answers, owing to its effective modeling of the KG paths and conversational history interactions simultaneously with a deep “prompt” manner. On the ConvRef benchmark, DANTE still outperforms in the majority of domains and approaches the best results of the SOTA method. Our model is also competitive from both sides of the metrics H@5 and MRR.

Fig. 4 presents the ranking results for the H@5 and H@10 ranking metrics. DANTE achieves the best result in the “Books” domain, with H@10 of 0.857 on the ConvRef benchmark, while achieving the lowest scores in the “Music” domain. The results for H@5 and H@10 were consistently favorable across the majority of domains, suggesting that DANTE exhibits a tendency to rank accurate paths at the forefront of the list, thus showcasing the resilience of our approach.

In conclusion, the use of a dialog graph enhanced prompt learning on contrastive learning to rank KG paths positively impacts the overall empirical performance of DANTE. Furthermore, as illustrated in the following section, a dialog graph with prompts plays a significant role in substantially improving the results.

4.6. Ablation study

To explore the effectiveness of DANTE, we performed various ablation studies from the perspective of the three sub-tasks and reported the results in Table 5.

Table 5
Effectiveness of dialog graph, prompted fluent generation, and joint learning. The first row (from top) presents the results of DANTE with all available modules. The second row removes the dialog graph module. The third row removes the fluent answer generation with the prefixed “prompt.” In the last row, we present the outcomes achieved when training the modules separately, thus highlighting the benefits of collaborative training.

Dataset	ConvQuestion		
Model	P@1	H@5	MRR
DANTE	0.352	0.595	0.483
w/o Dialog Graph	0.290	0.489	0.398
w/o Fluent Generation	0.324	0.531	0.426
Train Separately	0.269	0.455	0.385

Effect of Dialog Graph. We first study the empirical advantage of the dialog graph that models the QA pairs, conversation interactions, and KG paths, which is designed to empower a deep, explicit, and semantic-rich representation for the downstream tasks with a “pre-train, prompt, predict” manner in our approach. Hence, we created a DANTE configuration (w/o Dialog Graph) to train the model without a dialog graph module (dialog graph and GAT). As a result, an obvious decrease is observed in the performance of DANTE (without a dialog graph), which demonstrated that DANTE is an effective method for KGs-based ConvQA tasks.

Effect of Fluent Answer Generation. Generation task with prefixed “prompt” inputs is also a key component. To demonstrate the effectiveness of fluent answer generation with prefixed “prompt” inputs, we conducted an ablation experiment by eliminating them and instead using standalone answers extracted directly from the knowledge graph (w/o Fluent Generation). As illustrated in Table 5, we obtained decreases of 2.8% for P@1, 6.4% for H@5, and 5.7% for

Table 6
Representative cases of DANTE on ConvQuestions. The numbers above the arrows denote different property paths over the KG.

Case	Conv. ID	Question	Topic	KG Paths		Answer Generation
1	11200-2	Where was he born?	ground_truth: soccer **prediction**: *soccer* *isRight*: ✔	gold path 1: *"Lionel Messi"* $\longrightarrow$ *"Rosario"*	*P*@1: ✔ *H*@5: ✔ *H*@10: ✔	reference: he was born with [ans] **prediction**: *he was born in [ans].* $Bleu_2$: 0.51
2	8963-4	What is the name of the second movie?	ground_truth: movies **prediction**: *movies* *isRight*: ✔	gold path 1: *"Grease"* $\xrightarrow{1}$ *"Grease2"* gold path 2: *"Grease"* $\xrightarrow{2}$ *"Grease2"*	*P*@1: ✔ *H*@5: ✔ *H*@10: ✔ *P*@1: ✗ *H*@5: ✔ *H*@10: ✔	reference: [ans] is the name of the second movie. **prediction**: *the name of the second movie is [ans].* $Bleu_2$: 0.79
3	8962-0	What network was Dexter on?	ground_truth: tv_series **prediction**: *tv_series* *isRight*: ✔	gold path 1: *"Dexter"* $\xrightarrow{1}$ *"showtime"* gold path 2: *"Dexter"* $\xrightarrow{2}$ *"showtime"* gold path 3: *"Dexter"* $\xrightarrow{3}$ *"showtime"*	*P*@1: ✗ *H*@5: ✔ *H*@10: ✔ *P*@1: ✔ *H*@5: ✔ *H*@10: ✔ *P*@1: ✗ *H*@5: ✔ *H*@10: ✔	reference: magazine first published the book is [ans] **prediction**: *[ans] first published the book.* $Bleu_2$: 0.71
4	8962-4	What was the main location of it?	ground_truth: tv_series **prediction**: *tv_series* *isRight*: ✔	gold path 1: *"Dexter"* $\xrightarrow{1}$ *"Miami"* gold path 2: *"Dexter"* $\xrightarrow{2}$ *"Miami"* gold path 3: *"Dexter"* $\xrightarrow{3}$ *"Miami"*	*P*@1: ✔ *H*@5: ✔ *H*@10: ✔ *P*@1: ✗ *H*@5: ✔ *H*@10: ✔ *P*@1: ✗ *H*@5: ✔ *H*@10: ✔	reference: the main location of it was [ans]. **prediction**: *[ans] was the actress.* $Bleu_2$: 0.01
5	9438-0	Who wrote Harry Potter?	ground_truth: books **prediction**: *books* *isRight*: ✔	gold path 1: *"Harry Potter"* $\xrightarrow{1}$ *"J.K.Rowling"* gold path 2: *"Harry Potter"* $\xrightarrow{2}$ *"J.K.Rowling"* gold path 3: *"Harry Potter Character"* $\longrightarrow$ *"J.K.Rowling"*	*P*@1: ✗ *H*@5: ✗ *H*@10: ✗ *P*@1: ✗ *H*@5: ✔ *H*@10: ✔ *P*@1: ✔ *H*@5: ✔ *H*@10: ✔	reference: [ans] wrote harry potter. **prediction**: *[ans] wrote harry potter .* $Bleu_2$: 1.0
6	10033-4	Who is the protagonist?	ground_truth: books **prediction**: *books* *isRight*: ✔	gold path 1: *"The Catcher in the Rye"* $\xrightarrow{1}$ *"Holden Caulfied"* gold path 2: *"The Catcher in the Rye"* $\xrightarrow{2}$ *"Holden Caulfied"* gold path 3: *"The Catcher in the Rye"* $\xrightarrow{3}$ *"Holden Caulfied"* gold path 4: *"The Catcher in the Rye"* $\xrightarrow{4}$ *"Holden Caulfied"* gold path 5: *"J.D. Salinger"* $\longrightarrow$ *"Holden Caulfied"*	*P*@1: ✗ *H*@5: ✗ *H*@10: ✗ *P*@1: ✗ *H*@5: ✔ *H*@10: ✔ *P*@1: ✔ *H*@5: ✔ *H*@10: ✔ *P*@1: ✗ *H*@5: ✔ *H*@10: ✔ *P*@1: ✗ *H*@5: ✔ *H*@10: ✔	reference: [ans] is the protagonist. **prediction**: *[ans] is the protagonist.* $Bleu_2$: 1.0

MRR, where the statistics are presented as the average of experiments conducted more than ten times. The empirical results demonstrate that fluent answer generation with prefixed "prompt" inputs can learn more comprehensive commonsense from PLM and provide additional context information to support DANTE to generate a more accurate representation for the contrastive ranking task. Hence, we conclude that this module positively impacts the KG path ranking and topic classification tasks.

Effect of Multi-task Joint Learning DANTE is trained in a multi-task manner, therefore, the effectiveness of the three sub-tasks and the multi-object is required to be explored. To explore the effectiveness of multi-task joint learning, we independently trained each module without any parameter sharing between the sub-modules, the corresponding outcomes are detailed in Table 5. DANTE (Train Separately) presents lower scores for all the metrics. Moreover, we observed that all the modules could fall into overfitting compared with joint training, thus illustrating that DANTE generates superior embedding representations that comprehensively capture all the tasks and can learn deeper cognitive information for downstream tasks, thereby avoiding the occurrence of overfitting.

4.7. Case study

We demonstrate six examples from different topics and conversation turns of the results on ConvQuestions, as listed in Table 6. For Cases 1 and 2, which are simple cases in the middle of the dialog turn, DANTE's predictions are perfect for supporting a dialog-evolving process. For Cases 3 and 4, which originated from conversation 8962, the focal entity at turn 1 is "Dexter" after four dialog interaction turns. It changes back to "Dexter" at turn 5. DANTE also selects and ranks the correct KG paths in order of high quality. Furthermore, in some complex cases, as in Case 5, the KG paths begin from two entities while pointing to the same target. DANTE ranks two of the three KG paths in the correct order. For Case 6, with five paths starting from different entities with different property paths, DANTE still ranks the most KG paths in a highly relevant order. It should be noted that the predicted topics are all correct, and fluent Answers perform with high Bleu scores, demonstrating that DANTE aggregates topic information from the right direction and learns a semantic-rich graph representation of KGs-based ConvQA.

5. Conclusions and future work

We propose a novel dialog-graph-enhanced prompt learning method, DANTE, which models a user's intention and conversation historical interaction records simultaneously and effectively. The dialog graph constructed represents the QA pairs, KG paths, and dialog context dependency in an explicit mode for selecting high-quality answers with a "pre-train, prompt, predict" training manner. The experimental results demonstrate that DANTE significantly improves the KGs-based ConvQA task performance on ConvQuestions and ConvRef compared with existing approaches. In our future work, we intend to explore whether DANTE can be extended to multi-modal and cross-topic KGs-based Conversation QA tasks by applying it to ConvMix [58]. Furthermore, we are also interested in exploiting RL from human feedback [59] to prompt more LLMs potentials as DANTE's continuous work.

CRediT authorship contribution statement

Jingyang Li: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Shengli Song:** Supervision, Resources, Project administration. **Sitong Yan:** Writing – review & editing, Writing – original draft, Visualization. **Guangneng Hu:** Formal analysis, Conceptualization. **Chengen Lai:** Visualization. **Yulong Zhou:** Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

References

[1] Philipp Christmann, Rishiraj Saha Roy, Abdalghani Abujabal, Jyotsna Singh, Gerhard Weikum, Look before you hop: Conversational question answering over knowledge graphs using judicious context expansion, in: Proceedings of the 28th ACM International Conference on Information and Knowledge Management, 2019, pp. 729–738.

[2] Endri Kacupaj, Barshana Banerjee, Kuldeep Singh, Jens Lehmann, Paraqa: A question answering dataset with paraphrase responses for single-turn conversation, in: The Semantic Web: 18th International Conference, ESWC 2021, Virtual Event, June 6–10, 2021, Proceedings 18, Springer, 2021, pp. 598–613.

[3] Joan Plepi, Endri Kacupaj, Kuldeep Singh, Harsh Thakkar, Jens Lehmann, Context transformer with stacked pointer networks for conversational question answering over knowledge graphs, in: The Semantic Web: 18th International Conference, ESWC 2021, Virtual Event, June 6–10, 2021, Proceedings, Springer, 2021, pp. 356–371.

[4] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al., Training language models to follow instructions with human feedback, Adv. Neural Inf. Process. Syst. 35 (2022) 27730–27744.

[5] Aaron Daniel Cohen, Adam Roberts, Alejandra Molina, Alena Butryna, Alicia Jin, Apoorv Kulshreshtha, Ben Hutchinson, Ben Zevenbergen, Blaise Hilary Aguera-Arcas, Chung ching Chang, Claire Cui, Cosmo Du, Daniel De Freitas Adiwardana, Dehao Chen, Dmitry (Dima) Lepikhin, Ed H. Chi, Erin Hoffman-John, Heng-Tze Cheng, Hongrae Lee, Igor Krivokon, James Qin, Jamie Hall, Joe Fenton, Johnny Soraker, Kathy Meier-Hellstern, Kristen Olson, Lora Mois Aroyo, Maarten Paul Bosma, Marc Joseph Pickett, Marcelo Amorim Menegali, Marian Croak, Mark Díaz, Matthew Lamm, Maxim Krikun, Meredith Ringel Morris, Noam Shazeer, Quoc V. Le, Rachel Bernstein, Ravi Rajakumar, Ray Kurzweil, Romal Thoppilan, Steven Zheng, Taylor Bos, Toju Duke, Tulsee Doshi, Vincent Y. Zhao, Vinodkumar Prabhakaran, Will Rusch, YaGuang Li, Yanping Huang, Yanqi Zhou, Yuanzhong Xu, Zhifeng Chen, LaMDA: Language models for dialog applications, 2022, arXiv.

[6] Munazza Zaib, Quan Z. Sheng, Wei Emma Zhang, Adnan Mahmood, Keeping the questions conversational: Using structured representations to resolve dependency in conversational question answering, 2023.

[7] Hongshen Chen, Xiaorui Liu, Dawei Yin, Jiliang Tang, A survey on dialogue systems: Recent advances and new frontiers, Acm Sigkdd Explor. Newslett. 19 (2) (2017) 25–35.

[8] Jinjie Ni, Tom Young, Vlad Pandelea, Fuzhao Xue, Erik Cambria, Recent advances in deep learning based dialogue systems: A systematic survey, Artif. Intell. Rev. 56 (4) (2023) 3055–3155.

[9] Yunshi Lan, Jing Jiang, Modeling transitions of focal entities for conversational knowledge base question answering, 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), 2021, pp. 3288–3297.

[10] Yunqi Qiu, Kun Zhang, Yuanzhuo Wang, Xiaolong Jin, Long Bai, Saiping Guan, Xueqi Cheng, Hierarchical query graph generation for complex question answering over knowledge graph, in: Proceedings of the 29th ACM International Conference on Information and Knowledge Management, 2020, pp. 1285–1294.

[11] Pierre Marion, Pawel Nowak, Francesco Piccinno, Structured context and high-coverage grammar for conversational question answering over knowledge graphs, in: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 2021, pp. 8813–8829, http://dx.doi.org/10.18653/v1/2021.emnlp-main.695.

[12] Magdalena Kaiser, Rishiraj Saha Roy, Gerhard Weikum, Reinforcement learning from reformulations in conversational question answering over knowledge graphs, in: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2021, pp. 459–469.

[13] Gangwoo Kim, Hyunjae Kim, Jungsoo Park, Jaewoo Kang, Learn to resolve conversational dependency: A consistency training framework for conversational question answering, 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), Association for Computational Linguistics, Online, 2021, pp. 6130–6141, http://dx.doi.org/10.18653/v1/2021.acl-long.478.

[14] Endri Kacupaj, Kuldeep Singh, Maria Maleshkova, Jens Lehmann, Contrastive representation learning for conversational question answering over knowledge graphs, in: Proceedings of the 31st ACM International Conference on Information & Knowledge Management, 2022, pp. 925–934.

[15] Chen Qu, Liu Yang, Minghui Qiu, Yongfeng Zhang, Cen Chen, W. Bruce Croft, Mohit Iyyer, Attentive history selection for conversational question answering, in: Proceedings of the 28th ACM International Conference on Information and Knowledge Management, 2019, pp. 1391–1400.

[16] Munazza Zaib, Wei Emma Zhang, Quan Z. Sheng, Adnan Mahmood, Yang Zhang, Conversational question answering: A survey, Knowl. Inf. Syst. 64 (12) (2022) 3151–3195, http://dx.doi.org/10.1007/s10115-022-01744-y.

[17] Bill Lin, Wangchunshu Zhou, Ming Shen, Pei Zhou, Chandra Bhagavatula, Choi Yejin, Xiang Ren, CommonGen: A constrained text generation challenge for generative commonsense reasoning, 2020, pp. 1823–1840, http://dx.doi.org/10.18653/v1/2020.findings-emnlp.165.

[18] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, Graham Neubig, Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing, ACM Comput. Surv. 55 (9) (2023) 1–35.

[19] Munazza Zaib, Wei Emma Zhang, Quan Z Sheng, Adnan Mahmood, Yang Zhang, Conversational question answering: A survey, Knowl. Inf. Syst. 64 (12) (2022) 3151–3195.

[20] Nikita Bhutani, Xinyi Zheng, Kun Qian, Yunyao Li, H. Jagadish, Answering complex questions by combining information from curated and extracted knowledge bases, in: Proceedings of the First Workshop on Natural Language Interfaces, 2020, pp. 1–10.

[21] Wanyun Cui, Yanghua Xiao, Haixun Wang, Yangqiu Song, Seung-won Hwang, Wei Wang, KBQA: Learning question answering over QA corpora and knowledge bases, Proc. VLDB Endow. 10 (5) (2017) 565–576, http://dx.doi.org/10.14778/3055540.3055549.

[22] Mandar Joshi, Eunsol Choi, Daniel Weld, Luke Zettlemoyer, TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension, in: Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Vancouver, Canada, 2017, pp. 1601–1611, http://dx.doi.org/10.18653/v1/P17-1147.

[23] Pranav Rajpurkar, Robin Jia, Percy Liang, Know what you don't know: Unanswerable questions for SQuAD, in: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), Association for Computational Linguistics, Melbourne, Australia, 2018, pp. 784–789, http://dx.doi.org/10.18653/v1/P18-2124.

[24] Amrita Saha, Vardaan Pahuja, Mitesh Khapra, Karthik Sankaranarayanan, Sarath Chandar, Complex sequential question answering: Towards learning to converse over linked question answer pairs with a knowledge graph, in: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 32, no. 1, 2018.

[25] Junzhuo Li, Deyi Xiong, Kafsp: Knowledge-aware fuzzy semantic parsing for conversational question answering over a large-scale knowledge base, in: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 2022, pp. 461–473.

[26] Laura Perez-Beltrachini, Parag Jain, Emilio Monti, Mirella Lapata, Semantic parsing for conversational question answering over knowledge graphs, in: Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics, Association for Computational Linguistics, Dubrovnik, Croatia, 2023, pp. 2507–2522.

[27] Endri Kacupaj, Joan Plepi, Kuldeep Singh, Harsh Thakkar, Jens Lehmann, Maria Maleshkova, Conversational question answering over knowledge graphs with transformer and graph attention networks, in: Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, Association for Computational Linguistics, Online, 2021, pp. 850–862, http://dx.doi.org/10.18653/v1/2021.eacl-main.72.

[28] Endri Kacupaj, Shyamnath Premnadh, Kuldeep Singh, Jens Lehmann, Maria Maleshkova, Vogue: Answer verbalization through multi-task learning, in: Machine Learning and Knowledge Discovery in Databases. Research Track: European Conference, ECML PKDD 2021, Bilbao, Spain, September 13–17, 2021, Proceedings, Part III 21, Springer, 2021, pp. 563–579.

[29] Xin Bi, Haojie Nie, Xiyu Zhang, Xiangguo Zhao, Ye Yuan, Guoren Wang, Unrestricted multi-hop reasoning network for interpretable question answering over knowledge graph, Knowl.-Based Syst. 243 (2022) 108515, http://dx.doi.org/10.1016/j.knosys.2022.108515.

[30] Xin Bi, Haojie Nie, Guoliang Zhang, Lei Hu, Yuliang Ma, Xiangguo Zhao, Ye Yuan, Guoren Wang, Boosting question answering over knowledge graph with reward integration and policy evaluation under weak supervision, Inf. Process. Manage. 60 (2) (2023) 103242, http://dx.doi.org/10.1016/j.ipm.2022.103242.

[31] Zilong Zheng, Wenguan Wang, Siyuan Qi, Song-Chun Zhu, Reasoning visual dialogs with structural and partial observations, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 6669–6678.

[32] Idan Schwartz, Seunghak Yu, Tamir Hazan, Alexander G. Schwing, Factor graph attention, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 2039–2048.

[33] Dan Guo, Hui Wang, Hanwang Zhang, Zheng-Jun Zha, Meng Wang, Iterative context-aware graph inference for visual dialog, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020, pp. 10055–10064.

[34] Meng Zhao, Lifang Wang, Zejun Jiang, Ronghan Li, Xinyu Lu, Zhongtian Hu, Multi-task learning with graph attention networks for multi-domain task-oriented dialogue systems, Knowl.-Based Syst. 259 (2023) 110069, http://dx.doi.org/10.1016/j.knosys.2022.110069.

[35] Jie Wu, Ian G. Harris, Hongzhi Zhao, GraphMemDialog: Optimizing end-to-end task-oriented dialog systems using graph memory networks, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 36, No. 10, 2022, pp. 11504–11512.

[36] Zhitong Yang, Bo Wang, Jinfeng Zhou, Yue Tan, Dongming Zhao, Kun Huang, Ruifang He, Yuexian Hou, Topkg: Target-oriented dialog via global planning on knowledge graph, in: Proceedings of the 29th International Conference on Computational Linguistics, 2022, pp. 745–755.
[37] Yu Chen, Lingfei Wu, Mohammed J. Zaki, GraphFlow: Exploiting conversation flow with graph neural networks for conversational machine comprehension, in: Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI '20, 2021.
[38] Michihiro Yasunaga, Hongyu Ren, Antoine Bosselut, Percy Liang, Jure Leskovec, QA-GNN: Reasoning with language models and knowledge graphs for question answering, in: Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Association for Computational Linguistics, 2021, pp. 535–546, http://dx.doi.org/10.18653/v1/2021.naacl-main.45, Online.
[39] Philipp Christmann, Rishiraj Saha Roy, Gerhard Weikum, Explainable conversational question answering over heterogeneous sources via iterative graph neural networks, in: SIGIR '23, Association for Computing Machinery, New York, NY, USA, 2023, pp. 643–653, http://dx.doi.org/10.1145/3539618.3591682.
[40] Edoardo Maria Ponti, Goran Glavaš, Olga Majewska, Qianchu Liu, Ivan Vulić, Anna Korhonen, XCOPA: A multilingual dataset for causal commonsense reasoning, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP, Association for Computational Linguistics, Online, 2020, pp. 2362–2376, http://dx.doi.org/10.18653/v1/2020.emnlp-main.185.
[41] Xu Han, Weilin Zhao, Ning Ding, Zhiyuan Liu, Maosong Sun, Ptr: Prompt tuning with rules for text classification, AI Open 3 (2022) 182–192.
[42] Timo Schick, Sahana Udupa, Hinrich Schütze, Self-diagnosis and self-debiasing: A proposal for reducing corpus-based bias in nlp, Trans. Assoc. Comput. Linguist. 9 (2021) 1408–1424.
[43] Timo Schick, Hinrich Schütze, Few-shot text generation with natural language instructions, in: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, 2021, pp. 390–402.
[44] Lukasz Kaiser, Ofir Nachum, Aurko Roy, Samy Bengio, Learning to remember rare events, in: 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings, OpenReview.net, 2017.
[45] Nazneen Fatema Rajani, Bryan McCann, Caiming Xiong, Richard Socher, Explain yourself! leveraging language models for commonsense reasoning, in: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, 2019, pp. 4932–4942.
[46] Chiaming Hsu, Changtong Zan, Liang Ding, Longyue Wang, Xiaoting Wang, Weifeng Liu, Fu Lin, Wenbin Hu, Prompt-learning for cross-lingual relation extraction, in: 2023 International Joint Conference on Neural Networks (IJCNN), 2023, pp. 1–9, http://dx.doi.org/10.1109/IJCNN54540.2023.10192002.
[47] Wen Zhang, Yushan Zhu, Mingyang Chen, Yuxia Geng, Yufeng Huang, Yajing Xu, Wenting Song, Huajun Chen, Structure pretraining and prompt tuning for knowledge graph transfer, in: Proceedings of the ACM Web Conference 2023, 2023, pp. 2581–2590.
[48] Lifu Tu, Jin Qu, Semih Yavuz, Shafiq Joty, Wenhao Liu, Caiming Xiong, Yingbo Zhou, Efficiently aligned cross-lingual transfer learning for conversational tasks using prompt-tuning, in: Yvette Graham, Matthew Purver (Eds.), Findings of the Association for Computational Linguistics: EACL 2024, Association for Computational Linguistics, St. Julian's, Malta, 2024, pp. 1278–1294, https://aclanthology.org/2024.findings-eacl.85.
[49] Debjoy Saha, Bishal Santra, Pawan Goyal, A study on prompt-based few-shot learning methods for belief state tracking in task-oriented dialog systems, 2022, arXiv preprint arXiv:2204.08167.
[50] Derek Chen, Kun Qian, Zhou Yu, Stabilized in-context learning with pre-trained language models for few shot dialogue state tracking, in: Andreas Vlachos, Isabelle Augenstein (Eds.), Findings of the Association for Computational Linguistics: EACL 2023, Association for Computational Linguistics, Dubrovnik, Croatia, 2023, pp. 1551–1564, http://dx.doi.org/10.18653/v1/2023.findings-eacl.115, https://aclanthology.org/2023.findings-eacl.115.
[51] Makesh Narsimhan Sreedhar, Christopher Parisien, Prompt learning for domain adaptation in task-oriented dialogue, in: Proceedings of the Towards Semi-Supervised and Reinforced Task-Oriented Dialog Systems (SereTOD), Association for Computational Linguistics, Abu Dhabi, Beijing (Hybrid), 2022, pp. 24–30.
[52] Xuansheng Wu, Kaixiong Zhou, Mingchen Sun, Xin Wang, Ninghao Liu, A survey of graph prompting methods: Techniques, applications, and challenges, 2023, arXiv preprint arXiv:2303.07275.
[53] Zemin Liu, Xingtong Yu, Yuan Fang, Xinming Zhang, GraphPrompt: Unifying pre-training and downstream tasks for graph neural networks, in: Proceedings of the ACM Web Conference 2023, 2023, pp. 417–428.
[54] Jiacheng Liu, Alisa Liu, Ximing Lu, Sean Welleck, Peter West, Ronan Le Bras, Yejin Choi, Hannaneh Hajishirzi, Generated knowledge prompting for commonsense reasoning, in: Smaranda Muresan, Preslav Nakov, Aline Villavicencio (Eds.), Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022, Association for Computational Linguistics, 2022, pp. 3154–3169, http://dx.doi.org/10.18653/v1/2022.acl-long.225.
[55] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, Language models are unsupervised multitask learners, 2019.
[56] Thomas N. Kipf, Max Welling, Semi-supervised classification with graph convolutional networks, 2016, arXiv preprint arXiv:1609.02907.
[57] Xirui Ke, Jing Zhang, Xin Lv, Yiqi Xu, Shulin Cao, Cuiping Li, Hong Chen, Juanzi Li, Knowledge-augmented self-training of a question rewriter for conversational knowledge base question answering, in: Findings of the Association for Computational Linguistics: EMNLP 2022, 2022, pp. 1844–1856.
[58] Philipp Christmann, Rishiraj Saha Roy, Gerhard Weikum, Conversational question answering on heterogeneous sources, in: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2022, pp. 144–154.
[59] Shane Griffith, Kaushik Subramanian, Jonathan Scholz, Charles L. Isbell, Andrea L. Thomaz, Policy shaping: Integrating human feedback with reinforcement learning, Adv. Neural Inf. Process. Syst. 26 (2013).