

ConTEsT: Temporal Knowledge Graph Question Answering via Contractive Dual-Space Subgraph Search

Quan Xu¹ Yanting Yin^{1*} Shiyong Li¹ Wensheng Zhang^{1, 2*}

¹School of Computer Science and Network Engineering, Guangzhou University, Guangzhou, Guangdong 510000, China

²Institute of Automation, Chinese Academy of Sciences (CAS), Beijing 100190, China

xuquan@e.gzhu.edu.cn yanting.yin@gzhu.edu.cn

lsying@e.gzhu.edu.cn zhangwenshengia@hotmail.com

Abstract

Temporal Knowledge Graph Question Answering (TKGQA) commonly relies on retrieval-aggregation-reasoning pipelines, which suffer from two critical challenges: *error propagation*, where early retrieval mistakes accumulate along reasoning paths, and *temporal conflicts* caused by post-hoc temporal filtering, leaving temporally inconsistent paths in the search space and undermining multi-hop reasoning reliability. To address these limitations, we propose **ConTEsT** (Contractive Temporal Exploration), a contractive subgraph search framework over structural and temporal spaces for TKGQA. It dynamically anchors question-relevant entities, performs selective subgraph exploration guided by path semantics and sub-question consistency, and applies Temporal Feasible Region Propagation (TFRP), which actively propagates temporal constraints during search and maintains a temporal feasible region. Experiments on MultiTQ show that ConTEsT is particularly effective for complex temporal reasoning, yielding a 27.1% relative improvement in Hits@1 over the strongest baseline on Multiple questions. Its gains are smaller on Single and Time-answer questions that depend more on direct evidence matching or exact temporal endpoints. Our code is available at <https://github.com/xq641/ConTEsT>.

1 Introduction

Temporal Knowledge Graph Question Answering (TKGQA) aims to answer natural language questions by reasoning over time-aware facts in Temporal Knowledge Graphs (TKGs). Compared with static knowledge graph question answering, it requires identifying relevant entities and relations while satisfying temporal operators such as *before*, *after*, *first*, and *last*, making it particularly challenging under implicit event references, multi-hop dependencies, and multiple temporal constraints.

*Corresponding authors.




Figure 1: Comparison between chain-style decomposition and ConTEsT. Chain-style reasoning commits to intermediate answers sequentially and may propagate early mistakes, while ConTEsT maintains a compact reasoning subgraph and refines it under structural and temporal constraints.

Recent studies have advanced TKGQA through multi-granularity temporal modeling, LLM-enhanced retrieval, abstract temporal reasoning, and recursive question decomposition (Chen et al., 2023; Qian et al., 2024; Chen et al., 2024b; Gao et al., 2024; Gong et al., 2025). Despite these advances, existing methods often decouple structural evidence construction from temporal constraint propagation, allowing early retrieval errors and temporally inconsistent paths to persist throughout multi-hop reasoning.

Addressing *error propagation* remains a highly challenging issue for current reasoning pipelines. Existing methods construct question-specific subgraphs (Sun et al., 2018), iteratively retrieve evidence (Sun et al., 2019), or recursively decompose complex questions (Gong et al., 2025). However, because traditional subgraph expansion largely relies on local structures, early retrieval mistakes can accumulate and amplify along reasoning paths (He et al., 2021), as shown in Figure 1(a). Without continuous semantic screening and active structural compression, irrelevant neighbors can easily saturate the candidate space, severely undermining the reliability of multi-hop reasoning.

Furthermore, another critical deficiency stems




Figure 2: Without TFRP vs. with TFRP. Without TFRP, temporal validation is delayed until final ranking, allowing time-conflicting paths to remain during reasoning. With TFRP, feasible time regions are updated during search, so temporally invalid evidence can be pruned earlier.

from *temporal conflicts* caused by post-hoc temporal filtering. Existing TKGQA methods typically treat time in a passive manner, either as embedding-level temporal features (Saxena et al., 2021) or as an isolated final-stage filtering signal (Qian et al., 2024). Although these strategies improve temporal sensitivity (Mavromatis et al., 2022), they fail to resolve temporal conflicts early in the reasoning process, leaving temporally inconsistent paths scattered across the search space (Chen et al., 2022). This passivity leads to a rapid growth of invalid structural combinations that appear semantically plausible but are temporally infeasible, as demonstrated in Figure 2(a).

To circumvent these limitations, we propose ConTE_{XT}, which formulates TKGQA as a unified dual-space search over structural and temporal spaces. The framework jointly optimizes the candidate space through two interactive components. First, it dynamically anchors question-relevant entities and performs selective subgraph exploration. Distinct from naive neighborhood expansion, ConTE_{XT} combines path semantics with sub-question consistency to introduce query-aligned evidence while continuously compressing the structural reasoning space. Second, TFRP mitigates path conflicts by propagating temporal constraints during exploration and maintaining a dynamic temporal feasible region. By using temporal constraints as

online search signals rather than passive post-filters, TFRP removes temporally incompatible evidence earlier in the reasoning process. Fixed search budgets and a maximum of three iterations keep the procedure bounded, while the practical stopping conditions are detailed in Section 3.3.

Our contributions are summarized as follows:

- We propose ConTE_{XT}, a unified dual-space search framework for TKGQA that jointly optimizes structural paths and temporal spaces.
- We introduce a semantic-driven structural contraction mechanism that tightly couples expansion with active searching, complemented by TFRP which suppresses temporally incompatible paths.
- Extensive experiments on MultiTQ demonstrate that ConTE_{XT} achieves an approximate 27.1% improvement on challenging multi-hop questions.

2 Preliminaries

A temporal knowledge graph (TKG) is denoted as $\mathcal{G} = (\mathcal{E}, \mathcal{R}, \mathcal{T}, \mathcal{F})$, where $\mathcal{E}$, $\mathcal{R}$, $\mathcal{T}$, and $\mathcal{F}$ denote the sets of entities, relations, timestamps, and temporal facts, respectively. Each temporal fact is represented as a quadruple $f = (e_s, r, e_o, t) \in \mathcal{F}$, where e_s and e_o are the subject and object entities, r is the relation, and t specifies the timestamp at which the fact holds.

TKGQA aims to answer a natural language question q based on the temporal facts in $\mathcal{G}$. The answer can be either an entity or a temporal expression, depending on the expected answer type expressed in the question. Notably, TKGQA requires a model to satisfy not only entity–relation matching but also temporal constraints such as *before*, *after*, *first*, and *last*.

We use C_q to denote the structural and temporal constraints derived from q , and $A(\mathcal{G}', q)$ to denote the candidate answers supported by a subgraph $\mathcal{G}' \subseteq \mathcal{G}$.

3 Method

3.1 Overview

The operational workflow of ConTE_{XT} follows an iterative search-and-contraction paradigm over joint structural and temporal spaces to approximate a compact, self-consistent reasoning subgraph.

As illustrated in Figure 3, the pipeline executes through three tightly coupled phases:

Question Parsing and Initialization: Given a query q and a graph $\mathcal{G}$, the model extracts structural and temporal constraints to initialize a localized reasoning subgraph (Figure 3, Left).

Dual-Space Joint Optimization: Starting from the initial state, the framework alternates three operations (Figure 3, Center): selective frontier expansion retrieves query-aligned evidence, TFRP updates fact-level temporal feasibility, and structural contraction removes facts that violate relation-object, path-connectivity, temporal-feasibility, or answer-role constraints.

Constrained Answer Selection: ConTEXT evaluates and ranks all remaining graph-supported candidates under the accumulated joint constraints to output the final answer (Figure 3, Right).

The complete procedure is formally summarized in Algorithm 1.

3.2 Question Parsing and Initialization

Given an input question q , ConTEXT uses a deterministic rule-based parser to extract a structured query:

$$Q_q = (E_a, E_r, \mathcal{R}, o, y, g), \quad (1)$$

where E_a and E_r denote anchor entities and reference events; $\mathcal{R}$ denotes candidate relations; and o , y , and g represent the temporal operator, answer type, and temporal granularity, respectively. Lexical patterns, benchmark-provided question and answer-type fields, entity linking, and graph-relation matching populate these fields; this stage does not invoke an LLM. Q_q induces joint structural constraints $\mathcal{C}_q^{str}$ and temporal constraints $\mathcal{C}_q^{tmp}$. If no explicit timestamp is provided, the reference time t_{ref} is estimated by resolving E_r against relation- and anchor-compatible facts in $\mathcal{G}$.

Instead of performing unrestricted k -hop neighbor expansion, ConTEXT utilizes Q_q to construct a focused initial subgraph $\mathcal{G}_0$. Candidate facts are ordered by a fixed priority π_{init} : facts connecting two anchors are considered first, followed by anchor-adjacent facts matching the parsed relation and object, relation-matching facts with compatible time hints, other relation-matching anchor facts, and finally a bounded one-hop fill-in. The initial subgraph $\mathcal{G}_0 = (\mathcal{V}_0, \mathcal{E}_0)$ is materialized under an initialization budget B :

$$\mathcal{E}_0 = \text{Top}_B(\mathcal{F}; \pi_{init}), \quad (2)$$

where $\mathcal{V}_0$ comprises all entities involved in $\mathcal{E}_0$. This localized subgraph defines the initial reasoning state $\mathcal{S}_0 = (\mathcal{G}_0, \Theta_0, \mathcal{C}_q)$, where $\Theta_0(f) = \{t_f\}$ initializes the active temporal region for each fact, serving as the starting frontier for the subsequent dual-space joint optimization, as shown in Figure 3(left).

Algorithm 1: ConTEXT

Input: Question q , temporal knowledge graph $\mathcal{G}$
Output: Answer a^*
// Rule-based query parsing and state initialization
1 $\mathcal{O}_q \leftarrow \text{RuleParse}(q, \mathcal{G})$
2 $(Q_q, \mathcal{C}_q, t_q^{ref}) \leftarrow \text{Normalize}(\mathcal{O}_q, \mathcal{G})$
3 $\mathcal{S}_0 = (\mathcal{G}_0, \Theta_0, \mathcal{C}_q, \mathcal{B}_0) \leftarrow \text{InitState}(Q_q, t_q^{ref}, \mathcal{G})$
// Iterative subgraph evolution
4 $i \leftarrow 0$
5 **while** $\neg \text{Converged}(\mathcal{S}_i)$ **do**
6 $q_i \leftarrow \text{GenSubQ}(q, \mathcal{S}_i)$
7 $\mathcal{C}_i \leftarrow \mathcal{N}(\mathcal{B}_i)$
8 $\mathcal{G}_i^+ \leftarrow \mathcal{G}_i \cup \text{Top}_K(\mathcal{C}_i; s_{exp}(f; q_i, \mathcal{S}_i))$
9 **foreach** $f \in \mathcal{G}_i^+$ **do**
10 $\Theta_{i+1}(f) \leftarrow \Theta_i(f) \cap \{t \in \mathcal{T} \mid \phi(t, \mathcal{C}_q^{tmp}) = 1\}$
11 $\tilde{\mathcal{G}}_{i+1} \leftarrow (\Pi_{time} \circ \Pi_{path} \circ \Pi_{rel})(\mathcal{G}_i^+, \Theta_{i+1}, \mathcal{C}_q)$
12 $\mathcal{G}_{i+1} \leftarrow \Pi_{role}(\tilde{\mathcal{G}}_{i+1}, Q_q)$
13 $\mathcal{B}_{i+1} \leftarrow \text{Frontier}(\mathcal{G}_{i+1}, \mathcal{G}_i)$
14 $\mathcal{S}_{i+1} \leftarrow (\mathcal{G}_{i+1}, \Theta_{i+1}, \mathcal{C}_q, \mathcal{B}_{i+1}); \quad i \leftarrow i + 1$
// Constrained answer selection
15 $\mathcal{G}_* \leftarrow \text{FinalSubgraph}(\mathcal{S}_i)$
16 $(A_q, \tilde{s}, \mathcal{F}_{top}) \leftarrow \text{ScoreCandidates}(\mathcal{G}_*, Q_q)$
17 $a' \leftarrow \text{LLMSelect}(q, A_q, \mathcal{F}_{top})$
18 $a^* \leftarrow \text{Guard}(a', A_q, \tilde{s}, Q_q)$
19 **return** a^*

3.3 Iterative Subgraph Evolution

The initial subgraph may contain incomplete evidence or locally plausible but globally inconsistent facts. ConTEXT therefore evolves the reasoning space through a coupled expansion-propagation-contraction process. This design allows the model to introduce new evidence when necessary, while continuously suppressing facts that violate semantic, structural, or temporal constraints.

Sub-question guided expansion. At iteration i , ConTEXT derives a local sub-question q_i from the original question q and the current state $\mathcal{S}_i$. The generator receives q , the parsed constraints, the current frontier and subgraph size, and summaries of the two most recent expansion and contraction steps. It returns a compact structured record containing a local search goal and relation, time, and answer-type hints. An LLM may produce this record; if it is unavailable or its output is invalid, a deterministic rule-based template provides the




Figure 3: Overview of the ConTEXT framework. ConTEXT parses the input question into structured constraints, constructs an initial reasoning subgraph, iteratively evolves the subgraph through expansion, temporal propagation, and contraction, and finally resolves graph-supported candidates under structural and temporal constraints.

same fields.

The generated record contains only search hints and cannot overwrite Q_q or introduce entities into the parsed anchor set; the deterministic relations, temporal constraints, and answer type retain separate priority in expansion. Thus, q_i is a soft search guide rather than an intermediate answer commitment. For each candidate fact adjacent to the current frontier $\mathcal{B}_i$, the expansion score is computed as

$$s_{exp}(f) = \lambda_s \text{sim}(f, q_i) + \lambda_c \text{coh}(f, S_i) + \beta_r I_r(f) + \beta_o I_o(f), \quad (3)$$

where sim combines token overlap with relation and time-hint agreement, while coh measures contact with the current frontier, existing entities, the parsed object, and the reference entity. The indicators I_r and I_o mark parsed-relation and object-compatible facts, respectively. The expanded subgraph is

$$G_i^+ = G_i \cup \text{Top}_K(\mathcal{C}_i; s_{exp}), \quad \mathcal{C}_i = \mathcal{N}(\mathcal{B}_i). \quad (4)$$

This step keeps expansion selective: new facts are introduced only when they are consistent with the current search intent and the evolving reasoning state.

Structure time joint contraction. After expansion and temporal propagation, ConTEXT contracts the subgraph to preserve a compact and self-consistent reasoning space:

$$G_{i+1} = (\Pi_{role} \circ \Pi_{time} \circ \Pi_{path} \circ \Pi_{rel})(G_i^+), \quad (5)$$

where Π_{rel} removes relation-object mismatches, Π_{path} removes facts disconnected from valid reasoning paths, Π_{time} removes facts with empty feasible time regions, and Π_{role} enforces answer-side role consistency. The frontier is updated as

$$\mathcal{B}_{i+1} = \text{Frontier}(G_{i+1}, G_i). \quad (6)$$

The iteration terminates when any of the following holds: the maximum of three steps is reached; the subgraph is empty; expansion selects no new fact; the contracted subgraph contains at most 120 facts after one step when a complete relation–entity scaffold is parsed (or after two steps otherwise); or its fact count is unchanged across two consecutive contraction steps. The explicit three-step cap prevents indefinite add–remove oscillation even when the fact identities change while their count remains stable. Through repeated expansion and contraction, ConTEXT approaches a compact reasoning subgraph that is structurally connected and temporally valid under the parsed constraints.

Bounded search and scalability. The search is explicitly bounded by the initialization budget, the per-step expansion budget, and the three-step iteration cap. Contraction removes relation-, path-, time-, and role-incompatible facts after each expansion, preventing newly retrieved evidence from accumulating without control. The graph-side search cost therefore grows with the number of inspected frontier neighbors within these fixed budgets. Nevertheless, frontier enumeration remains dependent on local entity degree, and scaling to substantially

larger or denser TKGs may require additional indexing or neighbor sampling.

TFRP. For each fact f , ConTE_xT maintains a feasible time region $\Theta_i(f)$, which records the timestamps still compatible with the temporal constraints. Given $c \in C_q^{tmp}$, TFRP updates this region as

$$\Theta_{i+1}(f) = \Theta_i(f) \cap \{t \in \mathcal{T} \mid \phi(t, c) = 1\}, \quad (7)$$

where $\phi(t, c)$ checks whether timestamp t satisfies the temporal constraint c . This update makes temporal validity part of the evolving reasoning state. When $\Theta_{i+1}(f) = \emptyset$, the fact becomes temporally infeasible and is removed during contraction.

3.4 Answer Selection

After convergence, ConTE_xT selects the final answer from the stabilized subgraph G_* . For each candidate a , its supporting facts are

$$\mathcal{F}_*(a) = \{f \in G_* \mid \rho(f, Q_q) = a\}, \quad (8)$$

where $\rho(\cdot)$ extracts an answer according to the parsed answer type and role constraint. The candidate set is $A_q = \{a \mid \mathcal{F}_*(a) \neq \emptyset\}$. Each candidate is scored by

$$\tilde{s}(a) = |\mathcal{F}_*^{op}(a)| + \eta b(a, o, t_q^{ref}) + \Delta(a; \mathbf{x}_a), \quad (9)$$

where the three terms measure operator-compatible evidence, temporal boundary preference, and lightweight reranking features. If the top candidate is confident, ConTE_xT returns $\arg \max_{a \in A_q} \tilde{s}(a)$; otherwise, an LLM is used only to select from A_q , and invalid outputs are replaced by the highest-ranked graph-supported candidate.

4 Experiments

4.1 Experimental Setup

Dataset. We evaluate ConTE_xT on MultiTQ (Chen et al., 2023), a large-scale TKGQA benchmark built from ICEWS05-15. MultiTQ contains questions with multiple temporal granularities and covers both single-constraint and multi-constraint reasoning. The test split includes entity answers and time answers, which allows us to evaluate whether a model can jointly identify relevant facts, compare temporal boundaries, and return the required answer type. Table 1 reports the category distribution used in our experiments.

Table 1: Statistics of question categories in the MultiTQ benchmark.

	Category	Train	Dev	Test
Single	Equal	135,890	18,983	17,311
	Before/After	75,340	11,655	11,073
	First/Last	72,252	11,097	10,480
Multiple	Equal Multi	16,893	3,213	3,207
	After First	43,305	6,499	6,266
	Before Last	43,107	6,532	6,247
Total		386,787	57,979	54,584

Baselines. Following previous TKGQA evaluations (Qian et al., 2024; Gong et al., 2025), we compare ConTE_xT with three groups of methods. The first group contains pre-trained language models and general LLM baselines, including BERT (Devlin et al., 2019), DistilBERT (Sanh et al., 2019), ALBERT (Lan et al., 2020), LLaMA2 (Touvron et al., 2023), and ChatGPT. The second group contains embedding-based TKGQA methods, including EmbedKGQA (Saxena et al., 2020), CronKGQA (Saxena et al., 2021), and MultiQA (Chen et al., 2023). The third group contains recent temporal reasoning methods, including ARI (Chen et al., 2024b), TimeR4 (Qian et al., 2024), and RTQA (Gong et al., 2025). We report Hits@1 and Hits@10, and further break down the results by question type and answer type.

Implementation Details. We implement ConTE_xT in Python and conduct all experiments on the official MultiTQ split. The deterministic parser extracts anchors, matched relations, reference events, temporal operators, answer roles, and granularity without an LLM. For each question, initialization uses one-hop facts and keeps at most 800 facts. Each of at most three evolution steps considers up to 600 frontier-neighbor facts and adds the top 30. The final candidate pack keeps 15 candidates, while the final prompt contains at most the top 10 with compact supporting evidence.

The main experiments use deepseek-v4-flash through an OpenAI-compatible API for two bounded roles: generating the structured local sub-question and resolving ambiguous final candidates. Rule-based sub-question generation and deterministic top-ranked selection serve as fallbacks. We set the temperature to 0 and limit sub-question and final-selection outputs to 256 and 1024 tokens, respectively. The final selector must copy an answer from the evidence pack; off-list or wrong-granularity outputs are replaced by the highest-ranked valid candidate. The default hyper-

Table 2: Performance comparison on the MultiTQ test set. The best and second-best results are marked in bold and underlined, respectively.

Model	Hits@1 (%)					Hits@10 (%)				
	Overall	Multiple	Single	Entity	Time	Overall	Multiple	Single	Entity	Time
BERT	8.3	6.1	9.2	10.1	4.0	44.1	39.2	46.1	53.1	22.2
DistilBERT	8.3	7.4	8.7	10.2	3.7	48.2	42.6	50.5	59.1	21.6
ALBERT	10.8	8.6	11.6	13.9	3.2	48.4	41.5	51.2	58.9	22.8
LLaMA2	18.5	10.1	22.0	23.9	5.5	–	–	–	–	–
ChatGPT	10.2	7.7	14.7	13.7	2.0	–	–	–	–	–
EmbedKGQA	20.6	13.4	23.5	29.0	0.1	45.9	43.9	46.7	64.8	0.1
CronKGQA	27.9	13.4	33.7	32.8	15.6	60.8	45.3	67.1	69.6	39.2
MultiQA	29.3	15.9	34.7	34.9	15.7	63.5	51.9	68.2	73.3	39.6
ARI	38.0	21.0	68.0	39.4	34.4	–	–	–	–	–
TimeR4	72.8	33.5	<u>88.7</u>	63.9	94.5	–	–	–	–	–
RTQA	<u>76.5</u>	<u>42.4</u>	90.2	<u>69.2</u>	<u>94.2</u>	<u>76.8</u>	<u>42.7</u>	<u>90.7</u>	<u>69.7</u>	94.2
ConTEsT	77.1	53.9	86.2	74.5	82.6	87.2	74.8	92.2	86.6	<u>88.5</u>

parameters are operator-aware ranking strength 1.00, multiple-boundary score margin 0.250, and peer-relation promotion gap 0.20. All reported results are measured by Hits@1 and Hits@10. Table 2 uses the full test split, whereas the diagnostic ablation and LLM-selector studies in Tables 3 and 4 use the same fixed 50% test subset for efficiency; their absolute Overall scores are therefore not directly comparable.

4.2 Main Results

Overall Effectiveness. The main comparison on MultiTQ demonstrates that ConTEsT achieves the strongest overall performance among all compared methods, as reported in Table 2. Compared with RTQA, the strongest baseline in Table 2, ConTEsT improves overall Hits@1 by 0.6 percentage points (77.1% vs. 76.5%) and overall Hits@10 by 10.4 percentage points (87.2% vs. 76.8%). This indicates that dual-space contractive search can more effectively control both structural evidence and temporal validity than retrieval-augmented or decomposition-based reasoning pipelines.

Strong Gains on Multi-hop Reasoning. The most notable improvement appears on Multiple questions, where ConTEsT increases Hits@1 from 42.4% to 53.9% compared with RTQA. This gain suggests that ConTEsT is particularly effective when a question requires multiple reasoning steps or involves more than one temporal constraint. A likely reason is that ConTEsT does not commit to a fixed chain of intermediate answers. Instead, it maintains an evolving reasoning subgraph and continuously removes evidence that violates structural or temporal constraints, thereby reducing the effect of early reasoning errors.

ConTEsT is designed primarily for complex temporal reasoning rather than single-fact matching. Its advantage is therefore more pronounced when questions require multi-event search, temporal ordering, and structural–temporal consistency, whereas the gain can be smaller for Single questions that are already solvable through a direct evidence path. The method should thus not be interpreted as uniformly superior across all question categories.

Improved Candidate Quality. The Hits@10 results further show that ConTEsT improves not only the final top-ranked prediction but also the quality of the candidate set. It achieves 74.8% Hits@10 on Multiple questions and 86.6% Hits@10 on entity answers, indicating that correct answers are more frequently preserved after iterative subgraph evolution. This result supports the effectiveness of semantic-driven expansion, structural contraction, and TFRP, which jointly suppress irrelevant or temporally infeasible paths before answer selection.

Remaining Gap on Time Answers. Although ConTEsT performs strongly on overall and entity-based reasoning, its gains are not uniform across temporal question types, and TimeR4 and RTQA remain more competitive on time-answer Hits@1. TFRP is most useful when multi-event ordering and structural–temporal consistency jointly eliminate irrelevant evidence. Questions with explicit temporal links, such as before/after and equality, allow incompatible facts to be removed early. In contrast, boundary-sensitive types such as after-first and before-last depend on precise reference-event grounding and boundary selection, so a local grounding error can propagate to the final answer. Exact-time questions further require timestamp normalization and preservation of boundary-

adjacent evidence. Over-constrained pruning may therefore remove a correct timestamp even when the surrounding relation is identified, revealing a precision–recall trade-off rather than an inherent preference for entity answers. These cases motivate more robust reference-event grounding, boundary-aware retention, and timestamp calibration.

4.3 Ablation Study

Component Contributions. The ablation results in Table 3 examine whether the performance of ConTE_XT comes from its dual-space search design rather than from a single scoring heuristic. Each variant removes one core component while keeping the same evaluation protocol and comparable search budget. Here, “w/o” denotes removing the corresponding module from the full framework.

Table 3: Ablation results on the fixed 50% MultiTQ test subset (Hits@1).

Model	Overall	Multiple	Single	Entity	Time
ConTE _X T	77.84	54.69	86.94	75.26	83.89
w/o expansion	77.66	54.38	86.82	75.02	83.85
w/o contraction	76.75	53.52	86.03	74.36	83.51
w/o TFRP	51.36	19.52	63.74	39.03	80.03

Expansion Safely Boosts Topology Recall. Expansion safely boosts topology recall: removing expansion leads to a slight but consistent decline, with Overall Hits@1 decreasing from 77.84% to 77.66%, Multiple questions from 54.69% to 54.38%, and Entity answers from 75.26% to 75.02%. The small drop indicates that the initial subgraph already covers many useful facts, while expansion still recovers complementary topology around the anchored evidence. Its gains are therefore concentrated on questions that require additional connected facts beyond the initial neighborhood. This supports selective expansion as a controlled recall mechanism rather than unrestricted neighborhood growth.

Contraction Purges Inconsistent Candidates to Fine-Tune Reasoning Precision. Contraction purges inconsistent candidates to fine-tune reasoning precision: removing contraction causes a smaller but consistent decline, with Overall Hits@1 decreasing from 77.84% to 76.75%, Multiple questions from 54.69% to 53.52%, and Time answers from 83.89% to 83.51%. This suggests that expansion and TFRP recover and temporally regulate most necessary evidence, while contraction sharpens the final reasoning space by filtering relation-object mismatches, disconnected paths, temporally

infeasible facts, and role-inconsistent candidates. Its effect is most visible when several semantically plausible facts compete but only part of them remains compatible with the evolved subgraph.

TFRP Actively Suppresses Temporal Conflicts.

Removing TFRP causes the most substantial degradation. Overall Hits@1 drops from 77.84% to 51.36%, while Multiple-question performance falls sharply from 54.69% to 19.52%. Entity-answer performance also decreases from 75.26% to 39.03%, showing that temporal conflicts can mislead entity selection, not only timestamp prediction. This confirms that temporal information should actively regulate the search process. Without TFRP, relation-compatible but temporally invalid facts remain in the intermediate reasoning space, which increases the number of misleading candidates and weakens multi-hop reasoning stability.

4.4 LLM Analysis

We further evaluate ConTE_XT with different LLMs only in the role of final constrained selector. The deterministic parser, subgraph search, and the LLM used for local sub-question generation are held fixed, so this experiment does not establish backbone insensitivity for sub-question generation. It also does not compare the parametric knowledge of the selector models: all receive the same structured candidate evidence pack and must select from the provided candidates. As shown in Table 4, gpt-4o obtains the best overall Hits@1, while deepseek-v4-flash, gemini-2.0-flash, and gpt-4o-mini remain close.

Table 4: ConTE_XT results when varying only the final LLM selector on the fixed 50% MultiTQ test subset (Hits@1).

Model	Overall	Multiple	Single	Entity	Time
deepseek-v4-flash	<u>78.53</u>	<u>53.98</u>	<u>87.73</u>	<u>75.25</u>	86.30
gemini-2.0-flash	78.07	53.86	87.13	74.96	85.39
gpt-4o	78.83	54.48	87.96	75.48	86.74
gpt-4o-mini	77.67	52.70	87.04	73.89	<u>86.56</u>

The overall range across the four final selectors is only 1.16 points, indicating limited sensitivity to the LLM used at this final stage. This stability comes from the separation between deterministic evidence control and language-model selection: rule-based query parsing, subgraph evolution, and candidate ranking establish the admissible answer space, while the final LLM resolves ambiguous comparisons inside a compact evidence pack. The small gap between gpt-4o and gpt-4o-mini suggests that a smaller model can remain competitive




Figure 4: Parameter sensitivity on overall Hits@1.

as the final selector once the reasoning space is restricted. The effect of changing the sub-question-generation backbone is outside the scope of this table and should be evaluated separately.

4.5 Parameter Sensitivity

To examine parameter sensitivity, we vary one coefficient at a time while keeping the remaining settings fixed, as shown in Figure 4. This one-factor design isolates the contribution of each constraint signal in ConTEXT.

Operator-aware ranking strength (OAR). OAR adjusts the influence of temporal-operator compatibility during candidate ranking. Raising OAR from 0.25 to 1.00 improves overall Hits@1 from 70.65% to 77.11%, indicating that the final answer is sensitive to whether candidates satisfy operators such as *before*, *after*, *first*, *last*, and *equality*. The curve also shows that further increasing this factor does not bring additional gains, suggesting that temporal ordering should guide the ranking process without dominating all structural evidence.

Multiple-boundary score margin (MBM). MBM strengthens the separation among candidates in questions that involve several temporal boundaries. Its improvement is more moderate than OAR but remains stable across the tested range, which is consistent with the role of MBM: it mainly helps cases where an answer must satisfy more than one boundary condition rather than simple single-operator queries. The slight decline beyond the default setting implies that overly strict boundary margins may discard facts that are still compatible with the intended reasoning path.

Peer-relation promotion gap (PRG). PRG refines the ranking of semantically close relation candidates after temporal feasibility has been established. A larger gap helps prevent neighboring relations with similar surface forms from sharing nearly iden-

tical scores, thereby making the top candidate more stable. However, the curve is relatively flat compared with OAR, showing that PRG is a secondary calibration factor rather than the main source of improvement.

4.6 Case Study

We analyze a MultiTQ test question labeled as Multiple: “After Tony Blair, which country was the first to express an interest in cooperating with Iraq?” This case requires two coupled reasoning steps. The model must first resolve the reference event (*Tony Blair, express intent to cooperate, Iraq, 2006-05-24*), and then search for the earliest country-level subject satisfying the same cooperation-intent relation with Iraq after this boundary. As illustrated in Figure 5, a flat retriever can collect many relation-compatible facts, including (*France, express intent to cooperate, Iraq, 2005-03-27*), (*Kuwait, express intent to cooperate, Iraq, 2005-06-09*), (*Japan, express intent to cooperate, Iraq, 2005-12-06*), and (*Iran, express intent to cooperate, Iraq, 2006-03-17*). Although these facts match the target relation and object, they precede the reference event and should not support the answer.




Figure 5: Case study of ConTEXT on a multiple-constraint MultiTQ question. ConTEXT first grounds the reference event, propagates the resulting temporal boundary, removes relation-compatible but temporally invalid facts, and selects the earliest valid answer fact from the evolved subgraph.

ConTEXT handles this case by explicitly separating reference grounding, temporal propagation, and answer selection. The reference fact fixes the feasible region to timestamps later than 2006-05-24. TFRP removes all candidate facts whose timestamps fall outside this region, while contraction keeps only facts that remain connected to the parsed

relation-object pattern and answer role. In the remaining subgraph, (*Iran, express intent to cooperate, Iraq, 2006-05-26*) is the earliest valid candidate, whereas later valid facts such as (*Japan, express intent to cooperate, Iraq, 2006-10-23*) and (*South Korea, express intent to cooperate, Iraq, 2007-02-08*) satisfy the boundary but lose under the *first* operator. The final answer is therefore Iran. This example explains why ConTExT is especially effective on Multiple questions: the method does not merely retrieve facts that mention the right relation, but constructs a compact reasoning subgraph in which the reference event, temporal boundary, answer role, and ordering operator must be jointly satisfied.

5 Conclusion

In this paper, we introduce ConTExT, a framework for temporal knowledge graph question answering that reformulates temporal reasoning as a constrained search process over a compact reasoning subgraph. Rather than retrieving independent facts and relying on free-form generation, ConTExT progressively constructs a graph-supported reasoning space through query parsing, focused subgraph construction, sub-question guided expansion, structure time joint contraction, and TFRP. The final answer is resolved from graph-supported candidates under explicit structural and temporal constraints. Experimental results on the evaluated TKGQA benchmark show that ConTExT effectively improves complex temporal reasoning, particularly for questions involving temporal ordering, multi-hop evidence, and multiple temporal constraints.

Limitations

Although ConTExT shows promising performance, three limitations remain. First, its gains are not uniform across question types: contractive search may be unnecessary for Single questions with direct evidence paths, while exact-time and boundary-sensitive questions depend on accurate reference-event grounding, boundary selection, and timestamp normalization; overly aggressive pruning may therefore discard correct temporal evidence. Second, although fixed budgets and contraction bound the retained reasoning space, frontier expansion still depends on local entity degree, and substantially larger or denser TKGs may require additional indexing or neighbor sampling. Third, the current evaluation is centered on MultiTQ and normalized

discrete timestamps with explicit comparison operators. Broader evaluation and more expressive modeling of vague, interval-valued, and duration-based temporal information remain important directions for future work.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grant No. 6250073163, the National Science and Technology Major Program of China under Grant No. 2025ZD0123400.

References

- Zhuo Chen, Zhao Zhang, Zixuan Li, Fei Wang, Yutao Zeng, Xiaolong Jin, and Yongjun Xu. 2024a. [Self-improvement programming for temporal knowledge graph question answering](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 14579–14594.
- Ziyang Chen, Dongfang Li, Xiang Zhao, Baotian Hu, and Min Zhang. 2024b. [Temporal knowledge question answering via abstract reasoning induction](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4872–4889.
- Ziyang Chen, Jinzhi Liao, and Xiang Zhao. 2023. [Multi-granularity temporal question answering over knowledge graphs](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11378–11392.
- Ziyang Chen, Xiang Zhao, Jinzhi Liao, Xinyi Li, and Evangelos Kanoulas. 2022. [Temporal knowledge graph question answering via subgraph reasoning](#). *Knowledge-Based Systems*, 251:109134.
- Sitao Cheng, Ziyuan Zhuang, Yong Xu, Fangkai Yang, Chaoyun Zhang, Xiaoting Qin, Xiang Huang, Ling Chen, Qingwei Lin, Dongmei Zhang, Saravan Rajmohan, and Qi Zhang. 2024. [Call me when necessary: LLMs can efficiently and faithfully reason over structured environments](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 4275–4295, Bangkok, Thailand. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

- Yifu Gao, Linbo Qiao, Zhigang Kan, Zhihua Wen, Yongquan He, and Dongsheng Li. 2024. [Two-stage generative question answering on temporal knowledge graph using large language models](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 6719–6734.
- Zhaoyan Gong, Juan Li, Zhiqiang Liu, Lei Liang, Hua-jun Chen, and Wen Zhang. 2025. [RTQA: Recursive thinking for complex temporal knowledge graph question answering with large language models](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 9853–9870, Suzhou, China. Association for Computational Linguistics.
- Gaole He, Yunshi Lan, Jing Jiang, Wayne Xin Zhao, and Ji-Rong Wen. 2021. [Improving multi-hop knowledge base question answering by learning intermediate supervision signals](#). In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, pages 553–561. Association for Computing Machinery.
- Qianyi Hu, Xinhui Tu, Cong Guo, and Shunping Zhang. 2025. [Time-aware ReAct agent for temporal knowledge graph question answering](#). In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 6028–6039.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. [ALBERT: A lite BERT for self-supervised learning of language representations](#). *arXiv preprint arXiv:1909.11942*.
- Yonghao Liu, Di Liang, Fang Fang, Sirui Wang, Wei Wu, and Rui Jiang. 2023. [Time-aware multiway adaptive fusion network for temporal knowledge graph question answering](#). In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 1–5. IEEE.
- Linhao Luo, Yuan-Fang Li, Gholamreza Haffari, and Shirui Pan. 2023. [Reasoning on graphs: Faithful and interpretable large language model reasoning](#). *arXiv preprint arXiv:2310.01061*.
- Costas Mavromatis, Prasanna Lakkur Subramanyam, Vassilis N. Ioannidis, Adesoji Adeshina, Phillip R. Howard, Tetiana Grinberg, Nagib Hakim, and George Karypis. 2022. [TempoQR: Temporal question reasoning over knowledge graphs](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 5825–5833.
- Xinying Qian, Ying Zhang, Yu Zhao, Baohang Zhou, Xuhui Sui, Li Zhang, and Kehui Song. 2024. [TimeR4: Time-aware retrieval-augmented large language models for temporal knowledge graph question answering](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6942–6952.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [DistilBERT, a distilled version of BERT: Smaller, faster, cheaper and lighter](#). *arXiv preprint arXiv:1910.01108*.
- Apoorv Saxena, Soumen Chakrabarti, and Partha Talukdar. 2021. [Question answering over temporal knowledge graphs](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, pages 6663–6676.
- Apoorv Saxena, Aditay Tripathi, and Partha Talukdar. 2020. [Improving multi-hop question answering over knowledge graphs using knowledge base embeddings](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4498–4507.
- Chao Shang, Guangtao Wang, Peng Qi, and Jing Huang. 2022. [Improving time sensitivity for question answering over temporal knowledge graphs](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8017–8026. Association for Computational Linguistics.
- Haitian Sun, Tania Bedrax-Weiss, and William W. Cohen. 2019. [PullNet: Open domain question answering with iterative retrieval on knowledge bases and text](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 2380–2390.
- Haitian Sun, Bhuwan Dhingra, Manzil Zaheer, Kathryn Mazaitis, Ruslan Salakhutdinov, and William W. Cohen. 2018. [Open domain question answering using early fusion of knowledge bases and text](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4231–4242, Brussels, Belgium. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, and 49 others. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.
- Tingyi Zhang, Jiaan Wang, Zhixu Li, Jianfeng Qu, An Liu, Zhigang Chen, and Hongping Zhi. 2024. [MusTQ: A temporal knowledge graph question answering dataset for multi-step temporal reasoning](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 11688–11699.

A Related Work

Temporal Knowledge Graph Question Answering (TKGQA) requires models to answer natural language questions over temporally annotated facts. Compared with conventional KGQA, TKGQA

must jointly consider entity–relation semantics, temporal operators, event ordering, and multi-hop dependencies. Existing studies are closely related to three aspects of our problem setting: controlling error propagation in multi-hop evidence construction, modeling temporal constraints during TKGQA, and using LLMs for complex temporal reasoning.

A.1 Error Propagation in Multi-hop Evidence Construction

Multi-hop KGQA methods usually construct an intermediate reasoning space before predicting the final answer. Early studies build question-specific subgraphs and perform neural reasoning over retrieved evidence. GRAFT-Net constructs a heterogeneous subgraph from knowledge bases and text, while PullNet further retrieves evidence iteratively to expand the question-specific subgraph (Sun et al., 2018, 2019). Later work introduces intermediate supervision to improve multi-hop reasoning stability and reduce spurious paths (He et al., 2021). These methods show that explicit evidence construction is important for multi-hop question answering, but their expansion process is mainly driven by local graph neighborhoods or learned retrieval scores. When irrelevant nodes are introduced early, later reasoning may still proceed over a noisy and enlarged search space.

Recent LLM-enhanced graph reasoning methods further explore planning, path discovery, and interaction with structured environments (Luo et al., 2023; Cheng et al., 2024). Such methods improve semantic flexibility and interpretability, but they often rely on step-wise decisions or generated reasoning paths. This makes the reasoning process vulnerable to accumulated errors when an early path, entity, or intermediate answer is incorrect. In TKGQA, this issue becomes more severe because an intermediate state must be not only structurally plausible but also temporally valid. RTQA explicitly addresses this problem by recursively decomposing complex temporal questions and aggregating multiple answer sources to mitigate error propagation (Gong et al., 2025). However, decomposition-based reasoning still depends on the correctness of intermediate sub-questions and answers, and it does not fully replace the need for a globally constrained reasoning space.

A.2 Temporal Constraint Modeling and Delayed Validation

A central challenge in TKGQA is how to incorporate time into question answering. Early TKGQA methods formulate the task as temporal knowledge graph completion or representation matching. CronKGQA projects question representations into temporal KG embedding spaces, while TempoQR augments question embeddings with contextual, entity-aware, and time-aware modules (Saxena et al., 2021; Mavromatis et al., 2022). Other studies improve time sensitivity through temporal fusion, temporal-aware scoring, or subgraph-based reasoning (Shang et al., 2022; Liu et al., 2023; Chen et al., 2022). These methods demonstrate that temporal information is essential for TKGQA, especially when questions involve implicit time scopes or time-dependent entity states.

More recent benchmarks and models further emphasize temporal granularity and complex temporal constraints. MultiTQ introduces multi-granularity temporal question answering, where questions may involve different time units and temporal operators (Chen et al., 2023). MusTQ further focuses on multi-step temporal reasoning over temporal knowledge graphs (Zhang et al., 2024). TimeR4 incorporates time-aware retrieval and reranking to reduce temporal hallucination in LLM-based TKGQA (Qian et al., 2024). Despite these advances, temporal information is often used as an embedding feature, a retrieval signal, or a final-stage ranking criterion. Such designs can filter obviously invalid answers, but temporally conflicting facts may remain in the intermediate reasoning space until late-stage selection. Therefore, complex TKGQA still requires more explicit temporal-state control during the reasoning process, rather than treating time only as a post-hoc validation signal.

A.3 LLM-enhanced Complex Temporal Reasoning

Large language models have recently been introduced into TKGQA to improve question interpretation, temporal reasoning, and evidence utilization. ARI divides temporal reasoning into knowledge-agnostic and knowledge-based stages, enabling LLMs to use abstract reasoning guidance and factual support (Chen et al., 2024b). Two-stage generative methods combine LLM-generated representations with graph information to improve temporal question answering (Gao et al., 2024). Prog-TQA

represents temporal constraints through executable programs, while TempAgent uses time-aware tools to handle temporal constraints in an agent-style reasoning process (Chen et al., 2024a; Hu et al., 2025). These methods show the benefit of LLMs in parsing complex questions and interpreting temporal semantics.

However, LLM-centered temporal reasoning also introduces new challenges. If the model freely generates intermediate reasoning steps, the process may become difficult to verify against graph evidence. If it repeatedly calls an LLM for decomposition, retrieval, and selection, the framework may incur high inference cost and remain sensitive to early mistakes. TimeR4 and RTQA are representative attempts to strengthen LLM-based TKGQA through time-aware retrieval and recursive decomposition, respectively (Qian et al., 2024; Gong et al., 2025). Nevertheless, the two limitations discussed above remain central: the reasoning process needs stronger control over evidence scope to reduce error propagation, and temporal constraints need to participate more actively in intermediate reasoning rather than being delayed to the final answer stage.