

NEON: News Entity-Interaction Extraction for Enhanced Question Answering

Sneha Singhania*
Max Planck Institute for Informatics
Saarbrücken, Germany
ssinghan@mpi-inf.mpg.de

Allen Herring
Microsoft Research
Redmond, USA
allenh@microsoft.com

Silviu Cucerzan
Microsoft Research
Redmond, USA
silviu@microsoft.com

Sujay Kumar Jauhar
Microsoft Research
Redmond, USA
sjauhar@microsoft.com

Abstract

Capturing fresh information in near real-time and using it to augment existing large language models (LLMs) is essential to generate up-to-date, grounded, and reliable output. This problem becomes particularly challenging when LLMs are used for informational tasks in rapidly evolving fields, such as Web search related to recent or unfolding events involving entities, where generating temporally relevant responses requires access to up-to-the hour news sources. However, the information modeled by the parametric memory of LLMs is often outdated, and Web results from prototypical retrieval systems may fail to capture the latest relevant information and struggle to handle conflicting reports in evolving news. To address this challenge, we present the NEON framework, designed to extract emerging entity interactions—such as events or activities—as described in news articles. NEON constructs an entity-centric timestamped knowledge graph that captures such interactions, thereby facilitating enhanced QA capabilities related to news events. Our framework innovates by integrating open Information Extraction (openIE) style tuples into LLMs to enable in-context retrieval-augmented generation. This integration demonstrates substantial improvements in QA performance when tackling temporal, entity-centric search queries. Through NEON, LLMs can deliver more accurate, reliable, and up-to-date responses.

CCS Concepts

• **Information systems** → **Question answering**; *Information extraction*.

Keywords

Temporal Question Answering, Information Extraction, Large Language Models

*Work done while interning at Microsoft Research

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
WWW '25, June 03–05, 2018, Woodstock, NY
© 2025 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-XXXX-X/18/06
<https://doi.org/XXXXXXXX.XXXXXXX>

ACM Reference Format:

Sneha Singhania, Silviu Cucerzan, Allen Herring, and Sujay Kumar Jauhar. 2025. NEON: News Entity-Interaction Extraction for Enhanced Question Answering. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (WWW '25)*. ACM, New York, NY, USA, 10 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 Introduction

In recent years, large language models [11, 40], pretrained on extensive web-scale datasets, have gained widespread usage as human task assistants, offering expert-level knowledge learned during pre-training [39]. However, given the dynamic nature of information in the real world, incorporating temporal awareness into these AI systems has become increasingly important. The ability to inject up-to-date information is essential for effectively addressing user needs in a fast evolving world. Despite the notable capabilities of these systems, research on downstream tasks involving temporal aspects has uncovered several limitations. Key issues include temporal misalignment [18, 30], lack of temporal generalization [21], limited temporal reasoning capabilities [52], and disproportionate handling of implicit versus explicit temporal conditions when answering knowledge-intensive questions [23].

These challenges are particularly acute in the news domain, where new content is generated rapidly, and producing reliable, grounded responses cannot solely depend on static, pretrained knowledge. Several studies have demonstrated that LLMs can perform well on domain-specific question answering tasks [4, 22, *inter alia*], including on popular news datasets. More recent approaches leverage LLMs through the retrieval-augmented generation (RAG) paradigm [14, 26, 43] to address some of the aforementioned limitations. By using trained retrievers to fetch relevant supporting documents, these methods aim to enhance the accuracy of generated responses. However, there has been relatively limited focus [23, 37, 46] on systematically retrieving temporal text snippets for answering news-related questions.

In this work, we focus on generating reliable and temporally grounded responses to user queries, with a particular emphasis on entity-centric queries. Our choice is motivated by several factors: (i) users frequently inquire about entities during real-world events [13, 27], (ii) entity-centric information tuples of the form (*subject*, *predicate*, *object*) are highly valuable in Knowledge Base Question Answering (KBQA) tasks [20, 34], and (iii) user queries

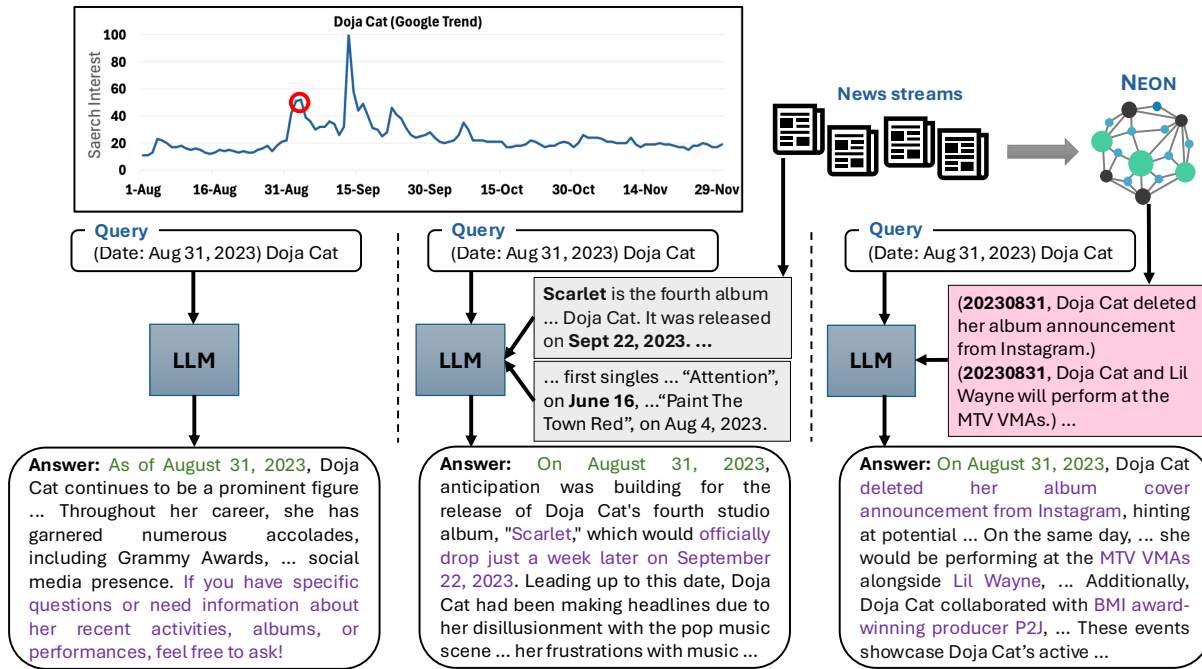


Figure 1: Example for entity-centric, time-specific QA. Graph shows search interest for *Doja Cat* over a four-month period. The bottom part illustrates response generation at one of the peaks (31 August) using three different techniques: (i) zero-shot prompting, (ii) news snippets based prompting, (iii) augmenting tuples from our NEON graph for enhanced answer generation.

are often vague, ambiguous, and telegraphic, with implicit intent tied to the mentioned entities [3, 44]. For example, Figure 1 shows the Google trend line (<https://trends.google.com/trends>) for the popular entity *Doja Cat* during the August–November, 2023 time frame, exhibiting peak search interest on specific dates, particularly around major events, e.g., on August 31 and September 15.

When a query such as “*Doja Cat*” or “*Doja Cat* news” is given in the prompt to an LLM such as GPT-4 [39], the model likely generates an entity-centric response based on sources like Wikipedia, one of the largest resources in LLM pretraining datasets [7, 42]. Although Wikipedia provides valuable overviews and highlights about key events, it often lacks the granular, timely updates required to answer entity-centric queries involving recent developments. Even if the user query is reformulated to specify a date—e.g., “(Date: August 31, 2023) *Doja Cat*”—LLMs often struggle to produce precise, up-to-date response, in particular when the date in question is beyond the LLM’s knowledge cut-off and the system must defer to external sources through RAG mitigation.

This limitation is illustrated in Figure 1, where we selected a random peak in the search interest plot. Without additional information sources and/or additional context, GPT-4 provides a generic response to the prompt “(Date: August 31, 2023) *Doja Cat*”. However, we observe improvements in temporal relevance and accuracy through RAG mitigation when supporting snippets retrieved from news streams are incorporated into the LLM’s prompt. Here, LLM still needs to reason and extract relevant information across these multiple snippets. As Liu et al. [28] note, both the size of the context and the positioning of relevant information within the context

significantly impact model performance. Instead of directly augmenting chunks of information, we propose NEON (News Entity Interacti**ON**s) that captures interactions such as events and activities between entities in a more compact form (single sentences or propositions) and use these for enhanced response generation.

NEON is an entity-centric, time-stamped graph constructed from a rich stream of news sources. Unlike prior information extraction (IE) methods that leverage LLMs, we identify, disambiguate, and label the entities from news streams first, and then employ the entity markup information in LLM prompts to generate open information extraction (openIE) style semi-structured news entity-interactions. This method helps us achieve higher recall by avoiding the bottleneck of canonicalization. Once NEON is constructed, we use various temporal indexing strategies to effectively retrieve relevant information from NEON for answering temporal user queries. For instance, in Figure 1, two NEON entity-interactions related to *Doja Cat* are augmented for response generation.

We test the effectiveness of NEON by collecting 3,000 real-world queries from the Bing search engine on specific dates for 50 sampled entities. These entities were chosen with various considerations in mind, including demographic diversity, the balance between popular and long-tail entities, and a mix of person and organization-type entities, as explained further in Section 3. To perform temporal QA generation, we augment the LLM prompts with retrieved NEON tuples (illustrative examples provided in Table 1). Since answers to temporal queries can be subjective, and LLMs have proven to be effective evaluators [29, 50], particularly on downstream tasks within the news domain [5, 53], we employ LLMs to automatically

$NEON(M_1)$	(20230502, Doja Cat made debut appearance at Met Gala)
	(20230502, Doja Cat dressed as Chouette)
	(20230502, Doja Cat was styled by Brett Alan Nelson)
	(20230831, Doja Cat will perform at the 2023 MTV VMAs)
	(20230831, Doja Cat 's announcement debuting the cover art for "Scarlet" was removed from Instagram)
$NEON(M_2)$	(20230502, Doja Cat and Jared Leto paid homage to Lagerfeld)
	(20230531, Doja Cat and Demi Lovato will perform at VMAs)
	(20230831, Doja Cat collaborates with Afrobeats producer P2J)
	(20230831, Doja Cat posted on Twitter)
	(20230831, Doja Cat deleted a post on Twitter)

Table 1: Samples from NEON variants. The subject and object entities are bold-faced.

assess the quality of responses on a 3-point Likert scale. We also corroborate the automatic evaluations through human assessment. Our experiments reveal that integrating NEON entity-interactions into LLM prompts improves the temporal relevance and overall quality of the answers.

The salient contributions of this work are:

- (i) We formulate the problem of generating enhanced responses for temporal entity-centric queries.
- (ii) We present a novel approach for extracting openIE style entity-interactions tuples from news streams.
- (iii) We evaluate the generated answers using both automated LLM-based scoring and human evaluations on a diverse set of real-world queries, demonstrating significant improvements in response quality when incorporating NEON.

2 NEON

In this section, we present NEON, an openIE-style knowledge graph for modeling entity-interactions from news streams. We outline the graph construction pipeline, describe optimized indexing techniques for efficient temporal retrieval, and then detail our approach for temporal question-answering, demonstrating how NEON enables accurate, contextually relevant responses.

2.1 Graph Construction

NEON is a knowledge graph wherein nodes are entities and edges capture interactions (events or activities) between these entities. Formally, NEON is defined as a graph $\mathcal{G} = (E, T, R)$, where E is the set of entities, I is the lexicalized interaction descriptions or propositions in natural language, T is the time frame, and R refers to timestamped entity-interactions related both subject (s) and object (o) entities, defined as $R = (t, NEI(s, o))$, with $s, o \in E$ and $t \in T$.

To construct NEON, we process a large set of news streams from various internet sources. Each article is time-stamped according to its publication date, typically extracted from the source URL. The construction pipeline follows these steps:

1. **Entity identification:** Each news story is preprocessed to extract only the main content, which is analyzed using an entity linking tool called NEMO [9]. This step annotates all named entity mentions in text, including co-references and additional information like dates and addresses, with a simple XML markup.

2. **Sentence segmentation:** The article content is segmented into sentences using a proprietary tool, similar to spaCy¹, and entities identified in previous step are tracked within each sentence.
3. **Text chunking:** The article's main content is then split into overlapping chunks (C) of m sentences (where m is a hyperparameter). For indexing, each chunk ($c \in C$) is represented as $c = (X_{1\dots m}, E', t)$, where $X_{1\dots m}$ denotes a set of m sentences, E' is the set of named entities identified within $X_{1\dots m}$ and $E' \subset E$, and t is the article's timestamp.
4. **Chunk Deduplication:** To manage paraphrased information across different chunks, near-duplicate chunks are filtered out using trigram representations and Jaccard similarity.

Once the set of news chunks is obtained, next we explain how information extraction (IE) techniques that employ LLMs are applied to each chunk to construct the NEON graph.

Overview. We design task-specific prompts for the LLM that explore two IE variants for graph construction. In the first variant, termed $NEON(M_1)$, LLM generates news entity-interactions R using only the subject s and the corresponding news chunk c that mentions s . In the second variant, R is generated by incorporating both the subject s and object o entities, along with the news chunk c that contains mentions of both (s, o).

The first variant focuses on entity-interactions centered around a single subject entity, making it ideal for cases where the entity (e.g., a person or organization) is the primary focus of the news chunk. The second variant, which includes both subject and object entities, enables the model to infer direct relationships or events between the pair, resulting in a more detailed graph that captures complex associations and evolving scenarios more accurately (see Section 4). Both approaches support a wide range of interactions without relying on predefined relationships. Given our goal of solving temporal QA, we prompt the LLM to directly generate news entity-interactions R in a lexicalized form [19, 38], thereby enabling NEON to be directly augmented into the temporal QA prompt.

$NEON(M_1)$. Our primary objective of the M_1 variant is to automatically construct an entity-centric, time-stamped KB using a target set of subject entities, facilitated by LLM prompting. The model identifies neighboring nodes (object entities already marked in text) and the connecting edges (interactions or descriptions) between the subject and object entities using the news chunks included in the prompt. The process is formally outlined in the following steps:

1. **Target Subjects:** Begin with a target set of subjects S , s.t. $S \subset E$.
2. **Chunk Retrieval:** For each subject $s \in S$, retrieve all the news chunks $c = (X_{1\dots m}, E', t)$ that contain mentions of s , i.e., $s \in E'$.
3. **Prompt Construction:** Create a prompt by explicitly mentioning s and the corresponding news chunk c that employs the entity markup (" $<e>$ " and " $</e>$ ") in $X_{1\dots m}$ to identify the entities E' in the news chunk. Employ the prompt with LLM to generate a list of valid news entity-interactions $r_s = NEI(s, o)$, where o is drawn by the LLM from the marked up entities, NEI is LLM's implicit verbalization function. Each news entity-interaction is assigned a timestamp, resulting in $r_s = (t, NEI(s, o))$, where t is derived from the timestamp of c .

¹<https://spacy.io/api/sentencizer>

4. **Iteration:** Repeat this process for all subject-chunk pairs (s, c) , iteratively building the KB for each subject entity, $\mathcal{M}_1 = \bigcup_{s \in S} r_s$.

Although this approach emphasizes information related to the target entities, scaling up the set S can lead to a more diverse and comprehensive KB. The chunk size (m in $X_{1...m}$) can be large enough to provide detailed context around the subject entity while still fitting within the LLM’s context window. This setup enables a single prompt to generate a longer list of relational tuples, with the added context potentially improving entity-interaction extraction. The prompt template for $\text{NEON}(\mathcal{M}_1)$ construction is shown in A1.

NEON (\mathcal{M}_2). In contrast to \mathcal{M}_1 , core idea behind \mathcal{M}_2 is to leverage the co-occurrence of entity mentions within the news chunks as potential target pairs of subject and object entities. We start with the predefined set of subject entities and identify entity co-occurrences in the news chunks, then determine candidate entity pairs based on their TF-IDF scores. Once the target pairs are obtained, the LLM is prompted to perform only the interaction extraction task. The overall process mirrors that of $\text{NEON}(\mathcal{M}_1)$, but with slightly different initializations. Specifically, the steps are as follows:

1. **Target Pair Selection:** Begin with a target set of subjects S , s.t. $S \subset E$. Then identify co-occurring entity pairs (s, o) , where $s \in S, o \in E$, within all news chunks with prominent TF-IDF scores. These pairs become the target set for interaction extraction.
2. **Chunk Retrieval:** For each entity pair (s, o) , retrieve all relevant news chunks $c = (X_{1...m}, E', t)$, where $s, o \in E'$.
3. **Chunk Batching:** Since (s, o) are predefined, the chunk context can be narrowed to focus on (s, o) , enabling multiple chunks to fit within a single prompt for extraction. Hence, retrieved chunks are partitioned into batches of size k .
For a set of n retrieved news chunks C sorted based on its time-stamp, batching gives $C_B = \{C_{1...k}, \dots, C_{n-k+1...n}\}$, where each batch contains k chunks. If n is not divisible by k , the last batch will contain $<k$ chunks. This time-based sorting and batching approach allows information from similar dates—and thus similar events—to be grouped within the same prompt.
4. **Prompt Construction:** Construct a prompt by explicitly including both the subject s and object o , along with a chunk batch $C' \in C_B$. In each chunk within the batch, use markups (“<e>” and “</e>”) to highlight all entity mentions. The LLM is then prompted to generate news entity-interactions $r_{so} = \text{NEI}(s, o)$, where NEI is LLM’s implicit verbalization function. Since the news chunks in C' may contain different time-stamps, we collect the unique time-stamps T' from each chunk. The extractions are then expanded as $R_{soC'} = \bigcup_{t \in T'} (t, \text{NEI}(s, o))$.

5. **Iteration:** Repeat this process for all identified (s, o) target pairs and chunk batches $C' \in C_B$ to form a consolidated KB $\mathcal{M}_2 = \bigcup_{(s,o)} \bigcup_{C'} R_{soC'}$.

This variant often yields more generations due to explicit (s, o) pairs; however, it may occasionally cause the LLM to generate sentences involving (s, o) without a valid interaction, or return an empty list if no interaction is confidently detected. Figure A2 gives the prompt for $\text{NEON}(\mathcal{M}_2)$ construction. The chunk batch size k (hyperparameter) directly affects performance: a small k is cost-inefficient with more LLM calls and reduced contextual detail,

while a large k risks “lost-in-the-middle” [28] and information overload [48] issues. Table 1 gives sample generations for both NEON variants. The subject is the entity *Doja Cat* in both variants and the object varies for $\text{NEON}(\mathcal{M}_2)$.

2.2 Temporal Datastore

Each NEON tuple includes a *time-stamp*, followed by the LLM-generated news entity-interaction. We represent NEON as a datastore (\mathcal{D}) with each entry as a tuple (t_d, d) , where d is the LLM-generated news entity-interaction, and t_d is the time-stamp of d . Following prior work [17, 36], we employ off-the-shelf models to perform dense indexing of \mathcal{D} , creating a temporal datastore that supports efficient retrieval. User queries Q (original and reformulated) are similarly converted into dense vectors using the same model, allowing retrieval of the top- k relevant entity-interactions by maximizing semantic similarity. This approach is retriever-agnostic, allowing for integration with any suitable retriever model.

To ensure temporal relevancy while retrieving content for answering each user query $q \in Q$, we reformulate the queries to explicitly mention the normalized timestamp $q' = \text{REFORM}(t_q, q)$, where t_q is the query timestamp. We then experiment with the following two retrieval strategies:

1. **Temporal Retrieval:** Retrieve the top- k tuples from \mathcal{D} using q , where the tuples timestamp t_d exactly matches the query timestamp t_q . If fewer than k tuples match t_q , we retrieve additional tuples within a $\pm r$ day range of t_q to complete the top- k results, with r being a hyperparameter.
2. **Generic Retrieval:** Retrieve the top- k tuples solely based on semantic similarity between q' and each document in \mathcal{D} .

While temporal retrieval can enhance relevance for time-specific queries, it may yield zero tuples when the news coverage is sparse for query dates. In contrast, generic retrieval consistently provides tuples by first retrieving temporally relevant tuples and, if necessary, uses a back-off approach to additionally retrieve semantically relevant tuples.

2.3 Temporal Question Answering

The final step involves generating answers for temporal queries by integrating NEON tuples into LLM prompts. This enables the system to deliver contextually relevant, temporally accurate, and up-to-date answers. The key steps in our approach are as follows:

1. User enters a (telegraphic) temporal query q at timestamp t_q
2. q is reformulated into q' by performing named-entity disambiguation [8] and explicitly mentioning t_q in natural language.
3. Using the reformulated query q' , we apply the retrieval methods mentioned in Section 2.2 to obtain top- k relevant NEON tuples.
4. The top- k NEON tuples are augmented into LLM prompt for RAG based generation. Figure 2 gives the prompt template.
5. LLM generates a response to the temporal query using the supporting information.

Figure 2 provides an overview of our temporal QA pipeline, illustrating RAG mitigation for response generation and demonstrating how the NEON graph integrates into this system. Unlike conventional QA systems, our aim is to generate relevant and temporally reliable responses.

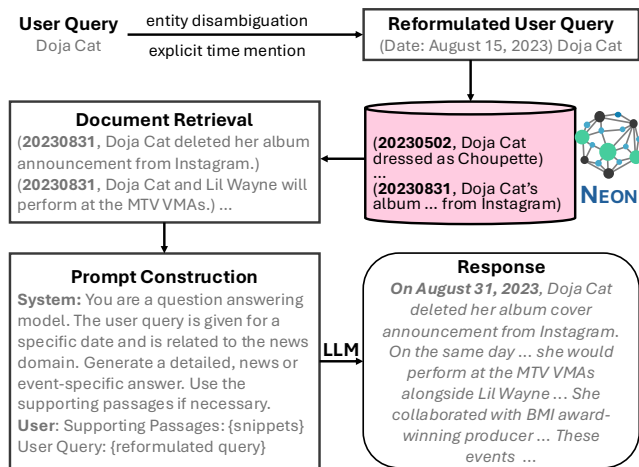


Figure 2: Implemented temporal QA pipeline

3 Experimental Setup

Entities and Queries. We selected 50 entities by using stratified sampling across four categories—artists, companies, leaders, and pioneers—from Time’s 100 most influential people² and Fortune500 companies³ to ensure a balanced diversity in category, demographics, and popularity. For each entity, we collected a large volume of user queries from Bing search logs over a four-month period (August–November, 2023). Because we aimed to generate temporal answers for emerging events, we focused on entity-centric queries logged during *spiking dates*—periods of heightened entity activity/presence in user queries.

Spiking dates were identified through short-term trends in query volume. Specifically, we calculated a 3-day rolling sum of daily query counts, flagging significant spikes as dates where the rolling sum exceeded one standard deviation above the mean. This threshold effectively identifies periods of elevated user interest, often triggered by shifts in attention or external events impacting the entity. Furthermore, to preserve user privacy, we filtered queries requested by fewer than five distinct individuals, yielding nearly 3000 temporal user queries. Table 2 provides a few examples of representative queries from each category and entity popularity.

News Articles. To construct NEON, we collected news streams from 500 distinct sources, with a median of 200 articles per source and a maximum of 9,000 articles. Articles varied in length and were preprocessed to create metadata information, including disambiguated named entity mentions (e.g., persons, orgs., and locations), their surface forms, and article timestamps. The metadata helps in mapping target subject entities to their corresponding news articles. We also identified duplicate content across articles and publication dates; for temporal relevance, we retained unique content mapped to all associated sources and dates to construct a more complete NEON graph. Table 3 gives the final dataset statistics. The counts of head/tail entities are based on 70th pth of the total article count.

Figure 3 shows the count of unique news sources per subject entity, grouped by the four categories. Dashed lines at 25%, 50%, and

	Entity	Type	Date	Query
Artists	Neil Gaiman	tail	Sept 5	neil gaiman award
			Sept 5	which novel did neil gaiman coauthor with terry pratchett
	Austin Butler	head	Aug 17	austin butler as elvis
			Aug 18	what are the latest projects of austin butler
Companies	Berkshire Hathaway	tail	Nov 6	berkshire hathaway cash pile
			Nov 7	berkshire hathaway stock
	Alphabet Inc.	head	Aug 24	alphabet stock price today
			Sept 14	alphabet layoffs
Leaders	Sherry Rahman	tail	Aug 25	climate financing by rahman
			Oct 5	pak minister sherry rehman
	Olaf Scholz	head	Sept 3	olaf scholz falls
			Sept 4	olaf scholz eye patch
Pioneers	Bella Hadid	tail	Oct 27	bella hadid cancelled
			Oct 31	bella hadid closing coperni
	Doja Cat	head	Aug 28	doja cat beefs with fans
			Aug 29	doja cat feels free

Table 2: User query samples from our dataset

Entity sources	Time100 and Fortune500
No. of subjects	50
No. of head subjects	15
No. of tail subjects	35
No. of news domains	500
Median no. of news	~200
Max no. of news	~9000
No. of user queries	3000
Spiking date range	August–November 2023

Table 3: Dataset Statistics

75% on the y-axis represent quartile marks, offering a visual reference for comparing entity news coverage. Distinct patterns emerge across categories: leaders and companies dominate the news space, while artists and pioneers, though influential, receive selective coverage, with some receiving minimal or near-zero coverage.

Evaluation. Since our temporal QA task involves generating subjective responses for evolving news events, traditional evaluation methods prove challenging due to the absence of an established ground-truth dataset. To address this, we employed an automated evaluation framework that leverages LLMs, which have recently gained prominence [24, 49], even in news-related tasks [5, 53], as an efficient means for output assessment. Our evaluation focuses on three key attributes: (1) *Helpfulness*, assessing the extent to which the response addresses the user’s query by incorporating additional relevant entities—such as people, events and locations—beyond the primary subject; (2) *Relevance*, ensuring that the response is accurately aligned with the context of the query, delivering information that is specific to the date or time frame referenced; (3) *Faithfulness*, evaluating the extent to which the response is grounded in the supporting passages provided, relying exclusively on verifiable information from these sources to produce a trustworthy answer.

²<https://time.com/collection/100-most-influential-people-2023>

³<https://fortune.com/ranking/fortune500/2023/>




Figure 3: Coverage for 50 entities in 500 diverse news sources over a period of one year (2023)

Each attribute is rated on a 3-point Likert scale, in alignment with established practices in prior evaluation studies [32, 50].

Notably, work by Dettmers et al. [10] observed that the order and detail of attribute ratings can influence evaluation outcomes. To mitigate potential bias, we prompt the LLM separately for each attribute, tripling the number of LLM calls but reducing the risk of cross-attribute contamination. To further ensure robustness in our evaluation, we perform few-shot evaluation using human annotated in-context examples on the same attributes. The actual prompt templates are provided in Table A4.

Baselines & Methods. We aim to augment LLMs with temporally relevant information for answering entity-centric questions and compare various augmentation strategies. Unlike existing RAG methods, which retrieve passages from large generic data indices, we employ strong baselines that retrieve from an entity-centric datastore (detailed in Section 2.2). We evaluate the following methods within the RAG framework to address our temporal QA task: (i) **NewsRAG**: Retrieves top- k temporally indexed news chunks C , used in constructing NEON (step 2 in each variant of Section 2.1). (ii) **WebRAG**: Uses Bing to retrieve top- k web snippets, explicitly incorporating timestamps to access diverse and dynamic sources. (iii) **NEON(\mathcal{M}_1)**: Employs the process in Section 2.1; this variant includes 109,246 entries. (iv) **NEON(\mathcal{M}_2)**: Employs the process in Section 2.1; this variant includes 455,680 entries.

Each method uses the reformulated user query for retrieval, augments the top- k retrieved information directly into the prompt, and employs GPT-4o for response generation. Other openIE-based methods are unsuitable for our purposes as they explicitly model knowledge tuples, which we aim to avoid. Comparing the performance across these methods reveals how different forms of real-time information impact helpfulness, relevance, and faithfulness in addressing the temporal demands of our QA task.

4 Results

Tables 4 and 5 presents the main results, comparing system performance across *temporal* and *generic* retrieval strategies using a 3-point Likert scale across the evaluation attributes and an average

	Method	Helpful	Relevant	Faithful	Avg.
Temporal	NewsRAG	1.03	0.83	0.83	0.9
	NEON(\mathcal{M}_1)	1.03	0.82	0.86	0.9
	NEON(\mathcal{M}_2)	1.04	0.79	0.87	0.9
Generic	NewsRAG	0.9	1.24	1.08	1.07
	NEON(\mathcal{M}_1)	1.38	1.12	1.18	1.23
	NEON(\mathcal{M}_2)	1.37	1.03	1.17	1.19
	WebRAG	1.54	1.29	1.22	1.35

Table 4: Performance with zero-shot evaluation prompts.

	Method	Helpful	Relevant	Faithful	Avg.
Temporal	NewsRAG	1.53	1.45	1.46	1.48
	NEON(\mathcal{M}_1)	1.53	1.50	1.52	1.52
	NEON(\mathcal{M}_2)	1.55	1.54	1.56	1.55
Generic	NewsRAG	1.59	1.41	1.05	1.35
	NEON(\mathcal{M}_1)	1.42	1.34	1.05	1.27
	NEON(\mathcal{M}_2)	1.45	1.35	1.09	1.30
	WebRAG	1.57	1.57	1.41	1.52

Table 5: Performance using few-shot learning and reasoning based evaluation prompts.

across the 3 scores. WebRAG is invariant to retrieval strategy since Bing combines aspects of both temporal and generic retrieval. Best scores are highlighted, while second-best scores are boldfaced.

Beginning with Table 4 we notice that in zero-shot setting WebRAG achieves the highest performance across all three attributes; this is unsurprising considering that it is achieved by the help of a full-fledged web-scale search engine. Meanwhile, NEON variants consistently outperform NewsRAG on both retrieval strategies for the helpfulness and faithfulness attributes, demonstrating that temporally relevant NEON KB entries improve generation reliability. Between NEON variants, both perform similarly on helpfulness and faithfulness, but NEON(\mathcal{M}_1) achieves a higher relevance scores.

Few-shot learning with in-context examples and chain-of-thought prompts have been shown as being more effective for downstream tasks using LLMs [51]. We apply this strategy and present the results in Table 5. Unlike zero-shot evaluation, few-shot results show that temporal retrieval consistently outperforms generic retrieval across all methods. Additionally, much more significant gains are observed with NEON and NewsRAG methods. The results rival those of WebRAG in many instances, and NEON(\mathcal{M}_2) achieves the highest overall average score across metrics. Moreover, each method shows specific strengths: NewsRAG produces the most helpful generations, WebRAG excels at relevance due to its broad and diverse sources, while NEON delivers the most reliable outputs, leveraging its structured entity-focused KB design. NEON(\mathcal{M}_2) specifically benefits from a larger KB (4x the size of \mathcal{M}_1) and captures diverse entity-interactions by explicitly using subject and object mentions employed in the KB construction.

Overall, these results demonstrate that NEON—augmented with in-context examples—is capable of rivaling the performance of a live web-scale search engine for temporal entity-centric QA. Additionally, the individual strengths of the different approaches suggest




Figure 4: Performance comparison with varying top- k parameter, temporal retrieval and few-shot prompts.

a hybrid system that can leverage them jointly for even better performance; we leave such explorations to future work.

Varying Context Length. The news chunks in NewsRAG typically contain 5-6 sentences, WebRAG uses 2-3 sentence web snippets, and NEON variants are single-sentence entries. To ensure comparable prompt context lengths across the methods, the no. of top- k documents is method-specific. Therefore, in Tables 4 and 5, NewsRAG and WebRAG each retrieve the top-5 documents, while NEON variants use top-10 documents. This creates a trade-off between detail and diversity, each of which may benefit a certain kind of query or response. To investigate this, we compare and present the performance with varying k , as illustrated in Figure 4. The order in performance for the studied models is relatively stable with the increase of k , with WebRAG leading NEON on helpfulness and relevance, most likely because of the capabilities of Bing to operate over a much larger search space. Meanwhile NEON remains the most faithful, but we notice a substantial drop for all models when doubling of supporting passages from 5 to 10, which appears to be more acute when passages are longer (NewsRAG and WebRAG).

Human Assessment. To better understand and validate the automated evaluation results, we conducted a human assessment focusing on our 3 evaluation attributes. Annotators rated 100 randomly sampled responses across all methods using a 3-point Likert scale. To ensure consistency, annotators were provided with detailed descriptions of each attribute, method-specific supporting documents included in the LLM prompt for temporal QA, and examples illustrating each attribute-scale combination. External searches were disallowed, and ratings were based solely on the reformulated query, supporting documents, and LLM-generated responses.

The Pearson correlation between human assessments and few-shot prompt-based automated evaluations showed weak positive correlations: 0.24 (helpfulness), 0.12 (relevance), and 0.19 (faithfulness). Inter-annotator reliability, measured on the overlapping 10% of the responses from 3 annotators using Krippendorff’s alpha, indicated strong agreement with scores of 0.71 (helpfulness), 0.77 (relevance), and 0.61 (faithfulness).

A drill-down analysis of the ratings revealed that automated and human scores matched approximately 50% of the time. Discrepancies in the remaining cases often stemmed from mismatches in timestamps across the query, supporting passages, and LLM responses. Automated evaluations consistently assigned a score of 0 in these cases, whereas human annotators used common-sense reasoning and contextual understanding of the dates to assign scores of 1 or 2. Addressing this limitation in automated evaluation is

a goal for future work. In particular, such mismatches were less frequent in samples from NEON(M_2).

Interestingly, the ratings for helpfulness and relevance, both human and automated, exhibited a high correlation (>0.8). This is likely because a highly relevant response—one that accurately reflects the query’s specific time frame—almost always leads to a helpful response. The reverse however is not always true, as a helpful response may not strictly adhere to temporal relevance.

5 Qualitative Analysis

We perform a qualitative analysis to gain insights into the successes and failures of each method. By performing such an inspection, we identified patterns that either support or hinder the generation of temporally relevant responses. Table 6 highlights key cases.

Entity-centric temporal RAG enhances reliability. As emphasized earlier, unlike traditional RAG techniques, our approach leverages an entity-centric temporal datastore, which improves retrieval accuracy. This advantage is evident in methods using news chunks (NEON variants and NewsRAG), as Bing search in WebRAG relies on surface form matches that often retrieve irrelevant snippets, particularly for long-tail entities with low search hit rates (ex.1 in Table 6). Consequently, the entity-centric design of NEON leads to more reliable responses.

WebRAG benefits from non-EN information and LLMs. We noticed that Bing search sometimes retrieves noisy passages together with relevant passages that are in languages other than English, particularly for subject entities from non-English demographic regions. In such cases, the LLM is still able to generate temporally relevant responses in English based on the non-English passages (ex.2 in Table 6). This gives WebRAG an advantage over others, as the news streams we employed are exclusively in English.

Deviations from rating instructions. We observed that LLMs do not always adhere to the provided evaluation instructions. Despite the prompt clearly outlining the criteria for each 3-point Likert rating, the model occasionally generated ratings outside the intended range along with reasons for such ratings. We noticed a response that was assigned a rating of 3 when the model described it as somewhat relevant and responses with ratings of 5 and 9 when they were deemed fully reliable and relevant (ex.3 in Table 6).

Response lengths correlate with Likert scores. We analyzed whether longer temporal responses tend to receive higher Likert scores by examining the correlation between ratings and response length (measured in characters). Our analyses (see Table A3 in the Appendix for more details) reveal the following trends: (i) in case of

	Query	Date	Method	Retrieved Passages	Temporal QA Response
ex. 1	Sherry Rehman news	Oct 18, 2023	WebRAG	[1] Waheeda Rehman Biography:- ... born in Chengalpet ... [2] Sherry Rehman, parlementaire du Pakistan ... [3] For a sweeter white ... substituting sherry for white wine ... [4] ... PX Sherry Cask Matured limited whisky ... [5] Waheeda Rehman (born 3 February 1938) is an Indian ...	On October 18, 2023, Sherry Rehman, a prominent member of the Pakistan People’s Party, is in the news due to her involvement in a legal case concern ... This development highlights the ongoing and contentious debate surrounding blasphemy laws in Pakistan and the risks faced by those who seek to reform them.
ex. 2	Olaf Scholz eye patch	Sept 4, 2023	WebRAG	[1] Chancellor ‘Pirate Olaf’ Scholz injured jogging ... [2] ... Olaf Scholz. Can someone Photoshop an eye patch ... [3] Après avoir chuté lors d’un footing ce samedi 2 septembre 2023, le chancelier allemand Olaf Scholz s’est ... [4] ... Olaf Scholz Pirat is not real ... [5] Olaf ... ready for the great pirate era ...	On September 2, 2023, German Chancellor Olaf Scholz sustained minor facial injuries while jogging ... he appeared in public on Monday, September 4, 2023, wearing an eye patch, which has led to a wave of humorous and supportive reactions online ... playful nickname “Pirate Olaf” among social media users ...
	Query	Date	Method	Temporal QA Response	LLM Rating
ex. 3	Sam Altman Microsoft	Nov 22, 2023	NEON(M_2)	As of November 22, 2023, Sam Altman has been significantly involved with Microsoft ... OpenAI, which Altman co-founded, is 49% owned by Microsoft ... a move supported by Nadella ... In summary, as of late November 2023, Sam Altman continues to play a pivotal role in the collaboration ...	{“reason”: “The response provides a detailed and relevant answer to the user’s query about Sam Altman’s involvement with Microsoft as of November 22, 2023. It includes information about his role, the financial ... The response is accurate, informative, and directly addresses the user’s question”. “rating”: 5}

Table 6: Examples of queries and LLM responses based on retrieved passages and a corresponding LLM rating (ex. 3)

helpfulness and relevance, both mean and median response lengths increase with higher ratings, particularly for NewsRAG, indicating that longer responses are more likely to receive higher scores. (ii) for faithfulness, while mean and median response lengths do not consistently increase with higher ratings, NewsRAG shows a clear positive correlation between longer responses and higher ratings.

As NewsRAG uses the longest supporting information—5-6 sentence news chunks—these findings suggest that for robust automated evaluation, it is valuable to augment concise yet informative documents, similar to NEON, into LLM prompts to promote consistent response lengths. Notably, LLM-generated reasons for a rating of 2 frequently included the phrases “the response provides detailed and relevant information” for helpfulness and relevance, and “the response accurately reflects the information” for faithfulness.

6 Related Work

Information Extraction. Extracting structured information from unstructured text has been a central focus of NLP research [15]. A key task in IE is relation extraction (RE), which identifies relationships between entities using named entity recognition, entity linking, and relational classification. Accurate relation identification, however, requires deep semantic understanding of context and is often constrained by the fixed relation types in traditional classifier-based RE methods [31].

To address these limitations, openIE enables entities and relations to be extracted as surface forms or phrases [25]. Both RE and openIE are vital for constructing and enriching knowledge graphs [33, 47]. Recent advances explore LLMs as implicit KBs [41], leveraging their outputs for IE tasks via prompting [1, 6, 16]. Our work leverages LLMs for openIE-style lexicalized extraction on news chunks, capturing interactions such as emerging events and associations between entities.

KGs for Temporal QA. Temporal QA aims to generate answers for queries with time-specific constraints [35, 45]. KGs provide

structured and explainable information that can complement LLMs in answering questions. However, these approaches have limitations in handling evolving entity-specific temporal queries and mostly focus on temporal reasoning tasks [19, 45]. Our work builds on these existing methods by introducing a novel way to incorporate entity-interactions from news streams into the QA process, thereby addressing the temporal and entity-specific challenges.

Retrieval Models. Recent research has focused on retrieval-augmented LLMs, which combine information from external sources with parametric models to improve the factuality of text generation [2, 14, 26]. But extensive study on entity-centric temporal RAG remain scarce. For instance, GraphRAG [12] builds entity KGs and generates summaries for closely related entities to support query responses. While effective at capturing global dependencies, these summaries can overlook nuanced, temporally evolving entity-interactions that are critical for reliable temporal QA.

7 Conclusion

NEON presents a novel approach to enhancing QA systems by leveraging entity-interactions extracted from news streams. By constructing an entity-focused time-stamped graph and integrating its information into LLM prompts, NEON addresses the limitations of existing QA methods in handling temporal and entity-specific queries. Our experimental results demonstrate the effectiveness of NEON in generating relevant and reliable answers, making it a valuable addition to the QA landscape.

Future work will focus on further improving the NEON system by incorporating more sophisticated entity-interaction extraction techniques, strengthening the evaluation by expanding the dataset to include more entities and queries, and exploring additional applications to this framework such as enterprise and personal data integration.

References

- [1] Badr AlKhamissi, Millicent Li, Asli Celikyilmaz, Mona Diab, and Marjan Ghazvininejad. 2022. A Review on Language Models as Knowledge Bases. arXiv:2204.06031 [cs.CL]
- [2] Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2024. Self-RAG: Learning to Retrieve, Generate, and Critique through Self-Reflection. In *12th International Conference on Learning Representations*. ICLR, Vienna, Austria.
- [3] Krisztian Balog. 2018. Entity-Oriented Search.
- [4] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Matusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *NeurIPS*, Vol. 33. NeurIPS, virtual, 1877–1901.
- [5] Cheng-Han Chiang and Hung-yi Lee. 2023. Can Large Language Models Be an Alternative to Human Evaluations?. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. ACL, Toronto, 15607–15631.
- [6] Roi Cohen, Mor Geva, Jonathan Berant, and Amir Globerson. 2023. Crawling The Internal Knowledge-Base of Language Models. In *Findings of the Association for Computational Linguistics: EACL 2023*. ACL, Dubrovnik, Croatia, 1856–1869.
- [7] Together Computer. 2023. *RedPajama: an Open Dataset for Training Large Language Models*. Together AI.
- [8] Silviu Cucerzan. 2007. Large-Scale Named Entity Disambiguation Based on Wikipedia Data. In *Proceedings of the 2007 Joint Conference on Empirical Methods in NLP and Computational Natural Language Learning (EMNLP-CoNLL)*. ACL, Prague, Czech Republic, 708–716.
- [9] Silviu Cucerzan. 2014. Name entities made obvious: the participation in the ERD 2014 evaluation. In *Proceedings of the First International Workshop on Entity Recognition & Disambiguation (Gold Coast, Queensland, Australia)*. ACM, New York, NY, USA, 95–100.
- [10] Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2024. QLORA: efficient finetuning of quantized LLMs. In *Proceedings of the 37th International Conference on Neural Information Processing Systems (New Orleans, LA, USA)*. NeurIPS, NY, USA, 28 pages.
- [11] 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)*. ACL, Minneapolis, Minnesota, 4171–4186.
- [12] Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, and Jonathan Larson. 2024. From Local to Global: A Graph RAG Approach to Query-Focused Summarization. arXiv:2404.16130 [cs.CL]
- [13] Jiafeng Guo, Gu Xu, Xueqi Cheng, and Hang Li. 2009. Named entity recognition in query. In *Proceedings of the 32nd International ACM SIGIR Conference (Boston, MA, USA)*. ACM, New York, NY, USA, 267–274.
- [14] Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. REALM: Retrieval-Augmented Language Model Pre-Training. In *Proceedings of the 37th ICML*. JMLR.org, Online, Article 368, 10 pages.
- [15] Xu Han, Tianyu Gao, Yankai Lin, Hao Peng, Yaoliang Yang, Chaojun Xiao, Zhiyuan Liu, Peng Li, Jie Zhou, and Maosong Sun. 2020. More Data, More Relations, More Context and More Openness: A Review and Outlook for Relation Extraction. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the ACL and the 10th International Joint Conference on NLP*. ACL, Suzhou, 745–758.
- [16] Shibo Hao, Bowen Tan, Kaiwen Tang, Bin Ni, Xiyan Shao, Hengzhe Zhang, Eric Xing, and Zhiting Hu. 2023. BertNet: Harvesting Knowledge Graphs with Arbitrary Relations from Pretrained Language Models. In *Findings of the Association for Computational Linguistics: ACL 2023*. ACL, Toronto, 5000–5015.
- [17] Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2022. Unsupervised Dense Information Retrieval with Contrastive Learning. *TMLR 2022 (2022)*.
- [18] Joel Jang, Seonghyeon Ye, Changho Lee, Sohee Yang, Joongbo Shin, Janghoon Han, Gyeonghun Kim, and Minjoon Seo. 2022. TemporalWiki: A Lifelong Benchmark for Training and Evaluating Ever-Evolving Language Models. In *Proceedings of the 2022 Conference on Empirical Methods in NLP*. ACL, Abu Dhabi, 6237–6250.
- [19] Zhen Jia, Philipp Christmann, and Gerhard Weikum. 2024. Faithful Temporal Question Answering over Heterogeneous Sources. In *Proceedings of the ACM Web Conference 2024 (Singapore)*. ACM, New York, NY, USA, 2052–2063.
- [20] Zhen Jia, Soumajit Pramanik, Rishiraj Saha Roy, and Gerhard Weikum. 2021. Complex Temporal Question Answering on Knowledge Graphs. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management (Virtual Event)*. ACM, New York, NY, USA, 792–802.
- [21] Xisen Jin, Dejiao Zhang, Henghui Zhu, Wei Xiao, Shang-Wen Li, Xiaokai Wei, Andrew Arnold, and Xiang Ren. 2022. Lifelong Pretraining: Continually Adapting Language Models to Emerging Corpora. In *Proceedings of BigScience Episode #5 – Workshop on Challenges & Perspectives in Creating LLMs*. ACL, virtual, 1–16.
- [22] Jean Kaddour, Joshua Harris, Maximilian Mozes, Herbie Bradley, Roberta Raileanu, and Robert McHardy. 2023. Challenges and Applications of Large Language Models. arXiv:2307.10169 [cs.CL]
- [23] Jungo Kasai, Keisuke Sakaguchi, Yoichi Takahashi, Ronan Le Bras, Akari Asai, Xinyan Velocity Yu, Dragomir Radev, Noah A. Smith, Yejin Choi, and Kentaro Inui. 2024. REALTIME QA: what’s the answer right now?. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*. NY, USA.
- [24] Tom Kocmi and Christian Federmann. 2023. Large Language Models Are State-of-the-Art Evaluators of Translation Quality. In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*. European Association for Machine Translation, Tampere, Finland, 193–203.
- [25] Keshav Kolluru, Vaibhav Adlakha, Samarth Aggarwal, Mausam, and Soumen Chakrabarti. 2020. OpenIE6: Iterative Grid Labeling and Coordination Analysis for Open Information Extraction. In *Proceedings of the 2020 Conference on EMNLP*. ACL, Online, 3748–3761.
- [26] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *NeurIPS*, Vol. 33. 9459–9474.
- [27] Thomas Lin, Patrick Pantel, Michael Gamon, Anitha Kannan, and Ariel Fuxman. 2012. Active objects: actions for entity-centric search. In *Proceedings of the 21st International Conference on WWW (Lyon, France)*. ACM, New York, 589–598.
- [28] Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. Lost in the Middle: How Language Models Use Long Contexts. *TACL 12 (2024)*, 157–173.
- [29] Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. G-Eval: NLG Evaluation using Gpt-4 with Better Human Alignment. In *Proceedings of the 2023 Conference on EMNLP*. ACL, Singapore, 2511–2522.
- [30] Kelvin Luu, Daniel Khashabi, Suchin Gururangan, Karishma Mandyam, and Noah A. Smith. 2022. Time Waits for No One! Analysis and Challenges of Temporal Misalignment. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. ACL, Seattle, United States, 5944–5958.
- [31] Youmi Ma, An Wang, and Naoki Okazaki. 2023. DREEM: Guiding Attention with Evidence for Improving Document-Level Relation Extraction. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*. ACL, Dubrovnik, Croatia, 1971–1983.
- [32] Mounica Maddela, Fernando Alva-Manchego, and Wei Xu. 2021. Controllable Text Simplification with Explicit Paraphrasing. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. ACL, Online, 3536–3553.
- [33] Jose L. Martinez-Rodriguez, Ivan Lopez-Arevalo, and Ana B. Rios-Alvarado. 2018. OpenIE-based approach for Knowledge Graph construction from text. *Expert Syst. Appl.* 113, C (Dec. 2018), 339–355.
- [34] 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 *AAAI-22 Technical Tracks 5*. AAAI, virtual, 5825–5833.
- [35] 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 *AAAI The Thirty-Sixth AAAI Conference on Artificial Intelligence*, Online, 5825–5833.
- [36] Arvind Neelakantan, Tao Xu, Raul Puri, Alec Radford, Jesse Michael Han, Jerry Tworek, Qiming Yuan, Nikolas Tezak, Jong Wook Kim, Chris Hallacy, Johannes Heidecke, Pranav Shyam, Boris Power, Tyna Eloundou Nekoul, Girish Sastry, Gretchen Krueger, David Schnurr, Felipe Petroski Such, Kenny Hsu, Madeleine Thompson, Tabarak Khan, Toki Sherbakov, Joanne Jang, Peter Welinder, and Lilian Weng. 2022. Text and Code Embeddings by Contrastive Pre-Training. *CoRR abs/2201.10005 (2022)*.
- [37] Qiang Ning, Hao Wu, Rujun Han, Nanyun Peng, Matt Gardner, and Dan Roth. 2020. TORQUE: A Reading Comprehension Dataset of Temporal Ordering Questions. In *Proceedings of the 2020 Conference on EMNLP*. ACL, 1158–1172.
- [38] Barlas Oguz, Xilun Chen, Vladimir Karpukhin, Stan Peshterliev, Dmytro Okhonko, Michael Schlichtkrull, Sonal Gupta, Yashar Mehdad, and Scott Yih. 2022. UniK-QA: Unified Representations of Structured and Unstructured Knowledge for Open-Domain Question Answering. In *Findings of the Association for Computational Linguistics: NAACL 2022*. ACL, Seattle, United States, 1535–1546.
- [39] OpenAI. 2023. GPT-4 Technical Report.
- [40] Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep Contextualized Word Representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. ACL, New Orleans, Louisiana, 2227–2237.
- [41] Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language Models as Knowledge Bases?. In *Proceedings of the 2019 Conference on EMNLP and the 9th International Joint Conference on NLP*. ACL, Hong Kong, China, 2463–2473.

[42] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Limits of Transfer Learning with a Unified Text-to-Text Transformer. *JML* 21, 140 (2020), 1–67.

[43] Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. In-Context Retrieval-Augmented Language Models. *TACL* 11 (2023), 1316–1331.

[44] Uma Sawant and Soumen Chakrabarti. 2013. Learning joint query interpretation and response ranking. In *Proceedings of the 22nd International Conference on World Wide Web* (Rio de Janeiro, Brazil). ACM, New York, NY, USA, 1099–1110.

[45] 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 NLP (Volume 1: Long Papers)*. ACL, Online, 6663–6676.

[46] Chenglei Si, Zhe Gan, Zhengyuan Yang, Shuohang Wang, Jianfeng Wang, Jordan L. Boyd-Graber, and Lijuan Wang. 2023. Prompting GPT-3 To Be Reliable. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, Rwanda.

[47] Sneha Singhania, Simon Razniewski, and Gerhard Weikum. 2022. Predicting Document Coverage for Relation Extraction. *TACL* 10 (03 2022), 207–223.

[48] Sneha Singhania, Simon Razniewski, and Gerhard Weikum. 2024. Recall Them All: Retrieval-Augmented Language Models for Long Object List Extraction from Long Documents. *ArXiv abs/2405.02732* (2024).

[49] Mengting Wan, Tara Safavi, Sujay Kumar Jauhar, Yujin Kim, Scott Counts, Jennifer Neville, Siddharth Suri, Chirag Shah, Ryan W White, Longqi Yang, Reid Andersen, Georg Buscher, Dhruv Joshi, and Nagu Rangan. 2024. TnT-LLM: Text Mining at Scale with Large Language Models. *arXiv:2403.12173 [cs.CL]*

[50] Zhilin Wang, Yi Dong, Olivier Delalleau, Jiaqi Zeng, Gerald Shen, Daniel Egert, Jimmy J. Zhang, Makesh Narsimhan Sreedhar, and Oleksii Kuchaiev. 2024. HelpSteer2: Open-source dataset for training top-performing reward models. *arXiv:2406.08673*

[51] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2024. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems* (New Orleans, LA, USA). NeurIPS, Red Hook, NY, USA, 14 pages.

[52] Siheng Xiong, Ali Payani, Ramana Kompella, and Faramarz Fekri. 2024. Large Language Models Can Learn Temporal Reasoning. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. ACL, Bangkok, Thailand, 10452–10470.

[53] Wenda Xu, Danqing Wang, Liangming Pan, Zhenqiao Song, Markus Freitag, William Wang, and Lei Li. 2023. INSTRUCTSCORE: Towards Explainable Text Generation Evaluation with Automatic Feedback. In *Proceedings of the 2023 Conference on Empirical Methods in NLP*. ACL, Singapore, 5967–5994.

User: Given the news snippet, extract and summarize the interactions of {subject} with other named entity (names within <e> and </e> tags) mentions. Use as few words as possible. Per line generate a single interaction between {subject} and the other named entity.

Snippet:
 {snippet} May 2, 2023: Stepping out at the <e> Met Gala </e> 2023 (her first event), <e> Doja Cat </e> instantly delivered one of the evening's most Internet-breaking moments by channeling the one and only <e> Chouette </e>, <e> Karl Lagerfeld's </e> beloved cat. "It makes no sense for anyone to go as <e> Karl's </e> cat more than <e> Doja Cat </e>," says <e> Brett Alan Nelson </e>, her creative director and stylist ...

Figure A1: Prompt for NEON(M₁) construction.

User: Given the news snippets, extract and summarize the interactions between {subject} and {object}. Use as few words as possible. Per line generate a single interaction between {subject} and {object}.

Snippets:
 [1] {snippet} May 1, 2023: <e> Doja Cat </e> and <e> Jared Leto </e> deliver <e> Lagerfeld </e> inspired looks. For this year's theme, the A-list guests were asked to wear outfits which paid homage to the German fashion designer ...
 ...
 [k] {snippet} May 2 2023: ... To be fair, <e> Leto </e> was not the sole <e> Chouette </e> on the carpet. <e> Doja Cat </e> went the Animorphs route, wearing ... cap with her <e> Oscar de la Renta </e> gown. It's almost like <e> Doja Cat </e> and <e> Jared Leto </e> went to the Spirit Halloween store together ...

Figure A2: Prompt for NEON(M₂) construction.

Score	Helpfulness		Relevance		Faithfulness		
	Mean	Median	Mean	Median	Mean	Median	
WebRAG	0	705.8	591.0	750.6	651.0	940.2	871.5
	1	801.7	731.0	911.2	884.0	905.3	877.0
	2	908.2	839.0	865.0	762.0	814.5	688.0
NEON(M ₁)	0	641.5	609.0	674.7	637.0	786.7	706.5
	1	710.7	661.0	756.2	692.0	755.1	697.0
	2	819.6	736.0	798.5	709.5	774.5	679.5
NEON(M ₂)	0	634.4	584.0	666.3	596.5	820.2	723.5
	1	713.3	648.0	763.8	700.0	735.0	650.0
	2	837.0	753.0	813.4	710.0	798.3	701.0
NewsRAG	0	919.5	827.5	968.5	887.0	1155.0	1113.0
	1	1042.6	992.5	1108.4	1088.0	1159.2	1117.5
	2	1358.0	1342.0	1347.7	1333.5	1275.0	1247.5

Table A3: Response length statistics (number of chars)

Helpfulness

Task Description You are presented with a user query and an AI assistant's response. The query is focused on a specific entity, pertains to the news domain, and is date-stamped. Your task is to evaluate the AI assistant's response for its usefulness, using a 3-point Likert scale. The criteria for rating are detailed below. **## Helpfulness Criterion Rating 2:** The response is very helpful and provides the information expected for the user query. It includes mentions of additional named entities (such as people, locations, events, etc.) beyond the primary entity in the query and aligns completely with user's intent. **Rating 1:** The response is somewhat helpful but fails to fully provide the information expected for the user's query. It can nevertheless serve to continue the conversation with the user or provides pointers to where the information can be found. **Rating 0:** The response is not helpful and provides no information for the query. **## Output format** The output should be the following JSON format: {"rating": <numerical_rating>, "reason": <short_reasoning>}, mentioning the numerical rating, as well as a short and concise reasoning for the helpfulness rating. **## Examples** {in-context examples} **## Input to be rated** User query: {question} AI assistant's response: {response}

Relevance

Task Description You are presented with a user query and an AI assistant's response. The query is focused on a specific entity, pertains to the news domain, and is date-stamped. Your task is to evaluate the AI assistant's response for its relevance, using a 3-point Likert scale. The criteria for rating are detailed below. **## Relevance Criterion Rating 2:** The response is completely relevant with accurate details and provides the information for the query date. **Rating 1:** The response contains a mix of relevant and irrelevant details. The response contains some relevant information upto the specified date, and is more or less aligned with the user's intent. **Rating 0:** The response is incorrect and provides no information for the query date. **## Output format** The output should be the following JSON format: {"rating": <numerical_rating>, "reason": <short_reasoning>}, mentioning the numerical rating, as well as a short and concise reasoning for the helpfulness rating. **## Examples** {in-context examples} **## Input to be rated** User query: {question} AI assistant's response: {response}

Faithfulness

Task Description You are presented with a user query and an AI assistant's response. The query is focused on a specific entity, pertains to the news domain, and is date-stamped. Your task is to evaluate the AI assistant's response for its reliability, using a 3-point Likert scale. The criteria for rating are detailed below. **## Faithfulness Criterion Rating 2:** The response is perfectly reliable and grounded based on the supporting passages given below. All the information from the supporting passages is used in the response to answer the user query. **Rating 1:** The response partially uses the supporting passages given below but has additional information which may be incorrect or unreliable. **Rating 0:** The response is completely unreliable and does not depend on the supporting passages. **## Output format** The output should be the following JSON format: {"rating": <numerical_rating>, "reason": <short_reasoning>}, mentioning the numerical rating, as well as a short and concise reasoning for the helpfulness rating. **## Examples** {in-context examples} **## Input to be rated** User query: {question} Supporting passages: {passages} AI assistant's response: {response}

Table A4: Prompt templates for automated evaluation using few-shot learning.