

LExTIME: A Benchmark for Temporal Ordering of Legal Events

Claire Barale¹ Leslie Barrett² Vikram Sunil Bajaj² Michael Rovatsos¹

¹School of Informatics, University of Edinburgh ²Bloomberg

{claire.barale, Michael.Rovatsos}@ed.ac.uk {lbarrett4, vbajaj11}@bloomberg.net

Abstract

Understanding temporal relationships and accurately reconstructing the event timeline is important for case law analysis, compliance monitoring, and legal summarization. However, existing benchmarks lack specialized language evaluation, leaving a gap in understanding how LLMs handle event ordering in legal contexts. We introduce LExTIME, a dataset designed to evaluate LLMs’ event ordering capabilities in legal language, consisting of 512 instances from U.S. Federal Complaints with annotated event pairs and their temporal relations¹. Our findings show that (1) LLMs are more accurate on legal event ordering than on narrative texts (up to +10.5%); (2) longer input contexts and implicit events boost accuracy, reaching 80.8% for implicit-explicit event pairs; (3) legal linguistic complexities and nested clauses remain a challenge. While performance is promising, specific features of legal texts remain a bottleneck for legal temporal event reasoning, and we propose concrete modeling directions to better address them.

1 Introduction

Understanding the temporal relationships between events is fundamental for natural language understanding. In the legal domain, temporal reasoning is important for tasks such as case law analysis, contract interpretation, and compliance monitoring. Accurately determining the timeline of events helps assess liability, obligations, and procedural validity. However, the unique linguistic characteristics of legal texts raise questions about the effectiveness of current temporal reasoning approaches.


Language models (LMs) face notable challenges with temporal reasoning due to the complex logic and ambiguities involved, and the necessity to process long contextual dependencies (Jain et al., 2023). While several benchmarks exist for temporal reasoning with large language models (LLMs),

¹<https://zenodo.org/records/17157439>

Context paragraph

On November 25, 2020, the district court **granted** the unions summary judgment motion and **denied** CompanyX. CompanyX subsequently **appealed** the matter to the United States Court of Appeals for the Sixth Circuit (no. 20-xxxx). on September 17, 2021, however, the company and the union **entered** a settlement agreement that **resolved** the appeal.

Extracted Events



LexTime instance	temporal relationship	Label
Does the event entering a settlement agreement follow companyX appeal ?	follow	Yes

Figure 1: Legal event ordering task in LExTIME: explicit and implicit events (MIDDLE) are extracted from a context paragraph (TOP) and combined into a yes/no question (BOTTOM).

none focus on specialized languages. Most current research on event extraction and ordering focuses on narratives, such as the EVENT STORYLINE CORPUS (Caselli et al., 2018), the FINE GRAINED TEMPORAL RELATIONS DATASET (Vashishtha et al., 2019), and the TRACIE dataset (Zhou et al., 2021a). As a result, LMs’ proficiency in expert types of language, such as legal language, is largely unknown. How do LMs address the unique features of legal language for event ordering?

In this work, we conduct the first study targeting temporal event ordering within legal language. Besides its implications for legal work, event ordering is a foundation for downstream legal NLP tasks, such as summarization, drafting, and question-answering. We start by assembling LExTIME, a novel dataset composed of 512 instances constructed from labor-related U.S. Federal Complaints (§2). We extract events relevant to both the case and its procedural timeline, providing a struc-

tured resource for studying temporal reasoning in legal texts. An example instance from LEXTIME is shown in Fig. 1.

We show that legal language used in LEXTIME exhibits specific lexical, syntactic, and discourse characteristics, making it a distinct genre (§3). Since event temporal ordering has been extensively studied in narrative contexts, we use TRACIE as a baseline for comparison due to its similarity in task design and focus on event ordering.

Next, we conduct a systematic evaluation of LMs’ ability to perform event ordering, defined as identifying temporal relations between event pairs (§4, 5). Our analysis focuses on key aspects relevant to the legal domain, such as the impact of context length, the distinction between explicit and implicit events, and a comparative study of narrative versus legal language. Following Zhou et al. (2021b), we define an implicit event as one whose trigger is not explicitly stated in the text but can be inferred from the surrounding context, such as *negotiated* in Fig. 1. A key indicator of the implicit event is that it can be fully entailed from and properly implied by the following explicitly stated events.

Our experiments assess Chain-of-Thought (CoT) prompting (Wei et al., 2022), a method previously successful in reasoning over complex and long inputs. However, we find that CoT prompting is ineffective for event ordering, leading to a 3.4% accuracy drop for LEXTIME compared to standard few-shot prompting, which yields the best accuracy.

Overall, we have three core findings: (1) LLMs show improved performance on legal event ordering in LEXTIME compared to narrative, achieving enhancements of up to +10.5% (GPT-4o, few-shot). This suggests that legal texts offer clearer temporal structures for models to use. (2) LLMs perform better with longer input contexts (+2.3% for GPT-4 Turbo, few-shot) and with event pairs that include implicit event types. The highest accuracy of 80.8% is achieved when queries contain a pair with one explicit and one implicit event (GPT-4 Turbo, few-shot), surpassing pairs of two explicit events by 2.6% and the overall dataset average by 3.2% (§5). (3) However, linguistic complexities commonly found in legal language contribute to model errors (§6). Features typical of the legal genre are overrepresented in misclassified instances, such as paraphrases (+89% in error samples) and events occurring in subordinate clauses (+84%).

In sum, we make the following contributions:

1. We introduce LEXTIME, the first dataset curated for legal event ordering, enabling the systematic study of temporal reasoning in legal texts.
2. We use LEXTIME to analyze the abilities of large LLMs on legal event ordering, a fundamental task in legal NLP. Our findings reveal that while LLMs achieve higher accuracy on legal texts than narrative texts, their performance is still constrained by linguistic complexities found in legal language.
3. We identify effective strategies for improving model accuracy, showing that longer context windows enhance accuracy for large models. In comparison, shorter context windows benefit smaller models, providing practical insights for optimizing model configurations in legal NLP tasks.

2 Datasets

LEXTIME is a dataset focused on temporal reasoning in legal language, derived from U.S. federal complaints between 2020 and 2024. We randomly sample complaints categorized under the Nature of Suit (NOS) codes beginning with 7, which correspond to labor-related cases². For comparison, we use TRACIE (Zhou et al., 2021b), which consists of short stories and has been curated for the same task, offering a natural choice.

2.1 Task Overview and LEXTIME Dataset

LEXTIME is a dataset built to measure the ability of LMs to perform temporal event ordering on a specific type of language: legal language. Each instance contains: (1) A context paragraph taken from a complaint document. (2) A pair of events, which can be either one implicit event and one explicit event or two explicit events. An event is an occurrence or action triggered by a verb or noun that takes place at a specific time and can be ordered relative to other events based on temporal relationships. (3) A query consisting of a pair of events and a temporal relationship between them, creating a triplet following this template: $\{event A, TR, event B\}$. (4) A binary label, *yes* or *no*. For instance, in the positive instance shown in Fig. 1, *entered a settlement agreement* and *CompanyX*

²Complaints are publicly accessible at <https://pacer.uscourts.gov/>.

appealed are two explicit events taken from the *context paragraph*, *follow* is the *TR* that links them in the *query*, and the answer (*label*) to the *query* according to the *context* is *yes*.

Pre-labeling. We used *Mistral-Large-Instruct-2407* to help construct instances and then manually reviewed all generated instances (prompts in Appendix C.2). We employed a prompt-chaining approach to minimize errors that could occur during the multiple steps involved (Wu et al., 2022).

We kept paragraphs with at least two *-ed* tokens (event triggers). We merged every three paragraphs for more context, resulting in an average *context paragraph* length of 172 tokens. Next, we extracted explicitly mentioned events using an initial prompt. We manually reviewed these to select the most significant case or procedural events. Intermediate cases with ambiguous temporal relations were resolved during this annotation stage through manual review. We then identified five implicit events per paragraph and generated sentences linking events with temporal relationships, indicating whether one event occurred *before*, *at the same time*, or *after* another. All possible TR variations are provided in the last column of Fig. 2. Each TR was reversed to create a contradictory query (e.g., *happens after* to *happens before*).

Fig. 2 compares the TR used in LEXTIME with Allen’s relations (Allen, 1983), which are foundational to temporal reasoning. We simplify the TR to three categories for practical and theoretical reasons: (1) Strict event order: in legal reasoning, important TRs primarily depend on precedence and causality (e.g., *The lawsuit was filed before the hearing*), and finer-grained relations like *Meets* are less relevant. (2) Ambiguous boundaries: Allen’s relations need precise interval boundaries, which are often vague. Simplifying these TRs improves interpretability and consistency and reduces annotation errors. (3) Consistency with TRACIE: we employ three key relationships to ensure alignment and facilitate comparison with Zhou et al. (2021b).

Label collection. We manually reviewed the 512 pre-labeled instances and corrected their labels if needed. In total, we found that 23.24% of the samples were wrongly labeled. Among those errors, 54.55% were an error on the temporal relationship, 37.87% on the event extraction, and 7.58% on both. We collect a balanced dataset of 257 instances labeled *yes* (TR is correct) and 255 instances labeled *no* (TR is incorrect).


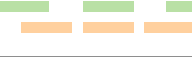

Representation	Allen's Relations	LexTime's Relations
	Before, Meets	Precedes, Happens before, Starts before, Arises before, Is followed by
	Overlaps, Finished by, Contains, Starts, Equals, Started by, During, Overlapped by	Is simultaneous, Happens simultaneously
	Finishes, Met by, After	Follows, Happens after, Starts after, Comes after

Figure 2: LEXTIME’s temporal relations and their correspondence to Allen’s interval relations (Allen, 1983).

2.2 TRACIE Dataset

We randomly selected 512 instances from the TRACIE dataset (Zhou et al., 2021b) to match the number of samples in LEXTIME. TRACIE contains short story paragraphs and implicit events. We modified the original TRACIE dataset to match the format of LEXTIME by converting instances to a query instead of the original entailment task format. We rerun all experiments as previous benchmarks using TRACIE have employed older models (Chu et al., 2024).

3 Specific Features of Legal Events

This section examines the linguistic and structural characteristics that differentiate legal language from narrative, using LEXTIME and TRACIE as representative datasets. We show that legal language constitutes a distinct genre characterized by specialized event structures and temporal representations.

Prior work has shown that legal language follows distinct linguistic conventions (Yunus et al., 2016), including lengthy nominalized expressions, frequent passive forms, and specialized vocabulary, making its syntactic structure more complex than that of general prose or news texts (Kuzman and Ljubešić, 2023). In an empirical study of the text in the U.S. Code, Martinez et al. (2024) found higher rates of complex syntactic structures relative to six baseline genres of English. Similarly, Martinez et al. (2022) noted that U.S. case law contains center-embedded clauses at double the rate of other text genres. Guided by the features identified by Yunus et al. (2016), we divide our analysis into: (1) a lexical analysis and (2) a syntax and discourse analysis (Fig. 3).

We employ a combination of methods, each selected for its suitability to a specific feature of interest: (1) named-entity recognition to extract dates; (2) a set of rules and keyword search to extract tem-




Figure 3: Comparison of linguistic features in LEXTIME and TRACIE. Values (y-axis) represent the percentage of sentences that contain each feature, indicating its frequency in the text.




Figure 4: Top-10 event triggers in LEXTIME by frequency (%), showing specialized legal vocabulary.

poral markers (Appendix A.2); (3) part-of-speech and dependency parsing to extract modal verbs, tense of the verbs, negations; (4) we further the dependency parsing with manual annotation for nominalized events, passive voice, citations, subordinate and conditional clauses.

3.1 Vocabulary

Domain-specific event triggers. An event trigger is a verb or noun that signals an event, such as *filed* in *filed a lawsuit*. In LEXTIME, event triggers are mostly legal terms (Fig. 4), and they vary significantly from those in TRACIE (Appendix A.1). None of the 10 most frequent event triggers overlap between the datasets, and we compute a low lexical overlap of 0.09 using all unique event triggers found in each dataset (Jaccard similarity). TRACIE shows a higher diversity of event triggers with a ratio of 0.148 and 701 unique triggers, compared to 0.035 and 403 unique triggers for LEXTIME (unique triggers divided by the total number of triggers).

Many terms in LEXTIME have specific meanings that differ from their general usage (semantic shift). For example, *charged* in legal terms means a formal accusation, while in everyday language, it can refer to *charging a phone*. In contrast, TRACIE includes event triggers that often describe everyday actions (e.g., *started, asked*).

Precise and explicit temporal cues. LEXTIME contains 10.28x more date mentions than TRACIE, reflecting reliance on precise and explicit temporal expressions (e.g., *on June 4, 2021*) and statutory deadlines (e.g., *within 30 days of filing*). In contrast, narrative texts use vague temporal expressions, as evidenced by fewer dates but a higher occurrence of frequency adverbs (+2.43%, Appendix A.3).

Higher density of temporal markers. LEXTIME contains a higher density of temporal markers than TRACIE (+3.99%), highlighting a structural and functional difference in the representation of events (Fig. 3).

Frequent use of modalities. LEXTIME has 11% more modal verbs than TRACIE, except for *could*, which shows a decrease of -1.96% (Appendix A.4). The use of modalities in legal writing indicates binding obligations (*must, shall*), permissions (*may, can*), and advisory language (*should*). In contrast, narratives employ modality to express uncertainty (*could*).

Frequent use of negation. Negation appears 4.73% more frequently in LEXTIME than in TRACIE. This reflects its structural and functional role in legal reasoning (restrictions, prohibitions, and exceptions). LEXTIME includes double negatives, characteristic of legal discourse. For example: *A decision shall not be invalid solely because of a procedural error*.

3.2 Syntax, Discourse, and Tenses

Events in subordinate or conditional clauses. Subordinate and conditional clause events are more frequent in LEXTIME than in TRACIE (+10.5% and +9.56%). This reflects complex hierarchical dependencies, where an event’s occurrence can depend on another event or condition.

Sentence length and event density. LEXTIME has longer sentences than TRACIE (Fig. 5), averaging 34.72 vs. 8.86 tokens, resulting in higher event density (4.01 vs. 1.60 events per sentence).




Figure 5: Boxplot showing sentence length distributions (in tokens). Median (*Med*) and quartiles (*Q1*, *Q3*) are noted, outliers are excluded for clarity.

Nominalization, passive voice, and paraphrasing. Legal writing often uses nouns instead of verbs to express events, for instance: *The revocation of the contract was issued yesterday*. Nominalized events appear in 52.50% of sentences (+50.4% compared to TRACIE). Legal texts prioritize nominalization to emphasize outcomes and obligations. They also use the passive voice (+24.76%) and paraphrase reported events (+6.56%) more frequently, enhancing their impersonal and authoritative tone. Passives create implicit event structures by omitting the agent. For example, in *The revocation of the contract*, the actor is unstated, making it difficult to identify who acted and when.

Citations. Legal texts frequently include citations (23.68% of sentences), anchoring arguments through references to past court decisions and statutes, often encoding past events.

Broader range of verb tenses and speculative events. In TRACIE, most events appear in the simple past (87.62%), while LEXTIME shows more tense variation, with 24.24% in the present simple and 5.57% in the present perfect (Appendix A.5). Additionally, speculative events are more common in LEXTIME (9.83%, +4.1%). Tense variation affects temporal reasoning, as some tenses signal rules and precedents. For instance, the present perfect marks past events with ongoing relevance (e.g., *The plaintiff has submitted an appeal*).

4 Benchmark Methodology

We conduct evaluations using two different prompt-based approaches: standard prompting and chain-of-thought (CoT) prompting. We perform experiments in zero-shot (ZS), one-shot (1S), and few-

shot (FS) settings (3 examples). Prompts and examples can be found in Appendix D.

Standard prompting. In a ZS setting, the model is prompted to answer a query containing the triplet $\{event A, TR, event B\}$ based on a context paragraph. In 1S and FS settings, models are provided with one and three examples (context paragraphs and question-answer pairs), respectively.

CoT prompting. Instructions for CoT prompts are the same as those for standard prompting. CoT prompts guide the LM to reason through steps before giving a final answer. We do not use CoT in a ZS setting, as we consider examples necessary. We manually annotate up to three examples for step-by-step reasoning in 1S and FS settings.

Evaluation. We evaluate our benchmark using accuracy. The evaluation compares long and short inputs to reflect a key characteristic of legal language (sentence and context length), which captures finer linguistic complexities such as conditional clauses, subordination, and event density. There are 230 paragraphs with fewer than 150 tokens (short) and 282 paragraphs with more than 150 tokens (long). Inspired by TRACIE and based on our linguistic analysis, we address the importance of implicit event reasoning in legal NLP. We evaluate the influence of implicit events considering 210 explicit-implicit pairs, compared to 288 explicit-explicit pairs (and 14 implicit-implicit pairs).

Models and experimental setup. We evaluate a range of popular LMs, including both open-source and proprietary models (Appendix B), spanning from smaller models (LLaMA 3.2_{1B}, FLAN-T5_{780M}) to larger ones (GPT-4). Proprietary models are accessed through their APIs, non-proprietary models are accessed through HuggingFace and used in lower precision formats (16-bit floating-point and 4-bit quantization).

5 Experimental Results

GPT-4o, GPT-4 Turbo outperform Mistral, LLaMA, and FLAN models across all evaluated prompts and settings.

Experimental results show that LLMs perform better on LEXTIME than on TRACIE, suggesting their ability to tackle some linguistic complexities of legal language. However, §6 will show that these complexities still limit performance.

Setup	All			Long Context			Short Context			Explicit-Explicit			Explicit-Implicit		
	ZS	1S	FS	ZS	1S	FS	ZS	1S	FS	ZS	1S	FS	ZS	1S	FS
GPT-4o	69.9	74.9	77.4	70.3	74.6	77.4	71.6	71.2	77.0	68.8	74.8	76.8	71.5	76.9	79.0
+CoT	-	68.6	72.6	-	69.0	72.6	-	68.7	73.2	-	69.3	72.7	-	69.2	74.0
GPT-4 Turbo	70.2	75.1	77.2	70.1	76.6	79.9	70.1	73.5	76.4	71.0	76.6	78.2	71.1	74.3	80.8
+CoT	-	77.6	74.3	-	77.7	71.3	-	70.7	75.2	-	76.0	73.1	-	78.7	75.0
Mistral _{123B}	61.9	70.1	73.9	62.2	69.6	70.8	63.3	71.6	77.8	63.5	71.5	74.7	62.6	69.6	73.6
+CoT	-	67.8	72.4	-	62.3	70.1	-	69.6	74.6	-	65.5	72.0	-	68.1	71.2
LLaMA 3.1 _{70B}	68.9	72.9	73.7	69.9	70.3	73.9	70.1	71.7	73.3	67.8	70.9	75.3	70.5	73.0	68.7
+CoT	-	66.1	69.7	-	64.3	67.3	-	69.7	72.8	-	64.4	68.5	-	68.7	70.2
LLaMA 3.1 _{8B}	51.9	60.5	55.9	51.2	59.1	51.1	51.7	59.1	58.4	57.5	56.7	53.7	52.6	62.6	57.3
+CoT	-	48.4	54.6	-	49.9	55.1	-	47.0	56.1	-	49.7	51.6	-	46.7	60.8
LLaMA 3.1 _{8B} (Base)	49.6	50.5	50.2	50.4	51.7	50.0	49.3	53.6	50.3	49.6	50.8	50.1	47.6	48.2	48.8
+CoT	-	52.3	50.8	-	49.9	53.2	-	52.6	50.9	-	53.0	51.9	-	51.7	52.4
LLaMA 3.2 _{3B}	52.2	53.5	51.0	51.7	52.3	47.2	55.9	55.1	51.0	54.7	51.2	51.3	53.5	52.6	51.7
+CoT	-	53.8	57.2	-	53.7	57.1	-	53.0	55.9	-	53.6	59.5	-	49.9	56.5
LLaMA 3.2 _{3B} (Base)	49.8	49.5	49.3	50.0	57.6	47.9	47.7	48.7	48.5	49.8	54.9	49.8	49.3	45.8	45.4
+CoT	-	48.5	47.5	-	48.4	47.5	-	51.1	50.3	-	49.4	49.3	-	47.8	48.6
LLaMA 3.2 _{1B}	49.9	47.7	48.9	48.7	48.6	51.1	51.6	48.8	51.2	49.6	48.4	48.7	49.8	46.7	52.3
+CoT	-	49.9	50.1	-	45.3	52.1	-	45.7	48.7	-	46.7	50.2	-	47.6	49.0
LLaMA 3.2 _{1B} (Base)	47.5	52.4	48.7	45.4	52.3	49.4	49.8	49.6	51.2	48.1	50.9	46.1	48.8	48.7	48.1
+CoT	-	51.0	51.0	-	50.5	38.5	-	51.1	51.7	-	50.5	50.5	-	52.2	49.3
Flan-T5 _{780M}	57.2	56.3	55.5	55.3	54.6	53.5	57.4	58.3	57.8	55.3	54.5	52.4	60.4	59.0	59.9
+CoT	-	58.2	55.9	-	56.0	54.3	-	59.6	57.8	-	55.3	54.2	-	61.4	59.0

Table 1: LEXTIME experimental results (accuracy) under zero-shot (ZS), one-shot (1S), and few-shot (FS) settings for standard and CoT prompting. All results are averaged on three runs. A gradient-based color scheme highlights higher accuracy scores, with darker shades indicating better performance.

Across the five studied subsets, GPT-4 Turbo achieves the highest accuracy on LEXTIME in four cases (up to 80.8% accuracy for explicit-implicit pairs of events), including three in an FS setting and one in a 1S CoT setting (all data). The best accuracy for the short-context subset is obtained by Mistral_{123B} in a FS setting with 77.8%. Overall, CoT appears among the top three results only three times across all categories (all data, long context, and explicit-implicit), suggesting its limited impact on performance in this task.

Compared to GPT-4o, Mistral_{123B} exhibits lower accuracy across all prompts, with performance differences ranging from 0.2% to 8%. LLaMA 3.1_{70B} also shows reduced accuracy, with a performance gap between 1% and 3.7%. Flan-T5_{780M} demonstrates a significant decline, with accuracy differences ranging from 10.4% to 21.9%. When considering the other LLaMA models collectively, their accuracy is lower by 14.4% to 28.7%.

Compared to GPT-4 Turbo, Mistral_{123B} shows lower accuracy across all prompts, with differences ranging from 1.9% to 9.8%. LLaMA 3.1_{70B} exhibits a decline in accuracy ranging from 1.3% to 11.5%, while Flan-T5_{780M} shows a more significant drop, with differences between 13% and 21.7%. The other LLaMA models collectively show accuracy reductions ranging from 14.6% to 29.2%.

5.1 Chain-of-Thought in Event Ordering




Figure 6: Performance gap (accuracy) with and without CoT prompting. Results are averaged from GPT-4, GPT-4 Turbo, LLaMA 3.1_{70B} and Mistral_{123B}.

CoT prompting does not improve accuracy for event ordering.

Prior research has demonstrated that CoT prompting can improve a model’s reasoning skills (Wei et al., 2022). As illustrated in Fig. 6, CoT prompting does not improve performance for LEXTIME or TRACIE. The best results are achieved using an FS approach with a standard prompt (75.6%). Employing CoT in FS reduces performance by 3.4%. In a 1S setting, we also observe a decrease in accuracy of 3.2%. Similar results are seen with the TRACIE dataset. Table 1 shows that these findings hold true regardless of context length and implicit events. This indicates that incorporating elicited reasoning steps did not improve the model’s ability to retrieve temporal relationships, possibly because of the increased prompt length when including mul-

multiple examples. Longer prompts may lead to attention dilution — as the prompt length increases, attention is distributed across more tokens, making it harder to distinguish key reasoning steps from less relevant information, as discussed in Shi et al. (2023).

5.2 Impact of Context Length

Larger LMs benefit from longer contexts, smaller LMs perform better with shorter context windows.

GPT-4 Turbo achieves the highest accuracy in long-context settings (79.9%), while Mistral_{123B} performs best in short-context settings (77.8%), both with a FS prompting setup.

For GPT-4o, GPT-4 Turbo, and LLaMA 3.1_{70B}, increasing the input context length leads to improved performance, with GPT-4 Turbo achieving up to a 3.5% gain on a FS setting. In contrast, all other evaluated models exhibit higher accuracy with shorter context lengths (below 150 tokens), with Mistral_{123B} demonstrating a 7% improvement and LLaMA 3.1_{8B} achieving up to a 7.3% increase. This observation suggests that larger models, such as GPT-4 Turbo, leverage the additional contextual information to capture temporal structures and event relationships, ultimately leading to more accurate predictions. In settings other than FS, we observe that shorter context lengths also enhance performance. Specifically, GPT-4 Turbo achieves a 3.9% improvement with shorter context windows in the CoT few-shot setting, and LLaMA 3.1_{70B} demonstrates a performance gain of 5.5%.

5.3 Impact of Implicit Events

GPT-4 Turbo achieves the highest overall accuracy at 80.8% for explicit-implicit pairs, outperforming the best result on the full dataset by +3.2%.

GPT-4 Turbo achieves the highest accuracy across both subsets (78.8% and 80.8%, FS). Several hypotheses may explain why queries containing one explicit and one implicit event achieve higher scores. One possible explanation is due to the dataset structure: implicit events often appear near explicit ones, which may make the temporal relation between the two events more constrained and easier for the model to infer. By contrast, two explicit events can occur in distant parts of the text or be separated by complex subordinate clauses and repeated paraphrases.

	ZS	1S	FS
GPT-4o	67.0 (↓ 2.9)	67.4 (↓ 7.5)	66.9 (↓ 10.5)
+CoT	-	61.8 (↓ 6.8)	64.0 (↓ 8.6)
GPT-4 Turbo	70.8 (↑ 0.6)	66.5 (↓ 8.6)	73.1 (↓ 4.1)
+CoT	-	73.6 (↓ 4.0)	65.9 (↓ 8.4)
Mistral _{123B}	63.9 (↑ 2)	61.1 (↓ 9)	64.6 (↓ 9.3)
+CoT	-	58.4 (↓ 9.4)	72.4 (= 0.0)
LLaMA 3.1 _{70B}	62.3 (↓ 6.6)	63.0 (↓ 9.9)	66.2 (↓ 7.5)
+CoT	-	62.0 (↓ 4.1)	65.5 (↓ 4.2)
LLaMA 3.1 _{8B}	53.5 (↑ 1.6)	54.9 (↓ 5.6)	52.7 (↓ 3.2)
+CoT	-	54.6 (↑ 6.2)	56.4 (↑ 1.8)
LLaMA 3.1 _{8B} (Base)	50.0 (↑ 0.4)	48.9 (↓ 1.6)	50.1 (↓ 0.1)
+CoT	-	49.5 (↓ 2.8)	51.8 (↑ 1.0)
LLaMA 3.2 _{3B}	53.4 (↑ 1.2)	53.5 (= 0.0)	53.5 (↑ 2.5)
+CoT	-	51.6 (↓ 2.2)	51.2 (↓ 6.0)
LLaMA 3.2 _{3B} (Base)	54.2 (↑ 4.4)	50.9 (↑ 1.4)	50.8 (↑ 1.5)
+CoT	-	50.8 (↑ 2.3)	48.2 (↑ 0.7)
LLaMA 3.2 _{1B}	50.3 (↑ 0.9)	50.2 (↑ 2.5)	51.1 (↑ 2.2)
+CoT	-	49.1 (↓ 0.8)	49.2 (↓ 0.9)
LLaMA 3.2 _{1B} (Base)	49.4 (↑ 1.9)	51.1 (↓ 1.3)	51.7 (↑ 3)
+CoT	-	50.9 (↓ 0.1)	45.6 (↓ 5.4)
Flan-T5 _{780M}	53.9 (↓ 3.3)	53.9 (↓ 2.4)	54.3 (↓ 1.2)
+CoT	-	54.6 (↓ 3.6)	54.9 (↓ 1.0)

Table 2: TRACIE experimental results (accuracy) under zero-shot (ZS), one-shot (1S), and few-shot (FS) settings for standard prompting and CoT. Top 3 results are in bold. Values in parentheses indicate the absolute difference between LEXTIME and TRACIE. ↑ signifies that accuracy is higher on TRACIE compared to LEXTIME. ↓ indicates that accuracy is lower on TRACIE.

GPT-4o, GPT-4 Turbo, and Flan-T5_{780M} yield better accuracy on explicit-implicit event pairs, improving by +2.2%, +2.6%, and +7.5%, respectively, in their FS settings. This suggests that these models are capable of inferring missing temporal information. In contrast, Mistral_{123B} and LLaMA 3.1_{70B} perform better on explicit-explicit pairs, achieving improvements of +1.1% and +6.6% in the FS settings. This suggests that these models rely on direct lexical and temporal cues.

When considering all LLaMA models except LLaMA 3.1_{70B}, base models exhibit a larger performance improvement on explicit-explicit pairs (+4.7%) compared to explicit-implicit pairs, while we observe a slight decrease with aligned models (-0.4%) — on average across all prompt types. Overall, for the aligned LLaMA models (excluding 70B), the average accuracy for explicit-implicit pairs is 52.6%, compared to 45.6% for base models. This suggests that alignment enhances implicit reasoning capabilities.

5.4 Comparison with TRACIE

The highest accuracy achieved on TRACIE at 73.6% (GPT-4 Turbo) is 4% lower than the highest accuracy on LEXTIME.

GPT-4o, LLaMA 3.1_{70B}, and FLAN-T5 con-




Figure 7: Error analysis for LEXTIME on 100 misclassified samples from the few-shot setting across GPT-4, GPT-4 Turbo and LLaMA 3.1_{70B}. The left chart shows the TR distribution in the queries, and the right plot categorizes errors by counting features found in the context paragraphs. **Takeaway:** Reasoning errors are disproportionately associated with the linguistic specificities of legal language identified in §3.

sistently perform better on LEXTIME across all prompt settings, with improvements of up to 10.5% (GPT-4o, FS). GPT-4 Turbo and Mistral_{123B} generally perform better on LEXTIME, except in the ZS setting. For the smaller LLaMA models (all except 70B), LEXTIME outperforms TRACIE by up to 6% in the CoT FS setting (LLaMA 3.2_{3B}). However, in the 1S CoT, TRACIE surpasses LEXTIME by 6.2%, indicating that task performance varies depending on the prompting strategy.

The accuracy difference suggests distinct challenges between legal and narrative event ordering. One possible explanation for the lower accuracy on TRACIE may be longer contexts in LEXTIME (Fig. 5), which provides more contextual information. A second hypothesis relates to less explicit and potentially more challenging language for temporal reasoning (Fig. 3) in TRACIE.

6 Errors Analysis

Building on the linguistic analysis in §3, we manually analyze 100 erroneous predictions by GPT-4, GPT-4 Turbo, and LLaMA 3.1_{70B}. Visualization of errors is shown in Fig. 7.

TR in queries. We analyze the distribution of errors by categorizing according to the three defined temporal relationship types. Errors are evenly distributed between *After* and *Before* TRs, with fewer *Simultaneous* errors. This pattern reflects the initial dataset composition (§2), indicating no bias toward any specific temporal relation type.

Reasoning errors. We find that reasoning errors are disproportionately associated with the linguistic specificities of legal language identified earlier (§3). Compared to the overall dataset (Fig. 3), errors are associated with a +89% increase in paraphrases, +84% in subordinate clauses, +63% in events not in

the past simple tense, +24% in passive voice, +12% in negations and future speculative constructions, +6% in modality markers, and +2% in conditional clauses. This suggests that these linguistic phenomena contribute to the difficulty of event ordering in legal texts.

We analyze the characteristics of context paragraphs that led to errors in event ordering. Several features appear in more than half of the analyzed error cases: subordinate clauses, non-past events, paraphrases, and passive constructions.

(1) Subordinate clauses embed events in hierarchical structures, complicating dependency relations between events. For example:

The applicant Mrs. X learned of the misuse of union funds only in March 2021, when Mr. Y, the vice president at the time, obtained copies of her weekly expense reports from 2015.

Here, the event queried (*obtained copies of her weekly expense reports*) occurs within a subordinate clause introduced by *when*, rather than in the main clause. If we wanted to order the two main events in this sentence (event A as *Mrs X learned of the misuse* and Event B as *Mr Y obtains copies of expense reports*), we note that event B in the subordinate clause happens before event A, but occurs later in the sentence. This creates a discrepancy between the order of the words and the actual order of the events.

(2) Events not expressed in the past simple tense are +63% more frequent in LEXTIME than in TRACIE. Tense variation increases temporal complexity, compared to texts where most events occur in the same tense. This is particularly the case when verb tense signals specific legal functions, such as rules, obligations, or precedents – patterns that are less

common in generic narratives. Consequently, we hypothesize that models trained primarily on non-legal corpora struggle to interpret these temporal cues.

(3) Legal texts frequently employ paraphrases to increase clarity or restate obligations, which can lead to multiple mentions of the same event.

(4) Passives often obscure the agent or temporal anchor of an event, making it harder for models to determine when an action occurred and who performed it.

While §3 highlighted nominalizations as a key mechanism for presenting events in legal texts, LEXTIME does not include nouns as event triggers. Nevertheless, noun events appear in the context paragraphs, which explains their presence in the error analysis Fig. 7.

Overall, the error patterns suggest that the linguistic phenomena most characteristic of legal discourse are precisely those that hinder event ordering accuracy with language models. This reinforces the claim that improving legal temporal reasoning requires domain-sensitive modeling.

7 Related Work

Benchmarking LLMs. Previous datasets and benchmarks (Hosokawa et al., 2023; Chen et al., 2021) study the performance of LLMs on temporal reasoning, showing that those tasks pose a challenge and that the limitations of LLMs in temporal reasoning are not well understood. They show that LMs generally struggle with temporal reasoning, as evidenced by Feng et al. (2023), who demonstrate that LLMs often take random guesses. Jain et al. (2023) show that LLMs perform better on event duration tasks than on event ordering tasks and that reasoning over longer contexts is more difficult, two aspects we study in LEXTIME.

Qiu et al. (2024) indicate that LLMs lag behind human performance and small-scale but specialized LMs. The performance gap is only marginally reduced by introducing in-context learning, fine-tuning, and CoT prompting. Their analysis shows that current LLMs struggle with having a consistent temporal model due to insufficient exposure during pretraining. We investigate the effectiveness of CoT prompting within the legal domain.

Temporal reasoning and understanding. Grounding events in a temporal context has a notable history in NLP. Pustejovsky et al. (2005) demonstrated that temporal grounding affects

question-answering model quality and defined an event as a situation that occurs or holds true at a specific time point or interval. Recent temporal reasoning benchmarks advanced our understanding of LMs’ capabilities (Chen et al., 2021; Hosokawa et al., 2023). However, existing datasets primarily focus on general or narrative texts (Zhou et al., 2021a). Temporal reasoning consists of subtasks such as event frequency, duration, or ordering (Zhou et al., 2019; Jain et al., 2023). Event ordering is particularly important for constructing timelines in the legal domain. Despite its significance, no prior work has systematically evaluated LMs’ ability to perform temporal event ordering in legal texts.

Existing datasets include several approaches: natural language inference and relation extraction (Thukral et al., 2021; Hosokawa et al., 2023; Zhou et al., 2021a), question-answering (QA), such as and MC-TACO (Zhou et al., 2019), TIMEDIAL for dialogue settings (Qin et al., 2021), TEMPREASON (Tan et al., 2023), and TIMEQA (Chen et al., 2021). Our work is inspired by those methods, framing the task as a simplified QA format with binary labels.

8 Conclusion

We present LEXTIME, the first benchmark for evaluating temporal event ordering in legal language, and systematically evaluate LLMs’ ability to reason over legal temporal structures. Our findings show that while LLM performance on LEXTIME is superior to their performance on a narrative text benchmark, the complexities of legal discourse, such as paraphrasing, constrain accuracy. We show that longer input contexts and implicit-explicit event pairs improve model accuracy, offering practical insights for optimizing LLM-based legal NLP applications. Experimental results suggest that further performance improvements could rely on exploiting identified errors and linguistic structure to align models with specialized domain language.

9 Limitations

Our study has several limitations. First, we do not provide a comprehensive evaluation of ambiguous cases, as these are classified and resolved during the annotation stage. Second, the temporal relations scheme is limited to three categories, whereas a more granular approach would be necessary to handle more scenarios and would capture richer temporal distinctions.

The generalizability of our findings is constrained by the specificity of the area of the law and by the size of LEXTIME (512 samples). While our benchmark represents an initial step in evaluating temporal reasoning in legal texts, it primarily focuses on event ordering. Future research should extend this evaluation to include other temporal reasoning tasks, such as event frequency and duration analysis, to achieve a more comprehensive understanding.

Acknowledgments

The work was completed while Claire Barale was a Bloomberg Data Science Ph.D. Fellow. Experiments have made use of the resources provided by the Edinburgh Compute and Data Facility (ECDF) and the Edinburgh International Data Facility (EIDF).

References

- James F Allen. 1983. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843.
- Tommaso Caselli, Ben Miller, Marieke van Erp, Piek Vossen, Martha Palmer, Eduard Hovy, Teruko Mitamura, David Caswell, Susan W Brown, and Claire Bonial. 2018. Events and stories in the news: Proceedings of the workshop. In *Events and Stories in the News*. Association for Computational Linguistics (ACL).
- Wenhu Chen, Xinyi Wang, William Yang Wang, and William Yang Wang. 2021. [A dataset for answering time-sensitive questions](#). In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1.
- Zheng Chu, Jingchang Chen, Qianglong Chen, Weijiang Yu, Haotian Wang, Ming Liu, and Bing Qin. 2024. [TimeBench: A comprehensive evaluation of temporal reasoning abilities in large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1204–1228, Bangkok, Thailand. Association for Computational Linguistics.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2024. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Yu Feng, Ben Zhou, Haoyu Wang, Helen Jin, and Dan Roth. 2023. [Generic temporal reasoning with differential analysis and explanation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12013–12029, Toronto, Canada. Association for Computational Linguistics.
- Taishi Hosokawa, Adam Jatowt, and Kazunari Sugiyama. 2023. [Temporal natural language inference: Evidence-based evaluation of temporal text validity](#). In *Advances in Information Retrieval: 45th European Conference on Information Retrieval, ECIR 2023, Dublin, Ireland, April 2–6, 2023, Proceedings, Part I*, page 441–458, Berlin, Heidelberg. Springer-Verlag.
- Raghav Jain, Daivik Sojitra, Arkadeep Acharya, Sriparna Saha, Adam Jatowt, and Sandipan Dandapat. 2023. [Do language models have a common sense regarding time? revisiting temporal commonsense reasoning in the era of large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6750–6774, Singapore. Association for Computational Linguistics.
- Taja Kuzman and Nikola Ljubešić. 2023. Automatic genre identification: a survey. *Language Resources and Evaluation*, pages 1–34.
- Eric Martinez, Francis Mollica, and Edward Gibson. 2022. Poor writing, not specialized concepts, drives processing difficulty in legal language. *Cognition*, 224.
- Eric Martinez, Francis Mollica, and Edward Gibson. 2024. Even lay people use legalese. *PNAS*, 121.
- OpenAI. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- James Pustejovsky, Robert Knippen, Jessica Littman, and Roser Saurí. 2005. Temporal and event information in natural language text. *Language resources and evaluation*, 39:123–164.
- Lianhui Qin, Aditya Gupta, Shyam Upadhyay, Luheng He, Yejin Choi, and Manaal Faruqui. 2021. [TIME-DIAL: Temporal commonsense reasoning in dialog](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7066–7076, Online. Association for Computational Linguistics.
- Yifu Qiu, Zheng Zhao, Yftah Ziser, Anna Korhonen, Edoardo Ponti, and Shay Cohen. 2024. [Are large language model temporally grounded?](#) In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7064–7083, Mexico City, Mexico. Association for Computational Linguistics.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.

Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed H. Chi, Nathanael Schärli, and Denny Zhou. 2023. **Large language models can be easily distracted by irrelevant context**. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 31210–31227. PMLR.

Qingyu Tan, Hwee Tou Ng, and Lidong Bing. 2023. **Towards benchmarking and improving the temporal reasoning capability of large language models**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14820–14835, Toronto, Canada. Association for Computational Linguistics.

Shivin Thukral, Kunal Kukreja, and Christian Kavouras. 2021. **Probing language models for understanding of temporal expressions**. In *Proceedings of the Fourth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 396–406, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Siddharth Vashishtha, Benjamin Van Durme, and Aaron Steven White. 2019. **Fine-grained temporal relation extraction**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2906–2919, Florence, Italy. Association for Computational Linguistics.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*, Red Hook, NY, USA. Curran Associates Inc.

Tongshuang Wu, Ellen Jiang, Aaron Donsbach, Jeff Gray, Alejandra Molina, Michael Terry, and Carrie J Cai. 2022. **Promptchainer: Chaining large language model prompts through visual programming**. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems, CHI EA '22*, New York, NY, USA. Association for Computing Machinery.

Kamariah Yunus, Radzuwan Ab Rashid, et al. 2016. Colligations of prepositions: Essential properties of legal phraseology. *International Journal of Applied Linguistics and English Literature*, 5(6):199–208.

Ben Zhou, Daniel Khashabi, Qiang Ning, and Dan Roth. 2019. **“going on a vacation” takes longer than “going for a walk”**: A study of temporal commonsense understanding. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference*

on Natural Language Processing (EMNLP-IJCNLP), pages 3363–3369, Hong Kong, China. Association for Computational Linguistics.

Ben Zhou, Kyle Richardson, Qiang Ning, Tushar Khot, Ashish Sabharwal, and Dan Roth. 2021a. **Temporal reasoning on implicit events from distant supervision**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1361–1371, Online. Association for Computational Linguistics.

Ben Zhou, Kyle Richardson, Qiang Ning, Tushar Khot, Ashish Sabharwal, and Dan Roth. 2021b. **Temporal reasoning on implicit events from distant supervision**. *Preprint*, arXiv:2010.12753.

Appendix

A Specific Features of Legal Events

A.1 TRACIE Event Triggers




Figure 8: Top-10 event triggers in TRACIE by frequency (%).

A.2 Keywords for Temporal Markers Extraction

- **General time:** *now, then, before, after, later, soon, earlier, previously, recently, already, immediately, eventually, finally, formerly, subsequently, yet, ever, until, till, since, while, when, whenever, as soon as, as long as, afterwards, meanwhile, in the meantime, up to now, up until now.*
- **Frequency Adverbs:** *always, usually, often, sometimes, occasionally, rarely, seldom, never, frequently, constantly, continuously, periodically, infrequently, repeatedly*

Adverbs, conjunctions or prepositions that don't appear in Fig. 9 were not found in the datasets.

A.3 Frequency of Temporal Cues

See Fig. 9.

A.4 Frequency of Modal Verbs

See Fig. 10.




Figure 9: Comparison of the frequency of temporal adverbs in LEXTIME (blue) and TRACIE (green). Values (y-axis) represent the percentage of sentences that contain each adverb, indicating its frequency in the text.




Figure 10: Comparison of the frequency of the modal verbs in LEXTIME (blue) and TRACIE (green). Values (y-axis) represent the percentage of sentences that contain each modal, indicating its frequency in the text.

A.5 Frequency of Verb Tenses

Fig. 11 shows the frequency of the tenses found in the context paragraph, for both datasets.




Figure 11: Distribution of the tenses of the verbs in LEXTIME and TRACIE. The y-axis is the frequency, in %. The y-axis scales differ between the charts for readability.

B Models

B.1 Impact of Model Size




Figure 12: Scaling effect of model size and overall temporal reasoning performance. The y-axis shows the average accuracy across all prompt types for each model. The x-axis (model size) is shown in the log scale. **Takeaway:** Results show a log-linearity between parameter size and performance.

B.2 Models Used

GPT-4o and GPT-4-turbo. GPT-4-turbo refers to *gpt-4-turbo-2024-04-09*, and GPT-4o refers to *gpt-4o-2024-08-06*. GPT-4o provides faster inference compared to turbo (OpenAI, 2023).

Llama 3.1 and Llama 3.2. We use the HuggingFace implementation of *Llama-3.1-8B*, *Llama-3.1-8B-Instruct*, *Llama-3.1-70B*, *Llama-3.1-70B-Instruct*, *Llama-3.2-1B*, *Llama-3.2-1B-Instruct*, *Llama-3.2-3B*, *Llama-3.2-3B-Instruct*. Llama 3.1 and 3.2 fine-tuned versions use supervised fine-tuning (SFT) and reinforcement learning from hu-

man feedback (RLHF), Rejection Sampling (RS), and Direct Preference Optimization (DPO, (Dubey et al., 2024)).

Mistral. We use the large version, 123B parameters, *Mistral-Large-Instruct-2407*.

FLAN-T5-large. We use the large version of FLAN-T5 (Chung et al., 2024). FLAN-T5 is built on top of T5 (Raffel et al., 2020) through instruction fine-tuning.

B.3 GPU Resources

This work was made possible by the use of: Nvidia Tesla V100-SXM2, Nvidia Tesla K80, Nvidia A100 40GB and 80GB, and Nvidia RTX A6000 40GB.

C Dataset Construction

C.1 Examples from LEXTIME

See Fig. 13.

C.2 Prompts for Dataset Generation

See Table 3.

D Prompts for Evaluation

See Table 4 for the evaluation prompts used for various settings, and Tables 5 and 6 for examples from LEXTIME and TRACIE, respectively.

Context	Query (Event A, TR, Event B)	Answer (Label)
Rather than pursue the grievance through the contracts grievance and arbitration procedure, the union directed hospital employees to refuse to pick up extra shifts beyond their normal scheduled work [...]	Does the hospital being forced to hire traveller nurses and pay an extraordinary premium precede the union members' concerted refusal to pick up extra shifts ?	No
On November 25, 2020, the district court granted the unions summary judgment motion and denied company X . Company X subsequently appealed the matter to the United States Court of Appeals [...]	Does the company and the union experiencing a dispute precede the district court denying Company X a summary judgement motion ?	Yes
In February of 2023, a union representative engaged in physically intimidating conduct with a hospital manager. In response to the incident, the hospital limited the union representatives access to the hospital [...]	Does the union disputing the hospital's response follow the union representative engaging in physically intimidating conduct ?	Yes

Figure 13: Examples from LEXTIME, showing three instances composed of a context paragraph, a query, and a binary label. Event B (explicit) is highlighted in green in the context and the query. Event A from the query occurs later in the context, beyond the visible portion.

Description	Prompt
Prompt 1: Get paragraphs	Your task is to find if {context_paragraph} contains events. The events should be verbs. If it does not contain relevant events, just answer with "None" (no explanation). If it contains relevant events, please answer "relevant" and list them as verbs only (no explanation).
Prompt 2: Get implicit events	Your task is to find 5 implicit events in {context_paragraph}. Explicit events have already been identified in {explicit_events}. Implicit events are events that are not explicitly written in the paragraph, but that the reader understands happened. Here is an example: {example}. Return the events as verbs, in a numbered list, followed by 2 explicit events that are the closest to the start and end time of the implicit event.
Prompt 3: Construct sentences	Context Paragraph: {context_paragraph} Implicit and Explicit Events: {instance} Your task is to construct pairs of events initially extracted from the {context_paragraph}: events have already been extracted and are contained in {instance}. One instance includes one implicit event and two explicit events. Construct 3 sentences that demonstrate the temporal relationship between the events. Use phrases like "Event A precedes Event B," "Event A is simultaneous with Event B", "Event A follows Event B" or "Event A happens after Event B" to indicate the temporal order. Here is an example: {example}.

Table 3: Prompts used for generating pre-labeled instances of LEXTIME.

Description	Prompt
Zero-shot	Given the context, the task is to answer the query with "yes" or "no". context: "{context}" query: "{query}" answer:
One-shot	Given the context, the task is to answer the query with "yes" or "no". example: {example1} context: "{context}" query: "{query}" answer:
Few-shot	Given the context, the task is to answer the query with "yes" or "no". example 1: {example1} example 2: {example2} example 3: {example3} context: "{context}" query: "{query}" answer:
CoT one-shot	Given the context, the task is to answer the query with "yes" or "no". Let's reason through it step-by-step: example: {example1_cot} context: "{context}" query: "{query}" answer:
CoT few-shot	Given the context, the task is to answer the query with "yes" or "no". Let's reason through it step-by-step: example 1: {example1_cot} example 2: {example2_cot} example 3: {example3_cot} context: {context} query: {query} answer:

Table 4: Prompt templates for evaluation, on all five types of prompts used: zero-shot, one-shot, few-shot, CoT one-shot, CoT few-shot. These templates are used for evaluation on both LEXTIME and TRACIE, examples used differ and are extracted from each dataset respectively.

Description	Examples
Example 1	<p>context: "Before assigning the plaintiff to work, the defendant hired employees, including the plaintiff. The plaintiff was hired as a trackman."</p> <p>query: "Does hiring employees happen before assigning the plaintiff to work?"</p> <p>answer: "yes"</p>
Example 2	<p>context: "ms. X advised ms. Y to issue the written discipline and to speak with mr. R (as well as the other roll tunnel operators) regarding the defaced guidelines. accordingly, ms. Y, along with John Doe (the union steward) approached mr. R to present him with the disciplinary warning and to discuss the defacement of the posted guidelines. mr. R was on his clamp truck and ms. Y advised that mr. R that she needed to speak with him. without stopping, mr. R told her that he did not have time to speak with her (his own supervisor) because he was working and drove by her."</p> <p>query: "Does ms. X learning about the defaced guidelines precedes her advising ms. Y to issue the written discipline?"</p> <p>answer: "yes"</p>
Example 3	<p>context: "rather than pursue the grievance through the contracts grievance and arbitration procedure, the union directed hospital employees to refuse to pick up extra shifts beyond their normal scheduled work. specifically, prior to the dispute over employees obligation to work holidays, employees would regularly and voluntarily pick up extra shifts beyond their normal scheduled work. but once the dispute over holiday work arose, employees stopped picking up shifts beyond their normal scheduled work. after the aforementioned dispute arose, a member of the union and a hospital employee provided to a member of the hospitals administration a text message stating, in part: per union unless you had originally picked up ot or st when it was first offered we should be giving back those hours for this coming week and they are asking us not to pick up anything through the end of may and the first 2 weeks of june because they still are not allowing people to take vacation on a holiday unless they have adequate staffing. upon information and belief, the references to ot and st in the text message mean overtime and straight time, respectively. the union members concerted refusal to pick up extra shifts is a violation of the contracts and has caused the hospital great hardship in achieving adequate staffing for the hospitals operations. as a result of the union and its members conduct, the hospital was forced to hire so-called traveler nurses and to pay an extraordinary premium, in the millions of dollars, for adequate staffing compared with the cost to the hospital had the union not directed its members to engage in this concerted effort and had the union members picked up extra shifts as they had in the past."</p> <p>query: "Does the hospital being forced to hire traveler nurses and pay an extraordinary premium precedes the union members' concerted refusal to pick up extra shifts?"</p> <p>answer: "no"</p>
Example 1 CoT	<p>context: "Before assigning the plaintiff to work, the defendant hired employees, including the plaintiff. The plaintiff was hired as a trackman."</p> <p>query: "Does hiring employees happens before assigning the plaintiff to work?"</p> <p>answer:" Reasoning: 1. The defendant hired employees, including the plaintiff. 2. The plaintiff was hired as a trackman. 3. The employees including the plaintiff were all assigned to work. The correct answer is: "yes"</p>
Example 2 CoT	<p>context: "ms. X advised ms. Y to issue the written discipline and to speak with mr. R (as well as the other roll tunnel operators) regarding the defaced guidelines. accordingly, ms. Y, along with John Doe (the union steward) approached mr. R to present him with the disciplinary warning and to discuss the defacement of the posted guidelines. mr. R was on his clamp truck and ms. Y advised that mr. R that she needed to speak with him. without stopping, mr. R told her that he did not have time to speak with her (his own supervisor) because he was working and drove by her."</p> <p>query: "Does ms. X learning about the defaced guidelines precedes her advising ms. Y to issue the written discipline?"</p> <p>answer:" Reasoning: 1. Ms. X learned about the defaced guidelines. 2. She then advised Ms. Y to issue the written discipline. Since Ms. X's learning about the defaced guidelines happened before advising Ms. Y, the correct answer is "yes".</p>
Example 3 CoT	<p>context: "rather than pursue the grievance through the contracts grievance and arbitration procedure, the union directed hospital employees to refuse to pick up extra shifts beyond their normal scheduled work. specifically, prior to the dispute over employees obligation to work holidays, employees would regularly and voluntarily pick up extra shifts beyond their normal scheduled work. but once the dispute over holiday work arose, employees stopped picking up shifts beyond their normal scheduled work. after the aforementioned dispute arose, a member of the union and a hospital employee provided to a member of the hospitals administration a text message stating, in part: per union unless you had originally picked up ot or st when it was first offered we should be giving back those hours for this coming week and they are asking us not to pick up anything through the end of may and the first 2 weeks of june because they still are not allowing people to take vacation on a holiday unless they have adequate staffing[.] upon information and belief, the references to ot and st in the text message mean overtime and straight time, respectively. the union members concerted refusal to pick up extra shifts is a violation of the contracts and has caused the hospital great hardship in achieving adequate staffing for the hospitals operations. as a result of the union and its members conduct, the hospital was forced to hire so-called traveler nurses and to pay an extraordinary premium, in the millions of case 3:23-cv-00466 document 1 filed 04/13/23 page 5 of 10 6 dollars, for adequate staffing compared with the cost to the hospital had the union not directed its members to engage in this concerted effort and had the union members picked up extra shifts as they had in the past."</p> <p>query: "Does the hospital being forced to hire traveler nurses and pay an extraordinary premium precedes the union members' concerted refusal to pick up extra shifts?"</p> <p>answer:" Reasoning: 1. The union members' concerted refusal to pick up extra shifts occurred as a result of the dispute. 2. This refusal led to the hospital being forced to hire traveler nurses and pay an extraordinary premium. Since the hospital was forced to hire traveler nurses and pay a premium as a result of the union members' refusal, the correct answer is "no".</p>

Table 5: Examples used for one-shot and few-shot setting for evaluation on LEXTIME.

Description	Examples
Example 1	<p>context: "My boss pulled me aside at work. I thought he was going to tell me Happy Birthday. Instead, he told me I wasn't needed today. I sulked home, not sure whether I was getting fired or not. When I arrived home, I walked into a surprise birthday party!"</p> <p>query: "Does the event i felt worried start after my boss told me i wasn't needed?"</p> <p>answer: "yes"</p>
Example 2	<p>context: "John got a new cat and was worried how it would take to the old one. He tried everything but the old cat would always act mean. One day the old cat was stuck in a box and the young cat helped it. Ever since they became best friends. "John was so happy with the result, he got another cat.""</p> <p>query: "Does the event john brought a cat home ends before the old cat stuck in a box?"</p> <p>answer: "yes"</p>
Example 3	<p>context: "Javier has always wanted to try instant ramen. He decides that today is his day to try it. He goes to the store to buy a few packets. On his way home Javier is in a car accident. Javier never got to cook his instant ramen."</p> <p>query: "Does the event javier hasn't had instant ramen before. ends before he go to the store to buy a few packets ?"</p> <p>answer: "no"</p>
Example 1 CoT	<p>context: "My boss pulled me aside at work. I thought he was going to tell me Happy Birthday. Instead, he told me I wasn't needed today. I sulked home, not sure whether I was getting fired or not. When I arrived home, I walked into a surprise birthday party!"</p> <p>query: "Does the event i felt worried start after my boss told me i wasn't needed?"</p> <p>answer:" Reasoning: 1. The boss informed the speaker that they weren't needed today. 2. After hearing this news, the speaker felt worried about their job status. 3. The feeling of worry is a response to the boss's statement. The correct answer is: "yes"</p>
Example 2 CoT	<p>context: "John got a new cat and was worried how it would take to the old one. He tried everything but the old cat would always act mean. One day the old cat was stuck in a box and the young cat helped it. Ever since they became best friends. "John was so happy with the result, he got another cat.""</p> <p>query: "Does the event john brought a cat home ends before the old cat stuck in a box?"</p> <p>answer:" Reasoning: 1. John brought a new cat home and was concerned about how the old cat would react. 2. The old cat initially acted mean towards the new cat. 3. At some point, the old cat got stuck in a box. 4. The event of bringing the new cat home occurs first and precedes the incident of the old cat getting stuck. Since the correct answer is "yes".</p>
Example 3 CoT	<p>context: "Javier has always wanted to try instant ramen. He decides that today is his day to try it. He goes to the store to buy a few packets. On his way home Javier is in a car accident. Javier never got to cook his instant ramen."</p> <p>query: "Does the event javier hasn't had instant ramen before. ends before he go to the store to buy a few packets ?"</p> <p>answer:" Reasoning: 1. The context states that Javier has always wanted to try instant ramen, indicating he has not had it before. 2. Javier decides to try instant ramen today. 3. He goes to the store to buy instant ramen, but he has not yet cooked or eaten it. 4. Therefore, the event of not having instant ramen continues until he purchases the packets and cooks, meaning it does not end before he goes to the store. Since, the correct answer is "no".</p>

Table 6: Examples used for one-shot and few-shot setting for evaluation on TRACIE.