

TComQA: Extracting Temporal Commonsense from Text

Lekshmi R Nair^{1,†}, Arun Sankar^{1,†} and Koninika Pal^{1,*}

¹Department of Computer Science and Engineering, Indian Institute of Technology Palakkad, India

Abstract

Understanding events necessitates grasping their temporal context, which is often not explicitly stated in natural language. For example, it is not a trivial task for a machine to infer that a museum tour may last for a few hours, but can not take months. Recent studies indicate that even advanced large language models (LLMs) struggle in generating text that require reasoning with temporal commonsense due to its infrequent explicit mention in text. Therefore, automatically mining temporal commonsense for events enables the creation of robust language models. In this work, we investigate the capacity of LLMs to extract temporal commonsense from text and evaluate multiple experimental setups to assess their effectiveness. Here, we propose a temporal commonsense extraction pipeline that leverages LLMs to automatically mine temporal commonsense and use it to construct TComQA, a dataset derived from SAMSum and RealNews corpora. TComQA has been validated through crowdsourcing and achieves over 80% precision in extracting temporal commonsense. The model trained with TComQA also outperforms an LLM fine-tuned on existing dataset of temporal question answering task.

Keywords

Temporal Commonsense, Question Generation, Question Answering, Large Language Models

1. Introduction

The automatic acquisition of knowledge has gained significant attention recently for developing AI systems capable of understanding and reasoning with human-like efficiency. Such a system needs information about typical properties of objects and concepts, which constitute commonsense knowledge that is rarely mentioned explicitly in natural language text. This work specifically focuses on automatically extracting temporal commonsense, which plays a crucial role in resolving time-related ambiguities and enabling contextually appropriate responses in conversation.

The major challenge in automatic acquisition of temporal commonsense is infrequent and imprecise mention of the explicit information related to the temporal arguments. For instance, when a user tells an AI assistant “Arrange me a babysitter for the vacation” vs “Arrange me a babysitter for the evening”, the bot should understand that the user needs someone available for multiple days in the former case, while the latter only requires someone available for 4-5 hours. Additionally, events often span a range of plausible durations, making it difficult to associated with a definitive timeframe. For example, the temporal context “going on a trip” can be associated with all the following valid temporal ranges, – a few days, or a few weeks, or a few months. Such variability introduces inherent ambiguity that can lead to a biased extraction models.

Many recent projects like Quasimode [1], WebChild [2], ConceptNet [3], COMET [4], etc. focus on the acquisition and creation of commonsense knowledge bases and aim to bridge the gap between existing knowledge about objects present in dominant knowledge bases like DBpedia, Freebase, Yago, etc., and their typical properties. Conversely, other works explore extracting temporal relations and temporal arguments from the text [5, 6] but do not always capture all temporal properties for events. Some works focus on acquisition of physical properties [7] like, size, weight, etc. and explore the distribution of quantitative attributes [8] for specific events to harness commonsense knowledge. However, these efforts do not cover different types of temporal properties introduced by Zhou et al. [9] – duration, frequency, stationary, ordering, and typical time. Moreover, recent literature in extracting temporal

IR-RAG@SIGIR'25: Information Retrieval's Role in RAG Systems, July 13–17, 2025, Padua, Italy

*Corresponding author.

[†]These authors contributed equally.

✉ 112001019@smail.iitpkd.ac.in (L. R. Nair); 112001008@smail.iitpkd.ac.in (A. Sankar); kpal@iitpkd.ac.in (K. Pal)



© 2025 Copyright © 2025 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).




Figure 1: Overview of extraction of TComQA

commonsense mainly rely on crowdsourcing, where annotators manually associate temporal ranges with events. This approach inherently limits the scalability of dataset creation and introduces biases based on individual annotators’ interpretations of temporal durations.

In this work, we address this challenges by proposing an automated extraction pipeline that leverages LLMs to infer temporal commonsense for all five temporal properties, defined by Zhou et al. [9], from short text. The pipeline first generates relevant questions for a target temporal properties and then utilize LLMs to extract corresponding answers from the context. For example, given the context “*Emma will be home soon and she will let Will know*”, we would like to generate a question targeting temporal property *typical time*: “*When will Emma be home?*” and find the answer “6PM” by exploiting the contextual information about typical workday routine. The contributions of this work are as follows:

- We propose a framework to extract temporal commonsense from the text by exploiting LLMs and introduce the TComQA dataset to enhance temporal Question-Answering task.
- We define two evaluation metrics to asses the quality and relevance of generated questions for a given context and specific temporal property.
- We conduct an extensive evaluation of each components of the proposed extraction pipeline. Furthermore, we validate the quality of the automatically generated TComQA dataset through crowdsourcing, demonstrating its high precision in extracting temporal commonsense.
- The codes and dataset are available *here*.

2. Extraction of temporal commonsense from text

The proposed framework for extracting temporal commonsense from text comprises three components: 1) a question generation module, 2) a question validity checker, and 3) an answer-generation module. Figure 1 illustrates the pipeline for proposed extraction framework that is used later to generate TComQA dataset.

2.1. Temporal Question Generation

Extracting temporal commonsense from text requires accurately identifying the temporal events from the context and the associated temporal properties. Hence, the question generation module takes a short text with one of the five types of temporal commonsense reasoning, introduced by Zhou et al. [9], to generate temporally relevant questions. We fine-tune the T5 [10] architecture with the MCTACO [9] dataset to create this module. This dataset provides the type of temporal commonsense reasoning required to generate answers for a question.

For instance, given a context: “*Bartender asked for ID when Emily entered the bar and after checking, she was let in*” and temporal property: ‘*Typical Time*’, we would like to generate a question: ‘*What is Emily’s age?*’. Answering such a question requires the temporal commonsense, and hence the answer becomes the extracted temporal commonsense for the event ‘entering the bar’.

2.2. Question Validation

The previous step generates questions for any given temporal properties. But, not all are appropriate for extracting the temporal value for the given context. Consider a context, “Emma will be home soon and she will let Bob know” and target temporal property: ‘frequency’. The question generator creates a relevant question for the event ‘coming home’ matching the target temporal property: “How often she is home after she is gone?”. But answering this question does not require any commonsense about the event, making it invalid for the commonsense extraction task. We observe that there is a contextual mismatch between the given context and the question when answering it for the target temporal property does not require any commonsense reasoning.

Hence, we propose two metrics for validating the questions generated from the question generation module for a given target temporal property. Using these metrics, we first filter the valid questions generated by the question generation module and consider them for creating the TComQA dataset.

2.2.1. Validation Using Lexical Similarity.

The proposed lexical validator (V_l) measures the lexical similarity between the input context C and the generated question Q . First, we filter noun, proper noun, and verb tokens from C and Q and lemmatize them. According to V_l , a question is considered valid if it contains a temporal marker (T_m) and there is at least one overlapping token between C and Q . Temporal markers are tokens, such as “how often,” “how long,” “what time,” “which year,” “before,” “afterward,” etc. (an exhaustive list can be found in the code), obtained from the SPACY library.

$$V_l(C, Q) = \begin{cases} 1 & \text{if } |C \cap Q| \geq 1 \wedge T_m \in Q \\ 0 & \text{otherwise} \end{cases}$$

2.2.2. Validation using Semantic Similarity.

Since the lexical validator relies on exact token matching between a context and a question, it fails when the questions are generated from context using semantically similar words or phrases. Hence, we propose a semantic validator (V_s) that matches phrases between the question Q and the context C . It first extracts noun and verb phrases from C and Q , denoted as C_p and Q_p using the SPACY library. The semantic validator defines a question as valid if cosine similarity between embeddings of any phrase from C and Q satisfies the given similarity threshold and the temporal markers are present in the question. Clearly, as threshold θ tends to 1, this semantic validator becomes stricter.

$$V_s(C, Q) = \begin{cases} 1 & \text{if } \exists_{i \in C_p, j \in Q_p} \text{sim}(i, j) \geq \theta \wedge T_m \in Q \\ 0 & \text{otherwise} \end{cases}$$

2.3. Generating Answers for Temporal Questions

To extract the temporal commonsense, we finally need to find the answers for the valid questions generated by the question generation module. Here, we investigate different LLM-based answer generation models that take a question with the associated context and the target temporal property and generate the typical value for the temporal property mentioned in the question.

Encoder-based Answer Generation Models. The first approach considers BERT [11] as the underlying LLM and fine-tunes it. As we need to generate numeric values for some of the temporal properties of an event, we also explore NumBERT [12], which emphasizes numbers and represents them using a special token while learning the pre-trained model. As these are encoder-based approaches, we append the MASK tokens of output size at the end of the input while finetuning these answer generation models. Clearly, this becomes a disadvantage of these models while finding the answer for a new question, because we need to know the length of the answer to mention the same number of [MASK] tokens in the input for predicting the answer.

Decoder-based Answer Generation Model. To overcome the drawback mentioned in the previous approaches, we adopt a decoder-based generative approach to find the answer. In this work, we fine-tuned Llama2-based [13] question-answering model for generating answers with dynamic length.

3. Evaluation

3.1. Datasets

We use the MCTACO dataset [9] to fine-tune the underlying LLMs for creating both question-generation and answer-generation modules, as this dataset includes context, temporal questions, plausible answers, and the temporal properties needed to derive the answer from the context. It contains 3K short contexts with an average of 17.8 tokens per context.

To create TComQA Dataset, we consider the SAMsum dataset [14], a chat summarization dataset of 14K contexts with an average of 20.6 tokens per context, and the RealNews dataset [15] of 5K contexts with an average of 31.22 tokens per context.

3.2. Experimental setup

3.2.1. Training setup for Question Generation Module

We append a context with the temporal property that will be used for finding the answer for the associated question. It is then passed to the t5-large pre-trained model to train our question generation module, and the associated question is considered as output. We set `max_seq_length`, batch size, and epoch as 128, 8, and 10 respectively, to fine tune the t5-large pre-trained model to train our question-generation module.

For training the encoder-based answer-generation models, we fine-tune the Bert_base_uncased and NumBERT using Adam optimizer with learning rate, batch size, and epochs set to 0.0001, 8, and 7 respectively. For fine-tuning NumBERT, we converted all numbers and numerical sequences according to the NumBERT directory. For the decoder-based model, we use Llama2 with 7 billion parameters. We fine-tune it using `paged_adamw_32bit` with a batch size of 4 and a learning rate of $2e-4$.

Training setup for Question Answering Module For the encoder-based approach, we fine-tune Bert_base_uncased and NumBERT models using Adam optimizer with learning rate, batch size, and epochs set to 0.0001, 8, and 7 respectively. For fine-tuning NumBERT, we converted all numbers and numerical sequences according to the NumBERT directory.

For the decoder-based model, we use Llama2 with 7 billion parameters. We trained it using `paged_adamw_32bit` with a batch size of 4 and a learning rate of $2e-4$. Here we pass context, questions, temporal property to the model with associated answer to be generated. For example, the training data is formatted as `< s > [INST] context, question, temporal property [INST] answer < /s >`. Now, Llama2-based model generates comparatively long answers, whereas answers in our training data are comparatively short. For capturing that characteristics, we omit the generated sequence of tokens after `< /s >` while retrieving the answers from the model.

3.3. Results and Discussion

Evaluation of Question Generation Model. We consider 95 context-question pairs from MCTACO as test dataset and evaluate generated questions using semantic similarity. We achieved a moderate similarity score of 0.65 by considering the cosine similarity between the embedding of the generated questions and the groundtruth questions. Here, embedding of a question is calculated by averaging the token embedding from the SPACY library. The inherent variability of language leads to multiple valid questions for a given context, which is reflected in the relatively low similarity score. For example, given the context “Our emphasis has been on having our potential clients know about us and deliver services to them, Dudovitz said. ”, the model generates the question – “What did Dudovitz do after

Table 1
Evaluation of answer generation module

Prop.	BERT		NumBERT		Llama2	
	DE	SS	DE	SS	DE	SS
Duration	56.7%	0.66	36.67%	0.42	90%	0.37
Typical Time	60%	0.39	56.67%	0.38	86.67%	0.44
Frequency	56.7%	0.63	60%	0.39	86.67%	0.56
Stationary	40.0%	0.42	13%	0.40	83.34%	0.38
Event order	16.67%	0.55	0%	0.38	76.67%	0.53
Total/Avg	46.02%	0.53	33.26%	0.39	84.67%	0.47

he spoke?” for the temporal property ‘Event Order’. Even though the generated question is valid and correctly align with the target temporal property, it is not semantically similar with the groupnth question for the context – “Do the clients learn about the company before becoming clients?”. As a result, we reach a lower semantic similarity score for this task. Hence, we also manually evaluated the generated questions and found all the them were valid, matching the target temporal property for the input context. This happens as all these test dataset consider correct temporal property for the question generation task.

Evaluation of Question Validation Module For the MCTACO test samples, the proposed lexical validator V_l confirms that 95% of generated questions are valid. On the other hand, the semantic validator V_s confirms the validity of 71% questions with $\theta = 0.5$, 49% with $\theta = 0.6$, and 37% with $\theta = 0.7$. As expected, higher θ rejects more questions. While comparing the valid questions with the manually evaluated valid questions, we found that the lexical validator is more aligned with manual validation. For the SAMsum dataset, V_l accepts 85% of the generated questions and V_s confirms the validity of 61% questions with $\theta = 0.5$.

In our extraction pipeline, we generate questions from a context for all five target temporal properties. So it is important to investigates further how well these two validators align with the human assessments. For this purpose, we sample 200 contexts from the SAMsum dataset and present the conext and the corresponding generated questions to assessors. Each context-question pair was evaluated by 5 judges, and they marked the question valid, invalid, or uncertain for the given context. We use majority votes to generate the groundtruth label for the generated question for a context. Based on the crowd-generated groundtruth, the V_l achieves 89.7% precision and 85.6% recall and V_s achieves 91.9% precision with a lower recall of 69.5%, as expected.

Evaluation of Question Answering Model. Table 1 presents the evaluation of all three answer-generation models tested on 150 question-answer pairs, 30 from each temporal types, from the MCTACO test dataset. in this experimental setup, generated answer can have different surface forms than the groundtruth answer for a question. For instance, ‘each year’ is similar to ‘every year’ or ‘rarely’ is semantically similar to ‘once’ for temporal property *frequency*. Furthermore, there are multiple valid answers for the temporal property ‘typical time’, e.g., the question, “what time of the day was the storm”, can have multiple valid answers, such as 3 PM and 4 PM. Hence, we consider evaluating the answer generation module using the semantic similarity (SS) between the embeddings of generated answers and groundtruth answers. We can see both BERT and Llama2-based models reach moderate SS-score for answer generation task. But due to answer variability mentioned before, semantic similarity (SS) does not directly reflect the preciseness of the model on correct answer generation. Hence, we validate the answer by domain experts (DE) given the reference of groundtruth answers, and we found Llama2-based generative model outperforms others, reaching 84% precision while generating answers.

Table 2

Evaluation of TComQA answers using crowdsourcing

Duration	Typical Time	Frequency	stationary	Event order
76.67%	90%	72.41%	77.77%	89.47%

Table 3

Evaluation of QA model fine-tuned with TComQA

Metrics	Duration	Typical Time	Frequency	stationary	Event order
SS	0.52	0.56	0.59	0.41	0.41
DE	86.7%	86.7%	93%	86.7%	76.6%

3.4. Generation and Evaluation of TComQA

Based on the evaluation, we consider the best-performing models for three components of our extraction pipeline. Therefore, we use the T5-based fine-tuned question generation module, then validates the generated questions using the lexical validator, and finally produces the answer using the Llama-2-based fine-tuned answer-generation model. Using the SAMsum and RealNews, we generated a total of 17K valid Question-answer pairs with associated temporal properties. We sampled 150 of them, 30 for each temporal property, and use crowdsourcing to evaluate them as no groundtruth answers are available. Each answer was assessed by 5 judges and labeled as valid, invalid, or uncertain. Groundtruth labels are generated using majority vote from crowdsourced data. Table 2 records the precision of the extracted answer for each temporal types, achieving overall 81.26% precision.

Further, we fine-tuned the answer-generation model with the TComQA using Llama2 and evaluated it on the MCTACO test set. We consider the semantic similarity between the groundtruth data and generated answer as the evaluation metric. Table 3 summarizes the performance of the fine-tuned model for all five temporal types. The evaluation shows improvement in the the answer generation task compared to the model created using the MCTACO dataset(cf. Table 1 last column). Specifically for duration, typical time, and frequency, it reaches 0.52, 0.56, and 0.59 semantic score respectively. This reflects that TComQA helps in temporal commonsense inference task.

4. Related Works

Here, we specifically focus on the recent works related to extraction, reasoning, and inference of temporal commonsense. Many studies aim to precisely structure temporal events, specifically finding the event duration [16, 6], extraction of relative timeline [17], and normalization of temporal expressions [18] for temporal events. Some works focus on event ordering, considering causal and temporal relations between events [19, 20]. Zhou et al. [9] categorize the temporal properties of events into five types of temporal commonsense and create the MCTACO dataset by enhancing crowdsourced question-answer pairs to cover all five types. Based on WikiHow pages, Lynden et al [21] extracted a large dataset on the duration of temporal events using crowdsourcing, demonstrating the effectiveness of inference tasks using LLMs.

With the advancement and availability of pre-trained LLMs, many studies have explored their power for commonsense inference tasks [22, 23, 15]. Jain et al. [23] present a survey of different LLMs on temporal inference task, discussing their limitations and potential future directions. Qin et al. [24] investigate LLMs on temporal reasoning over discourse and curated a crowdsourced dialogue dataset for this purpose. They found that LLMs are less efficient at correlating temporal contexts in conversation compared to humans. Zhou et al. [15] propose a pattern-based approach to extract temporal arguments covering three different temporal properties and use sequential language modelling to generate time-specific language model that enhances the performance of temporal commonsense inference task over BERT. Murata et al. [25] extend COMET by appending time between events that further improve the commonsense inference task and provide better embedding representations for events. In this work, to

overcome the scaling issue of crowdsourcing in temporal commonsense acquisition task, we propose an extraction pipeline for temporal commonsense by leveraging the power of LLMs with the prompts about temporal property.

5. Conclusion

We developed a temporal commonsense extraction pipeline using different large language models (LLMs). The framework identifies the temporal event and provides appropriate questions for a temporal property along with the plausible answers for the temporal event. With this framework, we created temporal commonsense QA dataset, TcomQA, from two distinct real-world datasets. Our approach eliminates the dependency on manual annotation, enabling large-scale, consistent, and efficient generation of temporal Automatically generated TComQA dataset overcomes the scaling issue in commonsense dataset generation. The evaluation of TcomQA demonstrated that our extraction pipeline successfully captures temporal commonsense from text with a precision score exceeding 80% and improve the temporal Question-answering task.

Declaration on Generative AI

During the preparation of this article, the authors have not used any GenAI tool to generate new content, such as text, images, tables, and code. During the preparation of this work, the author(s) used GPT-4 in order to: Grammar and spelling check. The authors have not used any GenAI tool to generate new content, such as text, images, tables, and code.

References

- [1] J. Romero, S. Razniewski, K. Pal, J. Z. Pan, A. Sakhadeo, G. Weikum, Commonsense properties from query logs and question answering forums, in: W. Zhu, D. Tao, X. Cheng, P. Cui, E. A. Rundensteiner, D. Carmel, Q. He, J. X. Yu (Eds.), Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM 2019, Beijing, China, November 3-7, 2019, ACM, 2019, pp. 1411–1420.
- [2] N. Tandon, G. de Melo, G. Weikum, Webchild 2.0 : Fine-grained commonsense knowledge distillation, in: M. Bansal, H. Ji (Eds.), Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, System Demonstrations, Association for Computational Linguistics, 2017, pp. 115–120.
- [3] R. Speer, C. Havasi, Conceptnet 5: A large semantic network for relational knowledge, in: I. Gurevych, J. Kim (Eds.), The People’s Web Meets NLP, Collaboratively Constructed Language Resources, Theory and Applications of Natural Language Processing, Springer, 2013, pp. 161–176.
- [4] A. Bosselut, H. Rashkin, M. Sap, C. Malaviya, A. Celikyilmaz, Y. Choi, COMET: commonsense transformers for automatic knowledge graph construction, in: A. Korhonen, D. R. Traum, L. Màrquez (Eds.), Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers, Association for Computational Linguistics, 2019, pp. 4762–4779.
- [5] J. Strötgen, M. Gertz, Heildtime: High quality rule-based extraction and normalization of temporal expressions, in: K. Erk, C. Strapparava (Eds.), Proceedings of the 5th International Workshop on Semantic Evaluation, SemEval@ACL 2010, Uppsala University, Uppsala, Sweden, July 15-16, 2010, The Association for Computer Linguistics, 2010, pp. 321–324.
- [6] A. Vempala, E. Blanco, A. Palmer, Determining event durations: Models and error analysis, in: M. A. Walker, H. Ji, A. Stent (Eds.), Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 2 (Short Papers), Association for Computational Linguistics, 2018, pp. 164–168.

- [7] H. Bagherinezhad, H. Hajishirzi, Y. Choi, A. Farhadi, Are elephants bigger than butterflies? reasoning about sizes of objects, in: D. Schuurmans, M. P. Wellman (Eds.), Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA, AAAI Press, 2016, pp. 3449–3456.
- [8] Y. Elazar, A. Mahabal, D. Ramachandran, T. Bedrax-Weiss, D. Roth, How large are lions? inducing distributions over quantitative attributes, in: A. Korhonen, D. R. Traum, L. Màrquez (Eds.), Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers, Association for Computational Linguistics, 2019, pp. 3973–3983.
- [9] B. Zhou, D. Khashabi, Q. Ning, D. Roth, "going on a vacation" takes longer than "going for a walk": A study of temporal commonsense understanding, in: K. Inui, J. Jiang, V. Ng, X. Wan (Eds.), Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019, Association for Computational Linguistics, 2019, pp. 3361–3367.
- [10] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, P. J. Liu, Exploring the limits of transfer learning with a unified text-to-text transformer, CoRR abs/1910.10683 (2019).
- [11] J. Devlin, M. Chang, K. Lee, K. Toutanova, BERT: pre-training of deep bidirectional transformers for language understanding, CoRR abs/1810.04805 (2018).
- [12] X. Zhang, D. Ramachandran, I. Tenney, Y. Elazar, D. Roth, Do language embeddings capture scales?, CoRR abs/2010.05345 (2020).
- [13] H. Touvron, et al., Llama 2: Open foundation and fine-tuned chat models, CoRR abs/2307.09288 (2023).
- [14] B. Gliwa, I. Mochol, M. Biesek, A. Wawer, Samsun corpus: A human-annotated dialogue dataset for abstractive summarization, CoRR abs/1911.12237 (2019).
- [15] B. Zhou, Q. Ning, D. Khashabi, D. Roth, Temporal common sense acquisition with minimal supervision, in: D. Jurafsky, J. Chai, N. Schluter, J. R. Tetreault (Eds.), Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020, Association for Computational Linguistics, 2020, pp. 7579–7589.
- [16] Q. Ning, H. Wu, D. Roth, A multi-axis annotation scheme for event temporal relations, in: I. Gurevych, Y. Miyao (Eds.), Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers, Association for Computational Linguistics, 2018, pp. 1318–1328.
- [17] A. Leeuwenberg, M. Moens, Temporal information extraction by predicting relative time-lines, in: E. Riloff, D. Chiang, J. Hockenmaier, J. Tsujii (Eds.), Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018, Association for Computational Linguistics, 2018, pp. 1237–1246.
- [18] K. Lee, Y. Artzi, J. Dodge, L. Zettlemoyer, Context-dependent semantic parsing for time expressions, in: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014, June 22-27, 2014, Baltimore, MD, USA, Volume 1: Long Papers, The Association for Computer Linguistics, 2014, pp. 1437–1447.
- [19] Q. Ning, Z. Feng, H. Wu, D. Roth, Joint reasoning for temporal and causal relations, in: I. Gurevych, Y. Miyao (Eds.), Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers, Association for Computational Linguistics, 2018, pp. 2278–2288.
- [20] P. Mirza, S. Tonelli, CATENA: causal and temporal relation extraction from natural language texts, in: N. Calzolari, Y. Matsumoto, R. Prasad (Eds.), COLING 2016, 26th International Conference on Computational Linguistics, Proceedings of the Conference: Technical Papers, December 11-16, 2016, Osaka, Japan, ACL, 2016, pp. 64–75.
- [21] S. J. Lynden, H. M. Yohannes, K. Kim, A. Jatowt, A. Matono, H. Yu, X. Liu, Y. Duan, Commonsense temporal action knowledge (cotak) dataset, in: I. Frommholz, F. Hopfgartner, M. Lee, M. Oakes, M. Lalmas, M. Zhang, R. L. T. Santos (Eds.), Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, CIKM 2023, Birmingham, United Kingdom, October

- 21-25, 2023, ACM, 2023, pp. 5361–5365.
- [22] M. Kimura, L. K. Pereira, I. Kobayashi, Toward building a language model for understanding temporal commonsense, in: H. Yan, Z. Yang, S. Ruder, X. Wan (Eds.), Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing, AACL/IJCNLP 2022 - Student Research Workshop, Online, November 20, 2022, Association for Computational Linguistics, 2022, pp. 17–24.
- [23] R. Jain, D. Sojitra, A. Acharya, S. Saha, A. Jatowt, S. Dandapat, Do language models have a common sense regarding time? revisiting temporal commonsense reasoning in the era of large language models, in: H. Bouamor, J. Pino, K. Bali (Eds.), Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023, Association for Computational Linguistics, 2023, pp. 6750–6774.
- [24] L. Qin, A. Gupta, S. Upadhyay, L. He, Y. Choi, M. Faruqi, TIMEDIAL: temporal commonsense reasoning in dialog, in: C. Zong, F. Xia, W. Li, R. Navigli (Eds.), Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021, Association for Computational Linguistics, 2021, pp. 7066–7076.
- [25] E. Murata, D. Kawahara, Time-aware COMET: A commonsense knowledge model with temporal knowledge, in: N. Calzolari, M. Kan, V. Hoste, A. Lenci, S. Sakti, N. Xue (Eds.), Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation, LREC/COLING 2024, 20-25 May, 2024, Torino, Italy, ELRA and ICCL, 2024, pp. 16162–16174.