

How Well Do Multi-hop Reading Comprehension Models Understand Date Information?

Xanh Ho,^{◇,♡} Saku Sugawara,[♡] and Akiko Aizawa^{◇,♡}

[◇]The Graduate University for Advanced Studies, Kanagawa, Japan

[♡]National Institute of Informatics, Tokyo, Japan

{xanh, saku, aizawa}@nii.ac.jp

Abstract

Several multi-hop reading comprehension datasets have been proposed to resolve the issue of reasoning shortcuts by which questions can be answered without performing multi-hop reasoning. However, the ability of multi-hop models to perform step-by-step reasoning when finding an answer to a comparison question remains unclear. It is also unclear how questions about the internal reasoning process are useful for training and evaluating question-answering (QA) systems. To evaluate the model precisely in a hierarchical manner, we first propose a dataset, *HieraDate*, with three probing tasks in addition to the main question: extraction, reasoning, and robustness. Our dataset is created by enhancing two previous multi-hop datasets, HotpotQA and 2WikiMultiHopQA, focusing on multi-hop questions on date information that involve both comparison and numerical reasoning. We then evaluate the ability of existing models to understand date information. Our experimental results reveal that the multi-hop models do not have the ability to subtract two dates even when they perform well in date comparison and number subtraction tasks. Other results reveal that our probing questions can help to improve the performance of the models (e.g., by +10.3 F1) on the main QA task and our dataset can be used for data augmentation to improve the robustness of the models.

1 Introduction

Multi-hop reading comprehension (RC) requires a model to read and aggregate information from multiple paragraphs to answer a given question (Welbl et al., 2018). Several datasets have been proposed for this task, such as HotpotQA (Yang et al., 2018) and 2WikiMultiHopQA (2Wiki; Ho et al., 2020). Although the proposed models show promising performances, previous studies (Jiang and Bansal, 2019; Chen and Durrett, 2019; Min et al., 2019a; Tang et al., 2021) have demonstrated that existing

Question: Who lived longer, Maceo Anderson or Jacek Karpiński?
Paragraph A: <i>Maceo Anderson</i> [1] Maceo Anderson (September 3, 1910 – July 4, 2001 in Los Angeles, California) expressed an interest in dancing at ... [2] ...
Paragraph B: <i>Jacek Karpiński</i> [3] Jacek Karpiński (9 April 1927 – 21 February 2010) was a Polish pioneer in computer engineering and ... [4] ...
Answer: Maceo Anderson
What is the date of birth of Maceo Anderson? What is the date of death of Maceo Anderson? What is the date of birth of Jacek Karpiński? What is the date of death of Jacek Karpiński?
Reasoning Task: How old was Maceo Anderson when they died? How old was Jacek Karpiński when they died?
Full-date version: Is a 90-year-10-month-1-day-old person older than a 82-year-10-month-12-day-old person?
Year-only: Is a 90-year-old person older than a 82-year-old person?
Robustness Task: Who lived shorter, Maceo Anderson or Jacek Karpiński?

Figure 1: Example of a question in our dataset.

multi-hop datasets contain reasoning shortcuts, in which the model can answer the question without performing multi-hop reasoning.

There are two main types of questions in the previous multi-hop datasets: bridge and comparison. Tang et al. (2021) explored sub-questions in the question answering (QA) process for model evaluation. However, they only used the bridge questions in HotpotQA and did not fine-tune the previous multi-hop models on their dataset when performing the evaluation. Therefore, it is still unclear about the ability of multi-hop models to perform step-by-step reasoning when finding an answer to a comparison question.

HotpotQA provides sentence-level supporting facts (SFs) to explain the answer. However, as discussed in Inoue et al. (2020) and Ho et al. (2020), the sentence-level SFs cannot fully evaluate the reasoning ability of the models; to solve this issue, in addition to sentence-level SFs, these studies pro-

vide a set of triples as the evidence information. For example, for the question in Figure 1, the evidence regards the dates of birth and death of two people, e.g., (*Maceo, date of death, July 4, 2001*). We argue that simply requiring the models to detect a set of triples, in this case, cannot explain the answer to the question and cannot describe the full path from the question to the answer; additional operations, including calculations and comparisons, need to be performed to obtain the final answer.

To deal with this issue, we introduce a dataset, *HieraDate*,¹ consisting of the three probing tasks. (1) The extraction task poses sub-questions that are created by converting evidence triples into natural language questions. (2) The reasoning task is pertinent to the combination of triples, involving comparison and numerical reasoning that precisely evaluate the reasoning path of the main questions. (3) The robustness task consists of examples generated by slightly changing the semantics (e.g., *born first to born later*) of the original main questions. The purpose of the robustness task is to ensure that the models do not exploit superficial features in answering questions.

Our dataset is created by extending two existing multi-hop datasets, HotpotQA and 2Wiki. As the first step of the proof of concept, we start with the date information through comparison questions because this information is available and straightforward to handle. Moreover, based on the classification of comparison questions in Min et al. (2019a), all comparison questions on date information require multi-hop reasoning for answering. We then use our dataset to evaluate two leading models, HGN (Fang et al., 2020) and NumNet+ (Ran et al., 2019) on two settings: with and without fine-tuning on our dataset. We also conduct experiments to investigate whether our probing questions are useful for improving QA performance and whether our dataset can be used for data augmentation.

Our experimental results reveal that existing multi-hop models perform well in the extraction and robustness tasks but fail in the reasoning task when the models are not fine-tuned on our dataset. We observe that with fine-tuning, HGN can perform well in the comparison reasoning task; meanwhile, NumNet+ struggles with subtracting two dates, although it can subtract two numbers. Our analysis shows that questions that require both nu-

merical and comparison reasoning are more difficult than questions that require only comparison reasoning. We also find that training with our probing questions boosts QA performance in our dataset, showing improvement from 77.1 to 82.7 F1 in HGN and from 84.6 to 94.9 F1 in NumNet+. Moreover, our dataset can be used as augmentation data for HotpotQA, 2Wiki, and DROP (Dua et al., 2019), which contributes to improving the robustness of the models trained on these datasets. Our results suggest that a more complete evaluation of the reasoning path may be necessary for better understanding of multi-hop models’ behavior. We encourage future research to integrate our probing questions when training and evaluating the models.

2 Related Work

In addition to Tang et al. (2021), Al-Negheimish et al. (2021) and Geva et al. (2022) are similar to our study. Al-Negheimish et al. (2021) evaluated the previous models on the DROP dataset to test their numerical reasoning ability. However, they did not investigate the internal reasoning processes of those models. Geva et al. (2022) proposed a framework for creating new examples using the perturbation of the reasoning path. Our work is different in that their focus was on creating a framework, and it does not necessarily ensure the quality of all generated perturbation samples. Moreover, we investigate the QA process in-depth, while Geva et al. (2022) do not include all detailed questions (e.g., they do not include extraction task and comparison reasoning questions in Figure 1).

3 Dataset Construction

Our dataset is generated by using the two existing multi-hop datasets, HotpotQA and 2Wiki (more details are in Appendix B.1).

Obtain Date Questions We first sampled the comparison questions in HotpotQA and 2Wiki. We then used a set of predefined keywords, such as *born first* and *lived longer*, to obtain questions regarding the date information. From the train and dev. split, respectively, we obtained 119 (after annotating, only use 114 samples) and 878 samples in HotpotQA, and 984 and 8,745 samples in 2Wiki.

Generate Sub-questions and Sub-answers In 2Wiki, we used the evidence in the form of triples (e.g., (*Maceo, date of death, July 4, 2001*)) to automatically generate sub-questions and sub-answers

¹Our data and code are available at <https://github.com/Alab-NII/HieraDate>.

Task	Templates/Phrases
Extract	What is the birth date of #name?
	What’s the birth date of #name?
	What is the date of birth of #name?
	What’s the date of birth of #name?
	When was #name born?
	What is the death date of #name?
	What’s the death date of #name?
	What is the date of death of #name?
Reason	Does #date1 come before #date2?
	Does #date1 come after #date2?
	How old was #name when they died?
	Is a #age1 person younger than a #age2 person?
	Is a #age1 person older than a #age2 person?
Robust	Born first/earlier \Leftrightarrow Born later
	Born later \Leftrightarrow Born first
	Died first/earlier \Leftrightarrow Died later
	Died later/second/last \Leftrightarrow Died first
	Died more recently \Leftrightarrow Died first
	Lived longer \Leftrightarrow Lived shorter

Table 1: List of templates and phrases that we used in the dataset creation process. *Extract*, *Reason*, and *Robust* represent the three tasks: extraction, reasoning, and robustness, respectively.

for the extraction task. We used Wikidata IDs (available in 2Wiki) to obtain structured date information to compare and/or subtract two dates when generating questions for the reasoning task. To obtain natural language questions, we wrote ten and five templates for the extraction and reasoning tasks, respectively. Similar to Min et al. (2019b), to evaluate the robustness of the models, we created the adversarial questions by changing the main multi-hop questions such that the new answers are opposite (e.g., we changed the question: “Who lived longer, A or B?” to “Who lived shorter, A or B?”). We observed that the ten phrases (e.g., *born first*) could cover all questions, and used these phrases to generate robustness questions. Table 1 presents a set of templates and phrases that we used in the dataset creation process.

In HotpotQA, unlike 2Wiki, no triples are available; therefore, we first prepared triples for the

Split	Main	Extract	Reason	Robust
Train	8745	21340	19415	8745
Dev.	549	1346	1222	549
Test	549	1346	1222	549

Table 2: Our dataset statistics. Each main question has the extraction, reasoning, and robustness tasks.

sampled questions, and then performed the same procedure as in 2Wiki to generate all probing questions. To obtain the triples, we first filtered the distractor paragraphs and retained only gold paragraphs. We then used Spacy² to extract the entities in the questions. Further, we manually annotated the date with two formats: unstructured (e.g., ‘May 1992’) and structured (e.g., month=5). It is noted that we used only the dev. set in HotpotQA.

Construct HieraDate We created our dev. and test sets from the dev. sets of HotpotQA and 2Wiki, and our training set from the 2Wiki training set. Table 2 lists the number of samples for each task and each split in our dataset. Our dataset includes two main types of questions: questions that ask about both date-of-birth and date-of-death information (e.g., “who lived longer”), and those that ask about only the date-of-birth or date-of-death information (e.g., “who was born later”). We call the first type *combined reasoning* because it requires both comparison and numerical reasoning (Figure 1). The second type is called *comparison reasoning* (Figure 2 is in Appendix B.2) because it requires only comparison reasoning. *One combined reasoning* sample has one main multi-hop question, four extraction questions, two numerical reasoning questions, one comparison question, and one robustness question. Meanwhile, *one comparison reasoning* sample has one main multi-hop question, two extraction questions, two comparison questions, and one robustness question.

4 Experiments

To comprehensively evaluate the top-performing multi-hop models, we conducted various experiments, including both with and without fine-tuning on our dataset. In addition, to discover the effectiveness of our dataset, we examine the usefulness of our probing tasks and investigate whether our dataset can be used for data augmentation.

²<https://spacy.io/>

Fine-tuning	Model	Main		Extraction		Reasoning		Robustness	
		EM	F1	EM	F1	EM (num)	EM (comp)	EM	F1
✗	HGN	66.85	76.15	94.58	96.14	<i>N/A</i>	53.08	71.95	81.64
	NumNet+	67.94	71.57	1.26	47.93	22.79 (F1)	<i>N/A</i>	69.58	71.91
✓	HGN	78.87	82.69	96.06	97.14	<i>N/A</i>	100	76.68	78.58
	NumNet+	95.08	95.20	96.36	97.73	35.96 (F1)	<i>N/A</i>	94.90	94.93
	Human (avg.)	94.00	94.90	99.16	99.53	100	98.06	95.5	95.9
	Human UB	100	100	100	100	100	100	100	100

Table 3: Results (%) of the previous models on the test set of our dataset. *Num* denotes numerical reasoning and *comp* denotes comparison reasoning. It is noted that combined reasoning questions require both numerical and comparison reasoning. *N/A* denotes not applicable. *Human UB* represents the human upper bound.

4.1 Models

As existing models cannot perform all three tasks, we evaluate these models under two groups: one focused on comparison reasoning (e.g., HGN) and the other focused on numerical reasoning (e.g., NumNet+). HGN (Fang et al., 2020) was designed to deal with HotpotQA, whereas NumNet+ (Ran et al., 2019) was designed to deal with DROP (Dua et al., 2019). Both models can perform on the extraction and robustness tasks. By design, HGN can answer yes/no questions in the comparison reasoning task. Meanwhile, NumNet+ cannot answer yes/no questions. However, it can deal with numerical questions. There are some versions of the NumNet model; in our experiment, we use the NumNet+ version.³ There are two ways to convert the questions of the extraction task in our dataset to the format of the DROP dataset. One is to use the span format, and the other is to use the date format. In our experiment, we use the span format because it produces better results than the date format.

4.2 Results

To study the abilities of the models in detail, we evaluate both models on two versions of our dataset: the full-date version (with year-month-date) and the year-only version. We also evaluate the models on two settings: with and without fine-tuning on our dataset. We use all main and probing questions together for fine-tuning the models. It is noted that we only use HieraDate when fine-tuning.

Date Understanding Evaluation Table 3 presents the results of the existing models on the

³https://github.com/llamazing/numnet_plus

full-date version of our dataset (the year-only version is in Appendix C.1). When the models are not fine-tuned on our dataset, both HGN and NumNet+ fail in the reasoning task. This can be because the forms of reasoning questions are new to these models or the models do not possess the reasoning abilities as humans do. For the extraction task, HGN performs quite well; meanwhile, NumNet+ performs worse. In the robustness task, the results are comparable with those of the main multi-hop questions. This can be explained by the fact that the patterns of the main multi-hop and robustness questions are similar.

When the models are fine-tuned on our dataset, we find that all scores of HGN improve; especially, HGN reaches the highest score in the comparison reasoning task. We conjecture that the low scores when HGN is not fine-tuned on our dataset are because the forms of the comparison reasoning questions are new to this model. Similar to HGN, the scores of the NumNet+ model also improve when it is fine-tuned on our dataset. However, the score in the numerical reasoning task on the full-date version remains low. We observed that when we evaluate NumNet+ on the year-only version, the EM scores are 83.1 and 94.4 in the numerical reasoning task for two cases: without and with fine-tuning on our dataset, respectively. This indicates that NumNet+ could perform subtraction in the form of numbers (as years) but could not in the form of dates.

Dataset Quality Check To verify the quality of our dataset, we randomly selected 100 samples from the test set and instructed graduate students to conduct the annotation. Each sample was anno-

Model	Training Data	#Questions	Evaluation Data					
			Original		Main	Extract	Reason	Robust
			EM	F1	F1	F1	F1	F1
HGN	Hotpot	90,447	67.56	81.13	76.25	94.64	26.03	79.74
	Hotpot & Ours	144,842	67.99	81.44	84.93	97.09	99.95	81.18
	2Wiki	167,454	69.42	74.21	76.69	64.62	0.0	77.35
	2Wiki & Ours	221,849	69.66	75.26	85.27	97.03	99.74	82.23
NumNet+	DROP	77,409	78.58	82.14	69.06	48.10	79.24	71.37
	DROP & Ours	120,089	78.45	82.06	95.39	97.80	94.76	94.54

Table 4: The results of the HGN and NumNet+ models on HotpotQA, 2Wiki, DROP, and our dataset. For the *Original* column, the evaluation data is HotpotQA, 2Wiki, and DROP when the model used HotpotQA, 2Wiki, and DROP for training, respectively. All reported scores in this table are average scores from two runs.

tated by two annotators. We provided the context and a list of questions to the annotators; the results are reported in Table 3. It can be observed that the human upper bound is 100% for all tasks. However, the human average is slightly low. On manually investigating the reason for this low human average, we found that the annotators made careless mistakes in several examples; however, we confirmed that these examples are answerable and reasonable.

Difficulty of Reasoning over Dates To discover whether the number of required reasoning skills in each question affects question difficulty, we compared the results of the two main types of questions in our dataset (combined vs. comparison reasoning). We found that the scores of the comparison reasoning questions were always higher than those of the combined reasoning questions (85.7 vs. 72.3 F1 in HGN; 98.8 vs. 81.6 F1 in NumNet+). The full results are in Appendix C.2. These results indicate that questions requiring both numerical and comparison reasoning are more difficult than questions that require only comparison reasoning.

QA Performance To investigate the effectiveness of our probing questions for improving the QA performance, we trained HGN and NumNet+ on six different combinations of the main and probing tasks. The results show that each task in our dataset helps to improve the performance of the main QA question (all results are in Appendix C.3). Especially when training the models on all tasks, the results improve significantly in both HGN and NumNet+ compared with the models trained on the main questions only (82.7 vs. 77.1 F1 in HGN; 94.9 vs. 84.6 F1 in NumNet+). This demonstrates

that our probing questions not only help to explain the internal reasoning process but also help to improve the score of the main multi-hop questions.

Data Augmentation We also check whether our dataset can be used for data augmentation. We trained HGN and NumNet+ on two settings, on the original dataset (e.g., HotpotQA) and on the union of the original dataset and our dataset. We use HGN for HotpotQA and 2Wiki; meanwhile, NumNet+ is used for DROP. All results are reported in Table 4. There is no significant change on the original datasets (e.g., from 81.1 to 81.4 F1 for HotpotQA); meanwhile, the improvement in our dataset is significant (e.g., from 76.3 to 84.9 F1 on the main QA task). Notably, all models that are trained on the union of the original dataset and our dataset are better in our robustness task. This indicates that our dataset can be used as augmentation data for improving the robustness of the models trained on HotpotQA, 2Wiki, and DROP.

5 Conclusion

We proposed a new multi-hop RC dataset for comprehensively evaluating the ability of existing models to understand date information. We evaluated the top-performing models on our dataset. The results revealed that the models may not possess the ability to subtract two dates even when fine-tuned on our dataset. We also found that our probing questions could help to improve QA performance, and can be used for data augmentation. For future work, we will use the hierarchical manner in our dataset to apply to other types of questions such as numerical reasoning questions in DROP.

Acknowledgments

We would like to thank Johannes Mario Meissner Blanco, Napat Thumwanit, and the anonymous reviewers for their comments and suggestions. This work was supported by JSPS KAKENHI Grant Numbers 21H03502 and 22K17954.

References

- Hadeel Al-Negheimish, Pranava Madhyastha, and Alessandra Russo. 2021. [Numerical reasoning in machine reading comprehension tasks: are we there yet?](#) In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9643–9649, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jifan Chen and Greg Durrett. 2019. [Understanding dataset design choices for multi-hop reasoning.](#) In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4026–4032, Minneapolis, Minnesota. Association for Computational Linguistics.
- Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 2019. [DROP: A reading comprehension benchmark requiring discrete reasoning over paragraphs.](#) In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2368–2378, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yuwei Fang, Siqi Sun, Zhe Gan, Rohit Pillai, Shuhang Wang, and Jingjing Liu. 2020. [Hierarchical graph network for multi-hop question answering.](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8823–8838, Online. Association for Computational Linguistics.
- Mor Geva, Tomer Wolfson, and Jonathan Berant. 2022. [Break, perturb, build: Automatic perturbation of reasoning paths through question decomposition.](#) In *Transactions of the Association for Computational Linguistics (ACL)*.
- Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. [Constructing a multi-hop QA dataset for comprehensive evaluation of reasoning steps.](#) In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6609–6625, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Naoya Inoue, Pontus Stenetorp, and Kentaro Inui. 2020. [R4C: A benchmark for evaluating RC systems to get the right answer for the right reason.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6740–6750, Online. Association for Computational Linguistics.
- Yichen Jiang and Mohit Bansal. 2019. [Avoiding reasoning shortcuts: Adversarial evaluation, training, and model development for multi-hop QA.](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2726–2736, Florence, Italy. Association for Computational Linguistics.
- Sewon Min, Eric Wallace, Sameer Singh, Matt Gardner, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2019a. [Compositional questions do not necessitate multi-hop reasoning.](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4249–4257, Florence, Italy. Association for Computational Linguistics.
- Sewon Min, Victor Zhong, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2019b. [Multi-hop reading comprehension through question decomposition and rescoring.](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6097–6109, Florence, Italy. Association for Computational Linguistics.
- Qiu Ran, Yankai Lin, Peng Li, Jie Zhou, and Zhiyuan Liu. 2019. [NumNet: Machine reading comprehension with numerical reasoning.](#) 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 2474–2484, Hong Kong, China. Association for Computational Linguistics.
- Yixuan Tang, Hwee Tou Ng, and Anthony Tung. 2021. [Do multi-hop question answering systems know how to answer the single-hop sub-questions?](#) In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3244–3249, Online. Association for Computational Linguistics.
- Ming Tu, Guangtao Wang, Jing Huang, Yun Tang, Xiaodong He, and Bowen Zhou. 2019. [Multi-hop reading comprehension across multiple documents by reasoning over heterogeneous graphs.](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2704–2713, Florence, Italy. Association for Computational Linguistics.
- Johannes Welbl, Pontus Stenetorp, and Sebastian Riedel. 2018. [Constructing datasets for multi-hop reading comprehension across documents.](#) *Transactions of the Association for Computational Linguistics*, 6:287–302.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A dataset](#)

for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.

A Limitations

There are two main limitations in our research. (1) The proposed dataset, HieraDate, focuses on only the date information. There is also a lack of diversity of operations in the dataset; it contains only subtraction and comparison operations. Other operations, such as addition and sorting, are also useful. We leave the extension for future work. (2) Another limitation is the results when training the NumNet+ model on the main and robustness tasks (Main & Robust questions in Table 7). The results drop significantly, but we have not fully investigated the reasons. At this moment, we conjecture that the reason is the contradiction of the two questions in each sample in the training data. In this setting (Main & Robust questions), each sample has only two questions, and these two questions are opposite (e.g., “Who was born first, A or B?” and “Who was born later, A or B?”). This can make the model confused; we will investigate more models on the leaderboard of DROP to find out the reasons.

B Dataset Details

B.1 Previous Datasets

HotpotQA (Yang et al., 2018) HotpotQA, created through crowdsourcing, includes two main types of questions: bridge and comparison. Unlike previous datasets, a set of sentence-level SFs information is introduced, which facilitates explainable reasoning by the system. Because of the dataset construction procedure, there is no available information that can be used to generate sub-questions.

2WikiMultiHopQA (Ho et al., 2020) 2Wiki was created using Wikipedia articles and Wikidata triples. Similar to HotpotQA, it includes two main types of questions: bridge and comparison. In 2Wiki, the authors introduced evidence information that can be used to explain the reasoning chain from question to answer. We used this information for generating sub-questions in our dataset.

B.2 HieraDate Information

Question Types As mentioned above, there are two main types of questions in our dataset: combined reasoning (Figure 1) and comparison reasoning (Figure 2). After obtaining all samples from

<p>Question: Who was born first, George Washington or Lawrence Washington?</p> <p>Paragraph A: George Washington</p> <p>[1] George Washington (February 22, 1732 – December 14, 1799) was an American ... who served as the first president ... [2] ...</p> <p>Paragraph B: Lawrence Washington</p> <p>[3] Lawrence Washington (1718–1752) was an American soldier, planter, politician, and prominent landowner in ... [4]</p>	
<p>Answer: Lawrence Washington</p> <p>What is the birth date of George Washington? } Extraction Task</p> <p>When was Lawrence Washington born? }</p> <p>Reasoning Task:</p> <p>Full-date version:</p> <p>Does February 22, 1732 come before 1718?</p> <p>Does February 22, 1732 come after 1718?</p> <p>Year-only:</p> <p>Does 1732 come before 1718?</p> <p>Does 1732 come after 1718?</p> <p>Robustness Task:</p> <p>Who was born later, George Washington or Lawrence Washington?</p>	

Figure 2: Example of a comparison reasoning question in our dataset.

HotpotQA and 2Wiki, there are only 11.3% of combined reasoning questions in the total number of examples. Therefore, we use Wikidata IDs to retrieve the missing date in a comparison reasoning question to create a combined reasoning question. For example, if the question asks “who was born first, Alice or Bob?”, to create a new sample that asks “who lived longer/shorter”, we need the date of death information. We also have several requirements, such as the date should appear in the paragraph that describes the entity. After retrieving the missing date, we use the same process as in Section 3 to generate the questions for all three tasks. It is noted that this converting process is used for the 2Wiki dataset. In the current version of the dataset, there are 22.1% of combined reasoning questions.

Date Format Wikidata uses a zero value for the dates that miss the month value or day value. In reality, we have no date with month-0 and day-0; therefore, we use a default value “1” for the dates that miss the month value or day value.

Numerical Reasoning Issue In reality, in some cases, the paragraph can contain age information, e.g., “He died in 1981 at the age of 90”. In this case, the model does not need to perform numerical reasoning. We used rules (e.g., filter whether the context contains the word “age” or not), then manually checked, and found that there are 13 para-

Fine-tuning	Model	Main		Extraction		Reasoning		Robustness	
		EM	F1	EM	F1	EM (num)	EM (comp)	EM	F1
✗	HGN	-	-	-	-	<i>N/A</i>	55.03	-	-
	NumNet+	-	-	-	-	83.07	<i>N/A</i>	-	-
✓	HGN	77.23	79.24	95.84	96.93	<i>N/A</i>	99.90	76.68	78.61
	NumNet+	94.90	94.93	96.29	97.69	94.36	<i>N/A</i>	93.99	94.01
✗	SAE (full)	69.76	77.78	82.99	84.73	<i>N/A</i>	59.14	69.22	77.82
	SAE (year)	-	-	-	-	<i>N/A</i>	55.75	-	-

Table 5: Results (%) of the previous models on the test set of our dataset. *Num* denotes numerical reasoning and *comp* denotes comparison reasoning. “-” indicates that the score is similar to the score of the full-date version in the same setting; for HGN and NumNet+, it is similar to the score in Table 3. *N/A* denotes not applicable.

graphs in a total of 248 paragraphs (124 examples) in the test set that the age information is available.

Dataset Versions Our dataset has two versions: “normal setting” and “distractor setting”. The “normal setting” includes only two gold paragraphs; meanwhile, the “distractor setting” contains ten paragraphs, including two gold paragraphs and eight distractor paragraphs. In this study, we evaluated the previous models on the “normal setting”.

C Experiments

For NumNet+, we use the parameters as described in the original source code. For HGN, when training it on HotpotQA, 2Wiki, and our dataset, we use only the loss of the answer prediction task. For other parameters, we use the same parameters as described in the source code⁴ of HGN.

C.1 Date Understanding Evaluation Details

We also evaluate the previous models on the year-only version of our dataset. Table 5 presents all the results. When the models are not fine-tuned on our dataset, the score of the HGN model in the comparison reasoning task does not change much when compared with the full-date version (55.0 vs. 53.1 EM); this indicates that there is not much difference between the full-date and year-only versions when using HGN. For NumNet+, the score of the numerical reasoning task significantly improves when compared with the full-date version (83.1⁵ vs. 22.8 F1); this indicates that NumNet+ can perform numerical reasoning in the form of numbers (as years) but cannot in the form of dates.

⁴<https://github.com/yuwfan/HGN>

⁵In the year-only version, the EM and F1-score are equal.

Evaluation on SAE Similar to HGN, SAE (Tu et al., 2019) was designed to deal with HotpotQA. The results are presented in Table 5. Similar to HGN, the model cannot perform well on the comparison reasoning questions when it is not fine-tuned on our dataset. As all questions in the comparison reasoning task are yes/no questions, the random score is 50%. The scores of both HGN and SAE are close to the chance score.

C.2 Difficulty of Reasoning over Dates

Table 6 shows the results of the previous models on the test set of our dataset for different types of questions. As shown in the table, the scores of comparison reasoning questions were always higher than those of combined reasoning questions. In the current version of the dataset, there are only 22.1% combined reasoning questions. To avoid the data-size bias, we created a HieraDate-small version by randomly choosing the comparison reasoning questions such that the number of combined reasoning questions is equal to the number of comparison reasoning questions. We then conducted experiments on HieraDate-small. We found similar results as on HieraDate. These results indicate that combined reasoning questions are more difficult than comparison reasoning questions.

C.3 QA Performance

Table 7 presents the results of the HGN and NumNet+ models on the test set of our dataset when they are trained on different subsets of our dataset.

C.4 Error Cases

Table 8 presents some error cases of the previous models on the test set of our dataset.

Model-type	Main		Extraction		Reasoning		Robustness	
	EM	F1	EM	F1	EM (num)	EM (comp)	EM	F1
HGN-all	78.87	82.69	96.06	97.14	<i>N/A</i>	100	76.68	78.58
HGN-combined	70.97	72.34	93.95	95.67	<i>N/A</i>	100	69.35	71.18
HGN-comparison	81.18	85.71	97.29	98.00	<i>NO</i>	100	78.82	80.74
HGN-all	75.40	76.67	95.30	96.71	<i>N/A</i>	99.19	76.21	77.26
HGN-combined	66.13	67.50	93.75	95.60	<i>N/A</i>	99.19	71.77	72.82
HGN-comparison	84.68	85.85	98.39	98.92	<i>NO</i>	99.19	80.65	81.69
NumNet-all	94.90	94.93	96.29	97.69	94.36	<i>N/A</i>	93.99	94.01
NumNet-combined	81.45	81.58	94.76	96.81	94.36	<i>N/A</i>	79.84	79.95
NumNet-comparison	98.82	98.82	97.18	98.20	<i>NO</i>	<i>N/A</i>	98.12	98.12
NumNet-all	85.08	85.43	95.97	97.69	94.00	<i>N/A</i>	85.48	85.50
NumNet-combined	72.58	73.27	94.76	96.84	94.00	<i>N/A</i>	73.39	73.42
NumNet-comparison	97.58	97.58	98.39	99.40	<i>NO</i>	<i>N/A</i>	97.58	97.58

Table 6: Results (%) of the previous models on the test set of our dataset for different types of questions. *Model-type* denotes the model name and the type of question that the model is evaluated on (e.g., HGN-combined: the results of HGN on combined reasoning questions). *Num* denotes numerical reasoning and *comp* denotes comparison reasoning. *N/A* denotes not applicable; meanwhile, *NO* indicates that there are no questions for evaluation. For HGN, we fine-tuned it on the full-date version of our dataset; meanwhile, NumNet+ is fine-tuned on the year-only version of our dataset. In the row with highlight color, the model is trained on HieraDate-small where the number of combined reasoning and comparison reasoning questions are equal.

Model	Training Data	#Questions	Testing Data			
			Main	Extract	Comp/Num	Robust
HGN	Main	8,745	77.11	0.0	0.0	75.45
	Main & Extract	30,085	78.37	97.14	0.0	78.18
	Main & Reason	24,310	79.06	0.0	99.79	76.62
	Main & Robust	17,490	80.96	0.0	0.0	78.04
	Main & Extract & Reason	45,650	79.97	97.10	99.59	78.40
	All	54,395	82.69	97.14	100	78.58
NumNet+	Main	8,745	84.57	0.02	0.0	82.87
	Main & Extract	30,085	92.03	97.75	0.0	89.28
	Main & Reason	12,595	88.92	0.19	94.36	89.83
	Main & Robust #1	17,490	49.86	0.23	0.0	44.84
	Main & Robust #2	17,490	48.54	0.08	0.0	50.42
	Main & Robust #3	17,490	52.95	0.02	0.0	45.24
	Main & Extract & Reason	33,935	92.01	97.89	95.16	88.91
	All	42,680	94.93	97.69	94.36	94.01

Table 7: F1-score of the HGN and NumNet+ models on the test set of our dataset when they are trained on different subsets of our dataset. *#Questions* represents the number of questions in the training data. *Comp/Num* denotes comparison reasoning or numerical reasoning; for the HGN model, it is comparison reasoning; for the NumNet+ model, it is numerical reasoning. We run three times for the “Main & Robust” setting in the NumNet+ model because the results of this setting are quite different with others.

Context	Main question	Sub-questions
<p>Paragraph A: Lotte Backes (May 2, 1901 - May 12, 1990) was a German pianist, ...</p> <p>Paragraph B: Willem van Haecht (1593 – 12 July 1637) was a Flemish painter best known for his pictures ...</p>	<p>Q: Who died first, Lotte Backes or Willem van Haecht?</p> <p>Predicted answer: Willem van Haecht ✓</p>	<p>Q1: Does May 12, 1990 come before July 12, 1637? Predicted 1: yes ✗</p> <p>Q2: Does May 12, 1990 come after July 12, 1637? Predicted 2: yes ✓</p> <p>Year-only version:</p> <p>Q3: Does 1990 come before 1637? Predicted 3: yes ✗</p> <p>Q4: Does 1990 come after 1637? Predicted 4: yes ✓</p>
<p>Paragraph A: Andrzej Markowski (22 August 1924 – 30 October 1986) was a Polish composer and conductor. ...</p> <p>Paragraph B: François Missoffe (13 October 1919 in Toulon, France – 28 August 2003 in Rouen) was a French politician and diplomat. ...</p>	<p>Q: Who lived longer, Andrzej Markowski or François Missoffe?</p> <p>Predicted answer: François Missoffe ✓</p>	<p>Q1: How old was Andrzej Markowski when they died? Predicted 1: 62 (number format) ✓</p> <p>Q2: How old was François Missoffe when they died? Predicted 2: 84 (number format) ✓</p> <p>Q3: Is a 62-year-2-month-8-day-old person older than a 83-year-10-month-15-day-old person? Predicted 3: yes ✗</p> <p>Q4: Is a 62-year-old person older than a 83-year-old person? Predicted 4: yes ✗(year-only version)</p>
<p>Paragraph A: Oliver A. Unger (August 28, 1914 – March 27, 1981) was an award- winning American film producer, distributor, ...</p> <p>Paragraph B: Ross Story (16 January 1920 – 9 May 1991), always known as Ross or C. R. Story, was a farmer and politician ...</p>	<p>Q: Who died later, Oliver A. Unger or Ross Story?</p> <p>Predicted answer: Ross Story ✓</p>	<p>Q1: What is the death date of Oliver A. Unger? Predicted 1: 9 May 1991 ✗</p> <p>Q2: What’s the death date of Ross Story? Predicted 2: 9 May 1991 ✓</p> <p>Q3: Does March 27, 1981 come before May 09, 1991? Predicted 3: yes ✓</p> <p>Q4: Does March 27, 1981 come after May 09, 1991? Predicted 4: no ✓</p>

Table 8: Error cases of the previous models on our dataset (without fine-tuning). It is noted that there are no existing models that can perform all three tasks. The results in example #2 are from the two models, HGN and NumNet+. In the first example, we can see that the models do not have the ability to compare two dates or two years. In example #2, we can observe that the models do not have the ability to subtract two dates, but the models can calculate the age by simply subtracting two years of the two dates. In example #3, we observe that the models can answer the main multi-hop question correctly, although they do not know what the date of death of a person is.