

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and
Information Systems

School of Computing and Information Systems

10-2021

NOAHQA: Numerical reasoning with interpretable graph question answering dataset

Qiyuan ZHANG

Lei WANG

Sicheng YU

Singapore Management University, scyu.2018@phdcs.smu.edu.sg

Shuohang WANG

Singapore Management University, shwang.2014@phdis.smu.edu.sg

Yang WANG

See next page for additional authors

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Databases and Information Systems Commons](#), and the [Numerical Analysis and Scientific Computing Commons](#)

Citation

ZHANG, Qiyuan; WANG, Lei; YU, Sicheng; WANG, Shuohang; WANG, Yang; JIANG, Jing; and LIM, Ee-peng.
NOAHQA: Numerical reasoning with interpretable graph question answering dataset. (2021). 1-15.
Available at: https://ink.library.smu.edu.sg/sis_research/7153

This Working Paper is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

Author

Qiyuan ZHANG, Lei WANG, Sicheng YU, Shuohang WANG, Yang WANG, Jing JIANG, and Ee-peng LIM

NOAHQA: Numerical Reasoning with Interpretable Graph Question Answering Dataset

Qiyuan Zhang^{1*} Lei Wang^{2*} Sicheng Yu^{2*} Shuohang Wang
Yang Wang³ Jing Jiang² Ee-Peng Lim²

¹University of Electronic Science and Technology of China

²Singapore Management University ³ Verily Life Sciences

qiyuanzhang97@gmail.com, {lei.wang.2019, scyu.2018}@phdcs.smu.edu.sg,
utleonwang@gmail.com, {shwang.2014, jingjiang, eplim}@smu.edu.sg

Abstract

While diverse question answering (QA) datasets have been proposed and contributed significantly to the development of deep learning models for QA tasks, the existing datasets fall short in two aspects. First, we lack QA datasets covering complex questions that involve answers as well as the reasoning processes to get the answers. As a result, the state-of-the-art QA research on numerical reasoning still focuses on simple calculations and does not provide the mathematical expressions or evidences justifying the answers. Second, the QA community has contributed much effort to improving the interpretability of QA models. However, these models fail to explicitly show the reasoning process, such as the evidence order for reasoning and the interactions between different pieces of evidence. To address the above shortcomings, we introduce NOAHQA, a conversational and bilingual QA dataset with questions requiring numerical reasoning with compound mathematical expressions. With NOAHQA, we develop an interpretable reasoning graph as well as the appropriate evaluation metric to measure the answer quality. We evaluate the state-of-the-art QA models trained using existing QA datasets on NOAHQA and show that the best among them can only achieve 55.5 exact match scores, while the human performance is 89.7. We also present a new QA model for generating a reasoning graph where the reasoning graph metric still has a large gap compared with that of humans, *e.g.*, 28 scores. The dataset and code are publicly available ¹.

1 Introduction

Question answering (QA) plays a core role in natural language understanding and is a proxy to evaluate the reading comprehension ability of intelli-

gent systems. Due to its profound significance, a

Passage:

A total of 1327 kilograms of wheat was harvested in a wheat field last year. P_1

This year, P_2
a total of 35 bags were collected, P_3
each weighing 53 kilograms. P_4

Conversation:

Q_1 : How much wheat was collected last year? QA_1
 A_1 : 1327 kilograms.

$Evidence: P_1$
 Q_2 : What about the year before last year? QA_2
 A_2 : Do not know.
 $Evidence: P_1, P_2, P_3, P_4$

Q_3 : How many bags are received this year? QA_3
 A_3 : 35 bags.
 $Evidence: P_2, P_3$

Q_4 : Does Each pack weigh 43 kilos? QA_4
 A_4 : No.
 $Evidence: P_4$

Q_5 : How much have you collected this year? QA_5
 A_5 : $35 * 53$
 $Evidence: P_4, QA_3$

Q_6 : In which year did you harvest more grain? QA_6
 A_6 : This year.
 $Evidence: QA_1, QA_5$

Q_7 : How much has it been collected until this year? QA_7
 A_7 : $1327 + A_5$
 $Evidence: QA_1, QA_5$

Reasoning Graph




Figure 1: A sample of NOAHQA dataset, which consists of a passage and several question-answer pairs. The supporting evidences and reasoning graph are provided for correctness and interpretability evaluation.

surge of datasets, *e.g.*, span-extraction (Rajpurkar et al., 2016), multiple-choice (Lai et al., 2017) and open-domain (Kwiatkowski et al., 2019), have been proposed recently. However, those datasets have limitations in numerical reasoning and interpretabil-

* The first three authors contributed equally. The order of authorship is decided through dice rolling.

¹<https://github.com/Don-Joey/NoahQA>

ity, which hinder the further advancement of QA community.

On the one hand, numerical reasoning is one of the intelligent skills of human beings. To endow such an ability to QA models, we need to provide some math word question answering dataset (Ling et al., 2017) for training these QA models. In particular, given a math question, a good QA model should select the correct answer among the multiple pre-defined answer options. Recently, DropQA (Dua et al., 2019) includes numerical questions into conventional span-extraction question answering. However, the mathematical forms of numerical questions found in DropQA are relatively simple, *e.g.*, most of the questions are only about addition and subtraction. Besides, DropQA only provides the final answer without a full expression to the answer.




Figure 2: Comparison between rationale/evidence (HotpotQA and CoQA), reasoning chain (2WikiMultiHopQA) and reasoning graph (NOAHQA).

On the other hand, existing QA datasets reveal drawbacks in interpretability (or explainability). As the leaderboards of QA datasets are overwhelmed by powerful language models, *e.g.*, BERT (Devlin et al., 2019), researchers turn their attention to explainable model and the collection of QA datasets likewise pursue this trend: models should not only give the answers, but also the explanation for the answers, *e.g.*, complex questions with discrete reasoning or numerical reasoning. The model with explanation is also friendly with digging and improving the system. To achieve the goal, CoQA (Reddy et al., 2019) and HotpotQA (Yang et al., 2018) provide the model with rational and evidences as additional supervisions. Nonetheless, how the model conducts reasoning is still vague. For example, how the model process several pieces of evidence? Are they processed in parallel or follow a specific order? R⁴C (Inoue et al., 2020), and 2WikiMultiHopQA (Ho et al., 2020) solve this problem to some extent by introducing a set of triplets or reasoning path, which is not suitable in the scenario where the reasoning process is

complicated. Taking the example in Figure 1 for illustration, the answer of Q_7 comes from the answer of Q_5 , and the answer of Q_5 comes from the passage and the answer of Q_3 .

To address the above shortcomings of existing QA datasets, we present NOAHQA, Numerical reasOning with interpretAble graph QA dataset. An overall comparison of the differences between NOAHQA and others is shown in Table 1. NOAHQA is constructed in a conversational and bilingual form with fruitful complex numerical questions demanding addition, subtraction, multiplication, division, and combination of parentheses. Meanwhile, NOAHQA provides annotated explanations in the form of reasoning graphs, namely a graph of reasoning steps, to explicitly represent the global reasoning process for each question. The comparison of reasoning graph and other explanation annotations is shown in Figure 2. Reasoning graph can be used as supervision in training as well as surrogate for evaluating the interpretability of QA models.

We apply strong baselines from existing datasets on NOAHQA and discover that the best baseline achieves 55.50 exact match scores, while human performance is 89.67. For evaluating of reasoning graph, we introduce an automatic evaluation method named DAG_{sim}, considering the structural and semantic similarity at the same time between the predicted and ground-truth reasoning graphs. To facilitate the research along with reasoning graph, we also contribute a new model named Reasoning Graph Network (RGNet) for generating the reasoning graph. Experiments with RGNet show that there is still a large gap behind human performance with 28.01 and 18.14 DAG_{sim} scores on the English and Chinese versions of NOAHQA, respectively.

2 Task

Following previous works on math word questions (Amini et al., 2019), our first task is still to generate answers for all questions in conversations. Besides, in order to get an explainable model, our second task is to generate the reasoning graph when answering questions. Formally, given a background passage P , a history conversation $QA_{1:(t-1)}$ and the next question Q_t , the task is to return a textual answer A_t to the next question Q_t and generate a reasoning graph \hat{G}_t^r . Next we will introduce the detailed notations.

Dataset	Conversational	Cross-lingual	Mathematics	Expression	Evidence	Explanation Annotations
CoQA (Reddy et al., 2019)	✓	✗	✗	✗	✓	Rationale (text span)
HotpotQA (Yang et al., 2018)	✗	✗	✗	✗	✓	Evidence (set of supporting facts)
R ⁴ C (Inoue et al., 2020)	✗	✗	✗	✗	✓	Derivation (set of triplets)
2WikiMultiHopQA (Ho et al., 2020)	✗	✗	✗	✗	✓	Reasoning Path (chain of triplets)
DropQA (Dua et al., 2019)	✗	✗	✓	✗	✗	-
Dream (Sun et al., 2019)	✗	✗	✓	✗	✗	-
MathQA (Amini et al., 2019)	✗	✗	✓	✓	✗	Operation Programs
Math23K (Wang et al., 2017a)	✗	✗	✓	✓	✗	-
NOAHQA	✓	✓	✓	✓	✓	Reasoning Graph

Table 1: Comparison of NOAHQA with existing datasets. One may argue that expression equals to evidence which is not always the case, *e.g.*, both of “two years” and “two miles” contain “two” but with totally different semantic meanings.

Textual Answer. Each sample in our dataset consists of a background passage P splitted into a sequence of segments $\{P_1, P_2, \dots, P_n\}$ by punctuations, and a conversation $QA_{1:(t-1)}$ with a series of question-answer pairs $\{QA_1, QA_2, \dots, QA_{t-1}\}$. Each answer A_i ($1 \leq i \leq t$) is associated with a set of first-order evidences, E_i , which can be some text segments in P and/or first-order evidences of previous question-answer pairs. These first-order evidences provide the information to derive the answer A_t for Q_t directly. Unlike most QA datasets (Rajpurkar et al., 2016), we only use exact match (EM) score as our evaluation metric for answer correctness, the F1 score is not adopted to bypass the occasion that two different but overlapped numbers may give high F1 scores, *e.g.*, 1203.4 and 1204.4.

Reasoning Graph. We now define the reasoning graph (RG) for question Q_t to be a directed acyclic graph (DAG) $G_t^r = \langle \mathcal{V}_t, \mathcal{L}_t^r \rangle$. For example in Figure 1, \mathcal{V}_t denotes the set of first order evidences required to derive A_t and \mathcal{L}_t^r denotes the set of directed edges between evidences from \mathcal{V}_t . For any next question Q_t , we treat it as the root node and apply the breadth-first search (BFS) to construct the RG. Specifically, BFS starts from the root node and visit all of the neighbor nodes, *i.e.*, the first-order evidence. If the present evidence is leaf node (segment in P), BFS stops. Otherwise, BFS will continue to explore its first-order evidence. Formally, for Q_t we denote the ground-truth and predicted RG as G_t and \hat{G}_t respectively.

To evaluate the quality of predicted reasoning graph \hat{G}_t , we propose DAG_{sim} , an automatic evaluation method considering the structural and semantic similarity between two DAGs. We first decompose the ground-truth graph G_t and the predicted graph \hat{G}_t into two sets of paths \mathcal{P} and $\hat{\mathcal{P}}$, respectively. A path consists of nodes from the root node to a leaf node. Then, we compute the matrix

of the best alignment scores $S \in \mathcal{R}^{|P| \times |\hat{P}|}$. Then, we use a bipartite matching algorithm over S to find the optimal matching set Π^* between \mathcal{P} and $\hat{\mathcal{P}}$. DAG_{sim} is defined as follows:

$$DAG_{sim} = \sum_{\pi^* \in \Pi^*} w_{\pi^*} s_{\pi^*}, \quad (1)$$

where s_{π^*} is the score from S for a pair of matching paths. w_{π^*} is a weight for the matching computed in terms of node frequency in the longer of two paths. The best alignment score between $p_i \in \mathcal{P}$ and $p_j \in \hat{\mathcal{P}}$ can be computed as below:

$$c(p_i, p_j) = \max_{A_k \in \mathcal{A}(p_i, p_j)} c(p_i, p_j, A_k), \quad (2)$$

where $\mathcal{A}(p_i, p_j)$ denotes all possible one-to-one alignments between p_i and p_j that do not violate chronological order. An alignment score is calculated as follows:

$$c(p_i, p_j, A_k) = \sum_{(v_{i,k}, v_{j,k}) \in A_k} a(v_{i,k}, v_{j,k}), \quad (3)$$

where $a(v_{i,k}, v_{j,k})$ is a semantic similarity score based on the text of the two nodes.

Most of traditional graph similarity methods, *e.g.*, Graph Edit Distance (GED), are hard to scale and ignore the semantic similarity between two nodes. Precision-recall of the expected edges ignores the semantic similarity of two text nodes. The proposed DAG_{sim} is more comprehensive and can run efficiently and consider structured and semantic similarity simultaneously.

3 Data Collection

In this section, we describe the collection process of NOAHQA which consists of raw passages preparation, conversation collection, evidence labeling (translation) and validation. We elaborate each step as follows.

Raw Passages Preparation. A math word problem (MWP) is composed of a short passage and a question that naturally requires the model to carefully understand the passage and take a few steps to solve. We collected questions from two classical open-source MWP datasets: the Chinese word problem dataset Math23K (Wang et al., 2017a) and the English word problem dataset MAWPS (Koncel-Kedziorski et al., 2016). Then we curated 19,098 MWPs from Math23K and 2,249 MWPs from MAWPS as the starting points for conversations.

Conversation Collection. We first hired undergraduates to create conversation collection, where each of them was provided with the annotation guidelines and examples. Finally, we chose qualified crowd-workers among them to complete the work. The guidelines are summarized as follows. *Conversation*: we require annotators to provide at least five conversation turns for each passage, except for the very few short articles. Written conversations should be concise and natural as in real occasion, *e.g.*, if two consecutive questions share the same subject, the subject of the latter question can be omitted. *Question-Answer pair*: we define six types of question-answer pairs (*i.e.*, “Extraction”, “Numerical Reasoning”, “Counterfactual”, “Comparison”, “Yes/No”, and “Unanswerable”) and provide annotators examples of these types so as to encourage them to create diverse questions. *Multi-step Reasoning*: to come up with questions with multi-step reasoning, we adopt questions in original MWPs as the reference to guide annotators to create conversations. It helps to enhance the coherence and build relationship between QA pairs in the conversation. *Answer*: the annotated answer should be either a minimum span of the text, an equation, or a fixed content, *e.g.*, “Yes / No / Do not know”. During the whole conversation generation process, we adopted the quality control mechanism of quantitative sampling, *i.e.*, for every 100 samples, we checked 20 randomly selected ones. Once any error was detected within the samples, we requested the corresponding annotator to review his 100 samples of conversation.

Evidence Labeling and Translation. We recruited another group of undergraduates to label the evidence for question. These annotators started with the first turn and labeled the spans from the passage or question-answer pairs in the previous turns as evidence, which directly supports the an-

swer to the current question. The rule for labelling is that: when labeling the evidence for the QA_t , (1) if $QA_{1:(t-1)}$ provide all information, we directly label the history turn(s) that is useful in reasoning this current question; (2) if $QA_{1:(t-1)}$ cannot offer information to answer this question, and there are another sentences in the passage that can provide this information, add the text segment P_j (even if P_j has already been used in the conversation history). This process is repeated until the evidence of the last turn is labeled. During evidence labeling, annotators marked the question-answer pair as incorrect once they find that the answer is not correct. Meanwhile, we translated the passages and annotated conversations from Chinese to English using Google Translation.

Error Correction. We invited a group of English Translation major graduates to review and correct grammatical and pragmatic errors in the translated text. This is necessary as we discovered that Google translation failed to provide high-quality translation for mathematics-related text. The common translation errors are listed in the Appendix. In the end, we have two annotators verifying 300 randomly selected samples and marking them as valid or invalid for both the English and Chinese test sets. Their Cohen’s kappa scores are 88.72 and 79.83, respectively.

Annotation De-biasing As suggested in previous papers (Clark et al., 2019; Kaushik and Lip-ton, 2018), the existing benchmarks on question answering have annotation biases, which makes designing models unnecessary. We discuss different biases and our counter-measures as follows. *Annotation Style Bias*: personal language style may affect conversation collection and evidence labeling. To prevent the dataset from simple repetitive style bias, we have 23 annotators involved in conversation collection, 3 in evidence labeling and 4 in error correction. *Question Bias*: when generating questions, annotators may prefer simple questions (*e.g.*, Extraction) over difficult ones (*e.g.*, Numerical Reasoning). We, therefore, use thresholds to restrict the proportions of different types of questions. *Reasoning Bias*: to prevent the conversation from a complete step-by-step reasoning of numerical problem, we slightly relaxed the scope of the generated questions by inserting “Counterfactual”, “Comparison”, and “Unanswerable” (*c.f.* Table 3). The above problems will be examined and corrected by the mechanism of quan-




Figure 3: Distribution of question prefixes (bigram) in NOAHQA.

titative sampling mentioned before.

4 Data Analysis

For simplicity and clarity, our data analysis is based on the English version of our dataset. NOAHQA contains 21,347 examples, each of which consists of a passage, a conversation, and corresponding evidence set. we randomly split the dataset into training, development, and test sets. The detailed statistics of the dataset are shown in Table 2. In the following, we quantitatively analyze the properties of questions, answers, and evidence in the NOAHQA dataset.

Question Analysis. We analyze the question types in conversation created by annotators and visualize the distribution of question types in Figure 3. As shown, questions beginning with “how” account for the vast majority, among them the questions asking about specific number are popular, *e.g.*, “how many” and “how much”. This is due to NOAHQA having “Extraction” and “Numerical Reasoning” as the two most frequent QA types. The number of questions beginning with “what” and “is” are ranked second and third which mainly corresponds to “Yes/No” and “Unanswerable” QA types.

Answer Analysis. Based on the annotated types of answers provided by annotators, we analyze NOAHQA to assess the distribution of the answers. As shown in Table 3, Most of the answers that can be extracted in passage directly, accounting for 46.90%. Second most of the answers are numeric values (26.22%) which require inferring the correct arithmetic expressions consisting of oper-




Figure 4: Distribution of first-order evidence. X-axis: the advance of the conversation, Y-axis: the evidence positions in segments and QA turns, The darker the position is, the more likely it is to be first-order evidence for the current question.

ations. Some of the answers in this type need external knowledge. For example, people who have never learned geometric knowledge do not know how to use π to solve circle-related problems. Beyond, we try to mix a few counterfactual questions (1.29%), where the conditions are changed, in the conversations to increase the difficulty of answering. The rest of the answers include “Yes/No” (13.76%), “Unanswerable” (6.47%), and “Comparison” (5.36%).

Evidence Analysis. Figure 4 shows the distribution of first-order evidence of answering questions as the conversation progresses before the 8th turn. The first three turns of questions are prone to finding relevant first-order evidence from the first few segments of the paragraph. The 4th and 5th questions tend to capture useful information from QA pairs close to them. The 6th and 7th questions not only utilize nearby QA pairs, but also refer to the passage. Interestingly, throughout the conversation, the segments involved in first-order evidence of all questions are the first four segments, very few questions refer to the later segments.

5 RGNet

NOAHQA requires the model to predict answers of different types, *i.e.*, a span or an arithmetic expression, and the corresponding reasoning graph (RG). Here we propose Reasoning Graph Network (RGNet) as our baseline model. Our framework consists of three main components: an encoding

	Train	Dev	Test	All
# Examples	17,077	2,135	2,135	21,347
Ave. / Max. # QA-Pairs / Example	5.08 / 10	5.09 / 10	5.09 / 9	5.08 / 10
Ave. / Max. # Segments / Passage	2.90 / 16	2.91 / 13	2.87 / 9	2.90 / 16
Ave. / Max. # Tokens / Passage	37.02 / 140	36.92 / 118	36.68 / 102	36.98 / 140
Ave. / Max. # Tokens / Question	8.78 / 48	8.77 / 40	8.74 / 39	8.77 / 48
Ave. / Max. # Tokens / Answer	1.57 / 28	1.57 / 21	1.58 / 24	1.57 / 28
Ave. / Max. # Evidences / Question	2.88 / 19	2.86 / 15	2.86 / 12	2.88 / 19

Table 2: Statistics of training, development, and test sets of NOAHQA.

Answer Type	Percentage
Extraction	46.90%
Numerical Reasoning	26.22%
w/o external knowledge	25.73%
w/ external knowledge	0.49%
Counterfactual	1.29%
Yes/No	13.76%
Unanswerable	6.47%
Comparison	5.36%

Table 3: Distribution of question-answer types in NOAHQA.

module, a reasoning module, and a prediction module. The overall structure of RGNet is shown in Figure 5.

Encoding Module. The input includes two channels: segments in the passage and ground-truth history QA pairs followed by the current question.² We use the encoding module in NAQANet (Yu et al., 2018; Dua et al., 2019) to get corresponding contextual representations.

Reasoning Module. We first construct a candidate graph with the all potential edges for RG. Then graph convolutional network (GCN) (Kipf and Welling, 2016) with $M = 3$ GCN layers is utilized. We denote two sets of representations for nodes (identifiers) and edges at m^{th} GCN layer as \mathbf{H}_n^m and \mathbf{H}_e^m , respectively. Specifically, \mathbf{H}_n^0 is extracted from the output of encoding module and \mathbf{H}_e^0 is initialized in the beginning of training. Given the representations at m^{th} GCN layer, representations at $(m+1)^{th}$ GCN layer can be obtained by:

$$\mathbf{H}_n^{m+1} = \text{Node_Update}(G'; \mathbf{H}_n^m, \mathbf{H}_e^m), \quad (4)$$

$$\mathbf{H}_e^{m+1} = \text{Edge_Update}(G'; \mathbf{H}_n^m, \mathbf{H}_e^m), \quad (5)$$

Then we apply an edge classifier over the output

²We add special tokens: *Yes*, *No*, *Unknown*, and operations, i.e., $\{+, -, \times, \div\}$, to the end of the input text to facilitate the prediction for corresponding types of questions.

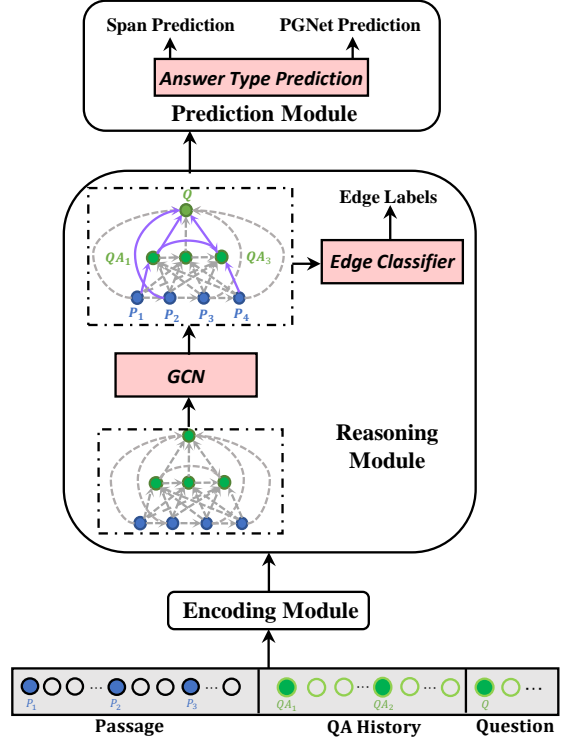


Figure 5: Framework of RGNet. Our model consists of an encoding module, a reasoning Module, and a prediction module.

of the last GCN layer, \mathbf{H}_e^M to determine which edge in candidate graph G' is existed in RG. Then we feed the combination of \mathbf{H}_n^M and contextual representations (from encoding module) into the prediction module.

Prediction Module. For simplicity, we group six types of answer into two sets: extractive span or generative sequence. Following NAQANet (Dua et al., 2019), we use an answer type prediction layer to decide the type of the answer. For extractive span type, we follow the standard implementation (Wang and Jiang, 2016) to find the start and end positions. For generative sequence, we adopt the pointer-generator network (PGNet) (See et al., 2017a).

Training. In RGNet, there are four objectives dur-

ing training, *i.e.*, answer type prediction, span prediction, sequence generation, edge classification. Thus, the final loss function is defined as:

$$Loss = XE_{type} + XE_{span} + XE_{seq} + XE_{edge}, \quad (6)$$

where XE denotes cross-entropy loss.

6 Experiment

6.1 Experimental Settings

Baselines. We evaluate NOAHQA on the models designed for four widely used datasets: 1) baselines from CoQA includes seq2seq implemented by OpenNMT (Klein et al., 2017), PGNet (See et al., 2017b), and FlowQA (Huang et al., 2018). These methods utilize ground-truth history answers to answer the current question; 2) baselines (denoted as HOTPOT) (Yang et al., 2018) introduced in HotpotQA dataset, which utilizes the supervised signals of evidence while answering the questions; 3) GTS (Xie and Sun, 2019a) from the MWP dataset Math23K. GTS translates an MWP text into an arithmetic expression through a goal-driven tree decoder; 4) NAQANet (Dua et al., 2019) and its strong variant NumNet+ (Ran et al., 2019) from DropQA. NAQANet can predict multiple types of answers. NumNet+ combines NAQANet with a numeric aware graph neural network. To establish human performance, we randomly sample 300 examples from the test set. For each example, we average evaluation scores of predictions from two annotators.

Implementation Details. In RGNet, we utilize RoBERTa (Liu et al., 2019) and XLM-R (Conneau et al., 2020) as encoding module for monolingual and cross-lingual experiments, respectively. Adam (Kingma and Ba, 2014) is selected as optimizer in training. In the prediction module, we use one layer GRU (Cho et al., 2014) for PGNet.

6.2 Results and Discussion

Main Results. The results of different models on our NOAHQA dataset are shown in Table 4. We use EM scores mentioned in Section 2 as the evaluation metric for answer correctness. We observe that GTS as a strong baseline in solving MWP performs worst because it can not generate text or extract spans from the text. PGNet outperforms seq2seq due to its stronger ability to produce tokens in the paragraph and historical context. Notably, FlowQA not only outperforms PGNet and seq2seq but also beat all of the other baselines due to its

	en (EM)		zh (EM)	
	Dev	Test	Dev	Test
Seq2seq	44.81	44.93	47.32	46.78
PGNet	47.33	47.15	49.85	49.37
FlowQA	56.78	55.50	55.26	54.02
HOTPOT	51.24	51.18	48.02	47.92
GTS	8.73	8.36	11.17	10.40
NAQANet	51.35	50.45	47.34	45.78
NumNet+	53.04	52.34	49.60	47.82
RGNet	63.04	61.69	64.90	62.94
w/o Pre	57.48	56.23	57.60	54.32
w/o Edge	62.67	60.12	64.23	62.07
w/o Edge&Pre	57.54	55.59	57.11	55.89
w/o GCN	60.03	58.74	62.23	60.89
w/o GCN&Pre	56.56	55.13	55.68	52.78
Human	-	89.67	-	92.76

Table 4: The experimental results of different models, including an ablation study for different RGNet variants. Pre = Pre-Trained Model (RoBERTa), Edge = Edge Classifier, GCN = Graph Convolutional Network.

special design for conversation. HOTPOT using intermediate evidence as supervision is comparable to PGNet. NumNet+ consistently outperforms NAQANet on two versions of NOAHQA due to the numeric comparison graph.³ Since RGNet is able to cover all types of questions, including GCN for reasoning, span prediction for locating text span, and PGNet for producing arithmetic expressions, it obtains overall highest performance. We also perform the one-sample t-test, and p-value = $2e-6 < 0.05$ indicates that the improvements of RGNet w/o Edge over FlowQA (the strongest baseline) are statistically significant.

From the ablative settings for RGNet, we discover that performance gain from pretraining is considerable. Meanwhile, GCN is also effective w/ or w/o supervision of ground-truth reasoning graph (edge classifier). Comparing full model and the model w/o Edge, we find that edge supervision contributes to the final performance, *e.g.*, 1.57 points on English test set. From an overall view, the best model performance from RGNet is still 27.98 points and 29.82 points behind humans on the English and Chinese test sets, respectively.

Reasoning Graph Evaluation. Table 5 presents the graph exact match (GEM) and similarity score (DAG_{sim}) for the predicted reasoning graphs where GEM measures whether two graphs are the same.

We can clearly observe that DAG_{sim} is higher than GEM as it considers both structural and seman-

³More experiment results under cross-lingual setting are in appendix.

	en		zh	
	GEM	DAG _{sim}	GEM	DAG _{sim}
Edge w/o Pre	43.14	63.17	47.15	66.92
Edge w/ Pre	49.21	69.83	59.32	79.56
Human	93.80	97.84	92.53	97.71

Table 5: Performance of RGNet w.r.t GEM score and DAG_{sim} score for reasoning graphs on test set.




Figure 6: Performance changes in conversation flow on the test set (zh).

tic similarities. Our model can achieve 49.21 GEM points and 69.83 DAG_{sim} scores for the English test set and 59.32 GEM points and 79.56 DAG_{sim} scores for the Chinese test set. It also suggests that there is still a lot of room for improvement in generating reliable reasoning graph.

Analysis of Different Answer Types. As shown in Table 6, we perform a break-down analysis of different answer types to various methods on the Chinese test set. GTS obtains the best result of predicting the arithmetic answers. FlowQA achieves the best results on text spans prediction. As seq2seq and PGNet can generate arithmetic expressions, they perform better on numerical questions.

Analysis of Conversation Flow. Figure 6 shows how the performance changes as the conversations progress. As the complexity of questions is increasing along the X-axis, the performance is also decreasing. Since more unanswerable questions appear after the 6th question, the performances increase. Notably, the performance of RGNet degrades more slowly than NumNet+.

Analysis of the Necessity of Passages and Historical QA Pairs. Min et al. (2019) has shown that HOTPOTQA dataset has artifacts to cheat models with superficial patterns and facts in multihop reasoning are redundant. Thus, to verify that evidence is necessary in NOAHQA, we perform experiments on RGNet w/ or w/o passages and historical QA pairs. As shown in Table 7, the significant drop in the results of model w/o passages or historical QA

pairs indicates that these content are indispensable in our dataset.

Zero-shot Transfer. We perform a zero-shot transfer experiment to investigate how different our NOAHQA from the complex QA dataset DropQA and the MWP dataset Math23K. We train NumNet+ model on DropQA and GTS model on Math23k and then test them on the English and Chinese test set of NOAHQA, respectively. As shown in Table 8, both models perform poorly, indicating that our NOAHQA is vastly different from DropQA and Math23K.

7 Related Work

Math Word Problems. In the past few years, there has been a growing number of datasets (Wang et al., 2017a; Miao et al., 2020; Patel et al., 2021) and methods that have been proposed for MWPs, including statistical machine learning methods (Mittra and Baral, 2016; Roy and Roth, 2018), semantic parsing methods (Liang et al., 2018), and deep learning methods (Wang et al., 2017b, 2018b,a, 2019; Xie and Sun, 2019b; Zhang et al., 2020; Qin et al., 2020; Wu et al., 2020), emerging in the field of solving MWPs.

Question Answering Datasets. This work mainly refers to conversational QA datasets (Reddy et al., 2019; Choi et al., 2018; Christmann et al., 2019), Multi-hop QA datasets (Talmor and Berant, 2018; Yang et al., 2018; Inoue et al., 2020; Ho et al., 2020; Chen et al., 2020), and discrete reasoning datasets (Dua et al., 2019; Sun et al., 2019). For explanation in QA, CoQA (Reddy et al., 2019) provide rationale to make models understandable under conversations. HotpotQA (Yang et al., 2018), R⁴C (Inoue et al., 2020), and 2WikiMultiHopQA (Ho et al., 2020), provide a set of evidence to support training models to learn reasoning across paragraphs. In concurrent, Dalvi et al. (2021) present a new dataset with explanations in the form of entailment trees, however, they do not consider numerical reasoning.

Interpretable Reasoning. There have been many works (Gontier et al., 2020; Saha et al., 2020; Wolfson et al., 2020) exploring interpretable reasoning recently. Related to our work, (Dalvi et al., 2021) generates explanations in the form of entailment trees, namely a tree of entailment steps from known facts.

	Seq2seq	PGNet	FlowQA	Hotpot	GTS	NAQANet	NumNet+	RGNet*	RGNet
Extract	60.19	61.21	79.21	68.96	-	63.58	63.61	65.34	73.20
Yes/No	78.34	78.64	81.23	76.86	-	81.83	84.21	85.11	90.96
Comparison	12.15	14.76	38.07	18.66	-	3.37	26.65	31.24	46.56
Arithmetic	19.45	23.34	-	-	37.80	1.02	1.16	24.12	33.73
Counterfactual	20.12	24.51	11.06	2.14	-	3.76	4.81	22.45	29.18
Unanswerable	79.51	78.48	86.23	85.45	-	88.96	88.94	88.95	93.89

Table 6: Performance of different answer types on the test set (zh). RGNet* represents RGNet w/o Pre.

	RGNet	w/o Passage	w/o History
Overall Performance	62.94	19.31	52.67
Yes/No	90.96	80.04	90.35
Comparison	46.56	21.74	23.68
Arithmetic	33.73	11.79	23.87
Counterfactual	29.18	11.22	22.89
Extract	73.20	1.12	61.21
Unanswerable	93.89	37.57	80.56
Reasoning Graph	59.32	9.01	38.92

Table 7: Experiments w.r.t. RGNet w/ or w/o passages or historical QA pairs on our NOAHQA dataset.

Datasets	Models	en	zh
DropQA (en)	NumNet+	7.53	-
Math23K (zh)	GTS	-	8.96

Table 8: Transfer results of other datasets on NOAHQA.

8 Conclusion

In this work, we present a new QA datasets with complex numerical questions and interpretable reasoning graph. We also introduce an automatic evaluation metric for the generated reasoning process. We finally present an initial model producing the reasoning process while answering questions. The experiments show that NOAHQA is challenging and will become an interesting direction in both numerical QA and explainable QA.

9 Ethics

Finally, we want to state the cost and wage issues of our dataset annotation. We recruited 30 annotators from Chinese universities. We pay CNY1400 per 1000 samples, and it takes 83.33 hours to annotate every 1000 samples. Therefore, the hourly salary is about CNY16.8, equivalent to USD2.52 per hour. Please note that the minimum average hourly wage of the Sichuan Province of China (which is where the recruited annotators are from) is CNY16.3 per hour in 2018. Therefore, our pay is above the minimum average hourly wage. In total, We spent USD4500 in this dataset.

References

- Aida Amini, Saadia Gabriel, Shanchuan Lin, Rik Koncel-Kedziorski, Yejin Choi, and Hannaneh Hajishirzi. 2019. MathQA: Towards interpretable math word problem solving with operation-based formalisms. 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 2357–2367.
- Wenhu Chen, Hanwen Zha, Zhiyu Chen, Wenhan Xiong, Hong Wang, and William Yang Wang. 2020. Hybridqa: A dataset of multi-hop question answering over tabular and textual data. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 1026–1036.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734.
- Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wentaoh Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. 2018. Quac: Question answering in context. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2174–2184.
- Philipp Christmann, Rishiraj Saha Roy, Abdalghani Abujabal, Jyotsna Singh, and Gerhard Weikum. 2019. Look before you hop: Conversational question answering over knowledge graphs using judicious context expansion. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pages 729–738.
- Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. 2019. Don’t take the easy way out: Ensemble based methods for avoiding known dataset biases. 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 4069–4082.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco

- Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451.
- Bhavana Dalvi, Peter Jansen, Oyvind Tafjord, Zhengnan Xie, Hannah Smith, Leighanna Pipatanangkura, and Peter Clark. 2021. Explaining answers with entailment trees. *arXiv preprint arXiv:2104.08661*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- 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.
- Nicolas Gontier, Koustuv Sinha, Siva Reddy, and Chris Pal. 2020. Measuring systematic generalization in neural proof generation with transformers. In *Advances in Neural Information Processing Systems*, volume 33, pages 22231–22242.
- 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.
- Hsin-Yuan Huang, Eunsol Choi, and Wen-tau Yih. 2018. Flowqa: Grasping flow in history for conversational machine comprehension. In *International Conference on Learning Representations*.
- Naoya Inoue, Pontus Stenertorp, 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.
- Divyansh Kaushik and Zachary C Lipton. 2018. How much reading does reading comprehension require? a critical investigation of popular benchmarks. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5010–5015.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander Rush. 2017. OpenNMT: Open-source toolkit for neural machine translation. In *Proceedings of ACL 2017, System Demonstrations*, pages 67–72, Vancouver, Canada.
- Rik Koncel-Kedziorski, Subhro Roy, Aida Amini, Nate Kushman, and Hannaneh Hajishirzi. 2016. MAWPS: A math word problem repository. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1152–1157, San Diego, California.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. Race: Large-scale reading comprehension dataset from examinations. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 785–794.
- Chao-Chun Liang, Yu-Shiang Wong, Yi-Chung Lin, and Keh-Yih Su. 2018. A meaning-based statistical English math word problem solver. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 652–662.
- Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. 2017. Program induction by rationale generation: Learning to solve and explain algebraic word problems. *ArXiv*, abs/1705.04146.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Shen-yun Miao, Chao-Chun Liang, and Keh-Yih Su. 2020. A diverse corpus for evaluating and developing English math word problem solvers. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 975–984, Online.
- Sewon Min, Eric Wallace, Sameer Singh, Matt Gardner, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2019. Compositional questions do not necessitate multi-hop reasoning. *arXiv preprint arXiv:1906.02900*.

- Arindam Mitra and Chitta Baral. 2016. Learning to use formulas to solve simple arithmetic problems. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2144–2153. Association for Computational Linguistics.
- Arkil Patel, Satwik Bhattamishra, and Navin Goyal. 2021. Are NLP models really able to solve simple math word problems? In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2080–2094.
- Jinghui Qin, Lihui Lin, Xiaodan Liang, Rumin Zhang, and Liang Lin. 2020. Semantically-aligned universal tree-structured solver for math word problems. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3780–3789.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392.
- 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.
- Siva Reddy, Danqi Chen, and Christopher D Manning. 2019. Coqa: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics*, 7:249–266.
- Subhro Roy and Dan Roth. 2018. Mapping to declarative knowledge for word problem solving. *Transactions of the Association for Computational Linguistics*, 6:159–172.
- Swarnadeep Saha, Sayan Ghosh, Shashank Srivastava, and Mohit Bansal. 2020. PProver: Proof generation for interpretable reasoning over rules. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 122–136.
- A. See, Peter J. Liu, and Christopher D. Manning. 2017a. Get to the point: Summarization with pointer-generator networks. *ArXiv*, abs/1704.04368.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017b. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083.
- Kai Sun, Dian Yu, Jianshu Chen, Dong Yu, Yejin Choi, and Claire Cardie. 2019. Dream: A challenge data set and models for dialogue-based reading comprehension. *Transactions of the Association for Computational Linguistics*, 7:217–231.
- Alon Talmor and Jonathan Berant. 2018. The web as a knowledge-base for answering complex questions. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 641–651.
- Lei Wang, Yan Wang, Deng Cai, Dongxiang Zhang, and Xiaojiang Liu. 2018a. Translating a math word problem to a expression tree. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1064–1069. Association for Computational Linguistics.
- Lei Wang, Dongxiang Zhang, Lianli Gao, Jingkuan Song, Long Guo, and Heng Tao Shen. 2018b. Mathdqn: Solving arithmetic word problems via deep reinforcement learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*, pages 5545–5552.
- Lei Wang, Dongxiang Zhang, Zhang Jipeng, Xing Xu, Lianli Gao, Bing Tian Dai, and Heng Tao Shen. 2019. Template-based math word problem solvers with recursive neural networks. In *Thirty-Third AAAI Conference on Artificial Intelligence*, pages 7144–7151.
- Shuohang Wang and Jing Jiang. 2016. Machine comprehension using match-istm and answer pointer. *arXiv preprint arXiv:1608.07905*.
- Yan Wang, Xiaojiang Liu, and Shuming Shi. 2017a. Deep neural solver for math word problems. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, Copenhagen, Denmark.
- Yan Wang, Xiaojiang Liu, and Shuming Shi. 2017b. Deep neural solver for math word problems. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 845–854. Association for Computational Linguistics.
- Tomer Wolfson, Mor Geva, Ankit Gupta, Matt Gardner, Yoav Goldberg, Daniel Deutch, and Jonathan Berant. 2020. Break it down: A question understanding benchmark. *Transactions of the Association for Computational Linguistics*, 8:183–198.
- Qinzhao Wu, Qi Zhang, Jinlan Fu, and Xuanjing Huang. 2020. A knowledge-aware sequence-to-tree network for math word problem solving. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7137–7146, Online. Association for Computational Linguistics.
- Zhipeng Xie and Shichao Sun. 2019a. A goal-driven tree-structured neural model for math word problems. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 5299–5305.

- Zhipeng Xie and Shichao Sun. 2019b. A goal-driven tree-structured neural model for math word problems. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 5299–5305. International Joint Conferences on Artificial Intelligence Organization.
- 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.
- Adams Wei Yu, David Dohan, Minh-Thang Luong, Rui Zhao, Kai Chen, Mohammad Norouzi, and Quoc V Le. 2018. Qanet: Combining local convolution with global self-attention for reading comprehension. *arXiv preprint arXiv:1804.09541*.
- Jipeng Zhang, Lei Wang, Roy Ka-Wei Lee, Yi Bin, Yan Wang, Jie Shao, and Ee-Peng Lim. 2020. Graph-to-tree learning for solving math word problems. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3928–3937.

Model	Train	en		zh	
		Dev	Test	Dev	Test
RoBERTa	en	63.04	61.69	–	–
	zh	–	–	64.90	62.94
XLM-R	en	64.42	62.26	57.00	54.13
	zh	48.63	45.78	68.84	66.76
	en+zh	66.37	63.89	69.24	67.00

Table 9: Performance comparison under different cross-lingual settings.

	RGNet	RGNet w/o Pre
train epoch	40	40
batch size	4	8
max length	512	512
hidden size	768	128
num hidden layers	24	–
num heads	24	–
learning rate schedule	–	–
learning rate	5e-6	1e-4
dropout	0.1	0.1

Table 10: Summary of hyperparameters derived from the defaults. Default hyperparameter sets are: RGNet, RGNet w/o Pre.

A Cross-Lingual Experiment

In Table 9, we conduct cross-lingual experiments on Chinese corpus and machine-translated English corpus. When training and testing in the same language, XLM-R pre-training achieved better results than normal RoBERTa; while in the case of monolingual training, the experimental effect of testing in another language was worse than that of testing in the same language. Finally, the best experimental results can be achieved by mixing the two languages corpus.

B Summary of Hyperparameters

In Table 10, We show the hyperparameters of the default model.

C Error Details

In addition to grammatical and pragmatic errors in machine translation, which occur in traditional generic text domain, our dataset, due to the properties of MWP and natural conversation, can generate some unique translation errors in machine translation, which deserve attention from cross-lingual research and are also the focus of manual error correction, as follows in Figure 1:

D Examples of Question-answer types

In Table 11, We show examples of different types of Question Answering.

E Detailed Examples

We show two detailed examples in Figure 8, including passage, conversations, and reasoning graphs.

F Interface of Conversation Collection

As depicted in Figure 9, we show the translated interface for annotating the conversation about an passage. We automatically create an excel worksheet for each passage (the white part is automatically imported and filled in by the machine), and the annotator needs to follow the guidelines to complete the collection (the blue part is the part that the annotator needs to fill in). Column 2 is the area where the annotator completes the conversation collection, and the annotator fills in the desired conversation according to the content of the passage in (row 2, column 2) and the annotating guidelines. The annotator also needs to fill in some attributes of the corresponding question answer pair, column 3-7, respectively, coreference relationship (between the current question and historical question), Matching information (whether the answer directly corresponds to the phrase in the passage), Question Type, the historical question answer pairs evidence labeling, and unanswerable. After generating the conversation, the annotator needs to make a multi-choice among these attributes that corresponding to QA pair, meanwhile, We give explanation examples and definitions above.

Error Types	Examples	
	Original Text	Machine Translation
Numerical	One-thirteenth	thirteenth
	15% discount on all leather shoes	0.85 discount on all leather shoes
	1.3 billion and 3 million	1.3 billion
	Grandma's age is more five years than seven times her age.	Grandma's age is 7 times more than 5 years old
	A canal has been built, 8/15 of which has been built in eight days.	Build a canal, which has been built in 8 days (8/15).
Non-Numerical	A:How tall is Xiaofang? B:117cm. A:What about XiaoMing?	A:How tall is Xiaofang? B:117cm. A:Where's XiaoMing?

Figure 7: Error Types and corresponding examples. "original text" represents that the original Chinese text is expressed in English accurately.

Answer Type	Example
Extraction	P: The canteen has 580 kilograms of coal. It burns for 6 days. It burns 36 kilograms of coal every day. How many kilograms are left? Q: How many kilograms does it burn per day? A: 36 kilograms.
Numerical Reasoning	
w/o external knowledge	Q: How much is the price of sandals? A: 19. $(10 + 10 \times 90\%)$
w/ external knowledge	Q: What is the circumference of the bottom surface? A: 4.71. $(\pi \times 1.5)$
Counterfactual	Q: If the survival rate increased to 90%, how many saplings do you need? A: 4666.67. $(4200 \div 90\%)$
Yes/No	Q: Is the BBK VCD price reduced? A: Yes.
Unanswerable	Q: What brand of soy sauce is it? A: Do not know.
Comparison	Q: Which is more expensive, the original price or the current price? A: Original price.

Table 11: Examples of question-answer types in NOAHQA.

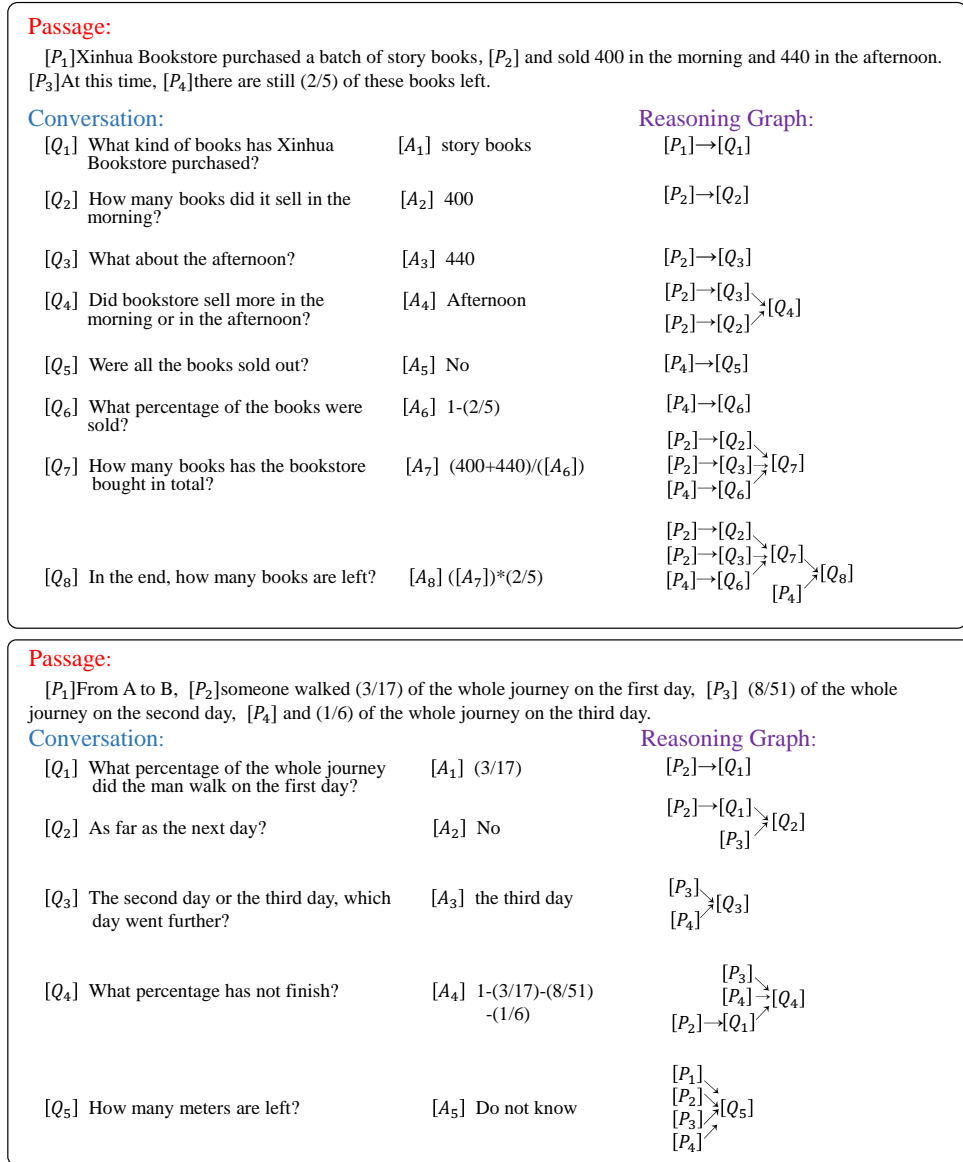


Figure 8: Detailed Examples.

	text	Coreference Relation	Matching information	Question Type	historical question answer evidence	Unanswerable
passage	A total of 1327 kilograms of wheat was harvested in a wheat field last year. This year, a total of 35 bags were collected, each weighing 53 kilograms.	1 Non Coreference: example question:Who does Xiao Ming love? 2 Explicit Coreference: example question 1:Who is this person in the article? example answer 1:Xiao Ming example question 2:Who does he love? 3 Vague Coreference: example question 1:When did Xiaoming fall in love with her? example answer 1:2019 example question 2:Where	passage: Xiao Ming was rescued by the Coast Guard. Question:Who saved Xiao Ming? Answer: Coast Guard.	1.Yes/No: Is it big? 2.Comprison: Which is bigger? 3.Arithmetic: How much is the addition of the two? 4.counterfactual: Change one condition in the question and ask a new question. 5.others	Use the answer of the previous question to answer this question at the moment. There are 1542 fruit trees in the orchard, including 147pear trees, 684peach trees and the rest are apple trees. Q1: how many apple trees are there? A1:711. Q2: which kind of tree is the most? A2: Apple trees	There are some questions that can't be answered.
num_list	{num_1:'1327', num_2:'35', num_3:'53'}					
question_id_1	How much wheat was collected last year?	Non Coreference				
answer_id_1	1327 kilograms.		yes			
question_id_2	What about the year before last year?	Vague Coreference				
answer_id_2	Do not know.		no			yes
question_id_3	How many bags are received this year?	Non Coreference				
answer_id_3	35 bags.		yes			
question_id_4	Does Each pack weigh 43 kilos?	Non Coreference				
answer_id_4	No.		yes	Yes/No		
question_id_5	How much have you collected this year?	Vague Coreference				
answer_id_5	num_2+num_3		no	arithmetic		3
question_id_6	in which year did you harvest more grain?	Vague Coreference				
answer_id_6	This year.		no	comprison	1, 5	
question_id_7	How much has it been collected until this year?	Explicit Coreference				
answer_id_7	num_1-(num_2+num_3)		no	arithmetic	1, 5	
question_id_8						
answer_id_8						
question_id_9						
answer_id_9						

Figure 9: Interface of Conversation Collection.