

NumNet: Machine Reading Comprehension with Numerical Reasoning

Qiu Ran^{1*}, Yankai Lin^{1*}, Peng Li¹, Jie Zhou¹, Zhiyuan Liu²

¹Pattern Recognition Center, WeChat AI, Tencent Inc, China

²Department of Computer Science and Technology, Tsinghua University, Beijing, China

Institute for Artificial Intelligence, Tsinghua University, Beijing, China

State Key Lab on Intelligent Technology and Systems, Tsinghua University, Beijing, China

{soulcaptran, yankailin, patrickpli, withtomzhou}@tencent.com

liuzy@tsinghua.edu.cn

Abstract

Numerical reasoning, such as addition, subtraction, sorting and counting is a critical skill in human’s reading comprehension, which has not been well considered in existing machine reading comprehension (MRC) systems. To address this issue, we propose a numerical MRC model named as NumNet, which utilizes a numerically-aware graph neural network to consider the comparing information and performs numerical reasoning over numbers in the question and passage. Our system achieves an EM-score of 64.56% on the DROP dataset, outperforming all existing machine reading comprehension models by considering the numerical relations among numbers.

1 Introduction

Machine reading comprehension (MRC) aims to infer the answer to a question given the document. In recent years, researchers have proposed lots of MRC models (Chen et al., 2016; Dhingra et al., 2017; Cui et al., 2017; Seo et al., 2017) and these models have achieved remarkable results in various public benchmarks such as SQuAD (Rajpurkar et al., 2016) and RACE (Lai et al., 2017). The success of these models is due to two reasons: (1) Multi-layer architectures which allow these models to read the document and the question iteratively for reasoning; (2) Attention mechanisms which would enable these models to focus on the part related to the question in the document.

However, most of existing MRC models are still weak in numerical reasoning such as addition, subtraction, sorting and counting (Dua et al., 2019), which are naturally required when reading financial news, scientific articles, etc. Dua et al. (2019) proposed a numerically-aware QANet

(NAQANet) model, which divides the answer generation for numerical MRC into three types: (1) extracting spans; (2) counting; (3) addition or subtraction over numbers. NAQANet makes a pioneering attempt to answer numerical questions but still does not explicitly consider numerical reasoning.

To tackle this problem, we introduce a novel model NumNet that integrates numerical reasoning into existing MRC models. A key problem to answer questions requiring numerical reasoning is how to perform numerical comparison in MRC systems, which is crucial for two common types of questions:

(1) **Numerical Comparison:** The answers of the questions can be directly obtained via performing numerical comparison, such as sorting and comparison, in the documents. For example, in Table 1, for the first question, if the MRC system knows the fact that “ $49 > 47 > 36 > 31 > 22$ ”, it could easily extract that the second longest field goal is 47-yard.

(2) **Numerical Condition:** The answers of the questions cannot be directly obtained through simple numerical comparison in the documents, but often require numerical comparison for understanding the text. For example, for the second question in Table 1, an MRC system needs to know which age group made up more than 7% of the population to count the group number.

Hence, our NumNet model considers numerical comparing information among numbers when answering numerical questions. As shown in Figure 1, NumNet first encodes both the question and passages through an encoding module consisting of convolution layers, self-attention layers and feed-forward layers as well as a passage-question attention layer. After that, we feed the question and passage representations into a numerically-aware graph neural network (NumGNN) to further

*indicates equal contribution

Question	Passage	Answer
What is the second longest field goal made?	... The Seahawks immediately trailed on a scoring rally by the Raiders with kicker <i>Sebastian Janikowski nailing a 31-yard field goal</i> ... Then in the third quarter <i>Janikowski made a 36-yard field goal</i> . Then <i>he made a 22-yard field goal</i> in the fourth quarter to put the Raiders up 16-0 ... The Seahawks would make their only score of the game with kicker <i>Olindo Mare hitting a 47-yard field goal</i> . However, they continued to trail as <i>Janikowski made a 49-yard field goal</i> , followed by RB Michael Bush making a 4-yard TD run.	47-yard
How many age groups made up more than 7% of the population?	Of Saratoga Countys population in 2010, <i>6.3%</i> were between ages of 5 and 9 years, <i>6.7%</i> between 10 and 14 years, <i>6.5%</i> between 15 and 19 years, <i>5.5%</i> between 20 and 24 years, <i>5.5%</i> between 25 and 29 years, <i>5.8%</i> between 30 and 34 years, <i>6.6%</i> between 35 and 39 years, <i>7.9%</i> between 40 and 44 years, <i>8.5%</i> between 45 and 49 years, <i>8.0%</i> between 50 and 54 years, <i>7.0%</i> between 55 and 59 years, <i>6.4%</i> between 60 and 64 years, and <i>13.7%</i> of age 65 years and over ...	5

Table 1: Example questions from the DROP dataset which require numerical comparison. We highlight the relevant parts in the passage to infer the answer.

integrate the comparison information among numbers into their representations. Finally, we utilize the numerically-aware representation of passages to infer the answer to the question.

The experimental results on a public numerical MRC dataset DROP (Dua et al., 2019) show that our NumNet model achieves significant and consistent improvement as compared to all baseline methods by explicitly performing numerical reasoning over numbers in the question and passage. In particular, we show that our model could effectively deal with questions requiring sorting with multi-layer NumGNN. The source code of our paper is available at <https://github.com/ranqiu92/NumNet>.

2 Related Work

2.1 Machine Reading Comprehension

Machine reading comprehension (MRC) has become an important research area in NLP. In recent years, researchers have published a large number of annotated MRC datasets such as CNN/Daily Mail (Hermann et al., 2015), SQuAD (Rajpurkar et al., 2016), RACE (Lai et al., 2017), TriviaQA (Joshi et al., 2017) and so on. With the blooming of available large-scale MRC datasets, a great number of neural network-based MRC models have been proposed to answer questions for a given document including Attentive Reader (Kadlec et al., 2016), BiDAF (Seo et al., 2017), Interactive AoA Reader (Cui et al., 2017), Gated Attention Reader (Dhingra et al., 2017), R-Net (Wang et al., 2017a), DCN (Xiong et al., 2017), QANet (Yu et al., 2018), and achieve promising results in most existing public MRC

datasets.

Despite the success of neural network-based MRC models, researchers began to analyze the data and rethink to what extent we have solved the problem of MRC. Some works (Chen et al., 2016; Sugawara et al., 2018; Kaushik and Lipton, 2018) classify the reasoning skills required to answer the questions into the following types: (1) Exact matching/Paraphrasing; (2) Summary; (3) Logic reasoning; (4) Utilizing external knowledge; (5) Numerical reasoning. They found that most existing MRC models are focusing on dealing with the first three types of questions. However, all these models suffer from problems when answering the questions requiring numerical reasoning. To the best of our knowledge, our work is the first one that *explicitly* incorporates numerical reasoning into the MRC system. The most relevant work to ours is NAQANet (Dua et al., 2019), which adapts the output layer of QANet (Yu et al., 2018) to support predicting answers based on counting and addition/subtraction over numbers. However, it does not consider numerical reasoning explicitly during encoding or inference.

2.2 Arithmetic Word Problem Solving

Recently, understanding and solving arithmetic word problems (AWP) has attracted the growing interest of NLP researchers. Hosseini et al. (2014) proposed a simple method to address arithmetic word problems, but mostly focusing on subsets of problems which only require addition and subtraction. After that, Roy and Roth (2015) proposed an algorithmic approach which could handle arithmetic word problems with multiple steps and operations. Koncel-Kedziorski et al. (2015) further

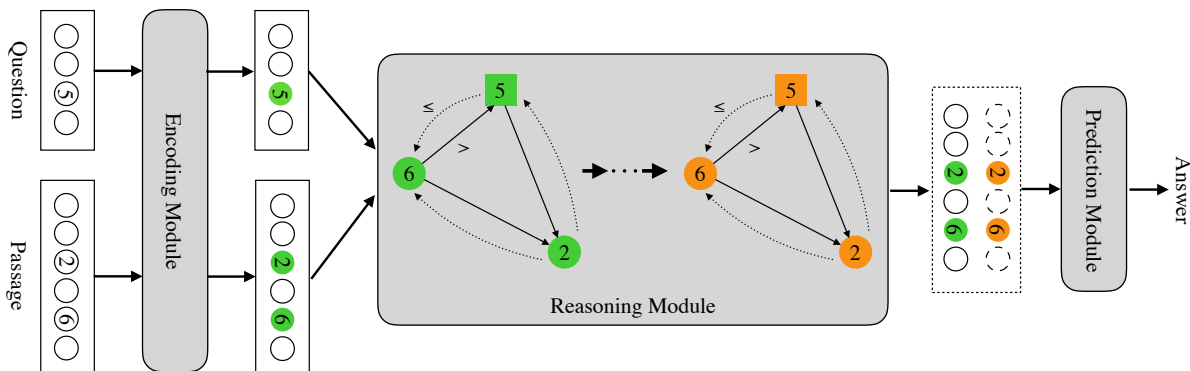


Figure 1: **The framework of our NumNet model.** Our model consists of an encoding module, a reasoning module and a prediction module. The numerical relations between numbers are encoded with the topology of the graph. For example, the edge pointing from “6” to “5” denotes “6” is greater than “5”. And the reasoning module leverages a numerically-aware graph neural network to perform numerical reasoning on the graph. As numerical comparison is modeled explicitly in our model, it is more effective for answering questions requiring numerical reasoning such as addition, counting, or sorting over numbers.

formalized the AWP problem as that of generating and scoring equation trees via integer linear programming. Wang et al. (2017b) and Ling et al. (2017) proposed sequence to sequence solvers for the AWP problems, which are capable of generating unseen expressions and do not rely on sophisticated manual features. Wang et al. (2018) leveraged deep Q-network to solve the AWP problems, achieving a good balance between effectiveness and efficiency. However, all the existing AWP systems are only trained and validated on small benchmark datasets. Huang et al. (2016) found that the performance of these AWP systems sharply degrades on larger datasets. Moreover, from the perspective of NLP, MRC problems are more challenging than AWP since the passages in MRC are mostly real-world texts which require more complex skills to be understood. Above all, it is nontrivial to adapt most existing AWP models to the MRC scenario. Therefore, we focus on enhancing MRC models with numerical reasoning abilities in this work.

3 Methodology

In this section, we will introduce the framework of our model NumNet and provide the details of the proposed numerically-aware graph neural network (NumGNN) for numerical reasoning.

3.1 Framework

An overview of our model NumNet is shown in Figure 1. We compose our model with encoding module, reasoning module and prediction module. Our major contribution is the reasoning module,

which leverages a NumGNN between the encoding module and prediction module to explicitly consider the numerical comparison information and perform numerical reasoning. As NAQANet has been shown effective for handling numerical MRC problem (Dua et al., 2019), we leverage it as our base model and mainly focus on the design and integration of the NumGNN in this work.

Encoding Module Without loss of generality, we use the encoding components of QANet and NAQANet to encode the question and passage into vector-space representations. Formally, the question Q and passage P are first encoded as:

$$Q = \text{QANet-Emb-Enc}(Q), \quad (1)$$

$$P = \text{QANet-Emb-Enc}(P), \quad (2)$$

and then the passage-aware question representation and the question-aware passage representation are computed as:

$$\bar{Q} = \text{QANet-Att}(P, Q), \quad (3)$$

$$\bar{P} = \text{QANet-Att}(Q, P), \quad (4)$$

where $\text{QANet-Emb-Enc}(\cdot)$ and $\text{QANet-Att}(\cdot)$ denote the “stacked embedding encoder layer” and “context-query attention layer” of QANet respectively. The former consists of convolution, self-attention and feed-forward layers. The latter is a passage-question attention layer. \bar{Q} and \bar{P} are used by the following components.

Reasoning Module First we build a heterogeneous directed graph $\mathcal{G} = (V; E)$, whose nodes

(V) are corresponding to the numbers in the question and passage, and edges (E) are used to encode numerical relationships among the numbers. The details will be explained in Sec. 3.2.

Then we perform reasoning on the graph based on a graph neural network, which can be formally denoted as:

$$M^Q = \text{QANet-Mod-Enc}(\mathbf{W}^M \bar{Q}), \quad (5)$$

$$M^P = \text{QANet-Mod-Enc}(\mathbf{W}^M \bar{P}), \quad (6)$$

$$U = \text{Reasoning}(\mathcal{G}; M^Q, M^P), \quad (7)$$

where \mathbf{W}^M is a shared weight matrix, U is the representations of the nodes corresponding to the numbers, $\text{QANet-Mod-Enc}(\cdot)$ is the ‘‘model encoder layer’’ defined in QANet which is similar to $\text{QANet-Emb-Enc}(\cdot)$, and the definition of $\text{Reasoning}(\cdot)$ will be given in Sec. 3.3.

Finally, as U only contains the representations of numbers, to tackle span-style answers containing non-numerical words, we concatenate U with M^P to produce numerically-aware passage representation M_0 . Formally,

$$M_0^{\text{num}}[i] = \begin{cases} U[I(i)] & \text{if } w_i^p \text{ is a number} \\ \mathbf{0} & \end{cases},$$

$$M'_0 = \mathbf{W}_0[M^P; M^{\text{num}}] + \mathbf{b}_0, \quad (8)$$

$$M_0 = \text{QANet-Mod-Enc}(M'_0), \quad (9)$$

where $[\cdot; \cdot]$ denotes matrix concatenation, $\mathbf{W}[k]$ denotes the k -th column of a matrix \mathbf{W} , $\mathbf{0}$ is a zero vector, $I(i)$ denotes the node index corresponding to the passage word w_i^p which is a number, \mathbf{W}_0 is a weight matrix, and \mathbf{b}_0 is a bias vector.

Prediction Module Following NAQANet (Dua et al., 2019), we divide the answers into four types and use a unique output layer to calculate the conditional answer probability $\text{Pr}(\text{answer}|\text{type})$ for each type :

- *Passage span*: The answer is a span of the passage, and the answer probability is defined as the product of the probabilities of the start and end positions.
- *Question span*: The answer is a span of the question, and the answer probability is also defined as the product of the probabilities of the start and end positions.
- *Count*: The answer is obtained by counting, and it is treated as a multi-class classification problem over ten numbers (0-9), which

covers most of the *Count* type answers in the DROP dataset.

- *Arithmetic expression*: The answer is the result of an arithmetic expression. The expression is obtained in three steps: (1) extract all numbers from the passage; (2) assign a sign (plus, minus or zero) for each number; (3) sum the signed numbers¹.

Meanwhile, an extra output layer is also used to predict the probability $\text{Pr}(\text{type})$ of each answer type. At training time, the final answer probability is defined as the joint probability over all feasible answer types, i.e., $\sum_{\text{type}} \text{Pr}(\text{type}) \text{Pr}(\text{answer}|\text{type})$. Here, the answer type annotation is not required and the probability $\text{Pr}(\text{type})$ is learnt by the model. At test time, the model first selects the most probable answer type greedily and then predicts the best answer accordingly.

Without loss of generality, we leverage the definition of the five output layers in (Dua et al., 2019), with M_0 and Q as inputs. Please refer to the paper for more details due to space limitation.

Comparison with NAQANet The major difference between our model and NAQANet is that NAQANet does not have the reasoning module, i.e., M_0 is simply set as M^P . As a result, numbers are treated as common words in NAQANet except in the prediction module, thus NAQANet may struggle to learn the numerical relationships between numbers, and potentially cannot well generalize to unseen numbers. However, as discussed in Sec. 1, the numerical comparison is essential for answering questions requiring numerical reasoning. In our model, the numerical relationships are explicitly represented with the topology of the graph and a NumGNN is used to perform numerical reasoning. Therefore, our NumNet model can handle questions requiring numerical reasoning more effectively, which is verified by the experiments in Sec. 4.

3.2 Numerically-aware Graph Construction

We regard all numbers from the question and passage as nodes in the graph for reasoning². The

¹As few samples require multiplication/division expression in the DROP dataset, we simply adapt the module proposed (Dua et al., 2019) and leave multiplication/division expression handling as future work.

²As a number in the question may serve as a critical comparison condition (refer to the second example in Table 1),

set of nodes corresponding to the numbers occurring in question and passage are denoted as \mathbf{V}^Q and \mathbf{V}^P respectively. And we denote all the nodes as $\mathbf{V} = \mathbf{V}^Q \cup \mathbf{V}^P$, and the number corresponding to a node $v \in \mathbf{V}$ as $n(v)$.

Two sets of edges are considered in this work:

- **Greater Relation Edge (\vec{E}):** For two nodes $v_i, v_j \in \mathbf{V}$, a directed edge $\vec{e}_{ij} = (v_i, v_j)$ pointing from v_i to v_j will be added to the graph if $n(v_i) > n(v_j)$, which is denoted as solid arrow in Figure 1.
- **Lower or Equal Relation Edge (\overleftarrow{E}):** For two nodes $v_i, v_j \in \mathbf{V}$, a directed edge $\overleftarrow{e}_{ij} = (v_j, v_i)$ will be added to the graph if $n(v_i) \leq n(v_j)$, which is denoted as dashed arrow in Figure 1.

Theoretically, \vec{E} and \overleftarrow{E} are complement to each other. However, as a number may occur several times and represent different facts in a document, we add a distinct node for each occurrence in the graph to prevent potential ambiguity. Therefore, it is more reasonable to use both \vec{E} and \overleftarrow{E} in order to encode the equal information among nodes.

3.3 Numerical Reasoning

As we built the graph $\mathcal{G} = (\mathbf{V}, \mathbf{E})$, we leverage NumGNN to perform reasoning, which is corresponding to the function $\text{Reasoning}(\cdot)$ in Eq. 7. The reasoning process is as follows:

Initialization For each node $v_i^P \in \mathbf{V}^P$, its representation is initialized as the corresponding column vector of \mathbf{M}^P . Formally, the initial representation is $\mathbf{v}_i^P = \mathbf{M}^P[I^P(v_i^P)]$, where $I^P(v_i^P)$ denotes the word index corresponding to v_i^P . Similarly, the initial representation \mathbf{v}_j^Q for a node $v_j^Q \in \mathbf{V}^Q$ is set as the corresponding column vector of \mathbf{M}^Q . We denote all the initial node representations as $\mathbf{v}^0 = \{\mathbf{v}_i^P\} \cup \{\mathbf{v}_j^Q\}$.

One-step Reasoning Given the graph \mathcal{G} and the node representations \mathbf{v} , we use a GNN to perform reasoning in three steps:

(1) **Node Relatedness Measure:** As only a few numbers are relevant for answering a question generally, we compute a weight for each node to by-pass irrelevant numbers in reasoning. Formally, the weight for node v_i is computed as:

$$\alpha_i = \text{sigmoid}(\mathbf{W}_v \mathbf{v}[i] + b_v), \quad (10)$$

we also add nodes for them in the graph.

where \mathbf{W}_v is a weight matrix, and b_v is a bias.

(2) **Message Propagation:** As the role a number plays in reasoning is not only decided by itself, but also related to the context, we propagate messages from each node to its neighbors to help to perform reasoning. As numbers in question and passage may play different roles in reasoning and edges corresponding to different numerical relations should be distinguished, we use relation-specific transform matrices in the message propagation. Formally, we define the following propagation function for calculating the forward-pass update of a node:

$$\tilde{\mathbf{v}}'_i = \frac{1}{|\mathcal{N}_i|} \left(\sum_{j \in \mathcal{N}_i} \alpha_j \mathbf{W}^{\mathcal{r}_{ji}} \mathbf{v}[j] \right), \quad (11)$$

where $\tilde{\mathbf{v}}'_i$ is the message representation of node v_i , \mathcal{r}_{ji} is the relation assigned to edge e_{ji} , $\mathbf{W}^{\mathcal{r}_{ji}}$ are relation-specific transform matrices, and $\mathcal{N}_i = \{j | (v_j, v_i) \in \mathbf{E}\}$ is the neighbors of node v_i .

For each edge e_{ji} , \mathcal{r}_{ji} is determined by the following two attributes:

- Number relation: $>$ or \leq ;
- Node types: the two nodes of the edge corresponding to two numbers that: (1) both from the question (q-q); (2) both from the passage (p-p); (3) from the question and the passage respectively (q-p); (4) from the passage and the question respectively (p-q).

Formally, $\mathcal{r}_{ij} \in \{>, \leq\} \times \{q-q, p-p, q-p, p-q\}$.

(3) **Node Representation Update:** As the message representation obtained in the previous step only contains information from the neighbors, it needs to be fused with the node representation to combine with the information carried by the node itself, which is performed as:

$$\mathbf{v}'_i = \text{ReLU}(\mathbf{W}_f \mathbf{v}_i + \tilde{\mathbf{v}}'_i + \mathbf{b}_f), \quad (12)$$

where \mathbf{W}_f is a weight matrix, and \mathbf{b}_f is a bias vector.

We denote the entire one-step reasoning process (Eq. 10-12) as a single function

$$\mathbf{v}' = \text{Reasoning-Step}(\mathcal{G}, \mathbf{v}). \quad (13)$$

As the graph \mathcal{G} constructed in Sec. 3.2 has encoded the numerical relations via its topology, the reasoning process is numerically-aware.

Multi-step Reasoning By single-step reasoning, we can only infer relations between adjacent nodes. However, relations between multiple nodes may be required for certain tasks, e.g., sorting. Therefore, it is essential to perform multi-step reasoning, which can be done as follows

$$v^t = \text{Reasoning-Step}(v^{t-1}), \quad (14)$$

where $t \geq 1$. Suppose we perform K steps of reasoning, v^K is used as U in Eq. 7.

4 Experiments

4.1 Dataset and Evaluation Metrics

We evaluate our proposed model on DROP dataset (Dua et al., 2019), which is a public numerical MRC dataset. The DROP dataset is constructed by crowd-sourcing, which asks the annotators to generate question-answer pairs according to the given Wikipedia passages, which require numerical reasoning such as addition, counting, or sorting over numbers in the passages. There are 77,409 training samples, 9,536 development samples and 9,622 testing samples in the dataset.

In this paper, we adopt two metrics including Exact Match (EM) and numerically-focused F1 scores to evaluate our model following Dua et al. (2019). The numerically-focused F1 is set to be 0 when the predicted answer is mismatched for those questions with the numeric golden answer.

4.2 Baselines

For comparison, we select several public models as baselines including **semantic parsing models**:

- Syn Dep (Dua et al., 2019), the neural semantic parsing model (KDG) (Krishnamurthy et al., 2017) with Stanford dependencies based sentence representations;
- OpenIE (Dua et al., 2019), KDG with open information extraction based sentence representations;
- SRL (Dua et al., 2019), KDG with semantic role labeling based sentence representations;

and **traditional MRC models**:

- BiDAF (Seo et al., 2017), an MRC model which utilizes a bi-directional attention flow network to encode the question and passage;
- QANet (Yu et al., 2018), which utilizes convolutions and self-attentions as the building blocks of encoders to represent the question and passage;

- BERT (Devlin et al., 2019), a pre-trained bidirectional Transformer-based language model which achieves state-of-the-art performance on lots of public MRC datasets recently;

and **numerical MRC models**:

- NAQANet (Dua et al., 2019), a numerical version of QANet model.
- NAQANet+, an enhanced version of NAQANet implemented by ourselves, which further considers real number (e.g. “2.5”), richer arithmetic expression, data augmentation, etc. The enhancements are also used in our NumNet model and the details are given in the supplemental material.

4.3 Experimental Settings

In this paper, we tune our model on the development set and use a grid search to determine the optimal parameters. The dimensions of all the representations (e.g., Q , P , M^Q , M^P , U , M'_0 , M_0 and v) are set to 128. If not specified, the reasoning step K is set to 3. Since other parameters have little effect on the results, we simply follow the settings used in (Dua et al., 2019).

We use the Adam optimizer (Kingma and Ba, 2015) with $\beta_1 = 0.8$, $\beta_2 = 0.999$, $\epsilon = 10^{-7}$ to minimize the objective function. The learning rate is 5×10^{-4} , L2 weight decay λ is 10^{-7} and the maximum norm value of gradient clipping is 5. We also apply exponential moving average with a decay rate 0.9999 on all trainable variables. The model is trained with a batch size of 16 for 40 epochs. Passages and questions are trimmed to 400 and 50 tokens respectively during training, and trimmed to 1,000 and 100 tokens respectively during prediction³.

4.4 Overall Results

The performance of our NumNet model and other baselines on DROP dataset are shown in Table 2. From the results, we can observe that:

(1) Our NumNet model achieves better results on both the development and testing sets on DROP dataset as compared to semantic parsing-based models, traditional MRC models and even numerical MRC models NAQANet and NAQANet+. The reason is that our NumNet model can make full use of the numerical comparison information over

³Trimming passages/questions introduces little impact because it only affects about 8% of the samples.

Method	Dev		Test	
	EM	F1	EM	F1
Semantic Parsing				
Syn Dep	9.38	11.64	8.51	10.84
OpenIE	8.80	11.31	8.53	10.77
SRL	9.28	11.72	8.98	11.45
Traditional MRC				
BiDAF	26.06	28.85	24.75	27.49
QANet	27.50	30.44	25.50	28.36
BERT	30.10	33.36	29.45	32.70
Numerical MRC				
NAQANet	46.20	49.24	44.07	47.01
NAQANet+	61.47	64.85	60.82	64.29
NumNet	64.92	68.31	64.56	67.97
Human Performance				
	-	-	94.09	96.42

Table 2: Overall results on the development and test set. The evaluation metrics are calculated as the maximum over a golden answer set. All the results except “NAQANet+” and “NumNet” are obtained from (Dua et al., 2019).

numbers in both question and passage via the proposed NumGNN module.

(2) Our implemented NAQANet+ has a much better performance compared to the original version of NAQANet. It verifies the effectiveness of our proposed enhancements for baseline.

4.5 Effect of GNN Structure

In this part, we investigate the effect of different GNN structures on the DROP development set. The results are shown in Table 3. The “Comparison”, “Number” and “ALL” are corresponding to the comparing question subset⁴, the number-type answer subset, and the entire development set, respectively⁵. If we replace the proposed numerically-aware graph (Sec. 3.2) with a fully connected graph, our model fallbacks to a traditional GNN, denoted as “GNN” in the table. Moreover, “- question num” denotes the numbers in the question is not included in the graph, and “- \leq type edge” and “- $>$ type edge” denote edges of \leq and $>$ types are not adopted respectively.

⁴We find that many comparing questions in the DROP dataset are biased, of which the answers are the former candidates in the questions. Hence, we employ crowdsourced workers to identify and rewrite all comparing questions to construct an enhanced development set. Specially, for those comparing questions containing answer candidates, we also ask the crowdsourced workers to swap the candidates manually to enlarge the dataset.

⁵Note that the “ALL” result is not the average of “Comparison” and “Number”. It is the performance on the entire development set which also includes questions of selection type, coreference resolution type, etc.

Method	Comparison		Number		ALL	
	EM	F1	EM	F1	EM	F1
GNN	69.86	75.91	67.77	67.78	61.90	65.16
NumGNN	74.53	80.36	69.74	69.75	64.54	68.02
- question num	74.84	80.24	68.42	68.43	63.78	67.17
- \leq type edge	74.89	80.51	68.48	68.50	63.66	67.06
- $>$ type edge	74.86	80.19	68.77	68.78	63.64	66.96

Table 3: Performance with different GNN structure. “Comparison”, “Number” and “ALL” denote the comparing question subset, the number-type answer subset, and the entire development set, respectively.

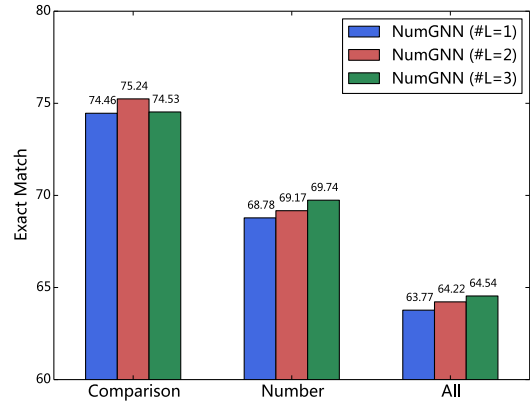


Figure 2: Effect of GNN layer numbers (#L).

As shown in Table 3, our proposed NumGNN leads to statistically significant improvements compared to traditional GNN on both EM and F1 scores especially for comparing questions. It indicates that considering the comparing information over numbers could effectively help the numerical reasoning for comparing questions. Moreover, we find that the numbers in the question are often related to the numerical reasoning for answering the question, thus considering numbers in questions in NumGNN achieves better performance. And the results also justify that encoding “greater relation” and “lower or equal relation” simultaneously in the graph also benefits our model.

4.6 Effect of GNN Layer Number

The number of NumGNN layers represents the numerical reasoning ability of our models. A K -layer version has the ability for K -step numerical inference. In this part, we additionally perform experiments to understand the values of the numbers of NumGNN layers. From Figure 2, we could observe that:

(1) The 2-layer version of NumNet achieves the best performance for the comparing questions. From careful analysis, we find that most compar-

Question & Answer	Passage	NAQANet+	NumNet
Q: Which age group is larger: under the age of 18 or 18 and 24? A: 18 and 24	The median age in the city was 22.1 years. <i>10.1%</i> of residents were under the age of 18; <i>56.2%</i> were between the ages of 18 and 24; 16.1% were from 25 to 44; 10.5% were from 45 to 64; and 7% were 65 years of age or older. The gender makeup of the city was 64.3% male and 35.7% female.	under the age of 18	18 and 24
Q: How many more yards was Longwell’s longest field goal over his second longest one? A: 26-22=4	... The Vikings would draw first blood with a <i>26-yard field goal</i> by kicker Ryan Longwell. In the second quarter, Carolina got a field goal with opposing kicker John Kasay. The Vikings would respond with another Longwell field goal (<i>a 22-yard FG</i>) ... In OT, Longwell booted the game-winning <i>19-yard field goal</i> to give Minnesota the win. It was the first time in Vikings history that a coach ...	26-19 = 7	26-22 = 4

Table 4: Cases from the DROP dataset. We demonstrate the predictions of NAQANet+ and our NumNet model. Note that the two models only output the arithmetic expressions but we also provide their results for clarity.

Question	Passage	Answer	NumNet
Which ancestral groups are at least 10%?	As of the census of 2000, there were 7,791 people, 3,155 households, and 2,240 families residing in the county. ... 33.7% were of <i>Germans</i> , 13.9% <i>Swedish</i> people, 10.1% <i>Irish</i> people, 8.8% United States, 7.0% English people and 5.4% Danish people ancestry ...	German; Swedish; Irish	Irish
Were more people 40 and older or 19 and younger?	Of Saratoga Countys population in 2010, <i>6.3%</i> were between ages of 5 and 9 years, <i>6.7%</i> between 10 and 14 years, <i>6.5%</i> between 15 and 19 years, ... , <i>7.9%</i> between 40 and 44 years, <i>8.5%</i> between 45 and 49 years, <i>8.0%</i> between 50 and 54 years, <i>7.0%</i> between 55 and 59 years, <i>6.4%</i> between 60 and 64 years, and <i>13.7%</i> of age 65 years and over ...	40 and older	19 and younger

Table 5: Typical error examples. Row 1: the answer is multiple nonadjacent spans; Row 2: Intermediate numbers are involved in reasoning.

ing questions only require at most 2-step reasoning (e.g., “*Who was the second oldest player in the MLB, Clemens or Franco?*”), and therefore the 3-layer version of NumNet is more complex but brings no gains for these questions.

(2) The performance of our NumNet model on the overall development set is improved consistently as the number of GNN layers increases. The reason is that some of the numerical questions require reasoning over many numbers in the passage, which could benefit from the multi-step reasoning ability of multi-layer GNN. However, further investigation shows that the performance gain is not stable when $K \geq 4$. We believe it is due to the intrinsic over smoothing problem of GNNs (Li et al., 2018).

4.7 Case Study

We further give some examples to show why incorporating comparing information over numbers in the passage could help numerical reasoning in MRC in Table 4. For the first case, we observe that NAQANet+ gives a wrong prediction, and we find that NAQANet+ will give the same prediction for the question “*Which age group is smaller: under the age of 18 or 18 and 24?*”. The reason is

that NAQANet+ cannot distinguish which one is larger for 10.1% and 56.2%. For the second case, NAQANet+ cannot recognize the second longest field goal is 22-yard and also gives a wrong prediction. For these two cases, our NumNet model could give the correct answer through the numeric reasoning, which indicates the effectiveness of our NumNet model.

4.8 Error Analysis

To investigate how well our NumNet model handles sorting/comparison questions and better understand the remaining challenges, we perform an error analysis on a random sample of NumNet predictions. We find that:

(1) Our NumNet model can answer about 76% of sorting/comparison questions correctly, which indicates that our NumNet model has achieved numerical reasoning ability to some extent.

(2) Among the incorrectly answered sorting/comparison questions, the most ones (26%) are those whose golden answers are multiple non-adjacent spans (row 1 in Table 5), and the second most ones (19%) are those involving comparison with an intermediate number that does not literally occur in the document/question but has to be de-

rived from counting or arithmetic operation (row 1 in Table 5).

4.9 Discussion

By combining the numerically-aware graph and the NumGNN together, our NumNet model achieves the numerical reasoning ability. On one hand, the numerically-aware graph encodes numbers as nodes and relationships between them as the edges, which is required for numerical comparison. On the other hand, through one-step reasoning, our NumGNN could perform comparison and identify the numerical condition. After multiple-step reasoning, our NumGNN could further perform sorting.

However, since the numerically-aware graph is pre-defined, our NumNet is not applicable to the case where an intermediate number has to be derived (e.g., from arithmetic operation) in the reasoning process, which is a major limitation of our model.

5 Conclusion and Future Work

Numerical reasoning skills such as addition, subtraction, sorting and counting are naturally required by machine reading comprehension (MRC) problems in practice. Nevertheless, these skills are not taken into account explicitly for most existing MRC models. In this work, we propose a numerical MRC model named NumNet which performs explicit numerical reasoning while reading the passages. To be specific, NumNet encodes the numerical relations among numbers in the question and passage into a graph as its topology, and leverages a numerically-aware graph neural network to perform numerical reasoning on the graph. Our NumNet model outperforms strong baselines with a large margin on the DROP dataset.

In the future, we will explore the following directions: (1) As we use a pre-defined reasoning graph in our model, it is incapable of handling reasoning process which involves intermediate numbers that not presented in the graph. How to incorporate dynamic graph into our model is an interesting problem. (2) Compared with methods proposed for arithmetic word problems (AWPs), our model has better natural language understanding ability. However, the methods for AWPs can handle much richer arithmetic expressions. Therefore, how to combine both of their abilities to develop a more powerful numerical MRC model is

an interesting future direction. (3) Symbolic reasoning plays a crucial role in human reading comprehension. Our work integrates numerical reasoning, which is a special case of symbolic reasoning, into traditional MRC systems. How to incorporate more sophisticated symbolic reasoning abilities into MRC systems is also a valuable future direction.

Acknowledgments

We would like to thank all anonymous reviewers for their insightful comments, and thank Yan Zhang for her help on improving the presentation of Figure 1.

References

- Danqi Chen, Jason Bolton, and Christopher D. Manning. 2016. [A thorough examination of the CNN/Daily Mail reading comprehension task](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2358–2367, Berlin, Germany.
- Yiming Cui, Zhipeng Chen, Si Wei, Shijin Wang, Ting Liu, and Guoping Hu. 2017. [Attention-over-attention neural networks for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 593–602, Vancouver, Canada.
- 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, Minneapolis, Minnesota.
- Bhuwan Dhingra, Hanxiao Liu, Zhilin Yang, William Cohen, and Ruslan Salakhutdinov. 2017. [Gated-attention readers for text comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1832–1846, Vancouver, Canada.
- 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.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Su-

- leyman, and Phil Blunsom. 2015. [Teaching machines to read and comprehend](#). In *Proceedings of Advances in Neural Information Processing Systems*, pages 1693–1701.
- Mohammad Javad Hosseini, Hannaneh Hajishirzi, Oren Etzioni, and Nate Kushman. 2014. [Learning to solve arithmetic word problems with verb categorization](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 523–533, Doha, Qatar.
- Danqing Huang, Shuming Shi, Chin-Yew Lin, Jian Yin, and Wei-Ying Ma. 2016. [How well do computers solve math word problems? Large-scale dataset construction and evaluation](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 887–896, Berlin, Germany.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611, Vancouver, Canada.
- Rudolf Kadlec, Martin Schmid, Ondřej Bajgar, and Jan Kleindienst. 2016. [Text understanding with the attention sum reader network](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 908–918, Berlin, Germany.
- 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, Brussels, Belgium.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of ICLR 2015 : International Conference on Learning Representations 2015*.
- Rik Koncel-Kedziorski, Hannaneh Hajishirzi, Ashish Sabharwal, Oren Etzioni, and Siena Dumas Ang. 2015. [Parsing algebraic word problems into equations](#). *Transactions of the Association for Computational Linguistics*, 3:585–597.
- Jayant Krishnamurthy, Pradeep Dasigi, and Matt Gardner. 2017. [Neural semantic parsing with type constraints for semi-structured tables](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1516–1526, Copenhagen, Denmark.
- 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, Copenhagen, Denmark.
- Qimai Li, Zhichao Han, and Xiao-Ming Wu. 2018. Deeper insights into graph convolutional networks for semi-supervised learning. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, pages 3538–3545.
- Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. 2017. Program induction by rationale generation: Learning to solve and explain algebraic word problems. *arXiv preprint arXiv:1705.04146*.
- 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, Austin, Texas.
- Subhro Roy and Dan Roth. 2015. [Solving general arithmetic word problems](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1743–1752, Lisbon, Portugal.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension. In *Proceedings of ICLR 2017: the 5th International Conference on Learning Representations*.
- Saku Sugawara, Kentaro Inui, Satoshi Sekine, and Akiko Aizawa. 2018. [What makes reading comprehension questions easier?](#) In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4208–4219, Brussels, Belgium.
- Lei Wang, Dongxiang Zhang, Lianli Gao, Jingkuan Song, Long Guo, and Heng Tao Shen. 2018. MathDQN: Solving arithmetic word problems via deep reinforcement learning. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)*.
- Wenhui Wang, Nan Yang, Furu Wei, Baobao Chang, and Ming Zhou. 2017a. [Gated self-matching networks for reading comprehension and question answering](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 189–198, Vancouver, Canada.
- 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, Copenhagen, Denmark.
- Caiming Xiong, Victor Zhong, and Richard Socher. 2017. Dynamic coattention networks for question answering. In *Proceedings of ICLR 2017: the 5th International Conference on Learning Representations*.

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. In *Proceedings of ICLR 2018: the 6th International Conference on Learning Representations*.