

GeoQA: A Geometric Question Answering Benchmark Towards Multimodal Numerical Reasoning

Jiaqi Chen^{1*}, Jianheng Tang^{2*}, Jinghui Qin¹, Xiaodan Liang^{2†},
Lingbo Liu^{1,3}, Eric P. Xing⁴, Liang Lin^{1,3}

¹Sun Yat-sen University, ²Shenzhen Campus of Sun Yat-sen University, ³Dark Matter AI Inc.,

⁴Mohamed bin Zayed University of Artificial Intelligence

{jadagechen, sqrt3tjh, xdliang328, liulingbo918}@gmail.com,

qinjinhg@mail2.sysu.edu.cn, eric.xing@petuum.com, linliang@ieee.org

Abstract

Automatic math problem solving has recently attracted increasing attention as a long-standing AI benchmark. In this paper, we focus on solving geometric problems, which requires a comprehensive understanding of textual descriptions, visual diagrams, and theorem knowledge. However, the existing methods were highly dependent on handcraft rules and were merely evaluated on small-scale datasets. Therefore, we propose a **Geometric Question Answering** dataset **GeoQA**, containing 5,010 geometric problems with corresponding annotated programs, which illustrate the solving process of the given problems. Compared with another publicly available dataset GeoS, GeoQA is 25 times larger, in which the program annotations can provide a practical testbed for future research on explicit and explainable numerical reasoning. Moreover, we introduce a Neural Geometric Solver (NGS) to address geometric problems by comprehensively parsing multimodal information and generating interpretable programs. We further add multiple self-supervised auxiliary tasks on NGS to enhance cross-modal semantic representation. Extensive experiments on GeoQA validate the effectiveness of our proposed NGS and auxiliary tasks. However, the results are still significantly lower than human performance, which leaves large room for future research. Our benchmark and code are released at <https://github.com/chenjudge/GeoQA>.

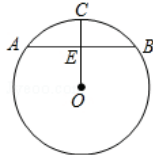
1 Introduction

In recent years, developing machine learning systems to solve math word problems (MWPs) automatically has attracted increasing attention due to its high academic value and the great application potential in smart education (Bajaj and Sharma,

*Equal contribution.

†Corresponding author.

As shown in the figure, in $\odot O$, AB is the chord, $OC \perp AB$, if the radius of $\odot O$ is 5 (N0) and $CE=2$ (N1), then the length of AB is ()



A. 2 B. 4 C. 6 D. 8
Answer: D. 8

Problem Type: Length Calculation
Knowledge Points: Vertical Diameter, Pythagorean Theorem
Problem Solving Explanations:
 $OE=OC-CE=5-2=3$. According to the Pythagorean Theorem,
 $AE = \sqrt{OA^2 - OE^2} = \sqrt{5^2 - 3^2} = 4$. Thus, $AB=2AE=8$.

Annotated Programs:

```
Minus | N0 | N1 | PythagoreanMinus | N0 | V0 | Double | V1
```

Step1: Minus(N0, N1) = 5 - 2 = 3 (V0)
Step2: PythagoreanMinus(N0, V0) = $\sqrt{5^2 - 3^2} = 4$ (V1)
Step3: Double(V1) = 2 * 4 = 8 (V2)

Figure 1: Illustration of a typical geometry problem with the annotated programs in our GeoQA dataset.

2018; Lin et al., 2018). Most of the existing methods focus on solving arithmetic and algebraic problems, including traditional machine learning approaches (Kushman et al., 2014; Zhou et al., 2015; Huang et al., 2016) and network-based models (Wang et al., 2017, 2018; Xie and Sun, 2019), while solving geometric problems has been rarely investigated (Seo et al., 2014, 2015a; Sachan et al., 2017). As a classic math problem, geometry dominates a large portion of secondary education. Due to its challenges and data characteristics, geometry problem can also serve as a multimodal numerical reasoning benchmark requiring joint reasoning over diagram and text.

As shown in Figure 1, a typical geometric question mainly consists of textual descriptions and geometric diagrams. Compared with math word problems, which only involve problem texts, geometric questions have posed the following new challenges. **First**, the additional problem diagrams provide essential information absent from the problem text,

such as the relative location of lines and points; thus, the solver should have the capability to parse the diagram. **Second**, to solve a geometry problem, we need to understand and align the semantics of text and diagram simultaneously. However, the problem text often includes some ambiguous references and implicit relations to diagram elements, which increases the difficulty of joint reasoning over text and diagram. **Third**, many geometric problems require extra theorem knowledge in the problem solving process. For example, in Figure 1, the Pythagorean Theorem is used to calculate the length of line AE .

Though some previous methods (Seo et al., 2014, 2015a; Sachan et al., 2017, 2020; Sachan, 2020) attempt to resolve the mentioned issues, the performance of their geometric problem solving systems is far away from satisfactory. They highly depended on limited handcraft rules and were only validated on small-scale datasets, making it hard to generalize to more complex and real-world cases. Besides, the solving process is sophisticated, which means it is difficult for a human to understand and examine its reliability.

To refresh the research on geometric problem solving and promote further study on multimodal numerical reasoning, we propose a large-scale real-world geometric question answering dataset called GeoQA, which contains 5,010 multiple-choice geometric problems collected from real math exams in Chinese middle school. Inspired by Amini et al. (2019), we additionally introduce a new domain-specific language to model precise operation programs corresponding to the geometry problem. These executable programs represent the numerical reasoning steps of geometry problems. Compared with the existing dataset GeoS and GeoS++ (Seo et al., 2015a; Sachan et al., 2017), our GeoQA is larger, more diverse, provides additional program annotation, thus serves as a promising benchmark to improve both generalization and interpretability of the multimodal numerical reasoning approaches.

Moreover, we propose the first deep learning-based approach for geometry problem solving, named as Neural Geometric Solver (NGS). It applies a co-attention mechanism to fuse the representation of text and diagram, and predicts the explainable programs based on the cross-modal representation. These sequential programs can be executed to obtain a final answer. Benefiting from the structured and explainable program prediction,

our NGS has both merits in superior performance by learning-based models compared to previous rule-based methods, as well as generating explainable numerical reasoning steps via the program sequence in favor of the model diagnosis. We further design three highly-relevant pretext tasks to enhance text-diagram semantic representation, including diagram jigsaw location prediction, diagram geometric element prediction, and knowledge point prediction. Extensive experiments are conducted on GeoQA benchmark, and the quantitative comparisons show the superiority of the proposed NGS and auxiliary tasks over other multimodal baselines.

In summary, our contributions are three-fold:

- We construct a large-scale dataset for geometric problem solving, which contains 5,010 Chinese geometric multiple-choice questions with rich domain-specific program annotations.
- A novel Neural Geometric Solver is proposed to solve geometric problems by generating symbolic programs based on the joint understanding of textual descriptions and diagrams.
- Multiple specialized auxiliary tasks are employed to effectively improve the semantic representation of text and diagrams. Experiments show the superiority of our NGS equipped with these auxiliary tasks.

2 Related Work

Geometry Problems Solving Developing automated systems to solve geometry problems has a long history in AI (Gelernter et al., 1960; Wen-Tsun, 1986; Chou et al., 1996; Ye et al., 2008). For example, Seo et al. (2014, 2015a) built the first automated system, GeoS, to solve SAT style geometry problems. GeoS used NLP and computer vision techniques (e.g., OCR) to parse a geometry problem’s text and diagram jointly as logic forms. However, this system highly depended on the manually designed logic forms and was only examined in a small dataset with 185 problems. Besides, the limited logic forms are hard to cover various geometry problems, leading to low generalization. To improve GeoS, Sachan et al. (2017); Sachan and Xing (2017) replaced these handcraft constraints with geometry axiomatic knowledge in the form of horn-clause rules, but their dataset and code are not released. To boost the generalization and interpretability of existing works, we propose a large-scale GeoQA benchmark, which is 25 times larger

than the only public dataset Seo et al. (2015a), and provides program annotation.

Multimodal Reasoning Visual question answering is a representative multimodal task that requires the model to have reasoning ability (Goyal et al., 2017; Yu et al., 2019). Johnson et al. (2017) built a new diagnostic VQA dataset (called CLEVR) with annotated functional programs. Based on this benchmark, some methods proposed an implicit reasoning framework to jointly encode multimodal information (Perez et al., 2017; Santoro et al., 2017). Moreover, several works (Yi et al., 2018; Mao et al., 2019) utilize domain-specific languages to perform explicit symbolic reasoning. However, these program languages only consider elementary operations, such as counting objects. They are not directly applicable to geometric problems, which require multiple steps of numerical calculation and involve theorem knowledge.

Self-supervised Auxiliary Task Self-supervised pretraining has gradually emerged (Doersch and Zisserman, 2017; Newell and Deng, 2020) as a effective technique to deal with label scarcity and improve model performance. To enhance visual features, most of these methods construct pseudo labels automatically and train on auxiliary tasks, including image jigsaw (Noroozi and Favaro, 2016; Ahsan et al., 2019), inpainting (Pathak et al., 2016), super resolution (Ledig et al., 2017), etc. Inspired by these works, we design two self-supervised auxiliary tasks and a supervised auxiliary task to enhance the reasoning ability of our NGS.

3 GeoQA Dataset

Due to the limited data scale and problem types, the existing geometric problem reasoning dataset (Seo et al., 2015a) can neither comprehensively reflect model’s reasoning ability, nor support the training of neural models. To propose a better benchmark for the evaluation of multimodal numerical reasoning and inspire applications in smart education, we collect a new dataset GeoQA. It contains 5,010 diverse real-world geometric problems in Chinese middle school exams, and each problem is additionally annotated by specific programs that describe the problem solving process. Besides, we also provide human performance on GeoQA, as shown in Table 3.

	Total	Train	Val	Test
Number	5010	3509	746	755
Angle	2745	1939	388	418
Length	1873	1303	287	283
Other	392	267	71	54
#Avg DS	108×140	108×140	107×141	107×140
#Avg QL	52.5	52.4	52.4	57.7
#Avg KP	2.10	2.10	2.07	2.14
#Avg ET	1.11	1.13	1.08	1.09
#Avg OP	1.98	1.99	1.92	1.98
#Avg PL	5.35	5.39	5.17	5.36

Table 1: Statistics of our GeoQA dataset. It contains three types of problems, including angle, length, and other problems. Besides, DS, QL, KP, and ET represent diagram size, question length, knowledge points, and element types, respectively. OP and PL represent operation number and program length.

3.1 Data Description

Generally, a geometry multiple-choice problem can be represented as a tuple (t, d, c, i) where t is the problem text in natural language, d is the problem diagram, $c = (c_1, c_2, c_3, c_4)$ represents the 4 numerical options for the problem, i is the answer index. Given the text t and diagram d , an algorithm is required to predict the correct answer $c_i \in c$. To collect as much useful information as possible, we also provide the natural language-based problem solving explanations e , problem type t , the related knowledge points k , and our annotated programs p for each problem. Therefore, a geometry problem can be represented as (t, d, c, i, e, t, k, p) . Fig. 1 shows an example of geometric problems.

Moreover, there are three problem types in our GeoQA, i.e., angle calculation, length calculation, and others which contain various types of problems such as area calculation. We adopt the corpus diversity metric proposed by Miao et al. (2020) to evaluate the diversity of GeoQA. The result is 0.47, which is relatively high compared with other math problem datasets, indicating that our dataset is diverse. The source data already contains manually tagged knowledge points in each problem, we design rule-based regular expressions to normalize the original knowledge points to 50 categories. We split our GeoQA into three subsets – train set, valid set, and test set, in a ratio of 7.0: 1.5: 1.5. The data statistics of our GeoQA are shown in Table 1.

3.2 Program Representation

The neural network has proved to be a powerful tool to address complex multimodal reasoning tasks. However, it still faces challenges when

Types	Programs
Basic	Equal, Double, Half
Arithmetic	Add, Minus, Multiply, Divide
Trigonometric	Sin, Cos, Tan, Arc-Sin, Arc-Cos
Theorem & Formula	Pythagorean Add/Minus, Proportion, Circle Area, Circle Perimeter, Cone Area
Constant	30, 60, 90, 180, 360, π , 0.618

Table 2: An overview of 18 operations of four different types and 7 constants in the defined program set.

conducting numerical calculation and providing explicit problem solving process, which are actually two crucial points in the task of geometric problem solving. To make better use of neural networks in the geometric problems solving process, inspired by Amini et al. (2019), we introduce a new domain-specific language to model the geometric problem solving process based on the GeoQA dataset. This program language can be directly executed to calculate the numerical answer based on the predefined operations and their arguments.

The program types contain operations OP , constants $Const$, problem variables N , and process variables V . As shown in Table 2, operations are divided into multiple categories, including Basic, Arithmetic, Trigonometric, Theorem, and Formula operations. Each operator involves $n(= 1, 2, 3)$ elements selected from constants and variables. Constants are predefined numbers that are frequently used in geometric problems, such as π and the degree of a Right Angle (90). The problem variables refer to all the number that appears in the current problem, and process variables are obtained during the execution process.

In addition to the common math operations, our programs also contain some operations representing the knowledge of theorems and formulas that is helpful to address geometric problems, such as the Pythagorean theorem and the area calculation formula of a circle. It is worth noting that many simple geometric formulas do not require additional definitions. For example, given a square with side length a , its area can be directly computed by $\text{Multiply}(a, a)$.

The interpretability of our program is reflected on the sequential process of the operations, the selected constants and variables, and the application of theorems and formulas. As shown in Figure 1, we can have a general understanding of the entire problem solving process after reading the program.

3.3 Collection and Annotation

We collect GeoQA from two online education websites¹. These problems are oriented grades 6-12, containing various types of problems with corresponding knowledge points and solving explanations. We organize more than ten well-trained college students with a relevant major to specifically annotate our programs by referring to the solving explanations. To ensure label quality and consistency, they are required to read the guideline of annotation standards and examples in advance. Each annotated program is double-checked by one of the authors, and the annotator with low accuracy would be disqualified. The annotated operations required to solve the problem are limited to a maximum of 4 steps, thus a small number of complex and hard problems are filtered.

4 Neural Geometric Solver

We propose Neural Geometric Solver (NGS) to address geometric problems by jointly understanding text, diagram, and then generating explainable programs. Moreover, we utilize some novel auxiliary tasks to enhance the understanding ability of our NGS. The overall architecture of our NGS is shown in Fig. 2.

4.1 The Architecture of NGS

4.1.1 Problem Encoder

Text Encoder Given a problem text $P = \{x_i\}_{i=1}^n$, each token x_i is first embedded into a word vector \mathbf{x}_i . A single-layer unidirectional LSTM (Hochreiter and Schmidhuber, 1997) is then applied to encode each word embedding \mathbf{x}_i into a hidden state \mathbf{h}_i . The sequence of the hidden state in LSTM are used to represent problem text P as $H_P = [\mathbf{h}_0; \dots; \mathbf{h}_n]$.

Diagram Encoder For representing a problem diagram, we apply the first three stages of a ResNet-101 (He et al., 2016) to extract it as the diagram feature, which can be formalized as a feature matrix $H_D \in \mathbb{R}^{m \times d}$. Moreover, we also apply multiple auxiliary tasks for pretraining the diagram encoder. Note that the parameters of the diagram encoder are fixed when training the overall NGS.

4.1.2 Joint Reasoning Module

Given the text feature H_P and the diagram feature H_D , it is crucial for solving geometric prob-

¹<http://www.zxxk.com/> and <http://www.jyeoo.com/>

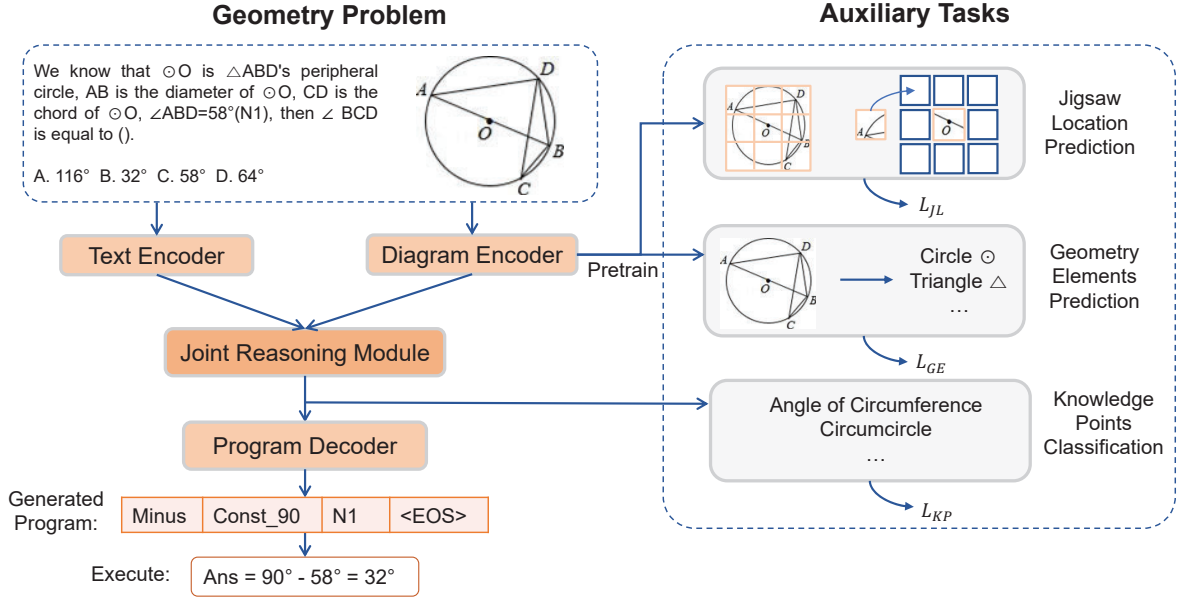


Figure 2: The overall architecture of our Neural Geometric Solver (left) in conjunction with auxiliary tasks (right). The problem text and diagram is encoded separately, then fed into a joint reasoning module together to obtain cross-modal fusion of text and diagram features. A decoder utilizes fused multimodal features to generate the interpretable programs. In addition, we propose three auxiliary tasks to enhance feature representation and facilitate multimodal reasoning.

lems to jointly fuse and align the cross-modal information. To this end, inspired by Yu et al. (2019), we adopt a co-attention module to conduct cross-modal joint reasoning with attention mechanism. The co-attention module consists of 12 self-attention (SA) units and 6 guided-attention (GA) units, which fully fuse and align the text representation H_P and the diagram representation H_D . H_P is first encoded by 6 self-attention units (i.e., original Transformer), and the final hidden state processed by the 6-th self-attention unit is used as guiding information. Then, the guiding information is fed into another stacked 6 self-attention units and 6 guided-attention units to achieve cross-modal semantic fusion and alignment. Finally, the co-attention module outputs a cross-modal representation $F_D = [f_1^D; \dots; f_n^D]$, which contains rich information over the problem text and diagram.

In this work, we find that text information is more fundamental than diagram information. Therefore, we further enhance the cross-modal representation with the help of textual information. Specifically, we concatenate H_P and F_D to acquire an enhanced reasoning module output F_R for decoding programs.

Besides, an attentional reduction network with a two-layer MLP is applied to aggregate feature F_D into \tilde{F}_D . Similarly, we concatenate \tilde{F}_D and the last

encoder state of the text encoder h_n , obtaining \tilde{F}_R as the final gathered multimodal feature vector.

4.1.3 Program Decoder

The program decoder module generates the programs sequentially under the guidance of multimodal information. Concretely, we use a LSTM decoder (Hochreiter and Schmidhuber, 1997) with attention (Bahdanau et al., 2014) over the reasoning module output F_R . Let $\{y_t\} (1 \leq t \leq T)$ be the target program to be generated and s_t be the hidden state of LSTM at time step t . \tilde{F}_R is fed into a linear layer to obtain the initial state s_0 . s_t is concatenated with the attention result and fed to a linear layer with the softmax function to predict the distribution of the next program token P_t .

During training, the generation loss L_g is the negative log-likelihood (NLL) of the target program:

$$\mathcal{L}_g(\theta) = \frac{1}{T} \sum_{t=1}^T \log P_t(y_t | \mathbf{x}, y_1, \dots, y_{t-1}; \theta),$$

where θ are the parameters of the entire NGS architecture except for the diagram encoder, \mathbf{x} is the input of both problem text and the extracted diagram feature. When testing, the decoder only observes the input text and diagram feature along with the program parts that have been generated.

4.1.4 Program Executor

After a beam of top N program sequences $\{g_1, \dots, g_n\}$ are generated from the program decoder, the executor computes them step by step. When executing the program, the token sequence is first divided into several parts based on the position of operators in the program. Once a complete operation program has been decoded, each operator in the program is executed sequentially to obtain a numerical result. The execution process fails if g_i has a grammar error (e.g., the number of arguments does not match the current operator), or the executed value does not match any options in the current problem. NGS adopts the first successfully executed program as the predicted solution and chooses the corresponding operation. If all N program sequences fail, the executor will report “no result” directly instead of guessing an option. Fig. 1 shows the detailed step-by-step execution of a final predicted program sequence.

4.2 Auxiliary Tasks

4.2.1 Self-supervised Diagram Auxiliary Task

Although our NGS can jointly fuse text feature and diagram feature with a co-attention mechanism, a powerful diagram encoder is needed to improve the problem understanding and answer accuracy. To obtain a high-quality diagram feature, we investigate two self-supervised auxiliary tasks, named as **Jigsaw Location Prediction** and **Geometry Elements Prediction**, to pretrain diagram encoder.

Jigsaw Location Prediction In the jigsaw location prediction task that enforces pixel-level perception, we first split a diagram as $m \times m$ blocks and select the center block as the target. Then, we shuffle other blocks randomly and train the diagram encoder to predict the correct relative location between these shuffled blocks and the target using a cross-entropy loss.

Geometry Elements Prediction For object-level understanding, we design a geometry elements prediction task that aims at training the diagram encoder to predict the geometry elements appearing in the diagram. A diagram usually contains multiple geometry elements which are also mentioned in the problem text and the solving explanation. We extract these geometry elements from text as the label and deploy an N -way classifier with binary cross-entropy (BCE) as the loss function to train the diagram encoder, where N is

the number of the possible geometry elements on the training set.

4.2.2 Knowledge Points Prediction

In addition to self-supervised diagram training, we also propose another auxiliary learning task called knowledge points prediction to enhance a problem’s overall representation by providing an extra training signal. We summarize about 50 knowledge points for our GeoQA, and label each problem with one or more knowledge points. We predict the knowledge points for each problem based on the gathered feature vector \tilde{F}_R outputted from the joint reasoning module. Different from the diagram training, the knowledge points prediction task is trained with NGS simultaneously. We also deploy a K -way classifier with binary cross-entropy (BCE) as the loss function to train the knowledge points prediction multi-label task, where K is the total number of the possible knowledge points on the training set.

5 Experiments

5.1 Experimental Setup and Training Details

We conduct experiments on GeoQA dataset, and adopt answer accuracy as the evaluation metric. Although there is another available geometric problem dataset (Seo et al., 2015a), the limited data scale (with only 67 training samples) makes it impossible to support neural network training. On the other hand, previous geometry problem solving systems require additional inputs (e.g., OCR and dependency parsing results) of the problem (Ye et al., 2008; Seo et al., 2015b) or not release their codes (Sachan and Xing, 2017; Sachan, 2020). Therefore, they are not comparable on our GeoQA dataset.

Implementation Details: In this work, we implement the proposed method with Pytorch (Paszke et al., 2017). The learning rate is $1e^{-3}$ and the batch size is set to 32. All models are trained around 100 epochs and optimized by Adam optimizer (Kingma and Ba, 2014). The beam size is typically set to 10. When pretraining the diagram encoder, we first fill the diagram with a white background to make it equal in length and width, and resize it to 224×224 . Then, we utilize the diagram feature extracted by the encoder to predict jigsaw location and geometry elements simultaneously and optimize the diagram encoder to obtain an informative diagram feature with a learning rate

Method		Total (%)	Angle (%)	Length (%)	Other (%)
Human	Text-Only	63.0	58.1	71.7	55.6
	Text-Diagram	92.3	94.3	90.5	87.0
W/O Program	FiLM (Perez et al., 2017)	31.7	34.0	29.7	24.1
	RN (Santoro et al., 2017)	38.0	42.8	32.5	29.6
	MCAN (Yu et al., 2019)	39.7	45.0	34.6	25.9
Text-Only	Seq2Prog (Amini et al., 2019)	52.3	62.4	42.1	27.8
	BERT2Prog (Devlin et al., 2018)	54.7	65.8	42.1	35.2
Text-Diagram	BERT2Prog + Diagram	50.3	63.4	33.2	38.9
	Seq2Prog + Diagram	52.6	63.6	39.2	37.0
	NGS (Ours)	56.7	67.5	44.5	37.0
	NGS-Auxiliary (Ours)	60.7	72.0	47.0	44.4

Table 3: The answer accuracy comparison on different test subsets of GeoQA dataset. “Human”, “W/O Program”, “Text-Only”, and “Text-Diagram” refer to the performance of human, not using the program, using text modal only, and conducting multimodal numerical reasoning on both text-diagram modals, respectively.

Method	BS	Acc(%)	NR(%)
Seq2Prog + Diagram	1	31.7	58.3
	10	52.6	20.4
	100	58.3	5.86
NGS-Auxiliary	1	45.6	42.6
	10	60.7	14.6
	100	64.5	2.95

Table 4: Performance comparison under different beam size settings. BS, Acc, and NR represent beam size, accuracy, and no result, respectively.

of $1e^{-5}$. Finally, the loss weight of the knowledge points classification task is set to 1, to promote the overall understanding of problems.

5.2 Experimental Results

We introduce three types of models here and test them on our GeoQA. The performance comparison with various methods on the different subset types of GeoQA is reported in Table 3.

Human Performance. We invite 10 students with a high score (top 1%) in the national university entry exam to answer these geometric problems. For each question, they first try to solve the problem with the text only and draw a diagram by themselves. Then, the actual diagram is given to answer the complete question. The total time for each question is limited to two minutes. When using text and diagram simultaneously, the human performance is improved from 63.2% to 92.3%, which indicates that humans are not good at solving problems using only text, but handling multimodal information successfully. The result shows that there is still a huge gap between the existing models and human experts, leaving large room for future research.

The effectiveness of program. “W/O Program” refers to not using programs and regarding GeoQA as a classification problem similar to VQA. Therefore, we conduct experiments on three models with multimodal reasoning capabilities, including FiLM, RN, and MCAN. However, their performance results show that these methods fail to reason about such complex geometric problems, achieving poor performance on GeoQA. These results prove the effectiveness and importance of our designed interpretable programs.

The necessity of multi-modality. “Text-Only” means that models only use text to generate program sequences since humans can also understand the intent of the question, draw diagram based on the text, and solve the problem. Motivated by Amini et al. (2019), we design a Sequence-to-Program (Seq2Prog) model using GRU encoder with an attention mechanism. Moreover, we replace the encoder with BERT and get a stronger BERT2Prog. By predicting our tailor-designed programs, these two methods are effective, while the performance is not satisfactory enough. These results show that multimodal reasoning ability is indispensable when solving geometric problems.

Multimodal numerical reasoning baselines. “Text-Diagram” refers to using text and diagram simultaneously. We concatenate the text embedding with the diagram feature extracted by ResNet (He et al., 2016) as the baseline methods, such as Seq2Prog + Diagram and BERT2Prog + Diagram. These feature fusion methods do not have the strong reasoning ability and fail to improve the program decoding. For instance, by adding diagram, the performance of BERT2Prog + Diagram declines from 54.7% to 50.3%, which may result

⊙O is a circle with a radius of 1, the distance from point O to the line L is 3. P is the moving point on L, and PQ is the tangent to the circle. PQRS is a square, and its minimum area is ().
A. 7 B. 8 C. 9 D. 10
Answer: B. 8

Knowledge Points: Pythagorean Theorem, Tangent, Square
Problem Solving Explanations:
When the square area is smallest, $PQ = \sqrt{PO^2 - QO^2} = 2\sqrt{2}$.
Thus, Area = $PQ \times PQ = 8$

Annotated Program: PythagoreanMinus | N1 | N0 | Multiply | V0 | V0
Baseline (No Result): Tan | N1 | CircleArea | N0
NGS (No Result): Add | N0 | N1 | Multiply | N0 | V0
NGS-Auxiliary (Right): PythagoreanMinus | N1 | N0 | Multiply | V0 | V0

A student saw a tree and their distance was 20m. The reflection of the top of the tree in the water was 5m away from him. The student's height is 1.7m, and the tree height is ()m.
A. 3.4 B. 5.1 C. 6.8 D. 8.5
Answer: B. 5.1

Knowledge Points: Similar Triangles, Distance
Problem Solving Explanations:
 $20 - 5 = 15$ m, By the nature of similar triangles, $5/15 = 1.7/H$.
 $H = 15 \div 5 \times 1.7 = 5.1$ m. Thus, height is 5.1m.

Annotated Program: Minus | N0 | N1 | Proportion | V0 | N1 | N2
Baseline (No Result): Add | N2 | N1 | Proportion | N0 | V0 | N2
NGS (No Result): Proportion | N2 | N0 | N1
NGS-Auxiliary (Wrong): Proportion | N2 | N1 | N0

Figure 3: Typical cases. No Result represents the answer executed by the programs is not in the options, and Wrong represents getting a wrong option. Baseline is a “Seq2Prog + Diagram” model. In the case on the left, NGS-Auxiliary successfully predicts the knowledge of the Pythagorean theorem through auxiliary tasks and get the right answer. For the case on the right, the problem is quite hard that current model cannot solve it.




Figure 4: Ablation study on different auxiliary components. ‘+’ represents we add the auxiliary component. NGS-Auxiliary means that adding all three auxiliary tasks together.

from the extra diagram that disturbs the text pre-training model. Note that multimodal pretraining models (Lu et al., 2019; Li et al., 2020) cannot be applied to geometric problems, since these models are based on Faster-RCNN to extract object-level features from natural images.

The effectiveness of our methods. Our proposed NGS shows a relatively-good performance compared to the various models mentioned above. When adding all three auxiliary tasks together to enhance NGS solver, our NGS-Auxiliary with multimodal reasoning ability becomes the existing best-performing method (60.7%) on GeoQA dataset. It also achieves the highest accuracy on all types of problems. For example, compared with Seq2Prog+Diagram, NGS-Auxiliary obtains an 8.4% performance improvement on the angle type problems. Compared with other “Text-Diagram” baselines, our model is effective when reasoning

on multimodal information.

The effect of different beam size. In general, we set the beam size to 10 for testing. In this section, we explore the influence of different beam size. After the searched sequence program is executed, there will be three situations: right answer, wrong answer, and no result. As shown in Table. 4, as the beam size is larger, we get higher accuracy and a lower proportion of no result. When beam size equals 1, the NGS-Auxiliary outperforms baselines significantly. Our model can achieve the highest accuracy of 64.5% when beam size is 100.

5.3 Ablation Study

As shown in Fig. 4, we conduct experiments to evaluate the contribution of different auxiliary tasks. We consider six different combinations: 1) only the NGS; 2) NGS + Geometry Elements (NGS+GE); 3) NGS + Jigsaw Location (NGS+JL); 4) NGS + Knowledge Points (NGS+KP); 5) NGS + diagram-based pretraining (NGS+JL+GE); 6) NGS with all three auxiliary tasks (NGS-Auxiliary). We can see that all three auxiliary tasks can promote the performance of NGS. The accuracy gains of GE, JL, KP, JL + GE, and combining all three tasks are 0.9%, 1.2%, 2.2%, 2.9%, 4.0%, respectively. These results show that all our self-supervised and auxiliary tasks can enhance the comprehensive understanding and multimodal reasoning ability of NGS.

5.4 Case Analysis

As shown in Fig. 3, we select two typical cases to demonstrate the programs generated by different models and some representative errors.

In the left case, the knowledge of the

Pythagorean theorem, tangent, and square are required for solving the problem. Baseline method and our NGS fail to generate the correct operations and get no result. However, our NGS-Auxiliary successfully predicts the use of knowledge in the proposed auxiliary task. And more importantly, the correct "PythagoreanMinus" program is generated, and the right answer is obtained.

The right one is a typical error case, in which model needs to understand a complex scene. Although NGS-Auxiliary has predicted the knowledge points of Proportion program correctly, all three models fail to predict the correct answer. A better multimodal method is required to handle this hard high-level reasoning task in the future.

6 Conclusion

In this work, we focus on the geometric problem and propose the first large-scale geometric question answering dataset "GeoQA", containing 5,010 problems with program annotation. Besides, we propose a deep neural baseline, named as Neural Geometric Solver (NGS), to solve a geometric problem by jointly reasoning over multimodal data and generating interpretable programs. We further propose multiple novel auxiliary tasks to enhance the semantic representation of text and diagram. Extensive experimental results and analyses show that our GeoQA is challenging, and our NGS-Auxiliary outperforms other methods on GeoQA.

7 Ethical Impact

We collected GeoQA from two online education websites, which is only used for academic research, and the copyright belongs to the original websites. This work may inspire research in the field of multimodal numerical reasoning.

Acknowledgements This work was supported in part by National Key R&D Program of China under Grant No.2020AAA0109700, National Natural Science Foundation of China (NSFC) under Grant No.U19A2073 and No.61976233, Guangdong Province Basic and Applied Basic Research (Regional Joint Fund-Key) Grant No.2019B1515120039, Shenzhen Fundamental Research Program (Project No.RCYX20200714114642083 and No.JCYJ20190807154211365), Zhijiang Lab's Open Fund (No.2020AA3AB14), and CSIG Young Fellow Support Fund.

References

- Unaiza Ahsan, Rishi Madhok, and Irfan Essa. 2019. Video jigsaw: Unsupervised learning of spatiotemporal context for video action recognition. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 179–189. IEEE.
- Aida Amini, Saadia Gabriel, Peter 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*, pages 2357–2367.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Richa Bajaj and Vidushi Sharma. 2018. Smart education with artificial intelligence based determination of learning styles. *Procedia computer science*, 132:834–842.
- Shang-Ching Chou, Xiao-Shan Gao, and Jing-Zhong Zhang. 1996. Automated generation of readable proofs with geometric invariants. *Journal of Automated Reasoning*, 17(3):325–347.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Carl Doersch and Andrew Zisserman. 2017. Multi-task self-supervised visual learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2051–2060.
- Herbert Gelernter, James R Hansen, and Donald W Loveland. 1960. Empirical explorations of the geometry theorem machine. In *Papers presented at the May 3-5, 1960, western joint IRE-AIEE-ACM computer conference*, pages 143–149.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6904–6913.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.

- 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. Association for Computational Linguistics.
- Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. 2017. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2901–2910.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Nate Kushman, Yoav Artzi, Luke Zettlemoyer, and Regina Barzilay. 2014. Learning to automatically solve algebra word problems. In *Proceedings of the 52th Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 271–281.
- Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. 2017. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690.
- Gen Li, Nan Duan, Yuejian Fang, Ming Gong, and Daxin Jiang. 2020. Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11336–11344.
- Jinjiao Lin, Haitao Pu, Yibin Li, and Jian Lian. 2018. Intelligent recommendation system for course selection in smart education. *Procedia Computer Science*, 129:449–453.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *arXiv preprint arXiv:1908.02265*.
- Jiayuan Mao, Chuang Gan, Pushmeet Kohli, Joshua B Tenenbaum, and Jiajun Wu. 2019. The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision. *International Conference on Learning Representations*.
- 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.
- Alejandro Newell and Jia Deng. 2020. How useful is self-supervised pretraining for visual tasks? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7345–7354.
- Mehdi Noroozi and Paolo Favaro. 2016. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision*, pages 69–84. Springer.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch.
- Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. 2016. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2536–2544.
- Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. 2017. Film: Visual reasoning with a general conditioning layer. *arXiv preprint arXiv:1709.07871*.
- Mrinmaya Sachan. 2020. Knowledge graph embedding compression. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2681–2691.
- Mrinmaya Sachan, Avinava Dubey, Eduard H Hovy, Tom M Mitchell, Dan Roth, and Eric P Xing. 2020. Discourse in multimedia: A case study in extracting geometry knowledge from textbooks. *Computational Linguistics*, 45(4):627–665.
- Mrinmaya Sachan, Kumar Dubey, and Eric Xing. 2017. From textbooks to knowledge: A case study in harvesting axiomatic knowledge from textbooks to solve geometry problems. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 773–784.
- Mrinmaya Sachan and Eric Xing. 2017. Learning to solve geometry problems from natural language demonstrations in textbooks. In *Proceedings of the 6th Joint Conference on Lexical and Computational Semantics (*SEM 2017)*, pages 251–261.
- Adam Santoro, David Raposo, David G Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy Lillicrap. 2017. A simple neural network module for relational reasoning. In *Advances in neural information processing systems*, pages 4967–4976.
- Min Joon Seo, Hannaneh Hajishirzi, Ali Farhadi, and Oren Etzioni. 2014. Diagram understanding in geometry questions. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*.
- Minjoon Seo, Hannaneh Hajishirzi, Ali Farhadi, Oren Etzioni, and Clint Malcolm. 2015a. Solving geometry problems: Combining text and diagram interpretation. In *Proceedings of the 2015 Conference on*

Empirical Methods in Natural Language Processing, pages 1466–1476.

Minjoon Seo, Hannaneh Hajishirzi, Ali Farhadi, Oren Etzioni, and Clint Malcolm. 2015b. Solving geometry problems: Combining text and diagram interpretation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1466–1476. Association for Computational Linguistics.

Lei Wang, Yan Wang, Deng Cai, Dongxiang Zhang, and Xiaojiang Liu. 2018. 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.

Yan Wang, Xiaojiang Liu, and Shuming Shi. 2017. 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.

Wu Wen-Tsun. 1986. Basic principles of mechanical theorem proving in elementary geometries. *Journal of automated Reasoning*, 2(3):221–252.

Zhipeng Xie and Shichao Sun. 2019. 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.

Zheng Ye, Shang-Ching Chou, and Xiao-Shan Gao. 2008. An introduction to java geometry expert. In *International Workshop on Automated Deduction in Geometry*, pages 189–195. Springer.

Kexin Yi, Jiajun Wu, Chuang Gan, Antonio Torralba, Pushmeet Kohli, and Joshua B Tenenbaum. 2018. Neural-symbolic vqa: Disentangling reasoning from vision and language understanding. *arXiv preprint arXiv:1810.02338*.

Zhou Yu, Jun Yu, Yuhao Cui, Dacheng Tao, and Qi Tian. 2019. Deep modular co-attention networks for visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6281–6290.

Lipu Zhou, Shuaixiang Dai, and Liwei Chen. 2015. Learn to solve algebra word problems using quadratic programming. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 817–822. Association for Computational Linguistics.