
Enhancing Neural Mathematical Reasoning by Abductive Combination with Symbolic Library

Anonymous Authors¹

Abstract

Mathematical reasoning recently has been shown as a hard challenge for neural systems. Abilities including expression translation, logical reasoning, and mathematics knowledge acquiring appear to be essential to overcome the challenge. This paper demonstrates that some abilities can be achieved through abductive combination with discrete systems that have been programmed with human knowledge. On a mathematical reasoning dataset, we adopt the recently proposed abductive learning framework, and propose the ABL-Sym algorithm that combines the Transformer neural models with a symbolic mathematics library. ABL-Sym shows 9.73% accuracy improvement on the interpolation tasks and 47.22% accuracy improvement on the extrapolation tasks, over the state-of-the-art approaches.

1. Introduction

Automatically solving natural language described mathematical problems has been shown very challenging, requiring natural language understanding, mathematical expression extraction, and complex symbolic reasoning. Existing deep learning-related methods mainly frame these problems as a machine translation task. A branch of the methods explicitly encode the structural relation and try to directly output the answers (Saxton et al., 2019; Schlag et al., 2019). These methods have a great expression ability, but are hard to generalize to unseen cases. Another branch learns a mapping from the problem description to a solution program (Wang et al., 2017; Amini et al., 2019), which explicitly encodes domain knowledge. These program-based methods rely heavily on human labeling, which is not only laborious, time-consuming, and error-prone. Besides, some problems are hard to be expressed in a program format,

¹Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

such as the varieties of probability problems (e.g., Three letters picked without replacement from `idiidauuuuuaiduuiiu`. What is prob of sequence `iaa?`).

Recently, the abductive learning (Dai et al., 2019) introduces a discrete logic module into a neural network with an integrated learning procedure. The logic module utilizes the logical consistency between the perception outputs and the logic background knowledge to optimize the perception module and the logic module jointly. This work demonstrates the possibility to produce a system with both the flexible perception power from neural networks and the generalization power from the programmed knowledge.

In this paper, we follow the abductive learning framework and propose a system that integrates the transformer networks and a mathematical symbolic library, ABL-Sym, for automatically solving math problems. ABL-Sym firstly runs a consistency check and correction procedure: it generates programs from natural language descriptions and uses a program executor to run the programs; if the program output is inconsistent with the answer, it employs a search routine to correct the program. ABL-Sym then learns from the problem descriptions and the corrected programs. ABL-Sym repeats the two steps to improve its model. We evaluate ABL-Sym on the mathematics dataset from (Saxton et al., 2019). The results show that ABL-Sym significantly outperforms the previous state-of-the-art approaches: it achieves 9.72% accuracy improvement on interpolation tasks, and 47.22% accuracy improvement on extrapolation tasks.

2. Background

2.1. The Mathematics Dataset

Saxton et al. (Saxton et al., 2019) introduced a mathematics dataset that contains a variety of math problems, including algebra, arithmetic, numerical comparison, numerical factorization, calculus, measurement, and probability. Each problem is a question-answer pair, where the question is like `Let $q(m) = m**3 + 2$. Let $r(c) = -4*c**3 - 9$. What is $18*q(f) + 4*r(f)$?` and the answer is like `2*f**3`. Although there may be many forms of answer sequences with the

same mathematical meaning, the evaluation criterion is character-by-character (i.e., each question is scored by either 0 or 1 according to whether the answer matches the correct answer character-by-character). The dataset is procedurally generated and consists of 56 modules, and each module provides 2M per-generated training samples and 10k interpolation samples. Extrapolation samples are also provided for an additional measure of algebraic generalization.

2.2. Sympy

Sympy (Meurer et al., 2017) is a mathematical symbolic computing library, which contains about 300+ mathematical functions. Although many mathematical engines can be used, we adopt Sympy because it can conveniently get all the appropriate mathematical functions, easily exclude non-mathematical functions, and support direct access to the docstrings of mathematical functions.

2.3. Related work

A mathematics dataset was released in (Saxton et al., 2019) that analyzes the reasoning and generalization ability of popular reasoning neural architectures such as recurrent neural architectures and attention-augmented architectures (i.e., Transformer (Vaswani et al., 2017)). The results show that the learned models did not do mathematical reasoning well, particularly for the extrapolation zone. (Schlag et al., 2019) incorporates the tensor-product representation technique within the Transformer to better support the explicit representation of relation structure. They achieved improved results than the vanilla Transformer architecture without introducing any domain knowledge.

Program format is a typical way to represent both of the discrete domain knowledge and the solution structure of mathematical problems. Amini et al. (Amini et al., 2019) released a dataset of math word problems that are densely annotated with programs by crowd-sourcing. Based on the dataset, they proposed a sequence-to-program model with automatic problem categorization. Comparing with their method, our approach applying to the dataset without annotated programs, and moreover, we use both the neural network and the discrete symbolic system for prediction.

Abductive learning (Dai et al., 2019) was recently proposed for connecting a perception module with an abductive logical reasoning module using consistency optimization. The perception module generates output, the reasoning module checks and corrects the logical consistency, and the consistency information is used to update the perception module to generate logically more consistent output. This constitutes a forward cycle. Our approach is inspired by the above abductive learning framework, while we are addressing a different domain.

3. ABL-Sym

In the following subsections, we introduce the program definition, the program correction, and the training procedure.

3.1. Programs

We define the program based on a domain-specific language (DSL) instead of arbitrary Turing-complete languages to reduce the search space of programs. Every word in the DSL is called an operator. All available operators form an operator space. The relationship between adjacent operators is appropriately restricted, such as *argc* operators must be followed by *math* operators, the number of optional variables must be no less than *argc*, and *argc* must be an available number of arguments to the mathematical operator.

3.1.1. OPERATOR SPACE

The operator space consists of about 400 operators, including mathematical operators, position-aware operators, and several auxiliary operators.

Mathematical Operators: We use Sympy as the program executor. In Sympy, there are about 300+ functions which are essential for solving math problems (e.g. *add*, *multiply*, *solve*, *diff*). We consider these functions as our mathematical operators.

Position Operators: Mathematical expressions in problem may appear anywhere. We tokenize the problem sentence with a simple tokenizer and use positional indexes to identify expressions. The tokenizer uses space to tokenize the sentence and uses tokens that are not in the ordinary word dictionary as expressions. The ordinary word dictionary consists of non-digit words and excludes common ordinal number words (e.g., *first*, *second*, *square*). In addition, we also exclude a-z single letters because they are often used to represent variables in math problems. After tokenizing the problem, the positional continuous expression tokens are merged into one token. We use *pos0*, *pos1*, *pos2*, ... as positional operators to represent the positions of related expressions.

Auxiliary Operators Functions in Sympy may have multiple parameters (e.g. *diff* function for obtaining derivative may have two usages: `diff(x**2+x`
`y, x)`, `diff(x**2+x*y, x, 2)`). We add *argc0*, *argc1*, *argc2*, *argc3* to the operator space in order to explicitly specify the number of function parameters. Some expressions in question do not conform to the input format of mathematical operators, and the output formats of some operators do not conform to the answer, so we add several additional format conversion operators and operator wrappers into the operator space.

3.1.2. PROGRAM EXECUTOR

We build a simple program executor based on Sympy to run programs. In a running, the program’s operators are executed sequentially, and intermediate results are saved in the environment through registry variables, which may be used by later operators. If an error is encountered during execution, execution will stop and return `none`, or if execution reaches the end, return execution `result`.

3.1.3. PROGRAMS SEARCH PROCEDURE

The program search space is too large to find the correct programs by random search. We design an abductive learning framework to search programs efficiently. Our framework performs multiple iterative searches. In the first iteration, we use a search-based method as a program generator to generate some programs. Then, the program executor runs the programs, and a consistency checker filters out the programs whose results are inconsistent with the answers. A neural network model is used to learn the mapping from the problem to the correct program. The learned model is then used to be a better program generator to start another iteration. In addition, we develop the following techniques to speed up the search process further.

Warm-up Operator Distribution In math problems, problems are often strongly related to mathematical terms (e.g. in the derivative problems, the terms `derivative`, `differentiate` often appear). Additionally, almost every mathematical function in Sympy has a docstring, which usually contains related mathematical terms. So we can build relationships between problems and mathematical operators. In this paper, we adopt (Arora et al., 2016) method to calculate the cosine similarity between the problem description and the docstring of an operator and then normalize by softmax to obtain the probability distribution of operators, which is used to generate the possible programs.

Curriculum Search Strategy According to whether the problem consists of simple problems, the problems in the Mathematics Dataset can be divided into simple problems and compositional problems (e.g. a compositional problem: Suppose $-2*v + 1873 = 4*x - 3*x$, $x = 2*v - 1863$. Let $u = -65 + 25$. Find the common denominator of $1/6$ and $v/(-920) - 8/u$.) Programs for simple problem can be found relatively easily by searching, but not for compositional problem. We observe that the compositional problem can be broken down into multiple parts, each of which is similar to a simple problem (e.g., the above problem can be broken down into three parts: Suppose $-2*v + 1873 = 4*x - 3*x$, $x = 2*v - 1863$ # Let $u = -65 + 25$ # Find the common denominator of $1/6$ and $v/(-920) - 8/u$).

Therefore, we use the neural network model learned from simple problems to generate possible programs for each part and organize them into complete programs. The program executor then executes the programs to get results, and the consistency checker then checks the results for correctness. The whole search process is time-consuming, so we only perform search on randomly generated 500k problems that meet the qualifying conditions, and use the learned model to generate the rest.

3.2. Neural Models

The neural network model we use is a modified version of the original Transformer (Vaswani et al., 2017), with a shared transformer encoder θ^{enc} and two separate transformer decoders θ_a^{dec} and θ_p^{dec} . We use the encoder θ^{enc} with hidden states \mathbf{h}^{enc} to encode the problem \mathbf{x} . The decoders θ_a^{dec} and θ_p^{dec} take the shared hidden states \mathbf{h}^{enc} and auto-regressively generates the answer sequence and program sequence respectively. During training, the decoders receive the shifted targets while during inference, we use the previously generated symbols with the highest probability. We treat the question and answer as a sequence of characters just like (Vaswani et al., 2017) and treat the question as a sequence of operators. The overall training loss is the weighted sum of the answer decoding loss and the program decoding loss:

$$\mathcal{L}(\theta^{enc}, \theta_a^{dec}, \theta_p^{dec}) = -\alpha_1 \log P(\mathbf{y}_a | \mathbf{x}; \theta^{enc}, \theta_a^{dec}) - \alpha_2 \log P(\mathbf{y}_p | \mathbf{x}; \theta^{enc}, \theta_p^{dec})$$

4. Experiments

We evaluate our framework on the mathematics dataset (Saxton et al., 2019). The reason we did not evaluate on other mathematical datasets (Kushman et al., 2014; Huang et al., 2016; Upadhyay & Chang, 2016; Wang et al., 2017; Ling et al., 2017; Amini et al., 2019) is because these datasets are either limited to narrow specific fields or demanded for manual annotated programs.

4.1. Settings

During the search, the maximum number of sampled programs for each problem is $N_w = 100k$ on the first iteration and $N_n = 1k$ on the other iterations. The number of iterations I is set to 5.

We extract a character-level vocabulary of 72 symbols and an operator-level vocabulary of 380 symbols, both including `START`, `END`, and `PADDING` symbols.

Our transformer-like model parameters θ^{enc} , θ_a^{dec} , θ_p^{dec} are set to an embedding size of 512, with 8 attentional heads, and intermediate feed-forward dimension of 2048. The

Table 1. Model accuracy averaged over all modules. A sample is correct if all characters of the target sequence have been predicted correctly. The column “>95%” counts how many of the 56 modules achieve over 95% accuracy.

	weights	steps	interpolation		extrapolation	
			acc	>95%	acc	>95%
Transformer (Saxton et al.)	30M	500k	76.00%	13	50.00%	1
TP-Transformer (Schlag et al.)	49.2M	700k	80.67%	18	52.48%	3
Transformer (ours)	44.2M	700k	76.41%	13	50.48%	2
TP-Transformer (ours)	49.2M	700k	79.82%	18	51.99%	3
ABL-Sym+Transformer (ours)	54.9M	700k	87.85%	29	73.41%	7
ABL-Sym+TP-Transformer (ours)	58.8M	700k	88.52%	33	77.26%	8

answer decoder θ_a^{dec} is with layers of 6 while the program decoder θ_p^{dec} is with layers of 2. We train our model via the Adam optimizer (Kingma & Ba, 2014) with a learning rate of 8×10^{-5} , $\beta_1 = 0.9$, $\beta_2 = 0.995$, $\epsilon = 10^{-9}$. We use a batch size of 1024, with absolute gradient value clipping of 0.1. We trained our model on one server with 8 V100 Nvidia GPUs for 12 days. During the search process, the parameters configuration of our program-generated model is the same as the above model.

At the inference, answers and programs are generated by sequential decoding. If the predicted program is none or fails to run successfully, the neural model answer is used as the final result.

4.2. Experimental Results

Table 1 presents the overall performance on the dataset. We can see that our model significantly outperforms the previous state-of-the-art by up to 7.8% absolute improvement on the interpolation test dataset and 24.8% absolute improvement on the extrapolation test dataset. Our program-augmented model dramatically improves the performance of the model, especially for generalizing the model to areas not previously seen. For a more detailed comparison, Fig. 1 shows the test performance on extrapolation modules.

Table 2 shows the performance of the 5 iterations of ABL-Sym together with the random search performance. ABL-Sym shows clearly better than random search. In the first iteration, it used an average of 20% fewer search times than the random search strategy but found 76% more programs, which mainly due to the warm-up strategy and curriculum search strategy. These strategies allow us to search more programs faster within the maximum search limit. After the first iteration, the model we learned as a better program generator generated better candidate programs, so we searched an additional 8% of the programs with negligible search times. Compared to the second iteration, the number of programs searched in the next few iterations increased by only a litter bit. This is because most programs that can be searched are also almost searched. Still, 57.7% programs were not found.

Table 2. The cost and the hit ratio of programs during iterating

Method	per-question searches	hit ratio
ABL-Sym (1 itr)	64.14k	33.2%
ABL-Sym (2 itr)	64.86k	40.1%
ABL-Sym (3 itr)	65.49k	41.3%
ABL-Sym (4 itr)	66.11k	42.0%
ABL-Sym (5 itr)	66.73k	42.3%
Random search	82.09k	18.9%




Figure 1. The extrapolation test performances of our implementations of Transformer, TP-Transformer(700K steps) and our ABL-Sym framework based on Transformer and TP-Transformer on the different modules.

ABL-Sym can find many programs of compositional or complex problems (e.g., Calculate the common denominator of 25/13728 and 121/1248. the program found is pos7 argc1 denom pos5 argc1 denom argc2 lcm), but random search strategy was failed.

5. Conclusion

In this work, we demonstrate that integrating discrete systems into neural systems is a feasible way to enhance the neural systems, particularly in the extrapolation ability. Notice that even human beings learn complex knowledge, e.g. mathematics, progressively from well organized textbooks. Well designed discrete systems may serve the role of textbooks for building a complex intelligent systems.

References

- 220
221 Amini, A., Gabriel, S., Lin, P., Koncel-Kedziorski, R., Choi,
222 Y., and Hajishirzi, H. Mathqa: Towards interpretable math
223 word problem solving with operation-based formalisms.
224 *arXiv preprint arXiv:1905.13319*, 2019.
225
- 226 Arora, S., Liang, Y., and Ma, T. A simple but tough-to-beat
227 baseline for sentence embeddings. 2016.
228
- 229 Dai, W.-Z., Xu, Q., Yu, Y., and Zhou, Z.-H. Bridging
230 Machine Learning and Logical Reasoning by Abductive
231 Learning. In Wallach, H., Larochelle, H., Beygelzimer,
232 A., d\textquotesingle Alché-Buc, F., Fox, E., and Garnett,
233 R. (eds.), *Advances in Neural Information Processing Sys-*
234 *tems* 32, pp. 2811–2822. Curran Associates, Inc., 2019.
- 235 Huang, D., Shi, S., Lin, C.-Y., Yin, J., and Ma, W.-Y. How
236 well do computers solve math word problems? large-scale
237 dataset construction and evaluation. In *Proceedings of*
238 *the 54th Annual Meeting of the Association for Computa-*
239 *tional Linguistics (Volume 1: Long Papers)*, pp. 887–896,
240 2016.
241
- 242 Kingma, D. P. and Ba, J. Adam: A method for stochastic
243 optimization. *arXiv preprint arXiv:1412.6980*, 2014.
244
- 245 Kushman, N., Artzi, Y., Zettlemoyer, L., and Barzilay, R.
246 Learning to automatically solve algebra word problems.
247 In *Proceedings of the 52nd Annual Meeting of the Asso-*
248 *ciation for Computational Linguistics (Volume 1: Long*
249 *Papers)*, pp. 271–281, 2014.
- 250 Ling, W., Yogatama, D., Dyer, C., and Blunsom, P. Pro-
251 gram induction by rationale generation: Learning to solve
252 and explain algebraic word problems. *arXiv preprint*
253 *arXiv:1705.04146*, 2017.
- 254 Meurer, A., Smith, C. P., Paprocki, M., Čertík, O., Kirpichev,
255 S. B., Rocklin, M., Kumar, A., Ivanov, S., Moore, J. K.,
256 Singh, S., et al. Sympy: symbolic computing in python.
257 *PeerJ Computer Science*, 3:e103, 2017.
258
- 259 Saxton, D., Grefenstette, E., Hill, F., and Kohli, P. Analysing
260 mathematical reasoning abilities of neural models. *arXiv*
261 *preprint arXiv:1904.01557*, 2019.
- 262 Schlag, I., Smolensky, P., Fernandez, R., Jojic, N., Schmid-
263 huber, J., and Gao, J. Enhancing the transformer with
264 explicit relational encoding for math problem solving.
265 *arXiv preprint arXiv:1910.06611*, 2019.
266
- 267 Upadhyay, S. and Chang, M.-W. Annotating derivations:
268 A new evaluation strategy and dataset for algebra word
269 problems. *arXiv preprint arXiv:1609.07197*, 2016.
- 270 Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones,
271 L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Atten-
272 tion is all you need. In *Advances in neural information*
273 *processing systems*, pp. 5998–6008, 2017.
274
- Wang, Y., Liu, X., and Shi, S. Deep neural solver for math
word problems. In *Proceedings of the 2017 Conference*
on Empirical Methods in Natural Language Processing,
pp. 845–854, 2017.