Math Word Problem Generation Using Transformers and Reinforcement Learning

D. Naga Jyothi[1], M. Naresh[2], and P. Guru Prasad[3]
[1] *Department of AI ML*
[2] *Chaitanya Bharathi Institute of Technology, Gandipet, Hyderabad, India.*

Abstract: **Manually crafting math word problems is a labour-intensive process that teachers do, and one can sense a growing need for automated systems. However, many of the present models will generate problems that are grammatically correct but semantically incoherent, not solvable, or not aligned with the educational objectives. Addressing these issues is the motivation behind our work that enhances an MWP generation model using transformer architecture and reinforcement learning. Having integrated the topic-expression transformer mechanism, our approach will be to align the problem context with appropriate mathematical operations: MWPs are generated that are linguistically sound and mathematically proper. Towards the future, we would focus on the increase of diversity and complexity of the generated problems and evaluation of model adaptability across different datasets. Finally, we shall end up with an application that is user-friendly to enable real-time generation and interaction with MWPs with improved relevance, solvability and effectiveness in the educational setting.**

Keywords: **Math Word Problems (MWPs) · Automated Problem Generation · Transformer Architecture · Reinforcement Learning · Natural Language Processing (NLP).**

1. INTRODUCTION

Generating math word problems is the task of expressing mathematical ideas in realistic and contextually meaningful problems that are meaningful and educational. Early template-based systems have created part of this framework and coherence, but these have often been inflexible because they lead to repetitive problems with little variation in structure or grammar. Maintaining and expanding these templates also required huge amounts of manual effort and were therefore not scalable. The focus has moved, over the past few years, toward more data-driven approaches, specifically deep learning models, to improve the linguistic and mathematical accuracy of problems that are generated. In fact, some of the limitations of template-based methods, which deep learning could potentially overcome, have now recently been faced by deep learning models in that the solutions produced by these deep learning models do not appropriately ensure that problems generated are consistently solvable and targeted toward specific learning objectives.

The intrinsic complexity of MWPs also creates difficulties in generation and automatic solving. Actually, such problems demand models to read natural language, then apply mathematical reasoning on it and to translate the under-standing into correct mathematical expressions. Early MWP solvers would typically use rule-based systems or statistical learning methods, but they functioned poorly in flexibility and generalization across different kinds of problems. Deep learning truly got the models close to becoming accurate and robust, but they can’t cope with long-range dependencies in problem texts. The lack of really good, high-quality, annotated datasets for training is another insurmountable barrier. It is within these considerations that data augmentation techniques can help alleviate the challenges these researchers are facing. They also considered using the transformer architecture and reinforcement learning, with the ability to help to capture long-range relationships in text while at the same time ensuring that problems generated meet both linguistic and educational standards.

The recent promise for solving MWP is the emergence of new developments, including using large language models trained on source code, known as Large Code Models (LCMs). This model can transform a problem statement in natural language to a form of solution in code that can be integrated with Intelligent Tutoring Systems (ITS) in order to solve a new MWP. This approach allows for the possibility of increasing the number of problems students could study, while providing a personalized mode of learning. Accuracy, however, is still a stumbling block for LCMs, especially as the complexity of the problem statement increased the chances that they translated it into a correct solution. Despite these problems, the emergence of LCMs and other advanced algorithms in machine learning, like transformers and reinforcement learning, will mark a bright future for MWP generation and solving-a task which will probably revolutionize ways we teach and practice mathematics.

2. LITERATURE SURVEY

[1] Wu et al. introduce MWPGen that generates MWPs by topics and mathematical expressions. The clever co-attention mechanism ensures that the model actually uses the specified topics and math expressions in a meaningful way. That is, they use even reinforcement learning, where the math problem solver provides feedback to improve the quality of generation, such that generated problems are both relevant and solvable. This directly addresses the common problem of generated MWPs that were lacking in coherence with their topic or equation.

[2] Qwen2.5-Math is a series of math special-purpose language models in- tended to augment mathematical reasoning. It includes constructing the highest- quality, math-specialized datasets that combine data from many sources: web content, code, encyclopedias, exam questions, and more-synthetic data generated by earlier Qwen models. These models also undergo iterative data synthesis and refinement due to the use of language models to evaluate and enhance the quality of the data: this process underlines the bi-directional nature in which LLMs can create training data and curate it. Qwen2.5-Math models also undergo continuous pre-training on such rich, math-centric datasets, which significantly enhances their capability to perform mathematical tasks, thereby possibly giving rise to high-quality MWPs.

Table 1. Comparison outlining each study's authors, techniques, advantages, and dis- advantages in Math Word Problem (MWP) generation and solving.

Authors	Methodology Used	Strengths	Limitations
Q. Wu, Q. Zhang, and X. Huang (2022)[1]	MWPGen, with a topic-expression co-attention mechanism and reinforcement learning to capture structural and semantic in- formation from expressions	Effectively links topic words to expressions, ensuring solvable and relevant problems	Dependent on the solver's quality; struggles with complex reasoning
A. Yang et al. (2024)	Qwen2.5-Math, a series of math-specific large language models using self-improvement techniques	Demonstrates advanced reasoning with Chain-of- Thought and Tool-Integrated Reasoning; supports both English and Chinese	High computational cost; may struggle with unseen problem types
A. Mitra, H. Khan- pour, C. Rosset, and A. Awadallah (2024)	Supervised Language Mod- els (SLMs) for grade school math, methodology not fully detailed	Aims to enhance grade school math tools; effectivefor early education	Specific details on the methodology not pro- vided; limitations un- clear
Q. Zhou and D. Huang (2019)	MAGNET, a neural network model for generating MWPs using fusion of equations and topics, with	Ensures relevance and high-quality problem generation; outperforms baselines	May struggle with more complex reasoning or multi-step problems

	entity-enforced loss		
Z. Wang, A. S. Lan, and R. G. Baraniuk (2021)	Pre-trained language mod- els with equation consistency constraint and context key- word selection for MWP generation	High mathematical consistency and language quality; model-agnostic	Focus on consistency may limit creativity; keyword selection could miss important context
P. Arnau-Gonza´lez, A. Serrano-Mamolar, S. Katsigiannis, T. Althobaiti, and M. Arevalillo-Herra´ez (2023)	LLM-based Python code generation for ITS, enabling automatic problem-solving	Automates MWP en- coding for ITS; enables easy problem addition	Limited accuracy in solving problems (39%) suggests the need for refinement
J. Qin, Z. Yang, J. Chen, X. Liang, and L. Lin (2024)	Template-Based Contrastive Distillation Pretraining (TCDP) combining mathematical logic and real-world knowledge	Incorporates both math logic and real- world knowledge; superior performance compared to state-of- the-art methods	Relies on quality of symbolic templates; may struggle with unseen problem types
Y. Wu and H. Nakayama (2024)	MILE, a neuro-symbolic solution for mathematical problems with new formula- representing techniques	Outperforms existing methods in accuracy, robustness, and generalization	Further investigation required to assess any potential limitations
S. Mandal and S. K. Naskar (2021)	AMWPS, combining machine learning and rule-based approaches for solving simple arithmetic MWPs	High accuracy on standard datasets (94.22%); effective for educational use	Limited to solving simple arithmetic problems; cannot handle complex scenarios
Y. Zhang, G. Zhou, Z. Xie, and J. X. Huang (2022)	HGEN, a hierarchical heterogeneous graph encoding method for MWP solving	Captures complex relationships and dependencies; out- performs Graph2Tree models	Computationally intensive; may struggle with large datasets or complex problems

[3] Mitra et al. explore how small language models (SLMs) might handle elementary math word problems, especially using GSM8K as a benchmark. The researchers create the Orca-Math-200K dataset, containing synthetic problems crafted by GPT-4 Turbo. Their experiments show that these smaller models, trained on high-quality synthetic data, can reach 87 percent accuracy on GSM8K, rivaling results usually expected from larger models. Instead of relying on resource-heavy ensemble methods, Mitra’s team demonstrates the value of iterative learning with synthetic data as a practical way to enhance smaller models’ reasoning.

[4] Zhou and Huang present MAGNET, designed to create math word problems from a given equation and set of keywords. The key innovation here is the equation-topic fusion mechanism, which combines information from both the math equation and relevant keywords, ensuring the generated problem is solvable and logically aligned with the input. This setup includes an entity-enforced loss, which keeps the model grounded to the entities from the equation, leading to consistent and contextually accurate MWPs.

[5] Wang and team focus on making sure the

math in generated MWPs actually matches the given equations. Their model uses a constraint-based approach that keeps the generated problem text consistent with the input equation, even employing a keyword-selection model to choose contextually relevant words automatically. By blending pre-trained language model capabilities with mathematical constraints, this approach helps generate meaningful and coherent MWPs that align well with the specified context.

[6] Arnau-González and colleagues focus on making it easier to include MWPs in intelligent tutoring systems (ITS). Instead of generating new problems, their approach translates problem statements into Python code that the ITS can use. This method simplifies adding new problems to ITS platforms but does not directly address the MWP generation challenge itself, instead focusing on the efficient integration of existing ones.

[7] Qin and colleagues take a slightly different angle, focusing on pre-training an MWP solver's comprehension abilities. They introduce a technique called template-based contrastive distillation pretraining (TCDP) to infuse a model with mathematical logic knowledge. By using formula templates and contrastive learning, the resulting model, Math Encoder, gains a deeper understanding of both the math and language in problems, ultimately boosting the solver's performance in downstream MWP tasks.

[8] Wu and Nakayama introduce MILE, a model that combines neural and symbolic methods for problem-solving. Using memory networks, MILE dynamically updates problem information throughout the solution process. They even add a formula mutation technique to expand the training data, a helpful addition for complex models like MILE. This combination of memory and symbolic reasoning makes MILE stand out, as it tackles math problems more effectively than previous methods.

[9] Mandal and Naskar explore a math solver named AMWPS designed to classify and solve single-operation arithmetic MWPs. The system categorizes problems based on keywords and verb analysis, ensuring it correctly interprets the problem type before solving it. While this system is effective for simple arithmetic problems, its focus is on solving rather than generating MWPs.

[10] Zhang and co-authors introduce HGEN, a model that uses a hierarchical graph to capture relationships within math problem texts. They argue that traditional text encoders miss the mark on complex relationships, so HGEN maps these relationships with a graph featuring word and quantity nodes. This approach uses multi-hop attention to capture long-range dependencies, making it particularly adept at handling problems with intricate mathematical relation- ships.

[11] Christ and colleagues present MATHWELL, a model fine-tuned to generate math word problems (MWPs) that are actually useful for K-8 students. It's built on LLaMa-2 with a massive 70-billion parameter model, but the real game-changer here is the EGSM dataset, which pairs MWPs with teacher- provided annotations. Teachers have flagged these examples for solvability, age- appropriateness, and accuracy, aiming to address issues of nonsensical or overly complex problems. This teacher-driven approach helps MATHWELL generate problems on par with GPT-4, while keeping the language clear and suitable for young learners. The authors argue that using a model without references to specific examples, as done here, is a viable way to keep generating fresh and suitable MWPs.

3. CHALLENGES AND EVOLUTION

Generating the Math Word Problems can be advanced from a simple rule-based system to a complex deep learning technique and currently advanced models such as the sequence-to-sequence and graph-to-tree architectures that have proceeded with such complex structures like trees, graphs, or attention mechanisms that more closely understand the relationship of a math problem. Tools such as reverse operation-based data augmentation (RODA) help create a broader variety of training data which makes the models more robust.

However, it still poses some challenges in passing over these barriers. With such developments in place, there is still difficulty in coming up with mathematically computable MWPs that are linguistically natural and clear. The models tend to do poorly if there is a need to make problems that demand a certain level of reasoning sophistication, and the metrics we apply to measure the quality of the generated problems remain in an embryonic stage.

Future research focuses more on how to improve reasoning using such models, the application of external knowledge to refine problem context, and making the models more explainable. It's very critical for educational tools, and generally it's very important, where teachers need to be able to trust that problems that the model generates are accurate and proper. If we can resolve these problems, MWP generators will prove worthwhile to include in intelligent tutoring systems as tools to help students learn much better as they would be subjected to personalized, well-crafted problems.

4. FUTURE DIRECTIONS

Integrating Qwen with advanced techniques, like heterogeneous graph encoding models, such as HGEN, appears to give the power to improve MWP generation. By appropriating the knowledge of Qwen in both natural language and mathematical concepts, there may be potential enhancement towards better representing relationships between different elements of a problem, such as quantities, operations, or contextual clues. This deeper understanding may well lead to more coherent, meaningful, and realistic MWPs that better approximate problem-solving at real-world levels. Qwen can also filter noisy or inaccurate augmented data out of her repertoire to improve higher-quality training examples. This would lead to even greater reliability and effectiveness in generating an MWP, one that is likely to be accurate and solvable.

Secondly, Qwen coupled with TCDP would enable the model to solve logic; hence, it would highly likely choose and use the mathematical templates the correct way. In tandem with this mathematical reasoning capability provided by Qwen, ideas of assured MWP correctness, linguistically as well as mathematically, would be established in line with specific learning objectives. Ultimately, these integrations will produce better and more pedagogically balanced MWPs that meet the needs of educators and learners in a more effective, personalized classroom learning experience.

5. CONCLUSIONS

This evolution in MWP solving has seen tremendous development from the rules and statistics to deeper neural network architectures such as sequence-to- sequence and graph-to-tree formats used for solving MWPs. Models employing tree structures, graph encoders and mechanisms of attention are better abstractions of relationships within the context of MWPs, while data augmentation techniques-based reverse operation RODA are known to help generate diverse, consistent training data. However, the challenges that arise in actually generating linguistically fluent and mathematically solvable MWPs are the following: current models struggle with complex reasoning chains, lack evaluative metrics on model quality, and require explicit information about the states. The future of research will be focused on more refined reasoning capabilities, inclusion of external knowledge sources, and better explanation power-a really prime requirement for educational applications. Other systems where MWP solvers can be added include intelligent tutoring platforms. This would therefore significantly improve educational experiences, but to fully realize the potential of MWP solving techniques, these challenges which remain must be addressed.

REFERENCES

[1] Q. Wu, Q. Zhang, and X. Huang, "Automatic Math Word Problem Generation With Topic-Expression Co-Attention Mechanism and Reinforcement Learning," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 1061–1072, 2022, doi: 10.1109/TASLP.2022.3155284.

[2] A. Yang et al., "Qwen2.5-Math Technical Report: Toward Mathematical Expert Model via Self-Improvement," Sep. 18, 2024, arXiv: arXiv:2409.12122. Accessed: Nov. 13, 2024. [Online].

Available:
http://arxiv.org/abs/2409.12122.

[3] A. Mitra, H. Khanpour, C. Rosset, and A. Awadallah, "Orca-Math: Unlocking the potential of SLMs in Grade School Math," Feb. 16, 2024, arXiv: arXiv:2402.14830. Accessed: Nov. 13, 2024. [Online]. Available:
http://arxiv.org/abs/2402.14830.

[4] Q. Zhou and D. Huang, "Towards Generating Math Word Problems from Equations and Topics," in *Proceedings of the 12th International Conference on Natural Language Generation*, K. van Deemter, C. Lin, and H. Takamura, Eds., Tokyo, Japan: Association for Computational Linguistics, Oct. 2019, pp. 494–503. doi:
https://doi.org/10.18653/v1/W19-8661.

[5] Z. Wang, A. S. Lan, and R. G. Baraniuk, "Math Word Problem Generation with Mathematical Consistency and Problem Context Constraints," Sep. 09, 2021, arXiv: arXiv:2109.04546. Accessed: Nov. 13, 2024. [Online]. Available:
http://arxiv.org/abs/2109.04546.

[6] P. Arnau-Gonza´lez, A. Serrano-Mamolar, S. Katsigiannis, T. Althobaiti, and M. Arevalillo-Herra´ez, "Toward Automatic Tutoring of Math Word Problems in Intelligent Tutoring Systems," *IEEE Access*, vol. 11, pp. 67030–67039, 2023, doi:
https://doi.org/10.1109/ACCESS.2023.329047810.1109/ACCESS.2023.3290478.

[7] J. Qin, Z. Yang, J. Chen, X. Liang, and L. Lin, "Template-Based Contrastive Distillation Pretraining for Math Word Problem Solving," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 35, no. 9, pp. 12823–12835, Sep. 2024, doi:
https://doi.org/10.1109/TNNLS.2023.3265173.

[8] Y. Wu and H. Nakayama, "MILE: Memory-Interactive Learning Engine for Neuro-Symbolic Solutions to Mathematical Problems," *IEEE Access*, vol. 12, pp. 134355–134365, 2024, doi:
https://doi.org/10.1109/ACCESS.2024.346115010.1109/ACCESS.2024.3461150.

[9] S. Mandal and S. K. Naskar, "Classifying and Solving Arithmetic Math Word Problems—An Intelligent Math Solver," *IEEE Transactions on Learning Technologies*, vol. 14, no. 1, pp. 28–41, Feb. 2021, doi:
https://doi.org/10.1109/TLT.2021.3057805 10.1109/TLT.2021.3057805.

[10] Y. Zhang, G. Zhou, Z. Xie, and J. X. Huang, "HGEN: Learning Hierarchical Het-erogeneous Graph Encoding for Math Word Problem Solving," *IEEE/ACM Trans- actions on Audio, Speech, and Language Processing*, vol. 30, pp. 816–828, 2022, doi:
https://doi.org/10.1109/TASLP.2022.314531410.1109/TASLP.2022.3145314.

[11] B. R. Christ, J. Kropko, and T. Hartvigsen, "MATHWELL: Generating Educational Math Word Problems Using Teacher Annotations," Sep. 27, 2024, arXiv: arXiv:2402.15861. Accessed: Nov. 13, 2024. [Online]. Available:
http://arxiv.org/abs/2402.15861http://arxiv.org/abs/2402.15861.