

Complex Mathematical Expression Recognition: Benchmark, Large-Scale Dataset and Strong Baseline

Weikang Bai^{1,*}, Yongkun Du^{1,*}, Yuchen Su¹, Yazhen Xie¹, Zhineng Chen^{2,†}

¹College of Computer Science and Artificial Intelligence, Fudan University, Shanghai, China

²Institute of Trustworthy Embodied AI, Fudan University, Shanghai, China

{wkbai24, ykdu23, ycsu23, yzxie24}@m.fudan.edu.cn, zhinchen@fudan.edu.cn

Abstract

Mathematical Expression Recognition (MER) has made significant progress in recognizing simple expressions, but the robust recognition of complex mathematical expressions with many tokens and multiple lines remains a formidable challenge. In this paper, we first introduce CMER-Bench, a carefully constructed benchmark that categorizes expressions into three difficulty levels: easy, moderate, and complex. Leveraging CMER-Bench, we conduct a comprehensive evaluation of existing MER models and general-purpose multimodal large language models (MLLMs). The results reveal that while current methods perform well on easy and moderate expressions, their performance degrades significantly when handling complex mathematical expressions, mainly because existing public training datasets are primarily composed of simple samples. In response, we propose MER-17M and CMER-3M that are large-scale datasets emphasizing the recognition of complex mathematical expressions. The datasets provide rich and diverse samples to support the development of accurate and robust complex MER models. Furthermore, to address the challenges posed by the complicated spatial layout of complex expressions, we introduce a novel expression tokenizer, and a new representation called Structured Mathematical Language, which explicitly models the hierarchical and spatial structure of expressions beyond LaTeX format. Based on these, we propose a specialized model named CMERNet, built upon an encoder-decoder architecture and trained on CMER-3M. Experimental results show that CMERNet, with only 125 million parameters, significantly outperforms existing MER models and MLLMs on CMER-Bench.

Code — <https://github.com/Baitlo/CMER>

1 introduction

Mathematical Expression Recognition (MER) has long been a fundamental task for visual-based document parsing and understanding. Unlike Optical Character Recognition (OCR) that primarily deals with plain text (Du et al. 2022; Zheng et al. 2024; Du et al. 2025a,b,c), MER poses unique challenges due to its two-dimensional spatial layout,

*These authors contributed equally.

†Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

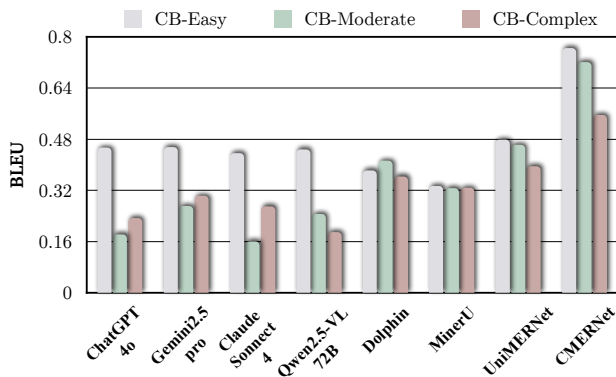


Figure 1: Comparing our CMERNet with existing models in complex mathematical expression recognition. CB means CMER-Bench, our constructed benchmark.

diverse symbols, complex syntax, and high structural ambiguity. In recent years, a variety of specialized methods have emerged to tackle this problem, such as (Deng et al. 2017; Yuan et al. 2022; Zhu et al. 2024; Wang et al. 2024a; Li et al. 2025), which have significantly advanced MER.

Existing approaches can be broadly classified into two paradigms: the first is the image-to-sequence interpretation, which treats MER as a direct image-to-LaTeX translation task. One representative work in this category is UniMER (Wang et al. 2024a), which leverages a custom visual encoder and an LLM decoder to autoregressively generate a LaTeX string from the input image. This work also proposes a dataset over one million scientific expressions for model training (Wang et al. 2024a). The second paradigm focuses on explicitly modeling the 2D structure of expressions, typically applied to handwritten MER. A prominent example is TDv2 by Wu et al.. Instead of generating a linear LaTeX string, TDv2 employs a tree-structured decoder that autoregressively constructs a syntax tree for the expression. At each step, it decouples the prediction into two sub-tasks: a Node Classification Module to identify the next symbol, and a Branch Prediction Module to determine its spatial relationship (e.g., superscript, below) to its parent node. In this manner, domain knowledge of expressions is incorporated for recognition.

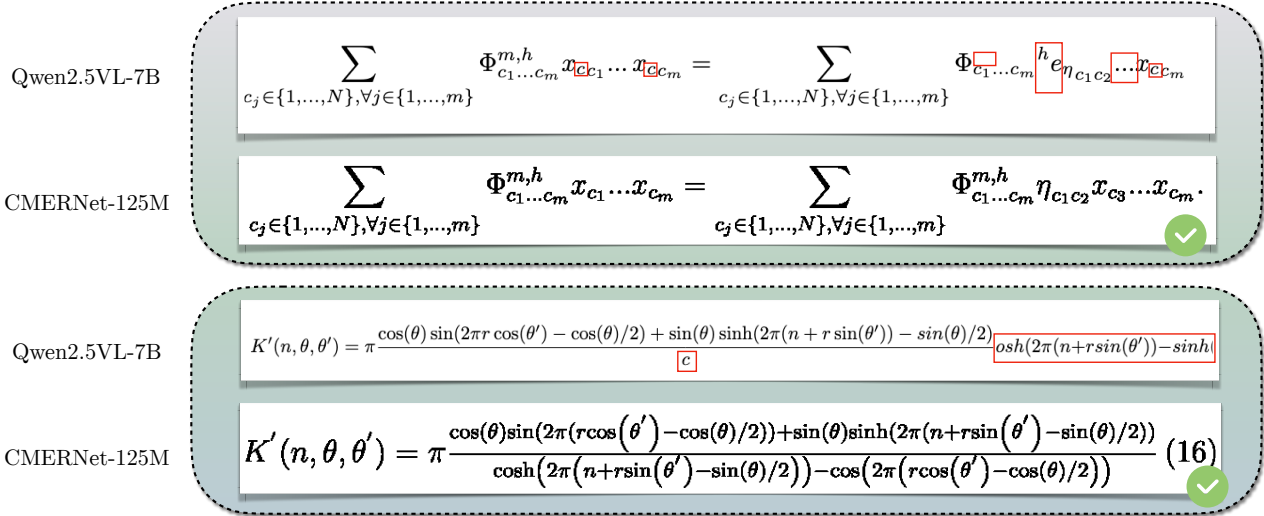


Figure 2: Visualization comparison of MER results between Qwen2.5VL and our CMERNet.

Additionally, recent advances built on Multimodal Large Language Models (MLLMs) have also made remarkable progress in MER (Wang et al. 2024a; Blecher 2021). However, due to the simple nature of current MER benchmarks (Deng et al. 2017; Wang et al. 2024a), both specialized methods and general-purpose MLLMs are biased towards mathematical expressions of limited complexity, typically involving dozens of tokens and a single-line layout. As shown in Fig. 2, these methods falter when confronted with complex, multi-line expressions commonly seen in scientific documents. Motivated by this observation, we argue that current MER models face great challenges in complex scenarios, which can be summarized as the following three fundamental challenges:

Lack of Challenging Benchmarks. Existing benchmarks predominantly focus on evaluating simple expressions. There is a lack of complex expressions containing hundreds of tokens, multiple lines, intricate spatial layout, diverse symbols, etc. These features are essential for a comprehensive assessment of MER models.

Training Data Scarcity and Skewness. MER currently suffers from a scarcity of large-scale training data. Existing datasets are not only limited in size but are also heavily skewed towards simple expressions. The distribution of expression length is highly imbalanced, with a critical underrepresentation of long and complex examples, which are critical to train robust MER models.

Data Representation Dilemma. Representing mathematical expressions faces a fundamental trade-off. On one hand, generating linear LaTeX string-based representation is straightforward. However, it complicates the encoding of the complex layout of mathematical expressions. On the other hand, syntax tree-based representation, while intuitive, poses significant integration challenges with the now dominant, autoregressive decoding models.

Facing these challenges, in this paper we present three

contributions to complex MER. Firstly, we introduce CMER-Bench, a benchmark with a clear difficulty stratification, designed to enable rigorous evaluation of model performance across Easy, Moderate, and Complex levels. Here Easy roughly corresponds to the difficulty of existing datasets. Secondly, we introduce two training datasets, CMER-3M and MER-17M, the former being the most complex and the latter the largest MER training dataset to our knowledge. Both are larger than existing MER training datasets to our knowledge. We have made it publicly available to support the development MER models for the community. Thirdly, we develop CMERNet, an encoder-decoder-based baseline with new tokenizer and representation. When trained on CMER-3M and evaluated on CMER-Bench, CMERNet achieves state-of-the-art (SOTA) performance across a majority of evaluation metrics, providing the community with a powerful model to comprehensively address the challenges of complex MER.

2 Related Work

Early deep learning-based methods treat MER as an image-to-sequence interpretation task, typically using an attention-based CNN encoder and an RNN decoder (Zhang et al. 2017; Deng et al. 2017). IM2LATEX (Deng et al. 2017) exemplifies this image-to-sequence approach and introduces the IM2LATEX-100K dataset to support MER model training. This architecture is later enhanced by Pix2tex (Blecher 2021), which replaces the CNN with a Vision Transformer (ViT) (Dosovitskiy et al. 2021) and the RNN with a Transformer decoder to capture rich global dependencies. To address the issue of lacking training data, Wang et al. introduces the large-scale UniMER-1M dataset with 1M instances. They also develop the UniMERNet model that achieves significant performance gains when trained on the dataset. These methods are not good at describing the spatial layout of expressions. The linear nature of LaTeX string

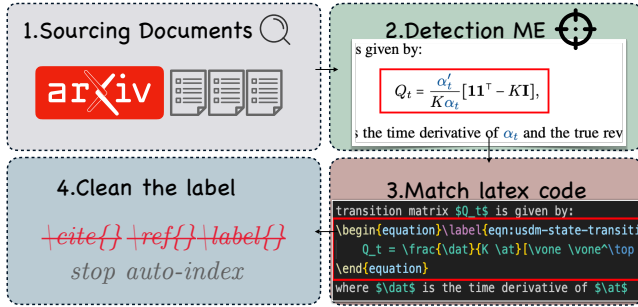


Figure 3: The pipeline of constructing CMER-3M and MER-17M.

Length	IM2LATEX	Pix2tex	UniMER	CMER-3M	MER-17M
0 – 20	<0.1K	3.6K	477.8K	85.9K	85.9K
21 – 150	43.4K	137.5K	342.6K	620.2K	5.5M
151 – 300	28.1K	70.6K	103.1K	858.1K	7.8M
301 – 450	3.6K	16.8K	62.6K	855.6K	3.0M
> 450	0.1K	5.8K	75.6K	641.2K	1.4M
Total	75.3K	233.8K	1.1M	3.1M	17.7M

Table 1: Statistics of current MER datasets and ours, categorized by the length of LaTeX sequences.

introduces ambiguity upon tokenization, making structural tokens like `{`, `}`, and `\frac` challenging to interpret.

Consequently, some research has shifted towards predicting more structured formats, such as syntax trees. TDv2 (Wu et al. 2022) introduces a tree-structured decoder that autoregressively constructs a formula’s syntax tree. It decouples each generation step into two parallel sub-tasks: to predict the symbol’s identity and to determine its spatial relationship to the parent. Another tree-based method, TAMER (Zhu et al. 2024), enhances the decoder by jointly modeling the symbol sequence and its hierarchical structure through tree-aware attention. However, both methods are mainly suitable for expressions with relatively simple layouts and fewer tokens, such as handwritten mathematical expressions (Mouchère et al. 2013, 2016; Xie et al. 2023).

3 Complex MER Datasets and Benchmark

3.1 CMER-3M and MER-17M Datasets

Existing public datasets for MER suffer from clear limitations, including insufficient data scale, mainly composed of structurally simple and single-line expressions, etc, restricting their usage as comprehensive MER training corpora. To overcome the issues, we introduce MER-17M, a large-scale dataset for MER, and its subset CMER-3M that emphasizes both distribution balance and data complexity. MER-17M is constructed through a multi-stage pipeline that begins with sourcing over 1 million scientific documents from repositories like ArXiv (Cornell University 2025). We extract mathematical expressions using advanced parsing tools to isolate LaTeX code snippets, clean those visual-semantic misalignment tags like `\cite`, which confuse the model. After that

Lines	IM2LATEX	Pix2tex	UniMER	CMER-3M	MER-17M
1	71.7K	221.8K	964.7K	505.7K	504.3K
2	0.3K	6.0K	45.6K	2.0M	15.6M
3	2.0K	3.6K	23.7K	407.4K	1.3M
4	0.4K	1.6K	13.1K	52.5K	143.3K
5	0.5K	0.8K	5.8K	25.2K	60.6K
> 5	0.3K	1.1K	8.9K	21.9K	43.5K

Table 2: Expression distribution based on the number of lines they occupied.

we render them into images. A rigorous filtering process ensures its quality: we discard malformed or overly simplistic expressions and apply diversity checks. As a result, we obtain over 17 million expression image-LaTeX pairs. We choose 2000 of them as the evaluation benchmark, and the collection of the rest is termed MER-17M.

Since the MER-17M distribution is not balanced, we further construct CMER-3M, a subset of MER-17M, which is more balanced across expression length. CMER-3M is featured in the following three key dimensions:

Unprecedented scale. CMER-3M largely surpasses prior datasets in instance quantity, containing over 3 million samples—nearly 3 times larger than UniMER (1M), the current largest, and much larger than IM2LATEX (100K) or Pix2tex (234K), as quantified in Table 1. This scale ensures many rare symbols and intricate patterns are included, building a strong foundation for robust model learning.

Balanced length distribution. Unlike existing datasets biased towards short-length (in LaTeX token) expressions, CMER-3M is more uniform in length distribution. As illustrated in Table 1, CMER-3M contains much more medium- and long-length expressions. Since length is the most important indicator describing the complexity of expressions, the existence of rich complex expressions makes complex MER studies no longer troubled by training data.

High structural complexity. As shown in Table 2, CMER-3M has more than 2.5M multi-line expressions, accounting for 83% of its samples. The ratio significantly exceeds existing datasets, especially in expressions with three or more lines. The statistics also indicate that MER-17M is rich in diverse and complex layouts, which is also a critical distinction to existing datasets.

3.2 CMER-Bench Benchmark

Complex MER is also hindered by the scarcity of comprehensive evaluation benchmarks. Existing test sets are heavily skewed towards simple, single or two-line, or synthetic expressions. These sets only partially evaluate the capability of existing models, where their ability in handling complex MER has not been fully assessed.

To address this issue, we introduce CMER-Bench, a benchmark meticulously curated to facilitate a comprehensive and fair assessment of MER models. While also derived from the multi-stage pipeline described in Sec. 3.1, the composition of CMER-Bench is determined through a rigorous hybrid process, combining automated difficulty scoring with

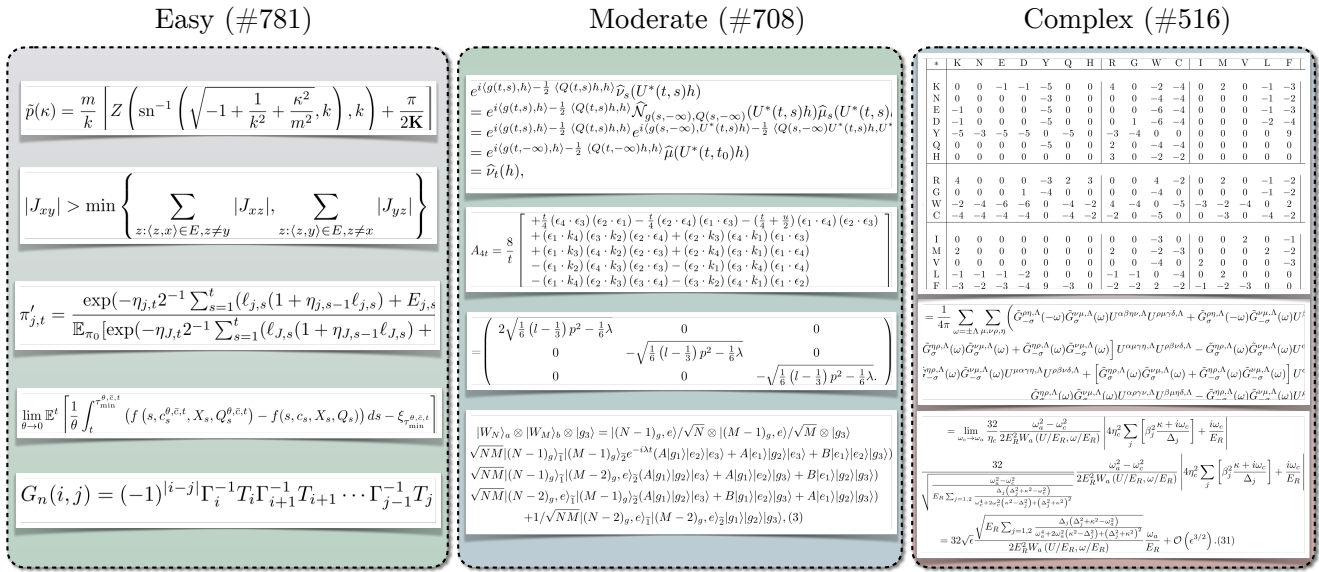


Figure 4: Showcase of CMER-Bench examples according to their difficulty tiers.

human expert verification to ensure both the quality and representativeness of selected samples. Following this criterion, 2000 samples are selected.

We divide CMER-Bench into three distinct subsets: Easy, Moderate, and Complex, indicating increasing recognition difficulties. Examples and numbers of each tier are shown in Figure 4. This classification enables us to observe the performance of different MER models along the difficulty axis, supporting more comprehensive model assessment.

4 CMERNet

To develop a general yet powerful MER model, we propose specialized tokenizer, expression representation, and modeling methods as follows.

4.1 Specialized Mathematical Tokenizer

As shown by studies in scientific text domain (Taylor et al. 2022), using custom tokenizer to parse the syntax of specified data can obtain better results. Mathematical expressions have their own syntax. However, no specialized tokenizer for MER has been proposed to date. Using general-purpose tokenizers often likely fragment critical mathematical commands (e.g., splitting $\sqrt{\quad}$ into \backslash, s, q, r, t), which disrupts semantic integrity and complicates model learning.

To address this, we develop a specialized Byte Pair Encoding (BPE) tokenizer. It parses the expression syntax following a two-step process. First, we curate a comprehensive vocabulary of common LaTeX environments and commands, such as $\sqrt{\quad}$, $\frac{\quad}{\quad}$, and $\begin{gathered}$. These were explicitly designated as special tokens, ensuring that they are treated as atomic units and preventing fragmentation. Second, we trained the BPE tokenizer on our CMER-3M dataset. This process yields a highly efficient mathematical tokenizer that preserves structural semantics well.

4.2 Structured Mathematical Language (SML)

To overcome the inherent limitation of linear LaTeX string-based representation, we introduce a novel representation called SML that explicitly injects structural information into the expression representation such that benefits learning process. Our goal is to create a target sequence that both preserves the rich, two-dimensional grammar of mathematical expressions, and seamlessly integrates these with common autoregressive decoder architectures. To achieve this, we have developed a two-step conversion scheme as follows:

Parsing the LaTeX string into the syntax tree. First, we parse the raw LaTeX string into a hierarchical syntax tree. This tree explicitly captures the nested relationships and the parent-child grammar between operators and operands, e.g., identifying the numerator and denominator within a fraction. This tree encodes the spatial structure of the expression.

Serializing the tree into a structured token sequence by additional grammar tokens. Then, we traverse this syntax tree to serialize it into a sequence of tokens. Note that this is not a flat representation like LaTeX string. Instead, the token sequence directly encodes the tree’s structure using additional designated tokens to signify parent-child relationships, node types, and sibling order.

The SML representation offers two advantages. First, it resolves the representation bottleneck of LaTeX by providing the model with a structure-aware target that makes the correspondence between visual layout and token semantics explicit. Second, because the final output is still a linear sequence of tokens—albeit a structurally meaningful one—it aligns seamlessly with the dominant autoregressive decoder. This makes our method highly scalable and allows it to directly benefit from advances in current leading models, facilitating our model design.

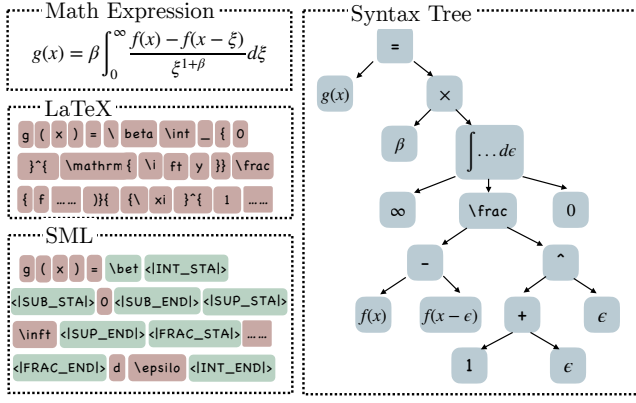


Figure 5: An example of syntax tree-based expression representation. Note that the raw LaTeX string can be represented as SML following the syntax tree.

4.3 Dynamic Resolution Pre-processing

Mathematical expressions have rich fine-grained details and often non-standard, portrait-oriented aspect ratios. Thus, they are highly susceptible to distortion when using conventional fixed-sized image pre-processing, e.g., downsampling. To address this, we introduce CMER-Fit, a dynamic resolution adjustment module inspired by NaViT (Dehghani et al. 2023) and SigLip2 (Tschannen et al. 2025).

CMER-Fit’s objective is to render each image using the highest resolution that adheres to a predefined token budget (e.g., 256 patches) required by the vision encoder. Given an original image of size $H_0 \times W_0$ and a patch size $p \times p$, we seek the greatest scale factor s^* such that the resulting patch grid does not exceed N_{\max} elements, H^* and W^* are the adjusted height and width, respectively.

$$\left\lceil \frac{s H_0}{p} \right\rceil \times \left\lceil \frac{s W_0}{p} \right\rceil \leq N_{\max} \quad (1)$$

$$H^* = p \left\lceil \frac{s^* H_0}{p} \right\rceil, \quad W^* = p \left\lceil \frac{s^* W_0}{p} \right\rceil \quad (2)$$

By using dynamic resolution resizing, we can explicitly measure the distortion caused by this resizing, which we define as Maximum Distortion Ratio (MDR) as follows.

$$\text{MDR} = 1 - \frac{H^* W^*}{H_0 W_0} \quad (3)$$

Note that we only consider downsampling-based resizing. That is, if $H_0 W_0 \leq p^2 N_{\max}$, this means $s^* \geq 1$ and corresponds to upsampling. In this case, no resizing occurs and $\text{MDR} = 0$.

Because each side is rounded up to a multiple of p , at most $p - 1$ surplus pixels are introduced along each axis. Hence

$$0 \leq \text{MDR} = 1 - \frac{H^* W^*}{H_0 W_0} < \frac{(p - 1)(H_0 + W_0)}{H_0 W_0}. \quad (4)$$

Setting $H_0 = W_0 \gg p$ gives the theoretical upper bound

$$\text{MDR}_{\max} < \frac{2(p - 1)}{H_0}. \quad (5)$$

Patch size p	Typical input $H_0 = W_0$	MDR_{\max}
16	1024	$< 2.93\%$
32	1024	$< 6.05\%$

Table 3: Maximum distortion ratio under different patch sizes with respect to 1024×1024 input.

Two practical examples shown in Table 3 illustrate the scale of this bound. It is observed that the distortion is negligible for high-resolution images.

CMER-Fit offers a parameter-free, theoretically sound mechanism for harmonizing token budget, memory overhead, and spatial fidelity. Its provably small worst-case distortion and negligible computational cost make it a robust choice for large-scale training and deployment under stringent hardware constraints.

4.4 CMERNet: A Strong Baseline

Building on the above techniques, we design CMERNet, a strong baseline model for complex MER. We develop specialized architecture to meet two key characteristics inherent in MER: capturing fine-grained visual details (e.g., dots, primes, small symbols) and understanding the complex, hierarchical 2D grammar of expressions. To this end, CMERNet is composed of three core components: a hybrid vision encoder, a feature-refining connector, and an autoregressive decoder. By incorporating SML, CMERNet could better understand mathematical expressions.

To effectively process both local details and global structure, our vision encoder adopts a CNN-Transformer hybrid design. Diverging from the traditional OCR model SVTRv2 (Du et al. 2025c), and more complex interleaved architectures like CvT (Wu et al. 2021) or task-specific models like DETR (Carion et al. 2020), our architecture explores the complementarity of well-established CNN and Transformer, leveraging their respective strengths in a two-step process:

Step 1: CNN for fine-grained feature capture. The vision encoder first employs a shallow CNN backbone composed of six stacked residual blocks. These convolutions are deliberately designed without downsampling in its early layers to preserve a high-resolution feature map. This is critical for capturing the precise location and appearance of fine-grained symbols like superscripts, subscripts, and diacritics, which are often lost with aggressive initial downsampling.

Step 2: Transformer for global context modeling. The CNN feature map is then passed to the Transformer part, where 12 standard Transformer layers are employed. This step is responsible for modeling the global context, allowing the model to understand the long-range dependencies and complex spatial relationships between all symbols in the expression.

Then, in the feature-refining connector, we adopt two MLP layers commonly used in MLLMs to align the image features extracted by vision encoder with the linguistic representations of mathematical formulas. Finally, an autoregressive decoder models these image features and decodes them into the SML representation. The decoder is built upon

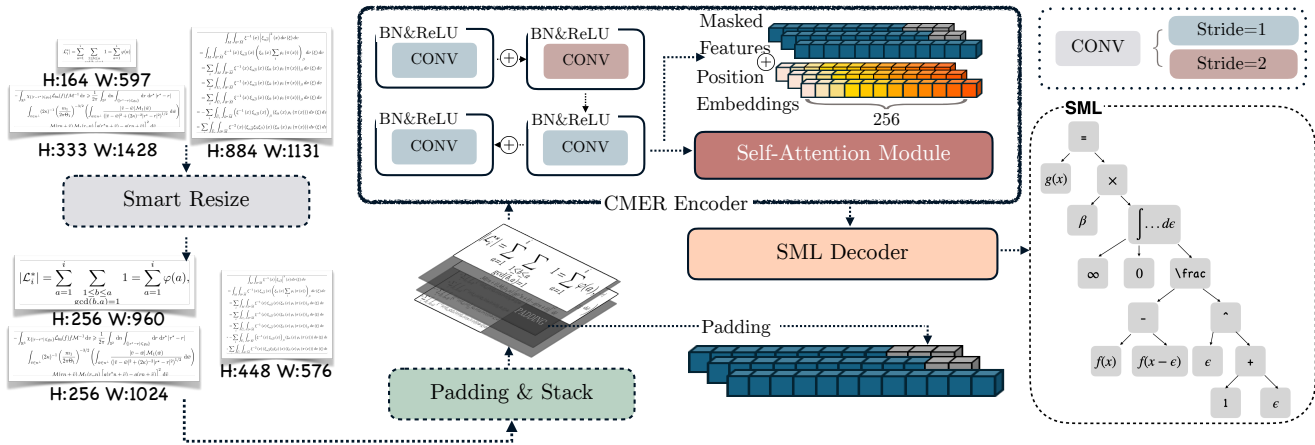


Figure 6: The architecture of CMERNet. Our model employs a hybrid vision encoder that sequentially combines a CNN for fine-grained feature extraction with a Transformer backbone for global context modeling. This design is tailored to robustly process complex mathematical expressions with both intricate local details and long-range structural dependencies.

a cross-attention-based Transformer architecture.

5 Experiments

5.1 Implementation Details

Training data. The model is trained on our CMER-3M and a collection of existing datasets. To enhance the model’s robustness and equip it with the capability to recognize handwritten mathematical expressions, in addition to CMER-3M, we augment our training data by incorporating samples from CROHME2013 (Mouchère et al. 2013), CROHME2016 (Mouchère et al. 2016), CROHME2019 (Mahdavi et al. 2019) and CROHME2023 (Xie et al. 2023), along with other publicly available mathematical formula corpora such as HME-100K (Yuan et al. 2022) and UniMER-1M (Wang et al. 2024a).

Training Details. We use the AdamW optimizer with default parameters ($\beta_1 = 0.9, \beta_2 = 0.999$). The training process spanned one full epoch over the dataset. We adopt a learning rate schedule with a linear warmup phase: the learning rate is increased from an initial value of 1×10^{-5} to a peak of 1×10^{-4} over the first 3,000 steps. Following the warmup, the learning rate is decayed to 1×10^{-9} using a cosine annealing schedule for the remainder of the training. All experiments are conducted on a single server equipped with eight NVIDIA RTX 4090 GPUs. We employ Distributed Data Parallel (DDP) to accelerate training, setting a per-GPU batch size of 8, which results in a global batch size of 64. The whole training takes around 3 days. We train CMERNet from scratch using PyTorch. The whole model, comprising approximately 125 million parameters, is trained end-to-end.

5.2 Evaluation Metrics

BLEU (Bilingual Evaluation Understudy). BLEU (Papineni et al. 2002) computes sequence similarity using n-gram precision between a prediction and reference, supplemented by a brevity penalty to penalize short hypotheses. Although

effective at capturing local correctness, its high sensitivity to the exact token sequence makes it less robust to synonymous LaTeX representation.

ROUGE (Recall-Oriented Understudy for Gisting Evaluation). We use ROUGE (Lin 2004), reporting ROUGE-1 (unigram recall), ROUGE-2 (bigram recall), and ROUGE-L (longest common subsequence). ROUGE-L specifically rewards in-order, non-contiguous matches. As a string-level metric, it shares BLEU’s susceptibility to the one-to-many mapping problem.

Levenshtein Distance. Also known as edit distance (Halдар and Mukhopadhyay 2011), this metric counts the minimum token-level insertions, deletions, and substitutions between prediction and reference. While intuitive (lower is better), its strict token-by-token comparison fails to recognize the structural or visual equivalence of synonymous LaTeX commands.

CDM (Canonicalization-based Distance Measure). Proposed by (Wang et al. 2025), CDM is a recent metric designed specifically to address one of many problems in MER. It first adds a color label in both predictions and labels, aiming to get the bounding box of each token in expressions. Then CDM apply matching algorithm to calculate f1 score and recall score. CDM can complement metrics such as BLEU to some extent.

BLEU and ROUGE are susceptible to length bias due to their n-gram match mechanism, which can obscure performance discrepancies, especially on complex expressions where similar scores may not reflect equal quality. We find that CDM, which evaluates structural element-wise similarity, can mitigate this issue. Therefore, we employ all these metrics, believing that this multi-faceted approach, balancing token overlap (BLEU/ROUGE) with structural fidelity, more accurately reflects true performance.

5.3 Result Analysis

We evaluate CMERNet against the baselines, including general-purpose MLLMs such as GPT-4o and Gemini-

Model	Difficulty Level			Performance Metrics					
	Easy	Moderate	Complex	ROUGE-1↑	ROUGE-2↑	ROUGE-L↑	BLEU↑	Avg Edit Dist. ↓	CDM ↑
UniMERNet (2024a)	✓			0.7464	0.6396	0.7388	0.4758	94.6978	0.938
		✓		0.68	0.5449	0.6703	0.4575	133.2556	0.866
			✓	0.7605	0.6126	0.7502	0.4093	483.34	0.68
Gemini-2.5-pro (2025)	✓			0.7402	0.6353	0.738	0.455	97.5134	0.862
		✓		0.5585	0.4029	0.5553	0.266	192.4393	0.862
			✓	0.6589	0.5424	0.6528	0.3089	1282.1165	0.856
ChatGPT-4o (2024)	✓			0.7166	0.5972	0.7135	0.4517	145.306	0.888
		✓		0.5418	0.3792	0.5372	0.185	200.6257	0.83
			✓	0.5977	0.486	0.5904	0.2082	1905.0465	0.635
Claude-Sonnet-4 (2025)	✓			0.7126	0.5996	0.7099	0.4341	131.3675	0.898
		✓		0.4905	0.3427	0.4864	0.1589	206.0297	0.747
			✓	0.6471	0.5356	0.6426	0.2623	1537.9671	0.671
Qwen-VL-72b (2025)	✓			0.7048	0.5979	0.7023	0.4477	133.525	0.796
		✓		0.5371	0.3739	0.5279	0.2433	200.2345	0.742
			✓	0.5932	0.4935	0.5872	0.1915	2107.3029	0.792
MinerUv2 (2024b)	✓			0.6238	0.4828	0.6218	0.333	144.1306	0.936
		✓		0.5716	0.4106	0.5676	0.3218	187.2429	0.866
			✓	0.6938	0.525	0.6863	0.3218	657.1495	0.63
Dolphin (2025)	✓			0.5911	0.4819	0.5832	0.3806	146.0141	0.901
		✓		0.6911	0.577	0.6689	0.4223	452.0213	0.886
			✓	0.6791	0.5488	0.6706	0.3805	776.9899	0.716
CMERNet (Ours)	✓			0.8626	0.6613	0.8611	0.7653	51.61	0.966
		✓		0.8266	0.6072	0.8181	0.7215	88.09	0.845
			✓	0.7750	0.5764	0.7453	0.5568	474.18	0.690

Table 4: Performance comparison of various models on CMER-Bench covering three difficulty levels. Our CMERNet demonstrates superior performance, especially on complex expressions. Best results in each column and difficulty level are in **bold**.

2.5-pro, state-of-the-art document parsing systems such as MinerU and Dolphin, and the dedicated MER method UniMERNet. This selection ensures a comprehensive comparison across different architectures and training objectives.

The results are listed in Table 4. The baselines perform poorly on the complex tier. All compared methods, especially general-purpose MLLMs like ChatGPT-4o and Gemini, suffer from a drastic performance drop compared to our CMERNet. The BLEU scores for Gemini-2.5-pro, ChatGPT-4o, and QWen-VL-72b are only 0.3089, 0.2082, and 0.1915 on the complex tier, respectively. This confirms their inadequacy in handling intricate, multi-line structures. Furthermore, the expert model UniMERNet and document parsing models (MinerUv2 and Dolphin) get slightly better results, mainly because they are dedicatedly trained by expression data. Nevertheless, CMERNet, by training using a large number of mathematic expressions, as well as appropriately handling the tokenizer and representation issue, excelling the second best method in Table 4 nearly 15%-29% in the BLEU score under different difficulty tiers. Note that the CDM metric does not exhibit so large gaps, this is because CDM mainly measures the bounding box-level matching, which hardly perceive subtle mistakes. Nevertheless, the results clearly validate the effectiveness of our CMERNet.

6 Conclusion

In this work, our aim is to address the challenge of complex MER, a task that has been largely overlooked previously. Observing that complex MER still lacks comprehensive evaluation benchmark and large-scale training dataset, we first introduce CMER-Bench, a benchmark composed of easy, moderate and complex expressions, which enables the detailed assessment of the MER ability of different models, and for the first time, covers the evaluations of complex expressions. Then, we construct CMER-3M and MER-17M, both excel existing datasets in instance quantity and diversity. We also prove that the dataset is more suitable to train robust MER models. Furthermore, we have investigated the expression tokenizer and representation, and develop a specialized MER model termed CMERNet. Experimental results demonstrate that CMERNet, despite only 125M parameters, has shown strong performance compared to general-purpose MLLMs, document parsing models and specialized MER model. The results basically validate the importance of introducing CMER-Bench, CMER-3M and MER-17M, and the effectiveness of our CMERNet. In future, we plan to conduct more experiments to further verify CMERNet, as well as exploring using the constructed datasets to strengthen the MER capability of popular MLLMs.

Acknowledgments

This work was supported by National Natural Science Foundation of China (No. 32341012).

References

- Anthropic. 2025. Models Overview: Claude Sonnet 4. Anthropic Documentation.
- Blecher, L. 2021. pix2tex: LaTeX-OCR – Using a ViT to Convert Images of Equations into LaTeX Code. <https://github.com/lukas-blecher/LaTeX-OCR>.
- Carion, N.; Massa, F.; Synnaeve, G.; et al. 2020. End-to-End Object Detection with Transformers. *arXiv preprint arXiv:2005.12872*.
- Cornell University. 2025. arXiv: e-Print Archive. <https://arxiv.org>.
- Dehghani, M.; Mustafa, B.; Djolonga, J.; et al. 2023. Patch n’ Pack: NaViT, a Vision Transformer for Any Aspect Ratio and Resolution. *arXiv preprint arXiv:2307.06304*.
- Deng, Y.; Kanervisto, A.; Ling, J.; and Rush, A. M. 2017. Image-to-Markup Generation with Coarse-to-Fine Attention. In *ICML*, 980–989.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; et al. 2021. An Image Is Worth 16x16 Words: Transformers for Image Recognition at Scale. *arXiv preprint arXiv:2010.11929*.
- Du, Y.; Chen, Z.; Jia, C.; Yin, X.; Li, C.; Du, Y.; and Jiang, Y.-G. 2025a. Context perception parallel decoder for scene text recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 47(6): 4668–4683.
- Du, Y.; Chen, Z.; Jia, C.; Yin, X.; Zheng, T.; Li, C.; Du, Y.; and Jiang, Y.-G. 2022. SVTR: Scene Text Recognition with a Single Visual Model. In *IJCAI*, 884–890.
- Du, Y.; Chen, Z.; Su, Y.; Jia, C.; and Jiang, Y.-G. 2025b. Instruction-guided scene text recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 47(4): 2723–2738.
- Du, Y.; Chen, Z.; Xie, H.; Jia, C.; and Jiang, Y.-G. 2025c. SVTRv2: CTC Beats Encoder-Decoder Models in Scene Text Recognition. In *ICCV*, 20147–20156.
- Feng, H.; Wei, S.; Fei, X.; et al. 2025. Dolphin: Document Image Parsing via Heterogeneous Anchor Prompting. *arXiv preprint arXiv:2505.14059*.
- Gemini Team. 2025. Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities. *arXiv preprint arXiv:2507.06261*.
- Haldar, R.; and Mukhopadhyay, D. 2011. Levenshtein Distance Technique in Dictionary Lookup Methods: An Improved Approach. *arXiv preprint arXiv:1101.1232*.
- Hurst, A.; et al. 2024. GPT-4o System Card. Technical report, OpenAI.
- Li, Z.; Wei, J.; Shen, Z.; Ma, C.; Wu, Y.; and Zhou, Y. 2025. PACM: Position-Aware Cross-Modality Decoder for Handwritten Mathematical Expression Recognition. In *ICDAR*, 96–114.
- Lin, C. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *ACL Workshop on Text Summarization*, 74–81.
- Mahdavi, M.; Zanibbi, R.; Mouchère, H.; Viard-Gaudin, C.; and Garain, U. 2019. ICDAR 2019 CROHME+TFD: Competition on Recognition of Handwritten Mathematical Expressions and Typeset Formula Detection. In *ICDAR*, 1533–1538.
- Mouchère, H.; Viard-Gaudin, C.; Zanibbi, R.; and Garain, U. 2016. ICFHR2016 CROHME: Competition on Recognition of Online Handwritten Mathematical Expressions. In *ICFHR*, 607–612.
- Mouchère, H.; Viard-Gaudin, C.; Zanibbi, R.; Garain, U.; Kim, D.; and Kim, J. 2013. ICDAR 2013 CROHME: Third International Competition on Recognition of Online Handwritten Mathematical Expressions. In *ICDAR*.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W. 2002. BLEU: A Method for Automatic Evaluation of Machine Translation. In *ACL*, 311–318.
- Taylor, R.; Kardas, M.; Cucurull, G.; et al. 2022. Galactica: A Large Language Model for Science. *arXiv preprint arXiv:2211.09085*.
- Team, Q. 2025. Qwen2.5-VL. <https://qwenlm.github.io/blog/qwen2.5-vl/>.
- Tschannen, M.; Gritsenko, A.; Wang, X.; et al. 2025. SigLIP 2: Multilingual Vision-Language Encoders with Improved Semantic Understanding, Localization, and Dense Features. *arXiv preprint arXiv:2502.14786*.
- Wang, B.; Gu, Z.; Liang, G.; Xu, C.; Zhang, B.; Shi, B.; and He, C. 2024a. UniMERNET: A Universal Network for Real-World Mathematical Expression Recognition. *arXiv preprint arXiv:2404.15254*.
- Wang, B.; Wu, F.; Ouyang, L.; Gu, Z.; Zhang, R.; Xia, R.; Shi, B.; Zhang, B.; and He, C. 2025. Image Over Text: Transforming Formula Recognition Evaluation with Character Detection Matching. In *CVPR*, 19681–19690.
- Wang, B.; Xu, C.; Zhao, X.; et al. 2024b. MinerU: An Open-Source Solution for Precise Document Content Extraction. *arXiv preprint arXiv:2409.18839*.
- Wu, C.; Du, J.; Li, Y.; Zhang, J.; Yang, C.; Ren, B.; and Hu, Y. 2022. TDv2: A Novel Tree-Structured Decoder for Offline Mathematical Expression Recognition. In *AAAI*, volume 36.
- Wu, H.; Xiao, B.; Codella, N.; et al. 2021. CvT: Introducing Convolutions to Vision Transformers. *arXiv preprint arXiv:2103.15808*.
- Xie, Y.; Mouchère, H.; Liwicki, F. S.; et al. 2023. ICDAR 2023 CROHME: Competition on Recognition of Handwritten Mathematical Expressions. In *ICDAR*, 553–565.
- Yuan, Y.; Liu, X.; Dikubab, W.; et al. 2022. Syntax-Aware Network for Handwritten Mathematical Expression Recognition. *arXiv preprint arXiv:2203.01601*.
- Zhang, J.; Du, J.; Zhang, S.; Liu, D.; Hu, Y.; Hu, J.; Wei, S.; and Dai, L. 2017. Watch, attend and parse: An end-to-end neural network based approach to handwritten mathematical expression recognition. *Pattern Recognition*, 196–206.
- Zheng, T.; Chen, Z.; Fang, S.; Xie, H.; and Jiang, Y.-G. 2024. Cdistnet: Perceiving multi-domain character distance for robust text recognition. *IJCV*, 132(2): 300–318.

Zhu, J.; Zhao, W.; Li, Y.; Hu, X.; and Gao, L. 2024. TAMER: Tree-Aware Transformer for Handwritten Mathematical Expression Recognition. *arXiv preprint arXiv:2408.08578*.