

Ryze: Evidence-Enriched Data Synthesis from Biomedical Papers

Yeqi Huang¹ Yue Chen¹ Yanwei Ye¹ Guan hao Su¹ Luo Mai¹

¹University of Edinburgh, United Kingdom

Demo: <https://ryze.12th.day>

Youtube Link: <https://youtu.be/5L2YShSaIQQ>

GitHub: <https://github.com/Chivier/Ryze>

Abstract

General-purpose VLMs remain unreliable for biomedical research because valid answers in scientific papers depend on evidence split across figures, tables, charts, captions, and referring text. Existing post-training pipelines are bottlenecked by costly expert annotation and by synthetic data that drops this evidence structure. We present Ryze, a fully automated system that converts raw biomedical papers into an *evidence-enriched* training set and a domain-specialized VLM. Ryze synthesizes QA pairs with complete supporting evidence (visual element, caption, extracted structure, and referring paragraphs), reduces layout and OCR errors via chart/table-aware extraction and LLM-based cleansing, and applies a progress-gated post-training strategy combining supervised fine-tuning with reinforcement learning. Starting from Qwen3-VL-8B, Ryze produces BioVLM-8B at under \$200, achieving 48.0% weighted accuracy on LAB-Bench—outperforming the base model by +12.6% and surpassing GPT-5.2 by +3.8%. We release Ryze as open source together with the trained BioVLM-8B model.

1 Introduction

General-purpose vision-language models (VLMs) can read and reason over everyday images and text, but they remain unreliable for biology and biomedical research workflows that depend on precise interpretation of scientific papers, especially figures, charts, tables, captions, and the surrounding prose that explains them. Recent post-training methods (supervised fine-tuning and reinforcement learning) can inject domain capability into a base model, yet in practice the dominant bottleneck is still domain-specific training data: obtaining high-quality biomedical QA pairs typically requires costly expert annotation, and existing public datasets often fail to match the needs of targeted scientific tasks.

This paper addresses a concrete systems problem: **how to turn a corpus of raw biology/biomedical papers into a specialized VLM at low cost, without human annotation, while preserving strong evidence-grounded reasoning.** The central challenge is not simply to generate QA pairs, but to preserve the contextual evidence that makes scientific answers valid. In scientific papers, this evidence is often split across visual elements (e.g., axes, legends, and multi-row table headers) and the surrounding text that interprets them. When training examples drop these components, models tend to learn shallow cues or memorized patterns rather than evidence-grounded reasoning.

While LLM-driven data synthesis has shown promise in general domains (Wang et al., 2023; Liu et al., 2023), applying such methods naively to complex scientific PDFs risks producing hallucinated contexts. Popular general-purpose tools—such as Meta Synthetic Data Kit (Meta, 2025), Marker (Paruchuri and Datalab Team, 2025a), and DeepSeek OCR v2 (DeepSeek-AI, 2025)—often have poor data efficiency and struggle to produce sufficiently high-quality training data to improve models for biology and biomedical use cases. Further, prior biomedical-optimized VLMs such as LLaVA-Med (Li et al., 2023) and PMC-VQA (Zhang et al., 2023) fine-tune on figure-caption pairs from PubMed Central but discard the referring prose and cross-element context that ground scientific reasoning, resulting in poor performance in precise in-depth analysis.

To close the gap, we present Ryze, a system that automatically converts raw biomedical papers into an evidence-enriched training set and an enhanced specialized VLM. Ryze improves post-training quality and efficiency through two mechanisms: (1) it uses *evidence-enriched prompting* to synthesize QA examples where each question is paired with the comprehensive supporting evidence from the source PDF, namely, the visual

element, its caption, extracted structure, and the referring paragraphs. Also, it applies *chart/table-aware OCR and structure extraction*, together with *LLM-based cleansing*, to reduce layout- and OCR-induced errors before training. (2) it adopts a *progress-gated post-training strategy*: BioVLM automatically detects when additional synthesis and supervised fine-tuning stop improving benchmark performance, and then switches to reinforcement learning to strengthen evidence-grounded reasoning for complex biomedical question answering.

Ryze shifts biomedical-related domain adaptation from an "expert-annotation bottleneck" to a reproducible pipeline that a small lab or team can run. Applying Ryze to Qwen3-VL-8B produces BioVLM-8B, a strong and compact model that outperforms public, much larger and more expensive models such as GPT-5.2, while being post-trained end-to-end for under \$200. On challenging LAB-Bench (1,967 samples across eight biology categories), BioVLM-8B achieves **48.0%** weighted accuracy, outperforming the strong GPT-5.2 by **3.8%**, while remaining deployable on a single consumer GPU or Apple Silicon devices. Under an equal-token budget (a metric to reflect the post-training data efficiency), Ryze’s synthesized SFT data also surpasses human-curated PubMedQA and MedQA by **+17.1** and **+14.7 %**, respectively.

Motivated by these results, we released Ryze for early access to biomedical scientists and received consistently positive feedback on its practicality for day-to-day research tasks. Although this paper focuses on biomedicine, Ryze is designed to generalize to scientific papers more broadly. Our current system already supports climate change, geoscience, and civil engineering beyond biomedicine, and we will report the corresponding results in the open-source repository once approved.

2 Ryze Design and Implementation

2.1 System Overview

Ryze is released as a reproducible, low-cost workflow system tailored for biomedical scientists. Consider a biology researcher with a large collection of open-access papers on gene regulation. The researcher feeds these PDFs into Ryze, together with a base VLM (Qwen3-VL-8B) and an evaluation benchmark (LAB-Bench). Without any manual annotation, Ryze automatically:

- 1) Extracts structured text, figures, and tables using chart-aware OCR;

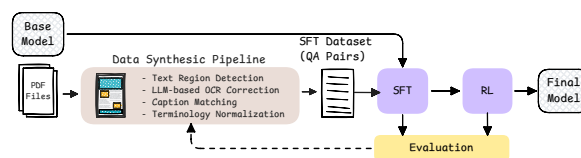


Figure 1: The Ryze pipeline overview. PDF files and a base model enter the Data Synthesis Pipeline, which trains the model via SFT then RL to produce the final domain-specialized model.

- 2) Synthesizes millions of tokens of domain QA data via agentic retrieval-augmented generation;
- 3) Post-trains the model with SFT followed by GRPO; and
- 4) Evaluates and iteratively improves via weakness-driven data augmentation.

The result is an enhanced model that can often surpass GPT-5.2 on biology benchmarks, while remaining lightweight enough to run locally on a single consumer GPU or an Apple device in the lab. Local deployment is also essential when some papers or internal reports are privacy-sensitive and cannot be uploaded to public model services (e.g., GPT-5.2, Gemini, or Claude).

Figure 1 overviews the above pipeline. Ryze consists of four interconnected modules operating in a closed loop: document processing (§2.2), data synthesis (§2.3), progress-gated training (§2.4), and evaluation with feedback (§2.5).

2.2 Chart-Aware Extraction and Cleansing

A common limitation for an automated data-synthesis pipeline is *OCR hallucination*: if the extraction stage misreads a chart axis, drops a table row, or severs the link between a figure and its caption, every downstream QA pair inherits that error. Layout-aware models such as LayoutLM (Xu et al., 2021) and Nougat (Blecher et al., 2023) advance document extraction but target form-like documents or plain Markdown conversion without preserving cross-referential chains. General-purpose tools such as Marker (Paruchuri and Datalab Team, 2025a) and DeepSeek OCR (DeepSeek-AI, 2025) were designed for broad document conversion, not for the dense, multi-column layouts of scientific PDFs with embedded equations, multi-panel figures, and merged-cell tables. In our early experiments, these tools produced frequent misrecognitions of gene names, chemical formulae, and chart annotations, where errors are propagated into training data and degraded downstream model accuracy.

Ryze mitigates this with a five-stage extraction-

and-cleansing pipeline: (1) **Layout detection** (Surya; Paruchuri and Datalab Team, 2025b) segments each page into text blocks, figures, tables, and captions, preserving spatial relationships that single-pass OCR discards. (2) Text regions are converted to structured Markdown preserving section hierarchy and paragraph boundaries. (3) **Cross-reference repair and content binding**: general-purpose OCR tools typically discard inline references (e.g., “Table 1”, “Figure 3”) during Markdown conversion, severing the link between body text and the visual elements it discusses. Ryze recovers these cross-references, resolves figure/table labels, and binds each visual element with its caption *and* the surrounding prose that interprets it, ensuring downstream synthesis can access the full evidential context. (4) Figures and tables receive **chart/table-aware extraction** (currently GLM-OCR; ZAI-org, 2025): charts are parsed into structured numerical descriptions; tables are converted to HTML with correct handling of merged cells and multi-row headers. (5) A **LLM-based cleansing stage** (currently Qwen3; Qwen Team, 2025) performs three passes—hallucination detection, domain-terminology repair (e.g., correcting “BRAF V6OOE” → “BRAF V600E”), and cross-element consistency verification—before data enters the synthesis pipeline. Because the preceding stages supply sufficient evidence and relevant context alongside each extracted element, we find that even compact models can perform cleansing reliably without introducing hallucinations.

The ablation in §3.3 quantifies the payoff: replacing our chart/table-aware extraction with generic alternatives degrades ChartQA accuracy by up to −7.8pp, and the gap widens monotonically with chart density (Table 3), confirming that scientific-domain-optimized extraction is necessary for high-fidelity evidence preservation.

2.3 Evidence-Enriched QA Synthesis

The synthesis pipeline transforms context-rich structured documents into the **evidence-enriched dataset**—high-quality QA training pairs that inherit the full structural and semantic fidelity of their source material.

Task-Aware Question Generation. Ryze draws question seeds from two complementary sources. From *raw papers*, the pipeline generates general domain questions (e.g., “What is the function of protein X?”) and questions raised in different papers (e.g., “Will X work on Mice?”). From the *evalua-*

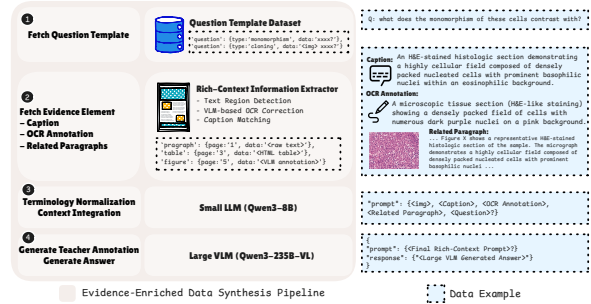


Figure 2: The evidence-enriched data synthesis pipeline. Source PDFs undergo text region detection, LLM-based OCR correction, caption matching, and terminology normalization. The resulting evidence-enriched dataset preserves the full structural and semantic fidelity of the source material for downstream QA generation.

tion benchmark, the system identifies the *reasoning skill categories* being tested (e.g., chart interpretation, protocol tracing, literature synthesis) rather than copying specific question templates or QA pairs from the dataset. These category-level cues are then extensively rephrased and diversified by a Qwen3-VL-235B model to produce new QA pairs that are not from any benchmark item, with all answers grounded entirely in the source papers. This process is analogous to curriculum-aware active learning: the benchmark informs *which reasoning skills to prioritize*, not how to answer specific questions. We emphasize that **no LAB-Bench questions, answers, rationales, outcomes, or scores enter generation, filtering, SFT, or RL**; all content derives from the source PDF corpus.

Evidence-Retrieval Answer Generation. Each question is accompanied by evidence retrieved from the structured document store, utilizing content binding from the document processing stage (§2.2) to enable precise extraction of passages, figures, and tables across document boundaries. All related elements in the paper are consolidated into the evidence set, including captions, OCR annotations (VLM annotations for charts and HTML format for tables), and associated paragraphs.

Active Data Augmentation. The development feedback loop (§2.5) identifies undercovered domain categories. For each such category, Ryze automatically searches for related open-access papers on PubMed, processes them through the document pipeline, and synthesizes additional QA pairs targeting the identified gaps. Evaluation data and scores are strictly excluded from all training; only the identified underperforming *categories* guide the search for new source papers.

2.4 Progress-Gated Training

Intuition. Data synthesis is the most expensive stage of the pipeline (§3.4), yet more data does not always help: beyond a saturation point, additional SFT examples yield diminishing returns while RL over the same budget can unlock substantial gains. Ryze therefore separates training into a *calibration* phase (SFT) that injects domain information and a *refinement* phase (GRPO) that strengthens evidence-grounded reasoning, with an automatic *progress gate* deciding when to switch.

Progress Gate. After each $\sim 1\text{M}$ -token increment of synthesized data, Ryze trains a checkpoint and evaluates on a held-out development split. When accuracy improvement stagnates over consecutive rounds, Ryze declares SFT saturation, freezes the dataset, and transitions to RL.

Calibration Phase. SFT calibrates the base model to domain-specific knowledge via LoRA (Zheng et al., 2024), alternating between text-only and vision QA batches so the model learns both domain terminology that appears in figures and textual descriptions that reference visual evidence.

Refinement Phase. Once the progress gate is triggered, Ryze converts the accumulated SFT data into an RL format and applies GRPO (Hu et al., 2024), which trains the model to generate coherent reasoning chains without relying on a separate reward model. GRPO has been widely adopted for mathematical reasoning (Shao et al., 2024), where comparable reasoning demands exist; similarly, complex biomedical experimental results require not only domain knowledge in biology but also rigorous figure and chart interpretation. Ryze integrates a progress-gated SFT-to-GRPO transition with evidence-enriched data for scientific document reasoning. During the SFT stage, it acquires common sense and fundamental biological concepts, leading to improvements in straightforward QA tasks such as SuppQA and DbQA, as shown in Table 1. During the RL stage, it achieves enhanced reasoning performance in domains that require more advanced analytical thinking. As shown in §3, SFT alone matches GPT-5.2 (43.7% vs. 44.2%), while GRPO provides the critical +4.3pp margin for overall superiority—confirming that calibration and refinement serve complementary roles.

Token Budget and Quality Control. Rather than pre-specifying a fixed data budget, Ryze incrementally grows the dataset in $\sim 1\text{M}$ -token in-

crements, training and evaluating after each round. When SFT performance plateaus, the system stops generation, converts the accumulated data into RL format, and switches to GRPO. Throughout, LLM-based quality filtering and deduplication ensure data integrity.

2.5 Evaluation and Feedback Loop

The evaluation module assesses model performance across benchmark categories and performs **limitation detection** by identifying domains where performance falls below a configurable threshold. Detected weaknesses update the QA template library with new question patterns targeting those domains, triggering a new round of data synthesis (§2.3). This closed-loop design supports iterative refinement until performance converges.

3 Evaluation

3.1 Evaluation Setup

We collect approximately 20,000 open-access biology papers from PubMed Central (PMC) as the source PDF corpus. We use Qwen3-VL-8B (Qwen Team, 2025) as the base model. We evaluate on LAB-Bench (Laurent et al., 2024), a biology research benchmark with 1,967 samples spanning eight diverse categories that test distinct reasoning skills—from literature comprehension (LitQA2) and figure/chart interpretation (FigQA, ChartQA) to sequence analysis (SeqQA), table reasoning (TableQA), and protocol understanding (ProtocolQA). Even state-of-the-art commercial models such as GPT-5.2 achieve only 44.2%, confirming the benchmark’s difficulty and discriminative power. Results are reported as sample-weighted averages.

We include **GPT-5.1 mini** and **GPT-5.2** as commercial reference points. These represent particularly strong baselines for biomedical tasks: OpenAI positions the GPT-5 family as a flagship for healthcare applications, reporting state-of-the-art results on HealthBench (OpenAI, 2025a)—a physician-curated benchmark with 5,000 clinical scenarios—and achieving 92.4% on GPQA Diamond graduate-level science questions including biology (OpenAI, 2025b). To validate that the evidence-enriched dataset is superior to existing alternatives, we train two baseline models using the same base model and training configuration: (1) **PubMedQA SFT** (Jin et al., 2019) and (2) **MedQA SFT** (Jin et al., 2021), both aligned to our token budget via trunca-

Table 1: LAB-Bench results (accuracy %). Qwen3-VL-8B (base) is the unmodified base model. BioVLM-8B models are trained from Qwen3-VL-8B. Best in **bold**.

Category	Qwen3 VL-8B	GPT 5.1m	GPT 5.2	PubMed SFT	MedQA SFT	BioVLM-8B (SFT only)	BioVLM-8B
Cloning	24.2	30.3	36.4	33.3	24.2	34.5	38.4
DbQA	31.2	40.6	41.7	14.0	20.0	44.7	48.9
FigQA	24.7	26.5	36.5	18.0	22.0	31.8	35.2
LitQA2	38.7	50.8	45.7	38.0	34.0	58.2	65.5
ProtocolQA	38.3	63.0	65.7	46.0	46.0	68.1	72.3
SeqQA	43.4	36.2	47.0	34.0	34.0	39.5	42.8
SuppQA	24.8	36.6	48.8	30.0	30.0	40.9	44.2
TableQA	34.0	34.0	36.9	22.0	30.0	40.3	45.6
W. Avg	35.4	39.1	44.2	26.6	29.0	43.7	48.0

tion/duplication.

All training configurations use identical token budgets: 8,051,591 tokens for SFT and 1,584,412 tokens for GRPO, measured by the Qwen3-8B tokenizer. All experiments are conducted on a server with an AMD EPYC 7313P CPU and 4× NVIDIA RTX A6000 (48 GB) GPUs.

3.2 Main Results

BioVLM-8B surpasses GPT-5.2 by +3.8% overall, winning 5 of 8 categories, with the largest margins in LitQA2 (+19.8), TableQA (+8.7), and DbQA (+7.2)—categories that demand deep literature comprehension and structured data reasoning. Compared to the unmodified Qwen3-VL-8B base model (35.4%), BioVLM-8B achieves an absolute gain of +12.6%. BioVLM-8B (SFT only) alone matches GPT-5.2 (43.7% vs. 44.2%), while the GRPO stage provides the critical margin for overall superiority. Equal-token human-curated datasets perform drastically worse (PubMedQA: 26.6%, MedQA: 29.0%), trailing BioVLM-8B (SFT only) by +17.1 and +14.7%, confirming that the *evidence-enriched dataset* is fundamentally superior to human-curated alternatives. GPT-5.2 retains advantages in FigQA (−1.3), SeqQA (−4.2), and SuppQA (−4.6), which involve heavier visual understanding and sequence analysis, which require larger parameter size or better model architecture.

Cross-Model Generalization. To verify that the evidence-enriched dataset generalizes beyond one architecture, we apply the same SFT data to three additional base models (Table 2). All models show consistent improvement, with the largest gain on Qwen3-VL-8B (+12.6%) which benefits from both text and vision components.

3.3 Ablation and Analysis

OCR Method Comparison. To validate the necessity of chart-aware extraction, we compare four

Table 2: Cross-model generalization: BioVLM-8B’s evidence-enriched dataset improves all tested base models on LAB-Bench (weighted accuracy %).

Base Model	Before SFT	After SFT	Δ
Qwen2.5-7B	33.1	35.1	+2.0
LLaMA-3.2	31.3	34.4	+3.2
Gemma-2	31.8	33.5	+1.7
Qwen3-VL-8B	35.4	43.7	+8.3

Table 3: OCR pipeline comparison (accuracy %). Our full document processing pipeline achieves the best performance, with larger gains on chart-intensive tasks.

Method	ArXivQA	SlideVQA	ChartQA
Without OCR	72.61	64.83	68.0
Marker	72.96	64.74	69.3
DeepSeek OCR	72.91	66.15	69.1
Ours	73.12	66.48	75.8

OCR configurations across three benchmarks of varying chart density (Table 3). Replacing GLM-OCR with generic alternatives degrades ChartQA by up to −7.8pp (75.8%→68.0%), with the gap widening as chart density increases: +7.8pp on ChartQA (pure charts), +1.65pp on SlideVQA (mixed), +0.51pp on ArXivQA (text-dominated). This gradient confirms that specialized chart/table extraction is *necessary* for high-quality evidence preservation.

Token Budget Ablation. We incrementally grow the SFT dataset in ∼0.8M-token steps and evaluate after each round (Figure 3). During the first 10 steps (0–7.2M tokens), performance plateaus at ∼33% with <2pp variation—the model acquires frequency-level familiarity with domain terminology but has not yet deeply internalized the knowledge. At 8M tokens, a sharp jump to 44.8% occurs, already matching GPT-5.2. Further tokens (8.8–10.4M) remain in the 42–44% range, confirming that knowledge consolidation requires sufficient data volume, and additional data beyond the sufficiency point yields diminishing returns.

Data Source Comparison. Under the same 8.05M-token SFT budget, PubMedQA achieves 26.6% and MedQA reaches 29.0%, while BioVLM-8B (SFT only) attains 43.7%—a margin of +17.1 and +14.7pp respectively (Figure 4). With GRPO, BioVLM-8B further climbs to 48.0%, surpassing GPT-5.2 by +3.8pp. This controlled comparison supports the evidence claim: preserving visual elements, captions, structures, and referring paragraphs is more effective than repurposing existing biomedical QA datasets.




Figure 3: Incremental SFT token budget vs. LAB-Bench accuracy. Dashed lines indicate GPT-5.1 mini and GPT-5.2 reference levels.




Figure 4: Data source comparison and per-category analysis on LAB-Bench. **Left:** Weighted accuracy by data source under an equal 8.05M-token SFT budget; dashed lines indicate GPT-5.1 mini and GPT-5.2 reference levels. **Right:** Per-category radar profile of GPT-5.1 mini, GPT-5.2, and BioVLM-8B.

Progress-Gated Training Analysis. Table 1 shows the progressive enhancement from commercial baselines through BioVLM-8B’s two training stages across all eight LAB-Bench categories. GRPO gains are largest on reasoning-intensive categories: LitQA2 (+7.3pp over SFT, +19.8pp over GPT-5.2), ProtocolQA (+4.2pp), and DbQA (+4.2pp). Categories where GPT-5.2 retains an edge—FigQA, SeqQA, SuppQA—involve heavier visual understanding and sequence analysis, suggesting future improvements in multimodal RL training. The radar chart in Figure 4 further highlights the per-category profile. This pattern confirms the “know the facts, then know the reasons” paradigm: SFT injects domain information that GRPO consolidates into knowledge through reasoning.

3.4 Cost Analysis

Table 4 breaks down the estimated cost of reproducing the full pipeline on cloud GPUs. The total GPU cost, including Qwen3-VL-235B QA synthesis and Qwen3 cleansing, is **under \$200** on RunPod rates,

Table 4: Estimated cost on cloud GPUs (RunPod rates: A100 80 GB \$1.19/hr; A6000 48 GB \$0.76/hr).

Stage	GPU Config	Hours	Cost
OCR + Cleansing	1×A6000	~24h	~\$18
QA Synthesis	8×A100	~15h	~\$143
SFT Training	4×A6000	~8h	~\$24
GRPO Training	4×A6000	~4h	~\$12
Total			< \$200

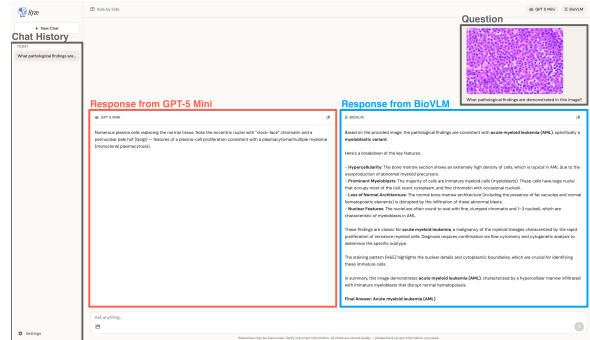


Figure 5: The Ryze demo interface, supporting side-by-side comparison between BioVLM-8B and commercial models such as GPT-5 Mini.

comparable to a single month of an API subscription and at least an order of magnitude cheaper than expert human annotation of similar scale.

We release the complete pipeline code (Apache 2.0), evidence-enriched dataset, and BioVLM-8B model weights at <https://github.com/Chivier/Ryze>. Figure 5 shows the Ryze web demo, where users can query BioVLM-8B and compare its responses with commercial models.

4 Conclusion

We present Ryze, a fully automated system that converts raw biomedical papers into an evidence-enriched dataset and a domain-specialized VLM. The core insight is that preserving the full contextual structure of scientific documents—figures, charts, tables, captions, and the prose that ties them together—produces training data that is fundamentally superior to both human-curated alternatives and context-agnostic automated approaches. On LAB-Bench, BioVLM-8B achieves 48.0% weighted accuracy, surpassing GPT-5.2 by +3.8% and outperforming equal-token human-curated datasets by over 17%, at an end-to-end cost under \$200. We release the complete pipeline, dataset, and model to support domain-specialized training across scientific fields.

Limitations

Our evaluation is limited to the biology domain; generalization to other scientific fields remains to be validated. GPT-5.2 retains advantages in categories involving visual understanding (FigQA) and sequence analysis (SeqQA, SuppQA), indicating room for improvement in multimodal reasoning. The scaling behavior of our progress-gated paradigm with larger models warrants further investigation.

We plan to address these limitations before presenting the demo at ACL, if accepted.

References

- Lukas Blecher, Guillem Cucurull, Thomas Scialom, and Robert Stojnic. 2023. Nougat: Neural optical understanding for academic documents. *arXiv preprint arXiv:2308.13418*.
- DeepSeek-AI. 2025. DeepSeek-OCR: Contexts optical compression. <https://github.com/deepseek-ai/DeepSeek-OCR>. GitHub repository.
- Jian Hu, Xibin Tao, Weixun Gao, and 1 others. 2024. OpenRLHF: An easy-to-use, scalable and high-performance RLHF framework. *arXiv preprint arXiv:2405.11143*.
- Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. 2021. What disease does this patient have? A large-scale open domain question answering dataset from medical exams. *Applied Sciences*, 11(14):6421.
- Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. 2019. PubMedQA: A dataset for biomedical research question answering. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*, pages 2567–2577.
- Jon M. Laurent, Joseph D. Burnett, Samson Rodman, and 1 others. 2024. LAB-Bench: Measuring capabilities of language models for biology research. *arXiv preprint arXiv:2407.10362*.
- Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao. 2023. LLaVA-Med: Training a large language-and-vision assistant for biomedicine in one day. *Advances in Neural Information Processing Systems*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. *Advances in Neural Information Processing Systems*.
- Meta. 2025. Synthetic data kit. <https://github.com/meta-llama/synthetic-data-kit>. Accessed: 2025.
- OpenAI. 2025a. HealthBench: Evaluating large language models towards improved human health. <https://openai.com/index/healthbench/>. Accessed: 2025.
- OpenAI. 2025b. Introducing GPT-5.2. <https://openai.com/index/introducing-gpt-5-2/>. Accessed: 2025.
- Vikas Paruchuri and Datalab Team. 2025a. Marker: A fast and accurate PDF to markdown converter. <https://github.com/VikParuchuri/marker>. GitHub repository. Built on Surya OCR models.
- Vikas Paruchuri and Datalab Team. 2025b. Surya: A lightweight document OCR and analysis toolkit. <https://github.com/VikParuchuri/surya>. GitHub repository.
- Qwen Team. 2025. Qwen3 technical report. <https://qwenlm.github.io/blog/qwen3/>. Accessed: 2025.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, Y.K. Li, Y. Wu, and Daya Guo. 2024. DeepSeekMath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. Self-instruct: Aligning language models with self-generated instructions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*.
- Yang Xu, Yiheng Xu, Tengchao Lv, Lei Cui, Furu Wei, Guoxin Wang, Yijuan Lu, Dinei Florencio, Cha Zhang, Wanxiang Che, and 1 others. 2021. LayoutLMv2: Multi-modal pre-training for visually-rich document understanding. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics*.
- ZAI-org. 2025. GLM-OCR: Accurate, fast, and comprehensive OCR. <https://github.com/zai-org/GLM-OCR>. GitHub repository.
- Xiaoman Zhang, Chaoyi Wu, Ziheng Zhao, Weixiong Lin, Ya Zhang, Yanfeng Wang, and Weidi Xie. 2023. PMC-VQA: Visual instruction tuning for medical visual question answering. *arXiv preprint arXiv:2305.10415*.
- Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyang Luo, Zhangchi Ma, and Yongqiang Ma. 2024. LlamaFactory: Unified efficient fine-tuning of 100+ language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pages 400–410.