Decoupled Physical Modeling and Execution for Physics Reasoning

Anonymous ACL submission

Abstract

Physics reasoning requires constructing a consistent model of the underlying physical system rather than relying solely on symbolic or formula-based manipulation. Although large language models have shown strong ability in solving math and coding problems, they still struggle with physics problems, as these problems entangle the physical modeling process with mathematical calculations. Humans approach physics by first building a representation of the system before performing calculations. Inspired by this, we introduce a unified framework that distills intermediate representations that explicitly encode the physical modeling process and adopt a two-stage post-training strategy, where supervised fine-tuning establishes structured modeling, and reinforcement learning with rubric-based feedback improves the quality of the modeling process. Experiments on multiple multimodal physics benchmarks show that our approach leads to consistent improvements in reasoning performance across different models and datasets. On PhysReason, PhyX and SeePhys benchmarks, physical modeling output performs GRPO by an average ~3%, showing that explicit physical modeling is an efficient strategy of improving physics reasoning for small LLMs.

1 Introduction

Solving physics problems requires more than selecting and applying equations. Human experts usually organize the problem into a structured physical representation that captures the relevant principles and relationships (Chi et al., 1981; Kohl and Finkelstein, 2008; Maloney, 1993). In this view, physics reasoning involves constructing and operating on an internal model of the system, rather than performing symbolic manipulation alone.

Multimodal Large Language Models (MLLMs) have demonstrated strong performance in a wide range of reasoning tasks, particularly in mathematical (Zhang et al., 2024a; Lu et al., 2024; Zhang et al., 2024b; Shi et al., 2024; Chen et al., 2025) and logical reasoning (Xu et al., 2024a,b; Hendrycks et al., 2021; Xu et al., 2026). However, in physics reasoning, MLLMs tend to implicitly conflate physical modeling and mathematical execution into a single reasoning process, leading to conceptual errors in physics problem-solving. In Figure 1, direct reasoning fails to identify the force arm by using the cylinder diameter radius relation incorrectly, leading to an overestimated angular acceleration and final cable speed.

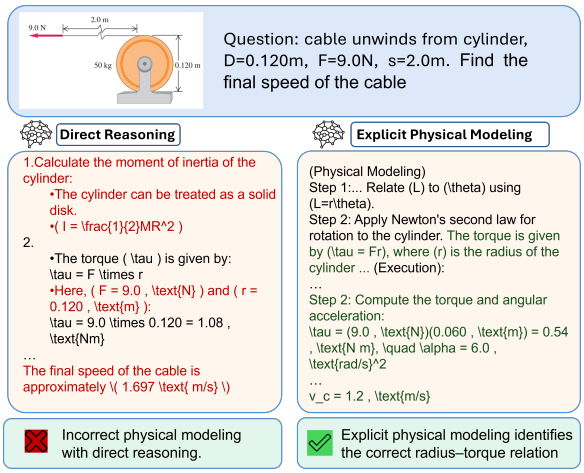

Figure 1: Physics reasoning errors often arise from incorrect physical modeling. Explicit modeling separates system representation from mathematical execution.

In this paper, we explicitly model physics reasoning as a two-stage process consisting of **physical modeling** and **mathematical execution**. We introduce structured physical modeling traces as an intermediate representation and train models through a two-stage post-training framework. First, supervised fine-tuning (SFT) teaches the model to generate explicit physical modeling before execution. Then, Group Relative Policy Optimization (GRPO) (Shao et al., 2024) with rubric-guided rewards further improves the quality and physical grounding of the generated reasoning process. Ex-

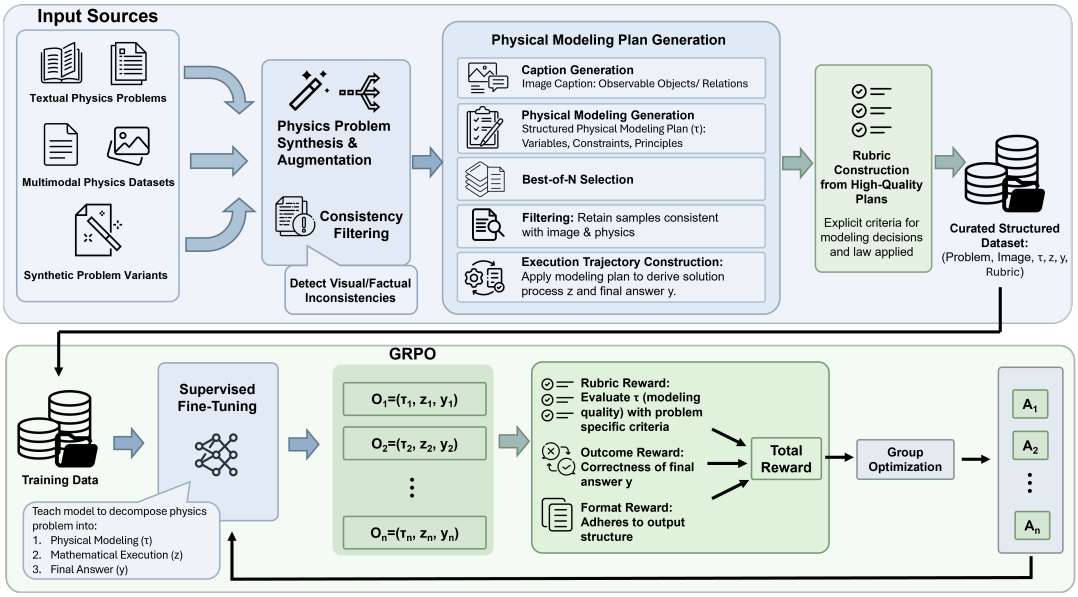

Figure 2: Overview of our framework. We construct structured physical modeling plans and physics-related rubrics, filter them into training data, and then train the model with supervised fine-tuning followed by GRPO using rubric-based rewards.

perimental results show that explicitly modeling the physical system consistently improves performance on open-ended multimodal physics reasoning benchmarks.

2 Structured Physical Modeling for Physics Reasoning

2.1 Overview and Task Formulation

We decompose the physics problem-solving into two stages: **physical modeling** and **mathematical execution**. In the **physical modeling** stage, the large language model applies relevant concepts and constraints to understand the system, while in the **mathematical execution** stage, these abstract physics concepts and constraints are formulated into specific equations for execution. Given the input problem x, the model first generates τ to construct the physical modeling, and then produces (z, y) based on this representation:

$$\pi_\theta(\tau, z, y \mid x) = \pi_\theta(z, y \mid \tau, x)\, \pi_\theta(\tau \mid x). \quad (1)$$

To formalize structured reasoning, we conduct a two-stage post-training strategy. We first perform supervised fine-tuning (SFT) as a cold start. We further refine this capability using Group relative Policy Optimization(GRPO) (Shao et al., 2024) with rubric-guided rewards, which improves modeling quality and encourages physically grounded and informative representations (see Figure 2).

2.2 Physics-Modeling Data Construction

Structured Modeling Trajectory Generation. We construct the multimodal problem pool from PhysReason (Zhang et al., 2025), PhyX (Shen et al., 2025), and SeePhys (Xiang et al., 2025), and the text-only pool from SCP-116K (Lu et al., 2025) and MegaScience (Fan et al., 2025). We further synthesize additional multimodal problems that preserve the underlying physical principles and concepts by using Gemini-2.5-Flash (Comanici et al., 2025). For each problem x_i, we generate a structured physical modeling trajectory τ_i by rejection sampling using Qwen3.5-Flash (Qwen Team, 2026). To improve visual grounding, we use Qwen3-VL-32B-Instruct (Bai et al., 2025a) to extract image captions. We then apply a two-stage verification process to retain high-quality modeling trajectories. Finally, we obtain the execution trajectory z_i by following the verified modeling trajectory, thereby deriving the correct solution. Details of the generation pipeline and dataset statistics are provided in Appendices C and D, respectively.

Rubric Derivation from Modeling Traces. We derive the rubrics from verified structured physical modeling trajectories. For each verified physical modeling trace, we generate a rubric set $\mathcal{C} = \{c_k\}_{k=1}^{|\mathcal{C}|}$ with 8–12 criteria, covering the following aspects: (1) **Visual Criteria**; (2) **Physical Modeling Criteria**; (3) **Strategic Planning Criteria**; (4)

Method	PhysReason		PhyX		SeePhys		PhysUni$_{\text{mc}}$		MMK-12$_{\text{phys.}}$	
	P@1	P@5	P@1	P@5	P@1	P@5	P@1	P@5	P@1	P@5
Qwen2.5-VL-3B										
Baseline	13.3	33.6	10.9	23.5	9.1	24.7	17.1	50.9	35.0	73.4
GRPO	17.3	36.8	12.6	27.0	12.4	29.8	19.1	52.7	42.4	82.4
Physics-Modeling + GRPO	**22.9**	**47.1**	**13.6**	**34.7**	**16.5**	**33.9**	**21.2**	**53.7**	**46.2**	**84.0**
Qwen2.5-VL-7B										
Baseline	25.0	47.8	17.0	35.7	14.0	28.9	19.3	53.7	45.4	79.4
GRPO	28.7	52.5	19.2	39.3	18.2	33.1	20.9	55.4	54.2	87.4
Physics-Modeling + GRPO	**32.1**	**58.2**	**21.4**	**40.0**	**19.8**	**34.7**	**23.5**	**57.8**	**56.8**	**89.6**
Qwen3-VL-8B										
Baseline	51.8	70.7	43.2	61.4	27.3	40.5	32.4	60.2	67.6	84.8
GRPO	**54.0**	73.9	**46.1**	64.1	30.6	45.5	34.9	63.4	73.8	90.0
Physics-Modeling + GRPO	53.3	**74.7**	45.1	**64.8**	**32.2**	**47.9**	**36.6**	**65.5**	**74.4**	**91.2**

Table 1: Main results on both in-domain and out-of-domain benchmarks. We compare the baseline model and GRPO variants across five physics reasoning datasets. For each benchmark, we report Pass@1 and Pass@5. Bold indicates the best result within each model block.

Method	PhysReason		PhyX		SeePhys	
	P@1	P@5	P@1	P@5	P@1	P@5
Qwen2.5-VL-7B						
Baseline	25.0	47.8	17.0	35.7	14.0	28.9
SFT Cold Start	27.2	54.4	18.2	36.0	15.7	28.9
Qwen3-VL-8B						
Baseline	51.8	70.7	43.2	61.4	27.3	40.5
SFT Cold Start	50.3	70.4	41.5	61.8	29.8	42.1

Table 2: Effect of Physical Modeling SFT Cold Start

Execution Criteria. Each criterion c_k consists of a textual specification d_k and an adaptive weight w_k, i.e., $c_k = (d_k, w_k)$. The weight reflects the criterion's relative importance for approaching the problem. Details are provided in Appendix C.3.

2.3 Two-Stage Post-Training

We adopt a two-stage post-training pipeline. We perform SFT on structured trajectories (τ, z, y), where τ is the physical modeling, z is the execution process, and y is the final answer:

$$\mathcal{L}_{\text{SFT}} = -\mathbb{E}_{(x,\tau,z,y)\sim\mathcal{D}} \left[\log \pi_\theta(\tau, z, y \mid x)\right]. \quad (2)$$

We use a text-to-multimodal curriculum: text-only data first establishes the modeling format, and multimodal data then teaches visual grounding.

We optimize the model with GRPO using a reward that combines final-answer correctness R_{outcome}, format validity R_{format}, and rubric-based evaluation R_{rubric} of the intermediate physical modeling:

$$R_{\text{total}} = R_{\text{outcome}} + R_{\text{format}} + R_{\text{rubric}}. \quad (3)$$

Details are provided in the Appendix E.

3 Results

Performance of Physical Modeling. Table 1 shows that the physical-modeling GRPO outperforms the base model and the standard GRPO method across different model scales and evaluation benchmarks. For relatively weak base models, such as Qwen2.5-VL-3B and Qwen2.5-VL-7B (Bai et al., 2025b), improvements in physics reasoning are significant. The gains are especially clear on PhysReason, where P@1/P@5 improves from 13.3/33.6 to 22.9/47.1 for the 3B model, and from 25.0/47.8 to 32.1/58.2 for the 7B model. For Qwen3-VL-8B (Bai et al., 2025a), the improvement is relatively marginal, as the base model already achieves strong performance across benchmarks (e.g., 51.8/70.7 on PhysReason and 43.2/61.4 on PhyX). Nevertheless, Physics-Modeling + GRPO still consistently improves the performance to 53.3/74.7 on PhysReason and 45.1/64.8 on PhyX. One possible reason is that the base model already tends to generate reflective reasoning traces, in which it explores multiple physics concepts before converging to an answer, resulting in an extremely long output response that ultimately leads to the correct answer.

Effect of SFT Cold Start. Table 2 shows the effect of SFT cold start across different models. For Qwen2.5-VL-7B, the proposed SFT cold start consistently improves performance on PhysReason, PhyX, and SeePhys. In contrast, the improvement on Qwen3-VL-8B is relatively inconsistent. To better understand this behavior, Figure 3 analyzes Qwen3-VL-8B under different token budgets. The

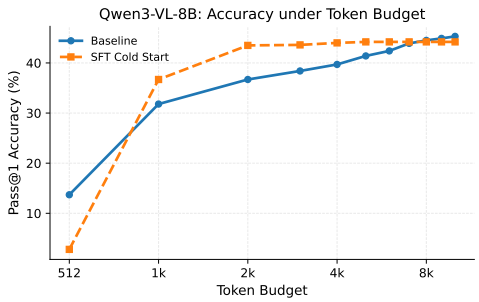

Figure 3: Pass@1 accuracy of Qwen3-VL-8B under different token budgets.

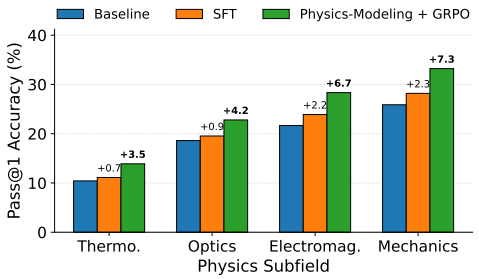

Figure 4: Performance across physics subfields for Qwen2.5-VL-7B.

baseline continues to improve as the token budget increases, suggesting that it often relies on longer reflective reasoning traces to reach the correct answer. By contrast, SFT cold start achieves stronger performance under smaller token budgets, indicating that physical modeling can make the reasoning process more concise. Further analysis is provided in Appendix H.

Ablation Studies. Table 3 presents the ablation analysis of our two main components. Removing either SFT cold start or the rubric reward consistently weakens performance, suggesting that the two components provide complementary benefits. In particular, SFT cold start provides a useful initialization via structured physical modeling, while rubric-based reinforcement learning further encourages the model to utilize these structures more effectively during reasoning. The full method achieves the best or tied-best performance across all benchmarks.

Subfield-wise Analysis. We further analyze performance across mechanics, electrodynamics, optics, and thermodynamics. As shown in Figure 4, Physics-Modeling + GRPO consistently improves performance across most sub-fields, with the largest gains observed in mechanics and electrodynamics. These sub-fields share a similar reasoning pattern, often requiring explicit reasoning about forces and interactions, which may benefit more from our proposed structured physical modeling. In contrast, the improvements in optics and thermodynamics are relatively smaller, which often involve different modeling patterns. Optics relies more on geometric ray configurations, while thermodynamics requires reasoning over changes between thermodynamic states and energy transfer across thermodynamic paths.

Method	PR	PhyX	SP	PUB	MMK
w/o SFT	28.1	18.2	15.7	20.5	52.8
w/o Rubric	30.4	20.9	19.8	21.9	54.0
Full	**32.1**	**21.4**	**19.8**	**23.5**	**56.8**

Table 3: Ablation study on Qwen2.5-VL-7B. PR, SP, PUB, and MMK denote PhysReason, SeePhys, PhysUniBench, and MMK-12.

Qualitative Analysis. The Carnot engine-refrigerator example (see Appendix I) illustrates how structured modeling supports reasoning in coupled physical systems. The model first analyzes the heat engine to determine the available work output, and then transfers this quantity to the refrigerator stage, enabling a coherent multi-step solution. This example highlights that the benefit of structured physical modeling lies in explicitly organizing key variables, constraints, and governing relations before numerical execution. Such intermediate representations provide a clear pathway that guides subsequent reasoning steps. We observe similar patterns across other subfields.

4 Conclusion

In this work, we study physics reasoning from the perspective of structured physical modeling. We formulate physics reasoning as a two-stage process that separates physical modeling from mathematical execution, and introduce an intermediate representation to explicitly capture the modeling stage. To support this formulation, we develop a data construction pipeline that provides structured supervision, supervised fine-tuning, and rubric-based reinforcement learning. The rubric signals are derived from high-quality modeling traces and provide fine-grained feedback on whether the constructed physical model is consistent and complete. Experiments across multiple multimodal physics benchmarks show that explicitly improving physical modeling leads to more reliable reasoning performance.

Limitations

In this study, we find that explicit physical modeling in post-training improves LLMs physics reasoning capability. However, our proposed method relies on the automatic pipeline to generate physical modeling traces and physics-specific criteria. The quality of the data depends on the underlying LLMs' reasoning ability. Although we apply multi-stage verification, the generated data still contains noise and flawed physical modeling. In addition, our formulation is designed for physics reasoning. The generalization to other domains remains unexplored. Finally, the proposed framework introduces additional compute cost due to the rubric-based RL requiring LLM-as-a-judge during training.

References

Rahul K. Arora, Jason Wei, Rebecca Soskin Hicks, Preston Bowman, Joaquin Quiñonero-Candela, Foivos Tsimpourlas, Michael Sharman, Meghan Shah, Andrea Vallone, Alex Beutel, Johannes Heidecke, and Karan Singhal. 2025. Healthbench: Evaluating large language models towards improved human health. *Preprint*, arXiv:2505.08775.

Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhifang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Mingsheng Li, and 45 others. 2025a. Qwen3-vl technical report. *Preprint*, arXiv:2511.21631.

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, and 8 others. 2025b. Qwen2.5-vl technical report. *Preprint*, arXiv:2502.13923.

Xinyan Chen, Renrui Zhang, Dongzhi Jiang, Aojun Zhou, Shilin Yan, Weifeng Lin, and Hongsheng Li. 2025. Mint-cot: Enabling interleaved visual tokens in mathematical chain-of-thought reasoning. *arXiv preprint arXiv:2506.05331*.

Michelene T. H. Chi, Paul J. Feltovich, and Robert Glaser. 1981. Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 5(2):121–152.

Gheorghe Comanici and 1 others. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *Preprint*, arXiv:2507.06261.

Ganqu Cui, Lifan Yuan, Zefan Wang, Hanbin Wang, Yuchen Zhang, Jiacheng Chen, Wendi Li, Bingxiang He, Yuchen Fan, Tianyu Yu, Qixin Xu, Weize Chen, Jiarui Yuan, Huayu Chen, Kaiyan Zhang, Xingtai Lv, Shuo Wang, Yuan Yao, Xu Han, and 6 others. 2025. Process reinforcement through implicit rewards. *Preprint*, arXiv:2502.01456.

Song Dai, Yibo Yan, Jiamin Su, Dongfang Zihao, Yubo Gao, Yonghua Hei, Jungang Li, Junyan Zhang, Sicheng Tao, Zhuoran Gao, and Xuming Hu. 2025. Physicsarena: The first multimodal physics reasoning benchmark exploring variable, process, and solution dimensions. *Preprint*, arXiv:2505.15472.

Nifu Dan, Yujun Cai, and Yiwei Wang. 2025. Symbolic or numerical? understanding physics problem solving in reasoning llms. *Preprint*, arXiv:2507.01334.

Run-Ze Fan, Zengzhi Wang, and Pengfei Liu. 2025. Megascience: Pushing the frontiers of post-training datasets for science reasoning. *arXiv preprint arXiv:2507.16812*.

Kaiyue Feng, Yilun Zhao, Yixin Liu, Tianyu Yang, Chen Zhao, John Sous, and Arman Cohan. 2025. Physics: Benchmarking foundation models on university-level physics problem solving. *Preprint*, arXiv:2503.21821.

Diana Galvan-Sosa, Gabrielle Gaudeau, Pride Kavumba, Yunmeng Li, Hongyi gu, Zheng Yuan, Keisuke Sakaguchi, and Paula Buttery. 2025. Rubrik's cube: Testing a new rubric for evaluating explanations on the cube dataset. *Preprint*, arXiv:2503.23899.

Zhiqi Gao, Tianyi Li, Yurii Kvasiuk, Sai Chaitanya Tadepalli, Maja Rudolph, Daniel J. H. Chung, Frederic Sala, and Moritz Münchmeyer. 2025. Test-time scaling techniques in theoretical physics – a comparison of methods on the tpbench dataset. *Preprint*, arXiv:2506.20729.

Anisha Gunjal, Anthony Wang, Elaine Lau, Vaskar Nath, Yunzhong He, Bing Liu, and Sean Hendryx. 2025. Rubrics as rewards: Reinforcement learning beyond verifiable domains. *Preprint*, arXiv:2507.17746.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. *Preprint*, arXiv:2009.03300.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *Preprint*, arXiv:2106.09685.

Raj Jaiswal, Dhruv Jain, Harsh Parimal Popat, Avinash Anand, Abhishek Dharmadhikari, Atharva Marathe, and Rajiv Ratn Shah. 2024. Improving physics reasoning in large language models using mixture of refinement agents. *Preprint*, arXiv:2412.00821.

Dulhan Jayalath, Shashwat Goel, Thomas Foster, Parag Jain, Suchin Gururangan, Cheng Zhang, Anirudh Goyal, and Alan Schelten. 2026. Compute as teacher:

Turning inference compute into reference-free supervision. *Preprint*, arXiv:2509.14234.

Fangkai Jiao, Chengwei Qin, Zhengyuan Liu, Nancy F. Chen, and Shafiq Joty. 2024. Learning planning-based reasoning by trajectories collection and process reward synthesizing. *Preprint*, arXiv:2402.00658.

Patrick B. Kohl and Noah D. Finkelstein. 2008. Patterns of multiple representation use by experts and novices during physics problem solving. *Phys. Rev. ST Phys. Educ. Res.*, 4:010111.

Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V. Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, Yuling Gu, Saumya Malik, Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Tafjord, Chris Wilhelm, Luca Soldaini, and 4 others. 2025. Tulu 3: Pushing frontiers in open language model post-training. *Preprint*, arXiv:2411.15124.

Dakuan Lu, Xiaoyu Tan, Rui Xu, Tianchu Yao, Chao Qu, Wei Chu, Yinghui Xu, and Yuan Qi. 2025. Scp-116k: A high-quality problem-solution dataset and a generalized pipeline for automated extraction in the higher education science domain. *Preprint*, arXiv:2501.15587.

Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. 2024. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. *Preprint*, arXiv:2310.02255.

Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Preprint*, arXiv:2209.09513.

David P. Maloney. 1993. Research on problem solving: Physics. In D. Gabel, editor, *Handbook of research in science teaching and learning*, pages 327–356. Macmillan.

Xinyu Pang, Ruixin Hong, Zhanke Zhou, Fangrui Lv, Xinwei Yang, Zhilong Liang, Bo Han, and Changshui Zhang. 2024. Physics reasoner: Knowledge-augmented reasoning for solving physics problems with large language models. *Preprint*, arXiv:2412.13791.

Mihir Parmar, Palash Goyal, Xin Liu, Yiwen Song, Mingyang Ling, Chitta Baral, Hamid Palangi, and Tomas Pfister. 2025. Plan-tuning: Post-training language models to learn step-by-step planning for complex problem solving. *Preprint*, arXiv:2507.07495.

Qwen Team. 2026. Qwen3.5: Towards native multimodal agents.

David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. 2023. Gpqa: A graduate-level google-proof q&a benchmark. *Preprint*, arXiv:2311.12022.

Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *Preprint*, arXiv:2402.03300.

Hui Shen, Taiqiang Wu, Qi Han, Yunta Hsieh, Jizhou Wang, Yuyue Zhang, Yuxin Cheng, Zijian Hao, Yuansheng Ni, Xin Wang, Zhongwei Wan, Kai Zhang, Wendong Xu, Jing Xiong, Ping Luo, Wenhu Chen, Chaofan Tao, Zhuoqing Mao, and Ngai Wong. 2025. Phyx: Does your model have the "wits" for physical reasoning? *Preprint*, arXiv:2505.15929.

Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. 2025. Hybridflow: A flexible and efficient rlhf framework. In *Proceedings of the Twentieth European Conference on Computer Systems*, EuroSys '25, page 1279–1297. ACM.

Wenhao Shi, Zhiqiang Hu, Yi Bin, Junhua Liu, Yang Yang, See-Kiong Ng, Lidong Bing, and Roy Ka-Wei Lee. 2024. Math-llava: Bootstrapping mathematical reasoning for multimodal large language models. *Preprint*, arXiv:2406.17294.

Gemini Team and 1 others. 2025. Gemini: A family of highly capable multimodal models. *Preprint*, arXiv:2312.11805.

Vijay Viswanathan, Yanchao Sun, Shuang Ma, Xiang Kong, Meng Cao, Graham Neubig, and Tongshuang Wu. 2025. Checklists are better than reward models for aligning language models. *Preprint*, arXiv:2507.18624.

Lintao Wang, Encheng Su, Jiaqi Liu, Pengze Li, Jiabei Xiao, Wenlong Zhang, Xinnan Dai, Xi Chen, Yuan Meng, Lei Bai, Wanli Ouyang, Shixiang Tang, Aoran Wang, and Xinzhu Ma. 2026. Physunibench: A multi-modal physics reasoning benchmark at undergraduate level. *Preprint*, arXiv:2506.17667.

Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyan Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhu Chen. 2024. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. *Preprint*, arXiv:2406.01574.

Kun Xiang, Heng Li, Terry Jingchen Zhang, Yinya Huang, Zirong Liu, Peixin Qu, Jixi He, Jiaqi Chen, Yu-Jie Yuan, Jianhua Han, and 1 others. 2025. Seephys: Does seeing help thinking?–benchmarking vision-based physics reasoning. *arXiv preprint arXiv:2505.19099*.

Fangzhi Xu, Qika Lin, Jiawei Han, Tianzhe Zhao, Jun Liu, and Erik Cambria. 2024a. Are large language models really good logical reasoners? a comprehensive evaluation and beyond. *Preprint*, arXiv:2306.09841.

Jundong Xu, Hao Fei, Liangming Pan, Qian Liu, Mong-Li Lee, and Wynne Hsu. 2024b. Faithful logical reasoning via symbolic chain-of-thought. *Preprint*, arXiv:2405.18357.

Jundong Xu, Hao Fei, Yuhui Zhang, Liangming Pan, Qijun Huang, Qian Liu, Preslav Nakov, Min-Yen Kan, William Yang Wang, Mong-Li Lee, and Wynne Hsu. 2026. Muslr: Multimodal symbolic logical reasoning. *Preprint*, arXiv:2509.25851.

Renrui Zhang, Dongzhi Jiang, Yichi Zhang, Haokun Lin, Ziyu Guo, Pengshuo Qiu, Aojun Zhou, Pan Lu, Kai-Wei Chang, Peng Gao, and 1 others. 2024a. Mathverse: Does your multi-modal llm truly see the diagrams in visual math problems? *arXiv preprint arXiv:2403.14624*.

Renrui Zhang, Xinyu Wei, Dongzhi Jiang, Ziyu Guo, Shicheng Li, Yichi Zhang, Chengzhuo Tong, Jiaming Liu, Aojun Zhou, Bin Wei, Shanghang Zhang, Peng Gao, Chunyuan Li, and Hongsheng Li. 2024b. Mavis: Mathematical visual instruction tuning with an automatic data engine. *Preprint*, arXiv:2407.08739.

Xinyu Zhang, Yuxuan Dong, Yanrui Wu, Jiaxing Huang, Chengyou Jia, Basura Fernando, Mike Zheng Shou, Lingling Zhang, and Jun Liu. 2025. Physreason: A comprehensive benchmark towards physics-based reasoning. *arXiv preprint arXiv:2502.12054*.

Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyan Luo, Zhangchi Feng, and Yongqiang Ma. 2024. Llamafactory: Unified efficient fine-tuning of 100+ language models. *Preprint*, arXiv:2403.13372.

Yang Zhou, Sunzhu Li, Shunyu Liu, Wenkai Fang, Kongcheng Zhang, Jiale Zhao, Jingwen Yang, Yihe Zhou, Jianwei Lv, Tongya Zheng, and 1 others. 2025. Breaking the exploration bottleneck: Rubric-scaffolded reinforcement learning for general llm reasoning. *arXiv preprint arXiv:2508.16949*.

A Related Work

LLMs for Physics Reasoning. Recent work has introduced a variety of benchmarks for evaluating physics reasoning across both textual and multimodal settings, covering a broad range of difficulty levels from high-school to advanced problems (Wang et al., 2026; Zhang et al., 2025; Feng et al., 2025; Xiang et al., 2025). Early works treat physics reasoning as part of general scientific benchmarks (Lu et al., 2022; Wang et al., 2024; Rein et al., 2023), while more recent work focuses on dedicated physics benchmarks that emphasize more challenging and diverse reasoning scenarios. In particular, multimodal physics reasoning requires models to jointly reason over textual and visual information to construct and solve physically grounded representations of the underlying system (Dai et al., 2025; Shen et al., 2025; Zhang et al., 2025; Xiang et al., 2025). Across these benchmarks, physics reasoning is typically characterized by a modeling-centric process, where the large language model must construct a structured representation of the physical system before performing symbolic or numerical computation (Dai et al., 2025). Prior works have explored approaches such as knowledge-augmented reasoning (Pang et al., 2024) or iterative refinement methods (Jaiswal et al., 2024) that aim to enhance intermediate reasoning quality. A recent study further suggests that scaling inference-time computation alone is often insufficient and that effective physics reasoning is typically based on more structured and consistent reasoning patterns (Gao et al., 2025; Dan et al., 2025). These observations highlight the importance of explicitly modeling the underlying physical system, rather than relying on implicit generation without structured modeling supervision.

Post-Training Methods for Reasoning. Reinforcement learning with verifiable rewards (RLVR) has shown strong performance in domains such as mathematics and coding by leveraging programmatic or externally verifiable rewards (Shao et al., 2024). Beyond outcome-level rewards, prior work explores richer feedback over reasoning trajectories (Cui et al., 2025; Lambert et al., 2025). Process supervision and planning-style training provides denser guidance over intermediate reasoning steps by encouraging structured decomposition of the solution process (Jiao et al., 2024; Parmar et al., 2025). More recently, rubric- or checklist-based evaluation has emerged as a way to assess open-ended reasoning tasks using multiple explicit criteria rather than relying solely on final-answer accuracy or holistic judgments (Arora et al., 2025; Galvan-Sosa et al., 2025). Beyond evaluation, such structured criteria are used as rubric-based rewards to evaluate the reasoning trajectory during policy optimization in reinforcement learning (Gunjal et al., 2025; Viswanathan et al., 2025; Zhou et al., 2025; Jayalath et al., 2026). Despite providing richer supervision than final-answer rewards alone, these approaches largely rely on generic criteria for assessing reasoning quality. In physics, however, effective problem-solving depends on constructing

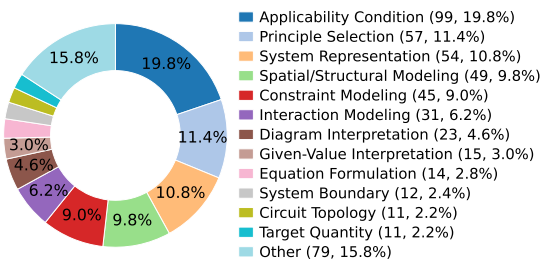

Figure 5: Distribution of physics reasoning error types from a preliminary analysis using a Qwen2.5-VL-7B.

an explicit physical representation of the system, thereby motivating supervision signals grounded in the underlying physical structure.

B Preliminary Error Analysis

To better understand the challenges of physics reasoning, we conduct a preliminary error analysis using Qwen2.5-VL-7B (Bai et al., 2025b) on a subset of 500 physics problems. We categorize the incorrect responses into different error types using Qwen3.5-Flash (Qwen Team, 2026). As shown in Figure 5, a large fraction of failures arises from incorrect or incomplete physical modeling. This observation motivates our focus on explicitly separating physical modeling from downstream computation.

C Data Generation and Curation Details

C.1 Data Source

PhysReason (Zhang et al., 2025) is a multimodal physics reasoning benchmark constructed from diverse sources including college entrance exams and international physics competitions. It consists of 1,200 carefully curated problems with diagram-grounded contexts, structured solutions, and step-level annotations, emphasizing long-horizon, multi-step reasoning. Notably, each problem often contains multiple sub-questions, which we decompose into independent samples to construct finer-grained training instances.

PhyX (Shen et al., 2025) focuses on visual physics reasoning across diverse scenarios. It contains 3,000 curated questions spanning six reasoning types, 25 sub-domains, and six core physics domains, including mechanics, electromagnetism, thermodynamics, optics, modern physics, and wave-related topics. Unlike purely symbolic benchmarks, PhyX emphasizes integrating visual understanding with physical principles.

SeePhys (Xiang et al., 2025) emphasizes vision-dependent physics reasoning grounded in diagrams. It contains 2,000 curated questions paired with 2,245 images, covering 7 physics domains and 21 heterogeneous diagram categories across knowledge levels from middle school to PhD qualifying exams. A substantial portion of the benchmark is vision-essential, requiring models to extract indispensable information directly from diagrams rather than relying solely on textual cues.

SCP-116K (Lu et al., 2025) is a large-scale dataset of high-quality problem–solution pairs designed for scientific reasoning across STEM domains. It contains 116,756 rigorously curated examples automatically extracted from heterogeneous educational resources, including textbooks and academic materials. The dataset is constructed via a generalized multi-stage pipeline with stringent filtering to ensure scientific rigor and appropriate educational levels. Unlike manually curated benchmarks, SCP-116K emphasizes scalable data construction, providing both the dataset and the extraction pipeline to facilitate future extensions and domain transfer.

MegaScience (Fan et al., 2025) is a large-scale scientific reasoning dataset consisting of 1.25 million instances of diverse STEM domains. It is designed to support complex multi-step reasoning and scientific problem solving.

We construct our training data from both textual and multimodal physics sources. For textual data, we sample 12,000 physics-related problem–solution pairs from SCP-116K and MegaScience. For multimodal data, we randomly sample instances from PhysReason, PhyX, and SeePhys. We first filter out non-English problems to ensure consistency. To establish a clean evaluation protocol, we reserve a held-out set of 1,000 problems, including 467 from PhysReason, 412 from PhyX, and 121 from SeePhys. These samples are excluded from all stages of data generation and training. The remaining data is used as input to our data generation pipeline for constructing structured physical modeling traces and rubrics .

C.2 Data Generation Pipeline

Problem Augmentation. To expand the multimodal physics problem pool, we use Gemini-2.5-Flash (Comanici et al., 2025) to synthesize additional problem instances from existing multimodal physics data. Specifically, we apply the following transformations:

- **Objective variation:** modifying the target quantity or question objective while keeping the physical setup unchanged;

- **Numerical perturbation:** altering numerical values (e.g., masses, distances, or constants) within physically valid ranges;

- **Linguistic rephrasing:** rewriting problem descriptions to introduce surface-level diversity without changing semantics.

We further filter out samples with visual conflicts or underlying physics inconsistencies using Gemini-2.5-Flash (Comanici et al., 2025).

Visual Captioning. To improve visual grounding, we first extract a caption from the image using Qwen3-VL-32B-Instruct (Bai et al., 2025a). The caption is constrained to describe only observable and physically relevant elements in the scene, including object identities, spatial relationships, geometric configurations, and annotated quantities (e.g., angles, lengths, or directions), without introducing additional assumptions or inferred conclusions.

Structured Physical Modeling. We construct the model input by combining the image, question, generated caption, and reference solution. The reference solution is used only to guide the modeling structure and does not directly determine the final answer during plan generation. We adopt a best-of-N strategy ($N = 5$) to generate multiple candidate physical modeling trajectories. Each trajectory is required to explicitly represent the physical system in a structured form, rather than a high-level summary. Specifically, each trajectory includes:

- **System Understanding:** the target physical quantity or system property to be solved;

- **Constraints:** conditions implied by the problem setup, diagram, or geometry, such as boundary conditions, conservation constraints, or kinematic relations;

- **Governing principles:** the physical laws or principles relevant to the problem, along with how they apply to the current system;

- **Equations and physical relationships:** symbolic formulations that instantiate the governing principles and explicitly connect variables through the underlying physical relationships.

Two-stage Quality Control. To ensure that the generated modeling plans provide reliable supervision for physical modeling, we apply a two-stage filtering procedure at both the modeling process and outcome levels.

Stage 1: Modeling-level filtering. We first evaluate each candidate modeling plan using Gemini-2.5-Flash, focusing on three aspects: (1) *internal coherence*, which checks whether the reasoning steps are logically consistent and free of contradictions; (2) *multimodal consistency*, which verifies that the modeling process is grounded in the visual scene and does not conflict with the image or problem description; and (3) *modeling specificity*, which requires the plan to explicitly represent key elements of the physical system, including variables, constraints, governing principles, and their relationships. Plans that are inconsistent, underspecified, or remain at a generic high-level description are discarded.

Stage 2: Execution-based validation. Among the physical modeling traces that pass the first stage, we further validate whether the modeling plan can support correct downstream reasoning. Specifically, we provide the question, image, and candidate modeling plan to Gemini-2.0-Flash(Team et al., 2025) and require the model to strictly follow the plan to derive the final answer. Only plans that lead to correct solutions are retained.

This two-stage procedure ensures that the retained plans are both physically grounded at the modeling stage and executable toward correct final answers.

C.3 Rubric Construction

Rubric Derivation from Physical Modeling Traces. For each problem, we derive a set of fine-grained evaluation criteria from the filtered modeling plan. Each criterion corresponds to a specific requirement of the modeling process, ensuring that the rubric is grounded in the underlying physical representation rather than generic reasoning patterns.

Design Principles. To ensure that the rubrics provide meaningful and reliable supervision, we enforce that each criterion is *atomic*, capturing a single aspect of the modeling process; *verifiable*, corresponding to a concrete and independently checkable requirement; and *physically grounded*, reflecting explicit physical concepts such as variables, constraints, governing principles, and their rela-

tionships. In addition, we explicitly avoid vague or generic criteria (e.g., "correct reasoning" or "proper simplification") to ensure that all criteria are non-trivial and informative.

Rubric Structure. Each rubric consists of 8-12 criteria that collectively capture the construction of a physical system representation. Specifically, the criteria cover key aspects of the modeling process, including:

- **Visual Criteria:** Verify the correctness of key information from physics diagrams (e.g., numerical values and spatial relationships);
- **Physical Modeling Criteria:** Evaluate the construction of the underlying physical modeling (e.g., concepts, constraints, and physics process);
- **Strategic Planning Criteria:** Evaluate the decomposition of a complex physics problem into appropriate subproblems;
- **Execution Criteria:** Evaluate the translation from abstract concepts to equations and symbolic representations, and intermediate answers.

This process converts modeling traces into structured, physically grounded rubrics that provide explicit supervision over the modeling process.

D Data Statistics

In Table 4, we provide detailed statistics of the constructed multimodal physical modeling trajectory dataset. The final dataset contains 8,880 physical trajectories, including 4,852 original samples and 4,028 synthesized samples, with relatively balanced coverage across major physics categories.

Data	Cat.	Orig.	Synth.	Total
PhysReason	Electro.	962	523	1485
	Mech.	567	476	1043
	Optics	353	207	560
	Thermo.	239	89	328
	Other	19	24	43
PhyX	Electro.	418	454	872
	Mech.	644	695	1339
	Optics	537	469	1006
	Thermo.	367	348	715
	Other	133	110	243
SeePhys	Electro.	227	249	476
	Mech.	212	254	466
	Optics	112	102	214
	Thermo.	43	25	68
	Other	19	3	22
Total	–	**4852**	**4028**	**8880**

Table 4: Category-level statistics of the constructed physical modeling trajectory dataset.

E Two-stage Post-training for Physical Modeling

E.1 SFT Cold Start for Structured Physical Modeling

We first fine-tune the model to generate structured reasoning trajectories (τ, z, y), where τ denotes physical modeling, z denotes execution, and y denotes the final answer:

$$\mathcal{L}_{\text{SFT}} = -\mathbb{E}_{(x,\tau,z,y)\sim\mathcal{D}} \left[\log \pi_\theta(\tau, z, y \mid x)\right]. \quad (4)$$

We use a curriculum from text-only to multi-modal data. The former establishes the structural pattern of modeling, while the latter teaches the models to ground physical modeling in visual information.

E.2 Physical Modeling-Guided Policy Optimization

We further optimize the model with rubric-based GRPO. Unlike standard GRPO, which primarily rewards final-answer correctness, our approach additionally rewards intermediate physical modeling and mathematical execution. Specifically, given a group of sampled responses $\{y_i\}_{i=1}^G$ for each problem x, we optimize

$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}_{x,\{y_i\}} \left[\frac{1}{G}\sum_{i=1}^{G} \hat{A}_i \log \pi_\theta(y_i \mid x)\right],$$

where $\hat{A}_i$ is computed from the normalized reward within the sampled group. The reward combines final-answer correctness with rubric-based scores for physical modeling and execution.

Rubric Reward. Given a sampled reasoning trajectory $(\tau, z, y)_i$ and a set of fine-grained rubrics $\mathcal{C} = \{(d_k, w_k)\}_{k=1}^{|\mathcal{C}|}$, we evaluate each criterion independently with an LLM-as-a-judge. The judge evaluates whether $(\tau, z, y)_i$ satisfies the criterion description d_k under the original question q and returns a binary score:

$$s_k(q, (\tau, z, y)_i) = \begin{cases} 1, & \text{criterion satisfied,} \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

The rubric reward is computed as the weighted average of criterion-level scores:

$$R_{\text{rubric}}(q, (\tau, z, y)_i) = \frac{\sum_{k=1}^{|\mathcal{C}|} w_k s_k(q, (\tau, z, y)_i)}{\sum_{k=1}^{|\mathcal{C}|} w_k}. \quad (6)$$

Outcome and Format Rewards. We use an outcome reward to evaluate final-answer correctness:

$$R_{\text{outcome}} = \begin{cases} 1, & \hat{y} = y, \\ 0, & \text{otherwise.} \end{cases} \tag{7}$$

We also use a format reward R_{format}, set to 0.5 if the output follows the required `<plan>`, `<solution>`, and `<answer>` structure, and 0 otherwise:

$$R_{\text{format}} = \begin{cases} 0.5, & \text{valid format,} \\ 0, & \text{otherwise.} \end{cases} \tag{8}$$

Total Reward. The total reward combines rubric, outcome, and format rewards:

$$R_{\text{total}} = R_{\text{rubric}} + R_{\text{outcome}} + R_{\text{format}}. \tag{9}$$

F Experimental Setup

Datasets. For text-only SFT, we construct a 12K-problem training set based on SCP-116K (Lu et al., 2025) and MegaScience (Fan et al., 2025). For multimodal SFT, we use PhysReason (Zhang et al., 2025), PhyX (Shen et al., 2025), and SeePhys (Xiang et al., 2025) as the problem pool, holding out 1,000 problems for evaluation. We synthesize new problems from the remaining pool using Gemini-2.5-Flash (Comanici et al., 2025), resulting in 8,880 multimodal training instances after synthesis and filtering. For GRPO, we further derive fine-grained rubrics from the verified structured modeling trajectories in the multimodal training set, enabling process-level reward of physical modeling.

Base Models. To evaluate our method, we fine-tune three open-source vision-language models of different scales and series: **Qwen2.5-VL-3B**, **Qwen2.5-VL-7B** (Bai et al., 2025b), and **Qwen3-VL-8B** (Bai et al., 2025a).

Implementation Details. For SFT, we adopt a curriculum training pipeline. We first perform text-only supervised fine-tuning for 2 epochs, followed by multimodal fine-tuning for 2 epochs on the constructed physics training set. We use LoRA (Hu et al., 2021) with rank 64 and $\alpha = 128$, and set the learning rate to 1×10^{-4}. SFT is implemented with `LLaMA-Factory` (Zheng et al., 2024). For GRPO, we continue training from the multimodal SFT checkpoint using only the multimodal physics training set. We use `verl` (Sheng et al., 2025) with LoRA rank 64, $\alpha = 128$, and a learning rate of 5×10^{-5} for 1 epoch. We evaluate open-ended physics reasoning using a hybrid evaluation framework, with details in Appendix G.

G Evaluation Details

Evaluating open-ended physics reasoning poses unique challenges due to the diversity of valid answer forms, including numerical values with units, symbolic expressions, and free-form textual responses. To ensure reliable and consistent evaluation, we adopt a hybrid evaluation framework that combines rule-based matching with LLM-assisted verification.

Answer Types. We categorize answers into four types: *Numerical*, *Symbolic*, *Textual*, and *Multiple-choice*. Each type is evaluated using a specialized strategy tailored to its structure.

Numerical Evaluation. For numerical answers, we first normalize both predicted and ground-truth responses into a canonical value–unit representation. This normalization handles diverse formats, including scientific notation, LaTeX expressions (e.g., fractions and square roots), and implicit unit expressions. We then evaluate correctness based on (i) numerical equivalence under a relative tolerance, and (ii) unit consistency. When units differ but are dimensionally equivalent, we convert values to a shared canonical form before comparison (e.g., $\text{eV} \leftrightarrow \text{J}$, $\text{deg} \leftrightarrow \text{rad}$). The evaluation also supports multiple-value answers by matching predicted and ground-truth value–unit pairs.

Symbolic Evaluation. For symbolic answers, we normalize LaTeX expressions to remove formatting variations and parse them into symbolic representations. We then evaluate equivalence using the `math-verify` toolkit. To further improve robustness, we apply additional normalization steps such as removing presentation-level commands and, when applicable, comparing only the right-hand side of equations.

Multiple-Choice Evaluation. For multiple-choice questions, we extract the predicted option (e.g., A–E) using pattern matching and directly compare it with the ground-truth label.

Fallback LLM-as-Judge. When rule-based evaluation fails (e.g., due to parsing errors, unsupported formats, or ambiguous answers), we employ an LLM-based judge as a fallback. The model is prompted to determine whether the predicted answer is equivalent to the ground truth and outputs a binary correctness decision. This mechanism ensures coverage for edge cases while preserving consistency with the primary evaluation criteria.

Method	Samples w/ Reflection	Ratio
Baseline	309	0.617
Modeling-SFT	5	0.009

Table 5: Frequency of reflective reasoning patterns in 500 sampled Qwen3-VL-8B generations.

This hybrid design balances precision and robustness: rule-based evaluation provides strict and interpretable correctness signals, while the LLM-based fallback handles the inherent variability of open-ended physics answers.

H Analysis of Reflective Reasoning Behavior

Reflection Keyword Detection. To quantify reflective reasoning behavior, we use a lightweight keyword-based diagnostic. A generated response is counted as containing reflection behavior if it includes at least one reflection-related expression, such as "wait", "hold on", "let me check", "rethink", or "verify". These keywords are intended to capture explicit self-checking or reconsideration patterns in the model's reasoning trace. We apply this diagnostic to 500 sampled Qwen3-VL-8B generations and report the proportion of responses containing such reflection keywords. As shown in Table 5, the baseline exhibits substantially more reflection-keyword patterns than Modeling-SFT, suggesting that the base model more frequently relies on explicit self-revision during reasoning.

Qualitative Analysis. Table 6 shows a problem asks for the final speed of a two-bucket Atwood system released from rest, where the $12.0\,\mathrm{kg}$ bucket starts $2.00\,\mathrm{m}$ above the floor and is connected to a $4.0\,\mathrm{kg}$ bucket by a lightweight rope over a frictionless pulley. Both models solve the problem correctly and obtain $\boxed{4.43\,\mathrm{m/s}}$. However, the baseline first produces an incorrect intermediate sign convention and then recovers through an explicit reflective correction. In contrast, SFT cold start begins by constructing the physical system. This suggests that structured modeling can reduce the need for explicit self-revision by making the relevant physical constraints available before mathematical execution.

Model	Reasoning Excerpt
Baseline	"By conservation of energy: $$(m_1 - m_2)gh + \frac{1}{2}(m_1 + m_2)v^2 = 0$$ Solve for v^2: $$v^2 = \frac{2(m_2 - m_1)gh}{m_1 + m_2}$$ Wait — this gives a negative value for v^2, which is impossible. That means we have the sign wrong. The net potential energy loss is $m_1gh - m_2gh$, which is positive, and this energy is converted to kinetic energy."
SFT cold start	"The system consists of two buckets connected by a lightweight, inextensible rope passing over a frictionless, massless pulley. Since the rope is inextensible, both buckets move with the same magnitude of acceleration a and reach the same final speed v. Determine Acceleration: Apply Newton's Second Law to the system. The net force is the difference between the weights of the two buckets. Relate this net force to the total mass to solve for acceleration a. $$a = \frac{(m_1 - m_2)g}{m_1 + m_2}$$ "

Table 6: Reasoning excerpts from an Atwood-machine example.

I Qualitative Analysis

Question

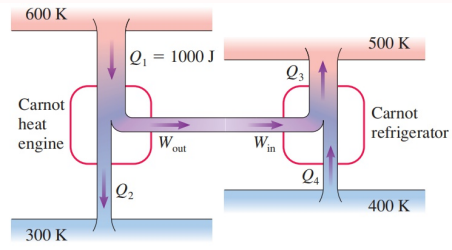

Figure shows a Carnot heat engine driving a Carnot refrigerator.
Determine Q_3.

Explicit Modeling+Execution

Objective:
The problem involves a Carnot heat engine driving a Carnot refrigerator. The first law of thermodynamics and the efficiency of Carnot engines and refrigerators are the primary theoretical frameworks. The diagram shows the heat transfer processes between the heat engine and the refrigerator.
Goal: Determine the amount of heat Q_3 extracted from the cold reservoir by the refrigerator.

Variables:

- $T_H = 600\,\text{K}$ | Hot reservoir temperature of the heat engine
- $T_C = 300\,\text{K}$ | Cold reservoir temperature of the heat engine
- $W_{out} = 1000\,\text{J}$ | Work output of the heat engine
- $Q_1 = 1000\,\text{J}$ | Heat absorbed from the hot reservoir by the heat engine
- $T_{\text{ref,hot}} = 500\,\text{K}$ | Hot reservoir temperature of the refrigerator
- $T_{\text{ref,cold}} = 400\,\text{K}$ | Cold reservoir temperature of the refrigerator
- W_{in} | Work input to the refrigerator
- Q_3 | Heat extracted from the cold reservoir by the refrigerator
- Q_4 | Heat rejected to the hot reservoir by the refrigerator

Constraints:

1. The heat engine and refrigerator operate as ideal Carnot machines.
2. The processes are reversible.

Steps:

1. Calculate the work output (W_{out}) of the Carnot heat engine using the formula

$$W_{out} = Q_1 - Q_2,$$

where Q_1 is the heat absorbed from the hot reservoir and Q_2 is the heat rejected to the cold reservoir. This is based on the first law of thermodynamics for a cyclic process.

2. Calculate the efficiency (η_{heat}) of the Carnot heat engine using the formula

$$\eta_{\text{heat}} = 1 - \frac{Q_2}{Q_1}.$$

Also, use the relationship between the efficiencies of the heat engine and the refrigerator, which are reciprocal:

$$\eta_{\text{ref}} = 1 - \frac{T_{\text{ref,cold}}}{T_{\text{ref,hot}}}.$$

3. Determine the work input (W_{in}) to the Carnot refrigerator using the formula

$$W_{in} = Q_3 - Q_4,$$

where Q_3 is the heat extracted from the cold reservoir and Q_4 is the heat rejected to the hot reservoir.

4. Use the efficiency of the Carnot refrigerator,

$$\eta_{\text{ref}} = \frac{W_{in}}{Q_3},$$

to find Q_3. Rearrange the formula to solve for Q_3:

$$Q_3 = \frac{W_{in}}{\eta_{\text{ref}}}.$$

5. Substitute the known values of W_{in} and η_{ref} into the equation from Step 4 to calculate the value of Q_3.

Execution:
Step 1: Calculate work output.
The work output of the Carnot heat engine is the difference between the heat absorbed from the hot reservoir and the heat rejected to the cold reservoir.

$$W_{out} = Q_1 - Q_2$$

From the diagram, $Q_1 = 1000\,\text{J}$. We need to find Q_2.

$$\eta_{\text{heat}} = 1 - \frac{Q_2}{Q_1}$$

For a Carnot engine,

$$\eta_{\text{heat}} = 1 - \frac{T_C}{T_H} = 1 - \frac{300\,\text{K}}{600\,\text{K}} = 1 - 0.5 = 0.5$$

So,

$$0.5 = 1 - \frac{Q_2}{1000\,\text{J}}$$

$$\frac{Q_2}{1000\,\text{J}} = 1 - 0.5 = 0.5$$

$$Q_2 = 0.5 \times 1000\,\text{J} = 500\,\text{J}$$

Therefore,

$$W_{out} = 1000\,\text{J} - 500\,\text{J} = 500\,\text{J}$$

956

Step 2: Calculate engine and refrigerator efficiencies.
The efficiency of the Carnot refrigerator is the reciprocal of the efficiency of the Carnot heat engine.

$$\eta_{\text{ref}} = 1 - \frac{T_{\text{ref,cold}}}{T_{\text{ref,hot}}} = 1 - \frac{400\,\text{K}}{500\,\text{K}} = 1 - 0.8 = 0.2$$

Step 3: Determine work input.
The work input to the Carnot refrigerator is the difference between the heat extracted from the cold reservoir and the heat rejected to the hot reservoir.

$$W_{in} = Q_3 - Q_4$$

Step 4: Solve for Q_3.
The efficiency of the Carnot refrigerator is also given by

$$\eta_{\text{ref}} = \frac{W_{in}}{Q_3}$$

Thus,

$$Q_3 = \frac{W_{in}}{\eta_{\text{ref}}}$$

Step 5: Calculate Q_3.
We know that

$$W_{in} = W_{out} = 500\,\text{J}$$

and

$$\eta_{\text{ref}} = 0.2$$

Therefore,

$$Q_3 = \frac{500\,\text{J}}{0.2} = 2500\,\text{J}$$

Answer:

$$\boxed{2500\,\text{J}}$$

957

Question

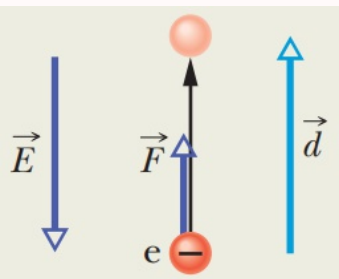

Electrons are continually being knocked out of air molecules in the atmosphere by cosmic-ray particles coming in from space. Once released, each electron experiences an electric force $\vec{F}$ due to the electric field $\vec{E}$ that is produced in the atmosphere by charged particles already on Earth. Near Earth's surface, the electric field has magnitude $E = 150\,\text{N/C}$ and is directed downward. What is the change ΔU in the electric potential energy of a released electron when the electric force causes it to move vertically upward through a distance $d = 520\,\text{m}$?

958

Explicit Modeling+Execution

Objective: Compute the change in electric potential energy of an electron moving upward in a uniform electric field. The solution relies on the relation between electric force, work done by the field, and potential energy change. **Variables:**

- $E = 150\,\mathrm{N/C}$: magnitude of the electric field
- $d = 520\,\mathrm{m}$: upward displacement
- $q_e = -1.602 \times 10^{-19}\,\mathrm{C}$: charge of electron
- ΔU: change in electric potential energy

Constraints:

- Electric field is uniform and directed downward
- Electron moves vertically upward

Steps:

1. Determine force direction using $\vec{F} = q_e\vec{E}$ (force is upward).
2. Compute work done by the electric field: $W = -q_eEd$.
3. Use $\Delta U = -W$ to obtain the potential energy change.

Execution
Step 1: Force on the electron:

$$F = q_eE = (-1.602 \times 10^{-19})(150) = -2.403 \times 10^{-17}\,\mathrm{N}$$

Step 2: Work done by the electric field:

$$W = -q_eEd = -(-1.602 \times 10^{-19})(150)(520) = 1.20 \times 10^{-14}\,\mathrm{J}$$

Step 3: Change in potential energy:

$$\Delta U = -W = -1.20 \times 10^{-14}\,\mathrm{J}$$

Answer:

$$\boxed{-1.20 \times 10^{-14}\,\mathrm{J}}$$

Question

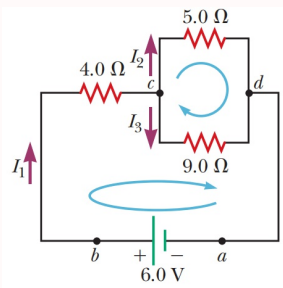

This example applies Kirchhoff's rules to analyze a multi-loop circuit with three currents and one battery. Find the currents I_1 in the circuit shown in Figure 18.14 by using Kirchhoff's rules. Find the current I_1 in the circuit shown in the figure by using Kirchhoff's rules.

Explicit Modeling+Execution

Objective:
The problem involves a DC circuit with multiple resistors and a voltage source. Kirchhoff's Voltage Law (KVL) and Kirchhoff's Current Law (KCL) will be used to analyze the circuit and determine the unknown currents.
Goal: Determine the values of the three unknown currents, I_1, I_2, and I_3, in the circuit.

Variables:

- $V = 6.0\,\text{V}$ | Voltage of the battery
- $R_1 = 4.0\,\Omega$ | Resistance of the first resistor
- $R_2 = 5.0\,\Omega$ | Resistance of the second resistor
- $R_3 = 9.0\,\Omega$ | Resistance of the third resistor
- I_1 | Current flowing out of the battery and through R_1
- I_2 | Current flowing through R_1 and R_2
- I_3 | Current flowing through R_2 and R_3

Constraints:

1. The circuit is a DC circuit, meaning the current and voltage are constant over time.
2. The wires connecting the components have negligible resistance.
3. The battery provides a constant voltage.

Steps:

1. Apply Kirchhoff's Current Law (KCL) at node 'c'. According to KCL, the sum of currents entering a node must equal the sum of currents leaving the node. In this case,

$$I_1 = I_2 + I_3.$$

   This equation relates the three unknown currents.

2. Apply Kirchhoff's Voltage Law (KVL) to the left loop (node 'a' to node 'b'). Starting from node 'a' and moving clockwise, the voltage drops across R_3 and then across R_1. Therefore,

$$V - I_3R_3 - I_1R_1 = 0.$$

   This equation relates the voltages and currents in terms of the battery voltage and the resistances.

3. Apply Kirchhoff's Voltage Law (KVL) to the right loop (node 'd' to node 'c'). Starting from node 'd' and moving counter-clockwise, the voltage drops across R_2 and then across R_1. Therefore,

$$I_2R_2 - I_1R_1 = 0.$$

   This equation relates the voltages and currents in terms of the resistances.

4. Solve the system of equations. We now have three equations with three unknowns (I_1, I_2, and I_3):

$$I_1 = I_2 + I_3,$$

$$V - I_3R_3 - I_1R_1 = 0,$$

$$I_2R_2 - I_1R_1 = 0.$$

Substitute the given values for V, R_1, R_2, and R_3, and solve the system of equations for I_1, I_2, and I_3.

Execution:
Step 1: Apply KCL at node 'c'.
Kirchhoff's Current Law (KCL) states that the total current entering a node is equal to the total current leaving the node. At node 'c', we have:

$$I_1 = I_2 + I_3$$

Step 2: Apply KVL to the left loop.
Kirchhoff's Voltage Law (KVL) states that the sum of the voltages around any closed loop is zero. Applying KVL to the left loop (a to b), we have:

$$V - I_3R_3 - I_1R_1 = 0$$

$$6.0 - I_3(9.0) - I_1(4.0) = 0$$

Step 3: Apply KVL to the right loop.
Applying KVL to the right loop (d to c), we have:

$$I_2R_2 - I_1R_1 = 0$$

$$I_2(5.0) - I_1(4.0) = 0$$

Step 4: Solve the system of equations.
We have the following equations:

1. $I_1 = I_2 + I_3$
2. $6.0 - 9.0I_3 - 4.0I_1 = 0$
3. $5.0I_2 - 4.0I_1 = 0$

From equation (3), we can express I_2 in terms of I_1:

$$5.0I_2 = 4.0I_1$$

$$I_2 = \frac{4.0}{5.0}I_1 = 0.8I_1$$

Substitute I_2 into equation (1):

$$I_1 = 0.8I_1 + I_3$$

$$I_3 = I_1 - 0.8I_1 = 0.2I_1$$

Substitute I_3 into equation (2):

$$6.0 - 9.0(0.2I_1) - 4.0I_1 = 0$$

$$6.0 - 1.8I_1 - 4.0I_1 = 0$$

963

$$6.0 - 5.8I_1 = 0$$

$$5.8I_1 = 6.0$$

$$I_1 = \frac{6.0}{5.8} = 1.0345\,\text{A}$$

Now, we can find I_2 and I_3:

$$I_2 = 0.8I_1 = 0.8(1.0345) = 0.8276\,\text{A}$$

$$I_3 = 0.2I_1 = 0.2(1.0345) = 0.2069\,\text{A}$$

Answer:

$$\boxed{I_1 = 1.03\,\text{A},\ I_2 = 0.83\,\text{A},\ I_3 = 0.21\,\text{A}}$$

964