


000 TRICKS OR TRAPS? 001 002 A DEEP DIVE INTO RL FOR LLM REASONING 003 004

005 **Anonymous authors**

006 Paper under double-blind review

009 ABSTRACT

012 Reinforcement learning (RL) for LLM reasoning has rapidly emerged as a prominent
013 research area, marked by a significant surge in related studies on both algorithmic
014 innovations and practical applications. Despite this progress, several critical
015 challenges remain, including the absence of standardized guidelines for applying
016 RL techniques and a fragmented understanding of their underlying mechanisms.
017 In addition, inconsistent experimental settings, variations in training data, and dif-
018 ferences in model initialization have led to conflicting conclusions, obscuring the
019 key characteristics of these techniques and creating confusion among practitioners
020 when selecting appropriate techniques. This paper systematically reviews widely
021 adopted RL techniques through rigorous reproductions and isolated evaluations
022 within a unified open-source framework. We analyze the internal mechanisms,
023 applicable scenarios, and core principles of each technique through fine-grained
024 experiments, including datasets of varying difficulty, model sizes, and architectures.
025 Based on these insights, we present clear guidelines for selecting RL techniques
026 tailored to specific setups and provide a reliable roadmap for practitioners navigat-
027 ing the RL for the LLM domain. Finally, we show that a minimalist combination
028 of two techniques can unlock the learning capability of critic-free policies with a
029 vanilla PPO loss. The results demonstrate that our simple combination consistently
030 improves performance, surpassing strategies such as GRPO and DAPO.



048 Figure 1: **Left:** The proliferation of RL optimization techniques, coupled with diverse initialized
049 models and data, has raised barriers to practical adoption. **Right:** We establish detailed application
050 guidelines via dissecting internal mechanisms of widely-used tricks, and introduce **Lite PPO**, a
051 minimalist two-technique combination that enhances learning capacity in critic-free policies with
052 vanilla PPO loss. The average accuracy is calculated across six mathematical benchmarks.

054
055
056
1 INTRODUCTION057
058
059
060
061
062
Recent breakthroughs in large language models (LLMs) such as OpenAI o1 (Wu et al., 2024) and
063
064
065
066
067
068
069
070
071
072
073
074
075
076
077
078
079
080
081
082
083
084
085
086
087
088
089
090
091
092
093
094
095
096
097
098
099
099
100
098
101
102
103
104
105
106
107
DeepSeek R1 (Shao et al., 2024) have positioned reinforcement learning (RL) as a key driver in
unlocking advanced reasoning capabilities in LLMs. This is particularly evident in challenging
reasoning tasks like mathematical problem solving (He et al., 2025b) and code generation (Zhuo
et al., 2025), where RL has demonstrated the potential to elevate LLM performance beyond what
pre-training alone can achieve. Such an emerging trend has sparked widespread interest within the
research community in the direction of "RL for LLM" (or RL4LLM).

However, this rapid progress is shadowed by a lack of usage guidelines (Huang et al., 2024b) and
in-depth understanding of underlying mechanisms. For instance, GRPO (Shao et al., 2024) advocates
group-level normalization for stability, while REINFORCE++ (Hu et al., 2025) favors batch-level.
Moreover, GRPO incorporates variance in normalization, yet Dr. GRPO (Liu et al., 2025b) explicitly
recommends removing variance normalization to prevent bias. Such contradictory and chaotic
phenomena underscore the fragmented understanding and inconsistent recommendations within the
RL4LLM community. A likely cause for the above phenomenon is that the experimental settings,
training data, and initialization of the existing work all have significant differences, which may also
cause deviations in the summary of the conclusions.

Apart from the confusion caused by the intrinsic differences of similar techniques, the numerous
and seemingly orthogonal techniques, including *Normalization*, *Clip*, and *Overlong Filtering*, also
increase the complexity of algorithm application in practice. Practitioners face non-trivial challenges
in identifying an effective combination from a wide range of techniques to unlock the learning
capacity of LLMs in specific scenarios. These ambiguities have naturally triggered a key requirement
of practitioners:

What scenarios are the existing techniques respectively suitable for? Is there a simple and generalizable
combination that can be used to enhance policy optimization?

Following established RL analysis practices (Andrychowicz et al., 2020; Engstrom et al., 2020;
Huang et al., 2024b), we systematically evaluate popular RL techniques through reproducible
experiments on a unified open-source framework. Our comprehensive setup spans datasets of varying
difficulty, different model sizes, and multiple model types, supported by **over 160 independent RL**
training experiments to ensure robust and statistically meaningful conclusions. We also delve into
the theoretical foundations, implementation, and recommended applications of each technique, as
summarized in Figure 1. Our findings show that most RL methods are highly sensitive to factors
like model type, data distribution, reward design, and hyperparameters. Notably, we demonstrate
that combining just two techniques—*advantage normalization* (*group mean*, *batch std*) and *token-level loss aggregation*—is sufficient to maximize the potential of critic-free policies with vanilla
PPO loss, outperforming mainstream RL4LLM approaches that rely on extra components. Our key
contributions are:

1. Removing the standard deviation when reward distributions are highly concentrated enhances
the stability and effectiveness of model training. (§4.1.2)
2. Group-level mean and batch-level standard deviation enable further robust normalization.
(§4.1.3)
3. Clip Higher promotes high-quality exploration for aligned models. (§4.2.1)
4. There appears to be a “scaling law” between the performance and the upper bound of the
clipping on the small-sized model. (§4.2.3)
5. Compared to sequence-level loss aggregation, token-level aggregation is effective on base
models but shows limited improvement on aligned models. (§4.3.1)
6. Overlong filtering enhances accuracy and clarity for short-to-medium reasoning tasks but
provides limited benefits for long-tail reasoning. (§4.4.1)
7. Two techniques may unlock learning capacity in critic-free policies based on vanilla PPO
loss. (§5)

108 **2 PRELIMINARIES**

109

110 A variety of practical techniques have been introduced to stabilize optimization, reduce variance, and
 111 accelerate the convergence of LLMs on reasoning tasks. We categorize commonly used techniques
 112 as follows: **(1) Baseline Design.** Baselines are crucial to reduce variance in the estimation of
 113 policy gradients, with recent advances including using the group-mean reward (Shao et al., 2024)
 114 or computing the baseline for each sample as the average gradient estimate from other samples in
 115 the group (Ahmadian et al., 2024; Kool et al., 2019). **(2) Clipping Strategies.** Clipping controls
 116 excessive policy updates by constraining rewards, advantages, or ratios, and the *Clip Ratio Higher*
 117 method enhances exploration by relaxing the upper bound in PPO’s ratio clipping (Yu et al., 2025). **(3)**
 118 **Normalization Strategies.** Normalization of rewards or advantages stabilizes gradient magnitudes,
 119 with representative methods including *Batch-level* (Hu et al., 2025), *Group-level* (Shao et al., 2024;
 120 Ahmadian et al., 2024), and *Reward Shift without Standard Deviation* (Liu et al., 2025b). **(4) Filtering**
 121 **Strategies.** Filtering excludes uninformative or undesirable samples before gradient computation,
 122 including *Overlong Filtering* for excessive lengths (Yu et al., 2025), *Error Max Clip Mask* and *Right*
 123 *Min Clip Mask* for extreme correctness or errors, and *Difficulty Mask* to exclude samples outside a
 124 target difficulty range (Yu et al., 2025; Zhang et al., 2025; Chu et al., 2025). **(5) Loss Aggregation**
 125 **Granularity.** The formulation of loss aggregation determines each token’s contribution to the overall
 126 objective, with common approaches including *Sequence-level Loss* and *Token-level Loss*, the latter
 127 computes per-token advantages to mitigate length bias. **(6) Additional Loss Functions.** Auxiliary
 128 losses can complement the primary objective and regularize training. *KL Loss* (Yu et al., 2025; Liu
 129 et al., 2025b) constrains divergence from a reference policy, while *SFT Loss* (Zhang & Zuo, 2025)
 130 incorporates supervised fine-tuning objectives to preserve alignment. **(7) Reward Design.** Shaping
 131 the reward function can guide desired output properties, with common strategies including *Length*
 132 *Penalty* to discourage overly long outputs, *Formatting Reward* to promote structured outputs, and
 133 *Length-Dependent Accuracy Reward* that combines correctness with output length.

134 **3 EXPERIMENTAL DESIGNS**

135


136 **3.1 EXPERIMENTAL SETUP**

137

138 **Training Algorithm:** We utilize the open-sourced ROLL framework¹ (Wang et al., 2025), an
 139 efficient and scalable platform specifically designed for reinforcement learning optimization in LLMs,
 140 to conduct all experiments. In addition, we adopt the PPO loss (Schulman et al., 2017), with advantage
 141 values computed using the REINFORCE algorithm (Sutton et al., 1999) as the unified RL baseline.
 142 To ensure consistency with prior research, we set the global batch size to 1024 by using a rollout
 143 batch size of 128 and sampling 8 responses per prompt, with a maximum response length of 8192
 144 tokens. The learning rate is set to $1e - 6$. For text generation, we use a top_p value of 0.99, a top_k
 145 value of 100, and a temperature of 0.99.

146 **Base Models:** To comprehensively evaluate reinforcement learning (RL) techniques across parame-
 147 ter scales, our experiments cover two model sizes: Qwen3-4B and Qwen3-8B. For each model size,
 148 we include both non-aligned pre-trained versions (Qwen3-4B-Base and Qwen3-8B-Base) and aligned
 149 versions, enabling assessment RL gains from different initialization conditions.

150 **Training Datasets:** To ensure reproducibility and
 151 fairness, we exclusively use open-source datasets for
 152 training, including *SimpleRL-Zoo-Data* (Zeng et al.,
 153 2025) and *Deepmath* (He et al., 2025b). To com-
 154 prehensively examine how problem difficulty (e.g., Easy,
 155 Medium, and Hard; detailed in Appendix B.3) affects
 156 the performance of the RL technique, we randomly
 157 sample 5,000 entries from the datasets. Figure 2
 158 visualizes the difficulty across the training dataset
 159 assessed by GPT-4o (Hurst et al., 2024).



151 Figure 2: Number of correct responses under
 152 8 rollout iterations across datasets.

153 ¹Open source RL framework: <https://github.com/alibaba/ROLL>




Figure 3: Test accuracy over training iterations is compared for different models trained on different datasets. The shaded regions represent the oscillation amplitude as $\text{mean} \pm (\text{std_multiplier} \times \text{std})$, with curves smoothed using an 11-step moving window and exponential smoothing ($\alpha = 0.8$).

Evaluation Benchmark: All the experiments are conducted on six math datasets: MATH-500 (Hendrycks et al., 2021), OlympiadBench (He et al., 2024), MinervaMath (Lewkowycz et al., 2022), and subsets of standardized examinations (AIME24-25, AMC23). These datasets span a broad complexity spectrum from basic arithmetic to competition-level mathematics, enabling a comprehensive evaluation of reasoning capabilities.

3.2 BASELINE RESULTS

Impact of Data Difficulty on Training Dynamics We investigate how data difficulty influences the training dynamics of Qwen3 models. Specifically, we analyze the training convergence patterns through accuracy trajectories and generalization gaps, across three tiers of complexity (*Easy*, *Medium*, *Hard*). The detailed learning curves in Figure 3 show that the model exhibits markedly different accuracy trajectories across training sets of different difficulty levels. When focusing on the differences in learning efficiency between the unaligned Base model and the aligned model under the same experimental setting, the aligned models exhibited substantially higher initial accuracy, but additional learning yielded only modest gains, with accuracy improving by roughly 2%. Additional details regarding the baseline’s training accuracy and response length can be found in Appendix C.

4 ANALYSIS

4.1 NORMALIZATION

Advantage normalization is a standard technique for stabilizing RL training in language models by reducing gradient variance (Zheng et al., 2023), yet implementations vary. GRPO (Shao et al., 2024) and RLOO (Ahmadian et al., 2024; Kool et al., 2019) use group-level normalization to promote intra-context competition, while REINFORCE++ (Hu et al., 2025) adopts batch-level normalization to mitigate overfitting and reward hacking in low-diversity settings. Formally, given a prompt x with K sampled responses and corresponding rewards $\{r_k\}_{k=1}^K$, and denoting r_i as the reward of the i -th response in a rollout batch of N prompts with K responses each, the group-level and batch-level normalized advantages are:

$$A_k^{\text{group}} = \frac{r_k - \text{mean}(\{r_j\}_{j=1}^K)}{\text{std}(\{r_j\}_{j=1}^K)} \quad A_i^{\text{batch}} = \frac{r_i - \text{mean}(\{r_j\}_{j=1}^{N \cdot K})}{\text{std}(\{r_j\}_{j=1}^{N \cdot K})} \quad (1)$$

4.1.1 ADVANTAGE NORMALIZATION IS SENSITIVE TO REWARD MECHANISMS

To systematically evaluate the impact of advantage normalization on PPO variants with a value function using the Monte Carlo return target, we conducted experiments under a unified training framework, exploring three settings: **no normalization**, **batch-level normalization**, and **group-level normalization**. When analyzing the performance in Figure 4, it can be concluded that both advantage normalization techniques can significantly influence the model’s convergence speed, performance stability, and final outcomes.




Figure 4: Accuracy comparison of Base models with different normalization techniques cross Easy and Hard datasets.


216 Specifically, on both model sizes, group-level normalization consistently achieves more stable training
 217 dynamics and higher final performance compared to both batch-level normalization and no normalization.
 218 Batch-level normalization exhibits high sensitivity to reward distribution skew, often leading
 219 to performance collapse under an imbalanced batch situation, where a few outlier samples dominate
 220 the advantage estimates.
 221

222 4.1.2 IMPACT OF THE STANDARD DEVIATION TERM IN ADVANTAGE NORMALIZATION

224 Takeaway 1

225 **Removing the standard deviation** when reward distributions are highly concentrated (e.g.,
 226 easy training dataset) enhances the stability and effectiveness of model training.
 227

228 We found when model responses within a prompt group yield highly similar rewards, e.g., when
 229 the responses are almost all correct or all incorrect, the resulting standard deviation becomes ex-
 230 tremely small. In such cases, dividing by this small standard deviation during normalization can
 231 excessively amplify gradient updates, causing the model to overemphasize tasks of extreme difficulty,
 232 a phenomenon similar to “difficulty bias” (Liu et al., 2025b).
 233



222 Figure 5: **Left:** Standard deviation variations during training on easy and hard datasets. **Middle and**
 223 **Right:** Test accuracy with and without batch-level standard deviation across easy and hard dataset.
 224

225 To test whether the standard deviation term is the critical factor driving differences in normalization
 226 performance, we employ the batch-level calculation, which exhibited unstable performance in the
 227 previous section, to calculate the mean of advantage, and conduct ablation experiments on the
 228 standard deviation term. This can be formalized as:
 229

$$A_k^{\text{std}} = r_k - \text{mean}(\{r_j\}_{j=1}^K). \quad (2)$$

230 Our experiments reveal that, when training on easy data, the policy quickly converges to consistent
 231 behaviors, resulting in a highly concentrated reward distribution and a rapid decline in standard
 232 deviation. In this scenario, using standard deviation-based advantage normalization can cause the
 233 denominator to become extremely small, excessively amplifying reward and advantage values. This
 234 leads to abnormally large gradients, destabilizes training, and may even cause gradient explosions.
 235 These findings empirically demonstrate that the standard deviation term plays a crucial role in
 236 advantage normalization.
 237

238 In summary, our experiments and analysis underscore that, in scenarios where reward distributions
 239 are highly concentrated, omitting the standard deviation from advantage normalization effectively
 240 prevents abnormal gradient amplification, thereby improving the stability and robustness of model
 241 training. However, for tasks characterized by inherently higher reward variance, either normalization
 242 approach is generally sufficient to maintain stable optimization.
 243

244 4.1.3 RECONSTRUCT A ROBUST NORMALIZATION TECHNIQUE

246 Takeaway 2

247 Calculating the mean at the local (group) level and the standard deviation at the global (batch)
 248 level enables more robust reward shaping.
 249

270
 271 Section 4.1.2 highlights the critical role of stan-
 272 dard deviation in advantage normalization. This
 273 raises the question: is there a more robust and
 274 effective combination of mean and standard de-
 275 viation for reward shaping? To explore this, we
 276 adopted the stable group-level mean calculation
 277 method demonstrated in section 4.1.1, paired
 278 with two approaches for computing the standard
 279 deviation: local (group-level) and global (batch-
 280 level). We then evaluated the performance of
 281 these combinations across two model sizes.
 282


283 The results, presented in Figures 6, reveal that global-level calculation exhibits a clear advantage. We
 284 attribute this to the batch-level standard deviation providing stronger normalization by effectively
 285 reducing gradient magnitudes, thereby preventing excessive policy updates. This approach aligns
 286 more effectively with the biased reward signals common in sparse rewards and coarse-grained
 287 advantage fitting, resulting in more stable and robust learning behavior. Furthermore, our experimental
 288 results support a claim from Hu et al. (2025) that batch-level normalization, or even subtracting the
 289 local mean and dividing by the batch standard deviation in certain scenarios, performs better.

290 4.2 CLIP-HIGHER

291 The Clip mechanism, while improving PPO training stability (Huang et al., 2024a), can cause issues
 292 in LLM-based text generation by overly suppressing low-probability tokens (Yu et al., 2025), leading
 293 to entropy collapse—where model output becomes overly deterministic and loses diversity (Jin
 294 et al., 2024). This reduction in entropy reinforces high-probability patterns and shrinks exploration,
 295 impairing performance on tasks requiring novel reasoning.

$$J_{DAPo}(\theta) = (r_{i,t}(\theta), 1 - \varepsilon_{low}, 1 + \varepsilon_{high}). \quad (3)$$

296 Here, a greater upper bound ε_{high} enables more exploration for low-probability tokens, mitigating
 297 entropy collapse. However, there's limited analysis on when and how to set this upper bound
 298 effectively. In this section, we address these gaps with targeted experiments and recommendations.



300
 301 Figure 7: Entropy comparison across different models with Clip-Higher. **A higher clip upper bound**
 302 **can mitigate the entropy drop in aligned models.**
 303

311 4.2.1 IN WHICH SETTINGS SHOULD WE CLIP HIGHER

313 Takeaway 3

314 For models with stronger fundamental reasoning abilities, increasing the clip higher parameter
 315 is more likely to facilitate exploration of better solution paths.

316
 317 As illustrated in Figure 7, experimental results indicate that the impact of increasing the upper
 318 clipping bound ε_{high} is model-dependent. For the base models, adjusting the upper clipping value
 319 yields minor effects on policy entropy. In contrast, aligned models exhibit a markedly different
 320 response: raising the upper clipping bound notably slows the entropy collapse, leading to consistent
 321 performance improvements in downstream evaluation metrics.

322 This disparity is mainly due to the base models' low clipping rate (around 0.003), which results in
 323 minimal policy updates and limited exploration because of their simple policy expressiveness. As a

324 result, increasing the clipping upper bound has little effect on learning outcomes. In contrast, aligned
 325 models demonstrate improved reasoning and generalization abilities due to advanced post-training
 326 techniques (Yang et al., 2025). Figure 18 in Appendix D, illustrates that, compared to base models,
 327 aligned models initially have fewer high-probability tokens. Therefore, a higher clipping upper
 328 bound reduces token probability disparity and mitigates entropy collapse. It expands the range of
 329 permissible policy updates, promoting more diverse action sampling and enhanced exploration during
 330 training. This mechanism maintains higher entropy while increasing the likelihood of finding optimal
 331 solutions, as shown by improved evaluation metrics.

332 333 4.2.2 ANALYZING THE EFFECTIVENESS OF CLIP-HIGHER FROM A LINGUISTIC PERSPECTIVE

334 Takeaway 4

336 **Traditional clipping** may constrain innovative reasoning structure generation. **Clipping**
 337 **higher** allows the model to explore a broader range of discourse reasoning structures.

339 Building on our token-level demonstration of Clip-Higher’s behavior in section 4.2.1, we now analyze
 340 its impact on reasoning logic through token-level linguistics. As shown in Figure 20 in Appendix D,
 341 an upper bound of 0.2 imposes strict constraints on policy updates, primarily affecting connective
 342 tokens like *therefore*”, *if*”, and “*but*” by limiting large probability deviations. This leads to fewer
 343 opportunities for the model to generate innovative or diverse logical structures. By increasing the
 344 upper bound to 0.28, the model gains more flexibility, resulting in fewer tokens being clipped and
 345 a shift in clipping focus toward high-frequency function words such as “*is*”, “*the*”, and “*,*”. This
 346 encourages richer reasoning paths while ensuring sentence stability through selective clipping of
 347 common function words.

348 349 4.2.3 HOW TO SET THE UPPER BOUND FOR ADVANTAGE CLIPPING


350 Takeaway 5

352 There appears to be a “scaling law” between the performance and the upper bound of the
 353 clipping on the **small-sized model**, which does not exist on **larger models**.

355 Section 4.2.1 shows that Clip-Higher improves
 356 aligned models. While most works use the
 357 default clip upper bound of 0.28 (Yu et al., 2025),
 358 we believe that different models have differ-
 359 ent preferences for this parameter. To verify
 360 this conjecture, we uniformly evaluate different
 361 clip upper bounds ranging from 0.2 to 0.32 on
 362 two model sizes. As shown in Figure 8, the
 363 4B model improves progressively with higher
 364 bounds, peaking at 0.32. In contrast, the 8B
 365 model performs best at 0.28, with no consistent
 366 gain beyond that.

367 368 4.3 LOSS AGGREGATION

369 The strategy of loss aggregation directly determines the contribution of each sample or token to the
 370 overall gradient during optimization (Liu et al., 2025a). Common strategies include token-level and
 371 sequence-level aggregation. The sequence-level aggregation adopted by GRPO (Shao et al., 2024)
 372 first averages the loss across all tokens within each sample, then averages these per-response losses
 373 across the batch, thereby assigning equal weight to each response regardless of its length. However,
 374 Yu et al. (2025) highlights a flaw in this method: longer responses possess a diminished influence
 375 per token on the total loss, hindering the model’s ability to learn effectively from longer, complex
 376 responses. This can reduce the model’s capacity to learn from long, complex answers, and may bias
 377 optimization toward brevity, since shorter correct responses receive larger gradient updates, while
 longer incorrect responses are insufficiently penalized (Liu et al., 2025b).



378 Figure 8: Test accuracy of aligned models (trained
 379 on medium data) with various clipping upper
 380 bounds.



378 4.3.1 DOES TOKEN-LEVEL LOSS AGGREGATION SUIT ALL SETTINGS?
379380 **Takeaway 6**
381382 Compared to sequence-level calculation, token-level loss proves to be more effective on Base
383 models, while showing limited improvement on Instruct models.
384385 To systematically evaluate loss aggregation
386 strategies, we compare sequence-level and
387 token-level loss aggregation on both base and
388 aligned versions of Qwen3-8B. As illustrated in
389 As shown in Figure 9, token-level loss improves
390 convergence and peak accuracy for base models
391 by balancing token contributions, particularly
392 on challenging data. In contrast, this advantage
393 disappears in aligned models, where sequence-
394 level aggregation achieves better convergence
395 speed and final performance. Analysis suggests
396 that aligned models already exhibit stable rea-
397 soning, making fine-grained token weighting unnecessary or harmful. In these cases, sequence-level
398 aggregation better preserves the structure and consistency of high-quality, aligned outputs. These
399 findings highlight that optimal aggregation is model-dependent: token-level suits base models;
400 response-level is preferable for instruction-tuned ones.
401

Figure 9: Accuracy comparison between sequence-level loss and token-level loss. Results are reported on both Medium and Hard Datasets.

402 4.4 OVERLONG FILTERING
403404 To improve efficiency, LLMs often use fixed maximum generation lengths for truncation (Chen et al.,
405 2025; Team et al., 2025). However, this may truncate multi-step reasoning early in training, causing
406 incomplete outputs to be mislabeled as negatives and introducing harmful noise. Overlong filtering
407 (Yu et al., 2025) mitigates this by masking rewards of over-length responses, preserving robustness
408 and reasoning quality (He et al., 2025a). Yet, its sensitivity to mask thresholds remains underexplored,
409 leaving optimal settings unclear.
410411 Figure 10: Total test accuracy and response length of Qwen3-8B-Base over training iterations under
412 different maximum generation lengths.
413414 4.4.1 WHEN TO USE THE OVERLONG FILTERING
415416 **Takeaway 7**
417418 Overlong filtering shows limited effectiveness on long-tail reasoning tasks; however, it can
419 enhance the accuracy and clarity of responses in medium and short-length reasoning tasks.
420421 The results in Figure 10 shows that overlong filtering improves learning at a short threshold (8k),
422 but benefits diminish at 20k. Response length analysis reveals the cause: under 20k, filtered models
423 generate longer outputs than the vanilla policy; at 8k, responses become shorter. As Figure 11
424 (Left) shows, in the 20k setting, frequent clipping occurs due to repetition or non-termination—signs
425 of degenerate generation—indicating the mask removes unproductive outputs. In contrast, the 8k
426 threshold also filters extended but valid reasoning, encouraging conciseness and reducing verbosity.
427




Figure 11: **Left:** Repeat ratios for correct (reward = 1) and incorrect (reward = 0) generations under varying maximum generation lengths. **Right:** Repeat ratios in truncated samples with vs. without overlong filtering. Repetition rate statistics are provided in Appendix E.1.

As illustrated in Figure 11 (Right), during RL training, the proportion of "repetitive but unable to terminate" overlong samples increases, indicating degraded EOS modeling and leading to output redundancy and termination failures. With the overlong mask, this proportion drops and remains low, enabling the model to better distinguish "generation completed" from "truncated" samples, avoiding invalid learning. More importantly, the mechanism helps policies accurately model termination, preventing them from mistakenly penalizing unfinished outputs as negative examples.

5 A SIMPLE COMBINATION: LITE PPO




Figure 12: Test accuracy of non-aligned models trained via three RL methods, i.e., Lite PPO (ours), GRPO (Shao et al., 2024) and DAPO (Yu et al., 2025).

Building on the in-depth mechanism analysis, we derive two practical guidelines for non-aligned models: (i) For small and medium-sized base models, advantage normalization (Section 4.1.3) significantly boosts performance by transforming sparse rewards into robust signals via group-mean and batch-std normalization. (ii) Token-level loss aggregation is also highly effective, as shown in Section 4.3.1, particularly for base model architectures.

Therefore, we integrate both techniques, called Lite PPO, into non-aligned models that use the vanilla PPO loss without the critic. As shown in Figure 17, Lite PPO outperforms GRPO and DAPO, a technique-heavy method with *Group-level Normalization*, *Clip-Higher*, *Overlong Reward Shaping*, *Token-level Loss*, *Dynamic Sampling*. Specifically, Lite PPO exhibits a stable upward trend on non-aligned models, while other policies collapse after peaking. This advantage arises from the normalization in Takeaway 2, which mitigates interference from homogeneous reward distributions in mixed datasets. Additionally, this gain stems from adopting token-level loss aggregation, which is more effective for base models.

6 CONCLUSION

We present a systematic evaluation of RL techniques for LLMs under a unified framework, addressing fragmentation in methodology and practice. By analyzing normalization, clipping, and filtering, we reveal that simplicity outperforms complexity: Lite PPO, combining only two core techniques, surpasses heavily engineered algorithms. Our findings highlight the importance of context-aware design and challenge the trend of over-complication in RL4LLM. We provide actionable guidelines for technique selection and advocate for standardized, reproducible practices that balance theoretical soundness with practical efficiency.

486 **7 ETHICS STATEMENT**
 487

488 This work adheres to the ICLR Code of Ethics. In this study, no human subjects or animal experimen-
 489 tation was involved. All datasets used, including Easy, Medium and Hard datasets, were sourced in
 490 compliance with relevant usage guidelines, ensuring no violation of privacy. We have taken care to
 491 avoid any biases or discriminatory outcomes in our research process. No personally identifiable infor-
 492 mation was used, and no experiments were conducted that could raise privacy or security concerns.
 493 We are committed to maintaining transparency and integrity throughout the research process.

494
 495 **8 REPRODUCIBILITY STATEMENT**
 496

497 We have made extensive efforts to ensure the reproducibility of our work. The baseline models we
 498 evaluated are detailed in Section 3.2 and Appendix D. We provide not only the specific models and
 499 API versions, but also the exact sampling parameters we used. The evaluation metrics are described
 500 in Section 3.1. We will release our benchmark dataset and evaluation code upon paper acceptance to
 501 facilitate reproduction and future research by the community.

502
 503 **REFERENCES**
 504

505 Arash Ahmadian, Chris Cremer, Matthias Gallé, Marzieh Fadaee, Julia Kreutzer, Olivier Pietquin,
 506 Ahmet Üstün, and Sara Hooker. Back to basics: Revisiting reinforce-style optimization for
 507 learning from human feedback in llms. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar
 508 (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics
 509 (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pp. 12248–12267.
 510 Association for Computational Linguistics, 2024. doi: 10.18653/V1/2024.ACL-LONG.662. URL
 511 <https://doi.org/10.18653/v1/2024.acl-long.662>.

512 Marcin Andrychowicz, Anton Raichuk, Piotr Stanczyk, Manu Orsini, Sertan Girgin, Raphaël Marinier,
 513 Léonard Huszenot, Matthieu Geist, Olivier Pietquin, Marcin Michalski, Sylvain Gelly, and Olivier
 514 Bachem. What matters in on-policy reinforcement learning? A large-scale empirical study. *CoRR*,
 515 abs/2006.05990, 2020. URL <https://arxiv.org/abs/2006.05990>.

516 Aili Chen, Aonian Li, Bangwei Gong, Binyang Jiang, Bo Fei, Bo Yang, Boji Shan, Changqing Yu,
 517 Chao Wang, Cheng Zhu, et al. Minimax-m1: Scaling test-time compute efficiently with lightning
 518 attention. *arXiv preprint arXiv:2506.13585*, 2025.

519 Xiangxiang Chu, Hailang Huang, Xiao Zhang, Fei Wei, and Yong Wang. GPG: A simple and strong
 520 reinforcement learning baseline for model reasoning. *CoRR*, abs/2504.02546, 2025. doi: 10.48550/
 521 ARXIV.2504.02546. URL <https://doi.org/10.48550/arXiv.2504.02546>.

523 Logan Engstrom, Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Firdaus Janoos, Larry Rudolph,
 524 and Aleksander Madry. Implementation matters in deep policy gradients: A case study on PPO
 525 and TRPO. *CoRR*, abs/2005.12729, 2020. URL <https://arxiv.org/abs/2005.12729>.

526 Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi
 527 Hu, Xu Han, Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu, and Maosong Sun.
 528 Olympiadbench: A challenging benchmark for promoting AGI with olympiad-level bilingual
 529 multimodal scientific problems. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.),
 530 *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume
 531 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pp. 3828–3850. Associa-
 532 tion for Computational Linguistics, 2024. doi: 10.18653/V1/2024.ACL-LONG.211. URL
 533 <https://doi.org/10.18653/v1/2024.acl-long.211>.

534 Jujie He, Jiacai Liu, Chris Yuhao Liu, Rui Yan, Chaojie Wang, Peng Cheng, Xiaoyu Zhang, Fuxiang
 535 Zhang, Jiacheng Xu, Wei Shen, et al. Skywork open reasoner 1 technical report. *arXiv preprint
 536 arXiv:2505.22312*, 2025a.

538 Zhiwei He, Tian Liang, Jiahao Xu, Qiuzhi Liu, Xingyu Chen, Yue Wang, Linfeng Song, Dian Yu,
 539 Zhenwen Liang, Wenxuan Wang, Zhuosheng Zhang, Rui Wang, Zhaopeng Tu, Haitao Mi, and Dong
 Yu. Deepmath-103k: A large-scale, challenging, decontaminated, and verifiable mathematical

540 dataset for advancing reasoning. *CoRR*, abs/2504.11456, 2025b. doi: 10.48550/ARXIV.2504.
 541 11456. URL <https://doi.org/10.48550/arXiv.2504.11456>.

542

543 Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob
 544 Steinhardt. Measuring massive multitask language understanding. In *9th International Conference
 545 on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net,
 546 2021. URL <https://openreview.net/forum?id=d7KBjmI3GmQ>.

547

548 Jian Hu, Jason Klein Liu, Haotian Xu, and Wei Shen. Reinforce++: An efficient rlhf algorithm with
 549 robustness to both prompt and reward models, 2025. URL <https://arxiv.org/abs/2501.03262>.

550

551 Yujing Hu, Weixun Wang, Hangtian Jia, Yixiang Wang, Yingfeng Chen, Jianye Hao, Feng Wu, and
 552 Changjie Fan. Learning to utilize shaping rewards: A new approach of reward shaping. *Advances
 553 in Neural Information Processing Systems*, 33:15931–15941, 2020.

554

555 Nai-Chieh Huang, Ping-Chun Hsieh, Kuo-Hao Ho, and I-Chen Wu. Ppo-clip attains global optimality:
 556 Towards deeper understandings of clipping. In Michael J. Wooldridge, Jennifer G. Dy, and Sriraam
 557 Natarajan (eds.), *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-
 558 Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth
 559 Symposium on Educational Advances in Artificial Intelligence, EAAI 2024, February 20-27, 2024,
 560 Vancouver, Canada*, pp. 12600–12607. AAAI Press, 2024a. doi: 10.1609/AAAI.V38I11.29154.
 561 URL <https://doi.org/10.1609/aaai.v38i11.29154>.

562

563 Shengyi Huang, Michael Noukhovitch, Arian Hosseini, Kashif Rasul, Weixun Wang, and Lewis
 564 Tunstall. The N+ implementation details of RLHF with PPO: A case study on tl;dr summarization.
 565 *CoRR*, abs/2403.17031, 2024b. doi: 10.48550/ARXIV.2403.17031. URL <https://doi.org/10.48550/arXiv.2403.17031>.

566

567 Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Os-
 568 trow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint
 569 arXiv:2410.21276*, 2024. URL <https://arxiv.org/abs/2410.21276>.

570

571 Ruinan Jin, Shuai Li, and Baoxiang Wang. On stationary point convergence of ppo-clip. In *The
 572 Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria,
 573 May 7-11, 2024*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=uznKlCpWjV>.

572

573 Wouter Kool, Herke van Hoof, and Max Welling. Buy 4 REINFORCE samples, get a baseline for
 574 free!, 2019. URL <https://openreview.net/forum?id=r11gTGL5DE>.

575

576 Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay V.
 577 Ramasesh, Ambrose Sloane, Cem Anil, Imanol Schlag, Theo Gutman-Solo, Yuhuai Wu, Behnam
 578 Neyshabur, Guy Gur-Ari, and Vedant Misra. Solving quantitative reasoning problems with language
 579 models. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (eds.),
 580 *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information
 581 Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December
 582 9, 2022*, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/18abbeef8cfe9203fdf9053c9c4fe191-Abstract-Conference.html.

583

584 Mingjie Liu, Shizhe Diao, Ximing Lu, Jian Hu, Xin Dong, Yejin Choi, Jan Kautz, and Yi Dong.
 585 Prorl: Prolonged reinforcement learning expands reasoning boundaries in large language models.
 586 *CoRR*, abs/2505.24864, 2025a. doi: 10.48550/ARXIV.2505.24864. URL <https://doi.org/10.48550/arXiv.2505.24864>.

587

588 Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min
 589 Lin. Understanding r1-zero-like training: A critical perspective. *CoRR*, abs/2503.20783, 2025b.
 590 doi: 10.48550/ARXIV.2503.20783. URL <https://doi.org/10.48550/arXiv.2503.20783>.

591

592 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy
 593 optimization algorithms. *CoRR*, abs/1707.06347, 2017. URL <http://arxiv.org/abs/1707.06347>.

594 John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional
 595 continuous control using generalized advantage estimation, 2018. URL <https://arxiv.org/abs/1506.02438>.
 596

597 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
 598 Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of
 599 mathematical reasoning in open language models. *arXiv preprint arXiv: 2402.03300*, 2024. URL
 600 <https://arxiv.org/abs/2402.03300v3>.
 601

602 Richard S. Sutton, David A. McAllester, Satinder Singh, and Yishay Mansour. Policy gra-
 603 dient methods for reinforcement learning with function approximation. In Sara A. Solla,
 604 Todd K. Leen, and Klaus-Robert Müller (eds.), *Advances in Neural Information Process-
 605 ing Systems 12, [NIPS Conference, Denver, Colorado, USA, November 29 - December 4,
 606 1999]*, pp. 1057–1063. The MIT Press, 1999. URL <http://papers.nips.cc/paper/1713-policy-gradient-methods-for-reinforcement-learning-with-function-approximation.pdf>.
 607

608 Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun
 609 Xiao, Chenzhuang Du, Chonghua Liao, et al. Kimi k1.5: Scaling reinforcement learning with
 610 llms. *arXiv preprint arXiv:2501.12599*, 2025.
 611

612 Weixun Wang, Shaopan Xiong, Gengru Chen, Wei Gao, Sheng Guo, Yancheng He, Ju Huang, Jiaheng
 613 Liu, Zhendong Li, Xiaoyang Li, Zichen Liu, Haizhou Zhao, Dakai An, Lunxi Cao, Qiyang Cao,
 614 Wanxi Deng, Feilei Du, Yiliang Gu, Jiahe Li, Xiang Li, Mingjie Liu, Yijia Luo, Zihe Liu, Yadao
 615 Wang, Pei Wang, Tianyuan Wu, Yanan Wu, Yuheng Zhao, Shuaibing Zhao, Jin Yang, Siran Yang,
 616 Yingshui Tan, Huimin Yi, Yuchi Xu, Yujin Yuan, Xingyao Zhang, Lin Qu, Wenbo Su, Wei Wang,
 617 Jiamang Wang, and Bo Zheng. Reinforcement learning optimization for large-scale learning: An
 618 efficient and user-friendly scaling library. *CoRR*, abs/2506.06122, 2025. doi: 10.48550/ARXIV.
 619 2506.06122. URL <https://doi.org/10.48550/arXiv.2506.06122>.
 620

621 Siwei Wu, Zhongyuan Peng, Xinrun Du, Tuney Zheng, Minghao Liu, Jialong Wu, Jiachen Ma,
 622 Yizhi Li, Jian Yang, Wangchunshu Zhou, Qunshu Lin, Junbo Zhao, Zhaoxiang Zhang, Wenhao
 623 Huang, Ge Zhang, Chenghua Lin, and Jiaheng Liu. A comparative study on reasoning patterns
 624 of openai’s o1 model. *CoRR*, abs/2410.13639, 2024. doi: 10.48550/ARXIV.2410.13639. URL
<https://doi.org/10.48550/arXiv.2410.13639>.
 625

626 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang
 627 Gao, Chengan Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu,
 628 Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin
 629 Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang,
 630 Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui
 631 Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang
 632 Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger
 633 Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan
 634 Qiu. Qwen3 technical report, 2025. URL <https://arxiv.org/abs/2505.09388>.
 635

636 Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Tiantian Fan, Gaohong
 637 Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi
 638 Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu, Jiaze Chen, Jiangjie Chen, Chengyi
 639 Wang, Hongli Yu, Weinan Dai, Yuxuan Song, Xiangpeng Wei, Hao Zhou, Jingjing Liu, Wei-Ying
 640 Ma, Ya-Qin Zhang, Lin Yan, Mu Qiao, Yonghui Wu, and Mingxuan Wang. DAPO: an open-source
 641 LLM reinforcement learning system at scale. *CoRR*, abs/2503.14476, 2025. doi: 10.48550/ARXIV.
 642 2503.14476. URL <https://doi.org/10.48550/arXiv.2503.14476>.
 643

644 Weihao Zeng, Yuzhen Huang, Qian Liu, Wei Liu, Keqing He, Zejun Ma, and Junxian He. Simplerl-
 645 zoo: Investigating and taming zero reinforcement learning for open base models in the wild. *CoRR*,
 646 abs/2503.18892, 2025. doi: 10.48550/ARXIV.2503.18892. URL <https://doi.org/10.48550/arXiv.2503.18892>.
 647

648 Jixiao Zhang and Chunsheng Zuo. GRPO-LEAD: A difficulty-aware reinforcement learning approach
 649 for concise mathematical reasoning in language models. *CoRR*, abs/2504.09696, 2025. doi: 10.
 650 48550/ARXIV.2504.09696. URL <https://doi.org/10.48550/arXiv.2504.09696>.
 651

648 Xiaojiang Zhang, Jinghui Wang, Zifei Cheng, Wenhao Zhuang, Zheng Lin, Minglei Zhang, Shaojie
649 Wang, Yinghan Cui, Chao Wang, Junyi Peng, Shimiao Jiang, Shiqi Kuang, Shouyu Yin, Chaohang
650 Wen, Haotian Zhang, Bin Chen, and Bing Yu. SRPO: A cross-domain implementation of large-scale
651 reinforcement learning on LLM. *CoRR*, abs/2504.14286, 2025. doi: 10.48550/ARXIV.2504.14286.
652 URL <https://doi.org/10.48550/arXiv.2504.14286>.

653 Rui Zheng, Shihan Dou, Songyang Gao, Yuan Hua, Wei Shen, Binghai Wang, Yan Liu, Senjie Jin,
654 Yuhao Zhou, Limao Xiong, Lu Chen, Zhiheng Xi, Nuo Xu, Wenbin Lai, Minghao Zhu, Haoran
655 Huang, Tao Gui, Qi Zhang, and Xuanjing Huang. Delve into PPO: Implementation matters for
656 stable RLHF. In *NeurIPS 2023 Workshop on Instruction Tuning and Instruction Following*, 2023.
657 URL <https://openreview.net/forum?id=rxEmiOEIfL>.

658 Terry Yue Zhuo, Minh Chien Vu, Jenny Chim, Han Hu, Wenhao Yu, Ratnadira Widyasari, Imam
659 Nur Bani Yusuf, Haolan Zhan, Junda He, Indraneil Paul, Simon Brunner, Chen Gong, James
660 Hoang, Armel Randy Zebaze, Xiaoheng Hong, Wen-Ding Li, Jean Kaddour, Ming Xu, Zhihan
661 Zhang, Prateek Yadav, and et al. Bigcodebench: Benchmarking code generation with diverse
662 function calls and complex instructions. In *The Thirteenth International Conference on Learning
663 Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. URL <https://openreview.net/forum?id=YrycTj11L0>.

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 A LLM USAGE STATEMENT
703
704705 Large Language Models (LLMs) were used to aid in the writing and polishing of the manuscript.
706 Specifically, we used an LLM to assist in refining the language, improving readability, and ensuring
707 clarity in various sections of the paper. The model helped with tasks such as sentence rephrasing,
708 grammar checking, and enhancing the overall flow of the text.709 It is important to note that the LLM was not involved in the ideation, research methodology, or
710 experimental design. All research concepts, ideas, and analyses were developed and conducted by
711 the authors. The contributions of the LLM were solely focused on improving the linguistic quality of
712 the paper, with no involvement in the scientific content or data analysis.713 The authors take full responsibility for the content of the manuscript, including any text generated or
714 polished by the LLM. We have ensured that the LLM-generated text adheres to ethical guidelines and
715 does not contribute to plagiarism or scientific misconduct.718 B DETAILED EXPERIMENTAL SETUP
719
720721 B.1 PARAMETERS
722
723724 We employ ROLL, a user-friendly and efficient open-source reinforcement learning framework, to
725 implement our pipeline. Subsequently, the key parameters observed during the training process are
726 presented as follows. See our code config file for more details on the parameters.729 B.2 PROMPT
730
731732 In this work, we incorporate the following instruction into the system prompt to encourage the model
733 to better demonstrate its reasoning process: **“Please reason step by step, and put your final answer
734 within \boxed{}.”** This setting is designed to guide the model to perform step-by-step reasoning
735 and explicitly present the final answer in the form of \boxed{}, thereby enhancing the clarity and
readability of the output.738 B.3 TRAINING DATASETS
739
740741 To ensure reproducibility and fairness, we exclusively use open-source datasets for training, including
742 *SimpleRL-Zoo-Data* (Zeng et al., 2025) and *Deepmath* (He et al., 2025b). To comprehensively
743 examine how problem difficulty affects the RL technique’s performance, we randomly sample from
744 the datasets, removing an excessive proportion of examples whose ground-truth label is simply “True”
745 or “False”. This adjustment addresses the **ostensible positive phenomenon**, where models produce
746 correct binary answers from erroneous reasoning chains, thereby introducing noisy supervision that
747 compromises training quality (please refer to Appendix E.2 for case studies).748
749
750
751
752
753
754
755

- Easy Data : We randomly sample 5,000 entries from SimpleRL-Zoo-Data-Easy, which consists of problems drawn from GSM8K and MATH-500-level1.
- Medium Data: We select the 5,000 easiest examples from the *DeepMath-103k* dataset, based on their assigned difficulty annotations.
- Hard Data: We randomly sample 5,000 entries from *DeepMath-103k*, with sampling probability proportional to each entry’s assigned difficulty level.

```

756
757     seed: 42
758     max_steps: 500
759     save_steps: 20
760     logging_steps: 1
761     eval_steps: 1
762
763     rollout_batch_size: 128
764     prompt_length: 1024
765     response_length: 8000
766
767     ppo_epochs: 1
768     adv_estimator: "reinforce"
769     init_kl_coef: 0.0
770     async_generate_level: 1
771
772     actor_train:
773         training_args:
774             learning_rate: 1.0e-6
775             weight_decay: 0
776             per_device_train_batch_size: 4
777             gradient_accumulation_steps: 32
778             # warmup_ratio: 0.1
779             warmup_steps: 50
780             num_train_epochs: 50
781
782         ...
783
784     actor_infer:
785         generating_args:
786             max_new_tokens: ${response_length}
787             top_p: 0.99
788             top_k: 100
789             num_beams: 1
790             temperature: 0.99
791             num_return_sequences: 8
792
793     ...

```

C DETAILED EXPERIMENTAL RESULTS

C.1 BASELINE




Figure 13: Test accuracy and response length of four model variants: Qwen3-4B-Base, Qwen3-8B-Base, Qwen3-4B, and Qwen3-8B across different data difficulty.




Figure 14: **Middle 2 rows:** Accuracy over training iterations of Base models. The first row presents results of Qwen3-4B-Base. The second row shows results of Qwen3-8B-Base. **Bottom 2 rows:** Accuracy over training iterations of aligned models. The first row presents results of Qwen3-4B, while the second row shows results of Qwen3-8B.

C.2 RESULTS OF ALIGNED MODEL

Following our analysis in section 4.2.1 that Clip Higher provides stronger benefits on aligned models, we instantiated a variant of LitePPO that combines mixture normalization and Clip Higher, trained on both the Easy and Hard datasets.

As shown in Tables 1 and 2, our simple recipe significantly improves the reasoning ability of the aligned model and consistently outperforms DAPO, particularly when trained on the hard dataset. This confirms that LitePPO remains efficient for strong, instruction-tuned models.

	MATH-500	OlympiadBench	MinervaMath	AIME24	AIME25	AMC23	Average
DAPO	90.57	64.82	43.43	47.92	37.54	79.69	60.66
Lite PPO	91.43	66.01	44.72	49.17	39.81	81.88	62.17

Table 1: Results using **Qwen3-8B (aligned model)** trained on the **Easy dataset**.

	MATH-500	OlympiadBench	MinervaMath	AIME24	AIME25	AMC23	Average
DAPO	93.02	69.15	45.73	52.50	39.17	85.31	64.17
Lite PPO	94.77	71.12	46.60	49.58	47.50	92.19	66.96

Table 2: Results using **Qwen3-8B (aligned model)** trained on the **Hard dataset**.

864 C.3 VALIDATING METHOD EXTENSIBILITY
865866 We have conducted additional experiments on Llama3-8B, we recorded the average score among
867 six benchmarks, i.e., AIME 24, AMC23, GSM8k, MATH-500, Minerva Math, OlympiadBench. The
868 results in Table 3 show that, compared to GRPO and DAPO, our Lite PPO demonstrates the best
869 average performance, indicating the portability of our method..
870871
872

	MATH-500	OlympiadBench	AIME 24	AMC23	GSM8k	Minerva Math	Average
GRPO	19.8	8.3	3.7	18.2	77.3	8.5	22.63
DAPO	25.1	7.6	4.6	18.4	77.1	9.2	23.67
Lite PPO	27.3	9.7	9.1	17.9	79.6	9.6	25.53

873 Table 3: Results on Llama3-8B.
874
875876 We have conducted additional experiments to evaluate the generalization performance of our method.
877 Specifically, we test the Qwen3-8b-Base trained by LitePPO on several out-of-domain reasoning tasks,
878 including coding (LCB-v5, LCB-v6), Interdisciplinary QA (GPQA), and Language Understanding
879 (MMLU-Pro). The results in Table 4 show that Lite PPO achieves the pass@1 score of 25.08 on
880 LCB-v5, 19.71 on LCB-v6, 46.63 on GPQA, and 62.38 on MMLU-Pro, demonstrating our method’s
881 excellent generalization ability.
882883
884

	LCB-v5	LCB-v6	GPQA	MMLU-Pro
GRPO	22.94	18.28	42.49	61.84
DAPO	24.01	19.00	45.08	60.60
LitePPO	25.08	19.71	46.63	62.38

885 Table 4: Pass@1 score on other reasoning modalities.
886
887888 D CASE STUDY OF CLIP HIGHER
889890 We present a comparison of token distributions between the base model and the aligned model.
891892 As shown in Figure 18, compared to the base model, the aligned model has very few preferred tokens
893 with high probability in the initial stage. Therefore, a higher clipping upper bound can effectively
894 bridge the probability gap between tokens and alleviate the entropy collapse. For these models,
895 raising the upper bound expands the permissible range of policy updates, which in turn facilitates
896 more diverse action sampling and enhances exploratory behavior during training. This mechanism
897 preserves higher entropy while simultaneously increasing the probability of identifying optimal
898 solutions, as evidenced by improved evaluation metrics.
899900 Building on our token-level demonstration of Clip-Higher’s behavior in section 4.2.1, we now analyze
901 its impact on reasoning logic through token-level linguistics. As illustrated in Figure 20, setting an
902 upper bound to 0.2 imposes stringent constraints on policy updates by limiting substantial probability
903 deviations for individual tokens. Under these stricter conditions, our analysis reveals that clipping
904 predominantly affects connective tokens such as “therefore”, “if”, and “but”. These tokens frequently
905 appear at the beginnings of sentences, serving as key semantic markers or transition words within
906 dialog generation. Such connectors often introduce new directions in reasoning. However, their
907 probability ratios between updated and old policies frequently exceed clipping thresholds, triggering
908 aggressive suppression in PPO optimization. While this traditional clipping ensures stability in
909 the overall token distribution, it may restrict the model’s capacity to generate innovative or diverse
910 argumentative reasoning structures by limiting flexibility in the use of discourse-level connectives.
911




Figure 15: Accuracy comparison of Base models with different standard deviation calculation. **Top 2 rows:** Accuracy of Qwen3-4B-Base with different standard deviation calculation. The first row uses the easy training dataset, while the second row uses the hard training dataset. **Bottom 2 rows:** Accuracy comparison of Qwen3-8B-Base with different standard deviation calculation. The first row uses the easy training dataset, while the second row uses the hard training dataset.




Figure 16: Test accuracy of sample-level loss and token-level loss on medium and extremely hard datasets.

E OVERLONG FILTER

E.1 REPEAT RATIO

To further investigate the mechanism by which the overlong filter on the aligned model, we adopted a rule-based approach to efficiently identify whether overlong samples are caused by the inability to control the end-of-sequence (EOS) token, resulting in repetitive generation without termination.




Figure 17: Test accuracy of non-aligned models trained via three RL methods, i.e., Lite PPO (ours), GRPO (Shao et al., 2024) and DAPO (Yu et al., 2025).




Figure 18: Predicted probability distributions of Qwen3-4B-Base (left) and Qwen3-4B (right) under two clipping upper bound $\in \{0.20, 0.28\}$.

Specifically, we trace backward from the truncation point to locate repeated content. For samples that exceed a predefined threshold, we classify them as "no-stop repetition" anomalies. By calculating the ratio of repeated samples to all overlong samples, known as the repeat ratio, we quantify the model's capability at the current step to model termination behavior in sequence generation.

E.2 EXAMPLES OF OSTENSIBLE POSITIVE PHENOMENA

As demonstrated in Figure 11 in the main text, we observe that models with weaker capabilities tend to continue generating content aimlessly even after correctly reasoning and providing the correct answer, until exceeding the output length limit. Such false positives, although receiving a reward




Figure 19: Predicted probability distributions of Qwen3-8B-Base (left) and Qwen3-8B (right) under two clipping upper bound $\in \{0.20, 0.28\}$.




Figure 20: **Left:** A case study under the same prompt across various clipping upper bounds. **Right:** The trigger differences of various upper bounds at the top 20 tokens with the highest clip frequencies.

of 1 through rule-based evaluation, introduce noise into the model during training. We present a representative case for illustration, as shown in Figure 21.

F DETAILED PRELIMINARIES

F.1 PROXIMAL POLICY OPTIMIZATION (PPO)

1068 Proximal Policy Optimization (PPO)(Schulman et al., 2017) is a widely used actor-critic algorithm
1069 grounded in the policy gradient framework. It improves the stability of policy learning by optimizing
1070 a clipped surrogate objective that restricts the divergence between the new and old policies during
1071 training. The PPO objective is:

$$\mathcal{J}_{\text{PPO}}(\theta) = \mathbb{E}_{[q \sim P(Q), o \sim \pi_{\theta_{\text{old}}}(O|q)]} \left[\frac{1}{|o|} \sum_{t=1}^{|o|} \min \left(\frac{\pi_{\theta}(o_t|q, o_{<t})}{\pi_{\theta_{\text{old}}}(o_t|q, o_{<t})} A_t, \text{clip} \left(\frac{\pi_{\theta}(o_t|q, o_{<t})}{\pi_{\theta_{\text{old}}}(o_t|q, o_{<t})}, 1-\epsilon, 1+\epsilon \right) A_t \right) \right], \quad (4)$$

1079 where π_θ and $\pi_{\theta_{\text{old}}}$ denote the current and old policy models, respectively. q and o represent the
 1080 sampled question and output sequence, with o_t as the t -th token in o . ϵ is a clipping hyperparameter

1080
 1081 Given an angle $\theta = \frac{2\pi}{5}$, we must prove the equation $2\cos(2\theta) + 2\cos(\theta) + 1 = 0$.
 1082 First, use the double-angle identity for cosine, which states that $\cos(2\theta) = 2\cos^2(\theta) - 1$:
 1083 Substitute this into the given equation:
 1084
$$2\cos(2\theta) + 2\cos(\theta) + 1 = 2(2\cos^2(\theta) - 1) + 2\cos(\theta) + 1.$$

 1085
 1086 Therefore, the initial equation $2\cos(2\theta) + 2\cos(\theta) + 1 = 0$ holds true for $\theta = \frac{2\pi}{5}$. Thus, the answer is:
 1087
$$\boxed{0}$$
.
 1088 We have just shown step by step that the equation $2\cos(2\theta) + 2\cos(\theta) + 1 = 0$ holds true for $\theta = \frac{2\pi}{5}$, with final affirmation that $\boxed{0}$ is the appropriate answer.
 1089
 1090 Do you have any other questions?
 1091 \\\user
 1092 Next, prove that $\cos(\frac{2\pi}{5})$ is a root of the quartic polynomial $x^4 + x^3 - 6x^2 - x + 1 = 0$
 1093
 1094

1095 Figure 21: An ostensible positive case, which cannot be terminated after the answer is given at the
 1096 end of inference.
 1097

1098
 1099 for stabilizing updates. A_t is the advantage at step t , typically estimated via Generalized Advantage
 1100 Estimation (GAE) (Schulman et al., 2018). The objective encourages the new policy to improve
 1101 advantage-weighted probabilities while constraining changes within a trust region.
 1102

1103 F.2 GROUP RELATIVE POLICY OPTIMIZATION (GRPO)

1104 Group Relative Policy Optimization (GRPO), proposed in DeepSeekMath (Shao et al., 2024), eliminates
 1105 the value function (critic) and instead estimates the advantage by normalizing rewards within a
 1106 group of sampled responses for the same prompt. Specifically, for a prompt x with G responses and
 1107 associated rewards $\{r_i\}_{i=1}^G$, the group-normalized advantage is given by:
 1108

$$1111 \hat{A}_{i,t} = \frac{r_i - \text{mean}(\{r_i\}_{i=1}^G)}{\text{std}(\{r_i\}_{i=1}^G)}. \quad (5)$$

1112 The effectiveness of the above normalization method can be understood from the perspective of
 1113 reward shaping. By emphasizing the differences among candidate outputs for the same prompt, it
 1114 effectively preserves the reliability of the gradient signal, even in sparse reward settings (Hu et al.,
 1115 2020). Instead of adding a KL penalty to the reward, GRPO directly regularizes by directly adding the
 1116 KL divergence between the trained policy and the reference policy to the loss. The overall surrogate
 1117 objective is:
 1118

$$1119 \mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}_{[q \sim P(Q), \{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(O|q)]} \quad (6)$$

$$1120 \quad \frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \left\{ \min \left(r_{i,t}(\theta) \hat{A}_{i,t}, \text{clip}(r_{i,t}(\theta), 1-\epsilon, 1+\epsilon) \hat{A}_{i,t} \right) - \beta D_{\text{KL}} [\pi_{\theta} \parallel \pi_{\text{ref}}] \right\},$$

1121 where $r_{i,t}(\theta) = \frac{\pi_{\theta}(o_{i,t}|q, o_{i,<t})}{\pi_{\theta_{\text{old}}}(o_{i,t}|q, o_{i,<t})}$, ϵ and β are hyper-parameters, and D_{KL} denotes the KL divergence
 1122 between the learned policy and a reference policy π_{ref} .
 1123

1124 F.3 DECOUPLED CLIP AND DYNAMIC SAMPLING POLICY OPTIMIZATION (DAPO)

1125 Decoupled Clip and Dynamic Sampling Policy Optimization (DAPO) (Yu et al., 2025) is a recent RL
 1126 method designed to address the unique challenges in LLM reasoning. For each question q with gold
 1127 answer a , DAPO samples a group of G outputs $\{o_i\}_{i=1}^G$ from the old policy, computes their rewards,
 1128 and maximizes the following surrogate objective:
 1129

$$\begin{aligned}
1134 \quad \mathcal{J}_{\text{DAPO}}(\theta) &= \mathbb{E}_{[(q,a) \sim \mathcal{D}, \{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot|q)]} \\
1135 \\
1136 \quad &\frac{1}{\sum_{i=1}^G |o_i|} \sum_{i=1}^G \sum_{t=1}^{|o_i|} \left\{ \min \left(r_{i,t}(\theta) \hat{A}_{i,t}, \text{clip}(r_{i,t}(\theta), 1-\epsilon_{\text{low}}, 1+\epsilon_{\text{high}}) \hat{A}_{i,t} \right) \right\}, \\
1137 \\
1138 \quad &\quad (7)
\end{aligned}$$

1139 where $\hat{A}_{i,t}$ is the group-normalized advantage. In addition, DAPO decouples the upper and lower
1140 clipping ranges ($\epsilon_{\text{low}}, \epsilon_{\text{high}}$) to better support exploration, dynamically filters out samples where all
1141 responses are correct or incorrect, aggregates losses at the token level, and applies special reward
1142 shaping for overlong or truncated responses.

1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187