From Uniform to Heterogeneous: Tailoring Policy Optimization to Every Token's Nature

Anonymous ACL submission

Abstract

Reinforcement Learning has emerged as the fundamental technique for enhancing reasoning in LLMs. However, existing algorithms apply uniform optimization to all tokens, ignoring their different roles in reasoning. To address this limitation, we introduce **H**eterogeneous **A**daptive **P**olicy **O**ptimization (HAPO), a token-aware algorithm that adapts optimization based on every token's entropy. Our algorithm includes four key components: (1) **Adaptive Temperature Sampling** that adjusts sampling temperature in real time, promoting exploration at high-entropy tokens while preserving coherence at low-entropy ones. (2) **Token-Level Group Average Advantage Estimation** that estimates advantages at token level, jointly accounting for sequence-length effects while preserving non-biased treatment.(3) **Differential Advantage Redistribution** that leverages entropy and importance ratios to adjust advantages for tokens with clear signals. (4) **Asymmetric Adaptive Clipping** that allows aggressive probability reduction for noisy tokens while enabling exploration for critical tokens. Through systematic investigation of entropy and training dynamics, we embed token-level treatment into every stage. Extensive experiments on mathematical reasoning, code, and logic tasks across multiple models demonstrate HAPO's consistent superiority over DAPO.

1 Introduction

Reinforcement Learning from Human Feedback (Ouyang et al., 2022) has emerged as a fundamental technique for enhancing the reasoning capabilities of LLMs. State-of-the-art models including OpenAI o1 (OpenAI, 2024), DeepSeek-R1 (DeepSeek-AI et al., 2025), and Qwen3 series (Yang et al., 2025) have achieved remarkable performance gains, underscoring how carefully designed reinforcement learning frameworks can empower LLMs to tackle tasks requiring reasoning and systematic problem solving.

However, existing algorithms (Schulman et al., 2017b; Shao et al., 2024; Yu et al., 2025) share a fundamental limitation: they employ a uniform optimization across all tokens, failing to distinguish between tokens that represent critical reasoning paths vs routine patterns. This conflicts with the heterogeneous nature of language generation, where tokens serve different roles in the reasoning.

Recent works have pioneered token heterogeneity by using entropy to distinguish tokens with different roles. DAPO with forking tokens (Wang et al., 2025b) proves that only a minority of high-entropy tokens guide the optimization process. Archer (Wang et al., 2025a) divides tokens into high/low entropy groups, relaxing clipping bounds for high-entropy tokens to encourage exploration. EDGE-GRPO (Zhang et al., 2025) modifies advantages based on sequence-level entropy, assigning higher advantages to responses that are confident. However, these methods remain coarse-grained. Archer's binary discretization creates discontinuities at arbitrary boundaries, causing tokens with marginally different entropy values to receive drastically different treatments. EDGE-GRPO operates at sequence level, failing to distinguish between heterogeneous tokens within the same sequence.

These limitations motivate us to develop a fine-grained algorithm that tailors optimization to every token rather than coarse categorical treatments. We first conduct a systematic empirical analysis of entropy and training dynamics, delving into the mechanisms and requirements for heterogeneous treatment across different training stages:

Rollout Generation: Exploration-Exploitation Imbalance. High-entropy tokens encoding critical reasoning decisions are rare. We find the temperature dilemma: low temperatures preserve accuracy but suppress critical tokens, while high temperatures increase their occurrence but introduce noise.

Advantage Calculation: Reward Attribution Granularity. Standard sequence-level advantage

estimation ignores intra-sequence token heterogeneity and becomes problematic under DAPO's token-mean averaging loss where longer negative samples dominate gradients. Moreover, our importance ratio analysis shows that tokens require different update magnitudes regardless of their entropy values, yet uniform and entropy-based reweighting fail to identify critical reasoning decisions.

Clipping Loss Computation: Clipping Constraint Mismatch. Our clipping analysis reveals paradoxical patterns: low-entropy tokens (mostly formatting symbols) hit left bounds preventing probability reduction, while high-entropy tokens (crucial reasoning elements) hit right bounds limiting exploration. Uniform clipping thus protects noise while constraining exploration.

Based on the discoveries above, we introduce **Heterogeneous Adaptive Policy Optimization (HAPO)**. Rather than using entropy for coarse categorization, HAPO leverages entropy as a continuous signal, enabling fine-grained optimization that respects token heterogeneity:

- **Adaptive Temperature Sampling**: We dynamically adjust sampling temperature based on token entropy—increasing it for high-entropy tokens to promote exploration while reducing it for low-entropy tokens to maintain coherence. This resolves the accuracy-exploration trade-off in training data.
- **Token-Level Group Average Advantage Estimation**: We estimate advantages at the token level across the group, respecting intra-sequence heterogeneity. This also ensures balanced optimization between positive and negative samples while preserving DAPO's token mean loss benefits for long-sequence learning.
- **Differential Advantage Redistribution**: We leverage entropy and importance ratios to modulate advantages—amplifying advantages for high-entropy tokens with extreme ratios while suppressing low-entropy ones near unity, enabling fine-grained attribution aligned with each token's optimization needs.
- **Asymmetric Adaptive Clipping**: We implement reversed asymmetric clipping boundaries. Low-entropy tokens receive expanded left boundaries to enable aggressive noise suppression, while high-entropy tokens receive expanded right boundaries to facilitate exploration at critical decision points.

HAPO is the first algorithm to achieve fine-grained token-level heterogeneous optimization, treating each token individually based on its continuous entropy signal. Our experiments demonstrate substantial improvements on mathematical reasoning, code, and logic tasks across multiple models.

2 Background

Given a language model π_θ. States consist of the prompt q and previously generated tokens $o_{<t} = (o_1, o_2, ..., o_{t-1})$. Actions are tokens o_t from vocabulary $\mathcal{V}$. Rewards $R(o)$ are typically provided at the sequence level for the complete response o.

Group Relative Policy Optimization. GRPO (Shao et al., 2024) eliminates the value function by employing group normalization. For each prompt q, GRPO samples G responses $\{o_i\}_{i=1}^G$ and optimizes:

$$\mathcal{L}^{\text{GRPO}} = E_{(q,o)}\left[\frac{1}{G}\sum_{i=1}^{G}\frac{1}{|o_i|}\sum_{t=1}^{|o_i|}\min\Big(r_{i,t}A_{i,t}, \text{clip}(r_{i,t}, 1-\epsilon, 1+\epsilon)A_{i,t}\Big) - \beta D_{\text{KL}}\left[\pi_\theta \| \pi_{\text{ref}}\right]\right] \quad (1)$$

The advantage is estimated at sequence-level within the group as $A_{i,t} = \frac{R_i - \text{mean}(\{R_j\}_{j=1}^G)}{\text{std}(\{R_j\}_{j=1}^G)}$.

Decoupled Clip and Dynamic Ampling Policy Optimization. DAPO (Yu et al., 2025) modifies GRPO with token-mean loss—averaging over tokens rather than sequences—to preserve gradient contributions from longer sequences. It also introduces clip-higher with asymmetric bounds:

$$\mathcal{L}^{\text{DAPO}} = E_{(q,o)}\left[\frac{1}{\sum_{i=1}^{G}|o_i|}\sum_{i=1}^{G}\sum_{t=1}^{|o_i|}\min\left(r_{i,t}A_{i,t}, \text{clip}(r_{i,t}, 1-\epsilon_{\text{low}}, 1+\epsilon_{\text{high}})\hat{A}_{i,t}\right.\right. \quad (2)$$

The asymmetric bounds $\epsilon_{\text{high}} > \epsilon_{\text{low}}$ (typically $\epsilon_{\text{high}} = 0.28$, $\epsilon_{\text{low}} = 0.2$) allow larger probability increases for high-entropy token exploration.

DAPO with Forking Tokens. (Wang et al., 2025b) demonstrates that a minority of high entropy tokens function as critical decision points. This motivates selective optimization focusing on high-entropy tokens:

$$\mathcal{L}^{\text{Ent}} = E_{(q,o)}\left[\frac{1}{\sum_{i=1}^{G}|o_i|}\sum_{i=1}^{G}\sum_{t=1}^{|o_i|}I[H_t \geq \tau_\rho]\cdot\min\left(r_{i,t}A_{i,t}, \text{clip}(r_{i,t}, 1-\epsilon_{\text{low}}, 1+\epsilon_{\text{high}})A_{i,t}\right)\right. \quad (3)$$

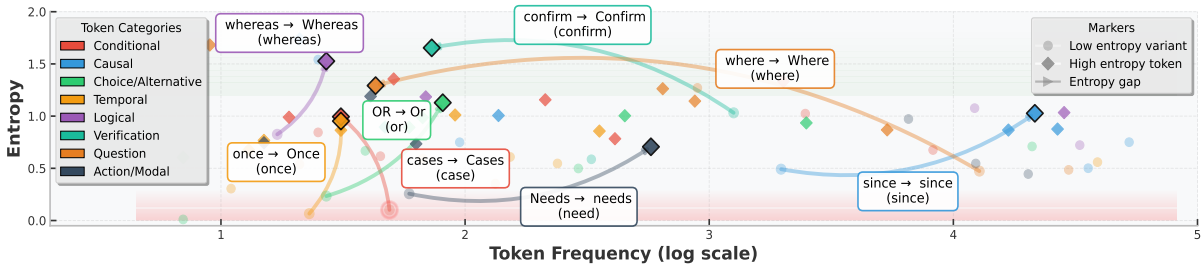

Figure 1: Dual-Entropy Token Frequency-Entropy Landscape

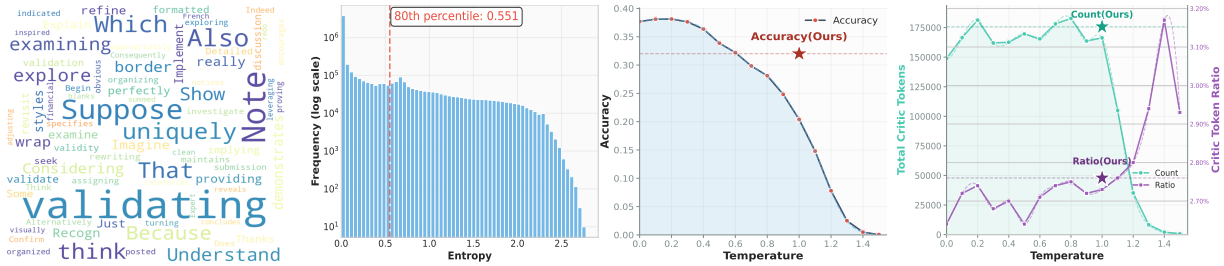

(a) Top100 High Entropy Tokens (b) Entropy Distribution (c) Accuracy (d) Critic Token Counts

Figure 2: Token entropy characteristics and temperature effects on critical tokens.

where H_t represents the entropy at position t, τ_ρ denotes the ρ-th percentile entropy threshold.

3 Token Heterogeneity in RLHF

3.1 Why Heterogeneity Is Necessary

The Dual-Entropy Phenomenon. We first analyzed entropy patterns during training. As shown in Figures 1, we discovered "Dual-Entropy Phenomenon": high-entropy tokens frequently have "twin siblings" in low-entropy region—tokens with same word stems but drastically different entropy. We show more visualizations in Appendix D.

This pattern spans multiple semantic categories and frequency ranges, explaining why low-entropy tokens are crucial: they contain important reasoning tokens that differ from high-entropy ones only in syntactic context. This reveals limitations of entropy-based binary discretization and uniform treatment, necessitating a nuanced approach that preserves low-entropy tokens' stabilizing influence without letting them dominate the learning signal.

3.2 Rollout Generation: Exploration-Exploitation Imbalance

Characterizing Tokens in Rollout Generation. We then analyze token distributions and occurrence frequencies in Figure 2a, 2b. We find that high-entropy tokens are rare. These tokens have low sampling probabilities but predominantly represent critical decision points. This creates a paradox: the tokens most valuable for learning are systematically under-sampled during standard rollout.

Temperature Effects on Critic Tokens. To address this, we investigate temperature's role in generating critical tokens. We evaluate rollout across various temperatures. Figure 2c, 2d reveal that while accuracy decreases with temperature, critical token occurrence increases before dropping at extremes. This demonstrates a fundamental trade-off: low temperatures preserve accuracy but suppress critical tokens, while high temperatures generate more meaningful tokens but reduce accuracy.

3.3 Advantage Calculation: Reward Attribution Granularity

Limitations of Sequence-Level Advantage Estimation. Current methods adopt GRPO's sequence-level normalization to estimate advantages—using a single reward to represent an entire sequence for estimation, disregarding length differences across sequences. Figure 3a reveals that negative samples are significantly longer than positive ones. GRPO divides loss by sequence length (Equation 1) to balance contributions but sacrifices individual token gradients. DAPO improves upon this by dividing by total tokens (Equation 2) to preserve token-level gradients—a natural approach for handling token heterogeneity. However, this creates a conflict: sequence-level advantages estimation with token-mean loss causes longer negative samples to dominate gradient, as shown in Figure 3b.

While DAPO's token-mean loss achieves better heterogeneity, sequence-level advantage estimation

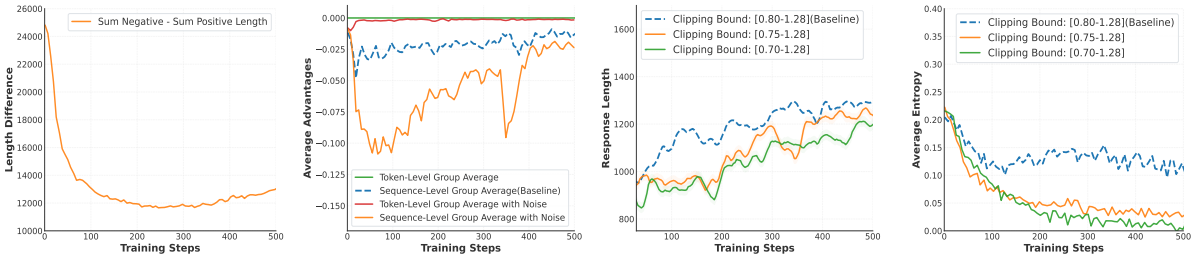

(a) Length Differences (b) Average Advantage (c) Response Length (d) Average Entropy

Figure 3: Bias of token-mean loss toward negative samples affects advantage and clipping.

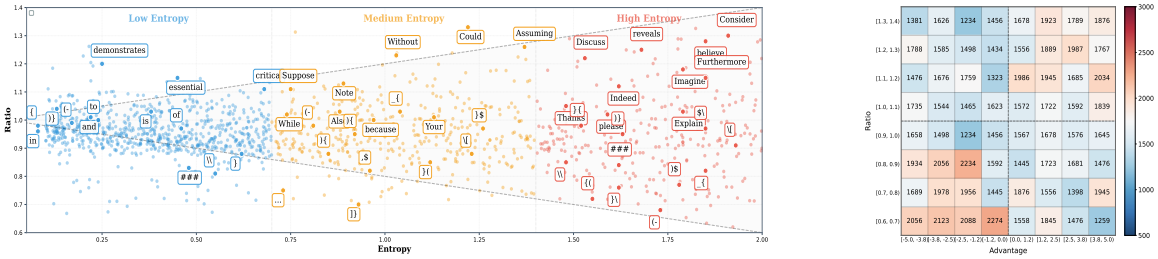

(a) Token Distribution: Entropy-Dependent Ratio Spread (b) Ratio-Advantage Heatmap

Figure 4: Token update patterns analysis via importance ratios

introduce negative length bias that impacts token-level optimization. First, when minor perturbations are applied to advantages, the ignorance of length in sequence-level advantage estimation amplifies these fluctuations into substantial deviations, causing training instability as shown in Figure 3b. Second, this bias conflicts with strategies like Archer, which relaxes left clipping bounds for high-entropy tokens. As shown in Figure 3c and Figure 3d, negative advantage dominance combined with relaxed left clipping bounds causes rapid probability degradation for high-entropy critical tokens, leading to entropy collapse and reduced response lengths.

Analyzing Token Update Patterns. Beyond sequence-level advantage estimation ignoring token heterogeneity, another limitation is that all tokens within a sequence receive same advantages. While EDGE-GRPO modulates advantages based on sequence entropy, they still lack token-level granularity and ignore actual update dynamics. We conduct finer-grained analysis by introducing importance ratios $r_{i,t} = \pi_\theta(a_{i,t}|s_{i,t})/\pi_{\text{ref}}(a_{i,t}|s_{i,t})$, which reflect model's existing optimization state and actual update needs for each token.

Figure 4a shows high-entropy tokens have broader ratio distributions, with most tokens deviating from 1.0 as the model preferentially updates them. By analyzing the distribution of top 20% high-entropy tokens across different advantage and ratio values in Figure 4b, we observe a strong correlation between advantages and ratios. Meanwhile, we find that low-entropy tokens' ratios cluster around 1.0, suggesting no clear update direction. The token examples in Figure 4a demonstrate that tokens with large ratio deviations are those requiring the most optimization—either critical tokens or redundant tokens. Entropy alone fails to capture update dynamics—tokens with similar entropy have vastly different optimization needs.

3.4 Clipping Loss Computation: Clipping Constraint

Analyzing Clipped Token Properties. Figures 5 reveal distinct clipping patterns: low-entropy tokens, mostly format and useless code tokens, are frequently left-clipped, while high-entropy tokens including critical reasoning tokens are frequently right-clipped. Figure 5d shows that even with DAPO's clip higher, high-entropy tokens remain frequently clipped, with frequency increasing with entropy. Existing clipping strategies prematurely exclude these tokens from gradient updates. Uniform clipping constrains both useless low-entropy tokens that need probability reduction and critical high-entropy tokens that require exploration. Methods like Archer that relax clipping bounds for high-entropy tokens while tightening them for low-entropy ones exacerbate this constraint.

4 Heterogeneous Adaptive Policy Optimization

Based on the challenges in Section 3, we propose HAPO that optimizes the RL pipeline according to

(a) Entropy-Clip Patterns (b) Right-Clipped (c) Left-Clipped (d) Critical Token Clipped

Figure 5: Token-Level clipping patterns and word cloud visualizations

token heterogeneity. We introduce our core principles and formalize how entropy serves as a continuous regulator, enabling smooth transitions between different tokens while maintaining simplicity.

4.1 Entropy-Based Continuous Regulation

As we use entropy as a continuous indicator, we first prepare the entropy to obtain a unified regulation signal. Following (Wang et al., 2025b), we normalize entropy within each batch and apply asymmetric scaling to map values to $[-1, 1]$:

$$h_{i,t} = \frac{\log(H_{i,t}) - Q_\rho(\log(H))}{\sigma(\log(H))} \tag{4}$$

$$\tilde{h}_{i,t} = \begin{cases} h_{i,t}/h_{\max} & \text{if } h_{i,t} > 0 \\ -h_{i,t}/|h_{\min}| & \text{if } h_{i,t} \leq 0 \end{cases} \tag{5}$$

where $H_{i,t}$ is the entropy at i, t, Q_ρ denotes the ρ-th quantile serving as the boundary between high and low entropy tokens, and σ is the standard deviation.

We apply logarithmic smoothing to handle the significant variance in entropy distribution. The asymmetric scaling ensures balanced regulation across both positive and negative ranges.

With this normalized entropy $\tilde{h}_{i,t}$, we implement adaptive modulation across multiple components:

$$\text{Parameter}_{i,t} = \text{Parameter}_{\text{base}} \cdot f(\tilde{h}_{i,t}) \tag{6}$$

where $f(\cdot)$ is designed for each specific component. For high-entropy tokens ($\tilde{h}_{i,t} > 0$), we typically amplify exploration and learning signals, while for low-entropy tokens ($\tilde{h}_{i,t} < 0$), we suppress unnecessary updates and maintain stability.

4.2 Adaptive Temperature Sampling

Now that we understand temperature's impact on token generation and accuracy, a balanced strategy is to introduce dynamic temperature. We apply high temperature to high-entropy tokens to encourage exploration while using low temperature for low-entropy ones to maintain coherence. Since rollout proceeds token-by-token without sequence-level context, we maintain the formulation of Equation 6 but dynamically adjust temperature based on current state, computing normalization on-the-fly:

$$T_{i,t} = T_{\text{base}} \cdot \left(1 + \frac{\log(H_{i,t}) - \hat{\rho}_{\log(H)}}{\hat{\sigma}_{\log(H)}} \cdot \tau\right) \tag{7}$$

$\hat{\rho}_{\log(H)}$ is the ρ-th quantile of entropy and serves as the direction of temperature adjustments. $\hat{\sigma}_{\log(H)}$ is the variance, and τ determines the maximum bounds of temperature adjustment. We compute $\hat{\rho}_{\log(H)}$ and $\hat{\sigma}_{\log(H)}$ using previously generated tokens rather than sequence-level normalization.

4.3 Token-Level Group Average Advantage Estimation

As DAPO computes loss through token-mean, a more compatible approach for advantage estimation is token-level averaging. We estimate advantages across all tokens in the group, distributing rewards to individual tokens before normalization:

$$A_{i,t} = \frac{a_{i,t} - \mu_{\text{tok}}}{\sigma_{\text{tok}}} \tag{8}$$

where $a_{i,t} = r_i \in \{0, 1\}$ is the token-level reward inherited from sequence i, $\mu_{\text{tok}} = \frac{1}{|\mathcal{T}|}\sum_{(i,t)\in\mathcal{T}} a_{i,t}$ is the mean across all tokens, $\sigma_{\text{tok}} = \sqrt{\frac{1}{|\mathcal{T}|}\sum_{(i,t)\in\mathcal{T}}(a_{i,t} - \mu_{\text{tok}})^2}$ is the standard deviation across all tokens, $\mathcal{T} = \{(i, t)\}$ represents all tokens across sequences in the group.

The key distinction from sequence-level advantage estimation lies in granularity. Instead of using a single reward per sequence, we assign rewards to individual tokens then involve all tokens in advantage estimation. This ensures $\sum_{(i,t)\in\mathcal{T}} \hat{A}_{i,t} = 0$, resolving the gradient bias in token-mean loss. More importantly, our method preserves length-dependent gradient scaling within reward categories—longer sequences contribute larger gradients for complex reasoning. This simultaneously maintains GRPO's unbiased gradient contributions and DAPO's token-level granularity advantages.

Table 1: Comparison between *HAPO* and existing methods on Math benchmarks. The models below are from 1000 steps of training. The best results are indicated by **boldface**.

Model	Method	Avg (1590)	AIME24(30)	AIME25(30)	AMC(83)	Math(500)	OlympiadBench(675)	Minerva (272)
Qwen2.5-Math-1.5B	Vanilla GRPO	36.83	20.12	15.67	71.23	65.91	29.84	18.23
	Vanilla DAPO	38.34	21.73	17.11	72.85	**67.33**	31.47	19.52
	DAPO w/ Forking Tokens	38.83	22.51	18.56	74.80	65.54	32.81	18.74
	Archer	37.06	19.62	17.24	72.17	64.21	29.98	19.14
	EDGE-GRPO	38.79	22.88	17.86	74.48	66.94	30.67	19.93
	HAPO(ours)	**40.62**	**25.33**	**20.12**	**75.69**	66.83	**33.86**	**21.91**
Qwen2.5-Math-7B	Vanilla GRPO	45.26	35.67	18.91	80.14	74.28	32.76	29.85
	Vanilla DAPO	46.97	37.24	20.24	81.99	76.72	34.35	31.28
	DAPO w/ Forking Tokens	47.43	38.45	21.90	80.55	75.38	35.52	32.79
	Archer	45.63	35.25	19.19	82.53	72.16	33.04	31.61
	EDGE-GRPO	45.89	38.85	21.83	83.23	69.78	32.43	29.24
	HAPO(ours)	**50.04**	**41.31**	**24.34**	**85.47**	**78.45**	**36.94**	**33.73**
Qwen3-8B	Vanilla GRPO	48.15	33.91	23.56	76.84	81.17	42.93	30.47
	Vanilla DAPO	50.00	35.84	25.24	78.37	83.54	44.87	32.13
	DAPO w/ Forking Tokens	50.21	36.22	25.85	80.26	82.75	**46.27**	29.91
	Archer	49.17	34.76	23.33	78.47	82.57	42.69	33.22
	EDGE-GRPO	50.06	35.15	24.12	80.47	83.33	43.35	33.91
	HAPO(ours)	**51.97**	**39.01**	**26.83**	**81.77**	**84.23**	45.26	**34.74**
Qwen3-14B	Vanilla GRPO	55.89	42.36	35.78	83.91	85.63	50.14	37.52
	Vanilla DAPO	58.06	44.82	37.53	85.52	88.25	52.28	39.97
	DAPO w/ Forking Tokens	60.00	46.71	39.12	87.31	88.90	**56.85**	40.84
	Archer	58.88	46.28	38.68	86.73	87.76	53.19	40.62
	EDGE-GRPO	59.16	46.94	39.01	87.15	89.12	52.45	40.31
	HAPO(ours)	**62.09**	**49.74**	**41.71**	**89.66**	**92.47**	56.43	**42.53**

Table 2: Comparison between *HAPO* and existing methods on LLaMA series models.

Model	Method	Avg (1590)	AIME24(30)	AIME25(30)	AMC(83)	Math(500)	OlympiadBench(675)	Minerva (272)
LLaMA3.2-3B Instruct	Vanilla GRPO	20.97	11.67	0.00	49.28	38.62	14.35	11.94
	Vanilla DAPO	22.99	13.93	0.00	52.11	41.95	16.63	13.32
	DAPO w/ Forking Tokens	23.86	12.45	0.00	55.96	43.34	16.89	14.54
	HAPO(ours)	**27.42**	**17.50**	**8.00**	**58.74**	**44.86**	**19.37**	**16.05**
LLaMA3.1-8B-Instruct	Vanilla GRPO	22.07	10.83	0.00	54.16	40.78	13.81	12.89
	Vanilla DAPO	24.21	13.33	0.00	57.11	44.05	16.03	14.76
	DAPO w/ Forking Tokens	25.30	14.50	0.00	58.46	46.55	15.89	16.39
	HAPO(ours)	**27.06**	**19.17**	**1.25**	**59.32**	**47.83**	**17.66**	**17.11**

4.4 Differential Advantage Redistribution

Our core idea remains amplifying advantages for high-entropy tokens while reducing them for low-entropy tokens. Different from existing methods, we additionally use importance ratios for regulation, ensuring updates based on each token's actual conditions to achieve token-level control.

To implement this, we define a neutral zone $[\gamma_L, \gamma_U]$, typically centered around ratio = 1. High-entropy tokens receive amplification only when their importance ratio falls outside the neutral zone, reflecting clear update trends. We enhance them accordingly. For low-entropy tokens, we apply suppression within the neutral zone to reduce the influence of tokens lacking update tendency. This aligns with our analysis, where ratios far from 1 typically represent relatively important tokens, and we assign them relatively larger advantages:

$$\hat{A}_{i,t} = \begin{cases} A_{i,t} \cdot (1 + \tilde{h}_{i,t}) & \text{if } C(\tilde{h}_{i,t}, r_{i,t}) \\ A_{i,t} & \text{otherwise} \end{cases} \tag{9}$$

$$C(\tilde{h}_{i,t}, r_{i,t}) = \begin{cases} r_{i,t} \notin [\gamma_L, \gamma_U] & \text{if } \tilde{h}_t > 0 \\ r_{i,t} \in [\gamma_L, \gamma_U] & \text{if } \tilde{h}_t \leq 0 \end{cases} \tag{10}$$

4.5 Asymmetric Adaptive Clipping

Unlike Archer (Wang et al., 2025a) that uniformly adjusts clipping bounds, we apply asymmetric adjustments respecting natural update tendencies of each token. By lowering the left boundary for low-entropy tokens, we allow these tokens to continue participating in the update process and are more likely to decrease substantially. By raising the right boundary for high-entropy tokens, we encourage exploration at critical decision points.

$$\epsilon_L(i,t) = \begin{cases} \epsilon_L^{\text{base}}(1 - \tilde{h}_{i,t}) & \text{if } \tilde{h}_{i,t} \leq 0 \\ \epsilon_L^{\text{base}} & \text{if } \tilde{h}_{i,t} > 0 \end{cases} \tag{11}$$

$$\epsilon_R(i,t) = \begin{cases} \epsilon_R^{\text{base}} & \text{if } \tilde{h}_{i,t} \leq 0 \\ \epsilon_R^{\text{base}}(1 + \tilde{h}_{i,t}) & \text{if } \tilde{h}_{i,t} > 0 \end{cases} \tag{12}$$

where $\epsilon_L^{\text{base}}, \epsilon_R^{\text{base}}$ are the base clipping bounds.

These continuous formulations unify our heterogeneous strategy, with all modulations smoothly varying as functions of normalized entropy $\tilde{h}$, ensuring a concise implementation. As entropy is already computed during sampling, HAPO introduces virtually no additional computational over-

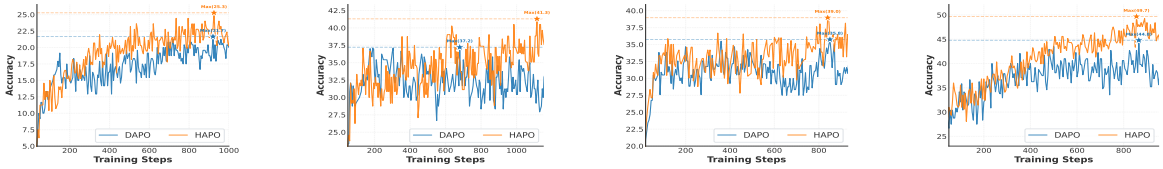

(a) Qwen2.5-Math-1.5B (b) Qwen2.5-Math-7B (c) Qwen3-8B (d) Qwen3-14B

Figure 6: Training dynamics of AIME24 results for HAPO and DAPO.

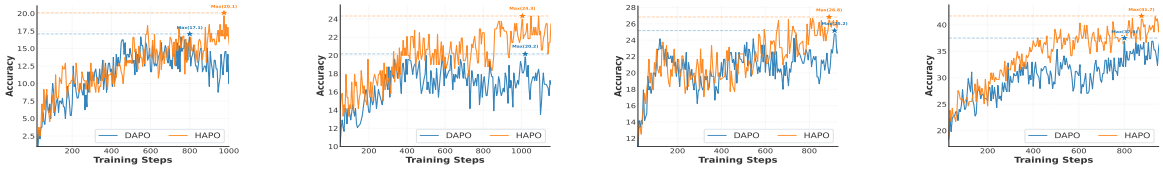

(a) Qwen2.5-Math-1.5B (b) Qwen2.5-Math-7B (c) Qwen3-8B (d) Qwen3-14B

Figure 7: Training dynamics of AIME25 results for HAPO and DAPO.

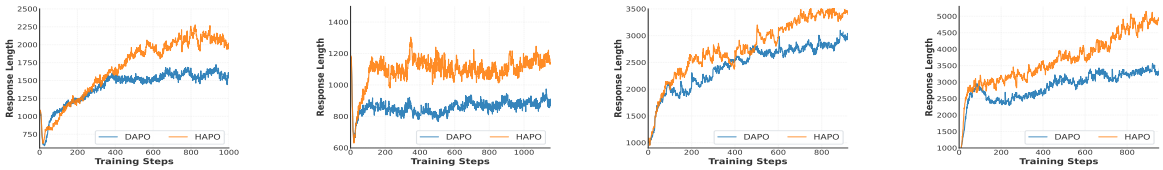

(a) Qwen2.5-Math-1.5B (b) Qwen2.5-Math-7B (c) Qwen3-8B (d) Qwen3-14B

Figure 8: Training dynamics of response length for HAPO and DAPO.

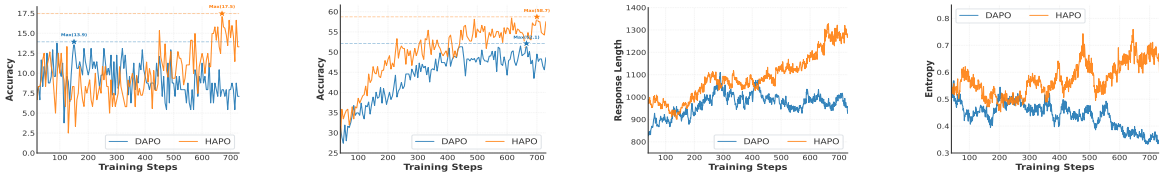

(a) AIME24 (b) AMC (c) Response Length (d) Entropy

Figure 9: Training dynamics of HAPO and DAPO on LLaMA3.2-3B-Instruct.

head. We present the overall framework of HAPO along with its pseudocode in Appendix B.

5 Experiments

Experimental Setup. We incorporate our token-level strategy into DAPO (Yu et al., 2025) within verl (Sheng et al., 2025). We maintain DAPO's settings: clip-higher of $\epsilon_{\text{high}} = 0.28$ and $\epsilon_{\text{low}} = 0.2$; overlong reward shaping with 10240-token maximum length and 4096-token cache. We use a training batch size of 512 and a mini-batch size of 32. Training proceeds with a learning rate of 10^{-6}.

For HAPO's token-level strategies, we follow (Wang et al., 2025b) and set the quantile ρ to 80%, encouraging exploration to the top 20% highest-entropy tokens. For Adaptive Temperature Sampling, we compute $\hat{\rho}_{\log(H)}$ and $\hat{\sigma}_{\log(H)}$ based on all tokens' entropy in the previous step. We set $T_{\text{base}} = 1.0$ and $\tau = 0.1$. For Differential Advantage Redistribution, we determine the neutral zone using the clipping ratios that correspond to each token's actual dynamics, setting it to $[1 - \frac{\epsilon_L}{2}, 1 + \frac{\epsilon_R}{2}]$. For Asymmetric Adaptive Clipping, we set $\epsilon_L^{\text{base}} = 0.2$ and $\epsilon_R^{\text{base}} = 0.28$. We provide hyperparameter ablations in Appendix C. Since entropy heterogeneity is prevalent across diverse models and tasks, our parameters exhibit strong generalizability.

We perform experiments across Qwen2.5-Math-1.5B,Qwen2.5-Math-7B (Yang et al., 2024) base models and Qwen3-8B,Qwen3-14B (Yang et al., 2025) base models, using DAPO-Math-17K (Yu et al., 2025) as the training dataset. We trained all models on 4 nodes with 32 A100 GPUs.

Evaluation. We evaluate all the models on mathematical benchmarks: AIME24, AIME25, AMC23, Minerva (Lewkowycz et al., 2022), MATH500 (Hendrycks et al., 2021), and OlympiadBench (He et al., 2024). For each question, we generate 8 responses under a decoding temperature $T = 0.5$, and report the average accuracy.

Main Results. We compare HAPO against vanilla GRPO, vanilla DAPO, DAPO with Forking Tokens, Archer, and EDGE-GRPO in Ta-

Table 3: Comparison between *HAPO* and existing methods on Code and Logic benchmarks.

Model	Method	Logic RL					LiveCodeBench		
		Avg	4ppl	5ppl	6ppl	7ppl	**Avg**	V5(24/8/1-25/2/1)	V6(25/2/1-25/5/1)
Qwen2.5-7B-Instruct-1M	Vanilla DAPO	87.13	92.23	90.16	85.72	80.42	29.08	28.22	29.94
	DAPO w/ Forking Tokens	90.15	93.44	92.13	90.26	84.75	31.46	29.72	33.19
	Archer	87.68	94.12	89.75	85.28	81.55	29.75	28.63	30.86
	EDGE-GRPO	88.92	92.35	91.13	89.58	82.63	30.40	29.27	31.53
	HAPO(ours)	**92.69**	**95.41**	**93.34**	**93.26**	**88.74**	**34.02**	**32.55**	**35.48**

Table 4: Component ablation of HAPO. We denote Adaptive Temperature Sampling, Token-Level Group Average Advantage Estimation, Differential Advantage Redistribution, and Asymmetric Adaptive Clipping as A, B, C, and D respectively for convenience.

A	**B**	**C**	**D**	**Avg (1560)**	AIME24 (30)	AIME25 (30)	AMC (83)	Math (500)	OlympiadBench(675)	Minerva (272)
				46.97	37.24	20.24	81.99	76.72	34.35	31.28
✓				48.85	39.77	23.01	84.20	76.88	**37.34**	31.92
	✓			48.56	39.04	23.77	83.59	77.19	35.02	32.76
		✓		48.28	38.63	22.07	84.19	76.58	36.00	32.19
			✓	48.02	38.38	20.73	82.45	78.56	36.49	31.51
		✓	✓	48.74	39.19	22.55	84.43	77.72	36.13	32.44
	✓	✓	✓	49.42	40.17	23.62	84.41	**79.10**	35.96	33.28
✓	✓	✓	✓	**50.04**	**41.31**	**24.34**	**85.47**	78.45	36.94	**33.73**

ble 1. On Qwen2.5-Math-7B, HAPO surpasses DAPO w/ Forking Tokens by 2.86 and 2.44 points on AIME'24 and AIME'25. When compared to entropy-based approaches, HAPO outperforms Archer by 2.80-4.41 points and EDGE-GRPO by 1.83-4.15 points, highlighting the effectiveness of our fine-grained heterogeneous treatment. Overall, HAPO outperforms these methods across models of different scales, demonstrating strong scalability. We visualize training dynamics in Figures 6, 7, 8, comparing HAPO and DAPO throughout the training process. Our method maintains longer response lengths and higher accuracy, indicating that HAPO successfully preserves model exploration capabilities while achieving better task performance.

6 Ablations

Results on Models Other Than Qwen. We employ same experiments on LLaMA models. We omit EDGE-GRPO and Archer since they yield no performance gains. Results are presented in Table 2. HAPO demonstrates more pronounced improvements on weaker models, suggesting that our fine-grained heterogeneous treatment effectively helps models capture critical information, further validating HAPO's scalability. In Figure 9, we show training dynamics of LLaMA3.2-3B. HAPO exhibits higher entropy and longer responses, enabling the model to maintain continuous exploration.

Generalization to Other Domains. We validate HAPO's generalizability in code and logic. We use code data from Archer (Wang et al., 2025a) and logic data from Logic-RL (Xie et al., 2025). We employ Qwen2.5-7B-Instruct-1M due to its general capabilities and the fact that it has not been pre-trained on these datasets. Evaluation on LiveCodeBench (Jain et al., 2024) and Logic in Table 3 shows that HAPO consistently outperforms DAPO w/ Forking Tokens. This demonstrates that token heterogeneity exists across different domains, and HAPO's approach naturally generalizes to diverse areas, confirming HAPO's broad applicability.

Component-wise Contributions. We examine each component's contribution in Table 4. All components contribute to HAPO's performance, with Adaptive Temperature Sampling proving particularly crucial as it governs token distribution and entropy. Token-Level Group Average Advantage Estimation significantly impacts advantage redistribution and clipping, corroborating our analysis. In fact, all components form an organic unity: Adaptive Temperature Sampling ensures sufficient high-entropy critical tokens, Differential Advantage Redistribution optimizes them with larger updates, and Asymmetric Adaptive Clipping encourages exploration. This synergy enables effective leverage of token heterogeneity throughout optimization.

7 Conclusion

We introduces Heterogeneous Adaptive Policy Optimization (HAPO), which addresses the limitation of uniform token treatment. HAPO implements token-level optimization throughout training, using entropy as a continuous signal. Experiments demonstrate consistent improvements over DAPO. HAPO establishes that effective RLHF requires adapting to token heterogeneity, opening new ways for developing more capable reasoning models.

Limitations

Despite HAPO's strong empirical performance, we acknowledge several limitations that provide opportunities for future refinement:

Entropy as Primary Heterogeneity Signal. While our entropy-based approach effectively identifies token heterogeneity, it may not capture all aspects of token importance. Tokens with moderate entropy could still be semantically crucial, particularly in structured or deterministic outputs. Future work could explore complementary signals like gradient magnitudes or attention weights.

Static Entropy Thresholds. HAPO uses a fixed 80th percentile to distinguish high/low entropy tokens throughout training, though entropy distributions evolve as training progresses. Dynamic threshold adaptation could potentially improve performance but would add complexity.

These limitations suggest directions for future refinement.

References

Shivam Agarwal, Zimin Zhang, Lifan Yuan, Jiawei Han, and Hao Peng. 2025. The unreasonable effectiveness of entropy minimization in llm reasoning. *Preprint*, arXiv:2505.15134.

Anthropic. 2024. Claude 3.5 sonnet.

Ganqu Cui, Yuchen Zhang, Jiacheng Chen, Lifan Yuan, Zhi Wang, Yuxin Zuo, Haozhan Li, Yuchen Fan, Huayu Chen, Weize Chen, Zhiyuan Liu, Hao Peng, Lei Bai, Wanli Ouyang, Yu Cheng, Bowen Zhou, and Ning Ding. 2025. The entropy mechanism of reinforcement learning for reasoning language models. *Preprint*, arXiv:2505.22617.

DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, and 181 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *Preprint*, arXiv:2501.12948.

Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. 2024. Kto: Model alignment as prospect theoretic optimization. *Preprint*, arXiv:2402.01306.

Zitian Gao, Lynx Chen, Joey Zhou, and Bryan Dai. 2025. One-shot entropy minimization. *Preprint*, arXiv:2505.20282.

Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu, and Maosong Sun. 2024. Olympiadbench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems. *Preprint*, arXiv:2402.14008.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset. *Preprint*, arXiv:2103.03874.

Jian Hu, Jason Klein Liu, Haotian Xu, and Wei Shen. 2025. Reinforce++: An efficient rlhf algorithm with robustness to both prompt and reward models. *Preprint*, arXiv:2501.03262.

Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. 2024. Livecodebench: Holistic and contamination free evaluation of large language models for code. *Preprint*, arXiv:2403.07974.

Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, Yuhuai Wu, Behnam Neyshabur, Guy Gur-Ari, and Vedant Misra. 2022. Solving quantitative reasoning problems with language models. *Preprint*, arXiv:2206.14858.

Zeju Li, Jianyuan Zhong, Ziyang Zheng, Xiangyu Wen, Zhijian Xu, Yingying Cheng, Fan Zhang, and Qiang Xu. 2025. Compressing chain-of-thought in llms via step entropy. *Preprint*, arXiv:2508.03346.

Jia Liu, ChangYi He, YingQiao Lin, MingMin Yang, FeiYang Shen, and ShaoGuo Liu. 2025. Ettrl: Balancing exploration and exploitation in llm test-time reinforcement learning via entropy mechanism. *Preprint*, arXiv:2508.11356.

Yu Meng, Mengzhou Xia, and Danqi Chen. 2024. Simpo: Simple preference optimization with a reference-free reward. *Preprint*, arXiv:2405.14734.

OpenAI. 2024. Learning to reason with llms.

Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. *Preprint*, arXiv:2203.02155.

Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. 2024. Direct preference optimization: Your language model is secretly a reward model. *Preprint*, arXiv:2305.18290.

John Schulman, Sergey Levine, Philipp Moritz, Michael I. Jordan, and Pieter Abbeel. 2017a. Trust region policy optimization. *Preprint*, arXiv:1502.05477.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017b. Proximal policy optimization algorithms. *Preprint*, arXiv:1707.06347.

ByteDance Seed, :, Jiaze Chen, Tiantian Fan, Xin Liu, Lingjun Liu, Zhiqi Lin, Mingxuan Wang, Chengyi Wang, Xiangpeng Wei, Wenyuan Xu, Yufeng Yuan, Yu Yue, Lin Yan, Qiying Yu, Xiaochen Zuo, Chi Zhang, Ruofei Zhu, Zhecheng An, and 255 others. 2025. Seed1.5-thinking: Advancing superb reasoning models with reinforcement learning. *Preprint*, arXiv:2504.13914.

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.

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*, page 1279–1297. ACM.

Jiakang Wang, Runze Liu, Fuzheng Zhang, Xiu Li, and Guorui Zhou. 2025a. Stabilizing knowledge, promoting reasoning: Dual-token constraints for rlvr. *Preprint*, arXiv:2507.15778.

Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xionghui Chen, Jianxin Yang, Zhenru Zhang, Yuqiong Liu, An Yang, Andrew Zhao, Yang Yue, Shiji Song, Bowen Yu, Gao Huang, and Junyang Lin. 2025b. Beyond the 80/20 rule: High-entropy minority tokens drive effective reinforcement learning for llm reasoning. *Preprint*, arXiv:2506.01939.

Tian Xie, Zitian Gao, Qingnan Ren, Haoming Luo, Yuqian Hong, Bryan Dai, Joey Zhou, Kai Qiu, Zhirong Wu, and Chong Luo. 2025. Logic-rl: Unleashing llm reasoning with rule-based reinforcement learning. *Preprint*, arXiv:2502.14768.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. Qwen3 technical report. *Preprint*, arXiv:2505.09388.

An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jianhong Tu, Jingren Zhou, Junyang Lin, Keming Lu, Mingfeng Xue, Runji Lin, Tianyu Liu, Xingzhang Ren, and Zhenru Zhang. 2024. Qwen2.5-math technical report: Toward mathematical expert model via self-improvement. *Preprint*, arXiv:2409.12122.

Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, and 16 others. 2025. Dapo: An open-source llm reinforcement learning system at scale. *Preprint*, arXiv:2503.14476.

Yu Yue, Yufeng Yuan, Qiying Yu, Xiaochen Zuo, Ruofei Zhu, Wenyuan Xu, Jiaze Chen, Chengyi Wang, TianTian Fan, Zhengyin Du, Xiangpeng Wei, Xiangyu Yu, Gaohong Liu, Juncai Liu, Lingjun Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Chi Zhang, and 8 others. 2025. Vapo: Efficient and reliable reinforcement learning for advanced reasoning tasks. *Preprint*, arXiv:2504.05118.

Weihao Zeng, Yuzhen Huang, Qian Liu, Wei Liu, Keqing He, Zejun Ma, and Junxian He. 2025. Simplerl-zoo: Investigating and taming zero reinforcement learning for open base models in the wild. *Preprint*, arXiv:2503.18892.

Xingjian Zhang, Siwei Wen, Wenjun Wu, and Lei Huang. 2025. Edge-grpo: Entropy-driven grpo with guided error correction for advantage diversity. *Preprint*, arXiv:2507.21848.

Appendix

A Related Work

A.1 Reinforcement Learning from Human Feedback

The evolution of reinforcement learning in language models represents a trajectory from basic preference alignment toward sophisticated reasoning capabilities. Foundational work in constrained optimization emerged through TRPO (Schulman et al., 2017a) and PPO (Schulman et al., 2017b), establishing principles for stable policy updates. Subsequent algorithmic innovations eliminated computational bottlenecks—GRPO (Shao et al., 2024) and REINFORCE++ (Hu et al., 2025) removed value network dependencies via group-based advantage computation. The offline optimization paradigm gained prominence through DPO (Rafailov et al., 2024), KTO (Ethayarajh et al., 2024), and SimPO (Meng et al., 2024), which circumvent reward model training entirely.

A fundamental transformation occurred with the advent of reasoning-centric models. OpenAI's o1 (OpenAI, 2024) pioneered effective multi-step reasoning through reinforcement learning at scale, catalyzing widespread development. DeepSeek-R1 (DeepSeek-AI et al., 2025) demonstrated reasoning emergence without supervised fine-tuning, while Claude3.5 (Anthropic, 2024), Qwen3 (Yang et al., 2025), and Seed-1.5-Thinking (Seed et al., 2025) expanded the reasoning frontier. Beyond these state-of-the-art models, various works explored complementary optimization strategies, including SimpleRLZoo (Zeng et al., 2025) and Open-Reasoner-Zero, which investigate alternative training paradigms. Meanwhile, algorithmic refinements continue to enhance existing methods—DAPO (Yu et al., 2025) introduced token-mean averaging and asymmetric clipping for extended sequences, while VAPO (Yue et al., 2025) developed variance-aware optimization strategies to improve training stability and convergence.

A.2 Entropy-Based Optimization in LLMs

Entropy serves as a crucial signal for model uncertainty and optimization guidance. (Cui et al., 2025) explored how entropy dynamics affect learning in reasoning models. Entropy-based optimization strategies have diversified significantly. (Agarwal et al., 2025; Gao et al., 2025) showed entropy minimization alone improves performance without labeled data. Step entropy methods (Li et al., 2025) identify superfluous reasoning steps by aggregating token-level entropy.

Recent studies have leveraged entropy as a key metric for identifying and treating critical tokens in reasoning tasks. (Wang et al., 2025b) established the "80/20 rule," demonstrating that 20% of high-entropy "forking tokens" fundamentally control reasoning diversity. Building upon token identification, subsequent works have developed entropy-based differential treatment strategies. Archer (Wang et al., 2025a) employs entropy to determine token-specific clipping boundaries, encouraging exploration for high-entropy tokens while maintaining stability for low-entropy ones. Similarly, ETTRL (Liu et al., 2025) and EDGE-GRPO (Wang et al., 2025a) modulate advantages based on token entropy, amplifying updates for high-entropy tokens and dampening those for low-entropy tokens to achieve an optimal exploration-exploitation balance.

B HAPO Overall Framework and Pseudocode

We present the complete implementation details of HAPO in this section. Figure 10 illustrates the overall framework of our method, showing how entropy-based continuous regulation is integrated across different components of the RL pipeline. The framework demonstrates the flow from rollout generation with adaptive temperature sampling, through token-level advantage estimation and redistribution, to the final policy update with asymmetric clipping.

Figure 12 provides the complete pseudocode of HAPO based on the verl implementation. The algorithm details the step-by-step process of: (1) computing entropy statistics for continuous regulation, (2) applying adaptive temperature during sampling, (3) performing token-level group average advantage estimation, (4) implementing differential advantage redistribution based on importance ratios and entropy values, and (5) executing asymmetric adaptive clipping for policy updates.

The pseudocode highlights HAPO's key innovation: using normalized entropy $\tilde{h}_{i,t} \in [-1, 1]$ as a continuous regulator that smoothly modulates all components without introducing significant computational overhead. The entropy computation required for sampling is reused throughout the pipeline, ensuring efficient implementation. The al-

Table 5: Ablation study on entropy quantile threshold ρ using Qwen2.5-Math-7B. The 80th percentile achieves optimal balance.

Quantile	**Avg (1560)**	AIME24 (30)	AIME25 (30)	AMC (83)	Math (500)	OlympiadBench(675)	Minerva (272)
90%	48.96	39.44	22.73	84.11	77.74	36.61	33.14
80%	**50.04**	**41.31**	**24.34**	**85.47**	78.45	36.94	**33.73**
70%	49.74	40.78	24.11	84.96	**78.83**	**37.05**	32.71
50%	48.45	38.92	23.67	82.41	76.98	36.42	32.35
30%	47.10	37.05	23.12	81.58	75.21	34.39	31.26

Table 6: Ablation study on temperature adjustment factor τ using Qwen2.5-Math-7B. Moderate adjustment ($\tau = 0.1$) yields best results.

τ	**Avg (1560)**	AIME24 (30)	AIME25 (30)	AMC (83)	Math (500)	OlympiadBench(675)	Minerva (272)
0.0	49.42	40.17	23.62	84.41	**79.10**	35.96	33.28
0.05	49.58	40.86	24.05	83.98	78.74	36.52	33.31
0.1	**50.04**	**41.31**	**24.34**	**85.47**	78.45	**36.94**	**33.73**
0.3	49.11	40.67	23.61	83.13	77.58	36.89	32.36
0.5	48.12	38.92	22.45	83.27	76.83	35.14	32.08

Table 7: Ablation study on neutral zone configuration using Qwen2.5-Math-7B. Using half of the clipping bounds provides optimal control.

Neutral Zone	**Avg (1560)**	AIME24 (30)	AIME25 (30)	AMC (83)	Math (500)	OlympiadBench(675)	Minerva (272)
Full Clip Bounds	49.57	40.77	24.13	84.86	78.23	36.48	32.93
1/2 Clip Bounds	**50.04**	41.31	**24.34**	**85.47**	78.45	**36.94**	**33.73**
Zero Bounds	49.71	**41.42**	23.51	85.23	**78.62**	36.17	33.32
No Zone	49.37	40.83	23.67	84.58	77.94	36.56	32.61

gorithm maintains simplicity while achieving fine-grained heterogeneous treatment, making it readily applicable to existing RL frameworks with minimal modifications.

C Additional Hyperparameter Ablations

We conduct extensive ablation studies to validate our hyperparameter choices and design decisions. All experiments in this section are performed on Qwen2.5-Math-7B following the same training protocol described in Section 5.

Entropy Quantile Threshold. We investigate the impact of the entropy quantile threshold ρ, which determines the boundary between high and low entropy tokens. Table 5 presents results across different quantile values on Qwen2.5-Math-7B.

The results confirm that the 80th percentile provides optimal performance. The 70th percentile achieves comparable results, suggesting robustness around this range. Setting ρ to 90% degrades performance as this restrictive threshold identifies too few high-entropy tokens (only top 10%), limiting exploration at critical points. Conversely, lower thresholds (50% and 30%) progressively worsen performance by classifying too many tokens as high-entropy, diluting the focused amplification effect. The 30% threshold particularly suffers, treating most tokens as high-entropy and applying unfocused modifications. This validates the 80th percentile as balancing critical token identification with sufficient coverage.

Temperature Adjustment Factor. The temperature adjustment factor τ controls the magnitude of dynamic temperature modulation. We evaluate its impact across a range of values in Table 6.

Our chosen value of $\tau = 0.1$ achieves the best overall performance. While $\tau = 0.05$ produces competitive results, this smaller adjustment factor limits the temperature modulation range too narrowly, preventing the mechanism from fully leveraging adaptive temperature benefits. The complete absence of temperature adjustment ($\tau = 0$) eliminates this exploration mechanism entirely. Conversely, larger values ($\tau \geq 0.3$ and $\tau \geq 0.5$) lead to excessive temperature modifications. Since high-entropy tokens inherently exhibit strong randomness, applying overly aggressive temperature increases to these already uncertain positions introduces extreme instability, disrupting coherent reasoning chains. This demonstrates the importance

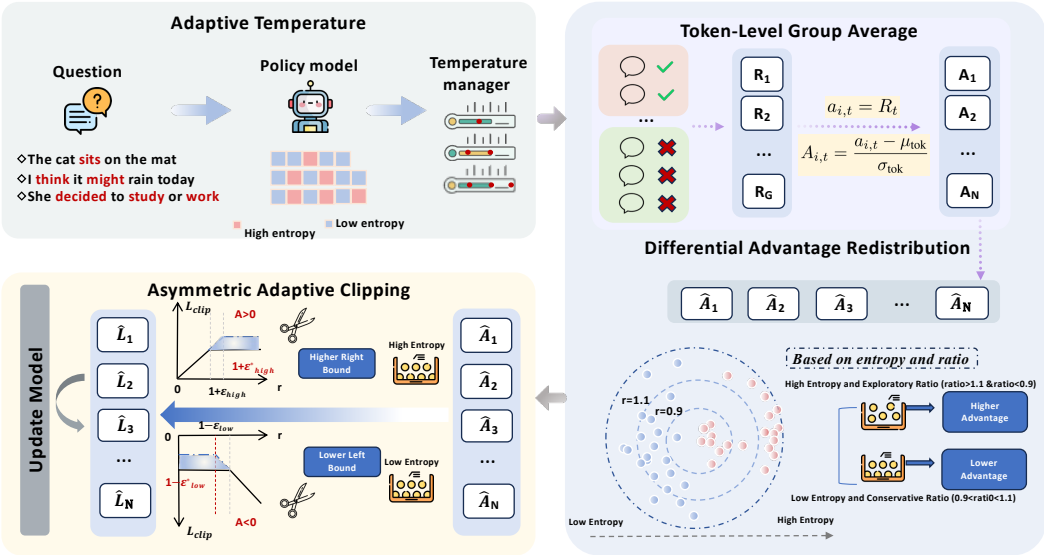

Figure 10: HAPO overall framework: four components address token heterogeneity. Adaptive Temperature Sampling adjusts temperature by token entropy. Token-Level Group Average Advantage Estimation normalizes advantages across tokens. Differential Advantage Redistribution modulates advantages using entropy and importance ratios. Asymmetric Adaptive Clipping applies entropy-conditioned boundaries for targeted noise suppression and exploration.

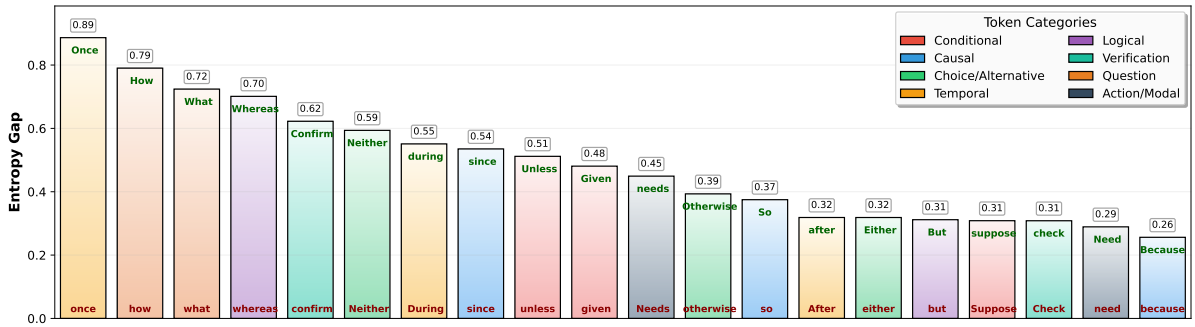

Figure 11: Dual-Entropy Token Frequency-Entropy Landscape

of balanced temperature modulation—sufficient to encourage exploration at critical points without transforming productive exploration into chaotic sampling.

Neutral Zone Configuration. The neutral zone $[\gamma_L, \gamma_U]$ determines when differential advantage redistribution is applied based on importance ratios. We examine various configurations based on the clipping bounds in Table 7.

Setting the neutral zone to half of the clipping bounds achieves optimal performance. The No Zone configuration, which applies no differential advantage redistribution, performs worst—without entropy-based modulation, the model treats all tokens equally regardless of their importance. Full Clip Bounds, which only suppresses low-entropy token advantages without amplifying high-entropy ones, shows marginal improvement but fails to encourage exploration at critical decision points. Zero Bounds performs better by only amplifying high-entropy token advantages. While this promotes exploration, it neglects to suppress trivial low-entropy tokens, leading to inefficient updates. **1/2 Clip Bounds** configuration achieves the best results by combining both mechanisms: it suppresses low-entropy tokens within the neutral zone while amplifying high-entropy tokens whose importance ratios indicate meaningful updates. This balanced approach ensures that computational resources focus on genuinely important tokens while reducing noise from trivial updates, demonstrating the necessity of bidirectional entropy-aware modulation.

These comprehensive ablation studies on Qwen2.5-Math-7B confirm that our hyperparam-

Require: Policy π_θ, dataset $\mathcal{D}$
Require: Entropy quantile: ρ
Require: Temperature params: $T_{\text{base}}, \hat{\rho}_{\log(H)}, \hat{\sigma}_{\log(H)}, \tau$ Clipping params: base bounds $\epsilon_L^{\text{base}}, \epsilon_R^{\text{base}}$
Ensure: Updated policy $\pi_{\theta'}$

▷ Sampling Rollout Sequences

for each token position i, t **do** ▷ Continuous Adaptive Temperature Sampling
  Compute entropy: $H_{i,t} \leftarrow -\sum_v \pi_\theta(v|s_{i,t}) \log \pi_\theta(v|s_{i,t})$
  Adaptive temperature: $T_{i,t} \leftarrow T_{\text{base}} \cdot \left(1 + \frac{\log(H_{i,t}) - \hat{\rho}_{\log(H)}}{\hat{\sigma}_{\log(H)}} \cdot \tau\right)$
end for
Assign rewards to tokens: $a_{i,t} \leftarrow r_i, \mathcal{T} \leftarrow \{(i,t)\}$ ▷ Token-Level Group Average Advantage Estimation
Compute token-level statistics: $\mu_{\text{tok}} \leftarrow \frac{1}{|\mathcal{T}|} \sum_{(i,t)} a_{i,t}, \sigma_{\text{tok}} \leftarrow \sqrt{\frac{1}{|\mathcal{T}|} \sum_{(i,t)} (a_{i,t} - \mu_{\text{tok}})^2}$
Normalize advantages: $A_{i,t} \leftarrow (a_{i,t} - \mu_{\text{tok}})/\sigma_{\text{tok}}$
$Q_\rho(\log(H)) \leftarrow \text{Quantile}_\rho(\{\log(H_{i,t}) : (i,t) \in \mathcal{T}\})$ ▷ Compute global entropy statistics
$\sigma(\log(H)) \leftarrow \sqrt{\frac{1}{|\mathcal{T}|} \sum_{(i,t)\in\mathcal{T}} (\log(H_{i,t}) - Q_\rho(\log(H)))^2}$
for each mini batch $\mathcal{B}$ **do** ▷ Training with Heterogeneous Treatment
  for each token i, t in $\mathcal{B}$ **do**
    Compute log-entropies: $h_{i,t} = \frac{\log(H_{i,t}) - Q_\rho(\log(H))}{\sigma(\log(H))}$
    Asymmetric scaling: $\tilde{h}_{i,t} \leftarrow \begin{cases} h_{i,t}/h_{\max} & \text{if } h_{i,t} > 0 \\ -h_{i,t}/|h_{\min}| & \text{if } h_{i,t} \leq 0 \end{cases}$
    ▷ Continuous Asymmetric Adaptive Clipping - Compute First
    if $\tilde{h}_{i,t} \leq 0$ **then**
      $\epsilon_L(i,t) \leftarrow \epsilon_L^{\text{base}}(1 - \tilde{h}_{i,t}), \epsilon_R(i,t) \leftarrow \epsilon_R^{\text{base}}$
    else
      $\epsilon_L(i,t) \leftarrow \epsilon_L^{\text{base}}, \epsilon_R(i,t) \leftarrow \epsilon_R^{\text{base}}(1 + \tilde{h}_{i,t})$
    end if
    $\gamma_L \leftarrow 1 - \frac{\epsilon_L(i,t)}{2}, \gamma_U \leftarrow 1 + \frac{\epsilon_R(i,t)}{2}$ ▷ Neutral Zone based on Dynamic Clipping
    ▷ Continuous Differential Advantage Redistribution
    Compute importance ratio: $r_{i,t} \leftarrow \pi_\theta(a_{i,t}|s_{i,t})/\pi_{\theta_{\text{old}}}(a_{i,t}|s_{i,t})$
    if $(\tilde{h}_{i,t} > 0$ and $r_{i,t} \notin [\gamma_L, \gamma_U])$ or $(\tilde{h}_{i,t} \leq 0$ and $r_{i,t} \in [\gamma_L, \gamma_U]))$ **then**
      $\hat{A}_{i,t} \leftarrow A_{i,t} \cdot (1 + \tilde{h}_{i,t})$
    else
      $\hat{A}_{i,t} \leftarrow A_{i,t}$
    end if
  end for
  ▷ Compute HAPO loss
  $\mathcal{L}^{\text{HAPO}}(\theta) = \left[\frac{1}{\sum_{i=1}^G |o_i|} \sum_{i=1}^G \sum_{t=1}^{|o_i|} \min\left(r_{i,t}(\theta)\hat{A}_{i,t}, \text{clip}(r_{i,t}(\theta), 1 - \epsilon_L(i,t), 1 + \epsilon_R(i,t))\hat{A}_{i,t}\right)\right]$
  Update parameters: $\theta' \leftarrow \theta + \eta \cdot \nabla_\theta \mathcal{L}^{\text{HAPO}}(\theta)$
end for
$\hat{\rho}_{\log(H)} \leftarrow Q_\rho(\log(H_t)), \hat{\sigma}_{\log(H)} \leftarrow \sigma(\log(H_t))$ ▷ Update entropy statistics for next step

Figure 12: HAPO algorithm overview. The method dynamically adjusts temperature, advantages, and clipping bounds based on normalized entropy $\tilde{h}_{i,t} \in [-1, 1]$, implementing heterogeneous treatment for tokens throughout the RL pipeline.

eter choices are well-calibrated and that each design decision contributes meaningfully to HAPO's strong performance. The continuous, entropy-based regulation framework proves essential for effectively leveraging token heterogeneity in reinforcement learning optimization.

D Example of Dual-Entropy Phenomenon

We show the 20 most frequent dual-entropy token pairs in Figure 11. These dual-entropy tokens are all critical tokens that determine reasoning paths, yet exhibit vastly different entropy values. We present concrete examples of the dual entropy phenomenon in Figure 13, 14, 15, 16, 17, 18. We use blue and red to represent the low-entropy and high-entropy parts of dual-entropy, respectively. These visualizations demonstrate how tokens with identical semantic meanings but different surface forms exhibit drastically different entropy values during model generation.

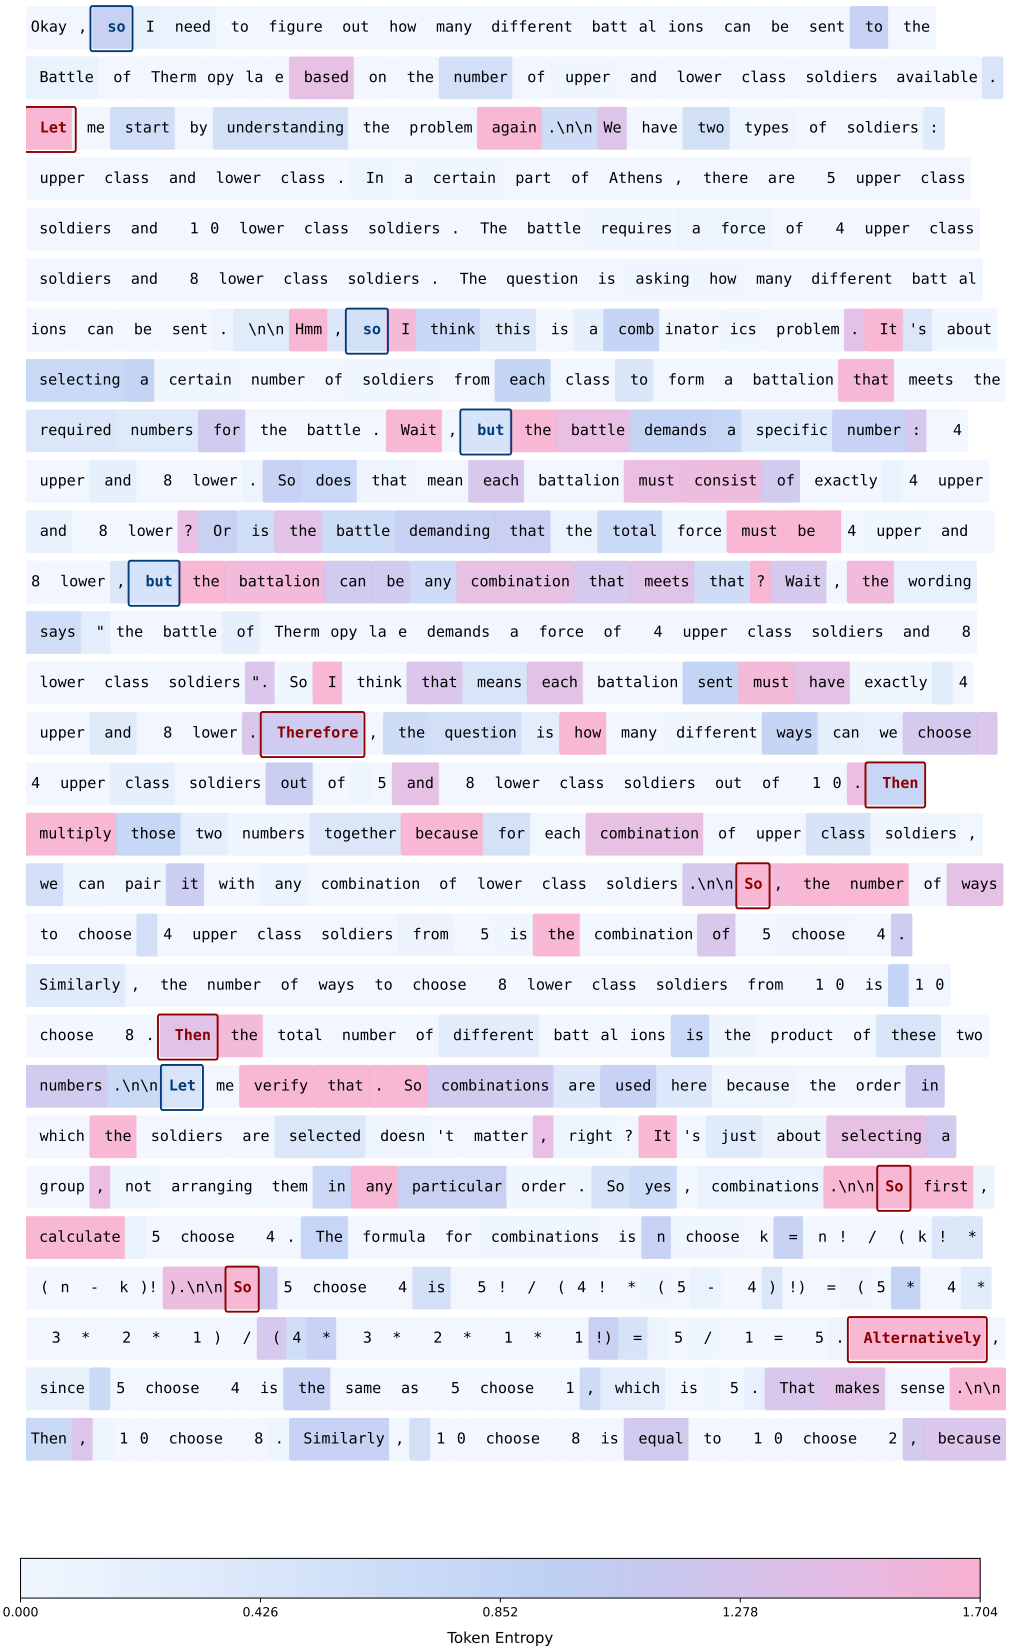

Figure 13: Dual-Entropy Example 1-part1

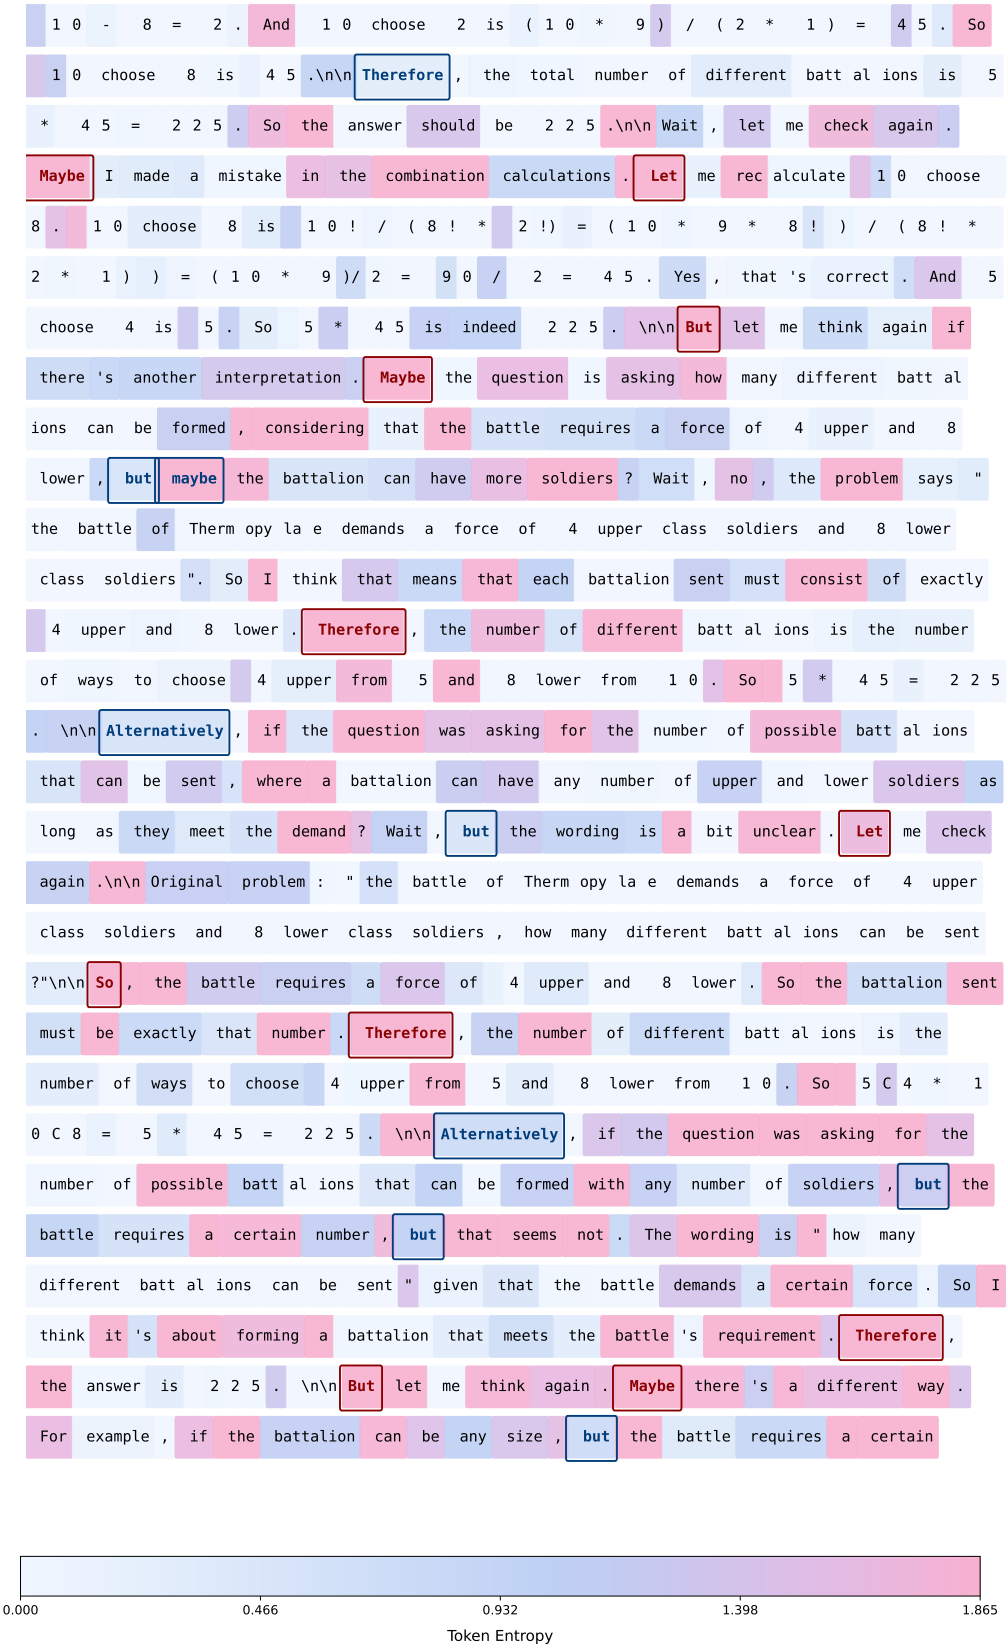

Figure 14: Dual-Entropy Example 1-part2

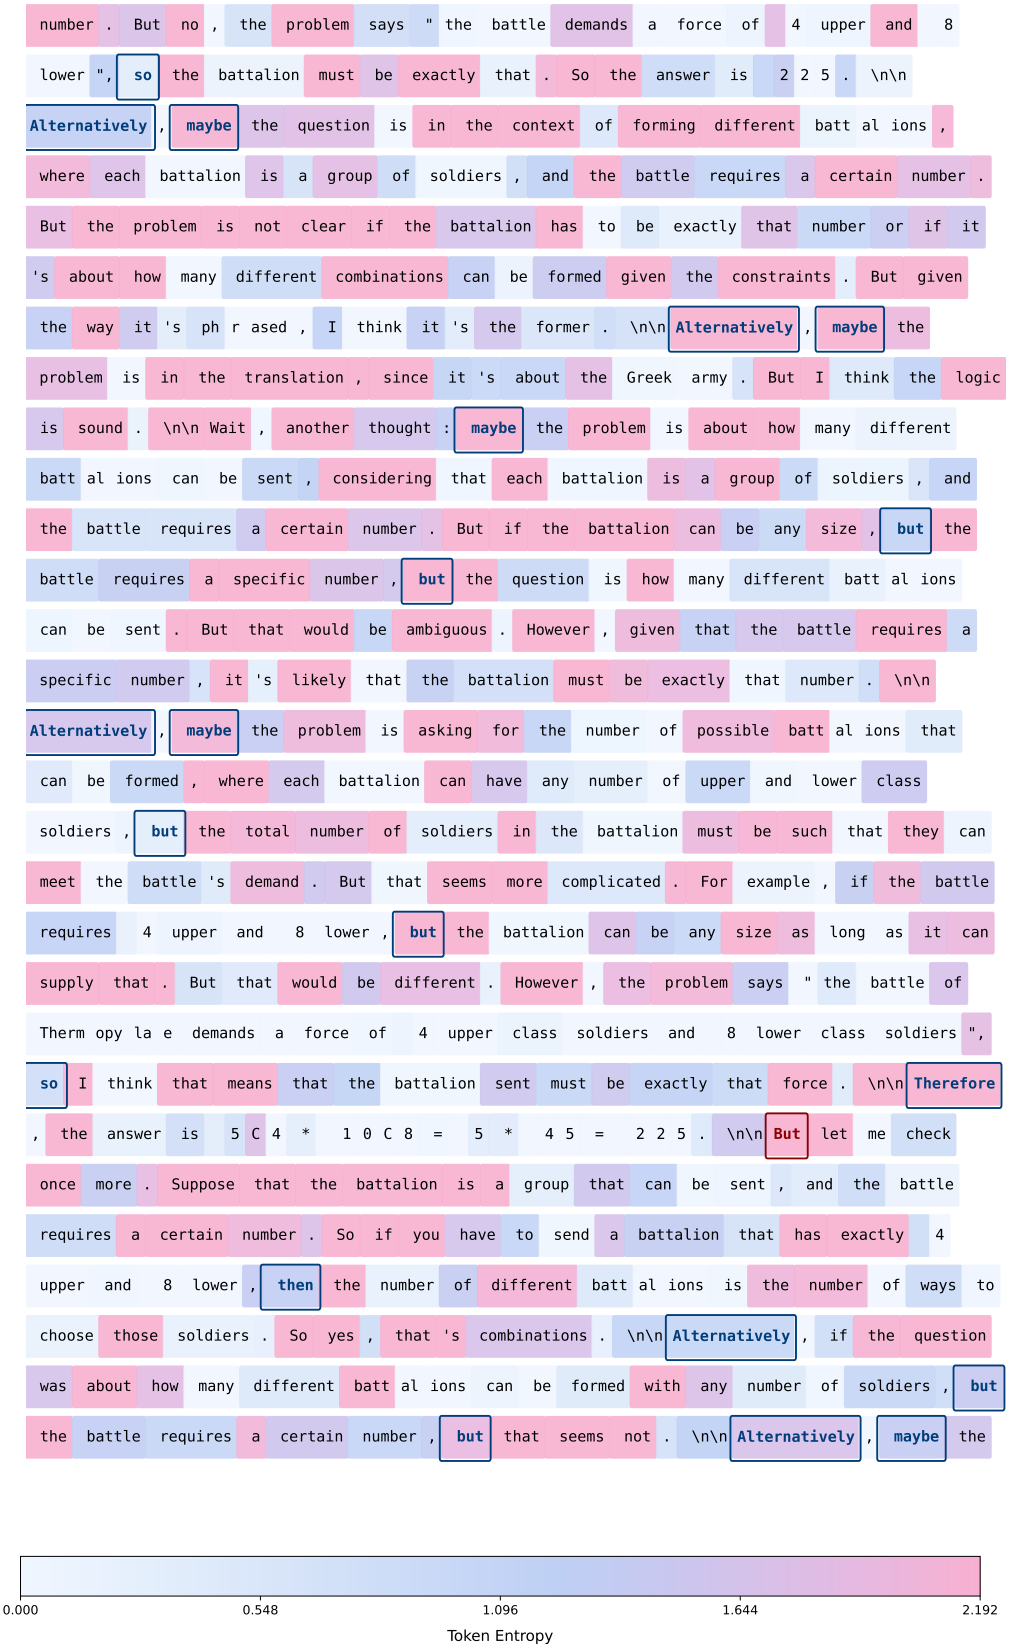

Figure 15: Dual-Entropy Example 1-part3

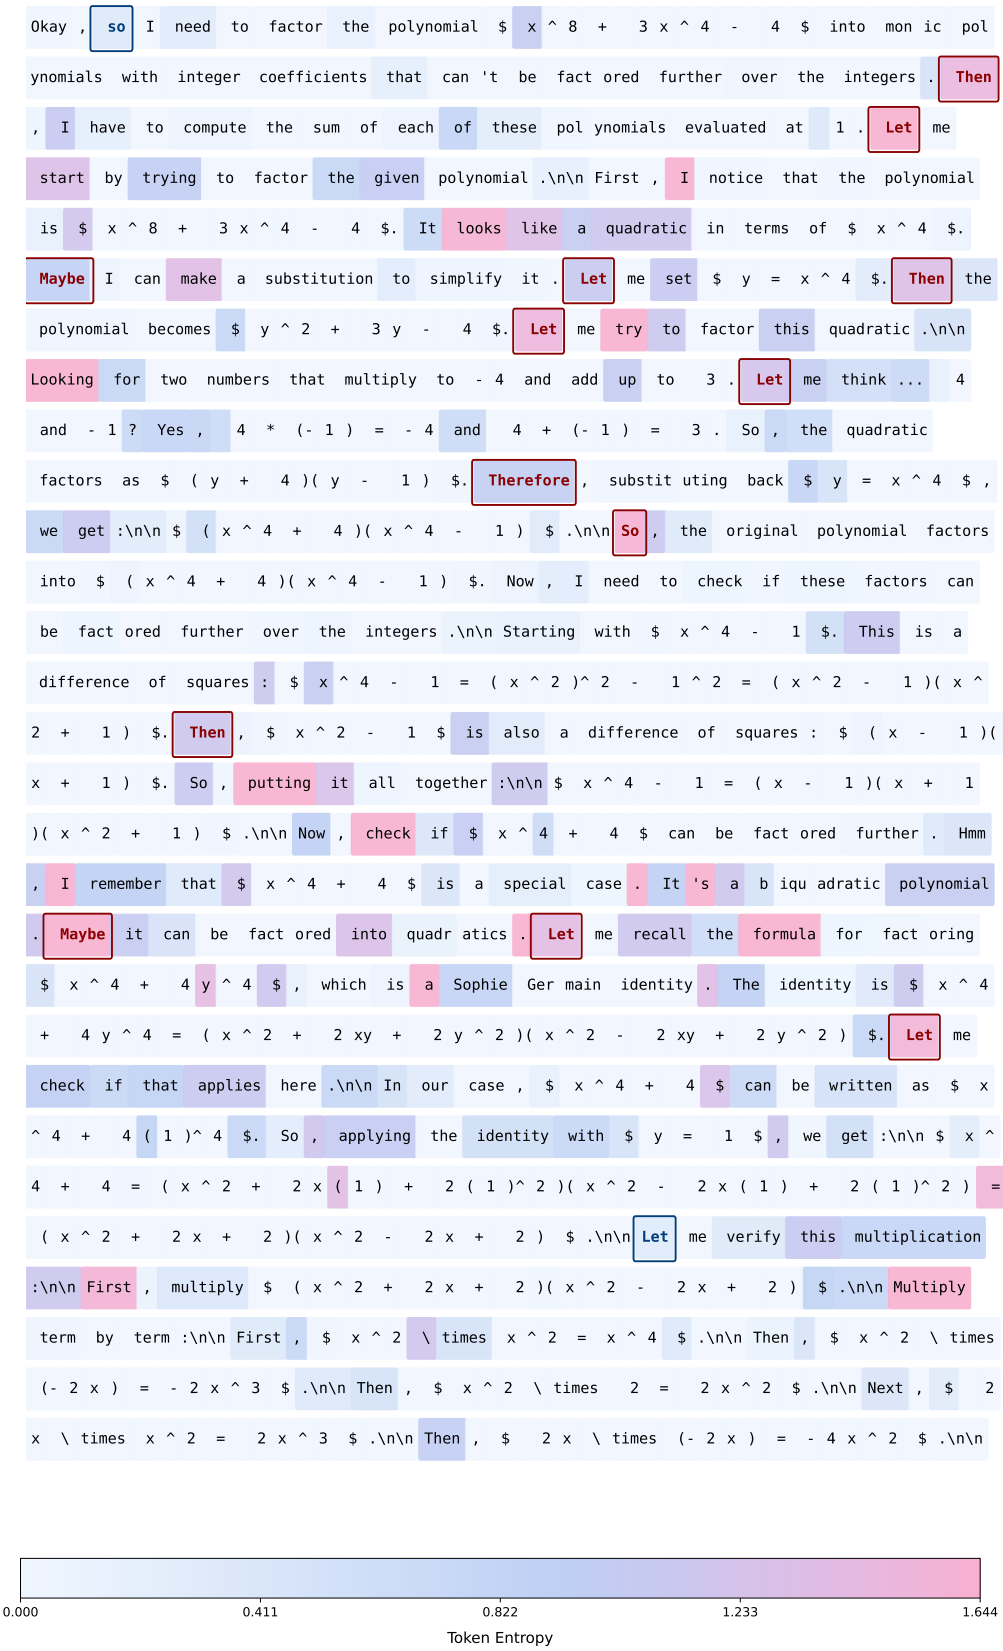

Figure 16: Dual-Entropy Example 2-part1

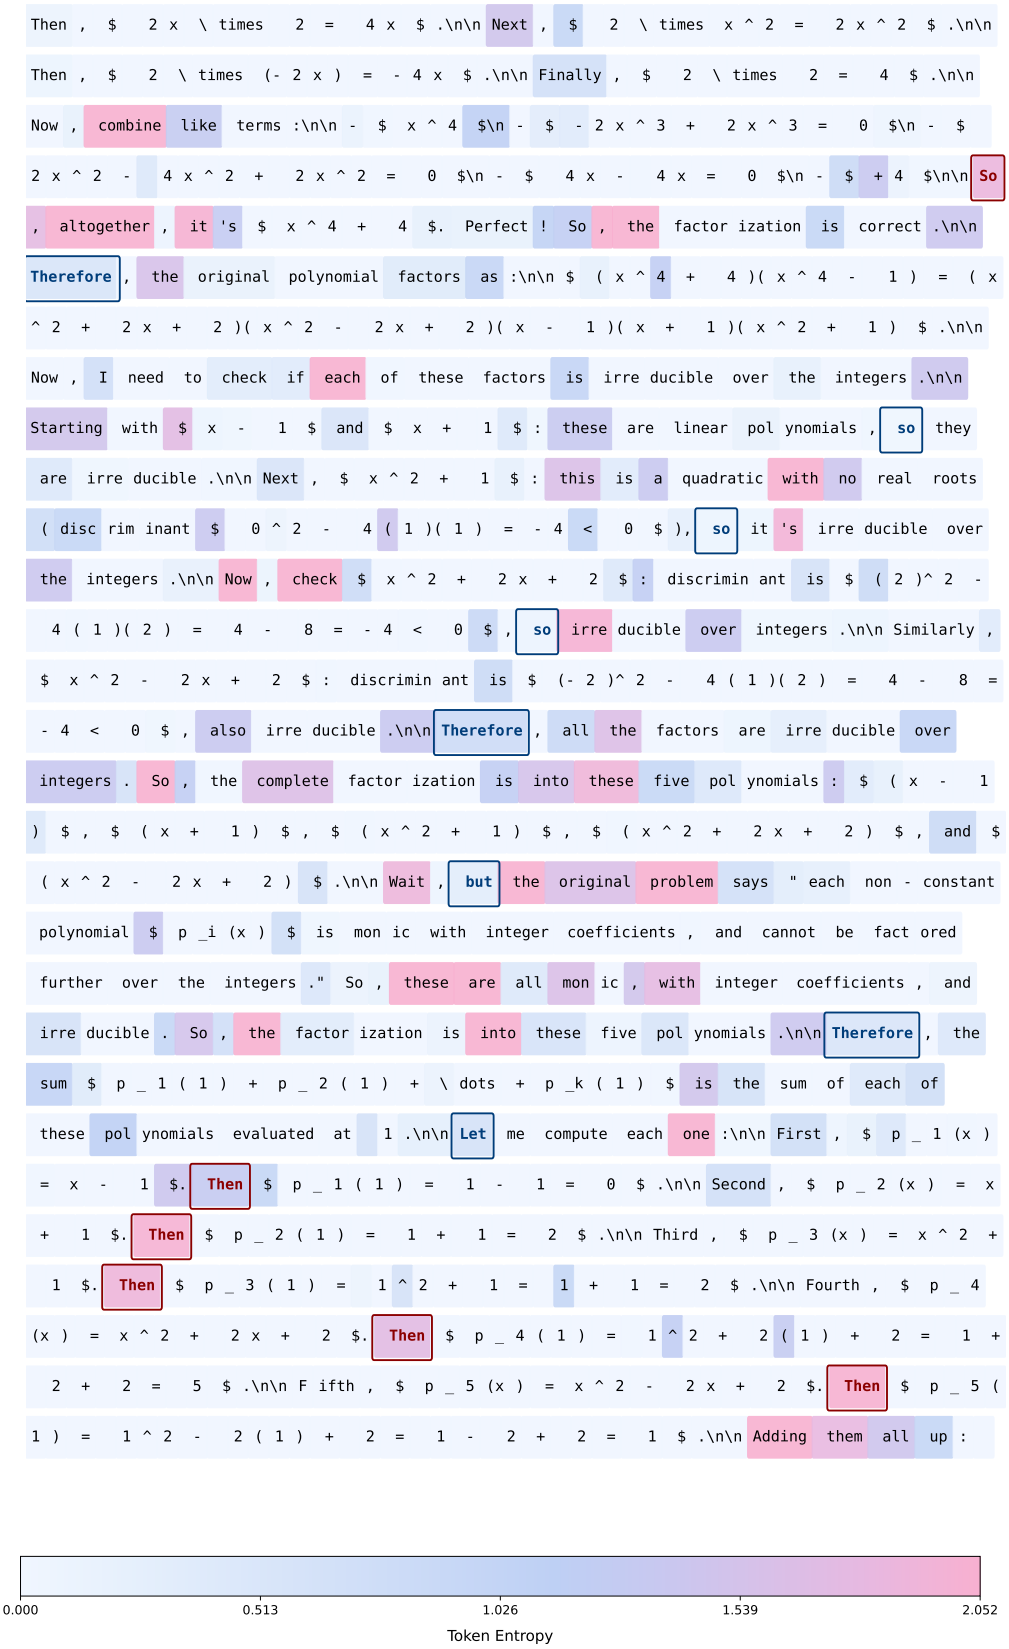

Figure 17: Dual-Entropy Example 2-part2

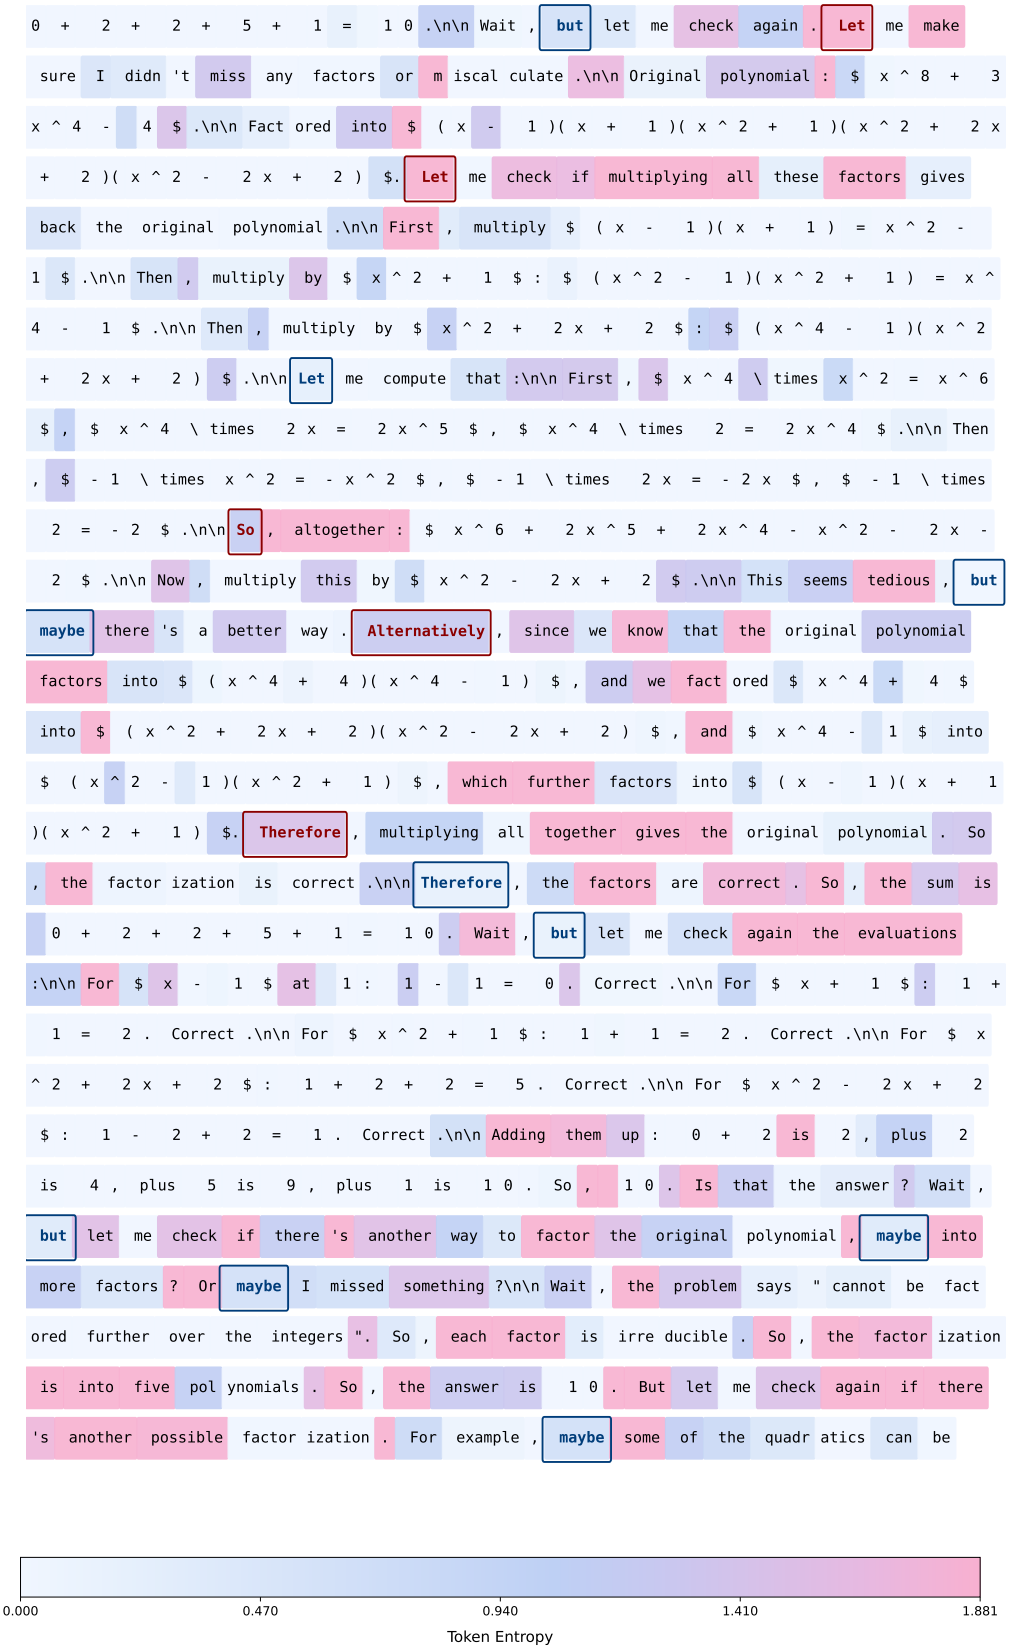

Figure 18: Dual-Entropy Example 2-part3