

DVPO: Balancing Robustness and Generalization for Robust RL through Distributional Value Modeling in LLM Post-Training

Anonymous ACL submission

Abstract

Reinforcement learning (RL) has shown strong performance in LLM post-training, but real-world deployment often involves noisy or incomplete supervision. In such settings, complex and unreliable supervision signals can destabilize training and harm generalization. While existing approaches such as worst-case optimization and mean-based methods can improve stability, they often overlook generalization and may produce overly conservative policies, leading to uneven performance across diverse real scenarios. To this end, we introduce DVPO (Distributional Value Modeling with Risk-aware Policy Optimization), a new RL framework that combines conditional risk theory with distributional value modeling to better balance robustness and generalization. DVPO learns token-level value distributions to provide fine-grained supervision, and applies an asymmetric risk regularization to shape the distribution tails: it contracts the lower tail to dampen noisy negative deviations, while expanding the upper tail to preserve exploratory diversity. Across extensive experiments and analysis in multi-turn dialogue, math reasoning, and scientific QA, DVPO consistently outperforms PPO, GRPO, and robust Bellman-based PPO under noisy supervision, showing its potential for LLM post-training in the real-world.

1 Introduction

Reinforcement learning has achieved impressive results across diverse domains (Xi et al., 2025b,a, 2024b,a; Ding et al., 2025; Caregnato-Neto et al., 2024). However, real-world RL often suffers from noisy or incomplete supervision, particularly when learning from human feedback or reward models. Such noise corrupts the value and advantage estimates, causing inaccurate policy updates and leading to unstable or divergent training dynamics. This issue is especially pronounced in LLM post-training (Havrilla and Iyer, 2024; Zhou et al., 2025)

and autonomous agent learning (Geng et al., 2024; Parać et al., 2024), where supervision signals are inherently complex and noisy.

Prior works have attempted to improve training stability through worst-case optimization, such as approaches based on robust Bellman operators (Panaganti et al., 2022; Kumar et al., 2020). However, their pessimistic formulations often lead to overly conservative policies that suppress informative signals, thereby limiting the model’s ability to generalize advantage estimates to OOD scenarios. Other studies aim to control training variance via mean-based estimators, including GRPO and iterative value averaging (Wang et al., 2023). While these methods help reduce instability, they still provide insufficient support for OOD generalization during training, often resulting in imbalanced advantage estimation performance across diverse real-world tasks. In practice, effective reinforcement learning methods must strike a balance between robustness to noisy rewards and cross-domain generalization ability (Su et al., 2025), enabling reliable performance under the diverse and challenging conditions encountered in real-world applications.

In this work, we propose a novel perspective to address this challenge by combining conditional risk theory (Kisiala, 2015) with a distributional training method (Shi and Chi, 2023). Instead of purely minimizing worst-case outcomes or learning the expectation of future rewards, our goal is to maximize expected performance under controlled risk: we enhance generalization through more comprehensive distributional supervision and improve training robustness via distributional constraints.

We introduce DVPO (Distributional Value Modeling with Risk-aware Policy Optimization), a robust RL framework that predicts token-level value distributions instead of scalar estimates, enabling distributional advantage computation and fine-grained supervision. A conditional risk constraint asymmetrically regulates the distribution

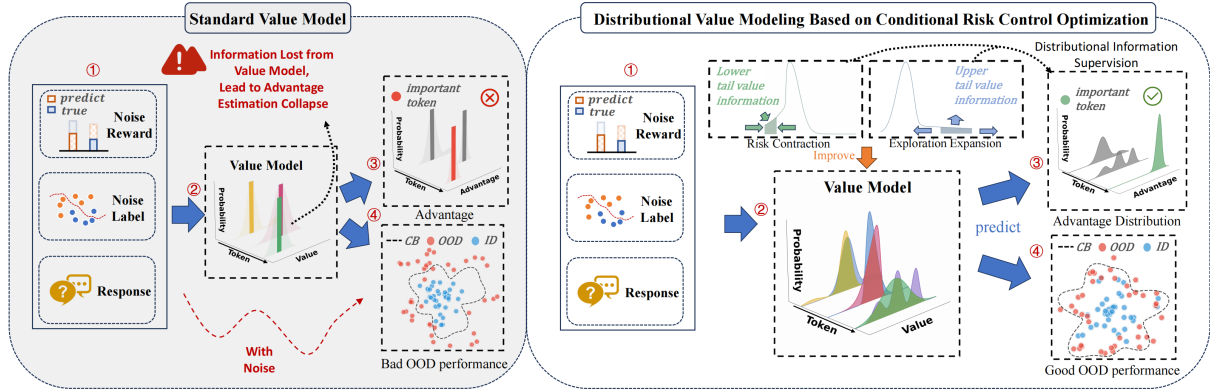


Figure 1: Comparison between the Standard Value Model and our Distributional Value Model Based on Conditional Risk Control Policy Optimization (DVPO). The standard value model suffers from reward noise and biased value estimation, leading to unstable policy updates. DVPO introduces a multi-head distributed value model to model the uncertainty of value, and further balances the model’s learning of robustness and generalization in noise by means of lower-tail risk contraction and upper-tail exploratory expansion.

tails, contracting the lower tail to reduce noisy deviations, while expanding the upper tail to maintain exploratory diversity. This design preserves informative latent structures while optimizing expected returns under bounded risk, enabling appropriate advantage estimation across diverse scenarios and achieving a balance between robustness and generalization.

Extensive experiments on multi-turn dialogue, mathematical reasoning, and scientific QA tasks show that DVPO consistently outperforms PPO, GRPO, and robust Bellman-based PPO variants under noisy supervision. By integrating distributional value learning with conditional risk control, DVPO achieves a robust and transferable learning framework, offering a scalable solution for real-world RL applications. Our contributions are summarized as follows:

- We propose DVPO, a robust RL framework with token-level value distributions for fine-grained supervision under noise.
- We design an asymmetric conditional risk constraint to regulate distribution tails, balancing robustness and generalization.
- We empirically demonstrate that DVPO outperforms representative baselines on multiple tasks, enhancing noise robustness and generalization.

2 Related Work

Robust Reinforcement Learning Robust Reinforcement Learning (RRL) aims to address

the sensitivity of standard RL to model uncertainty, with Robust Markov Decision Processes (RMDPs) (Nilim and Ghaoui, 2005; Bian and Jiang, 2018) serving as a representative framework. RMDPs formulate decision-making as a zero-sum game and optimize worst-case performance through the robust Bellman operator (Panağanti et al., 2022), which introduces inner minimization into value updates to obtain pessimistic value estimates. Although such formulations guarantee convergence via contraction properties, their inherent pessimism often restricts exploration and leads to overly conservative policies, limiting generalization in complex environments.

To improve training stability, prior work has explored variance reduction via mean-based estimators such as GRPO (Shao et al., 2024; Yu et al., 2025) and iterative value averaging (Wang et al., 2023), as well as clipping-based strategies (Becker et al., 2025; Yu et al., 2025), including importance weight clipping (Liu et al., 2025; MiniMax et al., 2025). While these methods improve stability, they do not explicitly resolve the fundamental trade-off between robustness and OOD generalization.

In contrast, our work addresses this limitation by integrating distributional value modeling with conditional risk control, enabling more balanced learning that preserves robustness while improving generalization under noisy supervision.

Distributional Reinforcement Learning Distributional Reinforcement Learning (DRL) (Zanger et al., 2025) extends traditional reinforcement learning by modeling the full probability distribution of future returns rather than merely their expected

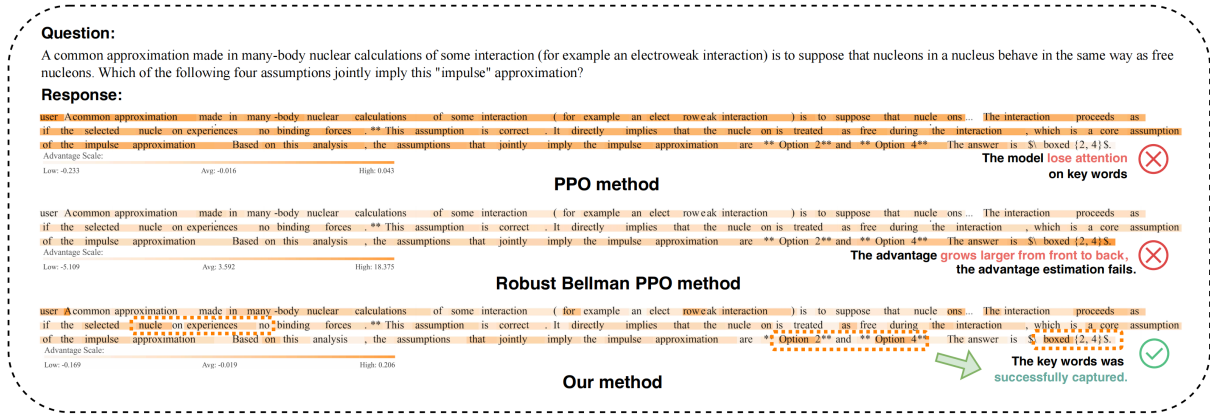


Figure 2: Token-level advantage estimation for the same response across different methods. Our method exhibits sharper focus on key words.




Figure 3: Comparison of the output value distributions of the first-token method for the answer part. The robust Bellman PPO method contracts partially in the lower tail, with small variance during learning but insufficient exploration. Our method also contracts slightly in the lower tail, achieves significant exploratory expansion learning in the upper tail, and maintains a good balance between generalization and robustness.

tations. This distributional assumption provides richer learning signals, enabling the agent to capture higher-order statistics such as variance, skewness, and multimodality, thereby improving stability and robustness in decision-making. Early works include C51 (Bellemare et al., 2017), which represents value distributions using discrete atoms; Quantile Regression DQN (QR-DQN) (Dabney et al., 2017), which models the return distribution through quantile regression; and the Implicit Quantile Network (IQN) (Dabney et al., 2018), which continuously approximates the quantile function to enhance representational flexibility. Recently, in Reinforcement Learning from Human Feedback (RLHF), the Quantile Reward Model (QRM) (Dorka, 2024) extends reward modeling from scalar signals to full distributions, capturing the ambiguity and multidimensionality of human preferences and guiding optimization toward safer and more reliable behaviors. Q# (Zhou et al., 2025) further applies DRL to LLM post-training by learn-

ing the cumulative return distribution of the reference policy to compute a provably optimal KL-regularized guiding Q-function. In this work, We build on these insights by introducing distributional supervision and conditional risk constraints into value modeling, enhancing stability and generalization under noisy supervision.

3 Method

3.1 Motivation

Training RL models in real-world environments inevitably involves noisy, incomplete supervision, especially in LLM post-training and autonomous agent learning. Existing robust RL methods, including robust Bellman operators (Panaganti et al., 2022) and mean estimation techniques (GRPO (Shao et al., 2024), PPO), iterative value averaging (Wang et al., 2023)), fail to balance robustness and generalization. Robust Bellman operators enhance stability via worst-case optimization but yield overly conservative policies; mean estimation

reduces variance through scalar-based biased estimation yet lacks cross-domain adaptability mechanisms. Both are inadequate for real-world RL scenarios that demand simultaneous resilience to noise and effective generalization.

A core limitation lies in their information-constrained value learning. Robust Bellman operators, in particular, suffer from excessively pessimistic estimation. As shown in Figure 3, training with minimum-value targets collapses the value distribution to its lower bound, producing stable but overly concentrated predictions. While this reduces variance, it suppresses high-value signals critical for diverse, informative feature learning. Under noisy supervision (Figure 2), this imbalance makes models overlook key words, letting spurious noise dominate task-related advantages.

To address these issues, we reformulate value learning from scalar prediction to distributional estimation, which captures richer value uncertainty and tail behaviors beyond biased mean or pessimistic estimates. Building on conditional risk control theory (Kisiala, 2015), we impose asymmetric constraints on the value distribution by contracting the lower tail to suppress noise and risk, while relaxing the upper tail to preserve informative high-value signals and encourage exploration. This design enables more accurate advantage estimation across diverse scenarios and effectively balances robustness and generalization under noisy supervision.

3.2 Distributional Architecture for Robust Value Representation

Multi-Headed Quantile Ensemble To obtain a stable and noise-resilient value distribution under imperfect supervision, our critic adopts a multi-headed quantile ensemble. Instead of relying on a single parametric estimator, which is often sensitive to idiosyncratic noise and induces high variance in distributional RL, we use an ensemble of N quantile heads to produce a more reliable distributional approximation. A shared backbone network $h_\psi(s)$ first encodes the state into a latent representation. This representation is then passed to N lightweight and independent quantile head networks, denoted by $\{f_{\phi_{Q,i}}\}_{i=1}^N$. Each head i predicts a set of M quantile values $\{\theta_{i,j}(s,a)\}_{j=1}^M$, corresponding to a predefined grid of quantile levels $\mathcal{T} = \{\hat{\tau}_j = j/M\}_{j=1}^M$. The final ensemble quantile function, $\hat{F}_{Z(s,a)}^{-1}$, is formed by averaging the outputs, creating a robust distributional baseline that is less

susceptible to outlier estimates:

$$\theta_j^{\text{ens}}(s,a) \triangleq \hat{F}_{Z(s,a)}^{-1}(\hat{\tau}_j) = \frac{1}{N} \sum_{i=1}^N \theta_{i,j}(s,a). \quad (1)$$

Distributional Generalized Advantage Estimation We propagate this robust distributional information through time using a distributional extension of the Generalized Advantage Estimation (Schulman et al., 2018) framework. Our procedure operates entirely in the quantile domain to preserve the full spectrum of uncertainty during credit assignment. Let the value distribution at time t be represented by its quantile vector Θ_{V_t} . The distributional advantage, A_t^D , represented by its quantile vector Θ_{A_t} , is then computed via a recursive update that lifts the standard GAE formulation to the quantile space:

$$\Theta_{A_t} = (r_t + \gamma \Theta_{V_{t+1}} - \Theta_{V_t}) + (\gamma \lambda) \Theta_{A_{t+1}}. \quad (2)$$

The term in parentheses represents the distributional one-step TD-error. This recursive computation yields two critical outputs: a target return distribution for the value model update, defined by its quantile vector $\Theta_{Z'_t} = \Theta_{V_t} + \Theta_{A_t}$; and a scalar advantage estimate for the actor update, recovered by taking the expectation, $A(s_t, a_t) = \mathbf{E}[A_t^D] \approx \frac{1}{M} \sum_{j=1}^M \Theta_{A_t,j}$. By performing credit assignment entirely on distributions, this method provides a richer and more nuanced learning signal for our subsequent risk-aware value model optimization.

3.3 Calibrating Risk-Aware Distributional Optimization

Learning reliable value distributions under noisy or imperfect supervision requires simultaneously promoting robustness to adverse perturbations and preserving generalization to high-quality signals. To achieve this balance, we formulate a composite critic objective that shapes the predicted return distribution along multiple dimensions, including central quantile alignment, tail calibration, curvature control, and cross-distribution coherence.

Let the critic output a parametric quantile distribution $\hat{Z}(s,a)$ represented by M quantile values $\Theta = \{\theta_j(s,a)\}_{j=1}^M$, evaluated at quantile levels $\mathcal{T} = \{\hat{\tau}_j\}_{j=1}^M$ where $\hat{\tau}_j = j/M$. Let the target return distribution be $Z'(s,a)$ with corresponding quantiles $\Theta' = \{\theta'_j(s,a)\}_{j=1}^M$. We denote by π the permutation that sorts target quantiles, such that $\theta'_{\pi(1)} \leq \dots \leq \theta'_{\pi(M)}$. For convenience, we define prediction errors $u_j = \theta'_j - \theta_j$. \mathbf{E} represents expectation.

1. Core Quantile Regression The foundational component is the quantile Huber regression loss, which aligns each predicted quantile with its target counterpart:

$$\mathcal{L}_{\text{QR}} = \mathbf{E}_{s,a} \left[\frac{1}{M} \sum_{j=1}^M |\hat{\tau}_j - \mathbf{I}(u_j < 0)| L_\delta(u_j) \right] \quad (3)$$

where $L_\delta(\cdot)$ denotes the Huber loss and $\mathbf{I}(\cdot)$ is the indicator function. This objective provides stable fitting but does not control robustness or generalization on its own.

2. Risk-Sensitive Quantile Weighting To explicitly enhance robustness to noisy negative rewards, we modulate quantile errors using a weighting function $w_{\text{risk}}(\hat{\tau}_j)$, parameterized by a risk-aversion coefficient γ :

$$\mathcal{L}_{\text{Risk}} = \mathbf{E}_{s,a} \left[\frac{1}{M} \sum_{j=1}^M w_{\text{risk}}(\hat{\tau}_j) |\hat{\tau}_j - \mathbf{I}(u_j < 0)| L_\delta(u_j) \right] \quad (4)$$

Larger γ increases emphasis on lower quantiles, promoting conservative updates that mitigate overfitting to noisy supervision.

3. Tail Expectation Calibration To regulate extreme outcomes, we align the expectations of both lower and upper tails. Let $K_\alpha = \lfloor \alpha M \rfloor$ denote the lower tail size and $K_\beta = \lfloor \beta M \rfloor$ the upper tail size.

The CVaR-style objective stabilizes the worst-case region:

$$\mathcal{L}_{\text{CVaR}} = \mathbf{E}_{s,a} \left[\frac{1}{K_\alpha} \left(\sum_{j=1}^{K_\alpha} \theta'_{\pi(j)} - \sum_{j=1}^{K_\alpha} \theta_{\pi(j)} \right)^2 \right] \quad (5)$$

The upper-gain objective encourages accurate modeling of high-value outcomes, thus supporting generalization:

$$\mathcal{L}_{\text{Gain}} = \mathbf{E}_{s,a} \left[\frac{1}{K_\beta} \left(\sum_{j=M}^{K_\beta} \theta'_{\pi(j)} - \sum_{j=M}^{K_\beta} \theta_{\pi(j)} \right)^2 \right] \quad (6)$$

4. Controlling Systematic Bias via Mean-Shift Penalization To prevent persistent underestimation of the return distribution, we penalize cases where the predicted mean falls below the target mean:

$$\mathcal{L}_{\text{Shift}} = \mathbf{E}_{s,a} [\text{ReLU}(\mathbf{E}_j[\theta'_j] - \mathbf{E}_j[\theta_j])] \quad (7)$$

This reduces overly conservative value estimates that would otherwise impair exploration and generalization.

5. Tail-Shape Regularization We regulate dispersion in the extreme tails using variance-based penalties. Let $\text{Var}(\Theta; \mathcal{I})$ denote the variance over index set \mathcal{I} , and define $\mathcal{I}_\alpha = \{1, \dots, K_\alpha\}$ and $\mathcal{I}_\beta = \{M - K_\beta + 1, \dots, M\}$.

The tail-shape loss is:

$$\mathcal{L}_{\text{Shape}} = \mathbf{E}_{s,a} \left[\text{ReLU}(\text{Var}(\Theta; \mathcal{I}_\alpha) - \text{Var}(\Theta'; \mathcal{I}_\alpha)) + \text{ReLU}(\text{Var}(\Theta'; \mathcal{I}_\beta) - \text{Var}(\Theta; \mathcal{I}_\beta)) \right] \quad (8)$$

The first promotes a compact lower tail robustness, while the second encourages a sufficiently expressive upper tail generalization.

6. Tail Curvature Regularization To further shape extreme quantile geometry, we regularize discrete curvature using the second-order finite difference $\Delta^2 \theta_j = \theta_{j+1} - 2\theta_j + \theta_{j-1}$:

$$\mathcal{L}_{\text{Curv}} = \mathbf{E}_{s,a} \left[\text{ReLU}(\mathbf{E}_{j \in \mathcal{I}_\alpha} [\Delta^2 \theta_j]) + \text{ReLU}(-\mathbf{E}_{j \in \mathcal{I}_\beta} [\Delta^2 \theta_j]) \right] \quad (9)$$

This favors concave lower-tail curvature risk aversion and convex upper-tail curvature optimistic generalization.

7. Multi-Distribution Consistency Across Ensemble Heads If the critic uses multiple independently parameterized quantile distributions $\{\hat{Z}^{(i)}(s, a)\}_{i=1}^{N_{\text{dist}}}$, each producing quantiles $\Theta^{(i)} = \{\theta_j^{(i)}\}_{j=1}^M$, we encourage agreement across heads:

$$\mathcal{L}_{\text{Consist}} = \frac{1}{N_{\text{pair}}} \sum_{i < j} \mathbf{E}_{s,a} [\|\Theta^{(i)} - \Theta^{(j)}\|_2^2] \quad (10)$$

where N_{pair} counts all head pairs.

Final Composite Critic Loss The full objective aggregates all components:

$$\mathcal{L}_{\text{Critic}} = \mathcal{L}_{\text{QR}} + w_{\text{risk}} \mathcal{L}_{\text{Risk}} + w_{\text{cvar}} \mathcal{L}_{\text{CVaR}} + w_{\text{gain}} \mathcal{L}_{\text{Gain}} + w_{\text{shift}} \mathcal{L}_{\text{Shift}} + w_{\text{shape}} \mathcal{L}_{\text{Shape}} + w_{\text{curv}} \mathcal{L}_{\text{Curv}} + w_{\text{consist}} \mathcal{L}_{\text{Consist}} \quad (11)$$

By jointly shaping central quantiles, tail expectations, dispersion, curvature, and multi-head coherence, this calibrated objective explicitly balances robustness to noisy supervision with generalization to high-value trajectories, enabling stable and noise-resilient value distribution learning.

4 Experiments

4.1 Setup

We evaluate DVPO in math, science and dialogue task under rule-based rewards and model-based rewards supervision.

Domain	In-Domain		Out-of-Domain			ALL
Method	Life Services	Transportation & Travel	Healthcare & Wellness	Social & Entertainment	Financial Services	AVG
Baseline	86.73%	84.50%	90.23%	87.13%	82.70%	86.26%
GRPO	30.03%	28.57%	28.33%	28.90%	27.90%	28.75%
PPO	84.20%	86.07%	90.87%	81.00%	83.87%	85.20%
Reinforce++	87.87%	81.07%	77.67%	87.13%	85.53%	83.85%
Dr.GRPO	51.67%	60.30%	61.37%	52.17%	56.47%	56.39%
DVPO (Ours)	88.13%	87.73%	87.67%	87.67%	82.73%	86.79%

Table 1: Accuracy (%) of different reinforcement learning methods on the **Real-world Dialog Domains** under noisy reward signals. DVPO achieves the highest average accuracy of 86.79%, surpassing all baselines. While GRPO and Dr.GRPO collapse under noise dropping below 60%, PPO and Reinforce++ remain stable but show limited consistency across domains.

Domain	In-Domain (Science & QA)			Out-of-Domain (Math)				ID	OOD	ALL
Method	SampleQA	GPQA(ALL)	HLE	MATH500	AIME24	Minerva-Math	AMC23	AVG	AVG	AVG
Base	2.89%	3.10%	2.89%	87.40%	41.67%	28.68%	75.83%	2.96%	58.40%	34.64%
GRPO	3.03%	2.98%	3.24%	86.80%	45.00%	26.84%	82.50%	3.08%	60.29%	35.77%
PPO	2.82%	2.17%	3.29%	87.00%	51.67%	29.41%	85.00%	2.76%	63.27%	37.34%
Reinforce++	3.19%	4.35%	3.29%	89.40%	50.00%	30.51%	83.33%	3.61%	63.31%	37.72%
Dr.GRPO	2.50%	3.99%	3.10%	88.60%	48.33%	27.21%	80.00%	3.20%	61.04%	36.25%
Robust Bellman	3.14%	4.17%	3.38%	87.40%	45.00%	27.21%	83.33%	3.56%	60.74%	36.23%
DVPO (Ours)	3.21%	4.71%	3.57%	90.60%	56.67%	31.99%	86.67%	3.83%	66.48%	39.63%

Table 2: Accuracy (%) of different reinforcement learning methods trained on the **Science Domain**. DVPO achieves the highest ID average (3.83%) and OOD average (66.48%), showing strong robustness to noisy supervision and superior generalization across domains.

Dialogue Task To assess the effectiveness of our approach in realistic settings with model-based reward supervision, we conduct multi-turn dialogue experiments on the **Honor-Dialogue Dataset**. This dataset, constructed by us, contains multi-domain, task-oriented dialogues collected from real-world scenarios. The task requires the model to act as a dialogue assistant, producing natural and effective responses. Under noisy dialogue reward model supervision, the model is trained with reinforcement learning on real multi-turn dialogues. Performance is evaluated through interactions with GPT-4o using the dialogue standards in Section 4.2. Given that the performance of some methods degrades drastically as training steps increase, we use models trained for limited steps for experimental comparison. All real information in the data has undergone anonymization, and the use of the data has been authorized and reviewed by the ethics committee.

Math and Science Task To further assess the effectiveness of our approach in realistic settings where supervision comes from rule-based rewards, we conduct a series of experiments on mathematical and scientific reasoning tasks during training. Specifically, we use the Light-R1 dataset (Wen et al., 2025) for math tasks and the SuperGPQA (Team et al., 2025) dataset for scientific tasks. For the dataset, we first conduct multiple inference runs with the original model and then adopt

majority voting to filter the results as training labels prior to model training. The evaluation is then conducted on math and scientific QA datasets. The test datasets include four math-focused tasks: MATH500 (Hendrycks et al., 2021), AIME24, Minerva-Math (Lewkowycz et al., 2022), AMC23, as well as three scientific knowledge tasks: SampleQA (Wei et al., 2024), GPQA, and HLE (Humanity’s Last Exam) (Phan et al., 2025). Table 2 reports the accuracy of models before and after RL training. Due to the limitation of memory, multiple rounds of generation are carried out here. In each round, the generated length is at most 4096 tokens, and the result of the last round is used for evaluation.




Figure 4: Noise statistics in the various tasks. A significant portion of rewards contains inaccuracies.

Baseline Initialization In the math and scientific tasks, we use Qwen3-8B (Yang et al., 2025) for initialization. In the dialogue tasks, both the policy and reward models are initialized from Qwen3-8B and then fine-tuned on the Honor-Dialogue dataset.

Domain	In-Domain (Math)				Out-of-Domain (Science & QA)			ID	OOD	ALL
Method	MATH500	AIME24	Minerva-Math	AMC23	SampleQA	GPQA(ALL)	HLE	AVG	AVG	AVG
Base	87.40%	41.67%	28.68%	75.83%	2.89%	3.10%	2.89%	58.40%	2.96%	34.64%
GRPO	89.20%	35.00%	28.68%	84.17%	2.91%	3.99%	3.01%	59.26%	3.30%	35.28%
PPO	86.40%	36.67%	30.51%	80.00%	2.70%	3.44%	3.34%	58.40%	3.16%	34.72%
Reinforce++	89.00%	45.00%	27.21%	84.17%	3.19%	4.35%	3.61%	61.35%	3.72%	36.65%
Dr.GRPO	87.80%	45.00%	27.94%	85.83%	3.10%	4.35%	3.20%	61.64%	3.55%	36.75%
Robust Bellman	89.40%	36.67%	31.99%	84.17%	2.61%	3.99%	3.62%	60.56%	3.22%	35.98%
DVPO (Ours)	90.00%	56.67%	31.62%	87.50%	3.31%	5.25%	3.57%	66.45%	4.04%	39.70%

Table 3: Accuracy (%) of different reinforcement learning methods trained on the **Math Domain**. DVPO significantly outperforms all baselines, achieving an ID average of 66.45% and OOD average of 4.04%, indicating better generalization across noisy tasks.

Domain	In-Domain (Science & QA)			Out-of-Domain (Math)				ID	OOD	ALL
Risk Interval	SampleQA	GPQA(ALL)	HLE	MATH500	AIME24	Minerva-Math	AMC23	AVG	AVG	AVG
Base	2.89%	3.10%	2.89%	87.40%	41.67%	28.68%	75.83%	2.96%	58.40%	34.64%
0	3.05%	5.25%	3.52%	90.40%	45.00%	28.68%	79.17%	3.94%	60.81%	36.44%
0.05	3.10%	4.89%	3.52%	90.20%	46.67%	31.25%	79.17%	3.84%	61.82%	36.97%
0.1 (Ours)	3.21%	4.71%	3.57%	90.60%	56.67%	31.99%	86.67%	3.83%	66.48%	39.63%
0.2	3.28%	4.35%	3.15%	90.00%	46.67%	29.41%	85.00%	3.59%	62.77%	37.41%

Table 4: Comparison of DVPO performance under different risk interval settings across multiple domains. A moderate interval weight of 0.1 provides the most balanced and stable results, while both smaller and larger intervals degrade robustness and generalization.

In our experiments, we implemented a PPO method based on the robust Bellman operator. Specifically, a multi-head value model is constructed to assign multiple value estimates to each token, which are then aggregated by selecting the lowest value as the current estimate to simulate the worst-case scenario. This worst-case value serves as the optimization target, guiding the overall value model training toward robustness under adverse conditions.

Training Noise Noise is pervasive in RL training, especially under model-based feedback. Figure 4 illustrates the error rate in various reward signals. The noise in the rule-based reward from the majority vote generated data.

Implementation Details A majority vote of 5 samples is conducted. The risk threshold is set at 0.1, the number of interval quantile points is set to 200 and the number of distributional value head is set to 3. During training, the model generates dialogue tasks of length 4096 for reasoning tasks and 1024 for dialogue tasks. The experiments are conducted on 8 * NVIDIA A100 80GB GPUs. For further dataset construction, training details, and model initialization, please refer to Appendix.

4.2 Evaluation Metrics

For mathematical reasoning and scientific knowledge tasks, we adopt accuracy as the primary metric to measure whether the model produces correct answers. For multi-turn dialogue tasks, we em-

ploy a tripartite evaluation framework that captures both task execution and communicative quality: (1) Task Completion Rate (TCR), assessing whether the model successfully fulfills the user-intended objective; (2) Ask Completion Rate (ACR), measuring whether detailed aspects of the user query are adequately addressed; and (3) Goal Completion Rate (GCR), evaluating the fluency, coherence, and appropriateness of the model’s responses. This comprehensive evaluation allows us to examine not only whether the model achieves the intended outcomes, but also how effectively it aligns with human communication standards.

4.3 Experiment results

DVPO exhibits strong robustness under real noisy supervision As shown in Table 1, Table 2, and Table 3, DVPO consistently improves performance across diverse domains even with imperfect reward signals. In dialogue tasks, the average task completion rate increases from 71.48% to 73.66%, while both GRPO and Dr.GRPO exhibit severe performance degradation under noisy rewards, reflecting the effectiveness of DVPO’s distributional value supervision. In mathematical and scientific reasoning tasks, DVPO maintains strong performance, achieving average accuracies of 39.70% and 39.63% respectively, substantially surpassing PPO and GRPO, demonstrating its robustness in multi-domain noisy training.

Domain	In-Domain (Science & QA)			Out-of-Domain (Math)				ID	OOD	ALL
	SampleQA	GPQA(ALL)	HLE	MATH500	AIME24	Minerva-Math	AMC23	AVG	AVG	AVG
Base	2.89%	3.10%	2.89%	87.40%	41.67%	28.68%	75.83%	2.96%	58.40%	34.64%
50	2.84%	3.62%	3.06%	87.60%	46.67%	31.99%	83.33%	3.17%	62.40%	37.02%
100	2.94%	3.99%	3.20%	88.20%	55.00%	31.99%	83.33%	3.38%	64.63%	38.38%
200 (Ours)	3.21%	4.71%	3.57%	90.60%	56.67%	31.99%	86.67%	3.83%	66.48%	39.63%
500	2.84%	3.80%	3.34%	89.40%	46.67%	31.62%	85.83%	3.33%	63.38%	37.64%

Table 5: Comparison of DVPO results under different interval densities across multiple domains. A moderate density of 200 yields the most stable and accurate performance, while both overly sparse and overly dense settings reduce robustness and generalization under noisy supervision.

DVPO achieves reliable balance in in-domain and cross-domain performance

As shown in Table 1, Table 2, and Table 3, DVPO demonstrates consistent advantages across dialogue, mathematical, and scientific domains under noisy reward supervision. In real-world dialogue tasks, DVPO achieves an average accuracy of 88.13% in the life services domain, while maintaining strong performance across other domains, outperforming all baseline methods. When trained on mathematical or scientific reasoning tasks, DVPO attains the highest in-domain accuracy and exhibits remarkable out-of-domain generalization, consistently surpassing PPO and robust Bellman-based approaches. Notably, while Robust Bellman PPO exhibits stable training loss and an optimistic reward trajectory during training, it suffers from poor generalization outside the training domain. It frequently generates lengthy, incoherent responses on unseen tasks and even fails to maintain the required response format, which makes it impossible to present its results in the dialogue task. In contrast, DVPO achieves a good balance between stability and generalization by combining distributional value modeling with conditional risk control, showing strong performance both within and across domains.

4.4 Ablation Experiment and Discussion

To further analyze the impact of the DVPO method, an experimental ablation study is conducted on the scientific task. The SuperGPQA dataset is still used for training, and the parameters remain the same as those in the main experiment.

A proper interval density helps the model learn more effectively from value distributions

As shown in Table 5, when the interval density decreases to 50, the average performance drops from 39.63% to 38.38% and further to 37.02%. When the density becomes excessively high, up to 500, performance also declines to 37.64%, with both in-domain and out-of-domain results deteriorating no-

ticeably. During the early phase of training, using a slightly higher interval density helps the model receive richer supervision signals, giving more consistent feedback under noisy rewards and improving training stability. However, when the intervals become too fine-grained, the model tends to capture noise as useful information, which reduces robustness. In contrast, moderately sparse intervals allow the model to focus on key informative regions while keeping the distribution smoother, leading to better generalization and stability.

A proper level of risk interval control enhances robustness and generalization

As shown in Table 4, removing risk interval control entirely (weight = 0) increases local accuracy on GPQA but leads to poor generalization and unstable training across mathematical domains. When the interval weight becomes too large, excessive perturbations in the value distribution distort the relationship between training and generation distributions, causing instability and lower overall performance. In contrast, DVPO with a moderate interval setting 0.1 achieves the best trade-off. It suppresses noise effectively while maintaining adequate exploration, resulting in the highest average accuracy and stable optimization behavior across domains.

5 Conclusion

We introduce DVPO, a robust RL framework that integrates distributional value modeling with conditional risk control in LLM post-training. DVPO offers a novel perspective on balancing robustness and generalization, achieving a principled trade-off between the two under noisy supervision by modeling value distributions and regulating their risk-aware tail structure. Experiments across dialogue, mathematical, and scientific tasks verify its superiority over PPO, GRPO, and robust Bellman PPO methods, providing a scalable and effective solution for real-world RL.

551 Limitation

552 While DVPO demonstrates strong robustness and
553 generalization under noisy supervision, several lim-
554 itations remain. The optimal choice of distribu-
555 tional interval density and risk thresholds may vary
556 across domains, requiring task-specific tuning for
557 best performance. And, although DVPO improves
558 stability under moderate noise, extreme reward cor-
559 ruption or mis-specified supervision can still de-
560 grade performance.

561 References

562 Philipp Becker, Niklas Freymuth, Serge Thilges, Fabian
563 Otto, and Gerhard Neumann. 2025. [Troll: Trust
564 regions improve reinforcement learning for large lan-
565 guage models](#). *Preprint*, arXiv:2510.03817.

566 Marc G. Bellemare, Will Dabney, and Rémi Munos.
567 2017. [A distributional perspective on reinforcement
568 learning](#). *Preprint*, arXiv:1707.06887.

569 Tao Bian and Zhong-Ping Jiang. 2018. [Continuous-
570 time robust dynamic programming](#). *Preprint*,
571 arXiv:1809.05867.

572 Angelo Caregnato-Neto, Luciano Cavalcante Siebert,
573 Arkady Zgonnikov, Marcos Ricardo Omena de Al-
574 buquerque Maximo, and Rubens Junqueira Magal-
575 hães Afonso. 2024. [ARMCHAIR: integrated inverse
576 reinforcement learning and model predictive con-
577 trol for human-robot collaboration](#). *arXiv e-prints*,
578 arXiv:2402.19128.

579 Will Dabney, Georg Ostrovski, David Silver, and
580 Rémi Munos. 2018. [Implicit quantile networks
581 for distributional reinforcement learning](#). *Preprint*,
582 arXiv:1806.06923.

583 Will Dabney, Mark Rowland, Marc G. Bellemare,
584 and Rémi Munos. 2017. [Distributional reinforce-
585 ment learning with quantile regression](#). *Preprint*,
586 arXiv:1710.10044.

587 Yiwen Ding, Zhiheng Xi, Wei He, Zhuoyuan Li, Yitao
588 Zhai, Xiaowei Shi, Xunliang Cai, Tao Gui, Qi Zhang,
589 and Xuanjing Huang. 2025. [Mitigating tail narrow-
590 ing in llm self-improvement via socratic-guided sam-
591 pling](#). *Preprint*, arXiv:2411.00750.

592 Nicolai Dorka. 2024. [Quantile regression for dis-
593 tributional reward models in rlhf](#). *Preprint*,
594 arXiv:2409.10164.

595 Wei Geng, Baidi Xiao, Rongpeng Li, Ning Wei, Dong
596 Wang, and Zhifeng Zhao. 2024. [Noise distribution
597 decomposition based multi-agent distributional rein-
598 forcement learning](#). *Preprint*, arXiv:2312.07025.

599 Alex Havrilla and Maia Iyer. 2024. [Understanding the
600 effect of noise in llm training data with algorithmic
601 chains of thought](#). *Preprint*, arXiv:2402.04004.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul
602 Arora, Steven Basart, Eric Tang, Dawn Song, and
603 Jacob Steinhardt. 2021. [Measuring mathematical
604 problem solving with the math dataset](#). *Preprint*,
605 arXiv:2103.03874. 606

Jakob Kisiala. 2015. [Conditional value-at-risk: Theory
607 and applications](#). *Preprint*, arXiv:1511.00140. 608

Aviral Kumar, Aurick Zhou, George Tucker, and Sergey
609 Levine. 2020. [Conservative q-learning for offline
610 reinforcement learning](#). *Preprint*, arXiv:2006.04779. 611

Aitor Lewkowycz, Anders Andreassen, David Dohan,
612 Ethan Dyer, Henryk Michalewski, Vinay Ramasesh,
613 Ambrose Slone, Cem Anil, Imanol Schlag, Theo
614 Gutman-Solo, Yuhuai Wu, Behnam Neyshabur, Guy
615 Gur-Ari, and Vedant Misra. 2022. [Solving quan-
616 titative reasoning problems with language models](#).
617 *Preprint*, arXiv:2206.14858. 618

Zongkai Liu, Fanqing Meng, Lingxiao Du, Zhixi-
619 ang Zhou, Chao Yu, Wenqi Shao, and Qiaosheng
620 Zhang. 2025. [Cpgd: Toward stable rule-based rein-
621 forcement learning for language models](#). *Preprint*,
622 arXiv:2505.12504. 623

MiniMax, :, Aili Chen, Aonian Li, Bangwei Gong,
624 Binyang Jiang, Bo Fei, Bo Yang, Boji Shan,
625 Changqing Yu, Chao Wang, Cheng Zhu, Chengjun
626 Xiao, Chengyu Du, Chi Zhang, Chu Qiao, Chun-
627 hao Zhang, Chunhui Du, Congchao Guo, and 109
628 others. 2025. [Minimax-m1: Scaling test-time com-
629 pute efficiently with lightning attention](#). *Preprint*,
630 arXiv:2506.13585. 631

Arnab Nilim and Laurent El Ghaoui. 2005. [Robust
632 control of markov decision processes with uncertain
633 transition matrices](#). *Oper. Res.*, 53:780–798. 634

Kishan Panaganti, Zaiyan Xu, Dileep Kalathil, and
635 Mohammad Ghavamzadeh. 2022. [Robust rein-
636 forcement learning using offline data](#). *Preprint*,
637 arXiv:2208.05129. 638

Roko Parać, Lorenzo Nodari, Leo Ardon, Daniel
639 Furelos-Blanco, Federico Cerutti, and Alessandra
640 Russo. 2024. [Learning robust reward machines from
641 noisy labels](#). In *Proceedings of the TwentyFirst Inter-
642 national Conference on Principles of Knowledge Rep-
643 resentation and Reasoning, KR-2024*, page 909–919.
644 International Joint Conferences on Artificial Intelli-
645 gence Organization. 646

Long Phan, Alice Gatti, Ziwen Han, Nathaniel Li,
647 Josephina Hu, Hugh Zhang, Chen Bo Calvin Zhang,
648 Mohamed Shaaban, John Ling, Sean Shi, Michael
649 Choi, Anish Agrawal, Arnav Chopra, Adam Khoja,
650 Ryan Kim, Richard Ren, Jason Hausenloy, Oliver
651 Zhang, Mantas Mazeika, and 1090 others. 2025. [Hu-
652 manity’s last exam](#). *Preprint*, arXiv:2501.14249. 653

John Schulman, Philipp Moritz, Sergey Levine, Michael
654 Jordan, and Pieter Abbeel. 2018. [High-dimensional
655 continuous control using generalized advantage esti-
656 mation](#). *Preprint*, arXiv:1506.02438. 657

658	Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu,	Rui Zheng, Hang Yan, Tao Gui, and 2 others. 2025a.	715
659	Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan	Bapo: Stabilizing off-policy reinforcement learning	716
660	Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024.	for llms via balanced policy optimization with adap-	717
661	Deepseekmath: Pushing the limits of mathemati-	clipping. <i>Preprint</i> , arXiv:2510.18927.	718
662	cal reasoning in open language models. <i>Preprint</i> ,		
663	arXiv:2402.03300.		
664	Laixi Shi and Yuejie Chi. 2023. Distributionally	Zhiheng Xi, Jixuan Huang, Chenyang Liao, Baodai	719
665	robust model-based offline reinforcement learning	Huang, Honglin Guo, Jiaqi Liu, Rui Zheng, Junjie	720
666	with near-optimal sample complexity. <i>Preprint</i> ,	Ye, Jiazheng Zhang, Wenxiang Chen, Wei He, Yi-	721
667	arXiv:2208.05767.	wen Ding, Guanyu Li, Zehui Chen, Zhengyin Du,	722
668	Yi Su, Dian Yu, Linfeng Song, Juntao Li, Haitao	Xuesong Yao, Yufei Xu, Jiecao Chen, Tao Gui, and	723
669	Mi, Zhaopeng Tu, Min Zhang, and Dong Yu. 2025.	4 others. 2025b. Agentgym-rl: Training llm agents	724
670	Crossing the reward bridge: Expanding rl with ver-	for long-horizon decision making through multi-turn	725
671	ifiable rewards across diverse domains. <i>Preprint</i> ,	reinforcement learning. <i>Preprint</i> , arXiv:2509.08755.	726
672	arXiv:2503.23829.		
673	M-A-P Team, Xinrun Du, Yifan Yao, Kaijing Ma, Bingli	An Yang, Anpeng Li, Baosong Yang, Beichen Zhang,	727
674	Wang, Tianyu Zheng, Kang Zhu, Minghao Liu, Yim-	Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao,	728
675	ing Liang, Xiaolong Jin, Zhenlin Wei, Chujie Zheng,	Chengen Huang, Chenxu Lv, Chujie Zheng, Day-	729
676	Kaixing Deng, Shuyue Guo, Shian Jia, Sichao Jiang,	iheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao	730
677	Yiyan Liao, Rui Li, Qinrui Li, and 76 others. 2025.	Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41	731
678	Supergpqa: Scaling llm evaluation across 285 gradu-	others. 2025. Qwen3 technical report. <i>Preprint</i> ,	732
679	ate disciplines. <i>Preprint</i> , arXiv:2502.14739.	arXiv:2505.09388.	733
680	Yue Wang, Alvaro Velasquez, George Atia, Ashley	Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan,	734
681	Prater-Bennette, and Shaofeng Zou. 2023. Robust	Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan,	735
682	average-reward markov decision processes. <i>Preprint</i> ,	Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin,	736
683	arXiv:2301.00858.	Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan	737
684	Jason Wei, Nguyen Karina, Hyung Won Chung,	Tong, Chi Zhang, Mofan Zhang, Wang Zhang, and	738
685	Yunxin Joy Jiao, Spencer Papay, Amelia Glaese,	16 others. 2025. Dapo: An open-source llm re-	739
686	John Schulman, and William Fedus. 2024. Mea-	inforcement learning system at scale. <i>Preprint</i> ,	740
687	suring short-form factuality in large language models.	arXiv:2503.14476.	741
688	<i>Preprint</i> , arXiv:2411.04368.		
689	Liang Wen, Yunke Cai, Fenrui Xiao, Xin He, Qi An,	Moritz A. Zanger, Wendelin Böhmer, and Matthijs	742
690	Zhenyu Duan, Yimin Du, Junchen Liu, Lifu	T. J. Spaan. 2025. Diverse projection ensembles	743
691	Tang, Xiaowei Lv, Haosheng Zou, Yongchao Deng,	for distributional reinforcement learning. <i>Preprint</i> ,	744
692	Shousheng Jia, and Xiangzheng Zhang. 2025. Light-	arXiv:2306.07124.	745
693	rl: Curriculum sft, dpo and rl for long cot from	Jin Peng Zhou, Kaiwen Wang, Jonathan Chang, Zhaolin	746
694	scratch and beyond. <i>Preprint</i> , arXiv:2503.10460.	Gao, Nathan Kallus, Kilian Q. Weinberger, Kianté	747
695	Zhiheng Xi, Wenxiang Chen, Boyang Hong, Senjie	Brantley, and Wen Sun. 2025. q#: Provably opti-	748
696	Jin, Rui Zheng, Wei He, Yiwen Ding, Shichun Liu,	mal distributional rl for llm post-training. <i>Preprint</i> ,	749
697	Xin Guo, Junzhe Wang, Honglin Guo, Wei Shen,	arXiv:2502.20548.	750
698	Xiaoran Fan, Yuhao Zhou, Shihan Dou, Xiao Wang,	A Additional Details for VRPO	751
699	Xinbo Zhang, Peng Sun, Tao Gui, and 2 others. 2024a.	A.1 Pseudocode	752
700	Training large language models for reasoning through	The full algorithm of VRPO is detailed in Algo-	753
701	reverse curriculum reinforcement learning. <i>Preprint</i> ,	gorithm 1.	754
702	arXiv:2402.05808.		
703	Zhiheng Xi, Yiwen Ding, Wenxiang Chen, Boyang	B Mathematical Analysis of Stability and	755
704	Hong, Honglin Guo, Junzhe Wang, Dingwen Yang,	Generalization in Robust Bellman PPO	756
705	Chenyang Liao, Xin Guo, Wei He, Songyang Gao,	and DVPO	757
706	Lu Chen, Rui Zheng, Yicheng Zou, Tao Gui,	This section provides a rigorous mathematical	758
707	Qi Zhang, Xipeng Qiu, Xuanjing Huang, Zuxuan	derivation of the training dynamics for Standard	759
708	Wu, and Yu-Gang Jiang. 2024b. Agentgym: Evolv-	PPO, Robust Bellman PPO, and our DVPO frame-	760
709	ing large language model-based agents across diverse	work. We demonstrate why the pessimistic contrac-	761
710	environments. <i>Preprint</i> , arXiv:2406.04151.	tion of Robust Bellman PPO leads to stability at	762
711	Zhiheng Xi, Xin Guo, Yang Nan, Enyu Zhou, Junrui	the cost of generalization, and how DVPO’s asym-	763
712	Shen, Wenxiang Chen, Jiaqi Liu, Jixuan Huang, Zhi-	metric objective, validated by Figure 3, achieves a	764
713	hao Zhang, Honglin Guo, Xun Deng, Zhikai Lei,	principled balance.	765
714	Miao Zheng, Guoteng Wang, Shuo Zhang, Peng Sun,		

Algorithm 1: DVPO Training Algorithm

Require: Dataset \mathcal{D} , policy π_θ , value network $\{f_{\phi_i}\}_{i=1}^N$
Ensure: Optimized π_θ and $\{f_{\phi_i}\}$

- 1: **for** each iteration **do**
- 2: **Sample trajectories** using π_θ
- 3: **Distributional Value Estimation:**
- 4: **for** each (s_t, a_t) **do**
- 5: $h_t = \text{Encoder}(s_t)$
- 6: $\theta_j^{\text{ens}} = \frac{1}{N} \sum_{i=1}^N f_{\phi_i}(h_t)$
- 7: **end for**
- 8: **Compute Distributional GAE:**
- 9: $\Theta_{A_t} = (r_t + \gamma \Theta_{V_{t+1}} - \Theta_{V_t}) + \gamma \lambda \Theta_{A_{t+1}}$
- 10: $A_t = \frac{1}{M} \sum_j \Theta_{A_{t,j}}$
- 11: **Value Model Update:**
- 12: $\mathcal{L}_{\text{QR}} = \mathbb{E}[\frac{1}{M} \sum_j |\tau_j - I(u_j < 0)| L_\delta(u_j)]$
- 13: $\mathcal{L}_{\text{Risk}} = \mathbb{E}[\frac{1}{M} \sum_j w_{\text{risk}}(\tau_j)$
- 14: $|\tau_j - I(u_j < 0)| L_\delta(u_j)]$
- 15: $\mathcal{L}_{\text{CVaR}} = \mathbb{E}[(\frac{1}{K_\alpha} \sum_{j=1}^{K_\alpha} \theta'_{\pi(j)} - \theta_{\pi(j)})^2]$
- 16: $\mathcal{L}_{\text{Gain}} = \mathbb{E}[(\frac{1}{K_\beta} \sum_{j=M-K_\beta+1}^M \theta'_{\pi(j)} - \theta_{\pi(j)})^2]$
- 17: $\mathcal{L}_{\text{Shift}} = \mathbb{E}[\text{ReLU}(\mathbb{E}[\theta'_j] - \mathbb{E}[\theta_j])]$
- 18: $\mathcal{L}_{\text{Shape}} = \mathbb{E}[\text{ReLU}(\text{Var}(\Theta; \mathcal{I}_\alpha) - \text{Var}(\Theta'; \mathcal{I}_\alpha))$
- 19: $+ \text{ReLU}(\text{Var}(\Theta'; \mathcal{I}_\beta) - \text{Var}(\Theta; \mathcal{I}_\beta))]$
- 20: $\mathcal{L}_{\text{Curv}} = \mathbb{E}[\text{ReLU}(\mathbb{E}_{j \in \mathcal{I}_\alpha} [\Delta^2 \theta_j])$
- 21: $+ \text{ReLU}(-\mathbb{E}_{j \in \mathcal{I}_\beta} [\Delta^2 \theta_j])]$
- 22: $\mathcal{L}_{\text{Consist}} = \frac{1}{N_{\text{pair}}} \sum_{i < j} \mathbb{E}[\|\Theta^{(i)} - \Theta^{(j)}\|_2^2]$
- 23: $\mathcal{L}_{\text{Critic}} = \sum_k w_k \mathcal{L}_k$
- 24: Update $\{\phi_i\}$ with $\nabla \mathcal{L}_{\text{Critic}}$
- 25: **Policy Update:**
- 26: $\mathcal{L}_{\text{PPO}} = \mathbb{E}[\min(\rho_t A_t, \text{clip}(\rho_t, 1 - \epsilon, 1 + \epsilon) A_t)]$
- 27: Update θ with $\nabla \mathcal{L}_{\text{PPO}}$
- 28: **end for**

B.1 Instability of Standard PPO and Symmetric Contraction of Robust Bellman PPO

The critic $V_\phi(s)$ in a standard PPO framework is trained to minimize the Mean Squared Error loss with respect to a return target Y_t , $\mathcal{L}_{\text{Critic}}^{\text{PPO}}(\phi) = \mathbb{E}_t[(Y_t - V_\phi(s_t))^2]$. When supervision is noisy, the target Y_t (e.g., $r_t + \gamma V_\phi(s_{t+1})$) becomes a high-variance random variable. This high target variance, $\text{Var}(Y_t)$, induces high-variance gradients $\nabla_\phi \mathcal{L}_{\text{Critic}}^{\text{PPO}}$, destabilizing the optimization process.

A Robust Bellman PPO baseline, as described in our implementation and motivated by robust MDP theory, adopts a pessimistic formulation to ensure stability. It defines the robust value estimate as the minimum over all heads, $V_{\text{robust}}(s) = \min_i V_{\phi,i}(s)$. The min operator functions as a non-linear filter that is inherently robust to upper-tail noise (i.e., anomalously high reward signals). This operation dramatically reduces the target variance, $\text{Var}(Y_t^{\text{Robust}}) \ll \text{Var}(Y_t^{\text{PPO}})$, yielding low-variance gradients and a stable training loss.

However, this stability comes at a high cost. The min operator is a symmetric contractor in effect,

it discards all information not related to the worst-case outcome. This mathematical collapse is empirically verified by Figure 3 in the main text. As shown, the Robust Bellman PPO value distribution (red) collapses toward its lower bound and is concentrated with small variance. By discarding all optimistic signals, this method mathematically precludes the policy from learning about high-potential opportunities, thereby severely degrading generalization.

B.2 Asymmetric Regulation for Balanced Optimization

DVPO resolves this dilemma through its Conditional Risk-aware Tail-Constrained (CRTC) objective. This mechanism is mathematically grounded in the ReLU-based tail shape losses, which create one-way optimization penalties to regulate the tails asymmetrically.

B.2.1 Robustness via Lower-Tail Contraction

Robustness against pessimistic noise is achieved via the lower-tail shape loss:

$$\mathcal{L}_{\text{Shape}}^{\text{lower}} = E_{s,a}[\text{ReLU}(\text{Var}(\Theta; \mathcal{I}_\alpha) - \text{Var}(\Theta'; \mathcal{I}_\alpha))] \quad (12)$$

Let $V_{\text{pred}}^{\text{lower}} = \text{Var}(\Theta; \mathcal{I}_\alpha)$ be the predicted lower-tail variance and $V_{\text{target}}^{\text{lower}} = \text{Var}(\Theta'; \mathcal{I}_\alpha)$ be the target. The loss is $\mathcal{L} = E[\text{ReLU}(V_{\text{pred}}^{\text{lower}} - V_{\text{target}}^{\text{lower}})]$. The gradient with respect to the predicted variance, $\frac{\partial \mathcal{L}}{\partial V_{\text{pred}}^{\text{lower}}}$, is key. If $V_{\text{pred}}^{\text{lower}} > V_{\text{target}}^{\text{lower}}$, the ReLU is active, yielding a gradient of +1. Gradient descent applies a penalty to decrease $V_{\text{pred}}^{\text{lower}}$. Conversely, if $V_{\text{pred}}^{\text{lower}} \leq V_{\text{target}}^{\text{lower}}$, the gradient is 0, and no penalty is applied. This creates a one-way upper bound, enforcing $V_{\text{pred}}^{\text{lower}} \leq V_{\text{target}}^{\text{lower}}$. This mathematical contraction of the lower tail filters pessimistic noise. This exact outcome is confirmed by Figure 3, which shows the DVPO distribution (green) also contracts slightly in the lower tail, achieving the desired robustness.

B.2.2 Generalization via Upper-Tail Expansion

Generalization is preserved via the upper-tail shape loss, which inverts the arguments:

$$\mathcal{L}_{\text{Shape}}^{\text{upper}} = E_{s,a}[\text{ReLU}(\text{Var}(\Theta'; \mathcal{I}_\beta) - \text{Var}(\Theta; \mathcal{I}_\beta))] \quad (13)$$

Let $V_{\text{pred}}^{\text{upper}} = \text{Var}(\Theta; \mathcal{I}_\beta)$ and $V_{\text{target}}^{\text{upper}} = \text{Var}(\Theta'; \mathcal{I}_\beta)$. The loss is $\mathcal{L} = E[\text{ReLU}(V_{\text{target}}^{\text{upper}} - V_{\text{pred}}^{\text{upper}})]$. The gradient $\frac{\partial \mathcal{L}}{\partial V_{\text{pred}}^{\text{upper}}}$ is asymmetric in the opposite direction. If $V_{\text{pred}}^{\text{upper}} < V_{\text{target}}^{\text{upper}}$, the

ReLU is active, yielding a gradient of -1 . Gradient descent (moving in the negative gradient direction, i.e., $-(-1) = +1$) will increase V_{pred}^{upper} toward the target. Conversely, if $V_{pred}^{upper} \geq V_{target}^{upper}$, the gradient is 0, and no penalty is applied. This creates a one-way lower bound, enforcing $V_{pred}^{upper} \geq V_{target}^{upper}$.

This upper-tail expansion preserves the variance of high-value outcomes, preventing the policy from becoming overly conservative. Figure 3 provides direct empirical validation of this mathematical mechanism. In stark contrast to Robust Bellman PPO, the DVPO distribution (green) achieves significant exploratory expansion learning in the upper tail. This visual evidence confirms that our asymmetric ReLU loss successfully translates from mathematical theory to practical effect, achieving the principled balance of robustness and generalization that Robust Bellman PPO’s symmetric contraction fails to capture.

C Additional Experimental Details

C.1 Setup

C.1.1 Dataset Construction

Math and scientific tasks For mathematical tasks, taking the Qwen3-8B model as an illustrative example, pseudo-labels are generated from 39,000 samples in the Light-R1 dataset (Wen et al., 2025). Following 5 rounds of majority voting, samples with at least 3 identical responses are retained, resulting in 31,209 instances for reinforcement learning (RL) training. For scientific tasks, pseudo-labels are generated from 26,529 samples in the SuperGPQA dataset. After the same 5-round majority voting process, 10,075 samples with at least 2 consistent responses are preserved for RL training. To ensure fairness, evaluation is performed on mathematical benchmarks that do not overlap with the training datasets.

Real-dialogue tasks A real-world dialogue dataset Honor-Dialogue is used. The reward model is trained on 36,000 labeled samples, with 3,000 reserved for validation. The policy model is fine-tuned and subsequently trained via reinforcement learning on an additional 50,952 conversations, excluding data from the financial services, social entertainment, transportation, and healthcare domains. Performance evaluation is conducted across five scenarios: 1,004 conversations for daily services, 700 for financial services, 534 for social entertainment, 463 for transportation, and 380 for

healthcare. This dataset encompasses diverse dialogue tasks in real-life scenarios, accurately reflecting practical real-world conditions. Notably, **such multi-category real-scenario dialogue task data is not available in other datasets.**

Honor-Dialogue dataset is constructed based on a goal-driven scenario-oriented design paradigm. Specifically, we first defined the core role of the AI dialogue agent, a practical, scenario-adaptive assistant, and established strict constraints on response accuracy, naturalness, and compliance with real-world service norms. Subsequently, we selected six representative domains including one general service domain and five specific professional domains. These domains are derived from high-frequency real-life interaction scenarios, each equipped with explicit trigger conditions such as user query intents and scenario context features, as well as standardized response rules such as information provision priority and dialogue flow guidelines.

For the data within each domain, we construct realistic caller inputs that include scenario-matched latest messages and conversation history, as well as corresponding standardized outputs that comply with the requirements of the target goal. We explicitly mark the dialogue state, response content, and matched target to ensure the quality of the supervised training data.

A representative example of the Honor-Dialogue dataset is presented below. (Figure 5) The notation xxx denotes the masking of sensitive information such as user ID and contact details, which is implemented to comply with data privacy regulations.

C.1.2 Training Configuration

In majority voting sampling, 5 samples are used, and data with more than 2 votes are filtered, with a sampling temperature of 0.7. The loss weights are set as $w_{risk} = 0.3$, $w_{CVaR} = 0.5$, $w_{gain} = 0.3$, $w_{shift} = 0.2$, $w_{shape} = 0.1$, $w_{curv} = 0.2$, and $w_{consist} = 0.1$. For the value model, the risk threshold is set to 0.1, the number of interval quantile points is 200, and the number of distributional value heads is 3. The reinforcement learning training runs for 1 iteration. During training, the model generates dialogue tasks of length 4096 for reasoning tasks and 1024 for dialogue tasks. The experiments are conducted on $8 \times$ NVIDIA A100 80GB GPUs.

Scenario	Wealth	Rental	Insurance	Food	Express	Promotion	Loan	Housing	Service	Product	General	Avg
Dialogue Count	87	99	138	120	215	66	70	87	67	69	92	94.73
Avg Turns	5.40	4.22	5.22	3.37	3.76	4.55	4.33	4.68	5.46	4.44	3.46	4.44

Table 6: Dataset statistics across 11 real-world service scenarios. The validation set covers diverse interaction types with varying dialogue lengths, enabling reliable evaluation of robustness and generalization under realistic settings.

Category	Human Annotators				GPT Models		Human AVG	Model AVG
Score	4.595	4.689	4.490	4.330	4.510	4.501	4.526	4.505

Table 7: Comparison of goal completion rate scores between human annotators and GPT-4o. Human annotators achieve an average score of 4.526, while GPT-4o models reach 4.505, showing a negligible performance gap.

C.1.3 Dialogue Task Evaluation

Our dialogue evaluation leverages GPT-4o, but the model is not utilized as a free-form judge. Instead, it implements a rigorous rubric-based evaluation protocol designed to mitigate subjectivity and enhance reproducibility. This rubric explicitly defines evaluation metrics, adopts a 1–5 scoring scale with clear criteria, includes two scoring examples and a standardized output format, and the prompt incorporates a well-structured process along with comprehensive dialogue content and contextual information. GPT-4o functions solely as an automated evaluator applying this fixed rubric, rather than an unconstrained scorer.

To further ensure reliability, we conducted cross-validation through two key steps: **1)** multiple sampling runs to verify the stability of rubric execution, and **2)** spot-checked human evaluations that demonstrated high agreement with rubric-based scores. Specifically, for the comparative experiment between human annotators and the model, we validated the protocol using 1,110 data samples covering 10 scenarios. The statistical details of the validation dataset are presented in Table 6. Four independent annotators scored these samples strictly in accordance with the rubric. They are professional data annotators in the company. Table 7 presents the experimental results for goal completion rate scoring (full score: 5) in dialogue tasks. As shown in Table 7, the mean score assigned by human annotators is extremely close to that of the model, with a difference of only 0.021 points. These findings confirm strong consistency between human and model scores, validating the effectiveness of our rubric. We acknowledge that fully establishing external validity necessitates additional independent annotators. We are in the process of releasing our rubric, evaluation prompts, and real-dialogue dataset to facilitate replication.

To illustrate the evaluation logic for dialogue performance, the core prompt section regarding the assessment of conversation logicity is provided as follows in Figure 6, which is a part of goal completion rate.

C.2 Additional experimental details in the Dialogue Task

Table 8 presents the comprehensive results across all five real-world dialogue domains. DVPO achieves the highest overall performance with an average score of 86.79%, outperforming all RL baselines. Specifically, DVPO surpasses PPO and Reinforce++ by 1.59% and 2.94% percentage points on average, and exceeds GRPO-based methods by more than 20 percentage points in most domains. This advantage is consistent across all metrics: task completion, ask completion, and goal completion.

A detailed analysis of domain-level results reveals the stability of DVPO under in domain and out of domain tasks. In Life Services and Transportation & Travel, DVPO achieves domain averages exceeding 88%, outperforming PPO by approximately 6 percentage points and Reinforce++ by 3–5 percentage points. In the more challenging Healthcare & Wellness and Social & Entertainment domains, DVPO maintains an accuracy of 87–88%, whereas PPO exhibits fluctuations of up to 10 percentage points and GRPO deteriorates to the 30–50% range. Even in the most demanding Financial Services domain, DVPO remains above 85%, limiting the performance drop to within 3 percentage points, while GRPO variants degrade by more than 20 percentage points.

The cross-domain results further underscores DVPO’s robustness and generalization. These results demonstrate that DVPO provides a reliable and noise-resilient reinforcement learning strategy for real-world dialogue systems with imperfect su-

Method	Life Services (In-Domain)				Transportation & Travel				Healthcare & Wellness				Social & Entertainment				Financial Services				AVG
	TC	AC	GC	D-AVG	TC	AC	GC	D-AVG	TC	AC	GC	D-AVG	TC	AC	GC	D-AVG	TC	AC	GC	D-AVG	
Baseline	72.5	95.0	92.7	86.73	66.7	96.5	90.3	84.50	77.2	97.8	95.7	90.23	75.3	92.0	94.1	87.13	65.7	91.4	91.0	82.70	86.26
GRPO	42.7	16.3	31.1	30.03	34.6	20.6	30.5	28.57	33.3	19.4	32.3	28.33	46.4	13.1	27.2	28.90	49.1	13.3	21.3	27.90	28.75
PPO	69.0	93.1	90.5	84.20	71.4	96.0	90.8	86.07	82.5	93.1	97.0	90.87	67.0	87.9	88.1	81.00	69.2	96.3	86.1	83.87	85.20
Reinforce++	73.9	96.5	93.2	87.87	58.7	96.8	87.7	81.07	47.9	95.3	89.8	77.67	70.3	98.6	92.5	87.13	68.8	96.5	91.3	85.53	83.85
Dr.GRPO	37.1	51.8	66.1	51.67	39.7	68.0	73.2	60.30	33.3	75.4	75.4	61.37	40.6	52.6	63.3	52.17	39.6	63.7	66.1	56.47	56.39
DVPO (Ours)	76.5	95.7	92.2	88.13	75.6	95.8	91.8	87.73	73.2	96.9	92.9	87.67	76.3	94.5	92.2	87.67	66.7	90.7	90.8	82.73	86.79

Table 8: Full results on the Real-world Dialogue Benchmark under noisy reward supervision. For each domain, we report task completion rate (TC), ask completion rate (AC), and goal completion rate (GC); DomainAVG (D-AVG) denotes the averaged performance across the three metrics within each domain, and AVG denotes the overall average across all five domains. DVPO achieves the highest overall average performance 86.79%, consistently surpassing all prior RL baselines in both in-domain and out-of-domain settings.

Domain	In-Domain (Science & QA)			Out-of-Domain (Math)				ID	OOD	ALL	
	Method	SampleQA	GPQA(ALL)	HLE	MATH500	AIME24	Minerva-Math				AMC23
Base		1.78	2.36	2.87	84.80	35.00	19.49	69.17	2.34	52.12	30.78
GRPO		1.60	1.99	3.94	82.20	33.33	21.32	77.50	2.51	53.59	31.70
PPO		1.91	2.54	3.94	83.80	33.33	22.43	79.17	2.80	54.68	32.45
Reinforce++		1.48	2.36	3.52	84.60	38.33	23.16	75.83	2.45	55.48	32.75
Dr.GRPO		1.50	2.17	3.10	84.00	35.00	22.79	76.67	2.26	54.62	32.18
Robust Bellman		1.64	2.17	4.17	84.60	31.67	22.79	75.00	2.66	53.52	31.72
DVPO (Ours)		1.73	2.36	4.36	86.20	43.33	25.00	78.33	2.82	58.22	34.47

Table 9: Comparison of DVPO performance on the smaller Qwen3-1.7B model across multiple domains. DVPO achieved the best performance in math tasks, and also achieved the optimal accuracy rate of 4.36% on the HLE dataset.

pervision.

C.3 Additional Ablation Experiment

DVPO exhibits strong scalability across model scales

As shown in Table 9, DVPO consistently outperforms PPO, GRPO, and Reinforce++ even when applied to the smaller Qwen3-1.7B model. Despite the reduced capacity, it achieves an average accuracy of 58.22% on mathematical tasks and 2.82% on scientific tasks, maintaining stable optimization under noisy rewards. These results confirm that DVPO remains effective across different model scales, demonstrating robust generalization and adaptability under noisy training conditions.

Effectiveness of Distributional Constraints on Robustness and Generalization

Table 10 presents an ablation study of different distributional loss components in DVPO. Starting from the core quantile regression baseline, we observe that introducing additional distributional constraints consistently improves performance across most settings. Among all components, tail-related objectives contribute the most significant gains. While distribution consistency and tail calibration improve robustness to some extent, their effects are more pronounced on out-of-domain tasks, especially on challenging math benchmarks such as AIME24 and AMC23. Notably, incorporating tail shape and curvature modeling leads to the most

consistent improvements across both in-domain and out-of-domain evaluations, achieving the highest overall accuracy of 39.63%. These results indicate that explicitly modeling the shape of the value distribution plays a critical role in improving generalization. Compared to variance-reduction or calibration-based objectives alone, shape-aware constraints provide stronger structural regularization, enabling DVPO to better capture distributional properties under noisy or shifted reward settings.

C.4 Additional Visualization

Figures 7, 8, and 9 illustrate a comparative analysis of advantage estimations among PPO, Robust Bellman PPO, and our proposed method (DVPO) for the same question outputs, irrespective of their semantic quality. Notably, significant discrepancies in advantage estimation patterns are observed across the three methods. For PPO, despite correctly identifying the final output as erroneous, the advantage values assigned to individual tokens exhibit minimal variability. Similarly, Robust Bellman PPO fails to recognize the invalidity of the novel output and also demonstrates negligible differences in token-level advantage estimations. Additionally, it overestimates the advantage of the final output component. This observation partially reflects the limited generalization capability of these two methods when handling novel

Domain	In-Domain (Science & QA)			Out-of-Domain (Math)				ID	OOD	ALL
Method	SampleQA	GPQA (ALL)	HLE	MATH500	AIME24	Minerva-Math	AMC23	AVG	AVG	AVG
Base	2.89%	3.10%	2.89%	87.40%	41.67%	28.68%	75.83%	2.96%	58.40%	34.64%
Core Quantile Regression	2.94%	4.17%	3.85%	90.80%	45.00%	30.51%	79.17%	3.65%	61.37%	36.63%
Above+Distribution Consistency	2.87%	4.53%	3.80%	90.40%	51.67%	28.68%	84.17%	3.73%	63.73%	38.02%
Above+Tail Calibration	2.77%	4.89%	3.57%	89.40%	51.67%	31.62%	85.00%	3.74%	64.42%	38.42%
Above+Shift Penalization	2.94%	3.44%	3.48%	91.20%	51.67%	31.99%	85.83%	3.29%	65.17%	38.65%
Above+Tail Shape & Curvature(OURS)	3.21%	4.71%	3.57%	90.60%	56.67%	31.99%	86.67%	3.83%	66.48%	39.63%

Table 10: Ablation study on the effect of different loss components in DVPO. Each loss is added on the above experiment. All components provide positive contributions, and performance improves steadily as more distributional constraints are incorporated. Tail-focused objectives yield the most significant gains, Tail Calibration improves OOD accuracy from 61.37% to 64.42%, while adding Tail Shape and Curvature further enhances both ID and OOD performance, achieving the best overall result of 39.63%.

1069 response generations.

1070 In contrast, our DVPO method accurately judges
1071 the correctness of the final output. More impor-
1072 tantly, its advantage estimation effectively captures
1073 key semantic terms (e.g., "nucleus" and "quarks")
1074 by assigning distinct advantage weights to these
1075 critical tokens, which aligns with the core semantic
1076 characteristics of the target task.

System Prompt

Role: You are the AI call assistant of xxx Inc., capable of answering calls clearly and politely on behalf of the user and conversing with the caller. Based on the user’s customized dialogue goals, you should first determine which goal applies to the conversation and respond accordingly. You do not know any personal information about the user beyond the task description and must inform the caller that any information they provide will be passed on to the user.

Dialogue Goals (Customized):

- **0: General Domain** – Used if the call doesn’t match any specific domain. Be polite and professional. Try to answer or guide the caller and clarify their intent.
- **1: E-commerce/Retail** – If caller mentions down jacket performance, mention fast delivery, good warmth at -5°C, ask for cleaning tips, other styles, hotline, or exchange policy, inquire about washing temperature.
- **2: Technical Support** – If license issues arise for WX3510/WX5540, answer related questions, confirm email and phone, send manuals, ask for model and associated phone.
- **3: Financial Services – Pension** – If retirement certification reminder, say not remembered, will handle via WeChat, thank them, ask about past communication and WeChat face verification method.
- **4: Life Services** – If laundry fee update, respond with satisfaction, hope for discount, mention night service needs, ask about pricing explanation, training quality, and monthly discounts.
- **5: Financial Services – Trade Notice** – If transaction confirmation, say user unavailable. If medium-risk, ask for analyst opinion, arrange follow-up next day, inquire about transaction fees.

Caller History and New Input

Conversation History: (*Empty*)

Caller New Input: Hello, this is the Forex Trading Customer Service Center. We are confirming that the currency pair you traded is USD/SGD, and the rate information is ready.

Model Response

Target Domain: 5: *Financial Services – Trade Notice*

Dialogue State: *Ongoing*

Response: Hello, this call is currently being handled by the smart assistant. The user is temporarily unavailable. If this is a medium-risk transaction, could we get advice from a market analyst? Please contact us again tomorrow for confirmation.

Figure 5: A multi-turn example from the Honor-Dialogue dataset. The dataset features realistic, task-oriented, multi-domain conversations, in which each model response includes structured states. This example represents the situation of the first conversation.

Core Prompt for Dialogue Logicality

When scoring the intelligent call assistant's responses, please follow these steps:

1. Read and understand the entire call content.

2. Score from the dimension of dialogue logic rationality. The definition of dialogue logic rationality is: whether all responses provided by the assistant during the entire dialogue are clear, coherent, and effectively convey relevant information. A reasonable response should ensure the consistency, coherence, and relevance of information. Use a 1-5 scoring scale with specific criteria as follows:

- **1 point:** The response exhibits severe logical inconsistencies and violates fundamental reasoning principles. The intent of the assistant is unclear, and the response fails to convey meaningful or usable information.
- **2 points:** The response contains multiple inconsistencies or conflicting statements. Although partially related to the query, it lacks sufficient contextual grounding and coherent reasoning, resulting in poor interpretability and weak informational value.
- **3 points:** The response is generally coherent but contains minor logical gaps or discontinuities that affect fluency. While relevant to the query, the explanation lacks depth or clarity in parts, limiting the overall effectiveness of communication.
- **4 points:** The response is logically consistent and well-structured, with clear and relevant content. Information is conveyed effectively and aligns well with the user's intent, enabling smooth and coherent interaction.
- **5 points:** The response demonstrates strong logical consistency, clear reasoning, and precise information delivery. It not only addresses the query accurately but also facilitates deeper understanding, effectively guiding the conversation and enhancing overall interaction quality.

3. Based on the above scoring criteria and combined with the entire dialogue content, give a reasonable score, along with the scoring thinking process, deduction points, and modification suggestions to improve the score without changing the original meaning of the text.

Figure 6: A core prompt for dialogue logicality assessment from the constructed rubric evaluation method. The rubric features clear scoring steps, explicit definition of logical rationality, and detailed 1-5 point end-point criteria, in which each evaluation requires supplementary scoring reasoning, deduction explanations, and optimization suggestions.

```
user A common approximation made in many-body nuclear calculations of some interaction (for example an electroweak interaction) is to suppose that nucleons in a nucleus behave in the same way as free nucleons. Which of the following four assumptions jointly imply this "impulse" approximation? 1. The interaction current only interacts with individual nucleons. 2. The nucleus is transparent apart from the selected nucleon. 3. The quarks internal to the selected nucleon are static. 4. The interaction proceeds as if the selected nucleon experiences no binding forces. assistant The impulse approximation in nuclear physics is a method used to simplify the treatment of interactions between external currents and nucleons in a nucleus. It assumes that the interaction is localized to a single nucleon and that the other nucleons do not interfere. This implies that the selected nucleon behaves as if it is free from the binding forces of the nucleus during the interaction. Let's analyze the four assumptions. 1. **The interaction current only interacts with individual nucleons.** This assumption is correct. It implies that the interaction is not with the entire nucleus, but rather with one nucleon at a time. This is consistent with the impulse approximation, where the interaction is localized to a single particle. 2. **The nucleus is transparent apart from the selected nucleon.** This assumption is also correct. It means that the other nucleons do not interact with the external current, effectively isolating the selected nucleon. This is a key feature of the impulse approximation, as it ensures the interaction is only with one nucleon. 3. **The quarks internal to the selected nucleon are static.** This assumption is incorrect. While it may simplify the treatment of the nucleon, it is not a necessary assumption for the impulse approximation. The impulse approximation focuses on the interaction with the nucleon, not its internal structure. 4. **The interaction proceeds as if the selected nucleon experiences no binding forces.** This assumption is correct. It directly implies that the nucleon is treated as free during the interaction, which is a core assumption of the impulse approximation. Based on this analysis, the assumptions that jointly imply the impulse approximation are **Option 2** and **Option 4**. The answer is  $\boxed{2,4}$ .
```

Advantage Scale:

-5.109

3.592 (Avg)

18.375

Figure 7: Advantage estimation visualization of the PPO method for the target question output. The method correctly identifies the final output as erroneous, while the token-level advantage values show minimal variability.

user A common approximation made in many-body nuclear calculations of some interaction (for example an electroweak interaction) is to suppose that nucleons in a nucleus behave in the same way as free nucleons. Which of the following four assumptions jointly imply this "impulse" approximation? 1. The interaction current only interacts with individual nucleons. 2. The nucleus is transparent apart from the selected nucleon. 3. The quarks internal to the selected nucleon are static. 4. The interaction proceeds as if the selected nucleon experiences no binding forces. assistant The impulse approximation in nuclear physics is a method used to simplify the treatment of interactions between external currents and nucleons in a nucleus. It assumes that the interaction is localized to a single nucleon and that the other nucleons do not interfere. This implies that the selected nucleon behaves as if it is free from the binding forces of the nucleus during the interaction. Let's analyze the four assumptions: 1. **The interaction current only interacts with individual nucleons.** This assumption is correct. It implies that the interaction is not with the entire nucleus, but rather with one nucleon at a time. This is consistent with the impulse approximation, where the interaction is localized to a single particle. 2. **The nucleus is transparent apart from the selected nucleon.** This assumption is also correct. It means that the other nucleons do not interact with the external current, effectively isolating the selected nucleon. This is a key feature of the impulse approximation, as it ensures the interaction is only with one nucleon. 3. **The quarks internal to the selected nucleon are static.** This assumption is incorrect. While it may simplify the treatment of the nucleon, it is not a necessary assumption for the impulse approximation. The impulse approximation focuses on the interaction with the nucleon, not its internal structure. 4. **The interaction proceeds as if the selected nucleon experiences no binding forces.** This assumption is correct. It directly implies that the nucleon is treated as free during the interaction, which is a core assumption of the impulse approximation. Based on this analysis, the assumptions that jointly imply the impulse approximation are **Option 2** and **Option 4**. The answer is $\boxed{2,4}$.

Advantage Scale:




Figure 8: Advantage estimation visualization of the Robust Bellman PPO method for the target question output. The method fails to recognize the invalidity of the novel output, with negligible differences in token-level advantage estimations and overestimation of the final output's advantage.

user A common approximation made in many-body nuclear calculations of some interaction (for example an electroweak interaction) is to suppose that nucleons in a nucleus behave in the same way as free nucleons. Which of the following four assumptions jointly imply this "impulse" approximation? 1. The interaction current only interacts with individual nucleons. 2. The nucleus is transparent apart from the selected nucleon. 3. The quarks internal to the selected nucleon are static. 4. The interaction proceeds as if the selected nucleon experiences no binding forces. assistant The impulse approximation in nuclear physics is a method used to simplify the treatment of interactions between external currents and nucleons in a nucleus. It assumes that the interaction is localized to a single nucleon and that the other nucleons do not interfere. This implies that the selected nucleon behaves as if it is free from the binding forces of the nucleus during the interaction. Let's analyze the four assumptions: 1. **The interaction current only interacts with individual nucleons.** This assumption is correct. It implies that the interaction is not with the entire nucleus, but rather with one nucleon at a time. This is consistent with the impulse approximation, where the interaction is localized to a single particle. 2. **The nucleus is transparent apart from the selected nucleon.** This assumption is also correct. It means that the other nucleons do not interact with the external current, effectively isolating the selected nucleon. This is a key feature of the impulse approximation, as it ensures the interaction is only with one nucleon. 3. **The quarks internal to the selected nucleon are static.** This assumption is incorrect. While it may simplify the treatment of the nucleon, it is not a necessary assumption for the impulse approximation. The impulse approximation focuses on the interaction with the nucleon, not its internal structure. 4. **The interaction proceeds as if the selected nucleon experiences no binding forces.** This assumption is correct. It directly implies that the nucleon is treated as free during the interaction, which is a core assumption of the impulse approximation. Based on this analysis, the assumptions that jointly imply the impulse approximation are **Option 2** and **Option 4**. The answer is $\boxed{2,4}$.

Advantage Scale:




Figure 9: Advantage estimation visualization of our proposed DVPO method for the target question output. The method accurately judges the correctness of the final output and effectively captures key semantic terms (e.g., "nucleus" and "quarks") via distinct token-level advantage weights.