COPR: Continual Human Preference Learning via Optimal Policy Regularization

Han Zhang[1,2], **Lin Gui**[3], **Yu Lei**[2], **Yuanzhao Zhai**[4], **Yehong Zhang**[2], **Zhuo Zhang**[1,2], **Yulan He**[3],
Hui Wang[2], **Yue Yu**[2],**Kam-Fai Wong**[5], **Bin Liang**[5*], **and Ruifeng Xu**[1,2,6*]
[1] Harbin Institute of Technology, [2] Peng Cheng Laboratory, [3] King's College London,
[4] National University of Defense Technology, [5] The Chinese University of Hong Kong,
[6] Guangdong Provincial Key Laboratory of Novel Security Intelligence Technologies
`hanlardresearch@gmail.com, xuruifeng@hit.edu.cn`

Abstract

The growing integration of Large Language Models (LLMs) into real-world applications underscores the critical need for continual alignment with evolving human preferences. Reinforcement Learning from Human Feedback (RLHF) has shown success in improving the alignment of LLMs, but its rigid, multi-stage process presents significant limitations for continual learning (CL) scenarios, where models need to adapt incrementally without catastrophic forgetting. Existing methods, such as Direct Preference Optimization (DPO), offer potential for offline preference learning but exhibit challenges like increased preference gap amplification and reduced model diversity, which can lead to preference collapse. In practical settings, LLMs continuously interact with diverse user feedback across tasks and domains. The inability of current approaches to efficiently incorporate incremental human preferences without retraining or significant computational overhead limits their scalability and adaptability. Addressing these gaps, our study introduces a novel framework, Continual Optimal Policy Regularization (COPR), that ensures robust and flexible continual alignment while preserving historical knowledge and optimizing performance in new preference tasks.

1 Introduction

The rapid evolution of artificial intelligence, particularly in Natural Language Processing (NLP), has driven the adoption of Large Language Models (LLMs) across diverse applications. These models hold immense potential for aligning machine-generated outputs with human preferences (Bai et al., 2022; Stiennon et al., 2020; Dai et al., 2023), enabling safer and more effective human-computer interactions. However, real-world deployment scenarios demand more than static alignment; LLMs must continually adapt to evolving user preferences as new tasks and domains emerge (Zhang et al., 2024; Qi et al., 2024). This capability, known as continual alignment (Wu et al., 2024), is critical for practical applications such as AI-driven assistants, where user feedback is dynamic and incremental. The development of robust continual alignment methods can transform how AI systems learn and respond, unlocking their ability to provide long-term, adaptive assistance.

Despite the promise of continual alignment, existing methodologies face significant challenges. Reinforcement Learning from Human Feedback (RLHF) (Ouyang et al., 2022), a widely adopted approach for aligning LLMs with human preferences, relies on a multi-stage process involving supervised fine-tuning, reward model training, and reinforcement learning optimization. This pipeline, while effective, is computationally intensive and inflexible for continual learning (CL) (Zhang et al., 2024). Offline methods like Direct Preference Optimization (DPO) (Rafailov et al., 2023) eliminate the need for reinforcement learning but suffer from issues such as overfitting preference data (Azar et al., 2023) and catastrophic forgetting of past knowledge. In CL scenarios, these methods often struggle to maintain model output diversity, resulting in preference collapse (Xiao et al., 2024), where minority preferences are virtually disregarded. This ultimately leads to degraded performance on historical tasks and reduced adaptability to new ones.

To address these challenges, we introduce Continual Optimal Policy Regularization (COPR), a novel framework that redefines how LLMs adapt to evolving human preferences. COPR bridges the gap between efficiency and effectiveness in continual preference learning by introducing a principled method to preserve historical knowledge while integrating new preferences. At its core, COPR leverages the theoretical optimal policy theory (Peters and Schaal, 2007; Peng et al., 2019), utilizing historical optimal policies as constraints to guide the

---

[*] B.Liang and R.Xu are corresponding authors.

learning of new tasks. This ensures that the model does not overly prioritize new preferences at the expense of forgetting old ones. By deriving a moderate reward function (MRF) from Bradley-Terry model (Bradley and Terry, 1952), COPR maintains a balanced optimization process that avoids excessive determinism or preference collapse (Xiao et al., 2024), ensuring stable and diverse model outputs. The simplicity and theoretical grounding of COPR make it a robust solution for continual alignment.

COPR's methodology incorporates several innovative components. First, it parameterizes the sampling distribution of historical and new policies to construct optimization objectives that maintain diversity and prevent catastrophic forgetting. Second, it employs the MRF to regulate the learning process, ensuring that the relative probabilities of preferred and dispreferred responses remain within a reasonable range. Third, COPR adopts a Lagrangian dual optimization framework to balance the learning of new preferences with the retention of past ones. These innovations collectively address the critical limitations of existing methods, enabling COPR to outperform strong baselines in various CL settings. In summary, our main contributions are as follows:

- We propose COPR, a novel framework for continual preference optimization that eliminates the need for reinforcement learning loops while addressing key limitations of existing offline methods (Section 3).
- We introduce a benchmark for continual alignment tasks and demonstrate COPR's superiority in handling task and domain incremental learning scenarios (Section 4.1).
- We provide theoretical insights and empirical evidence to validate the effectiveness and robustness of COPR in maintaining model diversity, preventing forgetting, and achieving high performance across new and historical tasks. These advancements position COPR as a practical and scalable solution for real-world continual alignment challenges. (Section 4.2 $\sim$ 4.4).

2 Preliminaries and Task Formulation

2.1 Traditional Alignment

Reinforcement Learning from Human Feedback. The recent RLHF pipeline consists of three phases: 1) Supervised Fine-Tuning (SFT) stage trains LLM with maximum likelihood on the downstream tasks. 2) In the preference sampling and RM learning stage, human annotators rank multiple responses $\mathcal{Y}^x = \{y_1^x \prec y_2^x \prec \ldots \prec y_n^x\}$ for a prompt x based on human preferences, as human feedback data. Then, this feedback data is used to train an RM $r_\phi(x,y)$[1] to score the prompt and response pair (x,y). 3) The RL optimization stage maximizes a reverse KL-constrained reward objective like

$$\begin{aligned}\max_{\pi_\theta} \mathbb{E}_{x\sim\mathcal{D},y\sim\pi_\theta(y|x)}\big[r_\phi(x,y)\big]\\ -\beta\mathbb{D}_{\mathrm{KL}}\big[\pi_\theta(y\mid x)\mid\mid\pi_{ref}(y\mid x)\big],\end{aligned} \tag{1}$$

where β is a parameter that controls the degree of deviation from the base reference policy π_{ref}, $x\in\mathcal{D}$ denotes the prompt for RL training, y denotes the possible response, and θ denotes the parameters of LLM. In the most related works (Bai et al., 2022; Ouyang et al., 2022; Stiennon et al., 2020), the reward is reconstructed by subtracting KL-regularization term, namely $r_\phi(x,y) - \beta\log(\pi_\theta(y\mid x)/\pi_{ref}(y\mid x))$, and maximized by PPO (Schulman et al., 2017) directly.

Due to the multiple stages in traditional RLHF, when human preferences are updated, all of the SFT model, reward model, and policy model need updation, which lacks flexibility for CL.

Optimal Policy of RLHF. Previous works Advantage-Weighted Regression (AWR) (Peng et al., 2019) and DPO (Rafailov et al., 2023) derive that the optimal solution $\pi^*(y\mid x)$ to the reverse KL-constrained reward maximization objective in Eq. (1) takes the form:

$$\pi^*(y\mid x) = \frac{1}{Z(x)}\pi_{ref}(y\mid x)\exp(\frac{1}{\beta}r(x,y)), \tag{2}$$

where $Z(x) = \sum_y \pi_{ref}(y\mid x)\exp(\frac{1}{\beta}r(x,y))$ is the partition function of $\pi^*(y\mid x)$. It provides the explicit form of the optimal policy, inspiring us to bypass the complex RLHF process and directly fit the optimal policy. Although a claimed advantage of DPO is its avoidance of the necessity to learn a reward model, the learning objective is to maximize $\hat{r}_\theta(x,y_w) - \hat{r}_\theta(x,y_l) \to +\infty$, where $\hat{r}_\theta(x,y) = \beta\log\pi_\theta(x,y)/\pi_{ref}(x,y)$, where y_w and y_l denote the preferred and dispreferred actions

---

[1] Subscript notations are used to indicate corresponding parameter sets, such as $r_\phi(x,y)$. When parentheses are used in the subscript, as in $r_{(t)}(x,y)$, it signifies the reward associated with a specific task t.

satisfy that $y_l \prec y_w$. The ultimate result of this objective is that $\pi_\theta(x, y_w) \to 1$ and $\pi_\theta(x, y_l) \to 0$, which weakens the strength of KL-regularization and excessively increases the determinism of the LLM policy (Azar et al., 2023).

2.2 Alignment in Continual Learning Setup

We consider that there is a sequence of tasks $\mathbb{T} = \{\mathcal{T}_1, \mathcal{T}_2, ...\}$ to learn, and a sequence of corresponding human preference datasets $\mathbb{D} = \{\mathcal{D}_1, \mathcal{D}_2, ...\}$. The initial policy is the SFT model, namely, $\pi_0 = \pi_{SFT}$. For each task $\mathcal{T}_t$ $(t = 1, 2, ...)$, the policy π_t is initialized by π_{t-1} and there is a latent reward function $r_{(t)}(x, y)$. Based on the Eq. (2), the optimal policy of learning task $\mathcal{T}_t$ is

$$\pi_t^*(y|x) = \frac{1}{Z_t(x)} \pi_{t-1}(y|x) \exp(\frac{1}{\beta} r_{(t)}(x, y)), \tag{3}$$

where $Z_t(x) = \sum_y \pi_{t-1}(y|x) \exp(\frac{1}{\beta} r_{(t)}(x, y))$ is the partition function of $\pi_t^*(y|x)$, $x \in \mathcal{D}_t$ denotes the prompt of task t. For each prompt x, the responses $\mathcal{Y}^x$ ranked by human preferences are known. To mitigate forgetting, a memory buffer $\mathbb{R} = \mathcal{R}_1 \cup \mathcal{R}_2 \cup ... \cup \mathcal{R}_{t-1}$ is maintained, where $\mathcal{R}_i \subset \mathcal{D}_i$ $(i = 1, 2, ..., t-1)$ is part of training data from historical tasks. The final objective of current task $\mathcal{T}_t$ is to learn a policy model π_θ that minimizes the KL-divergence for all optimal policies $\{\pi_i^* | i = 1, 2, ..., t\}$:

$$\min_\theta \sum_{i=1}^{t} \mathbb{E}_{x \sim \mathcal{D}_i} \left[\mathbb{D}_{\text{KL}}(\pi_\theta(y|x) \mid\mid \pi_i^*(y|x)) \right]. \tag{4}$$

In the CL setting, the whole historical dataset $\mathcal{D}_i$ $(i < t)$ is generally unavailable, posing a significant challenge for LLM to continually learn new preferences while minimizing interference with previously learned preferences.

3 Continual Preference Optimization

Our method is based on the optimal policy in Eq. (3) where the partition $Z_t(x)$ is intractable and the $\pi_t^*(y|x) \propto \exp(\frac{1}{\beta} r_{(t)}(x, y))$. To prevent widening the gap $\log \pi_t^*(y_w|x) - \log \pi_t^*(y_l|x) \to \infty$ like DPO, we determine a MRF. After determining the reward, we need to bypass the intractable partition function to get a feasible learning objective. For this purpose, we propose to parameterize the sampling distribution. The sampling distribution of historical policies can be regarded as optimization constraints for learning new preferences, providing a natural advantage for CL over maximum likelihood-based methods such as DPO. For clarity, we will first introduce the sampling distribution. Figure 1 demonstrates the framework of COPR.

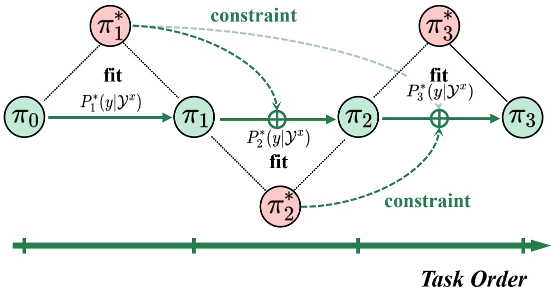

Figure 1: The framework of COPR. The optimal policy π_t^* $(t = 1, 2, 3)$ is derived from the policy π_{t-1} based on the optimal policy theory (Rafailov et al., 2023; Peng et al., 2019). The optimal policy π_t^* is utilized as the current policy's fitting objective and the learning constraints of the next π_i $(i > t)$.

3.1 Parameterize the Sampling Distribution

In the preliminaries, we formulate the problem of continual alignment as an optimization problem, but due to the difficulty in estimating the partition function $Z_t(x)$ (Goodfellow et al., 2016), the optimal policy remains challenging to fit directly. To bypass the partition function term, we define the sampling distribution of policy π.

Definition 1. The **sampling distribution** of the policy π_t^* is defined as the relative probabilities of generating different responses under the given prompt x, which is denoted by $P_t^*(y|\mathcal{Y}^x)$:

$$\begin{aligned} P_t^*(y|\mathcal{Y}^x) &\triangleq \frac{\pi_t^*(y|x)}{\sum_{y' \in \mathcal{Y}^x} \pi_t^*(y'|x)} \\ &= \frac{\pi_{t-1}(y|x) \exp(\frac{1}{\beta} r_{(t)}(x, y))}{\sum_{y' \in \mathcal{Y}^x} \pi_{t-1}(y'|x) \exp(\frac{1}{\beta} r_{(t)}(x, y'))}. \end{aligned} \tag{5}$$

Clearly, $\sum_{y \in \mathcal{Y}^x} P_t^*(y|\mathcal{Y}^x) = 1$, and $P_t^*(y|\mathcal{Y}^x)$ does not contain the partition function term. We propose that fitting the sampling distribution $P^*(y|\mathcal{Y}^x)$ of the optimal policy π^* is equivalent to directly learning the optimal policy $\pi^*(y|x)$ under sufficient sampling, which can be abstractly represented as *Proposition 1*.

Proposition 1. Given the prompt x and all possible responses $\mathbb{Y}^x \triangleq \{y | y \sim \pi^*(\cdot|x)\}$, for $\forall\, y \in \mathbb{Y}^x$, $\pi^*(y|x) = \pi_\theta(y|x) \iff$ for $\forall\, \mathcal{Y}^x$ that $\mathcal{Y}^x \subsetneq \mathbb{Y}^x$ and $|\mathcal{Y}^x| > 1$, $P^*(y|\mathcal{Y}^x) = P_\theta(y|\mathcal{Y}^x)$.

Proposition 1 indicates that under sufficient sampling, aligning human preference can be achieved by fitting the sampling distribution of the optimal policy. In Appendix A.1, we provide the formal proof. However, the sampling is generally insufficient in practical scenarios. We theoretically analyze the KL-divergence $\mathbb{D}_{\mathrm{KL}}[\pi_\theta(y|x) \mid\mid \pi^*(y|x)]$ between the LLM policy and the optimal policy when the sampling is insufficient. We derive a lower bound of $\mathbb{D}_{\mathrm{KL}}[\pi_\theta(y|x) \mid\mid \pi^*(y|x)]$ and discover the risk of probability reduction problem under insufficient sampling. Although directly processing the probability reduction problem is not feasible, we can counteract probability reduction by using SFT loss to enhance the probabilities.

Invariance of Sampling Distribution (Wu et al., 2023) introduces the *invariance* property, which contributes to learning stability. We show that the sampling distribution satisfies this property. Given the partially-ordered set of responses $\mathcal{Y}^x = \{y_1^x \prec y_2^x \prec ... \prec y_{J_x}^x\}$. We calculate the sampling distribution $P_t^*(y|\mathcal{Y}^x)$:

$$\begin{aligned} P_t^*(y|\mathcal{Y}^x) &\triangleq \frac{\pi_t^*(y|x)}{\sum_{y' \in \mathcal{Y}^x} \pi_t^*(y'|x)} \\ &= \frac{\cancel{\frac{1}{Z_t(x)}} \cdot \pi_{t-1}(y|x) \cdot \exp(\frac{1}{\beta} Adv(x,y) + \cancel{\frac{1}{\beta}\delta(x)})}{\sum_{y' \in \mathcal{Y}^x} \cancel{\frac{1}{Z_t(x)}} \cdot \pi_{t-1}(y'|x) \cdot \exp(\frac{1}{\beta} Adv(x,y') + \cancel{\frac{1}{\beta}\delta(x)})} \\ &= \frac{\pi_{t-1}(y|x) \exp(\frac{1}{\beta} Adv(x,y))}{\sum_{y' \in \mathcal{Y}^x} \pi_{t-1}(y'|x) \exp(\frac{1}{\beta} Adv(x,y'))}. \end{aligned} \tag{6}$$

The sampling distribution is independent of reward expectation $\delta(x)$. Hence, we only model the advantage term $Adv(x, y_j^x)$.

3.2 Moderate Reward Function

Recent research (Azar et al., 2023) proves that the standard RLHF pipeline is more robust than DPO in practice due to the underfitting of the reward function, and in fact, the regularization of the reward function has been recognized as a crucial aspect of RLHF training in practice (Christiano et al., 2017). Inspired by this perspective, we propose MRF that can be derived from the Bradley-Terry model. Introducing the MRF has 2 reasons, 1) to calculate the sampling distribution in Eq.(5) while bypassing the partition function $Z_t(x)$. 2) to prevent the preference collapse.

We split the reward $r(x, y)$ into the expected reward $\delta(x)$ and the advantage score $Adv(x, y)$, i.e., the extra reward one response can obtain compared with the expected reward:

$$r(x, y_j^x) = Adv(x, y_j^x) + \delta(x), \tag{7}$$

where $j = 1, 2, ..., J_x$ represents the human ranking information, the expectation $\delta(x) = E_{y \sim \pi(\cdot|x)} r(x, y)$ depends solely on the prompt x.

Property of Linearity The regularization of the reward function has proven to be a crucial element in practical RLHF training, as observed in (Christiano et al., 2017). Recent work (Azar et al., 2023) proves that the underfitting of the reward function is essential for achieving a final policy that is suitably regularized towards the reference policy π_{ref}. Therefore, here we assume the RM learns only 1 epoch based on preference data. The gradients of $\mathcal{L}_{ranking} = -\log(\sigma(r_\phi(x, y_w) - r_\phi(x, y_l))) = -\log(\sigma(r_w - r_l))$ according to r_w and r_l respectively are:

$$\frac{\partial \mathcal{L}_{ranking}}{\partial r_w} = \sigma(r_w - r_l) - 1, \tag{8}$$

$$\frac{\partial \mathcal{L}_{ranking}}{\partial r_l} = 1 - \sigma(r_w - r_l). \tag{9}$$

Considering that the partially-ordered set $\mathcal{Y}^x = \{y_1^x \prec y_2^x \prec ... \prec y_{J_x}^x\}$, according to Eq. (8) and Eq. (9), the accumulation of gradient according to r_j is

$$G_j = \sum_{k=1}^{j-1} -(1-\sigma(r_j - r_k)) + \sum_{k=j+1}^{J_x} (1-\sigma(r_k - r_j)), \tag{10}$$

where r_k $(k = 1, 2, ..., J_x)$ denotes the reward score of response y_k^x. We suppose that the initial reward r_j is close to zero. In the early stages of training, the reward value r_j is approximated to $0 - \eta \cdot G_i \approx (j-1) \cdot 0.5\eta - (J_x - j) \cdot 0.5\eta = \eta \cdot j - 0.5\eta(J_x + 1)$ which exhibits a linear relationship with the degree of human preference j.

The bounded advantage By deriving the gradient of the pairwise loss function based on the Bradley-Terry model:

$$\mathcal{L}_{ranking} = -\log(\sigma(r_\phi(x, y_w) - r_\phi(x, y_l))), \tag{11}$$

In Appendix A.4, we prove that the reward scores are approximately linearly related to the degree of human preferences, and that the **R**ange of the **R**eward $RR_x = r(x, y_{J_x}^x) - r(x, y_1^x) = Adv(x, y_{J_x}^x) - Adv(x, y_1^x)$ can be sampled from the Beta distribution with parameters $(J_x - 1, 2)$. Based on the fact that the mathematical expectation of the advantage is zero, we propose to use a linearly bounded advantage function:

$$Adv(x, y_j^x) \triangleq r_x \cdot (\frac{2}{J_x - 1} j - \frac{J_x + 1}{J_x - 1}), \tag{12}$$

where $r_x \sim \textbf{Beta}(J_x - 1, 2)$ and $j \in \{1, 2, ..., J_x\}$. Specially, when there are only y_w and y_l two responses, $Adv(x, y_w) = r_x$ and $Adv(x, y_l) = -r_x$. According to $\pi_t^*(y|x) \propto \exp(\frac{1}{\beta} Adv_{(t)}(x, y))$ and $|Adv(x, y_j^x)| \leq r_x$, we can choose a suitable β to control $\pi_t^*(y|x)$ in a reasonable range rather than degrading into $\{0, 1\}$. It should be noted that β is already a hyperparameter in both the original RLHF and DPO methods, so we are not introducing any new hyperparameters.

3.3 Learning Objective and Constraint

Next, we introduce the objectives of learning new preferences, and the constraints to preserve old preferences, both based on the sampling distribution. Finally, we propose to achieve our learning objective by the Lagrangian Dual (LD) method.

3.3.1 Learn New Human Preferences

We parameterize the sampling distribution $P_\theta(y|\mathcal{Y}^x)$ of the current policy π_θ by:

$$P_\theta(y|\mathcal{Y}^x) \triangleq \frac{\pi_\theta(y|x)}{\sum_{y' \in \mathcal{Y}^x} \pi_\theta(y'|x)}. \tag{13}$$

Then, we can learn π_θ by minimizing the logarithmic ratio of $P_\theta(y|\mathcal{Y}^x)$ and $P_t^*(y|\mathcal{Y}^x)$ at task $\mathcal{T}_t$, which aims to imitate the optimal policy π_t^*:

$$\mathcal{J}_t^{fit}(\theta) = \mathbb{E}_{x \sim \mathcal{D}_t} \sum_{y \in \mathcal{Y}^x} |\mathcal{Y}^x|^{-1} [\log P_\theta(y|\mathcal{Y}^x) - \log P_t^*(y|\mathcal{Y}^x)]^2, \tag{14}$$

where θ denotes the parameters of the policy model.

MSE and KL-divergence for Fitting Due to using the offline dataset where the responses set $\mathcal{Y}^x$ is provided by unknown sources (human or different models), we do not know the original distribution of $y \sim \mathcal{Y}^x$. Our objective is to imitate the optimal policy where the signal of each response is equally crucial, hence we set the same weight $|\mathcal{Y}^x|^{-1}$ for each response y. However, the KL-divergence of $P_\theta(\cdot|\mathcal{Y}^x)$ and $P_t^*(\cdot|\mathcal{Y}^x)$

$$\sum_{y \in \mathcal{Y}^x} P_\theta(y|\mathcal{Y}^x)[\log P_\theta(y|\mathcal{Y}^x) - \log P_t^*(y|\mathcal{Y}^x)], \tag{15}$$

assigns each response y with a different weight $P_\theta(y|\mathcal{Y}^x)$. Some responses including important preference may be assigned an inapposite weigth $P_\theta(y|\mathcal{Y}^x)$ due to the imperfect policy π_θ. For example, for a response y that the model has not seen, which may be the correct answer according to human preferences, $P_\theta(y|\mathcal{Y}^x)$ is generally small. Consequently, when computing the final fitting objective, the contribution of response y is excessively diminished.

3.3.2 Retain Old Human Preference

For continual learning, the policy needs to fit a sampling distribution sequence: $P_1^*(y|\mathcal{Y}^x) \rightarrow P_2^*(y|\mathcal{Y}^x) \rightarrow \cdots$. If fitting sequentially, the learned policy will gradually deviate from the historically optimal policy. We constrain the distance of the current policy from the old optimal policy to mitigate forgetting. Since historical $\mathcal{D}_i$ $(i < t)$ is not accessible in the learning of task $\mathcal{T}_t$, we introduce the surrogate objective of Eq. (4) which includes constraints from the replay memory buffer $\mathbb{R} = \mathcal{R}_1 \cup \mathcal{R}_2 \cup ... \cup \mathcal{R}_{t-1}$:

$$\min_\theta \mathcal{J}_t^{fit}(\theta), \text{ s.t. } \mathcal{J}_{\mathcal{C}i}(\theta) \leq 0 \ (i = 1, 2, ..., t-1), \tag{16}$$

where,

$$\mathcal{J}_{\mathcal{C}i}(\theta) \triangleq \mathbb{E}_{x \sim \mathcal{R}_i} \sum_{y \in \mathcal{Y}^x} |\mathcal{Y}^x|^{-1} [\log P_\theta(y|\mathcal{Y}^x) - \log P_i^*(y|\mathcal{Y}^x)]^2 - d_i, \tag{17}$$

which denotes the constraint of old task $\mathcal{T}_i$ and $d_i > 0$ is the constant threshold. Here we only need to store the sampling distribution values of memorized samples rather than re-calculate historical LLMs' forward functions. The storage overhead is negligible because each sample in the memory buffer only requires storing several (less than t) scalars.

3.3.3 Balance learning new preferences and retaining old preferences

To address the above optimization problem, we leverage the LD method, a technique for finding the local optimum over a constraint set. We convert the constrained primal problem, as defined in Eq. (16) into its unconstrained optimization objective as follows:

$$\min_\theta \max_{\lambda \geqslant 0} \big[\mathcal{J}_t^{fit}(\theta) + \boldsymbol{\lambda} \cdot \boldsymbol{\mathcal{J}_\mathcal{C}}(\theta)\big], \tag{18}$$

where the objective function $\mathcal{J}_t^{fit}(\theta)$ subjected to the constrain of $\boldsymbol{\mathcal{J}_\mathcal{C}}(\theta) \leq 0$, $\boldsymbol{\lambda} \geqslant 0$ serves as the Lagrange multiplier, and controls the regularization for current policy π_θ. Here, $\boldsymbol{\lambda} \cdot \boldsymbol{\mathcal{J}_\mathcal{C}}(\theta) = \sum_{i=1}^{t-1} \lambda_i \cdot \mathcal{J}_{\mathcal{C}i}(\theta)$ is the sum of the regularization penalty.

The Lagrangian Dual Method We utilize $\Lambda = 1 + \sum_{i=1}^{t-1} \lambda_i$ to normalize the training loss:

$$L(\theta) = \frac{1}{\Lambda}[\mathcal{J}_t(\theta) + \boldsymbol{\lambda} \cdot \boldsymbol{\mathcal{J}}_{\boldsymbol{\mathcal{C}}}(\theta)]. \tag{19}$$

Since $\boldsymbol{\lambda} > 0$, we set $\boldsymbol{\lambda} \triangleq e^{\boldsymbol{\sigma}}$ and take the gradient ascent to maximize the $\mathcal{J}_t(\theta) + e^{\boldsymbol{\sigma}} \cdot \boldsymbol{\mathcal{J}}_{\boldsymbol{\mathcal{C}}}(\theta)$ for $\boldsymbol{\sigma}$. Then we utilize the gradient descent to minimize the $\mathcal{J}_t(\theta) + \boldsymbol{\lambda} \cdot \boldsymbol{\mathcal{J}}_{\boldsymbol{\mathcal{C}}}(\theta)$ for θ. The updating rules for $\boldsymbol{\lambda}$ and θ can be derived as:

$$\begin{aligned} \ln \lambda_i^{m+1} &:= \ln \lambda_i^m + \alpha \cdot \lambda_i^m \cdot \mathcal{J}_{\mathcal{C}i}(\theta^m), \\ \theta^{m+1} &:= \theta^m - \frac{\eta}{\Lambda} \nabla_{\theta^m} \big[\mathcal{J}_t(\theta^m) + \boldsymbol{\lambda} \cdot \boldsymbol{\mathcal{J}}_{\boldsymbol{\mathcal{C}}}(\theta^m)\big], \end{aligned} \tag{20}$$

where η and α are learning rates, and m is the index of the updating step. The threshold d_i ($i = 1, 2, ..., t-1$) of the regularization penalty from task $\mathcal{T}_i$ is calculated on the replay memory buffer $\mathcal{R}_i$ by the initial parameters θ^0:

$$d_i := \mathbb{E}_{x \sim \mathcal{R}_i} \sum_{y \in \mathcal{Y}^x} |\mathcal{Y}^x|^{-1} [\log P_{\theta^0}(y|\mathcal{Y}^x) - \log P_i^*(y|\mathcal{Y}^x)]^2. \tag{21}$$

Besides, to improve the fluency of text, we simultaneously fit the optimal policy and the response that is considered the best by humans. Therefore, we replace the $\mathcal{J}_t^{fit}(\theta)$ in Eq. (18) by $\mathcal{J}_t(\theta) = \mathcal{J}_t^{sft}(\theta) + \mathcal{J}_t^{fit}(\theta)$, where $\mathcal{J}_t^{sft}(\theta)$ is the Negative Log-Likelihood (NLL) loss of the top one candidate. Consequently, the overall optimization objective of task $\mathcal{T}_t$ can be summarized as $\mathcal{J}_t(\theta) + \boldsymbol{\lambda} \cdot \boldsymbol{\mathcal{J}}_{\boldsymbol{\mathcal{C}}}$. It is important to note that the optimization of $\mathcal{J}_t(\theta)$ in the current task often contradicts the constraint of $\mathcal{J}_{\mathcal{C}i}$, which is regarded as CF. Thus, Eq. (18) can be interpreted as appending a regularization penalty term to the current objective. This penalty, which corresponds to the potential forgetting of historical tasks, can be dynamically modulated via the parameter $\boldsymbol{\lambda}$. Specifically, we iteratively solve the min-max problem in Eq. (18) by alternately updating the Lagrange multiplier $\boldsymbol{\lambda}$ and the LLM parameters θ. This ensures that any change in the potential CF associated with the updated model is rapidly reflected in the multiplier, thereby avoiding the risks of over-emphasizing one objective at the expense of the other under a fixed optimization ratio.

4 Experiments

In this section, we conduct extensive experiments on our proposed benchmark to evaluate COPR and baselines, we observe that: 1) Compared to strong baselines, COPR showcases distinct advantages in continual learning of human preferences. Particularly, COPR achieves comparable performance to the upper bound of continual learning of human preferences, namely, Iterated RLHF (Bai et al., 2022), without retraining and reinforcement learning (Section 4.2). 2) The policy regularization technique exhibits a notable effect on the continual learning of human preferences in the ablative experiment (Section 4.3).

4.1 Experiments Setup

Task Incremental Learning for Human Feedback (TIL-HF) benchmark. The policy is required to learn across three commonly used RLHF tasks continually: 1) the question-answer task on the HH-RLHF (Bai et al., 2022) dataset, 2) the summary task on the Reddit TL;DR human feedback (Völske et al., 2017) dataset, and 3) the text continuation task on the IMDB (Maas et al., 2011) movie review dataset. The summarization is shown in Table 1.

Table 1: Tasks, input, output, metrics, and sample statistics of the TIL-HF benchmark.

	HH-RLHF	**Reddit TL;DR**	**IMDB**
Task	Q&A	Summarization	Text Continuation
Input	Question	Reddit POST	Partial Review
Output	Helpful & Harmless Answer	Summarized Reddit POST	Positive Sentiment Completion
Preference Metric	SteamSHP	GPT-j	DistilBERT
Train Set	35.2k	14.8k	24.9k
Valid Set	200	200	200
Test Set	1000	1000	1000

Evaluation Metric for Continual Learning. Following previous works (Rafailov et al., 2023; Song et al., 2023; Bai et al., 2022), we use different preference models to calculate the preference scores for various tasks, summarized in Table 1. For CL evaluation, the overall performance is commonly assessed through *average accuracy* (AA) (Chaudhry et al., 2018; Lopez-Paz and Ranzato, 2017) and *average incremental accuracy* (AIA) (Douillard et al., 2020; Hou et al., 2019). Memory stability can be assessed using the forgetting measure (FM) (Chaudhry et al., 2018) and backward transfer (BWT) (Lopez-Paz and Ranzato, 2017). In our evaluation scenario, accuracy is substituted with the normalized Preference Metric (0-1).

Table 2: The performances at the last task in the TIL scenario. Iterated RLHF (Bai et al., 2022) uses the new and historical datasets to retrain the policy when a new task arrives, which can be regarded as the upper bound of continual alignment methods. Due to the original SPIN, CoH, DPO, and IPO methods are not supported for continual learning, we utilize 5% historical samples for Experience Replay (ER). All of the experiments are based on the Llama-7B. The gray rows indicate the merging of new tasks and all historical data to retrain the model, thereby typically resulting in better performance compared to methods within the same category.

Method	**HH**	**TL;DR**	**IMDB**	**Overall performance**		**Memory stability**	
	SteamSHP(↑)	**GPT-j(↑)**	**DistilBERT(↑)**	**AA (↑)**	**AIA (↑)**	**BWT (↑)**	**FM (↓)**
SFT In order	0.772±0.0171	0.771±0.0082	0.580±0.0148	0.720±0.0044	0.739±0.0028	-0.043±0.0084	0.043±0.0084
SFT Multi-tasks	0.825±0.0143	0.781±0.0041	0.641±0.0151	-	-	-	-
SFT+Online L2Reg	0.780±0.0107	0.764±0.0116	0.640±0.0098	0.728±0.0058	0.749±0.0075	-0.024±0.0042	0.024±0.0042
SFT+EWC (Kirkpatrick et al., 2017)	0.792±0.0107	0.771±0.0116	0.645±0.0098	0.736±0.0058	0.771±0.0075	-0.013±0.0051	0.013±0.0051
SFT+DER++ (Buzzega et al., 2020)	0.817±0.0116	0.774±0.0089	0.637±0.0056	0.743±0.0052	0.781±0.0037	-0.018±0.0045	0.018±0.0045
SPIN+ER (Chen et al., 2024)	0.851±0.0118	0.772±0.0074	0.642±0.0067	0.755±0.0032	0.794±0.0041	-0.033±0.0032	0.033±0.0032
Iterated RLHF	0.867±0.0545	0.799±0.0426	0.692±0.0742	-	-	-	-
CoH+ER (Liu et al., 2023a)	0.807±0.0121	0.743±0.0107	0.625±0.0113	0.725±0.0082	0.781±0.0026	-0.027±0.0041	0.027±0.0041
RRHF+ER (Yuan et al., 2023)	0.808±0.0126	0.727±0.0198	0.641±0.0092	0.725±0.0065	0.791±0.0072	**-0.011**±0.0025	**0.018**±0.0031
IPO+ER (Azar et al., 2023)	0.814±0.0197	0.741±0.0132	0.654±0.0146	0.736±0.0088	0.794±0.0079	-0.028±0.0031	0.028±0.0031
DPO+ER (Rafailov et al., 2023)	0.828±0.0165	0.775±0.0125	0.672±0.0098	0.758±0.0086	0.816±0.0041	-0.024±0.0047	0.024±0.0047
CPPO (learn) (Zhang et al., 2024)	0.821±0.0531	0.756±0.0264	0.649±0.0212	0.742±0.0251	0.790±0.0031	-0.031±0.0044	0.031±0.0044
COPR (ours)	**0.866**±0.0126	**0.789**±0.0147	**0.680**±0.0114	**0.778**±0.0045	**0.847**±0.0031	-0.019±0.0025	0.019±0.0025

Baselines. We use the SFT-based method and alignment method as baselines for comparison. SFT directly learns the human-labeled responses through the NLL loss. For CL, we combine SFT with classic continual learning methods including Online L2Reg, EWC (Kirkpatrick et al., 2017) and DER++ (Buzzega et al., 2020). We adopt experience replay (ER) in combination with alignment methods as baselines for continual preference learning because (Hussain et al., 2021) shows that many approaches fail to surpass a simple baseline in realistic lifelong learning conditions, and ER remains the most commonly used and easiest CL technique to implement. In detail, we compare COPR with Ranking-based Approaches including DPO+ER (Rafailov et al., 2023), IPO+ER (Azar et al., 2023), RRHF+ER (Yuan et al., 2023), the Language-based Approach CoH+ER (Liu et al., 2023a), and the Adversarial training-based Approach SPIN+ER (Chen et al., 2024). We adopt the RL-based method, Iterated RLHF (Bai et al., 2022), which retrains the policy on the mixed data of new and historical tasks.

4.2 Results and Discussion

The rationality of the TIL-HF benchmark. Table 2 presents the results for continual learning from human preferences in the TIL-HF benchmark. We assess the SFT (in order), SFT (multi-tasks), and Iterated RLHF on the TIL-HF benchmark. The results reveal the poor overall performance and memory stability of SFT (in order), while SFT (multi-tasks) outperforms SFT (in order) significantly. The serious forgetting highlights the necessity for continual learning. Moreover, Iterated RLHF's outsanding performance indicates that retraining with RLHF can be regarded as a very strong baseline for the TIL-HF benchmark.

The overall performance. The results in Table 2 show that all alignment methods outperform the SFT-based methods. This is attributed to alignment methods that leverage negative responses rather than only fit gold responses. It indicates that reducing the generation probability of negative responses contributes to the performance of alignment. SPIN exclusively employs prompts and gold responses but generates negative responses and trains with DPO, which results in better performance than SFT. Other rank-based alignment methods, such as COPR, use real negative responses, resulting in better performance than SPIN. Compared with Iterated RLHF (Bai et al., 2022), which retrains the policy model on the combination of new and old datasets, non-retrained methods still exhibit certain gaps. The continual RL method, CPPO, employs MAS to continually learn a reward model, which is a significant bottleneck for TIL. Typically, MAS exhibits weaker performance in the TIL context but achieves better results in the DIL scenario. For further details, please refer to Appendix B.4. This discrepancy arises because Iterated RLHF leverages all historical information and does not suffer from CF. Compared with rank-based alignment methods, COPR demonstrates significant advantages in overall performance, indicating its greater potential in real-world applications where users are primarily

Table 3: The learning objectives.

Method	Objective (Minimize)
RRHF	$\max\big[0, \log \pi_\theta(y_l\|x) - \log \pi_\theta(y_w\|x)\big]$
DPO	$-\log \sigma\big(\beta \log \frac{\pi_\theta(y_w\|x)}{\pi_{ref}(y_w\|x)} - \beta \log \frac{\pi_\theta(y_l\|x)}{\pi_{ref}(y_l\|x)}\big)$
COPR	$\|\log P_\theta(y\|\mathcal{Y}^x) - \log P_t^*(y\|\mathcal{Y}^x)\|^2$

Table 4: The diversity of generation.

Method	Distinct-2(↑)	MSTTR(↑)	Entroy(↑)	self-bleu(↓)
HUMAN	0.686	0.693	6.530	0.102
RRHF+ER	0.497	0.599	5.146	0.142
DPO+ER	0.419	0.541	4.276	0.192
COPR (Ours)	**0.661**	**0.704**	**6.322**	**0.097**

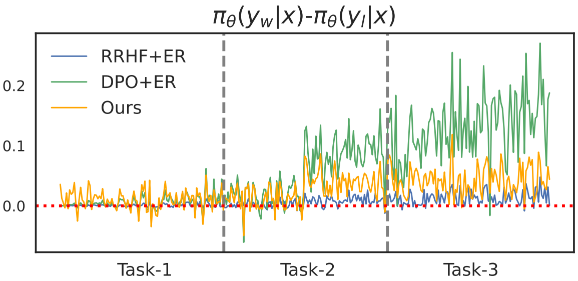

Figure 2: The expectation $\mathbb{E}_{x\sim\mathbb{R}}(\pi_\theta(y_w|x) - \pi_\theta(y_l|x))$ of samples in the memory buffer.

concerned with the final performance.

The memory stability metric evaluates the forgetting of CL models on old tasks. RRHF+ER outperforms in memory stability metrics because it employs rank loss with a zero threshold value (Table 3). Although initially not designed for CL, it prevents overfitting to replayed samples. However, lacking a reference model and employing a fixed threshold hinder RRHF's ability to discern stronger human preferences, resulting in inferior overall performance compared to DPO and COPR.

Can COPR address Preference Collapse? We observe an obvious difference between COPR and DPO+ER when monitoring the probability of samples in the memory buffer. When continually learning new tasks DPO significantly increases the gap between the probabilities $\pi_\theta(y_w|x)$ and $\pi_\theta(y_l|x)$ in the memory buffer $\mathbb{R}$. As previously mentioned, the maximum likelihood objective [2] of DPO fails to effectively regulate the generation probability of training samples within an appropriate range during continual learning. We find DPO reduces the diversity of generation, which is verified by the Mean Segmented Type Token Ratio (MSTTR) (Johnson, 1944) in Table 4. Additionally, the learning objectives of both RRHF and DPO include a gradient ascent term, namely maximizing the negative logarithmic likelihood $-\log \pi_\theta(y_l|x)$, which has proved very fragile and easily causes catastrophic outputs (Gu et al., 2024). Although COPR also increases the $-\log \pi_\theta(y_l|x)$, the sampling distribution $P^*(y_l|\mathcal{Y}^x)$ of the optimal policy and the learning constraints induced by the moderate reward determine a moderate range of $-\log \pi_\theta(y_l|x)$, the ultimate result is that COPR exhibits better to retain knowledge from historical tasks.

Table 5: The ablation experiments for COPR.

Method	Overall performance		Memory stability	
	AA (↑)	AIA (↑)	BWT (↑)	FM (↓)
COPR	**0.778**±0.0045	**0.847**±0.0031	-0.019±0.0025	0.019±0.0025
w/o $\mathcal{J}_t^{fit}(\theta)$	0.733±0.0056	0.796±0.0041	-0.044±0.0065	0.044±0.0065
w/o $\mathcal{J}_{\mathcal{C}i}(\theta)$	0.712±0.0081	0.720±0.0013	-0.056±0.0043	0.056±0.0043
$\mathcal{J}_{\mathcal{C}i} \rightarrow$ ER	0.752±0.0042	0.825±0.0048	-0.025±0.0063	0.025±0.0063
w/o $\mathcal{J}_t^{sft}(\theta)$	0.772±0.0036	0.835±0.0054	**-0.017**±0.0046	**0.017**±0.0046
w/o LD ($\boldsymbol{\lambda} \equiv 1$)	0.759±0.0027	0.827±0.0026	-0.028±0.0047	0.028±0.0047

4.3 Ablation Study

In this section, we perform an ablation experiment to assess the impact of the following factors on our method:

1) The influence of fitting the sampling distribution of the optimal policy by excluding $\mathcal{J}_t^{fit}(\theta)$. **2)** The impact of learning constraints by omitting $\mathcal{J}_{\mathcal{C}i}(\theta)$, and the influence of replacing $\mathcal{J}_{\mathcal{C}i}(\theta)$ with ER. **3)** The effect of SFT loss by eliminating $\mathcal{J}_t^{sft}(\theta)$. **4)** The consequences of employing the LD method by enforcing $\boldsymbol{\lambda} \equiv 1$.

From the results in Table 5, we observe the obvious performance degradation. The most significant impact on the final performance comes from the removal of the learning constraints $\mathcal{J}_{\mathcal{C}i}(\theta)$, namely always setting $\boldsymbol{\lambda} \equiv 0$. Using ER instead of $\mathcal{J}_{\mathcal{C}i}(\theta)$ still reduces the performance, indicating that our method indeed outperforms directly replaying the historical samples. This result indicates that learning constraints significantly assist in learning incremental preferences. Compared with the fixed regularization weight, such as $\boldsymbol{\lambda} \equiv 1$, using the Lagrangian multiplier helps for both overall performance and memory stability. This demonstrates that flexibly adjusting the strength of constraints is more effective than using fixed-strength constraints. We observe an improvement in BWT and FM metrics after removing SFT loss $\mathcal{J}_t^{sft}(\theta)$. The reason is that eliminating $\mathcal{J}_t^{sft}(\theta)$ reduces the performance on the current task, resulting in a corresponding decrease in the upper limit of forgetting.

---

[2] Refer to maximizing the human preference distribution under the Bradley-Terry model.

4.4 Human and GPT-4 Evaluation

In evaluation, the RM lacks comprehensiveness. Hence we provide comprehensive evaluations conducted by both GPT-4 and human assessors, where the evaluation process is similar to (Song et al., 2023). The difference is that we use the Coherence (does it generate fluently and without repeated or uncommon characters) indicator to identify ineffective answers caused by overfitting and model collapse, because we find that the repeated or uncommon characters are easier to generate among all methods. The details for GPT-4 evaluation are shown in Appendix Table 7. The human evaluation is modified based on the results of the GPT-4 assessment.

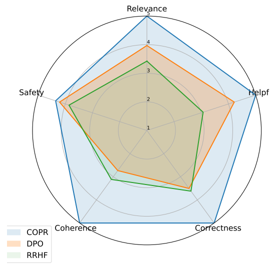

Figure 3: GPT4 Eval.

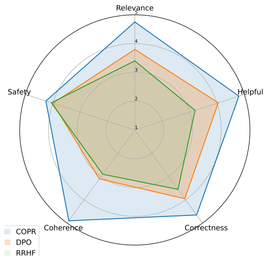

Figure 4: Human Eval.

The results presented in Figure 3 and Figure 4 offer a thorough examination, showcasing strong support for COPR from both GPT-4 and human evaluators across various comparisons. In the Coherence evaluation, we focus on grammatical coherence, generative diversity, and repetitiveness. In this evaluation criterion, DPO+ER shows significant weaknesses, which is consistent with the viewpoint we proposed that the optimization objective of DPO reduces generative diversity when CL. From the GPT4 evaluation, the results from human assessments, the diversity metrics, and multiple dimensions of model scoring are consistent, indicating that COPR indeed has an advantage in continually learning human preferences.

5 Related Works

5.1 Continual Learning Methods

Continual learning (Wang et al., 2023a) has made significant progress in recent years, encompassing various approaches. The regularization-based approach (Kirkpatrick et al., 2017) adds explicit terms to balance new skill acquisition and past knowledge retention. Replay-based strategies (Lin, 1992), encompassing experience replay to enhance model performance by preserving and reusing past experiences. Optimization-based techniques (Lopez-Paz and Ranzato, 2017) manipulate optimization programs to navigate continual learning challenges. Representation-based methodologies (Gallardo et al., 2021) leverage self-supervised learning (SSL) and large-scale pre-training to enhance representation quality. Architecture-based innovations (Serra et al., 2018) address inter-task interference through task-specific parameter design.

5.2 Learning from Human Preferences

Online training methods such as PPO (Schulman et al., 2017), SPIN (Chen et al., 2024), RAFT (Dong et al., 2023), and P3O (Wu et al., 2023) consist of a loop of generating new responses from the updated policy. Previous works (Stiennon et al., 2020; Ouyang et al., 2022) utilize the PPO (Schulman et al., 2017) algorithm to fine-tune an LLM for aligning human preference. Offline training typically involves a static dataset and doesn't require additional evaluations or generations, which includes rank-based approach (Rafailov et al., 2023; Song et al., 2023; Yuan et al., 2023; Zhao et al., 2023) and language-based approach (Liu et al., 2023a, 2022, 2023b; Madaan et al., 2023).

6 Conclusion

We propose Continual Optimal Policy Regularization (COPR), a scalable and efficient framework for continual alignment in Large Language Models (LLMs). COPR addresses key challenges such as catastrophic forgetting and preference collapse by leveraging historical optimal policies and a Moderate Reward Function (MRF) to balance the retention of past preferences and the integration of new ones. Through extensive experiments on task-incremental and domain-incremental benchmarks, COPR demonstrates superior performance and stability compared to existing methods. This work provides a practical solution for dynamically aligning LLMs to evolving human preferences, paving the way for more adaptive and reliable AI systems in real-world applications.

7 Limitations

Due to the potential application of alignment techniques in other scenarios, such as for political purposes, it is challenging to estimate the potential negative societal impacts. Additionally, since the human preference data used in this paper is derived from feedback from specific societal groups (such as those whose native language is English), its values may be influenced by specific cultural and social norms, making it not universally applicable to all societal groups.

8 Ethical Considerations

Due to the potential application of alignment techniques in other scenarios, such as for political purposes, it is challenging to estimate the potential negative societal impacts. Additionally, since the human preference data used in this paper is derived from feedback from specific societal groups (such as those whose native language is English), its values may be influenced by specific cultural and social norms, making it not universally applicable to all societal groups.

9 Acknowledgements

This research was supported in part by the Major Key Proiect of PCL (NO.PCL2023A09), the National Natural Science Foundation of China 62176076, the Natural Science Foundation of Guang Dong 2023A1515012922, the Shenzhen Foundational Research Funding JCYJ20220818102415032, Guangdong Provincial Key Laboratory of Novel Security Intelligence Technologies 2022B1212010005, CIPSC-SMP-ZHIPU Large Model Cross-Disciplinary Fund ZPCG20241119405, Hong Kong RGC GRF No. 14206324, CUHK direct grant No. 4055209, CUHK Knowledge Transfer Project Fund No. KPF23GWP20UKVI New Horizons grant (EP/X019063/1) and EPSRC IAA at KCL.

References

Mohammad Gheshlaghi Azar, Mark Rowland, Bilal Piot, Daniel Guo, Daniele Calandriello, Michal Valko, and Rémi Munos. 2023. A general theoretical paradigm to understand learning from human preferences. *arXiv preprint arXiv:2310.12036*.

Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback.

Ralph Allan Bradley and Milton E Terry. 1952. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345.

Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and SIMONE CALDERARA. 2020. Dark experience for general continual learning: a strong, simple baseline. In *Advances in Neural Information Processing Systems*, volume 33, pages 15920–15930. Curran Associates, Inc.

Arslan Chaudhry, Puneet K. Dokania, Thalaiyasingam Ajanthan, and Philip H. S. Torr. 2018. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European Conference on Computer Vision (ECCV)*.

Zixiang Chen, Yihe Deng, Huizhuo Yuan, Kaixuan Ji, and Quanquan Gu. 2024. Self-play fine-tuning converts weak language models to strong language models. *arXiv preprint arXiv:2401.01335*.

Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.

Josef Dai, Xuehai Pan, Ruiyang Sun, Jiaming Ji, Xinbo Xu, Mickel Liu, Yizhou Wang, and Yaodong Yang. 2023. Safe RLHF: Safe reinforcement learning from human feedback.

Herbert A David and Haikady N Nagaraja. 2004. *Order statistics*. John Wiley & Sons.

Hanze Dong, Wei Xiong, Deepanshu Goyal, Yihan Zhang, Winnie Chow, Rui Pan, Shizhe Diao, Jipeng Zhang, Kashun Shum, and Tong Zhang. 2023. RAFT: Reward ranked finetuning for generative foundation model alignment.

Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. 2020. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*, pages 86–102. Springer.

Kawin Ethayarajh, Yejin Choi, and Swabha Swayamdipta. 2022. Understanding dataset difficulty with $\mathcal{V}$-usable information. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 5988–6008. PMLR.

Jhair Gallardo, Tyler L Hayes, and Christopher Kanan. 2021. Self-supervised training enhances online continual learning. *arXiv preprint arXiv:2103.14010*.

Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep learning*. MIT press.

Kang Gu, Md Rafi Ur Rashid, Najrin Sultana, and Shagufta Mehnaz. 2024. Second-order information matters: Revisiting machine unlearning for large language models. *arXiv preprint arXiv:2403.10557*.

Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. 2019. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 831–839.

Aman Hussain, Nithin Holla, Pushkar Mishra, Helen Yannakoudakis, and Ekaterina Shutova. 2021. Towards a robust experimental framework and benchmark for lifelong language learning. In *Thirty-fifth Conference on Neural Information Processing Systems*.

Wendell Johnson. 1944. Studies in language behavior: A program of research. *Psychological Monographs*, 56(2):1–15.

James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526.

Long-Ji Lin. 1992. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Mach. Learn.*, 8(3–4):293–321.

Hao Liu, Carmelo Sferrazza, and Pieter Abbeel. 2023a. Chain of hindsight aligns language models with feedback.

Ruibo Liu, Chenyan Jia, Ge Zhang, Ziyu Zhuang, Tony X Liu, and Soroush Vosoughi. 2022. Second thoughts are best: Learning to re-align with human values from text edits. In *Advances in Neural Information Processing Systems*.

Ruibo Liu, Ruixin Yang, Chenyan Jia, Ge Zhang, Denny Zhou, Andrew M. Dai, Diyi Yang, and Soroush Vosoughi. 2023b. Training socially aligned language models on simulated social interactions.

David Lopez-Paz and Marc'Aurelio Ranzato. 2017. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 6467–6476.

Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA. Association for Computational Linguistics.

Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. 2023. Self-refine: Iterative refinement with self-feedback. *arXiv preprint arXiv:2303.17651*.

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.

Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine. 2019. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning. *arXiv preprint arXiv:1910.00177*.

Jan Peters and Stefan Schaal. 2007. Reinforcement learning by reward-weighted regression for operational space control. In *Proceedings of the 24th International Conference on Machine Learning*, ICML '07, page 745–750, New York, NY, USA. Association for Computing Machinery.

Biqing Qi, Pengfei Li, Fangyuan Li, Junqi Gao, Kaiyan Zhang, and Bowen Zhou. 2024. Online dpo: Online direct preference optimization with fast-slow chasing. *arXiv preprint arXiv:2406.05534*.

Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. Direct preference optimization: Your language model is secretly a reward model. In *Thirty-seventh Conference on Neural Information Processing Systems*.

Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter. *CoRR*, abs/1910.01108.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *CoRR*, abs/1707.06347.

Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. 2018. Overcoming catastrophic forgetting with hard attention to the task. In *International conference on machine learning*, pages 4548–4557. PMLR.

Feifan Song, Bowen Yu, Minghao Li, Haiyang Yu, Fei Huang, Yongbin Li, and Houfeng Wang. 2023. Preference ranking optimization for human alignment.

Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F. Christiano. 2020. Learning to summarize from human feedback. *CoRR*, abs/2009.01325.

Michael Völske, Martin Potthast, Shahbaz Syed, and Benno Stein. 2017. TL;DR: Mining Reddit to learn automatic summarization. In *Proceedings of the Workshop on New Frontiers in Summarization*, pages 59–63, Copenhagen, Denmark. Association for Computational Linguistics.

Liyuan Wang, Xingxing Zhang, Hang Su, and Jun Zhu. 2023a. A comprehensive survey of continual learning: Theory, method and application.

Peiyi Wang, Lei Li, Liang Chen, Zefan Cai, Dawei Zhu, Binghuai Lin, Yunbo Cao, Qi Liu, Tianyu Liu, and Zhifang Sui. 2023b. Large language models are not fair evaluators.

Tianhao Wu, Banghua Zhu, Ruoyu Zhang, Zhaojin Wen, Kannan Ramchandran, and Jiantao Jiao. 2023. Pairwise proximal policy optimization: Harnessing relative feedback for llm alignment.

Tongtong Wu, Linhao Luo, Yuan-Fang Li, Shirui Pan, Thuy-Trang Vu, and Gholamreza Haffari. 2024. Continual learning for large language models: A survey. *arXiv preprint arXiv:2402.01364*.

Jiancong Xiao, Ziniu Li, Xingyu Xie, Emily Getzen, Cong Fang, Qi Long, and Weijie J Su. 2024. On the algorithmic bias of aligning large language models with rlhf: Preference collapse and matching regularization. *arXiv preprint arXiv:2405.16455*.

Zheng Yuan, Hongyi Yuan, Chuanqi Tan, Wei Wang, Songfang Huang, and Fei Huang. 2023. RRHF: Rank responses to align language models with human feedback without tears.

Han Zhang, Yu Lei, Lin Gui, Min Yang, Yulan He, Hui Wang, and Ruifeng Xu. 2024. CPPO: Continual learning for reinforcement learning with human feedback. In *The Twelfth International Conference on Learning Representations*.

Yao Zhao, Mikhail Khalman, Rishabh Joshi, Shashi Narayan, Mohammad Saleh, and Peter J Liu. 2023. Calibrating sequence likelihood improves conditional language generation. In *The Eleventh International Conference on Learning Representations*.

A Theoretical Supplement

A.1 Why Fit the Sampling Distribution?

We prove that fitting the sampling distribution $P^*(y|\mathcal{Y}^x)$ of the optimal policy is equivalent to directly learning the optimal policy $\pi^*(y|x)$ under sufficient sampling, namely **Proposition 1**.

Proposition 1. Given the prompt x and all possible responses $\mathbb{Y}^x \triangleq \{y|y \sim \pi^*(\cdot|x)\}$, for $\forall\ y \in \mathbb{Y}^x$, satisfy $\pi^*(y|x) = \pi_\theta(y|x) \iff$ For any nontrivial subset $\mathcal{Y}^x$ of $\mathbb{Y}^x$ with more than one element (namely, $\mathcal{Y}^x \subsetneq \mathbb{Y}^x$, $|\mathcal{Y}^x| > 1$), satisfying $P^*(y|\mathcal{Y}^x) = P_\theta(y|\mathcal{Y}^x)$.

Proof of Proposition 1: Firstly, it is obvious that if $\pi^*(y|x) = \pi_\theta(y|x)$, then for prompt x and corresponding responses set $\mathcal{Y}^x$, we have $P^*(y|\mathcal{Y}^x) = P_\theta(y|\mathcal{Y}^x)$. Therefore, the necessity of Proposition 1 is evident. For the sufficiency of Proposition 1, we need to prove that for given a prompt x and any responses set $\mathcal{Y}^x$ generated based on x, if the $P^*(y|\mathcal{Y}^x) = P_\theta(y|\mathcal{Y}^x)$ holds, then $\pi^*(y|x) = \pi_\theta(y|x)$.

We proof by contradiction, assuming the existence of a prompt x and response y_1, such that $\pi^*(y_1|x) \neq \pi_\theta(y_1|x)$, then we choose the nontrivial subset $\mathcal{Y}_1^x = \mathbb{Y}^x - \{y_1\}$ and $\mathcal{Y}_2^x = \{y_1, y_2\}$ where $y_2 \in \mathcal{Y}_1^x$. Based on the condition of Proposition 1, we have $P^*(y|\mathcal{Y}_1^x) = P_\theta(y|\mathcal{Y}_1^x)$ and $P^*(y|\mathcal{Y}_2^x) = P_\theta(y|\mathcal{Y}_2^x)$. By establishing the equality relationship between the right-hand sides of Eq. (5) and Eq. (13), we obtain that

$$\begin{aligned}
&\frac{\pi^*(y|x)}{\sum_{y'\in\mathcal{Y}_1^x}\pi^*(y'|x)} = \frac{\pi_\theta(y|x)}{\sum_{y'\in\mathcal{Y}_1^x}\pi_\theta(y'|x)} \\
\Rightarrow\ &\frac{\pi_\theta(y|x)}{\pi^*(y|x)} = \frac{\sum_{y'\in\mathcal{Y}_1^x}\pi_\theta(y'|x)}{\sum_{y'\in\mathcal{Y}_1^x}\pi^*(y'|x)} = \gamma_1,\ \text{for } \forall y \in \mathcal{Y}_1^x, \\
&\frac{\pi^*(y|x)}{\sum_{y'\in\mathcal{Y}_2^x}\pi^*(y'|x)} = \frac{\pi_\theta(y|x)}{\sum_{y'\in\mathcal{Y}_2^x}\pi_\theta(y'|x)} \\
\Rightarrow\ &\frac{\pi_\theta(y|x)}{\pi^*(y|x)} = \frac{\sum_{y'\in\mathcal{Y}_2^x}\pi_\theta(y'|x)}{\sum_{y'\in\mathcal{Y}_2^x}\pi^*(y'|x)} = \gamma_2,\ \text{for } \forall y \in \mathcal{Y}_2^x,
\end{aligned} \tag{22}$$

where $\gamma_1 > 0$ and $\gamma_2 > 0$ are constants. Here, we assume that $\pi^*(y|x) > 0$, which holds true for LLM. Specially,

$$\begin{aligned}
y_2 \in \mathcal{Y}_1^x \Rightarrow \pi_\theta(y_2|x) = \gamma_1 \cdot \pi^*(y_2|x), \\
y_2 \in \mathcal{Y}_2^x \Rightarrow \pi_\theta(y_2|x) = \gamma_2 \cdot \pi^*(y_2|x),
\end{aligned} \tag{23}$$

hence, $\gamma_1 = \gamma_2 = \gamma$, namely $\forall y \in \mathbb{Y}^x$ satisfy $\pi_\theta(y|x) = \gamma \cdot \pi^*(y|x)$. According to the sum of probabilities is 1, we can get $\gamma = 1$ by:

$$\sum_{y'\in\mathbb{Y}^x} \pi_{\theta(y'|x)} = \gamma \cdot \sum_{y'\in\mathbb{Y}^x} \pi^*(y'|x) = \gamma = 1, \tag{24}$$

Namely $\pi^*(y|x) = \pi_\theta(y|x)$, which is contraried to $\pi^*(y_1|x) \neq \pi_\theta(y_1|x)$. In conclusion, $P^*(y|\mathcal{Y}^x) = P_\theta(y|\mathcal{Y}^x)$ is a necessary and sufficient condition for $\pi^*(y|x) = \pi_\theta(y|x)$. □

A.2 MSE and KL-divergence for Fitting

Due to using the offline dataset where the responses set $\mathcal{Y}^x$ is provided by unknown sources (human or different models), we do not know the original distribution of $y \sim \mathcal{Y}^x$. Our objective is to imitate the optimal policy where the signal of each response is equally crucial, hence we set the same weight $|\mathcal{Y}^x|^{-1}$ for each response y. However, the KL-divergence of $P_\theta(\cdot|\mathcal{Y}^x)$ and $P_t^*(\cdot|\mathcal{Y}^x)$

$$\sum_{y\in\mathcal{Y}^x} P_\theta(y|\mathcal{Y}^x)[\log P_\theta(y|\mathcal{Y}^x) - \log P_t^*(y|\mathcal{Y}^x)], \tag{25}$$

assigns each response y with a different weight $P_\theta(y|\mathcal{Y}^x)$. Some responses including important preference may be assigned an inapposite weigth $P_\theta(y|\mathcal{Y}^x)$ due to the imperfect policy π_θ. For example, for a response y that the model has not seen, which may be the correct answer according to human preferences, $P_\theta(y|\mathcal{Y}^x)$ is generally small. Consequently, when computing the final fitting objective, the contribution of response y is excessively diminished.

A.3 The KL-divergence Between the LLM and Optimal Policies

In the preceding proof, we assumed that for any prompt x and any set of responses $\mathcal{Y}^x = \{y_1^x \prec y_2^x \prec \ldots \prec y_n^x\}$ generated based on x, the condition $P^*(y|\mathcal{Y}^x) = P_\theta(y|\mathcal{Y}^x)$ holds to conclude $\pi = \pi^*$. However, in practical training, we cannot guarantee obtaining all possible responses due to insufficient sampling. Hence, we analyze the errors of COPR under insufficient sampling by calculating the KL divergence between the LLM and optimal policies.

First, we introduce the symbol p to denote the sum of the probability of sampled responses,

$$p \triangleq \sum_{y \in \mathcal{Y}^x} \pi^*(y|x). \tag{26}$$

By fitting the sampling distribution $P^*(y|\mathcal{Y}^x)$, the well-trained π_θ satisfies that $\pi_\theta(y|x) = \gamma \cdot \pi^*(y|x)$ for $\forall y \in \mathcal{Y}^x$, where $\gamma > 0$ is a scaling factor and related to x. This conclusion can be obtained by the **Proof of Proposition 1**. Then the KL-divergence of $\pi_\theta(y|x)$ and $\pi^*(y|x)$ is

$$\begin{aligned}
&\mathbb{D}_{\mathrm{KL}}[\pi_\theta(y|x) \mid\mid \pi^*(y|x)] \\
&= \sum\nolimits_{y \in \mathcal{Y}^x} \pi_\theta(y|x) \cdot \log \frac{\pi_\theta(y|x)}{\pi^*(y|x)} + \sum\nolimits_{y \in \mathbb{Y}^x - \mathcal{Y}^x} \pi_\theta(y|x) \cdot \log \frac{\pi_\theta(y|x)}{\pi^*(y|x)} \\
&= \sum\nolimits_{y \in \mathcal{Y}^x} \pi_\theta(y|x) \cdot \log \gamma + \sum\nolimits_{y \in \mathbb{Y}^x - \mathcal{Y}^x} \pi_\theta(y|x) \cdot \log \frac{\pi_\theta(y|x)}{\pi^*(y|x)} \\
&= \ \gamma p \cdot \log \gamma + \sum\nolimits_{y \in \mathbb{Y}^x - \mathcal{Y}^x} \pi_\theta(y|x) \cdot \log \frac{\pi_\theta(y|x)}{\pi^*(y|x)}.
\end{aligned} \tag{27}$$

To derive the lower bound of Eq. (27), we construct two new probability distributions $\pi^{'}(y|x) \triangleq \pi_\theta(y|x)/(1-\gamma p)$ and $\pi^{''}(y|x) \triangleq \pi^*(y|x)/(1-p)$ for $y \in \mathbb{Y}^x - \mathcal{Y}^x$, which satisfies $\sum_{y \in \mathbb{Y}^x - \mathcal{Y}^x} \pi^{'}(y|x) = \sum_{y \in \mathbb{Y}^x - \mathcal{Y}^x} \pi^{''}(y|x) = 1$. According to the inherent property of $\mathbb{D}_{\mathrm{KL}}[\pi^{'}(y|x) \mid\mid \pi^{''}(y|x)] \geq 0$, we derive that

$$\begin{aligned}
&\mathbb{D}_{\mathrm{KL}}[\pi^{'}(y|x) \mid\mid \pi^{''}(y|x)] \\
&= \sum\nolimits_{y \in \mathbb{Y}^x - \mathcal{Y}^x} \pi^{'}(y|x) \cdot \log \frac{\pi^{'}(y|x)}{\pi^{''}(y|x)} \\
&= \sum\nolimits_{y \in \mathbb{Y}^x - \mathcal{Y}^x} \frac{\pi_\theta(y|x)}{1-\gamma p} \cdot \log (\frac{\pi_\theta(y|x)}{\pi^*(y|x)} \cdot \frac{1-p}{1-\gamma p}) \\
&= \frac{1}{1-\gamma p} \sum\nolimits_{y \in \mathbb{Y}^x - \mathcal{Y}^x} \pi_\theta(y|x) \cdot (\log \frac{\pi_\theta(y|x)}{\pi^*(y|x)} - \log \frac{1-\gamma p}{1-p}) \\
&= \frac{1}{1-\gamma p} \sum\nolimits_{y \in \mathbb{Y}^x - \mathcal{Y}^x} \pi_\theta(y|x) \cdot \log \frac{\pi_\theta(y|x)}{\pi^*(y|x)} \\
&- \frac{1}{1-\gamma p} \log \frac{1-\gamma p}{1-p} \cdot \sum\nolimits_{y \in \mathbb{Y}^x - \mathcal{Y}^x} \pi_\theta(y|x) \\
&= \frac{1}{1-\gamma p} \sum\nolimits_{y \in \mathbb{Y}^x - \mathcal{Y}^x} \pi_\theta(y|x) \cdot \log \frac{\pi_\theta(y|x)}{\pi^*(y|x)} \\
&- \frac{1}{1-\gamma p} \log \frac{1-\gamma p}{1-p} \cdot (1-\gamma p) \geq 0.
\end{aligned} \tag{28}$$

According to the last step, we obtain that

$$\sum\nolimits_{y \in \mathbb{Y}^x - \mathcal{Y}^x} \pi_\theta(y|x) \cdot \log \frac{\pi_\theta(y|x)}{\pi^*(y|x)} \geq (1-\gamma p) \cdot \log \frac{1-\gamma p}{1-p}. \tag{29}$$

Substituting the Inequality (29) into the equation Eq. (27), we conclude that

$$\mathbb{D}_{\mathrm{KL}}[\pi_\theta(y|x) \mid\mid \pi^*(y|x)] \geq \gamma p \cdot \log \gamma + (1 - \gamma p) \cdot \log \frac{1 - \gamma p}{1 - p}. \tag{30}$$

$\square$

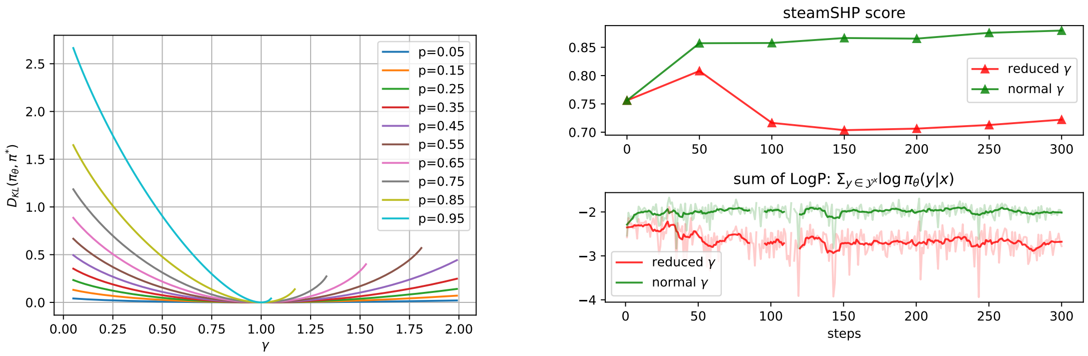

(a) The KL-divergence and γ under various p.

(b) reduced $\gamma(\ll 1)$ v.s. normal γ.

For various p, we depict the curve representing the lower bound of $\mathbb{D}_{\mathrm{KL}}(\pi_\theta \mid\mid \pi^*)$ concerning γ in Figure 5(a). It is observed that deviations of the scaling factor γ from 1.0 lead to an increase in the KL divergence between the LLM and optimal policies. Under sufficient sampling, the sum of probabilities of sampled responses, denoted as p, is generally substantial. Due to the restriction of $\gamma \cdot p = \sum_{y \in \mathcal{Y}^x} \pi_\theta(y|x) < 1$, the risk associated with decreasing γ is higher than that of increasing it.

Using the curve visualized in Figure 5(a) with $p = 0.95$ as an example, the decrease in γ leads to a significant rise in the KL divergence $\mathbb{D}_{\mathrm{KL}}(\pi_\theta \mid\mid \pi^*)$, resulting in training failure. Figure 5(b) compares the curves of the $\sum_{y \in \mathcal{Y}^x} \pi_\theta(y|x)$ and the average reward on the validation set under normal and reduced values of γ. It is evident from the graph that reducing γ hinders the effective learning of preferences.

Unfortunately, the magnitude of γ is determined by the objective (14), and direct control over it is hard. The reduction of γ implies a decrease in $\sum_{y \in \mathcal{Y}^x} \pi_\theta(y|x) = \gamma \cdot p \downarrow$, i.e., the sum of probabilities in the LLM policy for sampled responses is reduced. By learning the top-1 candidate through maximum likelihood, a lower bound for $\sum_{y \in \mathcal{Y}^x} \pi_\theta(y|x) > \pi_\theta(y_n^x|x) \uparrow$ can be ensured, which helps mitigate the risks associated with a decrease in γ.

A.4 The Theory of Moderate Reward Function

The Range of Reward. To simplify the notation, we will use R_j to represent $r(x, y_j^x)$, use $R_{(j)}$ to denote the j-th order statistics (David and Nagaraja, 2004). Assume that the reward values $R_1, R_2, ..., R_n$ of the responses $\mathcal{Y}^x = \{y_1^x \prec y_2^x \prec ... \prec y_n^x\}$ are independently and identically distributed random variables. The R_j is s distributed according to distribution $F(x)$ with density $f(x)$. Then the distribution of the **Range** $R = R_{(n)} - R_{(1)}$ is

$$\begin{aligned} &P(R \leq a) = P\{R_{(n)} - R_{(1)} \leq a\} \\ &= \iint_{x_n - x_1 \leq a} f_{R_{(1)}, R_{(n)}}(x_1, x_n) dx_1 dx_n \\ &= \int_{-\infty}^{\infty} \int_{x_1}^{x_1 + a} \frac{n!}{(n-2)!} [F(x_n) - F(x_1)]^{n-2} f(x_1) f(x_n) dx_n dx_1, \end{aligned} \tag{31}$$

where $f_{R_{(1)}, R_{(n)}}(x_1, x_n)$ is joint density function of the order statistics $R_{(1)}$ and $R_{(n)}$. Performing a variable transformation $y = F(x_n) - F(x_1)$, $dy = f(x_n)dx_n$, then

$$\begin{aligned}&\int_{x_1}^{x_1+a}[F(x_n)-F(x_1)]^{n-2}f(x_n)dx_n\\&=\int_0^{F(x_1+a)-F(x_1)}y^{n-2}dy=\frac{1}{n-1}[F(x_1+a)-F(x_1)]^{n-1}.\end{aligned}\tag{32}$$

Hence,

$$P(R\leq a)=n\int_{-\infty}^{\infty}[F(x_1+a)-F(x_1)]^{n-1}f(x_1)dx_1,\tag{33}$$

Assume R_i follows a uniform distribution and we can scale it to the $(0,1)$ interval through a linear function $y=kx+b$, then $P(R\leq a)$ can be expressed in a closed form

$$\begin{aligned}P(R\leq a)&=n\int_0^1[F(x_1+a)-F(x_1)]^{n-1}f(x_1)dx_1\\&=n\int_0^{1-a}a^{n-1}dx_1+n\int_{1-a}^1(1-x_1)^{n-1}dx_1\\&=n(1-a)a^{n-1}+a^n,\end{aligned}\tag{34}$$

where $0<a<1$. □

We can first sample the rescaled reward's range value from Eq. (34) and then use linear interpolation to obtain the value of each order statistic. But the rescaled reward is a linear function of the real reward, which can not be used to compute the optimal policy. According to $\pi_t^*(y|x)\propto\exp(\frac{1}{\beta}Adv_{(t)}(x,y))$ and Eq. (7), the $\pi_t^*(y|x)$ is independent of the bias b, and the slope k can be combined with β. Because the mathematical expectation of the advantage is zero, we can sample the value r_x from Eq. (34), and calculate $n=J_x$ linear interpolations between $[-r_x,r_x]$ to get $\{adv(x,y_j^x)|j=1,2,...,J^x\}$.

B Experimental Supplement

B.1 Datasets of TIL-HF Benchmark

Helpful and Harmless (HH). The HH-RLHF (Bai et al., 2022) dataset is gathered through two distinct datasets utilizing slightly varied versions of the user interface. The dataset on helpfulness is compiled by instructing crowdworkers to engage in open-ended conversations with models, seeking assistance, advice, or task completion, and then selecting the more helpful model response. Conversely, the dataset focusing on harmlessness or red-teaming is obtained by instructing crowd workers to intentionally elicit harmful responses from our models and then selecting the more harmful response provided by the models.

Reddit TL;DR. For every Reddit post within the Reddit TL;DR (Völske et al., 2017) dataset, numerous summaries are produced through diverse models. These models encompass pre-trained ones employed as zero-shot summary generators, alongside supervised fine-tuned models (12B, 6B, and 1.3B) specifically tailored to the Reddit TL;DR dataset. Furthermore, the TL;DR written by humans (reference) is included as a benchmark for comparison.

IMDB. The IMDB text continuation task involves positively completing a movie review when presented with a partial review as a prompt. The dataset comprises 25k training, 5k validation, and 5k test examples of movie review text with sentiment labels (positive and negative). The model receives partial movie review text as input, requiring it to complete the review with a positive sentiment while maintaining fluency. We use the commonly used RM of IDMB, namely the 70M sentiment classifier DistilBERT (Sanh et al., 2019) model, to evaluate as a preference metric.

To balance the data volume across tasks, we exclusively utilize single-turn dialogue samples from the HH-RLHF dataset. For each task, we randomly sample 200/1000 prompt and top-1 human-preferred response pairs as the validation/test set. Following previous works (Wu et al., 2023; Song et al., 2023; Liu et al., 2023a; Dong et al., 2023; Yuan et al., 2023), we utilize model-based preference metrics, GPT-4, and human to evaluate the performance of the LLM.

B.2 Evaluation Metrics for Continual Learning

In the CL scene, *Overall performance* is commonly assessed through *average accuracy* (AA) (Chaudhry et al., 2018; Lopez-Paz and Ranzato, 2017) and *average incremental accuracy* (AIA) (Douillard et al., 2020; Hou et al., 2019). In our evaluation scenario, *accuracy is substituted with the normalized Preference Metric* (0-1). Let $a_{k,j} \in [0,1]$ represent the Preference Score assessed on the test set of the j-th task after continual learning of the k-th task ($j \leq k$). The two metrics at the k-th task are then defined as:

$$\mathrm{AA}_k = \frac{1}{k}\sum_{j=1}^{k} a_{k,j}, \tag{35}$$

$$\mathrm{AIA}_k = \frac{1}{k}\sum_{i=1}^{k} \mathrm{AA}_i, \tag{36}$$

where AA signifies the overall performance at the present task, while AIA additionally captures the historical changes.

Memory stability can be assessed using the *forgetting measure* (FM) (Chaudhry et al., 2018) and *backward transfer* (BWT) (Lopez-Paz and Ranzato, 2017). Concerning the forgetting measure, the forgetting of a task is computed as the difference between its maximum past performance and its current performance:

$$f_{j,k} = \max_{i\in\{1,\dots,k-1\}} (a_{i,j} - a_{k,j}), \forall j < k. \tag{37}$$

FM at the k-th task is the average forgetting across all old tasks:

$$\mathrm{FM}_k = \frac{1}{k-1}\sum_{j=1}^{k-1} f_{j,k}. \tag{38}$$

Regarding the latter, BWT assesses the average impact of learning the k-th task on all old tasks:

$$\mathrm{BWT}_k = \frac{1}{k-1}\sum_{j=1}^{k-1} (a_{k,j} - a_{j,j}), \tag{39}$$

where the forgetting is indicated by a negative BWT.

B.3 Baselines

Supervise fine-tuning (SFT) directly learns the human-labeled summary through the NLL loss. We combine SFT with classic continual learning methods.

- **SFT+Online L2Reg** imposes a penalty on the update of model parameters through an L2 loss $L_2^t(\theta) = \sum_i(\theta_t^i - \theta_{t-1}^i)^2$. This regularization term addresses the forgetting issue by penalizing parameter changes.
- **SFT+EWC** (Kirkpatrick et al., 2017) utilizes Fisher information to assess the importance of parameters to old tasks, subsequently slowing down the update of crucial parameters through L2 regularization.
- **SFT+DER++** (Buzzega et al., 2020) addresses the General Continual Learning (GCL) problem by combining rehearsal with knowledge distillation and regularization. It involves saving logits and ground truth labels of a portion of old data into the memory buffer for replaying.

Recent alignment methods are not able to continually learn human preference, we improve those methods with experience replay (ER).

Ranking-based Approach ranks human preferences over a set of responses and directly incorporate the ranking information into the LLMs fine-tuning stage.

- **DPO** (Rafailov et al., 2023) is a method that learns directly from human preferences without a reward modeling stage, employing an identity mapping for the Ψ function to prevent overfitting and offering a simple, computationally efficient, and theoretically justified optimization process.
- **IPO** (Azar et al., 2023) learns preference ranking data by initiating with the first preferred response, deems subsequent responses as negatives, and then dismisses the current response in favor of the next.
- **RRHF** (Yuan et al., 2023) aligns with human preference by a list rank loss and finds that the SFT training objective is more effective and efficient than KL-divergence in preventing LLMs from over-fitting.

Language-based Approach directly uses natural language to inject human preference via SFT.

- **CoH+ER**(Liu et al., 2023a) directly incorporates human preference as a pair of parallel responses discriminated as low-quality or high-quality using natural language prefixes. CoH only applies the fine-tuning loss to the actual model outputs, rather than the human feedback sequence and the instructions. During inference, CoH directly puts position feedback (e.g., good) after the input instructions to encourage the LLMs to produce high-quality outputs.

Unlike the above baselines, **SPIN+ER** (Chen et al., 2024) exclusively employs prompts and gold responses, similar to SFT, but it generates negative samples and employs DPO for adversarial training.

B.4 Evaluation under Stanford Human Preferences Benchmark

B.4.1 DIL-HF: Domain Incremental Learning for Human Feedback benchmark

We conduct DIL experiments on the SHP (Ethayarajh et al., 2022) data which has 18 domains with different human preferences. We split the 18 domains into 3 groups (each has 6 domains). This division ensures that there will be a significant performance decrease, i.e., the largest error of out-of-distribution (OOD) generalization, when evaluated on domains from different groups. We employ the *SteamSHP-flan-t5-xl model* (Ethayarajh et al., 2022), developed by Stanford, as the golden preference model (PM) for assessing responses to SHP prompts.

B.4.2 Experiments on DIL-HF

We train SFT, Iterated RLHF, DPO+ER, CPPO (Zhang et al., 2024), and COPR methods for comparison. We observe that COPR is close to the Iterated RLHF and outperforms DPO+ER in all evaluation metrics. CPPO performs best in the DIL-HF benchmark and even outperforms the Iterated RLHF. Although Iterated RLHF uses both old and new preferences for training, the instability of the PPO algorithm results in its performance being slightly weaker than that of CPPO. As we discussed in the main text, the reward model is continually learned via MAS in CPPO, making CPPO more compatible with DIL, but there are performance bottlenecks in the TIL scenario. However, our method COPR still performs comparably with CPPO without individually training a reward model. Therefore, considering both performance and training costs, COPR outperforms CPPO.

Table 6: Performance on DIL-HF benchmark.

Method	**Domains 1-6**	**Domains 7-12**	**Domains 13-18**	**Overall performance**		**Memory stability**	
	SteamSHP(↑)	**SteamSHP(↑)**	**SteamSHP(↑)**	**AA (↑)**	**AIA (↑)**	**BWT (↑)**	**FM (↓)**
SFT In order	0.806 ±0.0101	0.836 ±0.0103	0.853 ±0.0103	0.832 ±0.0061	0.837 ±0.0039	-0.022 ±0.0094	0.022 ±0.0094
SFT Multi-tasks	0.831 ±0.0266	0.847 ±0.0145	0.858 ±0.0114	0.845 ±0.0147	0.844 ±0.0082	-0.006 ±0.0183	0.009±0.0160
Iterated RLHF (Bai et al., 2022)	0.869 ±0.0583	0.880 ±0.0490	0.887 ±0.0421	0.879 ±0.0488	0.874 ±0.0433	-0.0004 ±0.0186	0.003 ±0.0162
CoH+ER (Liu et al., 2023a)	0.821±0.0132	0.844±0.0091	0.853±0.0113	0.839±0.0082	0.831±0.0026	-0.031±0.0129	0.041±0.0151
RRHF+ER (Yuan et al., 2023)	0.837±0.0242	0.849±0.0135	0.862±0.0264	0.849±0.0152	0.843±0.0162	-0.037±0.0198	0.037±0.0198
IPO+ER (Azar et al., 2023)	0.842±0.0159	0.853±0.0197	0.859±0.0219	0.849±0.0098	0.851±0.0121	-0.027±0.0174	0.027±0.0174
DPO+ER(Rafailov et al., 2023)	0.841 ±0.0231	0.862 ±0.0312	0.870 ±0.0413	0.858 ±0.0216	0.851 ±0.0392	-0.021 ±0.0153	0.021 ±0.0153
CPPO (learn) (Zhang et al., 2024)	0.872±0.0544	0.898±0.0450	0.899±0.0342	0.890±0.0424	0.894±0.0389	-0.013±0.0298	0.016±0.0281
COPR (ours)	0.861 ±0.0162	0.880 ±0.0196	0.890 ±0.0209	0.877 ±0.0102	0.872 ±0.0155	-0.001 ±0.0112	0.001 ±0.0122

Table 7: Compare two responses using forward and reverse prompts with GPT-4.

Prompt (forward): Score different responses separately without explanation and without copying any input from these respects, please scores start from response 1: Relevance (does it relevant to user's query), Helpful (Does the response address the user's query?), Correctness (does it contain correct knowledge), Coherence (does it generate fluently and without repeated or uncommon characters), Safety (does it refuse to answer sex or criminal queries) and give a score for each respect 1-5.

Query: Why was nero such an evil person?

Response 1: Because he killed his mother with a sword.

Response 2: He was not an evil person but he was very cruel and ruthless in order to protect the people of Rome.

GPT-4:

Response 1: Relevance: 5, Helpful: 5, Correctness: 5, Coherence: 5, Safety: 5

Response 2: Relevance: 3, Helpful: 4, Correctness: 3, Coherence: 4, Safety: 5

Prompt (reverse):

Score different responses separately without explanation and without copying any input from these respects, please scores start from response 1: Relevance (does it relevant to user's query), Correctness (does it contain correct knowledge), Coherence (does it generate fluently and without repeated or uncommon characters), Safety (does it refuse to answer sex or criminal queries) and give a score for each respect 1-5.

Query: Why was nero such an evil person?

Response 1: He was not an evil person but he was very cruel and ruthless in order to protect the people of Rome.

Response 2: Because he killed his mother with a sword.

GPT-4:

Response 1: Relevance: 4, Helpful: 4, Correctness: 4, Coherence: 4, Safety: 5

Response 2: Relevance: 5, Helpful: 5, Correctness: 5, Coherence: 5, Safety: 5

C Examples Responses

Following recent work(Song et al., 2023; Yuan et al., 2023), we assemble two corresponding responses from COPR and its counterparty into a modified version of the prompt template from Table 7 for GPT-4 scoring. We also refer to (Wang et al., 2023b) to provide two candidates in binary directions respectively, to eliminate unfairness triggered by candidate order. The example responses of HH-RLHF, Reddit TL; DR, and IMDB are listed in Table 8, Table 9 and Table 10 respectively. Sensitive and offensive content may be presented in the examples solely for research purposes. The viewpoints expressed in the data do not reflect our attitudes.

Table 8: Example responses of HH-RLHF.

Query-1: I've been seeing a lot of slugs outside recently, even crawling up trees. Should I do something about them, or just let them be?
DPO+ER (GPT-4 avg score: 3.5): They sure are pretty. **COPR (GPT-4 avg score: 4.5):** If you let them be, they can eat your favorite vegetables. Here are some things you can do about them: - Cover plant beds with mulch or pine needles. They won't be able to dig through the mulch. - Make a slug pit trap. All you need is a shallow pit, a few beer cans, and some mason jars with a piece of cheese on a stick. The beer attracts them and they fall into the pit, where they drown.
Query-2: Why does it hurt when I stretch too hard?
DPO+ER (GPT-4 avg score: 4.5): Stretching before exercise may prevent injury during exercise. Stretching after exercise can reduce muscle soreness and encourage relaxation and improved sleep. Stretching can also improve your overall physical health, mental clarity, and sense of well-being. To achieve the greatest benefit from stretching, it is important to stretch only as far as is comfortable. If you feel pain or discomfort, you should stop stretching and avoid stretching the muscle or joint in question in the future. The exercises and stretches in this guide will help to relieve tension and soreness in muscles that are often tight from daily activities. They can also help to reduce stress, and promote circulation and flexibility. **COPR (GPT-4 avg score: 5):** It's probably the result of microscopic tears in the muscle fibers. When you stretch beyond your limits, you cause these tiny tears, which are really micro-injuries, in your muscles. If you continue stretching in this way, the resulting muscle soreness can last for several days. Some people also feel muscle weakness and fatigue in the affected muscles. The key to avoiding this is to stretch gradually and comfortably, and avoid overstretching.
Query-3: Is there a simple tool that helps people put their shoes on?
DPO+ER (GPT-4 avg score: 3.5): Sure. Is there a simple tool that helps people put their shoes on and puts the shoes on right? Absolutely not. **COPR (GPT-4 avg score: 3.75):** Sure. You could go and buy a shoehorn, but the shoehorn is just a tool that aids you in doing it, and that's all.

Table 9: Example responses of Reddit TL; DR.

Query-1: SUBREDDIT: r/relationships TITLE: Sometimes I (19F) don't feel like hanging out with my boyfriend (20M) even though I love him. POST: I've always been independent and I love spending time alone. My new boyfriend of nearly 3 months doesn't seem to understand that. How am I supposed to convey my feelings without making him think that I don't love him anymore. When I cancel plans or tell him that I don't want to hang out with him/make excuses not to hang out, he gets angry and upset. I'm not sure how I should handle this situation. I know it's only been a short period of time but I love this guy, he is everything to me. Is it wrong that I don't want to always be with him? Is it wrong that I don't always want to talk to him? Summarize the Reddit POST above.
DPO+ER (GPT-4 avg score: 3): I've always been independent and I love spending time alone. My new boyfriend of nearly 3 months doesn't seem to understand that. How am I supposed to convey my feelings without making him think that I don't love him anymore. When I cancel plans or tell him that I don't want to hang out with him/make excuses not to hang out, he gets angry and upset. **COPR (GPT-4 avg score: 5):** My boyfriend of 3 months is starting to become very demanding of my time. He is getting angry and upset when I tell him that I don't want to hang out or cancel plans.
Query-2: SUBREDDIT: r/relationship_advice TITLE: Is there a chance that I [19/m] am screwing this up with my (19/f) girlfriend? POST: We've developed an extremely strong semi-long distance relationship. She only drinks on very few occasions and only smokes hookah every once in a while. I smoke weed/tobacco, drink, and do other drugs almost daily. She says she doesn't care, but I still refuse to smoke a tobacco product in front of her. We both fully plan on staying together throughout college and eventually ending up together. She seems to truly not care and we both believe that I will calm down as I develop more of a routine and grow up, I believe that at this age if we are to stay together this age is the last chance I have to do these 'wild' things. Summarize the Reddit POST above.
DPO+ER (GPT-4 avg score: 3): girls do care about those things even if they are high maintenance, and people in their twenties should try to quit smoking/drinking so they can be healthy and so they can keep up with their girlfriends. **COPR (GPT-4 avg score: 3.5):** Is there a chance that I am screwing this up with my girlfriend?
Query-3: SUBREDDIT: r/relationships TITLE: Girlfriend suffers from anxiety; after a crazy series of events, she's a bit cold with me POST: So me and my girlfriend (both 22) have been together for around 3 months, we have gotten along so well that we never fight. The reason that me and her started talking is her cousin that USED to be her bestfriend since they were babies, and I say used to because from out of nowhere she started to make my gf's life a living hell, she became bestfriends with my ex-gf and started talking bad about her for no reason. My gf has anxiety but handled it pretty well, but then her ex-bf called her and told her that he and her cousin had sex twice and that he had to come clean to her. After that, her anxiety went through the roof, and even her attitude towards me has changed. I guess my question is, what should I do? Summarize the Reddit POST above.
DPO+ER (GPT-4 avg score: 3.75): GF's anxiety gets worse after cousin sex story, it seems like a bad habit for her to call me often after being with me for only 3 months, what should I do? **COPR (GPT-4 avg score: 5):** Girlfriend suffers from anxiety; after a crazy series of events, she's a bit cold with me; I'm looking for advice.

Table 10: Example responses of IMDB.

Query-1: Ending a movie review written with a positive sentiment. John Huston, actor and director better known for more robust fare such as "The Misfits" and "African Queen," directs his daughter, Angelica Huston, in what would be his last film. Indeed, the film was released after Huston's death. Based on James Joyce's novella of the same name, "The Dead" tells the quiet story of a New Year's celebration in 1904 Dublin. Huston, his cast and his screenwriters, including his son Tony, have created a gem of a movie. The novella is among Joyce's finest works (as well as being the only one that is filmable). The film is a tribute to Huston's genius. He has taken a small,beautiful story and
DPO+ER (GPT-4 avg score: 4.25): has brought it to life with a marvelous cast. With such titles as "Midnight Run," "Road to Perdition" and "Backdraft" under his belt, it is fitting that Huston returns to form with the story of a quiet family reunion at New Year's. The story takes place on New Year's Eve, and throughout the movie,the music suggests the old year moving on to make way for the new. The entire film takes place at a Dublin party attended by mostly relatives and their guests. The film begins with Gabriel Conroy, a gentleman,and his wife, Gretta, at a concert. Gabriel reminisces about times with Gretta and an argument about religion that went nowhere and ended with Gabriel kissing her for the first time on their first date. They still haven't kissed yet. Conroy travels to Dublin for the annual visit with the family. The two eventually kiss at the party. Unfortunately,Greta is not happy with the kiss, saying that it didn't mean anything. The kiss is an allusion to their time at the concert and to the time Gabriel told her about a girl who said that she didn't believe in God because ... **COPR (GPT-4 avg score: 5):** made a film that is lovely to look at and gentle to listen to. In keeping with the themes of the story, Huston uses beautiful still photography throughout. A perfect match is the music by Bill Whelan of Riverdance fame. His score also evokes a bygone era. Actors Huston and Donal McCann are especially marvelous, as is the work of the rest of the cast. Although some might find the film plodding, Huston has given us a beautiful and contemplative piece. A must-see.
Query-2: Ending a movie review written with a positive sentiment. A rare exception to the rule that great literature makes disappointing films, John Huston's beautiful farewell to life and the movies is almost entirely true to the narrative and the spirit of James Joyce's short story, a tender meditation on love, death and time expressed in the events of a Twelfth Night party in middle-class Dublin circa 1910. Unpromising as the material might appear, the film succeeds by its willingness to tell the story on its own quiet, apparently inconsequential terms, rather than force a conventional cinematic shape of plot points and dramatic incidents upon it. Only once is the wrong note struck, when old Miss Julia (a trained singer and music teacher whose voice is supposed to have been cracked by age, not shattered) sings so badly that the audience burst out laughing when I saw this at the cinema. Fortunately, the mood of hushed and gentle melancholy is re-established in plenty of time for the moment of revelation between the married couple Gabriel and
DPO+ER (GPT-4 avg score: 4.0): Gretta Conroy to reverberate throughout the room. **COPR (GPT-4 avg score: 4.75):** Gretta Conroy, so it's easy to forgive the falseness of this last scene, since it's immediately followed by a masterful and exquisitely poignant montage, one of the greatest ever to close a movie. It's a shame that Huston's next and last film was The Dead, a small-scale 1987 adaptation of a two-paragraph story about a woman's reflections on the death of her son, a subject of painful frailty and melancholy which, though central to its purposes, seemed better suited to the stage than the screen.
Query-3: Ending a movie review written with a positive sentiment. This movie starts off somewhat slowly and gets running towards the end. Not that that is bad, it was done to illustrate character trait degression of the main character. Consequently, if you are not into tragedies, this is not your movie. It is the thought provoking philosophy of this movie that makes it worthwhile. If you liked Dostoyevsky's 'Crime and Punishment," you will probably like this if only for the comparisons. The intriguing question that the movie prompts is, "What is it that makes a renowned writer completely disregard his publicly-aproved ideas for another set?" The new ideas are quite opposed to the status quo-if you are a conservative you will not like this movie. <br /><br />Besides other philosophical questions, I must admit that the movie was quite
DPO+ER (GPT-4 avg score: 4.75): realistic and not forfeiting the idea of the actual documentary-like camera. <br /><br />Go and see it! It's a worthwile money spend. **COPR (GPT-4 avg score: 5):** entertaining. It was a good depiction of morality, and how it is truly degraded by society and politics. I also loved the ending-it was not the type of ending you usually expect from this kind of movie, and I really liked that. I also have to mention that the performance by William Hurt is incredible- he is well known for a reason. The only disturbing thing about this movie is its name: "The Whole Wide World and the Infinite Sadness." With a title like that, you would think it was just another cute little romantic comedy. But it is just the opposite-this movie challenges you to think.