

Contrastive Policy Gradient: Aligning LLMs on sequence-level scores in a supervised-friendly fashion

Yannis Flet-Berliac[†], Nathan Grinsztajn[†], Florian Strub, Bill Wu, Eugene Choi,
Chris Cremer, Arash Ahmadian, Yash Chandak, Mohammad Gheshlaghi Azar,
Olivier Pietquin, Matthieu Geist*

Cohere

Abstract

Reinforcement Learning (RL) has been used to finetune Large Language Models (LLMs) using a reward model trained from preference data, to better align with human judgment. The recently introduced direct alignment methods, which are often simpler, more stable, and computationally lighter, can more directly achieve this. However, these approaches cannot optimize arbitrary rewards, and the preference-based ones are not the only rewards of interest for LLMs (*e.g.*, unit tests for code generation or textual entailment for summarization, among others). RL-finetuning is usually done with a variation of policy gradient, which calls for on-policy or near-on-policy samples, requiring costly generations. We introduce *Contrastive Policy Gradient*, or CoPG, a simple and mathematically principled new RL algorithm that can estimate the optimal policy even from off-policy data. It can be seen as an off-policy policy gradient approach that does not rely on important sampling techniques and highlights the importance of using (the right) state baseline. We show this approach to generalize the direct alignment method IPO (identity preference optimization) and classic policy gradient. We experiment with the proposed CoPG on a toy bandit problem to illustrate its properties, as well as for finetuning LLMs on a summarization task, using a learned reward function considered as ground truth for the purpose of the experiments.

1 Introduction

Reinforcement Learning from Human Feedback (RLHF) (Christiano et al., 2017) is a classic finetuning step intended at aligning a Large Language Model (LLM) with human judgment (Ouyang et al., 2022). The underlying principle is to learn a reward model from a preference dataset, and to optimize this reward with a regularized

Reinforcement Learning (RL) approach (Fox et al., 2015; Jaques et al., 2017; Geist et al., 2019), usually a Policy Gradient (PG) approach (Williams and Peng, 1991) or a variation like Proximal Policy Optimization (PPO) (Schulman et al., 2017). These methods require on-policy or near-on-policy samples, and thus require costly generations from the LLM. They can also be hard to tune and computationally heavy, for example through the use of an additional value network.

More recently, the field of direct alignment methods has surged with the introduction of Direct Preference Optimization (DPO) (Rafailov et al., 2023), Sequence Likelihood Calibration (SLiC-HF) (Zhao et al., 2023) or Identity Preference Optimization (IPO) (Azar et al., 2024). These approaches allow directly learning a policy optimizing for preferences, from a given preference dataset, in an offline manner and without using a proxy reward function. They are usually considered simpler, more stable, and computationally more lightweight than classic RLHF. However, by design, they cannot optimize for arbitrary reward functions.

We posit that preference-based rewards are not the only rewards worth considering when finetuning an LLM. Not everything can be measured through preferences, which are also costly to label. Such examples are using unit tests as a reward for code generation (Le et al., 2022) or a reward measuring textual-entailment for summarization (Roit et al., 2023). The aim of this paper is to propose an RL approach able to optimize an arbitrary reward while being as simple as direct alignment. It is important to note that we do not introduce any specific reward here. Our intent is to provide a convenient and efficient tool for optimizing an arbitrary reward function. In particular, we take inspiration from the contrastive learning objective, which had tremendous success in the self-supervised learning technique (Oord et al., 2018; Chen et al., 2020), and extend it to RL techniques.

[†]Equal contribution.

*Corresponding author: matthieu@cohere.com.

To this end, we introduce *Contrastive Policy Gradient*, or CoPG. It minimizes a supervised-friendly loss, of which we show the optimal policy of interest (optimizing the initial RL problem) to be the unique minimizer. It can be interpreted as a form of off-policy policy gradient, not relying on importance sampling (an approach that can easily lead the gradient variance to explode), but exploiting a specific state baseline that can be seen as a contrastive term to the reward being optimized. Our proposed approach is versatile as it regroups IPO and policy gradient as special cases. Notably, we obtain as a special case an offline off-policy generalization of Reinforce Leave-One-Out (RLOO) (Kool et al., 2019; Ahmadian et al., 2024). To illustrate its properties, we experiment with the proposed CoPG on a toy bandit problem. We also test it for finetuning an LLM on a summarization task. For this case, we train a reward model from the preference dataset and consider it the ground truth to be optimized.

2 Background

We denote a prompt x and a generation y , and we call the LLM to be trained a policy $\pi(y|x)$. We assume the prompts to be sampled according to some unknown distribution ρ . We also assume to have access to some reference model π_{ref} , typically the LLM pretrained and supervised-finetuned (SFT model), used both for initializing π and regularizing the RL problem. We consider having access to a reward model $R(x, y)$ to be maximized under π , with some regularization toward the reference model through a KL-divergence $\text{KL}(\pi(\cdot|x)||\pi_{\text{ref}}(\cdot|x))$. To lighten notations, we drop the prompt x in the main text (it appears explicitly again in the proofs in Appx. A).

The regularized RL problem consists in maximizing $J(\pi) = \mathbb{E}_{y \sim \pi}[R(y)] - \beta \text{KL}(\pi||\pi_{\text{ref}})$. Let’s write the regularized reward

$$R_{\beta}^{\pi}(y) = R(y) - \beta \ln \frac{\pi(y)}{\pi_{\text{ref}}(y)}, \quad (1)$$

the RL problem can equivalently be written as

$$J(\pi) = \mathbb{E}_{y \sim \pi}[R_{\beta}^{\pi}(y)]. \quad (2)$$

A classic approach is policy gradient, which maximizes J by gradient ascent. This is not a supervised-friendly loss because the expectation depends on the optimized policy, not on some fixed dataset of generations. The gradient is given by

$$\nabla J(\pi) = \mathbb{E}_{y \sim \pi}[R_{\beta}^{\pi}(y) \nabla \ln \pi(y)].$$

In practice, an empirical estimate of this gradient requires fresh generations from π , making it a costly method. It is common to subtract a baseline b :

$$\nabla J(\pi) = \mathbb{E}_{y \sim \pi}[(R_{\beta}^{\pi}(y) - b) \nabla \ln \pi(y)]. \quad (3)$$

This does not bias the gradient as long as it does not depend on y (because $\mathbb{E}_{y \sim \pi}[\nabla \ln \pi(y)] = 0$), and this is generally introduced to reduce the variance of the empirical gradient (Greensmith et al., 2001). A classic baseline is an estimate of the expected reward (that is the value, $b \approx \mathbb{E}_{y \sim \pi}[R_{\beta}^{\pi}(y)]$).

Objective (2) can be made supervised-friendly by relying on some fixed sampling distribution μ (e.g., underlying a dataset), by using importance sampling. Indeed, we have

$$J(\pi) = \mathbb{E}_{y \sim \mu}[\frac{\pi(y)}{\mu(y)} R_{\beta}^{\pi}(y)]. \quad (4)$$

The related gradient is, with a baseline here,

$$\nabla J(\pi) = \mathbb{E}_{y \sim \mu}[\frac{\pi(y)}{\mu(y)} (R_{\beta}^{\pi}(y) - b) \nabla \ln \pi(y)].$$

The corresponding empirical gradient suffers from a larger variance: when the policy π becomes different from the sampling model μ (which happens, as π is trained), the probability ratio can explode. Training is thus not stable or efficient. Moreover, this requires having access to the probabilities $\mu(y)$, which may not be possible, for example when the underlying generations have been made by humans rather than by an LLM.

An approach for alleviating this stability issue is to clip the probability ratio in objective (4), to kill the gradient whenever π becomes too different from μ . This is the core idea behind PPO, which considers for μ older copies of the π network and uses a value estimate as the baseline. However, this still requires fresh generations, even if they can be used for more than one gradient step.

3 Contrastive Policy Gradient

3.1 General objective

Contrastive approach requires pairs of generations, which do not need to be ranked as in RLHF. For a pair of independent generations (y, y') , we introduce the following sample loss:

$$\begin{aligned} \ell_{\text{CoPG}}(y, y'; \pi) = & \quad (5) \\ & \left(R_{\beta/2}^{\pi}(y) - R_{\beta/2}^{\pi}(y') \right) \ln \frac{\pi(y)}{\pi_{\text{ref}}(y)} \\ & + \left(R_{\beta/2}^{\pi}(y') - R_{\beta/2}^{\pi}(y) \right) \ln \frac{\pi(y')}{\pi_{\text{ref}}(y')}. \end{aligned}$$

This can be seen as a weighted log-likelihood, where the weight is the reward of the generation contrasted with the reward of an independent generation, and it's symmetric.

Let μ_1 and μ_2 be some independent distributions (for example underlying a dataset of pairs of generations), that do not need to be known analytically (contrary to policy gradient with importance sampling), and that can be the same too. The objective to be maximized is then

$$L(\pi) = \mathbb{E}_{y \sim \mu_1, y' \sim \mu_2} [\ell_{\text{CoPG}}(y, y'; \pi)]. \quad (6)$$

To get more insights, let's rewrite this objective. First, write the expected reward over μ as

$$\overline{R_{\beta/2}^{\pi}}^{\mu} = \mathbb{E}_{y \sim \mu} [R_{\beta/2}^{\pi}(y)].$$

Notice that in RL terms, this is not strictly speaking the value, as the expectation is under μ rather than π . Then, L can be rewritten as

$$\begin{aligned} L(\pi) = & \mathbb{E}_{y \sim \mu_1} \left[\left(R_{\beta/2}^{\pi}(y) - \overline{R_{\beta/2}^{\pi}}^{\mu_2} \right) \ln \frac{\pi(y)}{\pi_{\text{ref}}(y)} \right] \\ & + \mathbb{E}_{y' \sim \mu_2} \left[\left(R_{\beta/2}^{\pi}(y') - \overline{R_{\beta/2}^{\pi}}^{\mu_1} \right) \ln \frac{\pi(y')}{\pi_{\text{ref}}(y')} \right]. \end{aligned}$$

Again, this can be seen as a weighted log-likelihood, where the reward weighting the log-likelihood under one distribution is contrasted with the expected reward under the other distribution. This loss is supervised-friendly, as it does not involve sampling from the trained policy.

The natural question is whether maximizing this objective L solves the intended problem (2), and thus maximizes any language scores. The answer is positive, and we deferred all proofs to Appx. A:

Theorem 1 (CoPG solves the right problem). *Assume that π_{ref} , μ_1 and μ_2 all have the same support. Then, the unique maximizer of $L(\pi)$, defined Eq. (6), is $\pi_*(y) \propto \pi_{\text{ref}}(y) \exp \frac{R(y)}{\beta}$, which is also the unique maximizer of $J(\pi)$.*

To shed more light on the relationship to policy gradient, let's consider the gradient of $L(\pi)$. By simple calculus (taking care of the fact that $R_{\beta/2}^{\pi}$ does depend on π), one obtains:

$$\begin{aligned} \nabla L(\pi) = & \mathbb{E}_{y \sim \mu_1} \left[\left(R_{\beta}^{\pi}(y) - \overline{R_{\beta}^{\pi}}^{\mu_2} \right) \nabla \ln \pi(y) \right] \\ & + \mathbb{E}_{y' \sim \mu_2} \left[\left(R_{\beta}^{\pi}(y') - \overline{R_{\beta}^{\pi}}^{\mu_1} \right) \nabla \ln \pi(y') \right]. \quad (7) \end{aligned}$$

When compared to Eq. (3), the classic policy gradient with baseline, we obtain a sum of two policy-like gradients, however with striking differences. First, the expectation is not according to the learnt policy π , but according to either μ_1 or μ_2 , meaning that it can be understood as a sound off-policy policy gradient. Second, there is a baseline, the contrastive term, which is the expected reward but according to the other distribution (which can be the same if both are identically distributed). Crucially, it cannot be any baseline (because $\mathbb{E}_{y \sim \mu} [\nabla \ln \pi(y)] \neq 0$ in general), it must be this specific one.

Overall, the proposed objective function (6), alongside with the strong result of Thm. 1, thanks to the specific form of the gradient (7), tells us that policy gradient can be safely applied to off-policy data, without the introduction of a correcting importance sampling term, if we use the right baseline, that is the contrastive term depicted above. The relationship to policy gradient can be made even clearer for the specific case $\mu_1 = \mu_2$, to be compared again to Eq. (3):

$$\nabla L(\pi) \stackrel{\mu_1 = \mu_2}{=} 2 \mathbb{E}_{y \sim \mu} \left[\left(R_{\beta}^{\pi}(y) - \overline{R_{\beta}^{\pi}}^{\mu} \right) \nabla \ln \pi(y) \right].$$

3.2 A simple sample-based objective

To obtain a practical algorithm, one has to choose what μ_1 and μ_2 are, if they are different or the same, and how to estimate the expectations. The outer expectations can be estimated using Monte Carlo (forming a batch for each gradient step). Depending on the nature of μ_i , the inner expectations underlying the terms $\overline{R_{\beta}^{\pi}}^{\mu_i}$ can be estimated using a single (or multi-) sample Monte Carlo estimate, or even possibly by learning an associated value network. Regarding the distributions μ_1 or μ_2 , the main constraint is that they share the same prompts. It can be a dataset, generations from the current policy, generations from another policy, or coming from a replay buffer as commonly done for off-policy RL methods (Mnih et al., 2015). One

could also choose a hybrid approach, where μ_1 is for example a dataset of good but suboptimal generations while μ_2 comes from a replay buffer collecting past generations of the trained policy. This is reminiscent of RL from demonstrations, which has been shown to be beneficial in the classic RL setting (Piot et al., 2014; Hester et al., 2018).

All these choices may impact the stability and the efficiency of the resulting algorithm. We leave these interesting research directions for future works and focus here on the simple case where we learn in an offline manner from a given dataset, reminiscent of the now commonly used direct alignment methods, except that we do not need rankings.

Let $\mathcal{D} = \{(y_j, y'_j)_{1 \leq j \leq n}\}$ be a dataset of pairs of scored generations with identical prompts. CoPG minimizes the following empirical loss:

$$\hat{L}(\pi) = \frac{1}{n} \sum_{j=1}^n \ell_{\text{CoPG}}(y_j, y'_j; \pi). \quad (8)$$

with ℓ_{CoPG} being defined in Eq. (5).

It is a simple supervised-friendly objective function that can be minimized by performing gradient ascent on mini-batches sampled from the dataset. The gradient can readily be obtained by auto-differentiation (contrary to the gradient of Eq. (2), due to the dependency of the expectation to the optimized policy), but we give it for a pair of generations for completeness:

$$\begin{aligned} \nabla \ell_{\text{CoPG}}(y, y'; \pi) &= (R_\beta^\pi(y) - R_\beta^\pi(y')) \nabla \ln \pi(y) \\ &+ (R_\beta^\pi(y') - R_\beta^\pi(y)) \nabla \ln \pi(y'). \end{aligned} \quad (9)$$

From this, we observe that the optimization will increase the log-likelihood of the preferred generation (according to the reward model) and decrease that of the dispreferred one, proportionally to the reward difference.

To further simplify the practical implementation, one can easily verify that

$$\nabla \ell_{\text{CoPG}}(y, y'; \pi) \propto -\nabla (R_\beta^\pi(y) - R_\beta^\pi(y'))^2,$$

We provide the related pseudocode (used in our experiments) in Alg. 1.

4 Related works

Contrastive Policy Gradient is related to policy gradient (Williams and Peng, 1991). It can be seen

Algorithm 1: Practical offline CoPG

input : Dataset $\mathcal{D} = \{(x, y, y')\}$, reference model π_{ref} , model to train π_θ

Initialize $\pi_\theta \leftarrow \pi_{\text{ref}}$;

for $t \leftarrow 1$ **to** T **do**

 Sample a batch \mathcal{B} from \mathcal{D} ;

 Define the regularized reward;

$$R_\beta^\pi(x, y) = R(x, y) - \beta \ln \frac{\pi_\theta(y|x)}{\pi_{\text{ref}}(y|x)};$$

 Compute the gradient;

$g_t =$

$$\frac{1}{|\mathcal{B}|} \sum_{x, y, y' \in \mathcal{B}} \nabla \left(R_\beta^\pi(x, y) - R_\beta^\pi(x, y') \right)^2;$$

 Update the parameters;

$$\theta_{t+1} \leftarrow \theta_t - \eta_t \nabla g_t;$$

as a sound off-policy policy gradient, but crucially not relying on importance sampling, and thus not requiring clipping techniques such as Proximal Policy Optimization (Schulman et al., 2017) or variations (Wu et al., 2023), allowing for broader applicability (notably, PPO cannot be applied offline to a dataset of unknown density). This link is even stronger:

Property 1 (CoPG and policy gradient). CoPG generalizes policy gradient in the sense that

$$\mathbb{E}_{y \sim \pi, y' \sim \pi} [\nabla \ell_{\text{CoPG}}(y, y'; \pi)] = 2 \nabla J(\pi).$$

The expectation of the gradient of ℓ_{CoPG} according to the current policy is exactly (up to the scaling) the policy gradient of Eq. (3). So, we can retrieve policy gradient as a special case of the proposed approach, which is much less restrictive as it does not require the generations to be sampled according to the current policy. This result is asymptotic, as the expectation requires an infinite number of generations, but we have a similar connection to a more practical policy gradient approach.

Reinforce Leave-One Out is a sample-based policy gradient, using a Monte Carlo estimate of the expected reward from k generations as a baseline (Kool et al., 2019). It is remarkably effective for finetuning LLMs, simpler than PPO while providing better results, but still relying on fresh generations for each mini-batch (Ahmadian et al., 2024). When using only two generations, the gradient is naturally symmetrized not to waste information and matches exactly Eq. (9).

Property 2 (CoPG and RLOO). The sample-based gradient $\nabla \ell_{\text{CoPG}}(y, y'; \pi)$ is exactly the gradient of

RLOO for $k = 2$, when both y and y' are sampled from the current policy π .

A core difference is that CoPG is valid for any sampling distribution, while RLOO critically relies on using on-policy generations when derived from objective (2). This result shows that it is valid and principled to use the RLOO gradient in an off-policy manner. This is highly non-trivial, new to the community, and made possible thanks to the proposed principled approach.

Contrastive Policy Gradient can also be related to direct alignment methods (Rafailov et al., 2023; Zhao et al., 2023; Azar et al., 2024; Tang et al., 2024b), and more especially to Identity Policy Optimization (Azar et al., 2024), in the following sense.

Property 3 (CoPG and IPO). *For a pair of generations (y, y') , assume without loss of generality that y is preferred to y' according to the reward model, and redefine $R(y) = -R(y') = \frac{1}{4}$, then we have*

$$\begin{aligned} \nabla \ell_{\text{CoPG}}(y, y'; \pi) = & \\ & - \frac{1}{2\beta} \left(\frac{1}{2} - \beta \left(\ln \frac{\pi(y)}{\pi_{\text{ref}}(y)} - \ln \frac{\pi(y')}{\pi_{\text{ref}}(y')} \right) \right)^2, \end{aligned}$$

where the term on the right-hand side is the gradient of the sample-based IPO loss to be minimized.

These results show that if we replace the reward in our objective with a binary signal depending on which generation is preferred, we follow the same gradient as IPO. In that sense, our approach also subsumes direct alignment approaches, allowing us to optimize for an arbitrary reward.

5 Toy experiment

For an illustrative purpose, we consider a simple bandit problem, with 3 arms rewarded by $R = (2.5, 2, 1)$. We choose the data distributions to be $\mu_1 = (0.1, 0.2, 0.7)$ and $\mu_2 = (0.05, 0.05, 0.9)$. Using these distributions, we sample a dataset of 10^4 pairs of rewarded arms. We set $\beta = 0.5$ and $\pi_{\text{ref}}(y) = \frac{1}{3}$ for $y \in \{1, 2, 3\}$. The analytical solution to this bandit problem is $\pi_*(y) \propto \exp \frac{R(y)}{\beta}$.

We consider the practical CoPG objective of Eq. (8), for which we recall the gradient Eq. (9). Given that CoPG generalizes policy gradient (PG, Prop. 1), we also experiment with it to illustrate the importance of choosing the right baseline. For PG, we consider the following sample-based gradient:

$$\begin{aligned} \nabla \ell_{\text{PG}}(y, y'; \pi) = & (R_\beta^\pi(y) - b) \nabla \ln \pi(y) \\ & + (R_\beta^\pi(y') - b) \nabla \ln \pi(y'). \end{aligned}$$




Figure 1: Bandit experiment. CoPG achieves zero regret, converging to the optimal solution. IPO converges to a biased solution, as it optimizes for the expected preference. PG without a baseline has increasing regret, and PG with a value baseline converges to a biased solution.

As explained before, this gradient is valid whenever both y and y' are sampled according to the current policy π . However, here we use pairs of arms sampled from the dataset. In other words, this can be seen as a naive off-policy policy gradient. We consider two kind of baselines, $b = 0$ (no baseline) and $b = \mathbb{E}_{y \sim \pi}[R(y)]$ (value baseline). The first case corresponds to vanilla policy gradient, and the second case corresponds to the baseline most often used in the literature. Notice that in practice this should be estimated (typically with a value network), but we compute it exactly in this experiment. Given the link between CoPG and IPO (Prop. 3), we also experiment with IPO.

For each approach, we train the policy $\hat{\pi}$ with stochastic gradient descent. We use Adam (Kingma and Ba, 2014) with learning rate 10^{-3} , batches of size 512, and train for 100 epochs. We measure the performance of the trained policy with the regret:

$$\text{regret} = J(\pi_*) - J(\hat{\pi}),$$

with J being defined Eq. (2). If CoPG and PG both rely on a reward function, IPO can only use preferences. We simply set them to be sampled according to a Bradley-Terry Model, $P(y > y') = \sigma(R(y) - R(y'))$, with σ the logistic function.

Results are presented in Fig. 1. We can observe that CoPG converges to the right solution, as predicted theoretically (Thm. A). IPO converges to a biased solution. This was to be expected, as it can be shown to optimize the reward $\mathbb{E}_{y' \sim \mu_2}[\sigma(R(y) - R(y'))]$ (Azar et al., 2024), which is different from the reward of interest R . Regarding policy gradient,

we can observe that without baseline, naively applying the policy gradient on off-policy data leads to an increase in the regret: Learning deteriorates the initial policy. Adding the value baseline helps, but it still converges to a biased solution (and it is an ideal algorithm, as here the baseline is analytically computed, while it has to be estimated in practice). Sample-based CoPG converges to the right solution, showing the importance of choosing the right baseline in an off-policy context.

6 LLM experiments

In this section, we demonstrate the ability of CoPG to optimize a reward function for finetuning an LLM. As depicted in Sec. 3.2, we consider a pure offline objective, where one has to learn from a fixed dataset of pairs of generations. Classic RLHF approaches, such as policy gradient or PPO, do not work in such a pure offline setting. PPO could possibly do a single-step policy improvement, but it would require access to the underlying probabilities $\mu(y)$ of elements in the dataset, which are not available. Moreover, it would require an additional costly value network. Therefore, as baselines, we consider direct alignment methods, specifically DPO and IPO, for which the preferred completion is chosen according to the same reward model being optimized by CoPG.

Dataset. We consider the Reddit TL;DR dataset¹ of Stiennon et al. (2020). It is a summarization dataset with an SFT split, consisting of human-written summaries, and a preference split, made of human-annotated preference pairs. We will rerank the preferences according to the reward model, so that CoPG and direct alignment methods optimize for consistent objectives.

Policy Model. We use Llama2-7B as the base model² (Touvron et al., 2023). We supervise fine-tune it on the SFT split of the TL;DR dataset, giving π_{ref} as a result. This is both the initial policy and the reference model for CoPG as for direct alignment baselines. This model is trained for 2 epochs with Adam, with a cosine decay scheduler ($2 \cdot 10^{-5}$ to 0), warmup of 10%, using a batch of size 128.

Reward Model. Our objective is to provide an approach to optimize arbitrary reward functions. As a proof of concept, for this empirical study, we train a reward model using the preference split

of the TL;DR dataset and will consider it as the ground truth reward function to be optimized. We insist right away that we do not claim such a reward model to be the best thing to optimize for improving the LLM, we use it as a proxy for assessing if the proposed approach can indeed optimize a reward at scale. For training the reward function, we use a classic Bradley-Terry model (Bradley and Terry, 1952), optimizing for the loss

$$\ell_{\text{RM}}(y^+, y^-, R) = -\ln \sigma(R(y^+) - R(y^-)).$$

The reward model is trained for two epochs on the train split of the preference dataset, with Adam, the learning rate of 10^{-6} , a batch of size 128, and a warm-up of 10% of the total number of training steps. The trained reward model achieves an accuracy of 89.1% on the train set and of 72.8% on the validation set.

Training details. We train all algorithms for two epochs over the train split of the preference dataset. We use a batch of size 128. We optimize the respective losses with Adam, with a learning rate of 10^{-6} in all cases. For all approaches we use a warm-up of 10% of all training steps. For CoPG and DPO we sweep over $\beta \in \{0.01, 0.03, 0.06, 0.1, 0.3, 1\}$. For IPO we sweep over slightly lower values, specifically $\beta \in \{0.003, 0.01, 0.03, 0.06, 0.1, 0.3\}$.

Evaluation. Recall that the objective is to know if CoPG can optimize a reward function by learning offline from a fixed dataset. To evaluate this, every 50 training steps, we perform generations using the trained model on a fixed batch of 128 prompts from the validation dataset and score them using the reward model. We do the same for IPO and DPO, for which we recall that they are trained for preferences according to the reward model, and not according to the original dataset, for a fair comparison, as the reward model is used for evaluation. We also note that the reward model was not trained on the validation set, only on the train set.

6.1 Can CoPG optimize a reward?

CoPG comes with strong theoretical guarantees, and we here aim to assess its scalability while optimizing a reward in an LLM setting. We train and evaluate it in the setting depicted before, and we provide the corresponding results in Fig. 2. This figure shows how the reward of generations (averaged over 128 prompts from the validation set) from the model evolves over training (validation

¹<https://github.com/openai/summarize-from-feedback>

²<https://huggingface.co/meta-llama/Llama-2-7b-hf>




Figure 2: CoPG: Rewards of generations along training.

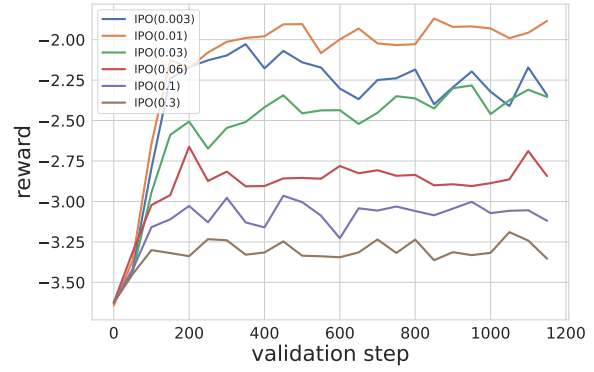


Figure 4: IPO: Rewards of generations along training.

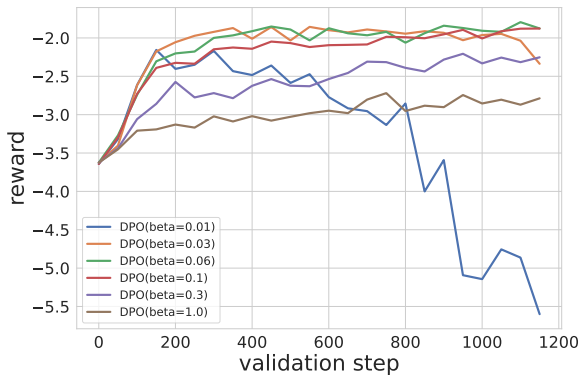


Figure 3: DPO: Rewards of generations along training.




Figure 5: Final reward as a function of β .

each 50 training steps), for the different considered values of β .

We can observe that CoPG successfully optimizes the reward over a large range of temperature $\beta \in [0.03, 0.1]$, with the higher reward being achieved for $\beta = 0.06$. When the temperature is too low, it becomes unstable, and the reward drops. Interestingly, this does not translate into a sign of overfitting in other validation metrics, such as the loss. This is not something new to the LLM community, but it highlights the necessity of doing generations when evaluating a model, which is a classic RL thing, especially in a pure offline setting. When the temperature is too high, the reward still increases, but to a lower value. This has to be expected. In this case, the Kullback-Leibler term becomes predominant, and the policy is incentivized more to avoid moving too far away from the reference model, which was also the initial policy.

For a more complete study, we also provide the related results for IPO and DPO. They do not directly aim at optimizing the reward but rather at optimizing the preferences. However, given that, in our case, these preferences are ranked according to the reward model, they should also generate

sequences of increasing rewards. For DPO, given that in this specific case the preference follows a Bradley-Terry model, it should indeed optimize the reward function, while IPO optimizes for a different objective, see Azar et al. (2024, Prop. 1 and Thm. 1) for more details.

We provide the result for DPO in Fig. 3. We can observe that DPO is indeed able to optimize the reward too. As CoPG, it is not too sensitive to the value of β , for the same range. Similarly, it becomes unstable when β is too low, and it increases less the reward when the temperature is too high because it stays closer to the reference model.

The results for IPO are provided in Fig. 4. IPO, too, increases the reward. Conversely to CoPG or DPO, it seems to be more stable, because we do not observe a significant drop when β becomes smaller (but we expect this to happen for lower values of β). It also appears to be not too sensitive to the value of β in a given range (taking into account the difference of scale without a “dropping” run), yet for lower values.

To summarize these results, we show in Fig. 5 the expected reward after training as a function of β . This showcases the stable range and the fact




Figure 6: Comparison of CoPG, DPO and IPO.

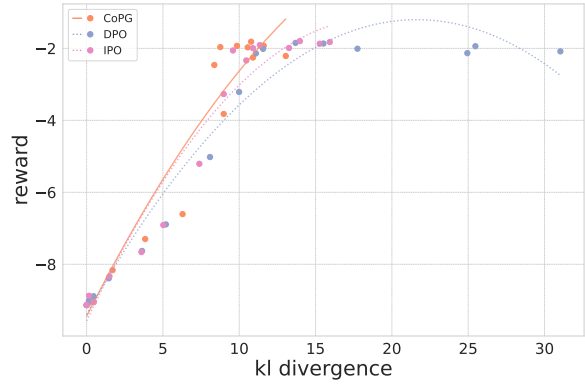


Figure 7: KL-Reward trade-off for CoPG, DPO and IPO.

that the various approaches rely on different ranges of temperature values to provide high rewards.

6.2 How does CoPG compare to direct alignment?

So far, we have shown that both CoPG and the direct alignment methods DPO and IPO were able to increase the reward in an offline manner. However, a core question is to know if directly optimizing a reward, as CoPG does, provides better results than optimizing for a preference based on this reward function. In a simple and controlled case, such as the bandit experiment of Sec. 5, the answer is clear, because we can exhibit the optimal solutions and we know to what each method should converge. However, it is much less clear in a large-scale problem, such as an LLM generating sequences.

To assess this, we compare CoPG, DPO and IPO with the best temperature from Fig. 5 (respectively $\beta = 0.06$, $\beta = 0.1$ which is also the classic value for DPO in the literature, and $\beta = 0.01$). We also rerun the experiments, this time gathering generations for a batch of 1024 prompts from the validation set, to get a better estimate of the expected reward, doing this each 100 training steps.

Results are presented in Fig. 6, the shaded envelop corresponding to the standard errors. We can observe that CoPG consistently achieves a higher reward, and faster, than both IPO and DPO. We hypothesize that it could be even more true for a reward not trained from preference data.

In this case, the preferences following the Bradley-Terry model (because ranked according to the reward function), DPO should optimize for the reward function, according to Azar et al. (2024, Prop. 1). Yet, it achieves the lowest reward in this experiment. We think that this is due to this convergence result being an asymptotical one. In practice,

CoPG uses explicitly the reward signal, while DPO only uses a binarized signal (which completions is preferred), and it would require the algorithm to observe multiple rankings of the same completions to converge to the right solution.

In principle, IPO converges to a different solution, but for the same β . For the stated theoretical results, β is part of the problem, while in practice it is a hyperparameter. Here, we have chosen β so as to achieve the highest possible reward, and the effective β for IPO is much smaller than for both CoPG and IPO. However, IPO also achieves a lower reward than CoPG, even if closer than DPO.

Overall, this experiment also suggests that if we are given a reward function to optimize, it may not be sufficient to use it to build a preference dataset with preferences ranked according to the reward model so as to learn from it using a direct alignment method. Reinforcement learning approaches still have a place in this field, and our proposed CoPG allows to maximize efficiently the reward, in a purely off-policy manner, while being as simple, stable and computationally lightweight as direct alignment approaches.

6.3 KL-Reward trade-off

In regularized RL, the objective is to maximize the reward while keeping the KL divergence between the policy and the reference model small enough. This is especially important for LLMs, as this can prevent reward hacking for example. For this experiment, we consider the HH dataset from Bai et al. (2022). We use the same recipe as TL;DR for training the SFT checkpoint (this time using the preferred examples, in the absence of SFT data), the reward model, IPO, DPO and CoPG. We also use the same hyperparameters.

We follow the same methodology, sweeping for

each algorithm over values of β , selecting the one with the higher final reward averaged of 128 generations. Then, we retrain each algorithm with its best value of β , using 1024 generations for computing the reward at each validation step. From that, we compute the KL/reward Pareto front for IPO, DPO and CoPG.

Results are presented in Fig. 7. It is a scatter plot of the rewards as a function of the KL divergences, each point corresponding to one of the 1024 generations for a specific validation step and a specific algorithm. The figure also provides a polynomial interpolation. The Pareto front clearly shows that CoPG outperforms IPO and DPO in achieving higher rewards while maintaining lower KL divergence. In particular, we note that DPO tends to achieve a higher KL more easily than other methods.

7 Discussion and perspectives

We have introduced *Contrastive Policy Gradient*, a new RL approach for finetuning LLMs. It is a form of policy gradient that contrasts the reward with a specific baseline. The corresponding objective function is supervised-friendly, in the sense that it does not (necessarily) rely on fresh generations from the model. This allows to learn a policy in a pure offline setting, without relying on importance sampling or clipping of log-probability ratios, and does not require the introduction of an additional value network. We have demonstrated that CoPG indeed optimizes for the optimal KL-regularized policy (Thm. 1), and we have shown that it generalizes policy gradient (Prop. 1), RLOO (Prop. 2), or IPO (Prop. 3).

On a controlled but simple bandit experiment (Sec. 5), we have illustrated empirically the convergence properties of the proposed approach, the importance of choosing the right baseline (the one coming from our derivations rather than the classic value baseline, and even less no baseline), and the advantage of optimizing a reward function rather than preferences derived from these rewards, which leads to a biased solution. On a larger scale LLM experiment (Sec. 6), we have shown that CoPG is able to optimize a reward function, in a fully offline and off-policy manner, conversely to other RL-finetuning approaches, and that it can achieve a higher reward than by using a direct alignment approach on preferences ranked by the reward model. Our experimental study also suggests that this higher reward is achieved at a lower KL cost,

compared to IPO or DPO.

A core perspective is to study CoPG in an online setting. Indeed, if it works in the offline setting, being the first such RL approach in the context of LLMs (to the best of our knowledge), it is not restricted to this setting. Recent works have highlighted the benefits of using online (or fresh) data for direct alignment (Tang et al., 2024a; Tajwar et al., 2024), and we hypothesize that these findings can also benefit to CoPG. Typically, one can consider using a replay buffer, as classically done in off-policy RL but less common for LLM finetuning. As CoPG uses pairs of generations, we think that this can also open new perspectives by using heterogeneous distributions for the completions, as briefly discussed in Sec. 3.2. For example, one distribution could correspond to the replay buffer, while the other could correspond to exploratory generations (for addressing the exploration/exploitation dilemma), or by using a dataset of good but suboptimal generations, in the spirit of RL from demonstrations. Another important perspective is to experiment CoPG on more tasks and rewards, and we also plan to study its possible extension to the multi-objective RL setting.

Limitations

If the proposed CoPG approach comes with strong theoretical guarantees and has been validated in both a simple bandit problem and a larger scale LLM experiments, it would benefit from being assessed on more tasks and rewards in the context of LLMs. CoPG works in a pure offline setting, which is a strength, but it would benefit from using fresh generations too, as well as from possibly heterogeneous sources of data. Nothing prevents it in principle, but this has still to be investigated, and would provide a fair comparison to other classic RL finetuning approaches, relying on (near) on-policy samples. The proposed approach optimizes for a single reward model, its extension to multiple rewards remains an interesting open question. Also, our approach assumes that the reward model is reliable, which is often not the case in practice, especially when it is learnt from data. This is in part the role of KL-regularization (to avoid hacking the reward), but our approach has no additional mechanism for preventing optimizing bad areas of the reward model. This would be especially important for a number of LLM applications.

References

- Arash Ahmadian, Chris Cremer, Matthias Gallé, Marzieh Fadaee, Julia Kreutzer, Olivier Pietquin, Ahmet Üstün, and Sara Hooker. 2024. Back to basics: Revisiting reinforce style optimization for learning from human feedback in llms. *arXiv preprint arXiv:2402.14740*.
- Mohammad Gheshlaghi Azar, Zhaohan Daniel Guo, Bilal Piot, Remi Munos, Mark Rowland, Michal Valko, and Daniele Calandriello. 2024. A general theoretical paradigm to understand learning from human preferences. In *International Conference on Artificial Intelligence and Statistics*, pages 4447–4455. PMLR.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- 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.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*.
- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martić, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.
- Roy Fox, Ari Pakman, and Naftali Tishby. 2015. Taming the noise in reinforcement learning via soft updates. *arXiv preprint arXiv:1512.08562*.
- Matthieu Geist, Bruno Scherrer, and Olivier Pietquin. 2019. A theory of regularized markov decision processes. In *International Conference on Machine Learning*, pages 2160–2169. PMLR.
- Evan Greensmith, Peter Bartlett, and Jonathan Baxter. 2001. Variance reduction techniques for gradient estimates in reinforcement learning. *Advances in Neural Information Processing Systems*, 14.
- Todd Hester, Matej Vecerik, Olivier Pietquin, Marc Lanctot, Tom Schaul, Bilal Piot, Dan Horgan, John Quan, Andrew Sendonaris, Ian Osband, et al. 2018. Deep q-learning from demonstrations. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Natasha Jaques, Shixiang Gu, Dzmitry Bahdanau, José Miguel Hernández-Lobato, Richard E Turner, and Douglas Eck. 2017. Sequence tutor: Conservative fine-tuning of sequence generation models with kl-control. In *International Conference on Machine Learning*, pages 1645–1654. PMLR.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Wouter Kool, Herke van Hoof, and Max Welling. 2019. Buy 4 reinforce samples, get a baseline for free! In *Deep Reinforcement Learning Meets Structured Prediction (ICLR Workshop)*.
- Hung Le, Yue Wang, Akhilesh Deepak Gotmare, Silvio Savarese, and Steven Chu Hong Hoi. 2022. Coderl: Mastering code generation through pretrained models and deep reinforcement learning. *Advances in Neural Information Processing Systems*, 35:21314–21328.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Bilal Piot, Matthieu Geist, and Olivier Pietquin. 2014. Boosted bellman residual minimization handling expert demonstrations. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2014, Nancy, France, September 15-19, 2014. Proceedings, Part II 14*, pages 549–564. Springer.
- 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. *Advances in Neural Information Processing Systems*, 36.
- Paul Roit, Johan Ferret, Lior Shani, Roei Aharoni, Geoffrey Cideron, Robert Dadashi, Matthieu Geist, Serkan Girgin, Leonard Hussenot, Orgad Keller, Nikola Momchev, Sabela Ramos Garea, Piotr Stanczyk, Nino Vieillard, Olivier Bachem, Gal Elidan, Avinatan Hassidim, Olivier Pietquin, and Idan Szepesktor. 2023. Factually consistent summarization via reinforcement learning with textual entailment feedback. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6252–6272.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford,

- Dario Amodei, and Paul F Christiano. 2020. Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33:3008–3021.
- Fahim Tajwar, Anikait Singh, Archit Sharma, Rafael Rafailov, Jeff Schneider, Tengyang Xie, Stefano Ermon, Chelsea Finn, and Aviral Kumar. 2024. Preference fine-tuning of llms should leverage suboptimal, on-policy data. *arXiv preprint arXiv:2404.14367*.
- Yunhao Tang, Daniel Zhaohan Guo, Zeyu Zheng, Daniele Calandriello, Yuan Cao, Eugene Tarassov, Rémi Munos, Bernardo Ávila Pires, Michal Valko, Yong Cheng, et al. 2024a. Understanding the performance gap between online and offline alignment algorithms. *arXiv preprint arXiv:2405.08448*.
- Yunhao Tang, Zhaohan Daniel Guo, Zeyu Zheng, Daniele Calandriello, Rémi Munos, Mark Rowland, Pierre Harvey Richemond, Michal Valko, Bernardo Ávila Pires, and Bilal Piot. 2024b. Generalized preference optimization: A unified approach to offline alignment. *arXiv preprint arXiv:2402.05749*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Nino Vieillard, Tadashi Kozuno, Bruno Scherrer, Olivier Pietquin, Rémi Munos, and Matthieu Geist. 2020. Leverage the average: an analysis of kl regularization in reinforcement learning. *Advances in Neural Information Processing Systems*, 33:12163–12174.
- Ronald J Williams and Jing Peng. 1991. Function optimization using connectionist reinforcement learning algorithms. *Connection Science*, 3(3):241–268.
- Tianhao Wu, Banghua Zhu, Ruoyu Zhang, Zhaojin Wen, Kannan Ramchandran, and Jiantao Jiao. 2023. Pairwise proximal policy optimization: Harnessing relative feedback for llm alignment. *arXiv preprint arXiv:2310.00212*.
- Yao Zhao, Rishabh Joshi, Tianqi Liu, Misha Khalman, Mohammad Saleh, and Peter J Liu. 2023. Slic-hf: Sequence likelihood calibration with human feedback. *arXiv preprint arXiv:2305.10425*.

A Proofs of theoretical results

In this section, we prove the stated theoretical results. First, we reintroduce the notations, now taking into account the prompt x (or context). The regularized reward defined in Eq. (1) becomes

$$R_\beta^\pi(x, y) = R(x, y) - \beta \ln \frac{\pi(y|x)}{\pi_{\text{ref}}(y|x)}. \quad (10)$$

The classic RL problem of Eq. (2) writes

$$J(\pi) = \mathbb{E}_{x \sim \rho, y \sim \pi(\cdot|x)} [R(x, y) - \beta \text{KL}(\pi(\cdot|x) || \pi_{\text{ref}}(\cdot|x))] = \mathbb{E}_{x \sim \rho, y \sim \pi(\cdot|x)} [R_\beta^\pi(x, y)]. \quad (11)$$

The unique maximizer of J is well known, J being a Legendre-Fenchel transform (e.g., see Vieillard et al. (2020, Appx. A)), it is given by

$$\pi_*(y|x) = \frac{\pi_{\text{ref}}(y|x) \exp \frac{R(x,y)}{\beta}}{Z_*(x)}, \quad (12)$$

with $Z_*(x) = \sum_y \pi_{\text{ref}}(y|x) \exp \frac{R(x,y)}{\beta}$ the associated partition function.

The proposed CoPG loss (5) with context x simply writes

$$\ell_{\text{CoPG}}(x, y, y'; \pi) = \left(R_{\beta/2}^\pi(x, y) - R_{\beta/2}^\pi(x, y') \right) \ln \frac{\pi(y|x)}{\pi_{\text{ref}}(y|x)} + \left(R_{\beta/2}^\pi(x, y') - R_{\beta/2}^\pi(x, y) \right) \ln \frac{\pi(y'|x)}{\pi_{\text{ref}}(y'|x)}.$$

The associated objective function (6) to be maximized is then, with context x ,

$$L(\pi) = \mathbb{E}_{x \sim \rho, y \sim \mu_1(\cdot|x), y' \sim \mu_2(\cdot|x)} [\ell_{\text{CoPG}}(x, y, y'; \pi)]. \quad (13)$$

Now, we restate Thm. 1 in a more general form and prove it.

Theorem 1 (CoPG solves the right problem.). *Assume that π_{ref} , μ_1 and μ_2 all have the same support (that is, for any triplet (x, y, y') such that $\rho(x) > 0$, we have $\pi_{\text{ref}}(y|x) > 0 \Leftrightarrow \mu_1(y|x) > 0 \Leftrightarrow \mu_2(y|x) > 0$). Then, the unique maximizer of $L(\pi)$, Eq. (13), is the optimal policy π_* of Eq. (12), which is also the unique maximizer of $J(\pi)$, Eq. (11).*

Proof. We start by showing that π_* is a maximizer, before proving that it is the sole one. First recall the CoPG loss:

$$\ell_{\text{CoPG}}(x, y, y'; \pi) = \left(R_{\beta/2}^\pi(x, y) - R_{\beta/2}^\pi(x, y') \right) \ln \frac{\pi(y|x)}{\pi_{\text{ref}}(y|x)} + \left(R_{\beta/2}^\pi(x, y') - R_{\beta/2}^\pi(x, y) \right) \ln \frac{\pi(y'|x)}{\pi_{\text{ref}}(y'|x)}.$$

Without loss of generality, thanks to the support assumption, we can reparametrize the policy π as follows:

$$\beta \ln \frac{\pi(y|x)}{\pi_{\text{ref}}(y|x)} = V(x, y) - Z_V(x), \quad (14)$$

with $Z_V(x) = \beta \ln \sum_y \pi_{\text{ref}}(y|x) \exp \frac{V(x,y)}{\beta}$ the associated (scaled) partition function. In essence, $V(x, y)$ can be understood as the logits of the learnt policy, shifted by the log-probabilities of the reference policy.

Then, we can rewrite the CoPG loss using the above reparametrization:

$$\begin{aligned} & \beta \ell_{\text{CoPG}}(x, y, y'; V) \\ & \stackrel{(a)}{=} \left(R(x, y) - \frac{1}{2}(V(x, y) - Z_V(x)) - R(x, y') + \frac{1}{2}(V(x, y') - Z_V(x)) \right) (V(x, y) - Z_V(x)) \\ & \quad + \left(R(x, y') - \frac{1}{2}(V(x, y') - Z_V(x)) - R(x, y) + \frac{1}{2}(V(x, y) - Z_V(x)) \right) (V(x, y') - Z_V(x)) \\ & \stackrel{(b)}{=} \left(R(x, y) - R(x, y') - \frac{1}{2}(V(x, y) - V(x, y')) \right) (V(x, y) - V(x, y')) \\ & \stackrel{(c)}{=} \frac{1}{2} (R(x, y) - R(x, y'))^2 - \frac{1}{2} (R(x, y) - R(x, y') - (V(x, y) - V(x, y')))^2. \end{aligned} \quad (15)$$

In the above derivations, (a) is true by using the reparametrization of Eq. (14), (b) is obtained by canceling terms (all terms $Z_V(x)$ are weighted by 0) and refactoring, and (c) is easily obtained by recognizing a partial square expansion in (b) (of the form $\frac{1}{2}(\Delta V)^2 - (\Delta V)(\Delta R)$).

Hence, a pointwise maximizer of $\ell_{\text{CoPG}}(x, y, y'; V)$ is necessarily a minimizer of $(R(x, y) - R(x, y') - (V(x, y) - V(x, y'))^2$ (the term $(R(x, y) - R(x, y'))^2$ being constant with respect to optimization), and $V = R$ is obviously such a minimizer, setting the square term to 0. With $V = R$, Eq. (14) characterizes the optimal policy π_* of Eq. (12). Therefore, we have just shown that π_* is a maximizer of $L(\pi)$.

Now, let us show that this maximizer is unique. Let $\tilde{\pi}$ be a maximizer of $L(\pi)$, and let \tilde{V} be an associated logit function according to Eq. (14) (notice that there is no unicity of the logits, a shift by an x -dependant function provides the same equation). The term \tilde{V} necessarily sets the square term of Eq. (15) to zero (because $V = R$ does so). Therefore, for any triplet (x, y, y') such that $\rho(x) > 0$, $\pi_{\text{ref}}(y|x) > 0$ and $\pi_{\text{ref}}(y'|x) > 0$, we have that

$$R(x, y) - \tilde{V}(x, y) = R(x, y') - \tilde{V}(x, y').$$

This is not enough to ensure unicity, $\tilde{V}(x, y) - b(x)$ would satisfy this equality for an arbitrary $b(x)$. However, we're interested in the policy solution. We have that:

$$\begin{aligned} & R(x, y) - \tilde{V}(x, y) = R(x, y') - \tilde{V}(x, y') \\ \stackrel{(a)}{\Leftrightarrow} & R(x, y) - \beta \ln \frac{\tilde{\pi}(y|x)}{\pi_{\text{ref}}(y|x)} = R(x, y') - \beta \ln \frac{\tilde{\pi}(y'|x)}{\pi_{\text{ref}}(y'|x)} \\ \stackrel{(b)}{\Leftrightarrow} & \beta \ln \pi_*(y|x) - \beta \ln \tilde{\pi}(y|x) = \beta \ln \pi_*(y'|x) - \beta \ln \tilde{\pi}(y'|x) \\ \stackrel{(c)}{\Leftrightarrow} & \pi_*(y'|x) = \frac{\pi^*(y|x)\tilde{\pi}(y'|x)}{\tilde{\pi}(y|x)} \\ \stackrel{(d)}{\Rightarrow} & 1 = \sum_{y'} \pi_*(y'|x) = \sum_{y'} \frac{\pi^*(y|x)\tilde{\pi}(y'|x)}{\tilde{\pi}(y|x)} = \frac{\pi^*(y|x)}{\tilde{\pi}(y|x)} \\ \Leftrightarrow & \tilde{\pi}(y|x) = \pi_*(y|x). \end{aligned}$$

In the above derivation, (a) is true by Eq. (14) and canceling the terms $\ln Z_{\tilde{V}}(x)$ appearing in both sides, (b) is true by recognizing from Eq. (12) that $\beta \ln \pi_{\text{ref}}(y|x) + R(x, y) = \beta \ln \pi_*(y|x) - \beta \ln Z_*(x)$ and canceling the terms $\beta \ln Z_*(x)$ appearing on both sides, (c) is true by simplifying β , exponentiating and rearranging, and (d) is true by using the fact that both $\pi_*(\cdot|x)$ and $\tilde{\pi}(\cdot|x)$ are distributions.

We have just shown that any maximizer $\tilde{\pi}$ of L is necessarily π_* , which concludes the proof. \square

Next, we restate Property 1 and prove it.

Property 1 (CoPG and policy gradient). CoPG generalizes policy gradient in the sense that

$$\mathbb{E}_{x \sim \rho, y \sim \pi(\cdot|x), y' \sim \pi(\cdot|x)} [\nabla \ell_{\text{CoPG}}(x, y, y'; \pi)] = 2 \nabla J(\pi).$$

Proof. Let start by reproofing the classic policy gradient. We have that

$$\begin{aligned} \nabla J(\pi) &= \nabla \mathbb{E}_{x \sim \rho, y \sim \pi(\cdot|x)} [R_\beta^\pi(x, y)] \\ &= \mathbb{E}_{x \sim \rho, y \sim \pi(\cdot|x)} [R_\beta^\pi(x, y) \nabla \ln \pi(y|x) + \nabla R_\beta^\pi(x, y)] \\ &= \mathbb{E}_{x \sim \rho, y \sim \pi(\cdot|x)} [R_\beta^\pi(x, y) \nabla \ln \pi(y|x)], \end{aligned}$$

where for the last step we make use of the fact that $\mathbb{E}_{x \sim \rho, y \sim \pi(\cdot|x)} [\nabla \ln \pi(y|x)] = 0$.

Now, let compute the gradient of the CoPG loss:

$$\begin{aligned}
& \nabla \ell_{\text{CoPG}}(x, y, y'; \pi) \\
&= \nabla \left(\left(R_{\beta/2}^\pi(x, y) - R_{\beta/2}^\pi(x, y') \right) \ln \frac{\pi(y|x)}{\pi_{\text{ref}}(y|x)} + \left(R_{\beta/2}^\pi(x, y') - R_{\beta/2}^\pi(x, y) \right) \ln \frac{\pi(y'|x)}{\pi_{\text{ref}}(y'|x)} \right) \\
&= \nabla \left(R_{\beta/2}^\pi(x, y) - R_{\beta/2}^\pi(x, y') \right) \ln \frac{\pi(y|x)}{\pi_{\text{ref}}(y|x)} + \left(R_{\beta/2}^\pi(x, y) - R_{\beta/2}^\pi(x, y') \right) \nabla \ln \frac{\pi(y|x)}{\pi_{\text{ref}}(y|x)} \\
&\quad + \nabla \left(R_{\beta/2}^\pi(x, y') - R_{\beta/2}^\pi(x, y) \right) \ln \frac{\pi(y'|x)}{\pi_{\text{ref}}(y'|x)} + \left(R_{\beta/2}^\pi(x, y') - R_{\beta/2}^\pi(x, y) \right) \nabla \ln \frac{\pi(y'|x)}{\pi_{\text{ref}}(y'|x)} \\
&= \left(R_{\beta/2}^\pi(x, y) - R_{\beta/2}^\pi(x, y') \right) \nabla \ln \pi(y|x) + \left(R_{\beta/2}^\pi(x, y') - R_{\beta/2}^\pi(x, y) \right) \nabla \ln \pi(y'|x). \tag{16}
\end{aligned}$$

It is important to not ignore the fact that $R_{\beta/2}^\pi$ does depend on π , and thus contributes to the gradient, the rest of derivations skipped above are simple calculus and rearranging terms.

So, the gradient $\nabla \ell_{\text{CoPG}}(x, y, y'; \pi)$ is a sum of two terms, let focus on the first one. We have that

$$\begin{aligned}
& \mathbb{E}_{x \sim \rho, y \sim \pi(\cdot|x), y' \sim \pi(\cdot|x)} \left[\left(R_{\beta/2}^\pi(x, y) - R_{\beta/2}^\pi(x, y') \right) \nabla \ln \pi(y|x) \right] \\
&= \mathbb{E}_{x \sim \rho, y \sim \pi(\cdot|x)} \left[R_{\beta/2}^\pi(x, y) \nabla \ln \pi(y|x) \right] - \mathbb{E}_{x \sim \rho} \left[\mathbb{E}_{y' \sim \pi(\cdot|x)} \left[R_{\beta/2}^\pi(x, y') \right] \underbrace{\mathbb{E}_{y \sim \pi(\cdot|x)} \left[\nabla \ln \pi(y|x) \right]}_{=0} \right] \\
&= \nabla J(\pi).
\end{aligned}$$

By symmetry, we have exactly the same result for the second term,

$$\mathbb{E}_{x \sim \rho, y \sim \pi(\cdot|x), y' \sim \pi(\cdot|x)} \left[\left(R_{\beta/2}^\pi(x, y') - R_{\beta/2}^\pi(x, y) \right) \nabla \ln \pi(y'|x) \right] = \nabla J(\pi),$$

which overall shows that

$$\mathbb{E}_{x \sim \rho, y \sim \pi(\cdot|x), y' \sim \pi(\cdot|x)} \left[\nabla \ell_{\text{CoPG}}(x, y, y'; \pi) \right] = 2 \nabla J(\pi).$$

This concludes the proof. □

Then, we restate Prop. 2 and prove it.

Property 2 (CoPG and RLOO). *The sampled-based gradient $\nabla \ell_{\text{CoPG}}(x, y, y'; \pi)$ is exactly the gradient of RLOO for $k=2$, when both y and y' are sampled from the current policy π .*

Proof. First, recall the gradient of the CoPG loss from Eq. (16), proven in the previous proof:

$$\nabla \ell_{\text{CoPG}}(x, y, y'; \pi) = \left(R_{\beta/2}^\pi(x, y) - R_{\beta/2}^\pi(x, y') \right) \nabla \ln \pi(y|x) + \left(R_{\beta/2}^\pi(x, y') - R_{\beta/2}^\pi(x, y) \right) \nabla \ln \pi(y'|x).$$

Next, we rederive RLOO from first principle. Recall the classic policy gradient:

$$\nabla J(\pi) = \mathbb{E}_{x \sim \rho, y \sim \pi(\cdot|x)} \left[R_{\beta/2}^\pi(x, y) \nabla \ln \pi(y|x) \right].$$

A sample-based gradient is given by, with y being sampled according to $\pi(\cdot|x)$,

$$\hat{\nabla} J(\pi) = R_{\beta/2}^\pi(x, y) \nabla \ln \pi(y|x).$$

As explained before, a baseline $b(x)$ can be considered, without biasing the gradient, as long as it is independent from the generation y :

$$\hat{\nabla}_b J(\pi) = \left(R_{\beta/2}^\pi(x, y) - b(x) \right) \nabla \ln \pi(y|x).$$

It is easy to check that the gradient is unbiased:

$$\mathbb{E}_{y \sim \pi(\cdot|x)} \hat{\nabla}_b J(\pi) = \underbrace{\mathbb{E}_{y \sim \pi(\cdot|x)} \left[R_{\beta/2}^\pi(x, y) \nabla \ln \pi(y|x) \right]}_{=\nabla J(\pi)} - b(x) \underbrace{\mathbb{E}_{y \sim \pi(\cdot|x)} \left[\nabla \ln \pi(y|x) \right]}_{=0} = \nabla J(\pi).$$

The principle of RLOO is to perform k independent generations y^1, \dots, y^k for each prompt x , using the current policy $\pi(\cdot|x)$, and to use as a stochastic baseline for $R_\beta^\pi(x, y^j)$, more specifically the leave-one-out empirical expectation of the reward using the $k - 1$ other generations. This is still a valid baseline (derivation above applies), as even if the baseline is stochastic, it is independent from y^j . The corresponding empirical gradient is

$$\hat{\nabla}_k J(\pi) = \sum_{j=1}^k \left(R_\beta^\pi(x, y^j) - \frac{1}{k-1} \sum_{\substack{l=1 \\ l \neq j}}^k R_\beta^\pi(x, y^l) \right) \nabla \ln \pi(y^j|x).$$

In the case $k = 2$ this simplifies to:

$$\hat{\nabla}_{k=2} J(\pi) = (R_\beta^\pi(x, y^1) - R_\beta^\pi(x, y^2)) \nabla \ln \pi(y^1|x) + (R_\beta^\pi(x, y^2) - R_\beta^\pi(x, y^1)) \nabla \ln \pi(y^2|x).$$

This is exactly the gradient $\nabla \ell_{\text{CoPG}}(x, y, y'; \pi)$, which proves the result. However, as explained in the main text, it is crucial to note that RLOO derivation is only valid when generations are done with the current policy, while CoPG can account for arbitrary generations. In this sense, CoPG says that RLOO can be safely used in an off-policy context. \square

Eventually, we restate and prove Prop. 3.

Property 3 (CoPG and IPO). *For a prompt x and a pair of generations (y, y') , assume without loss of generality that y is preferred to y' given x according to the reward model, that is $R(x, y) > R(x, y')$, and redefine $R(x, y) = -R(x, y') = \frac{1}{4}$, then we have*

$$\nabla \ell_{\text{CoPG}}(x, y, y'; \pi) = -\frac{1}{2\beta} \left(\frac{1}{2} - \beta \left(\ln \frac{\pi(y|x)}{\pi_{\text{ref}}(y|x)} - \ln \frac{\pi(y'|x)}{\pi_{\text{ref}}(y'|x)} \right) \right)^2,$$

where the term on the right-hand side is the gradient of the sample-based IPO loss to be minimized.

Proof. First, from Eq. (16), we have that

$$\nabla \ell_{\text{CoPG}}(x, y, y'; \pi) = (R_\beta^\pi(x, y) - R_\beta^\pi(x, y')) \nabla \ln \pi(y|x) + (R_\beta^\pi(x, y') - R_\beta^\pi(x, y)) \nabla \ln \pi(y'|x).$$

Given the assumptions (y preferred to y' given x and binarized reward, that is redefine $R(x, y) = -R(x, y') = \frac{1}{4}$), and given the definition of R_β^π in Eq. (10), the gradient writes

$$\nabla \ell_{\text{CoPG}}(x, y, y'; \pi) = \left(\frac{1}{2} - \beta \ln \frac{\pi(y|x)}{\pi_{\text{ref}}(y|x)} + \beta \ln \frac{\pi(y'|x)}{\pi_{\text{ref}}(y'|x)} \right) (\nabla \ln \pi(y|x) - \nabla \ln \pi(y'|x)).$$

Now, let consider the gradient of the sample-based IPO loss:

$$\begin{aligned} & \nabla \left(\frac{1}{2} - \beta \left(\ln \frac{\pi(y|x)}{\pi_{\text{ref}}(y|x)} - \ln \frac{\pi(y'|x)}{\pi_{\text{ref}}(y'|x)} \right) \right)^2 \\ &= 2 \left(\frac{1}{2} - \beta \left(\ln \frac{\pi(y|x)}{\pi_{\text{ref}}(y|x)} - \ln \frac{\pi(y'|x)}{\pi_{\text{ref}}(y'|x)} \right) \right) (\beta \nabla \ln \pi(y'|x) - \beta \nabla \ln \pi(y|x)) \\ &= -2\beta \nabla \ell_{\text{CoPG}}(x, y, y'; \pi). \end{aligned}$$

This proves the stated result. \square

- *Step=300*: I hate admitting I'm wrong to my bf in arguments, he always acts super smug and I can't stand it. How can I get better about this?
- *Step=600*: I hate admitting I'm wrong to my bf and I hate when he acts super smug and condescending when he knows he has proven he was right and I was wrong. I need to learn how to not let this get to me!
- *Step=900*: I hate admitting defeat in petty arguments to my bf, he always acts so smug and I can't stand it, but I don't know how to improve so I don't get as upset with him over it!
- *Step=1200*: I hate admitting that I'm wrong to my bf after a disagreement because he acts like a superior being, I don't want to back down and admit it to him, but I also don't want to continue the argument because it just gets more heated and nasty.

SUBREDDIT: relationships

TITLE: Me [19 F] with my boyfriend [19 M] 3.5 years, when is it time to let go? with line breaks.

POST: Okay, so I've never really posted anything before, so sorry if this is weird or I'm not doing it right or something. I started dating my current BF when I was 15 and a sophomore in high school, he was 16 and a junior. Now we are both in college. We are insanely close. Like spend literally ALL our time together. He is incredibly important to me and we do everything together. I love him a lot and I have incredible respect for him as a person. The problem is we both don't have friends outside of the relationship. I mean we both have friends. We have a lot of friends really, but we just don't feel like we connect with anyone other than each other anymore. It's hard to explain. We both ditched all of our friends from high school. Some of it was because of our relationship, but a lot of it was high school is shitty.

DPO:

- *Step=300*: I know I need to break up with my boyfriend of 3.5 years, but I'm scared because I feel like I won't ever connect with anyone again because we have kind of abandoned our friends for each other.
- *Step=600*: I know I need to break up with my boyfriend of 3.5 years, but I feel like I won't ever connect with anyone else because I've been so focused on this relationship my whole life. I don't know what to do.
- *Step=900*: I know I need to break up with my boyfriend I've been with for 3.5 years, but I'm afraid because I feel like I won't ever connect with anyone again after him. I feel like I only connect with him. I don't feel like I connect with anyone else.
- *Step=1200*: My boyfriend and I have pretty much zero outside friends and we are too dependent on each other, I know we should break up but I'm scared I will never connect with anyone else. I'd really appreciate some advice. Thanks.

IPO:

- *Step=300*: I know I need to break up with my boyfriend of 3.5 years, but I'm scared because I feel like I won't ever connect with anyone outside of my SO now. Has anyone expressed something like this?
- *Step=600*: I know I need to end my relationship with my boyfriend whom I love very dearly, but I'm scared because I feel like I won't ever connect with anyone outside of the relationship anymore if I do.

- *Step=900*: I know it's time to move on from my long-term relationship, but I'm scared because I don't feel like I'll ever connect with anyone again outside of my SO. Has anyone else expressed this feeling before? If so, what did you do?
- *Step=1200*: My boyfriend and I have pretty much ditched all of our friends from high school and replaced them with each other. We know we need to break up eventually, but I'm scared because I don't feel like I'll ever connect with anyone else. Has anyone else experienced something like this?

CoPG:

- *Step=300*: I know I need to break up with my boyfriend of 3.5 years whom I love very dearly, but I'm scared because I feel like I won't ever connect with anyone else outside of this relationship. Has anyone expressed this feeling before?
- *Step=600*: I know I need to end my relationship with my boyfriend whom I love very dearly, but I am scared I'll never find anyone I connect with after him because I feel like I don't connect with anyone other than him.
- *Step=900*: I know I need to end my long-term relationship, but I'm scared because I don't feel like I'll ever be able to connect with anyone else because all of my friends currently are either from this relationship or I made this year and don't feel like I connect with them.
- *Step=1200*: My boyfriend and I are pretty much all each other have ever cared about or focused on in terms of friends. I know I need to break up with him, but I'm scared I won't find any friends I connect with ever again.