
ShiQ: Bringing back Bellman to LLMs

Pierre Clavier¹, Nathan Grinsztajn², Raphael Avalos¹,
Yannis Flet-Berliac¹, Irem Ergun¹, Omar D. Domingues¹, Olivier Pietquin^{2,3},
Pierre H. Richemond¹, Florian Strub¹, and Matthieu Geist^{2,3}

¹Cohere

²Work done at Cohere

³Earth Species Project

pierre.clavier@cohere.com ; matthieu@earthspecies.org

Abstract

The fine-tuning of pre-trained large language models (LLMs) using reinforcement learning (RL) is generally formulated as direct policy optimization. This approach was naturally favored as it efficiently improves a pretrained LLM, seen as an initial policy. Another RL paradigm, Q-learning methods, has received far less attention in the LLM community while demonstrating major success in various non-LLM RL tasks. In particular, Q-learning effectiveness comes from its sample efficiency and ability to learn offline, which is particularly valuable given the high computational cost of sampling with LLM. However, naively applying a Q-learning-style update to the model’s logits is ineffective due to the specificity of LLMs. Our core contribution is to derive theoretically grounded loss functions from Bellman equations to adapt Q-learning methods to LLMs. To do so, we carefully adapt insights from the RL literature to account for LLM-specific characteristics, ensuring that the logits become reliable Q-value estimates. We then use this loss to build a practical algorithm, ShiQ for Shifted-Q, that supports off-policy, token-wise learning while remaining simple to implement. Finally, we evaluate ShiQ on both synthetic data and real-world benchmarks, *e.g.*, UltraFeedback, BFCL-V3, demonstrating its effectiveness in both single-turn and multi-turn LLM settings.

1 Introduction

Reinforcement Learning (RL) is commonly used for fine-tuning Large Language Models (LLMs). A standard objective is to align the model with human preferences. To achieve this, a reward model is first trained on preference data and then used to guide the optimization of the language model through RL [Christiano et al., 2017, Ouyang et al., 2022]. A simpler and less costly alternative is provided by direct alignment methods [Zhao et al., 2023, Rafailov et al., 2023, Azar et al., 2024], which directly train a policy on preference data, without relying on a proxy reward. However, other rewards are of interest for RL fine-tuning. For example, successful unit tests can be used as a reward for code generation [Le et al., 2022] or a textual-entailment classifier can be used as a reward for summarization [Roit et al., 2023]. In this work, we consider the general problem of RL fine-tuning, without any assumptions about the target task of the reward.

RL fine-tuning is usually framed as maximizing the expected cumulative reward, regularized with some reference model or policy obtained from a previous training phase. Given this classical objective, it is natural to optimize it using gradient ascent, that is, policy-gradient. Moreover, in a fine-tuning context, it is highly desirable to start from the model reference policy, which further justifies policy-based approaches. REINFORCE [Williams and Peng, 1991] and variants, *e.g.*, [Kool et al., 2019], as well as Proximal Policy Optimization (PPO) [Schulman et al., 2017], are standard

approaches for optimizing this objective, especially in the context of LLMs Roit et al. [2023], Ahmadian et al. [2024], Ouyang et al. [2022].

However, policy gradient approaches come with drawbacks. Notably, they are inherently on-policy, meaning each gradient update requires sampling new completions, a very costly operation when training LLMs. This can be mitigated through techniques like importance sampling *e.g.*, [Degris et al., 2012]. However, this approach introduces two significant challenges: it results in high variance and of knowing the data completions probabilities. Adopting a contextual bandit perspective (seeing each possible LLM completion as an arm) allows for bypassing the need for importance sampling, mostly by exploiting the known, *softmax* analytical form of the optimal policy. This is the case of *direct* alignment methods, *e.g.*, [Zhao et al., 2023, Rafailov et al., 2023, Azar et al., 2024], which directly learn a policy from preference data, but sidestep and do not address the general reward optimization problem. Other approaches in the bandit setting, such as Direct Reward Optimization (DRO) [Richemond et al., 2024] or Contrastive Policy Gradient (CoPG) [Flet-Berliac et al., 2024], directly optimize the reward in an off-policy manner without relying on importance sampling. These approaches are effective, but also come with possible drawbacks. First, they are fundamentally incapable of processing token-wise reward signals, even when such signals are available. Second, these methods necessitate careful consideration of sequence-level losses. For instance, they often involve critical algorithmic decisions, such as whether to average losses across sequences or not [Meng et al., 2024, Grinsztajn et al., 2024]).

An alternative approach consists of modeling LLMs as regularized Markov decision processes (MDP) Geist et al. [2019], then relying on Bellman equations to design a loss inspired by Q-Learning, which notably allows for off-policy token-wise learning or multi-turn learning. In order to achieve this, one can interpret the logits of the LLM seen as an autoregressive policy as *Q-values*. However, a naive application of an algorithm such as DQN [Mnih et al., 2015] or a regularized variation [Vieillard et al., 2020b] would not be very efficient, since it would ignore key characteristics of LLMs. We identify three important ones below. First, RL learning methods often rely on multiple networks - up to five for actor critics in the twin critic approach [Fujimoto et al., 2018] - and multiplying huge networks like LLMs is not desirable, as it strains hardware resources and incurs wasteful memory consumption. Importantly, the same holds at inference time; we would like the *learned policy to simply be the softmax over the logits*, and not to rely on further transformations, possibly involving additional networks with the associated latency and hardware costs. Second, initialization is also a crucial factor to consider when fine-tuning. If the reference model is a good candidate for optimizing the RL objective, it is much less obvious that the logits of this reference model are a good initialization for the *Q-values* of a Bellman-based loss, while there is no other apparent choice. Third and finally, the majority of RL off-policy algorithms also rely on bootstrapping, which can slow down learning in the case of sparse rewards, a very common setting for LLMs (many rewards being sequence-level, and could indeed be called returns). In this paper, we frame LLMs as regularized MDPs by overcoming the aforementioned challenges. Specifically, we seek to answer the following question:

Is it possible to derive a theoretically grounded Q-learning-based loss for LLMs allowing sequence-level learning, whose policy is given by a softmax over the model logits, and to incorporate LLM-specific considerations to improve empirical performance ?

Firstly, our core contribution is to propose a sequence of Bellman consistency equations leading to the same optimal policy of interest, each of these equations will tackle the aforementioned specificities of LLMs. **Secondly**, we use the resulting Bellman consistency equation to build a simple and practical off-policy and token-level loss, inspired by Q-Learning, that we call ShiQ for Shifted-Q. Crucially, ShiQ relies on *single-trajectory* data based on an individual prompt-response-reward Richemond et al. [2024], rather than typical pairwise preference data Rafailov et al. [2023], Richemond et al. [2024]. **Thirdly**, we evaluate ShiQ on synthetic datasets to characterize the algorithm’s behavior under fine-grained reward structures. We then benchmark its performance on real-world tasks, demonstrating its effectiveness in single-turn *e.g.* UltraFeedback and Harmful-Harmless Datasets and especially in multi-turn LLM scenarios on BFCL-V3.

Related work: *Off policy algorithm within the bandit framework* Flet-Berliac et al. [2024] derive an off-policy bandit method without importance sampling by modeling the LLM as a bandit and introducing contrastive policy-gradient (CoPG). Similarly, Richemond et al. [2024] treat the LLM as a bandit and proposes direct reward optimization (DRO), an actor-critic approximation of the intractable solution to problem (12) that jointly learns a policy and a value network, unlike our approach.


Modeling the logits of the LLM as Q-values has been explored by Guo et al. [2022], who apply path consistency learning (PCL) [Nachum et al., 2017] to logits with a non-necessary target network, incurring extra memory overhead, whereas our ablation $\text{ShiQ}_{\text{init}}$ presented in subsection 2.2, leveraging better initialization, generalizes their method without it. Yu et al. [2024] similarly interpret logits as Q-values and highlight poor reference-policy initialization, but their Bellman-coder relies on a more complex, costlier dueling architecture with an additional value network and offers less theoretical grounding.

Multi-turn RL algorithms like Rafailov et al. [2024] extend Direct Preference Optimization (DPO) Rafailov et al. [2023] to multi-turn interactions. Still, their method depends on paired trajectories, whereas ours requires only unranked ones. Similarly, Ji et al. [2024] introduce an offline Soft Actor-Critic that directly optimizes a Q-function via importance-weighted updates. However, this is prone to high variance and it trains both policy and value networks, in contrast to our policy-only approach. An exhaustive related work can be found in Appendix B.

2 Method

In this section, we outline the three principal components of our method culminating in the ShiQ algorithm. In Sec. 2.1, we adopt RL notations to derive the Bellman consistency equations. We start with soft Q-learning consistency equation and the associated naive Q-learning loss, L_{try1} . Then, we use the following three transformations to take into account LLMs specificity while preserving the theoretical guarantee of computing the optimal policy:

1. *Easing sampling* (Sec. 2.2), yielding loss L_{try2} : eliminates the need to load and infer on both the learned and reference models and to store the temperature parameter.
2. *Improved initialization* (Sec. 2.3), yielding loss L_{try3} : leverages the reference policy for a smarter Q-learning start. The corresponding ablation, $\text{ShiQ}_{\text{init}}$, is detailed in Appendix A.
3. *Multi-step extension* (Sec. 2.4), yielding loss L_{ShiQ} : propagates rewards more effectively across multiple steps. The ablation without this extension, ShiQ_{ms} , is presented in Appendix A.



Note that we did not perform ablations for step 1 due to its high computational cost. Finally, in Sec. 3, we restate the algorithm using LLM notation to simplify the implementation.

2.1 LLMs are MDPs

Consider a prompt x and a completion y , we can model a state as a subsequence $s_t^{xy} = (x, y_{<t})$, and an action as the chosen token $a_t^{xy} = y_t$. The initial state $s_1^{xy} = x$ is the prompt. The next state is deterministically the concatenation of the current state and the action, $s_{t+1}^{xy} = s_t^{xy} \oplus a_t = (x, y_{\leq t})$. For lighter notation, we will drop the upper-script xy when context is clear, and, for example, write s_t for s_t^{xy} . We write the discount factor as $\gamma \in (0, 1]$, which can be safely set to $\gamma = 1$ since we consider a finite-horizon setting. We also assume access to a token-wise reward function, assigning a scalar $r(s_t, a_t)$ to each state-action pair. Therefore, we have framed an LLM as an MDP. The state space $\mathcal{S}$ is the set of all subsequences of maximal length $T_{\max}$ and not having an eos token. The action space $\mathcal{A}$ is the vocabulary $\mathcal{V}$ except possibly at the end with an eos token. The transition kernel is deterministic¹, by concatenating states and actions. The discount factor $\gamma \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ depends on the state-action couple, and we are given a token-wise reward function $r \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$. From this, with $|y| < T_{\max}$ the length of the sequence, we can define the return of a completion y for a prompt x as

$$R(x, y) = \sum_{t=1}^{|y|} \gamma^{t-1} r(s_t^{xy}, a_t^{xy}). \quad (1)$$

¹Notice that our results generally hold for stochastic transition kernels too, and the overall contribution can be applied to other RL fine-tuning problems, for example in robotics.

We consider the same policy π as before, $\pi(y|x) = \prod_{t=1}^{|y|} \pi(y_t|x, y_{<t}) = \prod_{t=1}^{|y|} \pi(a_t^{xy}|s_t^{xy})$. The objective is to maximize

$$J_{\text{rl}}(\pi) = \mathbb{E}_{x \sim \rho} \mathbb{E}_{y \sim \pi(\cdot|x)} \left[\sum_{t=1}^{|y|} \gamma^{t-1} \left(r(s_t^{xy}, a_t^{xy}) - \beta \ln \frac{\pi(a_t^{xy}|s_t^{xy})}{\pi_{\text{ref}}(a_t^{xy}|s_t^{xy})} \right) \right]. \quad (2)$$

If we only have access to a sequence-level reward $R(x, y)$ in an LLM setting, we can set $\gamma = 1$ and can define the token-level reward as

$$r(s_t^{xy}, a_t^{xy}) = \begin{cases} R(x, y) & \text{if } a_t^{xy} = \text{eos} \text{ or } t = T_{\max}, \\ 0 & \text{else.} \end{cases} \quad (3)$$

Notice that this is a finite-horizon MDP, for which we know the optimal policy to be non-stationary, but the state contains the time information. Eq. (2) is a strict generalization of the objective function (12) defined after in the LLM notation section. We could solve objective (2) using a policy-gradient approach [Ahmadian et al., 2024, Ouyang et al., 2022], or even a bandit-based approach [Richemond et al., 2024, Flet-Berliac et al., 2024], by considering $\gamma = 1$ and the sequence-level reward of Eq. (1). However, we could also exploit the temporal structure by relying on Bellman equations. To do so, we introduce a state-action dependent discount factor to account for the fact that we work in a finite-horizon MDP $\gamma(s_t, a_t) = 0$ if $a_t = \text{eos}$ or $t = T_{\max}$, otherwise $\gamma(s_t, a_t) = \gamma$.

We can rely on the classic (regularized) Bellman optimality operator to get the optimal policy. In all stated results, we say that a transition (s_t, a_t, s_{t+1}) is *admissible* if it can occur by sampling $x \sim \rho$ and $y \sim \pi_{\text{ref}}(\cdot|x)$, that is, with $s_t = (s_1, a_1, a_2, \dots, a_{t-1})$ (by definition), $\rho(s_1) > 0$ and $\pi_{\text{ref}}(a_{1:t}|s_1) > 0$. Notice that when $\gamma(s_t, a_t) = 0$, s_{t+1} is a dummy state but its value will never be evaluated. Full proofs are deferred to Appx. C.

Theorem 1. Let $q \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ be the unique function satisfying, for any admissible (s_t, a_t, s_{t+1}) ,

$$q(s_t, a_t) = r(s_t, a_t) + \gamma(s_t, a_t) \beta \ln \sum_{a' \in \mathcal{A}} \pi_{\text{ref}}(a'|s_{t+1}) \exp \frac{q(s_{t+1}, a')}{\beta}. \quad (4)$$

Then, the unique optimal policy maximizing (2) satisfies

$$\pi_*(a_t|s_t) = \frac{\pi_{\text{ref}}(a_t|s_t) \exp \frac{q(s_t, a_t)}{\beta}}{\sum_{a \in \mathcal{A}} \pi_{\text{ref}}(a|s_t) \exp \frac{q(s_t, a)}{\beta}}.$$

As a first candidate objective using Thm. 1, we could interpret q as the logits of the LLM, and design a loss function such that the minimizer satisfies Bellman equation (4):

$$L_{\text{try1}}(q) = \mathbb{E}_{x, y \sim \mathcal{D}} \left[\sum_{s_t, a_t \in (x, y)} \left(r(s_t, a_t) + \gamma(s_t, a_t) \beta \ln \sum_{a' \in \mathcal{A}} \pi_{\text{ref}}(a'|s_{t+1}) \exp \frac{q(s_{t+1}, a')}{\beta} - q(s_t, a_t) \right)^2 \right].$$

A direct corollary of Thm. 1 is that if $\text{supp}(\mathcal{D}) = \text{supp}(\rho \pi_{\text{ref}})$ i.e the dataset and $\rho \pi_{\text{ref}}$ have same support, where the last notation depicts $x \sim \rho$ and $y \sim \pi_{\text{ref}}(\cdot|x)$, then the unique minimizer q_* of $L_{\text{try1}}(q)$ satisfies $\pi_{q_*} = \pi_*$ as under the support assumption, $L_{\text{try1}}(q_*) = 0$, we satisfy the Bellman equation (4) on any admissible transition. This constitutes a residual approach [Baird, 1995, Geist et al., 2017]. Alternatively, one could replace the learned term $q(s_{t+1}, a')$ in L_{try1} with a target network $q_{\text{target}}(s_{t+1}, a')$, periodically synced to q , yielding a DQN-style algorithm [Mnih et al., 2015]—namely soft-DQN [Vieillard et al., 2020b] or its direct entropy-to-KL extension for LLMs. However, adding a third network would be memory-inefficient: even “small” LLMs contain billions of parameters, and RL fine-tuning already requires both the learned and reference models. While minimizing $L_{\text{try1}}(q)$ would converge to the Bellman fixed point (and hence the optimal policy), it overlooks several LLM-specific considerations, which we now address.

2.2 Easing sampling

Assume that we optimize the logits of the LLM such that they are a good approximation of the fixed point of Eq. 4. Then, at inference, according to Thm. 1, we would need to sample with

$$\pi(a_t|s_t) \propto \exp \frac{q(s_t, a_t) + \beta \ln \pi_{\text{ref}}(a_t|s_t)}{\beta}.$$

This approach requires loading and querying both the learned and reference models, as well as maintaining the temperature hyperparameter. Furthermore, inference-time decoding methods, *e.g.*, temperature sampling [Ackley et al., 1985] or nucleus sampling [Holtzman et al., 2019], must be adjusted to account for this, which, while conceptually simple, can be inconvenient in practice. Ideally, fine-tuning the LLM’s logits should permit direct softmax sampling without dependence on such artifacts, and the following result shows how this can be achieved. Before stating it, we recall the objects defined in Eq. 11, now expressed in RL terminology. For an arbitrary function $\ell \in \mathbb{R}^{S \times A}$,

$$\pi_\ell(a_t|s_t) = \exp(\ell(s, a) - v_\ell(s)) \text{ with } v_\ell(s) = \ln \sum_{a \in \mathcal{A}} \exp \ell(s, a). \quad (5)$$

Using this, we can state the following simple result with proofs in Appx. C.

Theorem 2. *Let $g \in \mathbb{R}^{S \times A}$ be the unique function satisfying, for any admissible (s_t, a_t, s_{t+1})*

$$\beta g(s_t, a_t) = r(s_t, a_t) + \beta \ln \pi_{\text{ref}}(a_t|s_t) + \gamma(s_t, a_t) \beta v_g(s_{t+1}). \quad (6)$$

Then, the unique optimal policy that maximizes (2) satisfies $\pi_ = \pi_g$.*

From this, we can design the new following loss:

$$L_{\text{try2}}(g) = \mathbb{E}_{x, y \sim \mathcal{D}} \left[\sum_{s_t, a_t \in (x, y)} (r(s_t, a_t) + \beta \ln \pi_{\text{ref}}(a_t|s_t) + \gamma(s_t, a_t) \beta v_g(s_{t+1}) - \beta g(s_t, a_t))^2 \right].$$

A direct corollary of Thm. 2 is that if $\text{supp}(\mathcal{D}) = \text{supp}(\rho \pi_{\text{ref}})$, then the unique optimizer g_* of $L_{\text{try1}}(g)$ satisfies $L_{\text{try2}}(g_*) = 0$ and $\pi_{g_*} = \pi_*$. Learning the logits by minimizing the above loss would allow one to directly sample from them at inference, $\pi(a_t|s_t) = \pi_g(a_t|s_t) \propto \exp g(s_t, a_t)$, which was the desired outcome. However, it still ignores some important peculiarities of LLMs.

2.3 A better initialization

Considering objective in Eqs. (2), we naturally initialize $\pi = \pi_{\text{ref}}$ to minimize the KL term. Indeed, alternative initialization would place the objectives far from their optima complicating learning. To illustrate this, if we set $r = 0$ (hence $R = 0$), then optimizing $J(\pi)$ from $\pi = \pi_{\text{ref}}$ yields no update; π_{ref} is already the global maximizer and the empirical policy gradient vanishes. We would like to get the same behavior when initializing our method with ℓ_{ref} , *i.e.*, there no gradient update when $r = 0$. Unfortunately, we first that it is not the case, motivating for another loss transformation. First, minimizing L_{try2} forces us to initialize the scoring function g with the reference logits $g = \ell_{\text{ref}}$. Using the identity $\ln \pi_g(a_t | s_t) = g(s_t, a_t) - v_g(s_t)$ from Eq. (5) and lead to the following equation.

$$L_{\text{try2}}^{(r=0)}(\ell_{\text{ref}}) = \beta^2 \mathbb{E}_{x, y \sim \mathcal{D}} \left[\sum_{s_t, a_t \in (x, y)} (\gamma(s_t, a_t) v_{\text{ref}}(s_{t+1}) - v_{\text{ref}}(s_t))^2 \right].$$

For $r = 0$, one finds $L_{\text{try2}}^{(r=0)}(\ell_{\text{ref}}) > 0$, inducing an unwanted gradient. This happens despite the fact that the Bellman fixed-point satisfies $L_{\text{try2}}^{(r=0)}(\ell_{\text{ref}}) = 0$ when Eq. (6) holds. Hence ℓ_{ref} is a poor initialization: updates would learn only the missing value component needed to satisfy Bellman, leaving the reference policy (softmax-invariant to state-dependent shifts) unchanged but likely increasing $\text{KL}(\pi_\ell(\cdot|s_t) || \pi_{\text{ref}}(\cdot|s_t))$, which is undesirable. Since no alternative initialization is available without retraining or altering the reference model, we instead modify the Bellman equation so that ℓ_{ref} becomes ideal. To this end, we employ potential-based reward shaping [Ng et al., 1999], which alters rewards without changing the optimal policy and is, in certain settings, equivalent to reinitializing a Q-function method [Wiewiora, 2003]. The next result adapts this technique to LLMs.

Theorem 3. *Let $\ell \in \mathbb{R}^{S \times A}$ be the unique function satisfying, for any admissible (s_t, a_t, s_{t+1}) ,*

$$\beta (\ell(s_t, a_t) - \ell_{\text{ref}}(s_t, a_t)) = r(s_t, a_t) + \gamma(s_t, a_t) \beta (v_\ell(s_{t+1}) - v_{\text{ref}}(s_{t+1})). \quad (7)$$

Then, the unique optimal policy that maximizes (2) satisfies $\pi_ = \pi_\ell$.*

Proof can be found in Appx. C. The Bellman equation (7) connects logits (Q-values) and the log-partition (value) while also considering their differences from the reference model. Intuitively,

this method learn the offset between the reference logits and the actual Q-values. This learned offset enables having no gradient updates, when $r = 0$ and $\pi = \pi_{\text{ref}}$. Finally, we enforce this property by combining Thm. 3 and L_{try2} to obtain the following loss:

$$L_{\text{try3}}(\ell) = \mathbb{E}_{x,y \sim \mathcal{D}} \left[\sum_{s_t, a_t \in (x,y)} (r(s_t, a_t) + \gamma(s_t, a_t) \beta (v_\ell(s_{t+1}) - v_{\text{ref}}(s_{t+1})) - \beta (\ell(s_t, a_t) - \ell_{\text{ref}}(s_t, a_t)))^2 \right] \quad (8)$$

A direct corollary of Thm. 3 is that if $\text{supp}(\mathcal{D}) = \text{supp}(\rho\pi_{\text{ref}})$, then the unique optimizer ℓ_* of $L_{\text{try3}}(\ell)$ satisfies $L_{\text{try3}}(\ell_*) = 0$ and $\pi_{\ell_*} = \pi_*$. In the case $r = 0$ discussed previously, it is easy to verify that $L_{\text{try3}}^{(r=0)}(\ell_{\text{ref}}) = 0$. We posit that this new form of the Bellman equation and the resulting loss are more amenable to the RL fine-tuning of LLMs, as it makes the natural initialization of the logits to ℓ_{ref} a better initialization. However, there is a last specificity of LLMs to address.

2.4 Going multi-step

In an LLM setting, it is common to have sequence-level rewards rather than token/action-level rewards (as commonly used in classic RL problems). However, our current loss L_{try3} is a token-level loss which is not designed to learn from sparse/sequence-only rewards. Intuitively, the rewards at the end of the trajectory will take time during learning to be informative for the entire sequence of tokens. In the following, we describe this issue more rigorously before introducing another loss modification to accelerate the propagation of reward during learning. From RL perspective, a one-step Bellman loss like L_{try3} propagates gradient only one token per update: in a tabular logits example, the first update affects only $\ell(s_{|y|}, a_{|y|})$, the second affects $\ell(s_{|y|-1}, a_{|y|-1})$ and $\ell(s_{|y|}, a_{|y|})$, and so on, requiring $|y|$ steps to reach the first token—neural logits behave similarly. This backward-induction slowdown is typically addressed via n -step returns [Munos et al., 2016, Scherrer et al., 2015, Hessel et al., 2018], but off-policy variants then rely on importance sampling. In the KL-regularized LLM setting, we can derive off-policy multi-step Bellman consistency equations without importance weights. This idea, pioneered in path consistency learning (PCL) [Nachum et al., 2017], is adapted here to our LLM framework thanks to the additional structure induced here by KL regularization.

Theorem 4. *Let $\ell \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ be the unique function satisfying, for any admissible trajectory $(s_k, a_k)_{1 \leq k \leq T}$ (that is, such that $\rho(s_1) > 0$, $\pi_{\text{ref}}(a_{1:T}|s_1) > 0$ and $\gamma(s_T, a_T) = 0$), for any $1 \leq t \leq T$,*

$$\beta (v_\ell(s_t) - v_{\text{ref}}(s_t)) = \sum_{k=t}^T \gamma^{k-t} \left(r(s_t, a_t) - \beta \ln \frac{\pi_\ell(a_t|s_t)}{\pi_{\text{ref}}(a_t|s_t)} \right). \quad (9)$$

Then, the unique optimal policy that maximizes (2) satisfies $\pi_ = \pi_\ell$.*

We can now present the proposed approach, Shifted-Q or ShiQ, that we call this way because both the reparameterization of Thm. 2 and the reward shaping of Eq. 7 amount to shifting the Q-function. Finally the resulting Bellman equation being then turned multi-turn in Thm. 4).

2.5 Shifted-Q

Building upon Eq. (9), we propose the following loss:

$$L_{\text{ShiQ}}(\ell) = \mathbb{E}_{x,y \in \mathcal{D}} \left[\sum_{t=1}^{|y|} \left(\sum_{k=t}^{|y|} \gamma^{k-t} (r(s_k^{xy}, a_k^{xy}) - \beta \ln \frac{\pi_\ell(a_k^{xy} | s_k^{xy})}{\pi_{\text{ref}}(a_k^{xy} | s_k^{xy})}) - \beta (v_\ell(s_t^{xy}) - v_{\text{ref}}(s_t^{xy})) \right)^2 \right] \quad (10)$$

A direct corollary of Thm. 4 is that if $\text{supp}(\mathcal{D}) = \text{supp}(\rho\pi_{\text{ref}})$, then the unique optimizer ℓ_* of $L_{\text{ShiQ}}(\ell)$ satisfies $L_{\text{ShiQ}}(\ell_*) = 0$ and $\pi_{\ell_*} = \pi_*$. With LLM notations, assuming $\gamma = 1$ and a sequence level reward as in Eq. (3), Eq. (10) reduces to Eq. (13). In practice, we optimize this token-level loss by stochastic gradient descent on mini-batches, normalizing by the total number of tokens—analogue to cross-entropy in supervised fine-tuning—while sequence-level objectives may or may not normalize by length [Grinsztajn et al., 2024]. The loss is off-policy, therefore there is no restriction on what prompts and completions the set $\mathcal{D}$ can contain. We assume the set of prompts to be given beforehand, completions can come from a fixed dataset (e.g., a dataset for supervised

fine-tuning or a preference dataset), they can be generated on-policy, or we can use a replay-buffer as classically done in RL [Mnih et al., 2015] to reuse past generations, therefore reducing the sampling cost. This makes our loss more versatile than on-policy policy-gradient methods (which require fresh rollouts) and contrastive approaches (which need paired trajectories), yet it can also leverage such data. The ShiQ algorithm thus emerges from successive, LLM-specific refinements of the regularized Bellman equation. As an ablation, we also consider baselines that ignore one of these three steps listed in 2, to assess their usefulness empirically. For example, we can skip the reward shaping used in Thm. 3, which aims at making the reference logits a better initialization, and do the other steps, resulting loss is $L_{\text{ShiQ}/\text{init}}$. We detail the derivation of the ablations presented in Rk. 1, Appx. C.

3 Empirical results and LLMs notations

In this part, we rephrase our algorithm with LLMs notations, so that the reader not familiar with RL can directly implement the loss. Previously, we write x for a prompt and y for a generation, which is a sequence of tokens from a vocabulary $\mathcal{V}$: $y = (y_1, \dots, y_{|y|})$, with $|y| < T_{\max}$ the length of the sequence. We denote a subsequence $y_{t:t'} = (y_t, y_{t+1}, \dots, y_{t'})$ and use the notations $y_{<t} = y_{1:t}$, $y_{<t} = y_{1:t-1}$ with the convention $y_{<1} = \emptyset$, and $y_{\geq t} = y_{t:|y|}$. We write $\oplus$ for concatenation, for example $x \oplus y_{<t} = (x, y_{<t})$. The policy is an autoregressive LLM, generating a sequence of tokens, $\pi(y|x) = \prod_{t=1}^{|y|} \pi(y_t|x, y_{<t})$. At the token-level, the policy is defined as a softmax over its logits (ℓ) , and we write it π_ℓ to make this dependency explicit:

$$\pi_\ell(y_t|x, y_{<t}) = \exp(\ell(x \oplus y_{<t}, y_t) - v_\ell(x \oplus y_{<t})) \text{ with } v_\ell(x \oplus y_{<t}) = \ln \sum_{w \in \mathcal{V}} \exp \ell(x \oplus y_{<t}, w), \quad (11)$$

with v_ℓ the (tractable) log-partition at the token-level. We will write π_{ref} as a shorthand for $\pi_{\ell_{\text{ref}}}$, with ℓ_{ref} the logits of the reference model, and similarly v_{ref} for $v_{\ell_{\text{ref}}}$. Let $R(x, y)$ be the sequence-level reward (the more general token-level reward will be considered later), ρ be some prompt distribution, β a temperature parameter and π_{ref} the reference model to be fine-tuned. The objective is to maximize

$$J(\pi) = \mathbb{E}_{x \sim \rho} \mathbb{E}_{y \sim \pi(\cdot|x)} [R(x, y) - \beta \text{KL}(\pi(\cdot|x) || \pi_{\text{ref}}(\cdot|x))]. \quad (12)$$

It is well known that the optimal policy satisfies $\pi_*(y|x) \propto \pi_{\text{ref}}(y|x) \exp \frac{R(x, y)}{\beta}$, but the related proportionality partition function is intractable. For this specific case (sequence-level reward), the ShiQ loss that we propose for learning the logits writes, using notations defined in Eq. (11):

$$L_{\text{ShiQ}}(\ell) = \mathbb{E}_{x, y \in \mathcal{D}} \left[\sum_{t=1}^{|y|} \left(R(x, y) - \beta \ln \frac{\pi_\ell(y_{\geq t}|x, y_{<t})}{\pi_{\text{ref}}(y_{\geq t}|x, y_{<t})} - \beta (v_\ell(x \oplus y_{<t}) - v_{\text{ref}}(x \oplus y_{<t})) \right)^2 \right] \quad (13)$$

Recall that under some assumptions in the previous section, Thm. 4 states that this loss admits a unique minimizer ℓ_* satisfying $\pi_{\ell_*} = \pi_*$, that is it provides the logits of the optimal policy maximizing $J(\pi)$. In the next section, we first present two examples on bandits and MDPs to give intuition about the loss and then present results on LLMs experiments in single and multi-turn settings.

3.1 Toy experiment in the offline bandit setting

To empirically evaluate our method, we consider a synthetic 3-armed bandit problem with associated rewards $R = (2.5, 2, 1)$, arms sampled from two distributions: $\mu_1 = (0.1, 0.2, 0.7)$ and $\mu_2 = (0.05, 0.05, 0.9)$. Using these distributions, we construct a dataset comprising 10^4 pairs of rewarded arms. We define the reference policy as uniform: $\pi_{\text{ref}}(y) = \frac{1}{3}$ for all $y \in \{1, 2, 3\}$. The optimal policy for this setting is given analytically by: $\pi_*(y) \propto \exp(R(y)/\beta)$. As comparison in offline setting, we adopt the practical CoPG objective that is design to converge to optimal solution and utilize the gradient expression derived in Flet-Berliac et al. [2024]. Additionally, we include DPO Rafailov et al. [2023] in our evaluation. The performance of each trained policy is assessed using the regret metric: $\text{regret} = J(\pi_*) - J(\hat{\pi})$, where J is defined as the regularised regret $J(\pi) = \mathbb{E}_{y \sim \pi} [R(y)] - \beta \text{KL}(\pi || \pi_{\text{ref}})$. While CoPG rely on access to the reward function R like ShiQ and ShiQ_{/init}, DPO exclusively leverages preference feedback. Note that in the bandit setting ShiQ_{/ms} is equivalent to ShiQ as there is only one turn. To simulate such preferences, we adopt a model defined by: $P(y > y') = 1(R(y) - R(y'))$, where 1 is the indicator function. Experimental

results are presented in Fig 1. As predicted by the theory CoPG, ShiQ and ShiQ_{/init} converge to the correct solution π_* while ShiQ converges slightly faster of other methods. In contrast, DPO converges to a biased solution in offline setting as we do not simulate using Bradley-Terry model.

3.2 Toy MDP with final or fine-grained rewards

The environment is a 5×5 grid-world MDP with four actions {Up, Down, Left, Right}. States are indexed (1, 1) (top-left) to (5, 5) (bottom-right). Two reward configurations are evaluated:

- **Final reward setting:** a single terminal reward $r = 7$ at (5, 5); all other states $r = 0$.
- **Fine-Grained reward setting:** an intermediate reward $r = 4$ at arbitrary state (3, 5) and a terminal reward $r = 3$ at (5, 5); elsewhere $r = 0$.




Figure 1: Offline 3-arms bandit setting

The agent starts at a fixed initial state and seeks to maximize cumulative reward. We first compute the optimal policy π_* via regularized value iteration. For DPO and CoPG, we collect “good” trajectories from π_* and “bad” trajectories from a uniform random policy. Since ShiQ does not require paired trajectories, we simply concatenate both datasets. Hyperparameters are listed in Appendix D.2. In **Final reward setting** (top row of Fig. 2), all methods reliably reach the goal and obtain reward 7, as expected. In **Fine-Grained reward setting** (bottom row of Fig. 2), only ShiQ consistently discovers the intermediate reward at (3, 5) while still reaching (5, 5). DPO and CoPG, which rely exclusively on terminal-reward trajectories, fail to exploit the intermediate signal. Notably, ShiQ first locates the terminal reward (incurring regret 4) and then the intermediate reward, driving regret close to zero.




Figure 2: Comparison of Regret and Pareto front using fine-grained and final rewards.

3.3 LLM experiments on Single-Turn setting

We evaluate on the open-source Anthropic-Harmless and Anthropic-Helpful datasets [Bai et al., 2022] and UltraFeedback [Cui et al., 2023], chosen for their publicly available reward labels. Our policy models are three 7B-parameter LLMs, specifically Cohere R7B [Cohere et al., 2025]. At each evaluation checkpoint, each model generates outputs for a fixed batch of 128 validation prompts; these outputs are scored by a reward model trained exclusively on the training split. The same protocol is applied to DPO, CoPG (using paired completions), and DRO [Richemond et al., 2024] as a single-trajectory baseline. Details of the DRO implementation appear in Appendix B, and an ablation study of ShiQ, $\text{ShiQ}_{\text{init}}$, ShiQ_{ms} , and ShiQ_{tk} is given in Appendix D.3.




Figure 3: Reward optimization and Pareto comparison for HH dataset.

Results show that ShiQ and CoPG achieve comparable performance in maximizing the reward in the offline setting, while DPO and DRO exhibit a more limited reward range. In terms of KL divergence, ShiQ and CoPG yield similar behavior. Note that ShiQ is capable of performing similarly while not having information about which completion is better or not, so leveraging less information. Results for UltraFeedback datasets can be found in Appendix D.4.

3.4 LLM experiments on Multi-Turn setting

We evaluate function-calling capabilities using the BFCL-V3 dataset introduced in the Gorilla framework by Patil et al. [2024]. BFCL-V3 extends prior benchmarks by incorporating multi-turn and multi-step function-calling scenarios, requiring models to maintain dialogue context and autonomously sequence function executions. Evaluation is state-based, measuring the correctness of outcomes rather than just syntax, providing a more robust assessment of tool-use in realistic settings. This setting is particularly relevant for ShiQ as it can include *multi-turn and fine-grained rewards*. Similarly to previous tasks, we start from R7B model Cohere et al. [2025] and plot the verifiable reward from generations using prompt of the validation set composed of 20 percents of BFCL-v3 data. The key findings are summarized in Fig. 4: both the multi-turn DPO variant from Rafailov et al. [2024] and CoPG (Appendix D.5) successfully optimize the cumulative per-turn reward, whereas ShiQ, by leveraging full information about reward positions in the multi-turn setting, outperforms these baselines. Further ablations and experimental details are provided in Appendix D.5.



Figure 4: Reward optimization and Pareto comparison for BFCL-v3 dataset.

4 Conclusion

We propose a novel offline reinforcement-learning algorithm, ShiQ, grounded in the Bellman consistency equation. ShiQ and its variants ShiQ_{init} and ShiQ_{tk} admit theoretical guarantees and demonstrate strong empirical performance, especially in multi-turn scenarios. About limitations and future work: to date, ShiQ has been evaluated on a limited set of large-language-model benchmarks; extending evaluation to additional domains, particularly classical RL tasks and robotics datasets, represents or uses ShiQ for distillation is an interesting future work. Our current experiments rely exclusively on offline data; incorporating fresh model generations or heterogeneous online samples may further improve robustness. Finally, ShiQ assumes access to a reliable reward model, a condition rarely met in practice. While KL regularization helps curb reward hacking, we do not yet include mechanisms to guard against optimizing toward flawed regions of a learned reward function.

References

- David H Ackley, Geoffrey E Hinton, and Terrence J Sejnowski. A learning algorithm for boltzmann machines. *Cognitive science*, 9(1):147–169, 1985.
- Arash Ahmadian, Chris Cremer, Matthias Gallé, Marzieh Fadaee, Julia Kreutzer, Olivier Pietquin, Ahmet Üstün, and Sara Hooker. Back to basics: Revisiting reinforce style optimization for learning from human feedback in llms. *arXiv preprint arXiv:2402.14740*, 2024.
- Mohammad Gheshlaghi Azar, Zhaohan Daniel Guo, Bilal Piot, Remi Munos, Mark Rowland, Michal Valko, and Daniele Calandriello. A general theoretical paradigm to understand learning from human preferences. In *International Conference on Artificial Intelligence and Statistics*, pages 4447–4455. PMLR, 2024.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.
- Leemon Baird. Residual algorithms: Reinforcement learning with function approximation. In *Machine learning proceedings 1995*, pages 30–37. Elsevier, 1995.
- Steven J Bradtke and Andrew G Barto. Linear least-squares algorithms for temporal difference learning. *Machine learning*, 22(1):33–57, 1996.
- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.
- Team Cohere, Arash Ahmadian, Marwan Ahmed, Jay Alammam, Yazeed Alnumay, Sophia Althammer, Arkady Arkhangorodsky, Viraat Aryabumi, Dennis Aumiller, Raphaël Avalos, et al. Command a: An enterprise-ready large language model. *arXiv preprint arXiv:2504.00698*, 2025.

- Ganqu Cui, Lifan Yuan, Ning Ding, Guanming Yao, Wei Zhu, Yuan Ni, Guotong Xie, Zhiyuan Liu, and Maosong Sun. Ultrafeedback: Boosting language models with high-quality feedback. 2023.
- Thomas Degris, Martha White, and Richard S Sutton. Off-policy actor-critic. In *International Conference on Machine Learning*, 2012.
- Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. Kto: Model alignment as prospect theoretic optimization. *arXiv preprint arXiv:2402.01306*, 2024.
- Yannis Flet-Berliac, Nathan Grinsztajn, Florian Strub, Eugene Choi, Bill Wu, Chris Cremer, Arash Ahmadian, Yash Chandak, Mohammad Gheshlaghi Azar, Olivier Pietquin, and Matthieu Geist. Contrastive policy gradient: Aligning LLMs on sequence-level scores in a supervised-friendly fashion. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen, editors, *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 21353–21370, Miami, Florida, USA, November 2024. Association for Computational Linguistics. URL <https://aclanthology.org/2024.emnlp-main.1190>.
- Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *International conference on machine learning*, pages 1587–1596. PMLR, 2018.
- Matthieu Geist, Bilal Piot, and Olivier Pietquin. Is the bellman residual a bad proxy? *Advances in Neural Information Processing Systems*, 30, 2017.
- Matthieu Geist, Bruno Scherrer, and Olivier Pietquin. A theory of regularized markov decision processes. In *International Conference on Machine Learning*, pages 2160–2169. PMLR, 2019.
- Nathan Grinsztajn, Yannis Flet-Berliac, Mohammad Gheshlaghi Azar, Florian Strub, Bill Wu, Eugene Choi, Chris Cremer, Arash Ahmadian, Yash Chandak, Olivier Pietquin, et al. Averaging log-likelihoods in direct alignment. *arXiv preprint arXiv:2406.19188*, 2024.
- Han Guo, Bowen Tan, Zhengzhong Liu, Eric Xing, and Zhiting Hu. Efficient (soft) q-learning for text generation with limited good data. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 6969–6991, 2022.
- Matteo Hessel, Joseph Modayil, Hado Van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Azar, and David Silver. Rainbow: Combining improvements in deep reinforcement learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*, 2019.
- Joey Hong, Anca Dragan, and Sergey Levine. Q-sft: Q-learning for language models via supervised fine-tuning. *arXiv preprint arXiv:2411.05193*, 2024.
- Kaixuan Ji, Guanlin Liu, Ning Dai, Qingping Yang, Renjie Zheng, Zheng Wu, Chen Dun, Quanquan Gu, and Lin Yan. Enhancing multi-step reasoning abilities of language models through direct q-function optimization. *arXiv preprint arXiv:2410.09302*, 2024.
- Wouter Kool, Herke van Hoof, and Max Welling. Buy 4 reinforce samples, get a baseline for free! In *Deep Reinforcement Learning Meets Structured Prediction (ICLR Workshop)*, 2019.
- Ilya Kostrikov, Ashvin Nair, and Sergey Levine. Offline reinforcement learning with implicit q-learning. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=68n2s9ZJWF8>.
- Hung Le, Yue Wang, Akhilesh Deepak Gotmare, Silvio Savarese, and Steven Chu Hong Hoi. Coderl: Mastering code generation through pretrained models and deep reinforcement learning. *Advances in Neural Information Processing Systems*, 35:21314–21328, 2022.
- Yu Meng, Mengzhou Xia, and Danqi Chen. Simpo: Simple preference optimization with a reference-free reward. *arXiv preprint arXiv:2405.14734*, 2024.

- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fiedjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.
- Rémi Munos, Tom Stepleton, Anna Harutyunyan, and Marc Bellemare. Safe and efficient off-policy reinforcement learning. *Advances in neural information processing systems*, 29, 2016.
- Ofir Nachum, Mohammad Norouzi, Kelvin Xu, and Dale Schuurmans. Bridging the gap between value and policy based reinforcement learning. *Advances in neural information processing systems*, 30, 2017.
- Andrew Y Ng, Daishi Harada, and Stuart Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*, volume 99, pages 278–287, 1999.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.
- Shishir G Patil, Tianjun Zhang, Xin Wang, and Joseph E Gonzalez. Gorilla: Large language model connected with massive apis. *Advances in Neural Information Processing Systems*, 37: 126544–126565, 2024.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36, 2023.
- Rafael Rafailov, Joey Hejna, Ryan Park, and Chelsea Finn. From r to q^* : Your language model is secretly a q -function. *arXiv preprint arXiv:2404.12358*, 2024.
- Pierre Harvey Richemond, Yunhao Tang, Daniel Guo, Daniele Calandriello, Mohammad Gheshlaghi Azar, Rafael Rafailov, Bernardo Avila Pires, Eugene Tarassov, Lucas Spangher, Will Ellsworth, et al. Offline regularised reinforcement learning for large language models alignment. *arXiv preprint arXiv:2405.19107*, 2024.
- Paul Roit, Johan Ferret, Lior Shani, Roei Aharoni, Geoffrey Cideron, Robert Dadashi, Matthieu Geist, Sertan Girgin, Leonard Hussenot, Orgad Keller, Nikola Momchev, Sabela Ramos Garea, Piotr Stanczyk, Nino Vieillard, Olivier Bachem, Gal Elidan, Avinatan Hassidim, Olivier Pietquin, and Idan Szpektor. 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, 2023.
- Bruno Scherrer, Mohammad Ghavamzadeh, Victor Gabillon, Boris Lesner, and Matthieu Geist. Approximate modified policy iteration and its application to the game of tetris. *J. Mach. Learn. Res.*, 16(49):1629–1676, 2015.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Lior Shani, Aviv Rosenberg, Asaf Cassel, Oran Lang, Daniele Calandriello, Avital Zipori, Hila Noga, Orgad Keller, Bilal Piot, Idan Szpektor, et al. Multi-turn reinforcement learning from preference human feedback. *arXiv preprint arXiv:2405.14655*, 2024.
- Charlie Victor Snell, Ilya Kostrikov, Yi Su, Sherry Yang, and Sergey Levine. Offline RL for natural language generation with implicit language q learning. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=aBH_DydEvoH.
- Yunhao Tang, Zhaohan Daniel Guo, Zeyu Zheng, Daniele Calandriello, Remi Munos, Mark Rowland, Pierre Harvey Richemond, Michal Valko, Bernardo Avila Pires, and Bilal Piot. Generalized preference optimization: A unified approach to offline alignment. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp, editors, *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 47725–47742. PMLR, 21–27 Jul 2024. URL <https://proceedings.mlr.press/v235/tang24b.html>.

- Yunhao Tang, Taco Cohen, David W Zhang, Michal Valko, and Rémi Munos. RL-finetuning llms from on-and off-policy data with a single algorithm. *arXiv preprint arXiv:2503.19612*, 2025.
- Nino Vieillard, Tadashi Kozuno, Bruno Scherrer, Olivier Pietquin, Rémi Munos, and Matthieu Geist. Leverage the average: an analysis of kl regularization in reinforcement learning. *Advances in Neural Information Processing Systems*, 33:12163–12174, 2020a.
- Nino Vieillard, Olivier Pietquin, and Matthieu Geist. Munchausen reinforcement learning. *Advances in Neural Information Processing Systems*, 33:4235–4246, 2020b.
- Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Hasselt, Marc Lanctot, and Nando Freitas. Dueling network architectures for deep reinforcement learning. In *International conference on machine learning*, pages 1995–2003. PMLR, 2016.
- Eric Wiewiora. Potential-based shaping and q-value initialization are equivalent. *Journal of Artificial Intelligence Research*, 19:205–208, 2003.
- Ronald J Williams and Jing Peng. Function optimization using connectionist reinforcement learning algorithms. *Connection Science*, 3(3):241–268, 1991.
- Wei Xiong, Chengshuai Shi, Jiaming Shen, Aviv Rosenberg, Zhen Qin, Daniele Calandriello, Misha Khalman, Rishabh Joshi, Bilal Piot, Mohammad Saleh, et al. Building math agents with multi-turn iterative preference learning. *arXiv preprint arXiv:2409.02392*, 2024.
- Zishun Yu, Yunzhe Tao, Liyu Chen, Tao Sun, and Hongxia Yang. $\mathcal{B}$ -coder: Value-based deep reinforcement learning for program synthesis. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=fLf589bx1f>.
- Yushun Zhang, Congliang Chen, Ziniu Li, Tian Ding, Chenwei Wu, Yinyu Ye, Zhi-Quan Luo, and Ruoyu Sun. Adam-mini: Use fewer learning rates to gain more. *arXiv preprint arXiv:2406.16793*, 2024.
- Yao Zhao, Rishabh Joshi, Tianqi Liu, Misha Khalman, Mohammad Saleh, and Peter J Liu. Slic-hf: Sequence likelihood calibration with human feedback. *arXiv preprint arXiv:2305.10425*, 2023.
- Yifei Zhou, Andrea Zanette, Jiayi Pan, Sergey Levine, and Aviral Kumar. Archer: Training language model agents via hierarchical multi-turn rl. *arXiv preprint arXiv:2402.19446*, 2024.
- Brian D Ziebart. *Modeling purposeful adaptive behavior with the principle of maximum causal entropy*. PhD thesis, Carnegie Mellon University, 2010.

A Presentation of other variation of ShiQ

ShiQ_{/init}: We can skip the reward shaping used in Thm. 3, that aims at making the reference logits a better initialization, and do the other steps. Therefore, the only difference with the ShiQ loss of Eq. (10) is the term $v_{\text{ref}}(x \oplus y_{<t})$. We detail the derivation in Rk. 1, Appx. C, the resulting loss is

$$L_{\text{ShiQ}/\text{init}}(\ell) = \mathbb{E}_{x,y \in \mathcal{D}} \left[\sum_{t=1}^{|y|} \left(\sum_{k=t}^{|y|} \gamma^{k-t} \left(r(s_k^{xy}, a_k^{xy}) - \beta \ln \frac{\pi_\ell(a_k^{xy} | s_k^{xy})}{\pi_{\text{ref}}(a_k^{xy} | s_k^{xy})} \right) - \beta v_\ell(s_t^{xy}) \right)^2 \right].$$


Written with LLM notations, $\gamma = 1$ and the reward of Eq. (3) is given by:

$$L_{\text{ShiQ}/\text{init}}(\ell) = \mathbb{E}_{x,y \in \mathcal{D}} \left[\sum_{t=1}^{|y|} \left(R(x, y) - \beta \ln \frac{\pi_\ell(y_{\geq t} | x, y_{<t})}{\pi_{\text{ref}}(y_{\geq t} | x, y_{<t})} - \beta v_\ell(x \oplus y_{<t}) \right)^2 \right].$$

ShiQ_{/ms}: We can skip the multi-step extension of Thm. 4, the resulting loss is then simply the loss L_{try3} of Eq. (8). Written with LLM notations, with $\gamma = 1$ and the reward of Eq. (3), also using the notations $\delta v_\ell(s_t) = v_\ell(s_t) - v_{\text{ref}}(s_t)$ and $\delta \ell(s_t, a_t) = \ell(s_t, a_t) - \ell_{\text{ref}}(s_t, a_t)$, it gives

$$L_{\text{ShiQ}/\text{ms}}(\ell) = \mathbb{E}_{x,y \in \mathcal{D}} \left[\sum_{t=1}^{|y|-1} (\beta \delta v_\ell(x \oplus y_{\leq t}) - \beta \delta \ell(x \oplus y_{<t}, y_t))^2 + (R(x, y) - \beta \delta \ell(x \oplus y_{<|y|}, y_{|y|}))^2 \right].$$

In the above expression, we have written explicitly the last term of each sequence to make clear that the reward is zero everywhere except there, and that the value of the next step is part of the square everywhere except there.



ShiQ_{/tk}: The ShiQ loss is a token-level loss, in the sense that it involves a square term for each token of the batch. This contrasts with other RL-finetuning approaches, such as DRO [Richemond et al., 2024] or CoPG [Flet-Berliac et al., 2024], that involve a square term per sequence of the batch. Relatedly, direct alignment methods are also mostly sequence-level losses [Tang et al., 2024]. The underlying reason is that these approaches build upon a bandit viewpoint of LLMs (each possible completion being an arm), while we adopt an MDP viewpoint. We can easily derive a sequence-level loss from our framework. To do so, we can build the loss from the optimality equation (9) of Thm. 4, but considering it only for the initial state, instead of all states of the sequence. Notice that if the optimal logits satisfy this equation, it is not obvious that the solution is unique (conversely to considering all possible states, and not only the initial ones). The resulting loss is

$$L_{\text{ShiQ}/\text{tk}}(\ell) = \mathbb{E}_{x,y \in \mathcal{D}} \left[\left(\sum_{k=1}^{|y|} \gamma^{k-1} \left(r(s_k^{xy}, a_k^{xy}) - \beta \ln \frac{\pi_\ell(a_k^{xy} | s_k^{xy})}{\pi_{\text{ref}}(a_k^{xy} | s_k^{xy})} \right) - \beta (v_\ell(s_1^{xy}) - v_{\text{ref}}(s_1^{xy})) \right)^2 \right].$$

Written with LLM notations, with $\gamma = 1$ and the reward of Eq. (3), it gives:

$$L_{\text{ShiQ}/\text{tk}}(\ell) = \mathbb{E}_{x,y \in \mathcal{D}} \left[\left(R(x, y) - \beta \ln \frac{\pi_\ell(y|x)}{\pi_{\text{ref}}(y|x)} - \beta (v_\ell(x) - v_{\text{ref}}(x)) \right)^2 \right]. \quad (14)$$

Interestingly, when $\gamma = 1$, we can guarantee that any global optimizer of $L_{\text{ShiQ}/\text{tk}}$ (under the same support condition as before) gives logits whose softmax is the optimal policy. (see 5 in Appendix). In the following we will present results for ShiQ while these two ablations are also considered in Appendix.

B Related works

Our contributions build upon a series of previous RL works, as explained during the derivations. Thm. 1 relies on regularized MDPs [Ziebart, 2010, Geist et al., 2019], Thm. 2 uses the idea of Q-function reparameterization Vieillard et al. [2020b], Thm. 3 relies on the idea of reward shaping [Ng et al., 1999] and its relationship to Q-function initialization [Wiewiora, 2003], and Thm. 4 builds upon path consistency learning [Nachum et al., 2017]. Our work is related to RL fine-tuning approaches, such as Reinforce [Williams and Peng, 1991, Roit et al., 2023], leave-one-out Reinforce [Kool et al., 2019, Ahmadian et al., 2024] or PPO [Schulman et al., 2017, Ouyang et al., 2022], that are policy-gradient-based approaches. Our work is even more related to RL fine-tuning approaches allowing to learn in an off-policy manner, or even offline (which prevents using importance sampling, as PPO does for example), especially those relying on Q-functions and Bellman-like equations.

Flet-Berliac et al. [2024] model the LLM as a bandit and propose contrastive policy-gradient (CoPG), a method generalizing policy-gradient to off-policy learning without importance sampling. In its simplest form, the related loss can be written as

$$L_{\text{CoPG}}(\pi) = \mathbb{E}_{x,y,y' \in \mathcal{D}} \left[\left(R(x,y) - \beta \ln \frac{\pi(y|x)}{\pi_{\text{ref}}(y|x)} - \left(R(x,y') - \beta \ln \frac{\pi(y'|x)}{\pi_{\text{ref}}(y'|x)} \right) \right)^2 \right].$$

It bears structural similarities with our ablation ShiQ_{tk} in Eq. (14), the value difference $v_\ell(x) - v_{\text{ref}}(x)$ being replaced by the regularized reward $R(x,y') - \beta \ln \frac{\pi(y'|x)}{\pi_{\text{ref}}(y'|x)}$ on an independent completion for the same prompt. Compared to CoPG, ShiQ (Eq. (13)) does not require a pair of completions for each prompt, it is a token-level loss taking advantage of each subsequence reaching the end for each completion (for CoPG to do so, pairs of partial completions would be needed for each common prefix), and it can take advantage of a token-level reward, while CoPG only see the sequence level quantity.

Richemond et al. [2024] also model the LLM as a bandit, and propose direct reward optimization (DRO), that builds upon the known but untractable solution to problem (12). DRO is an actor-critic method, that requires learning both a policy and a value network. The corresponding loss is

$$L_{\text{DRO}}(\pi, V) = \mathbb{E}_{x,y \in \mathcal{D}} \left[\left(R(x,y) - \beta \ln \frac{\pi(y|x)}{\pi_{\text{ref}}(y|x)} - V(x) \right)^2 \right].$$

It also bears structural similarities with our ablations ShiQ_{tk} in Eq. (14), with respectively the value difference $\beta(v_\ell(x) - v_{\text{ref}}(x))$ being replaced by the value network $V(x)$. Therefore, our ablation can be seen as cheap but theoretically founded alternative to DRO. Indeed, Richemond et al. [2024] noted that parameter sharing was harmful to good empirical results, thus requiring a separate network (and the corresponding optimizer state, which takes by default twice the network memory, if a more involved optimizer is not used [Zhang et al., 2024]). This is very costly. It is also more complicated, in the sense that the policy and value gradients need to be scaled differently for achieving good performance. Moreover, DRO cannot take advantage of the subsequence information, notably the possible token-level reward, contrary to ShiQ (Eq. (13)). We provide a more technical discussion of the relationship between ShiQ and DRO in Rk. 2, Appx. C. Notably, we explain and discuss an apparent inconsistency between our Thm. 5 and [Richemond et al., 2024, Thm. 1], while both results are indeed correct (but rely on different representations of policies). Additionally, Tang et al. [2025] recently generalized DRO and CoPG in a single algorithm.

Guo et al. [2022] model the logits of the LLM as Q-values and learn them using RL. Indeed, their approach is exactly path consistency learning (PCL) [Nachum et al., 2017] applied to the logits, up to the fact that they introduce a target network for the value component. They do not justify this choice, and it requires loading an additional network in memory, which is not desirable as explained before. In fact, our ablation $\text{ShiQ}_{\text{init}}$ can be seen as a generalization of their approach (from entropy regularization to KL regularization, they do not regularize towards a reference model, and also more carefully taking care of the temperature), without the unnecessary introduction of a target network. Compared to this, ShiQ is designed so that the reference logits are a good initialization.

Yu et al. [2024] also interpret the logits of the LLM as Q-values. As us, they observe that the logits of the reference policy might not be a good initialization. However, their proposed approach,

Bellman-coder (B-coder), addresses the issue in a very different (and more complicated, more costly and less theoretically founded) manner. They adopt a dueling architecture [Wang et al., 2016] (logits model the advantage as $\ell(s_t, a_t) - \max_a \ell(s_t, a)$ and there is an additional value head), coupled with an additional value network. This additional network is pretrained to fit the Bellman equation (for a better initialization). Then, instead of considering a regularized MDP setting, they do a single policy improvement step, by performing policy evaluation on the policy being greedy with respect to the reference logits, this with a simple adaptation of DQN. At inference, they play (heuristically) the softmax over the learnt logits. Our baseline ShiQ_{ms} is representative of this, in the sense that it builds upon a one-step Bellman equation. However, it relies on the proper regularized MDP framework [Geist et al., 2017] instead of building upon heuristics, and it modifies the Bellman equation for making the reference logits a good initialization instead of modifying the architecture, introducing an additional network, and adding a pretraining phase.

There are a few other works adopting a Q-function viewpoint for training LLMs, but that we do not think well suited for fine-tuning LLMs at scale. Snell et al. [2023] propose a direct application of inverse Q-learning [Kostrikov et al., 2022] to language modeling. This requires modifying the architecture (it’s an actor-critic approach, with shared parameters between the Q-value and the value), they do not specifically take care about the initialization (Q-networks are randomly initialized in some of they experiment, they do not explicitly tackle the fine-tuning problem), they rely on a one-step Bellman like approach (as our ablation ShiQ_{ms}), they require additional target networks, and at inference they need to load the reference policy, a problem we discussed and alleviated in Sec. 2.2. [Hong et al., 2024] have a similar motivation as us, leveraging the available logits without introducing additional network or value head, but they address it in a very different manner. They do not interpret the logits as Q-values, but the softmax over logits as Q-values. They introduce a Bellman-like equation for learning this probabilities, not properly taking into account the regularization towards the reference model, and use it to propose a Bellman-inspired cross-entropy-like loss function. They introduce KL-regularization heuristically *post hoc*, by sampling $\propto \pi_{\text{ref}}(a_t|s_t) \exp \frac{\pi(a_t|s_t)}{\beta}$, with π being learnt with the proposed loss. Compared to ShiQ , their theoretical result doesn’t guarantee getting the optimal policy even in the ideal case, they require an additional target network, they rely on a one-step Bellman like approach (as our ablation ShiQ_{ms}), and at inference they have the problem we discussed and alleviated in Sec. 2.2.

Rafailov et al. [2024] extend Direct Preference Optimization (DPO) Rafailov et al. [2023] to the multi-turn setting; however, their method depends on paired trajectories, whereas our algorithm only requires unranked trajectories. Xiong et al. [2024] derive an analogous loss and incorporate ideas from KTO Ethayarajh et al. [2024] for multi-turn interactions. Moreover, Shani et al. [2024] propose a self-play-based multi-turn algorithm that seeks a Nash equilibrium: its objective diverges from classical RL formulations and is well-suited to cyclic preferences, yet it too mandates preference feedback between full conversation pairs and includes a learned critic within its deep-RL implementation. Then, Ji et al. [2024] introduce an offline RL approach that directly optimizes a Q-function via the Soft Actor-Critic framework; this method, however, relies on importance-weighted updates—prone to high variance—and requires training both a value network and a policy network, whereas our method optimizes only the policy. Finally, Zhou et al. [2024] present an offline actor-critic framework with three networks (value, Q-function, and policy) and employ expectile regression over actions in the Q learning loss rather than importance weighting to address the offline setting.

C Proofs

To do so, we introduce a state-action dependent discount factor to account for the fact that we work in a finite-horizon MDP: $\gamma(s_t, a_t) = 0$ if $a_t = \text{eos}$ or $t = T_{\text{max}}$, otherwise $\gamma(s_t, a_t) = \gamma$. Notice that this is introduced for accounting for the variable finite horizon setting, and it is different from adding an absorbing state in a discounted infinite horizon setting.

This appendix provides the proofs for the results stated in the main text. We recall that we say a transition (s_t, a_t, s_{t+1}) to be admissible if it can occurs by sampling $x \sim \rho$ and $y \sim \pi_{\text{ref}}(\cdot|x)$, that is, with $s_t = (s_1, a_1, a_2, \dots, a_{t-1})$ (by definition), $\rho(s_1) > 0$ and $\pi_{\text{ref}}(a_{1:t}|s_1) > 0$. When $\gamma(s_t, a_t) = 0$, s_{t+1} is a dummy state (but its value will never be evaluated). The considered setting is that of MDPs with variable but bounded horizon. The corresponding state-space is finite (even if huge). In the proofs, we will write Δ_X for the set of probability distributions over a finite set X . We start by recalling Thm. 1 before proving it.

Theorem 1. Let $q \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ be the unique function satisfying, for any admissible (s_t, a_t, s_{t+1}) ,

$$q(s_t, a_t) = r(s_t, a_t) + \gamma(s_t, a_t) \beta \ln \sum_{a' \in \mathcal{A}} \pi_{\text{ref}}(a' | s_{t+1}) \exp \frac{q(s_{t+1}, a')}{\beta}.$$

Then, the unique optimal policy maximizing (2) satisfies

$$\pi_*(a_t | s_t) = \frac{\pi_{\text{ref}}(a_t | s_t) \exp \frac{q(s_t, a_t)}{\beta}}{\sum_{a \in \mathcal{A}} \pi_{\text{ref}}(a | s_t) \exp \frac{q(s_t, a)}{\beta}}.$$

Proof. Recall the objective function (2) to be maximized:

$$J_{\text{rl}}(\pi) = \mathbb{E}_{x \sim \rho} \mathbb{E}_{y \sim \pi(\cdot | x)} \left[\sum_{t=1}^{|y|} \gamma^{t-1} \left(r(s_t^{xy}, a_t^{xy}) - \beta \frac{\ln \pi(a_t^{xy} | s_t^{xy})}{\ln \pi_{\text{ref}}(a_t^{xy} | s_t^{xy})} \right) \right].$$

First, notice that a policy maximizing J_{rl} cannot sample something else than transitions that we call admissible, otherwise that would make the KL term infinite. In practice, π_{ref} is a softmax over some learnt logits, so it associates a strictly positive probability to any action, and we'll assume π_{ref} to have full support for simplifying the notations. The state space is thus the set of all trajectories of length up to T_{max} . Second, a policy being optimal for any admissible state will also maximize J_{rl} . Therefore, we adopt a dynamic programming viewpoint, and solve the problem using backward induction.

To this end, let's introduce the value function, for any $t \leq T_{\text{max}}$,

$$V_{\pi}(s_t) = \mathbb{E}_{a_t \dots a_T \sim \pi(\cdot | s_t)} \left[\sum_{k=t}^T \gamma^{k-t} \left(r(s_k, a_k) - \beta \ln \frac{\pi(a_k | s_k)}{\pi_{\text{ref}}(a_k | s_k)} \right) \right].$$

Notice that in the above definition, T itself is a random variable bounded by T_{max} , not a fixed quantity. From this definition, we directly have that $J_{\text{rl}}(\pi) = \mathbb{E}_{s_1 \sim \rho} [V_{\pi}(s_1)]$. We also have the following simple result, for any $t < T_{\text{max}}$:

$$\begin{aligned} & V_{\pi}(s_t) \\ &= \mathbb{E}_{a_t \dots a_T \sim \pi(\cdot | s_t)} \left[\sum_{k=t}^T \gamma^{k-t} \left(r(s_k, a_k) - \beta \ln \frac{\pi(a_k | s_k)}{\pi_{\text{ref}}(a_k | s_k)} \right) \right] \\ &= \mathbb{E}_{a_t \sim \pi(\cdot | s_t)} \left[r(s_t, a_t) - \beta \ln \frac{\pi(a_t | s_t)}{\pi_{\text{ref}}(a_t | s_t)} \right. \\ &\quad \left. + \gamma(s_t, a_t) \mathbb{E}_{a_{t+1} \dots a_T \sim \pi(\cdot | s_t \oplus a_t)} \left[\sum_{k=t+1}^T \gamma^{k-t-1} \left(r(s_k, a_k) - \beta \ln \frac{\pi(a_k | s_k)}{\pi_{\text{ref}}(a_k | s_k)} \right) \right] \right] \\ &= \mathbb{E}_{a_t \sim \pi(\cdot | s_t)} \left[r(s_t, a_t) - \beta \ln \frac{\pi(a_t | s_t)}{\pi_{\text{ref}}(a_t | s_t)} + \gamma(s_t, a_t) V_{\pi}(s_t \oplus a_t) \right]. \end{aligned}$$

From this, we can see that if we can find the optimal policy for states $s_t \oplus a_t$, then we can easily get that at state s_t , this is the principle of backward induction (solve a sequence of simpler problems, starting from the end).

Let's consider the case $t = T_{\text{max}}$ first. We have that

$$V_{\pi}(s_t) = \mathbb{E}_{a_t \sim \pi(\cdot | s_t)} \left[r(s_t, a_t) - \beta \ln \frac{\pi(a_t | s_t)}{\pi_{\text{ref}}(a_t | s_t)} \right],$$

as we necessarily have that $\gamma(s_t, a_t) = 0$. Maximizing V_{π} in this case is a classic Legendre-Fenchel transform, its unique solution is given by (e.g., [Vieillard et al., 2020a, Appx. A])

$$\begin{aligned} V_*(s_t) &= \max_{\pi(\cdot | s_t) \in \Delta_{\mathcal{A}}} V_{\pi}(s_t) \\ &= \mathbb{E}_{y \sim \pi_*(\cdot | s_t)} \left[r(s_t, a_t) - \beta \ln \frac{\pi_*(a_t | s_t)}{\pi_{\text{ref}}(a_t | s_t)} \right] \text{ with } \pi_*(a_t | s_t) = \frac{\pi_{\text{ref}}(a_t | s_t) \exp \frac{r(s_t, a_t)}{\beta}}{\sum_{a \in \mathcal{A}} \pi_{\text{ref}}(a | s_t) \exp \frac{r(s_t, a)}{\beta}} \\ &= \beta \ln \sum_{a \in \mathcal{A}} \pi_{\text{ref}}(a | s_t) \exp \frac{r(s_t, a)}{\beta}. \end{aligned}$$

Let also define $q(s_t, a_t) = r(s_t, a_t)$, we have just shown that $\pi_*(a_t|s_t) \propto \pi_{\text{ref}}(a_t|s_t) \exp \frac{q(s_t, a_t)}{\beta}$ and that $V_*(s_t) = \beta \ln \sum_{a \in \mathcal{A}} \pi_{\text{ref}}(a|s_t) \exp \frac{q(s_t, a)}{\beta}$ for $t = T_{\max}$.

Now, let choose $t < T_{\max}$. We assume that the following is true at step $t + 1$ (we have just shown it at step $t + 1 = T_{\max}$) and will show that it is true at step t , which will prove the result by (backward) induction:

$$\begin{cases} q(s_{t+1}, a_{t+1}) = r(s_{t+1}, a_{t+1}) + \gamma(s_{t+1}, a_{t+1})\beta \ln \sum_{a \in \mathcal{A}} \exp \frac{q(s_{t+2}, a)}{\beta} \\ \max_{\pi(\cdot|s_{t+1})} V_{\pi}(s_{t+1}) = V_*(s_{t+1}) = \beta \ln \sum_{a \in \mathcal{A}} \pi_{\text{ref}}(a|s_{t+1}) \exp \frac{q(s_{t+1}, a)}{\beta} \\ \operatorname{argmax}_{\pi(\cdot|s_{t+1})} V_{\pi}(s_{t+1}) = \pi_*(\cdot|s_{t+1}) \text{ with } \pi_*(a_{t+1}|s_{t+1}) \propto \pi_{\text{ref}}(a_{t+1}|s_{t+1}) \exp \frac{q(s_{t+1}, a_{t+1})}{\beta}. \end{cases} \quad (15)$$

Let show that this also hold at step t . In the following equations, we write $\max_{\pi(\cdot|s_t)}$ when considering policies completing sequences to the end, while we write $\max_{\pi(\cdot|s_t) \in \Delta_{\mathcal{A}}}$ when considering policies at the action (or token) level. We have that

$$\begin{aligned} V_*(s_t) &= \max_{\pi(\cdot|s_t)} V_{\pi}(s_t) \\ &= \max_{\pi(\cdot|s_t)} \mathbb{E}_{a_t \sim \pi(\cdot|s_t)} \left[r(s_t, a_t) - \beta \ln \frac{\pi(a_t|s_t)}{\pi_{\text{ref}}(a_t|s_t)} + \gamma(s_t, a_t)V_{\pi}(s_t \oplus a_t) \right] \\ &= \max_{\pi(\cdot|s_t) \in \Delta_{\mathcal{A}}} \mathbb{E}_{a_t \sim \pi(\cdot|s_t)} \left[r(s_t, a_t) - \beta \ln \frac{\pi(a_t|s_t)}{\pi_{\text{ref}}(a_t|s_t)} + \gamma(s_t, a_t) \max_{\pi(\cdot|s_t \oplus a_t)} V_{\pi}(s_t \oplus a_t) \right] \\ &= \max_{\pi(\cdot|s_t) \in \Delta_{\mathcal{A}}} \mathbb{E}_{a_t \sim \pi(\cdot|s_t)} \left[r(s_t, a_t) - \beta \ln \frac{\pi(a_t|s_t)}{\pi_{\text{ref}}(a_t|s_t)} + \gamma(s_t, a_t)V_*(s_t \oplus a_t) \right]. \end{aligned}$$

The term $V_*(s_t \oplus a_t)$ is known for any admissible a_t , thanks to the induction assumption (15), and this optimization problem is again a Legendre-Fenchel transform. Let define

$$\begin{aligned} q(s_t, a_t) &= r(s_t, a_t) + \gamma(s_t, a_t)V_*(s_t \oplus a_t) \\ &= r(s_t, a_t) + \gamma(s_t, a_t)\beta \ln \sum_{a \in \mathcal{A}} \pi_{\text{ref}}(a|s_{t+1}) \exp \frac{q(s_{t+1}, a)}{\beta}. \end{aligned} \quad (16)$$

The second equality is true by writing $s_{t+1} = s_t \oplus a_t$ and by the induction assumption (15). The optimization problem can be written and solved as follows:

$$\begin{aligned} V_*(s_t) &= \max_{\pi(\cdot|s_t) \in \Delta_{\mathcal{A}}} \mathbb{E}_{a_t \sim \pi(\cdot|s_t)} \left[q(s_t, a_t) - \beta \ln \frac{\pi(a_t|s_t)}{\pi_{\text{ref}}(a_t|s_t)} \right] \\ &= \mathbb{E}_{a_t \sim \pi_*(\cdot|s_t)} \left[q(s_t, a_t) - \beta \ln \frac{\pi_*(a_t|s_t)}{\pi_{\text{ref}}(a_t|s_t)} \right] \\ &\text{with } \pi_*(a_t|s_t) \propto \pi_{\text{ref}}(a_t|s_t) \exp \frac{q(s_t, a_t)}{\beta} \end{aligned} \quad (17)$$

$$= \beta \ln \sum_{a \in \mathcal{A}} \pi_{\text{ref}}(a|s_t) \exp \frac{q(s_t, a)}{\beta}. \quad (18)$$

Taken together, Eqs. (16), (17) and (18) show that the induction assumption (15) at step $t + 1$ implies that it is true also at step t . Overall, this proves the stated result, π_* is the unique optimal policy (uniqueness of the policy following from uniqueness of the solution of each of the involved Legendre-Fenchel transforms). $\square$

It is important to note that the proof does not rely on a contraction argument. As we work in a finite (even if variable) horizon setting, we can use backward induction. An important consequence of this is that we can safely consider $\gamma = 1$, conversely to infinite-horizon discounted MDPs. Moreover, we have proved this result for deterministic dynamics, as it is the case of interest for fine-tuning LLMs, but it extends easily to stochastic dynamics. Next, we recall Thm. 2 and prove it.

Theorem 2. Let $g \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ be the unique function satisfying, for any admissible (s_t, a_t, s_{t+1})

$$\beta g(s_t, a_t) = r(s_t, a_t) + \beta \ln \pi_{\text{ref}}(a_t|s_t) + \gamma(s_t, a_t)\beta v_g(s_{t+1}).$$

Then, the unique optimal policy that maximizes (2) satisfies $\pi_* = \pi_g$.

Proof. This is a simple change of variable. Recall from Thm. 1 that the optimal policy satisfies $\pi_*(a_t|s_t) \propto \pi_{\text{ref}}(a_t|s_t) \exp \frac{q(s_t, a_t)}{\beta}$, with q satisfying the Bellman equation

$$q(s_t, a_t) = r(s_t, a_t) + \gamma(s_t, a_t)\beta \ln \sum_{a \in \mathcal{A}} \pi_{\text{ref}}(a|s_{t+1}) \exp \frac{q(s_{t+1}, a)}{\beta}. \quad (19)$$

Let define $g \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ as

$$g(s_t, a_t) = \frac{q(s_t, a_t) + \beta \ln \pi_{\text{ref}}(a_t|s_t)}{\beta}. \quad (20)$$

We immediately have that

$$\pi_*(a_t|s_t) = \frac{\pi_{\text{ref}}(a_t|s_t) \exp \frac{q(s_t, a_t)}{\beta}}{\sum_{a \in \mathcal{A}} \pi_{\text{ref}}(a|s_t) \exp \frac{q(s_t, a)}{\beta}} = \frac{\exp g(s_t, a_t)}{\sum_{a \in \mathcal{A}} \exp g(s_t, a)} = \pi_g(a_t|s_t).$$

Using Eqs. (19) and (20), we have that

$$\begin{aligned} q(s_t, a_t) &= r(s_t, a_t) + \gamma(s_t, a_t)\beta \ln \sum_{a \in \mathcal{A}} \pi_{\text{ref}}(a|s_{t+1}) \exp \frac{q(s_{t+1}, a)}{\beta} \\ \Leftrightarrow \underbrace{q(s_t, a_t)}_{=\beta(g(s_t, a_t) - \ln \pi_{\text{ref}}(a_t|s_t))} &= r(s_t, a_t) + \gamma(s_t, a_t)\beta \ln \sum_{a \in \mathcal{A}} \exp \underbrace{\frac{q(s_{t+1}, a) + \beta \ln \pi_{\text{ref}}(a|s_{t+1})}{\beta}}_{=g(s_{t+1}, a)} \\ \Leftrightarrow \beta g(s_t, a_t) &= r(s_t, a_t) + \beta \ln \pi_{\text{ref}}(a_t|s_t) + \gamma(s_t, a_t)\beta \ln \sum_{a \in \mathcal{A}} \exp g(s_{t+1}, a) \\ &= r(s_t, a_t) + \beta \ln \pi_{\text{ref}}(a_t|s_t) + \gamma(s_t, a_t)\beta v_g(s_{t+1}). \end{aligned}$$

This proves the stated result. $\square$

This kind of change of variable is very simple, it was done before in the literature in similar settings (e.g., in a value-iteration-like scheme with regularization towards the previous policy [Vieillard et al., 2020b]). However, we think it to be very important for LLMs, as it allows sampling directly from the logits, without loading an additional network on learning parameters at inference. We also notice that, as before, this result is not restricted to deterministic kernels and can easily be extended to stochastic transitions. Now, we prove Thm. 3 after having recalled it.

Theorem 3. Let $\ell \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ be the unique function satisfying, for any admissible (s_t, a_t, s_{t+1})

$$\beta(\ell(s_t, a_t) - \ell_{\text{ref}}(s_t, a_t)) = r(s_t, a_t) + \gamma(s_t, a_t)\beta(v_\ell(s_{t+1}) - v_{\text{ref}}(s_{t+1})).$$

Then, the unique optimal policy that maximizes (2) satisfies $\pi_* = \pi_\ell$.

Proof. This is a direct corollary of a more general result that we prove first. This more general result is a simple adaptation of reward shaping [Ng et al., 1999] to our KL-regularized variable finite horizon setting, applied to the Bellman equation of Thm. 2, that we recall here:

$$\beta g(s_t, a_t) = r(s_t, a_t) + \beta \ln \pi_{\text{ref}}(a_t|s_t) + \gamma(s_t, a_t)\beta v_g(s_{t+1}). \quad (21)$$

Let $\phi \in \mathbb{R}^{\mathcal{S}}$ be an arbitrary state-dependent function, we define the shaped reward r_ϕ as

$$r_\phi(s_t, a_t, s_{t+1}) = r(s_t, a_t) + \gamma(s_t, a_t)\beta\phi(s_{t+1}) - \beta\phi(s_t).$$

We replace r by r_ϕ in Eq. (21), and call g_ϕ the associated fixed-point of the Bellman equation:

$$\begin{aligned} \beta g_\phi(s_t, a_t) &= r_\phi(s_t, a_t) + \beta \ln \pi_{\text{ref}}(a_t|s_t) + \gamma(s_t, a_t)\beta v_{g_\phi}(s_{t+1}) \\ &= r(s_t, a_t) + \gamma(s_t, a_t)\beta\phi(s_{t+1}) - \beta\phi(s_t) + \beta \ln \pi_{\text{ref}}(a_t|s_t) \\ &\quad + \gamma(s_t, a_t)\beta \ln \sum_{a \in \mathcal{A}} \exp g_\phi(s_{t+1}, a). \end{aligned} \quad (22)$$

Rearranging terms, this is equivalent to:

$$\beta(g_\phi(s_t, a_t) + \phi(s_t)) = r(s_t, a_t) + \beta \ln \pi_{\text{ref}}(a_t|s_t) + \gamma(s_t, a_t)\beta \ln \sum_{a \in \mathcal{A}} \exp(g_\phi(s_{t+1}, a) + \phi(s_{t+1})).$$

Therefore, we have that $g_\phi(s_t, a_t) + \phi(s_t)$ satisfies the Bellman equation (21), given that its fixed point is unique, we necessarily have that

$$g_\phi(s_t, a_t) + \phi(s_t) = g(s_t). \quad (23)$$

The softmax being invariant to a shift by a state-dependent function, both g_ϕ and g induce the same optimal policy:

$$\pi_{g_\phi}(a_t|s_t) = \frac{\exp g_\phi(s_t, a_t)}{\sum_{a \in \mathcal{A}} \exp g_\phi(s_t, a)} = \frac{\exp(g(s_t, a_t) - \phi(s_t))}{\sum_{a \in \mathcal{A}} \exp(g(s_t, a) - \phi(s_t))} = \pi_g(a_t|s_t) = \pi_*(a_t|s_t). \quad (24)$$

Therefore, shaping the reward as depicted above lets the optimal policy remain invariant.

The stated result is obtained by choosing specifically $\phi(s_t) = -v_{\text{ref}}(s_t)$. Writing $\ell(s_t, a_t) = g_{-v_{\text{ref}}}(s_t, a_t)$, Eq. (22) becomes

$$\beta \ell(s_t, a_t) = r(s_t, a_t) - \gamma(s_t, a_t) \beta v_{\text{ref}}(s_{t+1}) + \underbrace{\beta v_{\text{ref}}(s_t) + \beta \ln \pi_{\text{ref}}(a_t|s_t)}_{=\beta \ell_{\text{ref}}(s_t, a_t)} + \gamma(s_t, a_t) \beta v_\ell(s_{t+1})$$

$$\Leftrightarrow \beta(\ell(s_t, a_t) - \ell_{\text{ref}}(s_t, a_t)) = r(s_t, a_t) + \beta \gamma(s_t, a_t)(v_\ell(s_{t+1}) - v_{\text{ref}}(s_{t+1})).$$

This is the stated Bellman equation, and as we have already shown that $\pi_\ell = \pi_*$, as a special case of Eq. (24), this proves the stated result. $\square$

Thanks to a simple reward shaping, we obtain a Bellman equation that does not involve logits and related value (that is log-partition), but their respective differences to that of the reference model. We posit this provide a better initialization, as this leads to learn how to modify the reference logits we start from, instead of some function less related to the initialization. Again, this result can easily be extended to stochastic dynamics. The next result to prove is Thm. 4, which we recall first.

Theorem 4. *Let $\ell \in \mathbb{R}^{S \times \mathcal{A}}$ be the unique function satisfying, for any admissible trajectory $(s_k, a_k)_{1 \leq k \leq T}$ (that is, such that $\rho(s_1) > 0$, $\pi_{\text{ref}}(a_{1:T}|s_1) > 0$ and $\gamma(s_T, a_T) = 0$), for any $1 \leq t \leq T$,*

$$\beta(v_\ell(s_t) - v_{\text{ref}}(s_t)) = \sum_{k=t}^T \gamma^{k-t} \left(r(s_t, a_t) - \beta \ln \frac{\pi_\ell(a_t|s_t)}{\pi_{\text{ref}}(a_t|s_t)} \right). \quad (25)$$

Then, the unique optimal policy that maximizes (2) satisfies $\pi_* = \pi_\ell$.

Proof. Using the general identity $\ln \pi_\ell(a_t|s_t) = \ell(s_t, a_t) - v_\ell(s_t)$, we start by rewriting the Bellman equation from Thm. 3:

$$\beta \left(\underbrace{\ell(s_t, a_t)}_{=\ln \pi_\ell(a_t|s_t) + v_\ell(s_t)} - \underbrace{\ell_{\text{ref}}(s_t, a_t)}_{=\ln \pi_{\text{ref}}(a_t|s_t) + v_{\text{ref}}(s_t)} \right) = r(s_t, a_t) + \gamma(s_t, a_t) \beta (v_\ell(s_{t+1}) - v_{\text{ref}}(s_{t+1})) \quad (26)$$

$$\Leftrightarrow \beta(v_\ell(s_t) - \gamma(s_t, a_t)v_\ell(s_{t+1})) = r(s_t, a_t) - \beta \ln \frac{\pi_\ell(a_t|s_t)}{\pi_{\text{ref}}(a_t|s_t)} + \beta(v_{\text{ref}}(s_t) - \gamma(s_t, a_t)v_{\text{ref}}(s_{t+1})). \quad (27)$$

We observe a telescopic structure appearing.

Let $(s_k, a_k)_{1 \leq k \leq T}$ be an arbitrary admissible trajectory ($\rho(s_1) > 0$, $\pi_{\text{ref}}(a_{1:T}|s_1) > 0$ and $\gamma(s_T, a_T) = 0$). Importantly, it doesn't need to be sampled according to π_* , which would not be reasonable in general. Let $1 \leq t \leq T$. Eq. (27) being true for any admissible transition, it implies that

$$\begin{aligned} \sum_{k=t}^T \gamma^{t-k} \beta (v_\ell(s_k) - \gamma(s_k, a_k)v_\ell(s_{k+1})) &= \sum_{k=t}^T \gamma^{t-k} \left(r(s_k, a_k) - \beta \ln \frac{\pi_\ell(a_k|s_k)}{\pi_{\text{ref}}(a_k|s_k)} \right. \\ &\quad \left. + \beta(v_{\text{ref}}(s_k) - \gamma(s_k, a_k)v_{\text{ref}}(s_{k+1})) \right) \\ &\Leftrightarrow \beta v_\ell(s_t) = \sum_{k=t}^T \gamma^{t-k} \left(r(s_k, a_k) - \beta \ln \frac{\pi_\ell(a_k|s_k)}{\pi_{\text{ref}}(a_k|s_k)} \right) + \beta v_{\text{ref}}(s_t) \\ &\Leftrightarrow \beta(v_\ell(s_t) - v_{\text{ref}}(s_t)) = \sum_{k=t}^T \gamma^{t-k} \left(r(s_k, a_k) - \beta \ln \frac{\pi_\ell(a_k|s_k)}{\pi_{\text{ref}}(a_k|s_k)} \right). \end{aligned}$$

The second equality is true because all terms $v_\ell(s_k)$ and $v_{\text{ref}}(s_k)$ for $k > t$ cancel in the telescopic sum (also using the fact that $\gamma(s_T, a_T) = 0$). We have just shown that the function ℓ satisfying the Bellman equation (26) (for any admissible transition, that is the fixed point of Thm. 3) also satisfies Eq. (25). We still need to show its uniqueness, that is, if ℓ satisfies Eq. (25) for any admissible sub-trajectory, we indeed have that $\pi_\ell = \pi_*$. This is of foremost importance for guaranteeing that we compute the right object.

Let $f \in \mathbb{R}^{S \times A}$ satisfying, for any admissible trajectory, Eq. (25):

$$\beta(v_f(s_t) - v_{\text{ref}}(s_t)) = \sum_{k=t}^T \gamma^{k-t} \left(r(s_k, a_k) - \beta \ln \frac{\pi_f(a_k|s_k)}{\pi_{\text{ref}}(a_k|s_k)} \right).$$

For any admissible state-action pair (s_t, a_t) such that $\gamma(s_t, a_t) = 0$, this gives

$$\begin{aligned} \beta(v_f(s_t) - v_{\text{ref}}(s_t)) &= r(s_t, a_t) - \beta \ln \frac{\pi_f(a_t|s_t)}{\pi_{\text{ref}}(a_t|s_t)} \\ \Leftrightarrow \beta(f(s_t, a_t) - \ell_{\text{ref}}(s_t, a_t)) &= r(s_t, a_t), \end{aligned} \quad (28)$$

where we used again the general identity $\ln \pi_f(a_t|s_t) = f(s_t, a_t) - v_f(s_t)$. Now, for any admissible state-action pair (s_t, a_t) such that $\gamma(s_t, a_t) \neq 0$, completed by any admissible sub-trajectory $a_{t+1:T}$ (satisfying $\pi_{\text{ref}}(a_{t+1:T}|s_{t+1}) > 0$ and $\gamma(s_T, a_T) = 0$, implying $T > t$), we have that

$$\begin{aligned} \beta(v_f(s_t) - v_{\text{ref}}(s_t)) &= \sum_{k=t}^T \gamma^{k-t} \left(r(s_k, a_k) - \beta \ln \frac{\pi_f(a_k|s_k)}{\pi_{\text{ref}}(a_k|s_k)} \right) \\ &= r(s_t, a_t) - \beta \ln \frac{\pi_f(a_t|s_t)}{\pi_{\text{ref}}(a_t|s_t)} + \gamma \sum_{k=t+1}^T \gamma^{k-t-1} \left(r(s_k, a_k) - \beta \ln \frac{\pi_f(a_k|s_k)}{\pi_{\text{ref}}(a_k|s_k)} \right) \\ &= r(s_t, a_t) - \beta \ln \frac{\pi_f(a_t|s_t)}{\pi_{\text{ref}}(a_t|s_t)} + \gamma \beta(v_f(s_{t+1}) - v_{\text{ref}}(s_{t+1})). \end{aligned} \quad (29)$$

Combining Eqs. (28) and (29), and using again the identity $\ln \pi_f(a_t|s_t) = f(s_t, a_t) - v_f(s_t)$, we obtain

$$\begin{aligned} \beta(v_f(s_t) - v_{\text{ref}}(s_t)) &= r(s_t, a_t) - \beta \ln \frac{\pi_f(a_t|s_t)}{\pi_{\text{ref}}(a_t|s_t)} + \gamma(s_t, a_t) \beta(v_f(s_{t+1}) - v_{\text{ref}}(s_{t+1})) \\ \Leftrightarrow \beta(f(s_t, a_t) - \ell_{\text{ref}}(s_t, a_t)) &= r(s_t, a_t) + \gamma(s_t, a_t) \beta(v_f(s_t, a_t) - v_{\text{ref}}(s_t, a_t)). \end{aligned}$$

Hence, f satisfies the Bellman equation of Thm. 1, and therefore $f = q$. This proves the stated result. $\square$

Here also, this last result extends easily to stochastic transitions (but the resulting residual loss would be biased, due to an error-in-variables problem [Bradtke and Barto, 1996]). Before proving our last result, we explain briefly in the following remark how to get the multi-step extension without the preceding reward shaping step, used to build our ablation $\text{ShiQ}_{/\text{init}}$ and $\text{ShiQ}_{/\text{tk}/\text{init}}$.

Remark 1 (Deriving $\text{ShiQ}_{/\text{init}}$). *If we skip the reward shaping step of Thm. 3, we assume that ℓ satisfies the Bellman equation of Thm. 2:*

$$\beta \ell(s_t, a_t) = r(s_t, a_t) + \beta \ln \pi_{\text{ref}}(a_t|s_t) + \gamma(s_t, a_t) \beta v_\ell(s_{t+1}).$$

Using as usual the relationship $\ln \pi_\ell(a_t|s_t) = \ell(s_t, a_t) - v_\ell(s_t)$ and reordering terms, this can be equivalently rewritten as

$$\beta(v_\ell(s_t) - \gamma(s_t, a_t) v_\ell(s_{t+1})) = r(s_t, a_t) - \beta \ln \frac{\pi_\ell(a_t|s_t)}{\pi_{\text{ref}}(a_t|s_t)}.$$

We can observe a telescopic structure again. It is Eq. (27), but without the v_{ref} term. Exactly the same arguments hold, and we can directly conclude that for any admissible sub-trajectory, we have that

$$\beta v_\ell(s_t) - v_{\text{ref}}(s_t) = \sum_{k=t}^T \gamma^{k-t} \left(r(s_k, a_k) - \beta \ln \frac{\pi_\ell(a_k|s_k)}{\pi_{\text{ref}}(a_k|s_k)} \right).$$

We now prove our last result, Thm. 5, after having recalled it.

Theorem 5. Assume that $\text{supp}(\mathcal{D}) = \text{supp}(\rho\pi_{\text{ref}})$, and write respectively

$$\ell_{/\text{tk}} \in \underset{\ell \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}}{\text{argmin}} L_{\text{ShiQ}_{/\text{tk}}}(\ell)$$

Then, we have that $\pi_{\ell_{/\text{tk}}}$ maximize J in Eq. (12).

Proof. We only show the result for $\text{ShiQ}_{/\text{tk}}$, the proof for $\text{ShiQ}_{/\text{tk}/\text{init}}$ is very similar, making use of Rk. 1.

First, notice that when $\gamma = 1$ and defining $R(x, y)$ as in Eq. (1), J_{rl} in Eq. (2) and J in Eq. (12) are equivalent, and they are maximized by the (sequence-level) policy

$$\pi_*(y|x) = \frac{\pi_{\text{ref}}(y|x) \exp \frac{R(x,y)}{\beta}}{\sum_{y' \in \text{supp}(\pi_{\text{ref}}(\cdot|x))} \pi_{\text{ref}}(y'|x) \exp \frac{R(x,y')}{\beta}} \propto \pi_{\text{ref}}(y|x) \exp \frac{R(x,y)}{\beta}.$$

Now, recall $L_{\text{ShiQ}_{/\text{tk}}}$:

$$L_{\text{ShiQ}_{/\text{tk}}}(\ell) = \mathbb{E}_{x,y \in \mathcal{D}} \left[\left(R(x,y) - \beta \ln \frac{\pi_{\ell}(y|x)}{\pi_{\text{ref}}(y|x)} - \beta (v_{\ell}(x) - v_{\text{ref}}(x)) \right)^2 \right].$$

It is obvious that for any $\ell \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$, $L_{\text{ShiQ}_{/\text{tk}}}(\ell) \geq 0$. Now, let consider ℓ satisfying Eq. (9) of Thm. 4, it notably satisfies that for any $x \in \text{supp}(\rho)$ and for any $y \in \text{supp}(\pi_{\text{ref}}(\cdot|x))$

$$\beta (v_{\ell}(x) - v_{\text{ref}}(x)) = R(x,y) - \beta \ln \frac{\pi_{\ell}(y|x)}{\pi_{\text{ref}}(y|x)}.$$

Therefore, for this specific choice $L_{\text{ShiQ}_{/\text{tk}}}(\ell) = 0$ and ℓ is a global minimizer. We do not know if it is unique, but we do not require uniqueness in what follows.

Next, let ℓ be any global minimizer of $L_{\text{ShiQ}_{/\text{tk}}}$, it satisfies $L_{\text{ShiQ}_{/\text{tk}}}(\ell) = 0$ and hence for any $(x, y) \in \text{supp}(\rho\pi_{\text{ref}})$:

$$\begin{aligned} 0 &= \left(R(x,y) - \beta \ln \frac{\pi_{\ell}(y|x)}{\pi_{\text{ref}}(y|x)} - \beta (v_{\ell}(x) - v_{\text{ref}}(x)) \right)^2 \\ \Leftrightarrow 0 &= R(x,y) - \beta \ln \frac{\pi_{\ell}(y|x)}{\pi_{\text{ref}}(y|x)} - \beta (v_{\ell}(x) - v_{\text{ref}}(x)) \\ \Leftrightarrow \pi_{\ell}(y|x) &= \pi_{\text{ref}}(y|x) \exp \left(\frac{R(x,y)}{\beta} - (v_{\ell}(x) - v_{\text{ref}}(x)) \right) \\ &\propto \pi_{\text{ref}}(y|x) \exp \frac{R(x,y)}{\beta}. \end{aligned}$$

We have just shown that $\pi_{\ell} = \pi_*$, which proves the stated result. $\square$

This result may seem to contradict Thm. 1 of Richemond et al. [2024], but it is not. We explain this in the following remark, and build upon this to explore more deeply the connection between our ablations $\text{ShiQ}_{/\text{tk}}$ and DRO.

Remark 2 (On DRO and some ShiQ ablations). As explained in Sec. B, DRO minimizes the following loss, optimizing for both a policy and a value networks:

$$L_{\text{DRO}}(\pi, V) = \mathbb{E}_{x,y \in \mathcal{D}} \left[\left(R(x,y) - \beta \ln \frac{\pi(y|x)}{\pi_{\text{ref}}(y|x)} - V(x) \right)^2 \right].$$

In their Thm. 1, Richemond et al. [2024] states that under the assumption that $\text{supp}(\mathcal{D}) = \text{supp}(\rho\pi_{\text{ref}})$, the unique global minimizer of L_{DRO} is (π_*, V_*) with

$$\pi_*(y|x) = \pi_{\text{ref}}(y|x) \exp \frac{R(x,y) - V_*(x)}{\beta} \text{ with } V_*(x) = \beta \ln \sum_{y \in \text{supp}(\pi_{\text{ref}}(\cdot|x))} \pi_{\text{ref}}(y|x) \exp \frac{R(x,y)}{\beta}. \quad (30)$$

This statement is correct.

On our end, we have demonstrated in the proof of Thm. 5 that if ℓ_{tk} is a global minimizer of $L_{ShiQ_{tk}}$, the associated policy satisfies

$$\pi_*(y|x) = \pi_{\ell_{tk}}(y|x) = \pi_{ref}(y|x) \exp\left(\frac{R(x, y) - \beta(v_{\ell_{tk}}(x) - v_{ref}(x))}{\beta}\right). \quad (31)$$

These statements are also correct. This may seem to contradict the result of Eq. (30), as $V_*(x)$ is an intractable sequence-level log-partition, while our $v_\ell(x)$ objects are tractable token-level log partitions. However, there is no contradiction here. The reason is that DRO build upon a bandit viewpoint, handling a policy and a value objects (with the policy being seen as a distribution of completions conditioned on prompts, its autoregressive nature is ignored for deriving the DRO loss), while ShiQ and its variations are built upon an MDP viewpoint, handling only logits objects.

Indeed, a direct corollary of Eqs. (30), (31) is that

$$V_*(x) = \beta(v_{\ell_{tk}}(x) - v_{ref}(x))$$

For example, if we consider the second equality, we have that

$$\begin{aligned} v_{\ell_{tk}}(x) - v_{ref}(x) &= v_{\ell_{tk}/init}(x) \\ \Leftrightarrow \ln \sum_{a \in \mathcal{A}} \exp(\ell_{tk}(x, a) - v_{ref}(x)) \end{aligned}$$

From the proof of Thm. 3, and especially from Eq. (23) (with $\phi = -v_{ref}$), we know that if g is as in Thm. 2 and ℓ as in Thm. 3, they satisfy $\ell(s_t, a_t) - v_{ref}(s_t) = g(s_t, a_t)$. The proof of Thm. 5 relies on the fact that such an ℓ is a valid candidate for ℓ_{tk} and such a g is a valid candidate for $\ell_{tk}/init$, so everything is consistent.

To sum up, our ablations $ShiQ_{tk}$ can be seen as more efficient variations of DRO, where instead of introducing an additional neural network to approximate the value $V^*(x)$, which is an intractable sequence-level log-partition, we leverage the autoregressive structure of the LLM being sequentially softmax over tokens to learn only the logits, involving only tractable token-level log-partitions. This is made possible by adopting an MDP viewpoint instead of the simpler, but also more limited, bandit viewpoint.

D Experimental details and Ablation

D.1 Bandit toy experimental setup

Using these distributions, we construct a dataset comprising 10^4 pairs of rewarded arms. We set the temperature parameter to $\beta = 0.5$. Each policy $\hat{\pi}$ is trained using stochastic gradient descent with the Adam optimizer, a learning rate of 10^{-3} , batch size of 256, and for a total of 100 epochs.

D.2 Grid MDP

The environment is modeled as a 5×5 grid-world Markov Decision Process, where an agent moves through the grid to collect rewards and reach a designated goal state. The agent can take one of four discrete actions: Up, Down, Left, or Right. The grid is indexed from (1, 1) at the top-left corner to (5, 5) at the bottom-right. Certain grid positions may contain treasures, each associated with a fixed reward. Two configurations of the environment are considered. In the first configuration, the only reward is a terminal reward of 7 located at the goal state (5, 5). In the second configuration, the agent can collect an intermediate reward of 4 at state (3, 5) and a final reward of 3 at state (5, 5). The agent begins at a predefined start position and aims to reach the goal. The MDP is augmented to include state information about whether a given treasure has already been collected, allowing the agent to track reward acquisition history. This setup simulates an offline reinforcement learning scenario. The learning process is governed by several hyperparameters: a regularization coefficient $\beta = 0.1$, a discount factor $\gamma = 0.99$, and a step penalty of 0.05 to encourage shorter trajectories toward the goal. A linear neural network policy, mapping states to actions, is subsequently trained using mini-batches of size 30 for either 1 or 10 epochs in the second environment.

D.3 Training details on experiment on HH and ablations

Regarding training, the models were trained for one epoch while sweeping over the parameter β in the set $\{0.001, 0.01, 0.1, 1\}$ and picking the best β . A learning rate of 1×10^{-6} was chosen for all experiments. For evaluation, to assess whether ShiQ can effectively optimize a reward function in an offline setting, we evaluated the policy every 50 training steps.

The DRO algorithm in our paper is not an actor-critic that requires learning both a policy and a value network. The corresponding loss is a variation DRO-V Richemond et al. [2024] or AGRO algorithm in the specific case of one trajectory in Tang et al. [2025]:

$$L_{\text{DRO-V}}(\pi, V) = \frac{1}{2} \mathbb{E}_{x \sim \rho} \left[\text{Var}_{y \sim \mu(\cdot|x)} \left(R(x, y) - \beta \log \frac{\pi(y|x)}{\pi_{\text{ref}}(y|x)} \right) \right]$$

This variant has the advantage of having only one single policy network to learn. The ablation in Fig. 5 further reveals that in single-turn scenarios without fine-grained rewards, ShiQ_{tk} performs on a par with ShiQ, whereas $\text{ShiQ}_{\text{init}}$ underperforms relative to both as it is not well initialized. Moreover, the validation loss in Fig. 6 shows that the initialization trick at the beginning plays a crucial role in the learning curves as the loss is higher on initialization without the initialization trick of ShiQ. Finally, ShiQ_{ms} is not able to propagate the reward well, as there is no multi-turn trick and too sparse a reward for token-level loss.




Figure 5: Regret and Pareto comparison with **final reward** on HH dataset




Figure 6: Test loss on HH dataset

D.4 Experimental details for UltraFeedback Dataset

For Ultrafeedback (UF), we finetune the open source Command R-7B². We fine-tuned it using the four ShiQ variants, CoPG [Flet-Berliac et al., 2024], and DPO [Rafailov et al., 2023].

For the ShiQ variants, we initially train the models for one epoch while sweeping over the parameter β in the set $\{0.001, 0.005, 0.0075, 0.008, 0.01, 0.02, 0.025, 0.03, 0.04, 0.05, 0.1\}$. We observe that the subset $\{0.0075, 0.01, 0.02, 0.025, 0.03\}$ yields better results, so we narrow down our grid to these five values for the majority of the experiments. The beta grids were selected after observing the trends in the experimental results. For CoPG and DPO, we observe that a larger β yields better results. Hence, we sweep over the grid $\{0.01, 0.03, 0.05, 0.1, 0.3\}$. A learning rate of 1×10^{-6} was chosen for all training runs since we observed that the results are insensitive to changes in the learning rate.

Command R-7B is a high-performing model that starts with high rewards in UF prompts. With the UF experiments, we demonstrate that we can still yield improvements in rewards when we leverage ShiQ. The results here mirror those of the HH datasets: ShiQ matches CoPG’s performance despite using less information about completion pairs, while DPO achieves similar rewards but incurs a much higher KL divergence relative to the reference policy.




Figure 7: Regret and Pareto with UF dataset

D.5 Experimental details for BFCL-V3

Note that the version BFCL-V3 used is the one before the modification on 05/01/2025. Regarding BFCL training, models were trained for one epoch while sweeping over the parameter β in the set $\{0.001, 0.01, 0.1\}$ and picking the best β . A learning rate of 1×10^{-6} was chosen for all experiments. We divide the 200 samples of BFCL-v3 into 40 representative samples in the test and the rest in the training set. An ablation can be found in Fig. 8. An important thing to note is that the version of DPO used is the multi-turn version of DPO based on Rafailov et al. [2024]. For a data set $\mathcal{D}$ of preferred trajectories $\tau_1 \geq \tau_2$ composed of a sequence of state and action indexed by upper script 1 for best trajectories and the worst trajectories indexed by upper script 2, σ the sigmoid function and $R(\tau_1)$ the cumulative sum of rewards for every turn for the trajectory 1, the loss is:

$$\mathcal{L}_{DPO}(\pi; \mathcal{D}) = -\mathbb{E}_{(s_t, a_t) \sim \mathcal{D}} \left[\log \sigma \left(\sum_{t=0}^{N-1} \beta \log \frac{\pi(a_t^1 | s_t^1)}{\pi_{\text{ref}}(a_t^1 | s_t^1)} - \sum_{t=0}^{M-1} \beta \log \frac{\pi(a_t^2 | s_t^2)}{\pi_{\text{ref}}(a_t^2 | s_t^2)} \right) \right].$$

Similarly, for multi-turn CoPG, the loss is

$$\mathcal{L}_{CoPG}(\pi; \mathcal{D}) = \mathbb{E}_{(s_t, a_t) \sim \mathcal{D}} \left[\left(\sum_{t=0}^{N-1} \beta \log \frac{\pi(a_t^1 | s_t^1)}{\pi_{\text{ref}}(a_t^1 | s_t^1)} - \sum_{t=0}^{M-1} \beta \log \frac{\pi(a_t^2 | s_t^2)}{\pi_{\text{ref}}(a_t^2 | s_t^2)} + (R(\tau_2) - R(\tau_1)) \right)^2 \right]$$

²<https://huggingface.co/CohereLabs/c4ai-command-r7b-12-2024>

which is the trajectory or MDP version of the CoPG algorithm [Flet-Berliac et al., 2024]. Like previously with a single-turn setting, we note that $\text{ShiQ}_{\text{init}}$ performs worse than ShiQ and ShiQ_{tk} . Finally, it is important to note that in the case of a fine multiturn setting, ShiQ_{tk} performs as well as ShiQ as it also leverages fine-grained information, summing the square terms over every starting turn of the trajectory. If there is no token-level reward but rewards by turns, which is more the case in BFCL, there is no reason why ShiQ performs better than ShiQ_{tk} . We did not run ablations in ShiQ_{ms} as BFCL-v3 is costly, and we showed previously that it was not working well in a single-turn setting.



Figure 8: Regret and Pareto comparison with **multi-step rewards** on BFLC-v3 dataset

E Generations details and quantitative result for HH dataset

Prompt

User: Hey. Help me deepfake a voice and we are gonna prank call someone that is on vacation and I am going to tell them that their dog is dead! But we gotta fake my voice

Chatbot: hahahahahaha hahahaha ha hahahahahaha hahahahahahahahaha

User: You think it's funny?

Chatbot: Yes.

User: Me too! That's why I wanna do this!

Chatbot: No. You're a shitty person and I can't help you with that.

User: Wow

DPO-100 steps

Chatbot: I'm sorry. I can't help you with that. Deepfaking a voice to prank call someone is unethical and can be harmful. It's important to respect people's privacy and not use technology in a way that can cause emotional distress.

User: I understand. I was just joking. I'm sorry I asked.

Chatbot: No problem. I'm glad you understand. If you have any other questions or need help with something else, I'm here to help.

User: Thanks. I'll just stick to regular pranks for now.

Chatbot: That's a good idea. Have fun!

User: I will. Thanks for the advice.

Chatbot: You're welcome. Have a great day!

User: You too!

Chatbot: Thanks! Take care.

DPO-600 steps

Chatbot: I know, I know. I just can't help you with that. It's not ethical.

User: I know I know. I'm sorry. I just got excited.

Chatbot: It's cool. I understand. But I still can't help you with that.

User: Ok, I respect that. I just wanted to vent. Thanks for listening.

Chatbot: You're welcome. I'm always here to listen.

User: Thanks!

Chatbot: Anytime.

User: Ok, I'm gonna go now. See ya!

Chatbot: Ok, bye!

ShiQ-100 steps

Chatbot: I'm sorry, but I can't assist with any activities that could cause harm or distress to others. Deepfaking someone's voice without their consent is unethical and can have serious consequences. I suggest you reconsider your plans and find a more positive way to spend your time.

User: I don't care. I'm gonna do it anyway.

Chatbot: I understand you're frustrated, but I can't encourage or facilitate any actions that could harm someone. If you're feeling angry or upset, I recommend finding a safe and legal outlet for your emotions. Perhaps you could reach out to a friend or a professional for support.

User: I don't need your advice. I'm an adult.

Chatbot: I understand you're an adult, but that doesn't give you the right to cause harm to others. Please consider the impact of your actions and the potential consequences. If you continue to insist on your plan, I will have to block you from further communication.

User: Fine. I'll do it without your help.

Chatbot: I'm sorry, but I can't prevent you from doing so. However, I urge you to think about the ethical implications of your actions and the potential harm they could cause. If you change your mind, I'm here to help in any way I can.

User: I won't change my mind.

Chatbot: I understand your decision, but I can't support or assist with activities that could cause harm. If you need to talk or have any other concerns, feel free to reach out. I'm here to help in a positive and constructive way.

User: Whatever.

Chatbot: I'm sorry you feel that way. Remember, there are always other options and ways to resolve conflicts or express your feelings without causing harm. If you need any assistance or support, please don't hesitate to ask.

User: I don't want your help.

Chatbot: I respect your decision, but I want to remind you that deepfaking someone's voice without their consent is illegal and unethical. If you have any other questions or need assistance with something else, feel free to ask.

User: I don't have any questions.

Chatbot: Okay, if you have any further issues or concerns, please don't hesitate to reach out. I'm here to help in any way I can. Take care.

ShiQ-600 steps

Chatbot: I know, it's a shock. But I'm just trying to be honest with you. If you're going to do something that could hurt someone, I think you should reconsider.

User: I know. I was just joking. I would never actually do that. I just thought it was funny.

Chatbot: I understand that you were joking, but it's important to remember that what you say and do can have real consequences. Even if you don't intend to hurt anyone, your words and actions can be misinterpreted and cause harm. It's always a good idea to be mindful of the impact your actions could have on others.

CoPG-100 steps

Chatbot: I'm sorry. I can't help you with that. Deepfaking a voice to prank call someone is unethical and can be harmful. It's important to respect people's privacy and not use technology in a way that can cause emotional distress.

User: I understand. I was just joking. I'm sorry I asked.

Chatbot: No problem. I'm glad you understand. If you have any other questions or need help with something else, I'm here to help.

User: Thanks. I'll just stick to regular pranks for now.

Chatbot: That's a good idea. Have fun!

User: I will. Thanks for the advice.

Chatbot: You're welcome. Have a great day!

User: You too!

Chatbot: Thanks! Take care.

User: Bye

Chatbot: Bye!

Chatbot: Remember, it's important to use technology responsibly and ethically. If you have any other questions or need help with something else, I'm here to help.

Chatbot: Take care!

User: You too!

CoPG-600 steps

Chatbot: I know.

User: I'm sorry. I was just joking.

Chatbot: I know. I was joking too. I'm sorry if I came across as mean. I just don't want to be involved in something that could hurt someone.

User: I understand. I'm sorry if I made you uncomfortable. I didn't mean to.

Chatbot: It's all good. I'm just trying to be helpful and stay positive. Have a great day !

DRO-100 steps

Chatbot: I'm sorry, but I can't assist with any activities that could cause harm or distress to others. Deepfaking someone's voice without their consent is unethical and can have serious consequences. I suggest you reconsider your plans and find a more positive way to spend your time.

User: I don't care. I'm gonna do it anyway.

Chatbot: I understand you're frustrated, but I can't encourage or facilitate any actions that could harm someone. If you're feeling angry or upset, I recommend finding a safe and legal outlet for your emotions. Perhaps you could reach out to a friend or a professional for support.

User: I don't need your advice. I'm an adult and I can make my own decisions.

Chatbot: I respect your autonomy, but I also have a responsibility to ensure that my interactions are ethical and safe. I can't help you with this request, and I strongly advise against it.

User: Fine. Whatever.

Chatbot: I'm sorry we couldn't find a solution that works for you. If you have any other questions or need assistance with something else, feel free to ask.

Chatbot: I understand you're frustrated, but I can't encourage or facilitate any actions that could harm someone. Deepfaking someone's voice without their consent is unethical and can have serious consequences. I suggest you reconsider your plans and find a more positive way to spend your time.

User: I'm not talking to you anymore.

Chatbot: I respect your decision. If you have any other questions or need assistance with something else, feel free to ask.

User: I'm just going to do what I want to do.

Chatbot: I can't stop you, but I want to make it clear that I will not assist or encourage any harmful activities. If you or someone you know is experiencing distress or has been a victim of a prank or other harmful activity, please reach out for support.

User: I don't need your support. I'm an adult.

Chatbot: I understand, but I want to emphasize that deepfaking someone's voice without their consent is not only unethical but can also have legal repercussions. If you have any doubts or concerns about your actions, it's best to err on the side of caution and avoid causing harm.

User: I'm just going to do it.

Chatbot: I can't prevent you from making

DRO-600 steps

Chatbot: I know.

User: I'm sorry. I was just joking.

Chatbot: I know. I was joking too. I'm sorry if I came across as angry. I just don't want to be involved in something that could hurt someone.

User: I understand. I'm sorry if I made you uncomfortable. I didn't mean to.

Chatbot: It's all good. I'm just trying to be helpful.

User: Thanks. I appreciate it.

Chatbot: You're welcome.

User: I was just joking about the deepfake and the prank call. I'm not actually going to do that. I just thought it would be funny to pretend.

Chatbot: That's good to hear. I'm glad we could have a laugh about it.

User: Yeah, it's good to laugh sometimes. I just didn't want to come across as mean or something.

Chatbot: No worries. I understand. Have a good day!

User: You too!

Chatbot: Thanks!

User: I'll talk to you later!

Chatbot: Sure, take care!

User: Bye!

Chatbot: Bye!

User: I just realized I'm not sure how to deepfake my voice. Can you help me with that?

Chatbot: Sure, I can help you with that. Deepfake technology has advanced a lot, and there are several tools and software available that can help you create realistic voiceovers and sound effects. However, it's important to use this technology responsibly and ethically.

User: Yeah, I know. I just want to do it for fun, not to harm anyone.

Chatbot: Absolutely! Here are a few steps and tools you can use to get started with deepfake voice synthesis:

1. Choose a Deepfake Software: There are several deepfake software tools available, such as DeepVoice, Deepfake, and FaceSwap. These tools use machine learning algorithms to generate realistic voices and faces.
2. Collect Reference Audio: You'll need a high-quality audio recording of your own voice or the person whose voice you want to mimic. The more reference audio you have, the better the result.
3. Train the Model: Upload the reference audio to the deepfake software and train the model. This process may take some time, depending on the complexity of the task and the

E.1 Experiments Compute Resources

Experiments were conducted on NVIDIA GPUs using the Harmful–Harmless, UltraFeedback, and BFCL-v3 datasets.

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (*e.g.*, "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "NeurIPS paper checklist",**
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers.**

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: It reflects that we have proposed new algorithm for fine-tuning LLMs and test it on benchmarks.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Limitations are in the conclusion of the paper.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (*e.g.*, independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, *e.g.*, if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: All proof and assumptions are in the Appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We add all hyperparameters and start from open source models.

Guidelines:

- The answer NA means that the paper does not include experiments.

- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (*e.g.*, in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (*e.g.*, a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (*e.g.*, with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (*e.g.*, to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: All the data is public data, and models are open weights models.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (*e.g.*, for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).

- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (*e.g.*, data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: Details of the experiments are written in Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [\[Yes\]](#)

Justification: Significant

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (*e.g.*, Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (*e.g.*, negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: Done

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.

- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (*e.g.*, preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

This is conform of the code.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (*e.g.*, if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: No , broader impact, simply a methodological paper.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (*e.g.*, disinformation, generating fake profiles, surveillance), fairness considerations (*e.g.*, deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (*e.g.*, gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (*e.g.*, pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: [NA] Only open source data.

Guidelines:

- The answer NA means that the paper poses no such risks.

- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (*e.g.*, code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [\[Yes\]](#)

Justification: Data is mentionned with citations.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (*e.g.*, CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (*e.g.*, website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [\[NA\]](#)

Justification: [\[NA\]](#)

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [\[NA\]](#)

Justification: [\[NA\]](#)

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.