

# 000 BEYOND IMITATION: RECOVERING DENSE REWARDS 001 002 FROM DEMONSTRATIONS 003 004

005 **Anonymous authors**

006 Paper under double-blind review

## 007 008 ABSTRACT 009

011 Conventionally, supervised fine-tuning (SFT) is treated as a simple imitation learning  
012 process that only trains a policy to imitate expert behavior on demonstration  
013 datasets. In this work, we challenge this view by establishing a fundamental  
014 equivalence between SFT and Inverse Reinforcement Learning. We prove that  
015 the SFT objective is a special case of Inverse Q-Learning, which implies that the  
016 SFT process does not just learn a policy, but also an implicit, dense, token-level  
017 reward model that explains the expert demonstrations. We then show how to  
018 recover this dense reward signal directly from the SFT model by formulating a  
019 baseline-relative reward function. The availability of such a dense reward model  
020 offers numerous benefits, providing granular credit assignment for each token  
021 generated. We demonstrate one key application by using these recovered rewards  
022 to further improve the policy with reinforcement learning. Our method, Dense-Path  
023 REINFORCE, consistently outperforms the original SFT models on instruction-  
024 following benchmarks. This work reframes SFT not merely as policy imitation but  
025 as a powerful reward learning mechanism, opening new possibilities for leveraging  
026 expert demonstrations.

## 027 1 INTRODUCTION

030 Large Language Models (LLMs) (Liu et al., 2024; Comanici et al., 2025; Achiam et al., 2023)  
031 have rapidly developed from research prototypes to general-purpose assistants that plan, reason, and  
032 generate helpful responses across domains. A significant driver of these capabilities is *post-training*  
033 on *demonstrations*—often called *Learning from Demonstrations* (LfD)—where a pretrained model  
034 is refined to follow expert responses (Ouyang et al., 2022; Chen et al., 2024). In practice, LfD  
035 is implemented almost exclusively as *Supervised Fine-Tuning* (SFT): teacher-forced maximum  
036 likelihood on expert tokens conditioned on prompts. Because SFT matches expert sequences, it is  
037 commonly framed as *imitation learning* (Xiao et al., 2024; Shaikh et al.; Sun, 2024) in which the  
038 model learns only to mimic expert behavior.

039 This paper argues that the imitation-only view is incomplete. We show that, under standard assumptions  
040 for token-level generation, SFT admits a precise interpretation through the lens of *Inverse*  
041 *Reinforcement Learning* (IRL) (Ng & Russell, 2000). Specifically, on the token Markov decision process  
042 (MDP) without discount, the token-level SFT objective is *equivalent* to optimizing the reduced  
043 objective of Inverse Soft-Q Learning (IQ-Learn) (Garg et al., 2021). In this view, SFT does more than  
044 fit a policy: it implicitly learns a *dense token-level reward* that rationalizes expert demonstrations,  
045 aligning SFT with the credit-assignment perspective of MaxEnt IRL and GAIL (Ziebart et al., 2008;  
Ho & Ermon, 2016).

046 The IQ-Learn perspective also yields a valid recipe for further improving an SFT policy. First, we  
047 prove a *dual-contraction* property of the IQ-Learn saddle: the error of the reward estimation is  
048 bounded by the policy’s occupancy error, so a reasonably accurate policy implies an even more  
049 stable reward estimation (near the saddle). Second, we show how to *recover a dense reward* directly  
050 from the trained SFT model. Using the soft-optimality identity and potential-based shaping (Ng  
051 et al., 1999), the teacher’s token log-probability decomposes as the task reward plus a telescoping  
052 potential value function. This implies two design choices. (i) We eliminate the value term via shaping,  
053 which keeps the token reward dense and avoids tricky value estimation. (ii) To avoid the length  
bias of raw log-likelihoods (non-positive by construction) and stabilize credit assignment, we use a

054 *baseline-relative* reward where baseline is a checkpoint during SFT training. This choice measures  
 055 *incremental performance* gained during SFT and empirically reduces variance. Together, these results  
 056 justify a simple reinforcement step that stays in the LfD setting: we optimize the SFT policy with  
 057 *token-level, undiscounted REINFORCE* (Williams, 1992; Ahmadian et al., 2024) using the dense  
 058 baseline-relative reward.

059 We evaluate this recipe on four pretrained LLMs and four public instruction-following benchmarks  
 060 using the same demonstration data for SFT and RL. Despite operating strictly in the LfD setting, the  
 061 resulting policy improves over the SFT model in head-to-head win rate and standardized multi-turn  
 062 scores, showing competitiveness with other LfD baselines such as SPIN (Chen et al., 2024) and  
 063 GSIL (Xiao et al., 2024).

064 Our primary contributions are as follows: (i) We establish formal equivalence between token-level  
 065 SFT and the reduced objective of IQ-Learn on the token MDP, reframing SFT as implicit dense reward  
 066 learning rather than pure imitation. (ii) We prove that near the IRL saddle, the reward estimation error  
 067 is bounded by the policy occupancy error, explaining why rewards recovered from an SFT policy can  
 068 be more stable than the policy itself. (iii) We construct meaningful token-level rewards through reward  
 069 shaping theory and the strategic selection of a reward baseline. (iv) We instantiate these insights in  
 070 a minimal reinforcement learning algorithm that uses token-level, undiscounted baseline-relative  
 071 reward as the learning objective. (v) Across four pretrained backbones and four instruction-following  
 072 evaluations, this method consistently improves over SFT and matches or exceeds other LfD baselines.

## 073 074 075 076 2 RELATED WORK

077 078 **Independent and contemporaneous (posted earlier) work: reward signals inside LLMs.** Before  
 079 this work, Li et al. (2025) also argues that LLMs contain useful token-level reward signals through the  
 080 lens of IRL. While we share the same theoretical framework on the equivalence of IRL and SFT, the  
 081 theoretical part of our work was developed *independently* before the publication of their work. Their  
 082 experimental focus is to extract sentence-level rewards through aggregation of token-level rewards,  
 083 often from instruction-tuned LLMs, and then use the aggregated reward for policy improvement and  
 084 reward-based pair-wise classification. Our setting and emphasis are different: we operate in LfD  
 085 with pretrained backbones and develop a shaping- and baseline-based reward construction that makes  
 086 **dense rewards workable for policy improvement in practice.**

087 **Imitation learning and LfD for LLMs.** Beyond direct cloning, several LfD approaches leverage  
 088 self-generated data to improve a policy without requiring explicit preference pairs. These methods  
 089 reframe the learning problem to go beyond the simple negative log-likelihood objective of SFT: SPIN  
 090 uses self-play fine-tuning to convert weaker models into stronger ones (Chen et al., 2024). (Li et al.,  
 091 2024) found that SPIN is a special case of IRL; however, they still focus on the gap between policy  
 092 and expert at the sample level. GSIL also uses both real demonstration data and self-generated model  
 093 data, but formulates the problem from an imitation learning perspective (Xiao et al., 2024). Our work  
 094 differs in both analysis and mechanism: we remain strictly in the LfD setting, but *re-interpret* SFT  
 095 through an IRL lens (SFT  $\equiv$  IQ-Learn on the token MDP).

096 **Preference-based post-training (RLHF, DPO family, GRPO).** Another line of work treats post-  
 097 training as optimization from *pairwise* human (or AI) preferences. PPO-based RLHF (Ouyang et al.,  
 098 2022) fits a reward model and then optimizes the policy with reinforcement learning. DPO (Rafailov  
 099 et al., 2023) replaces explicit reward learning and online rollouts with a direct, classification-style  
 100 objective. Recent GRPO-style methods explore preference optimization without an explicit critic:  
 101 *group relative* policy optimization has been used in scaling efforts to stabilize on-policy updates via  
 102 group-normalized advantages (Shao et al., 2024). These methods require preference data or verifiable  
 103 rewards and thus are outside our scope.

104 **Connection of reinforcement learning and SFT.** Xiao et al. establishes a theoretical connection  
 105 between reinforcement learning and imitation learning, revealing that RLHF implicitly performs  
 106 imitation learning on the preference data distribution. Qin & Springenberg (2025) unifies SFT with  
 107 RL through importance sampling. These studies are somewhat related to our work, but they primarily  
 focus on the relationship between RL and SFT, whereas we analyze SFT from the perspective of IRL.

108 

### 3 PRELIMINARIES

109

110 This section introduces the minimal background needed to follow our methodology and proofs.  
111 We formalize the token-level MDP for autoregressive generation, recall the entropy-regularized  
112 optimality equations, restate MaxEnt IRL in an occupancy form, explain the  $Q$ -space reduction used  
113 by IQ-Learn.

114  
115 **Problem setup and notation.** We model generation as a finite-horizon token MDP  $(\mathcal{S}, \mathcal{A}, f, \rho_0)$   
116 with deterministic concatenation  $f(s, a) = s|a$  and horizon  $H$ . A state  $s_t$  is the prompt plus the  
117 tokens generated so far, the action  $a_t$  is the next token, and an LLM induces a policy  $\pi(a | s)$ . We  
118 write the (state–action) *occupancy measure* of policy  $\pi$  as

119  
120 
$$\rho_\pi(s, a) = \sum_{t=0}^{H-1} \Pr_\pi(s_t = s, a_t = a), \quad \langle \rho_\pi, r \rangle := \sum_{s, a} \rho_\pi(s, a) r(s, a).$$
121

122 For any real-valued function  $Q : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ , define the soft value  $V(s) = \log \sum_a \exp Q(s, a)$  and  
123 the Boltzmann policy  $\pi_Q(a | s) \propto \exp Q(s, a)$  (temperature fixed to 1 throughout).

124  
125 **Soft-optimality equations.** In entropy-regularized control (Haarnoja et al., 2017), optimizing  
126  $\mathbb{E}_{a \sim \pi(\cdot | s)}[Q^*(s, a)] - \beta H(\pi(\cdot | s))$  over  $\pi(\cdot | s)$  yields the familiar logit form of the optimal policy  
127 and value:

128  
129 
$$\pi^*(a | s) = \exp\left(\frac{1}{\beta}(Q^*(s, a) - V^*(s))\right), \quad V^*(s) = \beta \log \sum_{a \in \mathcal{A}} \exp\left(\frac{1}{\beta}Q^*(s, a)\right). \quad (1)$$

130 That is,  $\pi^*(\cdot | s)$  is the Boltzmann distribution over  $Q^*(s, \cdot)$  and  $V^*(s)$  is the corresponding log-  
131 partition. A full derivation is provided in Appendix A.2.

132  
133 **MaxEnt IRL in occupancy space.** Maximum-entropy IRL seeks a reward  $r$  that rationalizes  
134 expert behavior by comparing expert and learner occupancies while keeping the policy stochastic via  
135 entropy Ziebart et al. (2008); Ho & Ermon (2016):

136  
137 
$$L(\pi, r) = \langle \rho_E - \rho_\pi, r \rangle - H(\pi) - \psi(r). \quad (2)$$

138 Here  $\psi$  is a convex regularizer on rewards (for identifiability/stability). The saddle point of (2)  
139 matches occupancies ( $\rho_{\pi^*} = \rho_E$ ) and produces a reward  $r^*$  unique up to potential-based shaping.

140  
141 **IQ-Learn: a  $Q$ -space reduction.** IQ-Learn re-parameterizes the IRL objective so that, after  
142 minimizing over  $\pi$ , one optimizes a *concave* functional of  $Q$  Garg et al. (2021). The policy minimizer  
143 is  $\pi_Q(a | s) = \exp(Q(s, a) - V(s))$ , and the reduced objective  $J^*(Q)$  aggregates the “soft-  
144 advantage”  $Q(s, a) - V(f(s, a))$  along expert trajectories. On a deterministic token tree ( $f(s, a) =$   
145  $s'$ ), telescoping arguments become particularly simple and will later allow us to show that *token-level*  
146 *SFT is equivalent to* maximizing  $J^*(Q)$  under a linear conjugate (Step 1).

147 

## 4 METHODOLOGY

148

149  
150 **High-level outline.** Our methodology follows three steps. **(S1)** We show that the token-level SFT  
151 objective is *equivalent to* the reduced IQ-Learn objective under a mild regularizer. **(S2)** Within the  
152 IRL/IQL framework, we prove that the *reward estimation error* is controlled by the *policy error* in  
153 occupancy space. **(S3)** We extract a baseline-relative, log-likelihood based dense reward (Chan et al.,  
154 2024) from the SFT model and show that any improvement on this proxy transfers to improvement  
155 under the true objective.

156 

### 4.1 STEP 1: SFT IS EQUIVALENT TO A SPECIAL CASE OF IQ-LEARN

157

158 **Statement.** Let  $J^*(Q)$  denote the reduced IQ-Learn objective after minimizing over  $\pi$  (Garg et al.,  
159 2021). On the token MDP with  $\gamma = 1$  and a linear conjugate (i.e., no extra reward regularization  
160 beyond convexity), maximizing  $J^*(Q)$  is *equivalent to* maximizing the teacher-forced log-likelihood  
161 on expert tokens:

$$\max_Q J^*(Q) \equiv \max_Q \mathbb{E}_{(s, a) \sim \rho_E} [\log \pi_Q(a | s)],$$

162 where  $\pi_Q(a | s) \propto \exp Q(s, a)$  and  $V(s) = \log \sum_a e^{Q(s, a)}$ .  
 163

164 **Intuition.** The reduction  $J^*(Q)$  aggregates a “soft-advantage” term of the form  $Q(s, a) - V(f(s, a))$   
 165 along expert trajectories. On a deterministic token sequence, the value contributions telescope  
 166 across time, and the identity  $\log \pi_Q(a | s) = Q(s, a) - V(s)$  converts the objective into the SFT  
 167 log-likelihood.

168 **Proposition 1** (SFT  $\equiv$  IQ-Learn with a linear conjugate). *On the token MDP with discount*  
 169 *rate  $\gamma = 1$ , maximizing  $J^*(Q)$  is equivalent to minimizing the token-level SFT loss*  $\mathcal{L}_{\text{SFT}}(\theta) =$   
 170  $\mathbb{E}_{(s, a) \sim \rho_E}[-\log \pi_\theta(a | s)]$ , where  $\pi_\theta(a | s) \propto \exp Q_\theta(s, a)$ .

171 *Proof.* See Appendix A.4 for a complete derivation via telescoping and the identity  $\log \pi_Q = Q - V$ .  
 172

173 **Takeaway.** SFT is not only policy imitation: it is *exactly* the  $Q$ -space objective of an IQ-Learn  
 174 instance on the token MDP. Consequently, SFT logits can be treated as a  $Q$ -function without leaving  
 175 the IRL/IQL lens, consistent with the token-level perspective in *From r to  $Q^*$*  (Rafailov et al.).  
 176

## 4.2 STEP 2: REWARD ERROR IS CONTROLLED BY POLICY ERROR (IRL VIEW)

178 We adopt the convex-analytic IRL objective (Ho & Ermon, 2016):  
 179

$$L(\pi, r) = \langle \rho_E - \rho_\pi, r \rangle - H(\pi) - \psi(r). \quad (3)$$

181 Let  $r^*$  be a reward at the IRL saddle. For any  $\pi$ , let the reward best response be  $\hat{r}(\pi) :=$   
 182  $\arg \max_r L(\pi, r)$ . Measure the *policy error* by  $\varepsilon_\pi := \|\rho_\pi - \rho_E\|_*$  and the *reward error* by  
 183  $\varepsilon_r := \|\hat{r}(\pi) - r^*\|$ , where  $\|\cdot\|$  and  $\|\cdot\|_*$  are dual norms.

184 **Theorem 2** (Dual contraction: reward error  $\leq$  policy error). *If  $\psi$  is  $\mu$ -strongly convex in  $\|\cdot\|$ , then*  
 185 *for any policy  $\pi$ ,*

$$\|\hat{r}(\pi) - r^*\| \leq \frac{1}{\mu} \|\rho_\pi - \rho_E\|_*.$$

188 *Proof.* By first-order optimality for the reward player,  $\nabla \psi(\hat{r}(\pi)) = \rho_E - \rho_\pi$  and  $\nabla \psi(r^*) = \rho_E - \rho_{\pi^*}$ .  
 189 At the saddle  $\rho_{\pi^*} = \rho_E$ , so  $\nabla \psi(r^*) = 0$  and hence  $\nabla \psi(\hat{r}(\pi)) - \nabla \psi(r^*) = \rho_E - \rho_\pi$ . Strong convexity  
 190 implies  $\mu$ -strong monotonicity of  $\nabla \psi$ ; applying Hölder’s inequality in dual norms yields the claim.  
 191 See Appendix A.6 for details.

192 **Takeaway.** Learning a reward is at least as stable as learning the policy near the saddle—precisely  
 193 the property we need before using the (SFT-derived) reward to further improve the policy.  
 194

## 4.3 STEP 3: FROM AN SFT-DERIVED DENSE REWARD TO POLICY IMPROVEMENT

197 **(A) Using SFT logits as a reward via potential shaping.** Combining the soft Bellman identity with  
 198  $\log \pi_{\text{SFT}}(a | s) = Q_{\text{SFT}}(s, a) - V_{\text{SFT}}(s)$  yields  
 199

$$\log \pi_{\text{SFT}}(a_t | s_t) = r(s_t, a_t) + (V_{\text{SFT}}(s_{t+1}) - V_{\text{SFT}}(s_t)), \quad (4)$$

201 so  $\log \pi_{\text{SFT}}$  is a shaped version of the task reward and shares the same optimal policies (Ng et al.,  
 202 1999). This lets us use SFT logits as dense token rewards without explicitly estimating values.

203 **(B) Why we eliminate  $V$  and choose REINFORCE.** For  $\gamma = 1$ , Step 1 guarantees the SFT  $\leftrightarrow$  IQ-  
 204 Learn equivalence; however, Monte-Carlo returns for early tokens are larger in magnitude than for  
 205 later tokens:

$$\sum_{k=t}^{H-1} \log \pi_{\text{SFT}}(a_k | s_k) = \sum_{k=t}^{H-1} r(s_k, a_k) - V_{\text{SFT}}(s_t) \quad (V_{\text{SFT}}(s_H) = 0),$$

209 so returns differ by a state-dependent constant  $-V_{\text{SFT}}(s_t)$ . Fitting a critic (as in PPO) to such  
 210 heteroskedastic targets is difficult, especially if  $V_{\text{SFT}}$  is noisy. Using REINFORCE avoids a critic  
 211 entirely; Appendix A.8 shows that the policy gradient with reward  $\log \pi_{\text{SFT}}$  equals that with reward  $r$   
 212 up to a baseline  $b_t(s_t) = V_{\text{SFT}}(s_t)$ .

213 **(C) A baseline-relative dense reward.** Directly maximizing  $\sum_t \log \pi_{\text{SFT}}(a_t | s_t)$  favors short  
 214 sequences (token log-probabilities are non-positive). We therefore use  
 215

$$\hat{r}(s, a) = \log \pi_{\text{SFT}}(a | s) - \log \pi_{\text{ref}}(a | s), \quad (5)$$

216 Response: Eliza's regular rate is \$10 per hour. She worked f  
 217 or 40 hours at this rate, so she earned  $40 \times 10 = 400$  dollars.  
 218 For the remaining 5 hours, she gets 1.2 times her regular rat  
 219 e. So she earns  $5 \times 1.2 \times 10 = 60$  dollars for overtime. Total e  
 220 arnings are  $400 + 60 = 460$  dollars. The answer is 460.  
 221  
 222

223 Response: Eliza's regular rate is \$10 per hour. She worked f  
 224 or 40 hours at this rate, so she earned  $40 \times 10 = 400$  dollars.  
 225 For the remaining 5 hours, she gets 1.2 times her regular rat  
 226 e. So she earns  $5 \times 2.2 \times 10 = 110$  dollars for overtime. Total  
 227 earnings are  $400 + 110 = 510$  dollars. The answer is 510.  
 228

223 Figure 1: Credit assignment in Dense-Path REINFORCE (Best viewed in color). We provide two  
 224 answers to a math question. The left is the correct response, and on the right is our modified response.  
 225 Each token is colored according to the baseline-relative dense reward as expressed in Eq. (5) (darker  
 226 red means higher reward), using the trained SFT model and SFT checkpoint. We see that the model  
 227 correctly identifies the erroneous number, without much change to the reward value of the other  
 228 tokens, which indicates the ability to do credit assignment.  
 229  
 230

231 where  $\pi_{\text{ref}}$  is a SFT checkpoint with half training samples. This cancels length bias, measures  
 232 incremental competence, and empirically reduces variance. Appendix A.9 bounds the return shift by  
 233  $\|V_{\text{SFT}} - V_{\text{ref}}\|_{\infty}$ .

234 **Illustrative example.** We provide two visualizations in Figure 1 to intuitively demonstrate how  $\hat{r}$   
 235 performs credit assignment at the token level. The reward is calculated by SFT-trained LLaMA-3.1-  
 236 8B and its checkpoint as the baseline. The original question is: “Eliza’s rate per hour for the first 40  
 237 hours she works each week is \$10. She also receives overtime pay at 1.2 times her regular hourly  
 238 rate. If Eliza worked for 45 hours this week, how much are her earnings for this week?” The left side  
 239 shows the correct answer, while the right displays our modified incorrect answer. When calculating  
 240 overtime pay, the incorrect answer erroneously added 1.2 times to the original amount, leading to an  
 241 incorrect result. Analysis reveals that multiplying 5 by 2.2 resulted in a low reward assigned to the  
 242 integer part “2”, indicating the proposed reward can identify this as an erroneous step. Furthermore,  
 243 although the subsequent calculations in the incorrect answer are correct, the final result remains  
 244 wrong, so the assigned reward is lower than that for the correct answer. Additionally, we observe that  
 245 the “5” in the third row receives a relatively high reward. This “5” does not actually appear in the  
 246 original question; it skips a calculation step (“45–40”) to derive overtime hours. Nevertheless,  $\hat{r}$  still  
 accurately identifies this as a valid step.

247 **(D) Safe improvement: transferring proxy gains to true gains.** Let  $\pi'$  be an update that increases  
 248 the proxy return by  $\Delta_{\hat{r}} := J_{\hat{r}}(\pi') - J_{\hat{r}}(\pi) \geq m$ . The performance-difference identity in occupancy  
 249 space gives

$$J_r(\pi') - J_r(\pi) \geq m - 2H \|r - \hat{r}\|_{\infty}, \quad (6)$$

250 since  $\|\rho_{\pi'} - \rho_{\pi}\|_1 \leq 2H$  for a length- $H$  token MDP. See Appendix A.7 for a complete proof.  
 251

252 **Takeaway.** (1)  $\log \pi_{\text{SFT}}$  is a shaped version of the task reward, so it is a valid dense token reward; (2)  
 253 An SFT checkpoint baseline stabilizes learning and removes the EOS pathology; (3) any optimizer  
 254 that increases the proxy return (REINFORCE in our case) *safely* improves the true objective once the  
 255 proxy is accurate enough.  
 256

## 257 5 EXPERIMENTS

### 258 5.1 EXPERIMENTAL SETUP

259 **Data.** We adopt **Open-Orca** and subsample **100k** (prompt, demonstration) pairs for SFT and for  
 260 the RL rollouts (same prompts; no new prompts are introduced in RL). Open-Orca is a large-scale  
 261 open dataset derived from FLAN-style sources augmented with synthetic expert demonstration from  
 262 LLMs (Mukherjee et al., 2023). Using the same pool of prompts ensures the effect of our dense,  
 263 baseline-relative reward does not come from newly introduced prompts.  
 264

265 **Backbones (pretrained only).** To ensure that learning signals from SFT-style demonstrations re-  
 266 main informative, we evaluate only on *foundation (pretrain)* checkpoints (not instruction-tuned). Con-  
 267 cretely, we use four sizes/families representative of current open models: LLaMA-3.1-8B (Dubey  
 268

---

270 **Algorithm 1** Dense-Path REINFORCE

271 **Require:** Expert dataset  $\mathcal{D}_E$ , base model  $\theta_{\text{base}}$ , total SFT steps  $N$ , horizon  $H$ , baseline fraction  
 272  $\alpha \in (0, 1)$  (default 0.5), discount  $\gamma \leftarrow 1$ , KL weight  $\lambda_{\text{KL}} \geq 0$

273 **Ensure:** Fine-tuned policy  $\pi_\phi$

274 1: **SFT stage:** Fine-tune  $\theta_{\text{base}}$  on  $\mathcal{D}_E$  for  $N$  steps; set teacher  $\pi_{\text{SFT}} \leftarrow \pi_{\theta_N}$ . Save the checkpoint  
 275 with half training steps as reference  $\pi_{\text{ref}} \leftarrow \pi_{\theta_{\lfloor \alpha N \rfloor}}$ .

276 2: **Initialize actor:**  $\pi_\phi \leftarrow \pi_{\text{SFT}}$ ; freeze  $\pi_{\text{SFT}}$  and  $\pi_{\text{ref}}$ .

277 3: **for** training iteration = 1, 2, ..., **do**

278 4:   Sample a batch of prompts  $\{x_i\}_{i=1}^B$ ; roll out trajectories  $\tau_i = (s_0, a_0, \dots, s_{T_i-1}, a_{T_i-1})$  using  
 279  $\pi_\phi$ .

280 5:   **for all** tokens  $(s_t, a_t)$  in each  $\tau_i$  **do**

281 6:     **Baseline-relative token reward (Eq. (5)):**  $\hat{r}_t \leftarrow \log \pi_{\text{SFT}}(a_t | s_t) - \log \pi_{\text{ref}}(a_t | s_t)$

282 7:   **end for**

283 8:   **Per-token returns:** For each trajectory  $i$ , compute  $G_t \leftarrow \sum_{k=t}^{T_i-1} \hat{r}_k$  for all  $t$ .

284 9:   **Total objective (token-level):**

285 
$$\mathcal{L}(\phi) = -\frac{1}{B} \sum_{i=1}^B \sum_{t=0}^{T_i-1} \log \pi_\phi(a_t | s_t) G_t$$

286 10:   **Gradient step:** Update  $\phi$  by Adam on  $\nabla_\phi \mathcal{L}(\phi)$ .

287 11: **end for**

---

291  
 292 et al., 2024), Qwen-2.5-7B (Yang et al., 2024), Mistral-7B-v0.1 (Jiang et al., 2024), and  
 293 Gemma-3-4B (Team et al., 2025).

294  
 295 **Baselines.** We compare with: (i) **SFT** (teacher-forced cross-entropy on the 100k set); (ii) **SPIN**  
 296 (self-play fine-tuning from demonstrations) (Chen et al., 2024); (iii) **GSIL** (self-imitation learning  
 297 on demonstrations) (Xiao et al., 2024); and (iv) **SR** (sentence-level REINFORCE): it uses the same  
 298 baseline-relative reward as our method but assigns the *entire trajectory return only at EOS*, i.e., a  
 299 sparse reward delivered once per sequence (conceptually close to PPO-style sparse credit assignment).  
 300 We also test the performance of PPO using sentence-level baseline-relative reward as reward signals,  
 301 but it doesn't show significant differences with REINFORCE. All baselines use the same prompts  
 302 and demonstrations.

303  
 304 **Our method.** We implement the REINFORCE variant described in §4.3 with token-level re-  
 305 turns (undiscounted,  $\gamma = 1$ ), and *baseline-relative dense rewards*  $\hat{r}(s, a) = \log \pi_{\text{SFT}}(a | s) - \log \pi_{\text{ref}}(a | s)$  (SFT checkpoint as  $\pi_{\text{ref}}$ ). We employ a modern RLHF stack based on **Open-**  
 306 **RLHF**'s REINFORCE++ implementation (KL regularization, clipping, and standard stability tricks)  
 307 (Hu et al., 2024; 2025).

308  
 309 **Evaluation.** We use four public instruction-following evaluations: **AlpacaEval**(Li et al., 2023),  
 310 **Arena-Hard** (Li et al.), **LIMA** prompts (Zhou et al., 2023), and **MT-Bench** (standardized 1–10  
 311 scoring) (Zheng et al., 2023). For *AlpacaEval*, *Arena-Hard*, and *LIMA*, we report *pairwise win rate*  
 312 *versus the SFT model* using **GPT-4o** as the judge (temperature 0; ties count as 0.5) (Achiam et al.,  
 313 2023). For *MT-Bench*, we report the standard 1–10 score using the official scripts. Following the  
 314 general test setting for instruction following tasks, decoding uses a temperature 0.7 with a fixed  
 315 max generation length. To minimize tuning bias, **all backbones share the same hyperparameters**  
 316 (Appendix Table 4); this avoids per-model over-tuning.

317  
 318 **5.2 MAIN RESULTS**

319  
 320 **Detailed analysis of Table 1.** (i) **LfD gains across backbones.** Across all four *pretrained* back-  
 321 bones, our token-level method (DPR) improves over the SFT policy on the three win-rate benchmarks  
 322 and MT-Bench scores, confirming that *dense, baseline-relative* rewards extracted from SFT logits  
 323 can further upgrade the policy without introducing new prompts. Typical gains over SFT range from  
 single digits on easier benchmarks to double digits on harder benchmarks (e.g., *Arena-Hard*).

324 **Table 1: Instruction following results across four pretrained backbones.** For **AlpacaEval**, **Arena-**  
 325 **Hard**, and **LIMA**, we report *GPT-4o win rate (%)* versus the SFT model. For **MT-Bench**, we report  
 326 the standard *1–10 score*. All methods train on the same 100k Open-Orca samples. Bold = best,  
 327 underline = second best, per model group.

Method	AlpacaEval	Arena-Hard	LIMA	MT-Bench
	GPT-4o Win Rate (%) ↑			Score (1–10) ↑
<i>LLaMA-3.1-8B</i>				
SFT	-	-	-	5.74
SPIN	55.2	53.3	53.0	5.81
GSIL	<u>58.1</u>	56.7	<u>61.0</u>	5.92
SR	57.9	<u>60.3</u>	60.8	<u>5.96</u>
DPR	<b>60.6</b>	<b>62.5</b>	<b>62.7</b>	<b>6.01</b>
<i>Qwen-2.5-7B</i>				
SFT	-	-	-	6.83
SPIN	55.5	51.4	<u>57.5</u>	6.98
GSIL	<u>56.2</u>	53.3	56.2	7.01
SR	55.9	<u>54.6</u>	54.0	<u>7.09</u>
DPR	<b>57.3</b>	<b>55.2</b>	<b>59.8</b>	<b>7.29</b>
<i>Mistral-v0.1-7B</i>				
SFT	-	-	-	5.23
SPIN	58.3	<u>55.0</u>	53.0	<u>5.45</u>
GSIL	<u>59.2</u>	54.8	<u>54.0</u>	5.43
SR	46.6	49.8	47.3	5.14
DPR	<b>61.0</b>	<b>60.7</b>	<b>59.3</b>	<b>5.65</b>
<i>Gemma-3-4B</i>				
SFT	-	-	-	5.32
SPIN	58.6	54.7	58.7	5.47
GSIL	60.3	57.1	60.8	<b>5.56</b>
SR	<u>65.6</u>	<u>58.0</u>	<u>64.5</u>	5.48
DPR	<b>66.7</b>	<b>58.9</b>	<b>66.8</b>	<u>5.54</u>

357 **(ii) Dense vs. sparse credit assignment.** Relative to *SR* (EOS-only return), DPR achieves systemati-  
 358 cally higher win rates and MT-Bench scores, supporting the hypothesis that *token-level* returns (with  
 359  $\gamma=1$ ) offer better credit assignment than sparse, trajectory-level returns. Notably on *Mistral-v0.1-7B*,  
 360 DPR has a large gap vs. SR on four benchmarks, indicating that per-token shaping is especially  
 361 beneficial when the base model underfits demonstrations.

362 **(iii) LfD baselines (SPIN/GSIL).** Compared with *SPIN* and *GSIL*, both LfD methods that also use  
 363 only demonstrations, DPR is competitive or superior on most benchmarks. The advantage is most  
 364 pronounced on *Arena-Hard*, which is known to better separate models and correlate with Arena  
 365 human preferences. This suggests that our reward extraction provides a stronger, more stable learning  
 366 signal than self-play or self-imitation on the same prompt/demonstration pool.

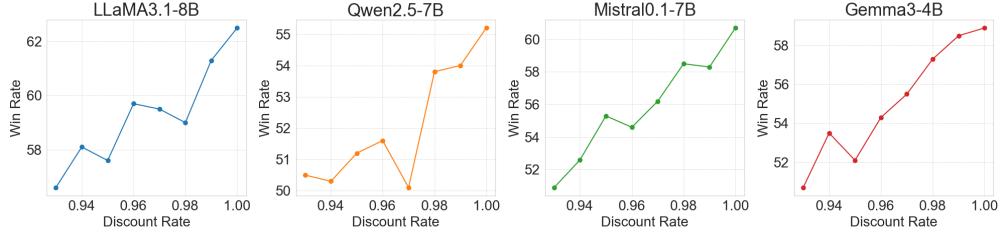
367 **(iv) MT-Bench improvements are consistent though modest.** On MT-Bench (1–10), DPR shows  
 368 small but consistent absolute gains over SFT across backbones (typically +0.2 to +0.5), in line  
 369 with the expectation that general multi-turn quality improves when local token decisions are better  
 370 rewarded.

### 372 5.3 ABLATION STUDY

374 **Findings.** **(a) Effect of eliminating  $V$ .** Compared to *w/DPR*, *w/V* drops on all backbones and  
 375 metrics (typically by 2–7 win-rate points), corroborating our theory that the potential term  $V$  induces  
 376 position-dependent return shifts that are hard to fit and unnecessary under  $\gamma=1$  (cf. §4.3 and Ap-  
 377 pendix A.8). **(b) Necessity of the baseline.** Removing the SFT checkpoint baseline (*wo/Baseline*)  
 causes large drops (often 10–15 win-rate points). This matches the EOS pathology: because token

378  
 379 Table 2: **Ablations on reward shaping and baseline.** *w/DPR*: our full method. *w/V*: do not eliminate  
 380 the potential term  $V$  (i.e., optimize with raw reward  $r(s_t, a_t) = \log \pi_{\text{SFT}}(a_t | s_t) + (V_{\text{SFT}}(s_t) -$   
 381  $V_{\text{SFT}}(s_{t+1}))$ , without using shaping to cancel  $V_{\text{SFT}}(s_t) - V_{\text{SFT}}(s_{t+1})$ . *wo/Baseline*: remove the  
 382 halfway SFT baseline (use only  $\log \pi_{\text{SFT}}$  as reward). Across backbones and benchmarks, *w/V*  
 383 consistently underperforms *w/DPR*, indicating that  $V$  is noisy and its position-dependent returns  
 384 harm stability; *wo/Baseline* degrades substantially, consistent with the EOS pathology and length  
 385 bias discussed in §4.3.

Variant	AlpacaEval $\uparrow$	Arena-Hard $\uparrow$	LIMA $\uparrow$	MT-Bench $\uparrow$
<i>LLaMA-3.1-8B</i>				
w/DPR	60.6	62.5	62.7	6.01
w/V	58.8	59.3	59.7	5.83
wo/Baseline	49.8	46.4	46.0	5.67
<i>Qwen-2.5-7B</i>				
w/DPR	57.3	55.2	59.8	7.29
w/V	55.0	52.9	58.0	7.12
wo/Baseline	46.6	44.9	45.7	6.59
<i>Mistral-7B-v0.1</i>				
w/DPR	61.0	60.7	59.3	5.65
w/V	53.9	51.8	52.3	5.47
wo/Baseline	44.5	40.3	42.7	5.14
<i>Gemma-3-4B</i>				
w/DPR	66.7	58.9	66.8	5.54
w/V	63.5	56.0	62.2	5.51
wo/Baseline	50.6	48.1	48.8	5.26



404  
 405 Figure 2: **The effect of reward discount-rate** ( $\gamma \in \{0.93, 0.94, \dots, 1.00\}$ ) across four backbones.  
 406 Performance (win rate vs. SFT, higher is better) peaks at the *undiscounted* setting  $\gamma=1.0$ . This  
 407 is consistent with our analysis: (i) the SFT $\leftrightarrow$ IQ-Learn equivalence is derived for  $\gamma=1$ ; (ii) with  
 408 discounting, early tokens are over-rewarded relative to later ones, weakening token-level credit  
 409 assignment.  
 410

411 log-probs are non-positive, shorter sequences spuriously obtain larger undiscounted returns without  
 412 the baseline correction; the baseline cancels this length bias and stabilizes updates.

#### 413 5.4 SENSITIVITY ANALYSES

414 **The effect of reward discount rate.** Undiscounted returns preserve the telescoping structure that  
 415 underpins our shaping equivalence and avoid compressing late-token contributions. Empirically, as  
 416 shown in Figure 2, moving from  $\gamma < 1$  to 1.0 improves the win rate consistently across models, with  
 417 larger gains for weaker backbones (e.g., Mistral-7B-v0.1) where late-token guidance matters more.

418 **The effect of baseline checkpoint selection.** As shown in Figure 3, across backbones, the performance  
 419 curve is roughly unimodal with a maximum near the checkpoint with around half of the total  
 420 training samples. This supports the interpretation of our reward as “incremental competence” gained  
 421 during SFT: too early, the baseline is not competitive enough; too late, the gap collapses and the  
 422 proxy reward diminishes.

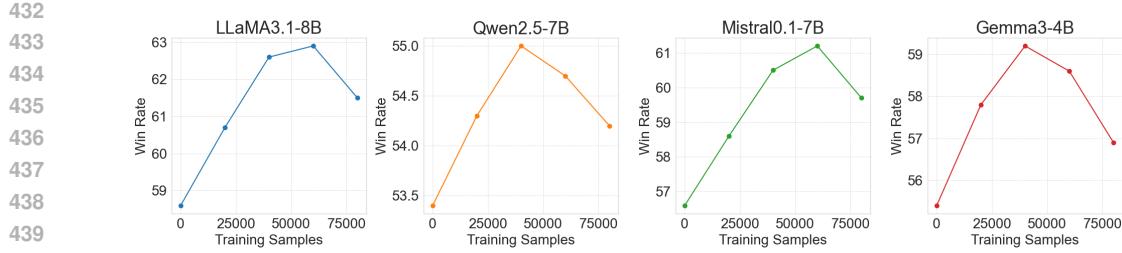
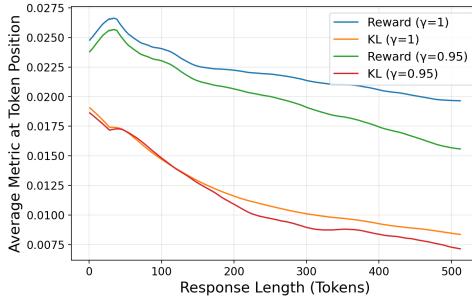
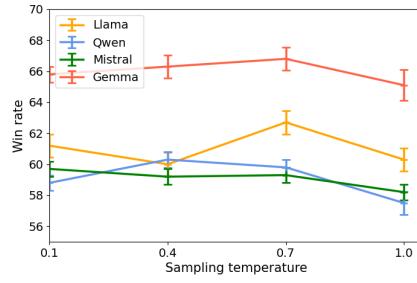


Figure 3: **Baseline checkpoint selection.** We vary the baseline  $\pi_{\text{ref}}$  along the SFT training trajectory (x-axis: SFT progress), keeping all else fixed. A baseline trained with around half of the total training samples yields the best results. Intuitively, an *early* baseline is too weak, over-inflating rewards and increasing variance; a *late* baseline is too close to the full SFT, shrinking  $\log \pi_{\text{SFT}} - \log \pi_{\text{ref}}$  and reducing signal-to-noise. The midpoint balances *magnitude* and *discriminativeness*, consistent with our bound in Appendix A.9.



(a) Visualization of the average KL divergence and reward of responses after DPR training.



(b) We vary the generation temperature of both DPR and the SFT baseline when evaluated on the LIMA benchmark.

**The effect of evaluation temperature.** As depicted in Figure 4b, taking the LIMA benchmark as an example, our algorithm demonstrates significant improvements over SFT across different sampling temperatures during evaluation, indicating its robustness to sampling temperature variations. Furthermore, we observe that although the win rate slightly decreases when the sampling temperature is set to 1, it remains markedly superior to the SFT model. This suggests that our model not only enhances sampling efficiency in high-confidence regions but also achieves notable improvements in other areas.

**Analysis of KL divergence and reward with respect to response length.** Previous studies have found that the majority of the contribution from post-training algorithms might be concentrated in the initial response tokens (Qi et al.). As the response length increases, the contribution of these algorithms may begin to diminish. Correspondingly, in our algorithm, this may be related to the discount rate, as a larger discount rate might exacerbate this phenomenon. To substantiate this observation, we compared the response rewards and KL divergence as a function of length when the discount rate was set to 1 and 0.95. As shown in the Figure 4a, the KL divergence decreases rapidly with increasing length. When the discount rate is 1, the model still retains a high reward within a limited KL budget. However, when the discount rate is 0.95, the model exhibits a more pronounced decline in reward. The results indicate that the phenomenon of rewards decreasing with length does indeed exist, but rewards without discounts can mitigate it to some extent.

## 6 CONCLUSION

This paper revisits LfD for LLMs through the lens of IRL. We show that the token-level SFT objective is *equivalent* to the reduced objective of IQ-Learning. In this view, SFT not only fits a policy but also encodes a dense token-level reward signal in its logits. Building on this equivalence, we propose DPR, a REINFORCE variant that uses dense baseline-relative rewards from the SFT model. Empirically, across four pretrained backbones and four public instruction-following benchmarks, DPR consistently surpasses the SFT baseline and is competitive with other LfD methods.

486 REFERENCES  
487

488 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman,  
489 Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report.  
490 *arXiv preprint arXiv:2303.08774*, 2023.

491 Arash Ahmadian, Chris Cremer, Matthias Gallé, Marzieh Fadaee, Julia Kreutzer, Olivier Pietquin,  
492 Ahmet Üstün, and Sara Hooker. Back to basics: Revisiting reinforce-style optimization for learning  
493 from human feedback in llms. In *Proceedings of the 62nd Annual Meeting of the Association for*  
494 *Computational Linguistics (Volume 1: Long Papers)*, pp. 12248–12267, 2024.

495 Alex J Chan, Hao Sun, Samuel Holt, and Mihaela Van Der Schaar. Dense reward for free in  
496 reinforcement learning from human feedback. In *Proceedings of the 41st International Conference*  
497 *on Machine Learning*, pp. 6136–6154, 2024.

498 Zixiang Chen, Yihe Deng, Huizhuo Yuan, Kaixuan Ji, and Quanquan Gu. Self-play fine-tuning  
499 converts weak language models to strong language models. In *International Conference on*  
500 *Machine Learning*, pp. 6621–6642. PMLR, 2024.

501 Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit  
502 Dhillon, Marcel Blstein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier  
503 with advanced reasoning, multimodality, long context, and next generation agentic capabilities.  
504 *arXiv preprint arXiv:2507.06261*, 2025.

505 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha  
506 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models.  
507 *arXiv e-prints*, pp. arXiv–2407, 2024.

508 Divyansh Garg, Shuvam Chakraborty, Chris Cundy, Jiaming Song, and Stefano Ermon. Iq-learn:  
509 Inverse soft-q learning for imitation. *Advances in Neural Information Processing Systems*, 34:  
510 4028–4039, 2021.

511 Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. Reinforcement learning with  
512 deep energy-based policies. In *International conference on machine learning*, pp. 1352–1361.  
513 PMLR, 2017.

514 Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *Proceedings of the*  
515 *30th International Conference on Neural Information Processing Systems*, pp. 4572–4580, 2016.

516 Jian Hu, Xibin Wu, Wei Shen, Jason Klein Liu, Zilin Zhu, Weixun Wang, Songlin Jiang, Haoran  
517 Wang, Hao Chen, Bin Chen, et al. Openrlhf: An easy-to-use, scalable and high-performance rlhf  
518 framework. *arXiv preprint arXiv:2405.11143*, 2024.

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

521 AQ Jiang, A Sablayrolles, A Mensch, C Bamford, DS Chaplot, Ddl Casas, F Bressand, G Lengyel,  
522 G Lample, L Saulnier, et al. Mistral 7b. arxiv 2023. *arXiv preprint arXiv:2310.06825*, 2024.

523 Jiaxiang Li, Siliang Zeng, Hoi-To Wai, Chenliang Li, Alfredo Garcia, and Mingyi Hong. Getting  
524 more juice out of the sft data: Reward learning from human demonstration improves sft for llm  
525 alignment. *Advances in Neural Information Processing Systems*, 37:124292–124318, 2024.

526 Tianle Li, Wei-Lin Chiang, Evan Frick, Lisa Dunlap, Tianhao Wu, Banghua Zhu, Joseph E Gonzalez,  
527 and Ion Stoica. From crowdsourced data to high-quality benchmarks: Arena-hard and benchbuilder  
528 pipeline. In *Forty-second International Conference on Machine Learning*.

529 Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy  
530 Liang, and Tatsunori B Hashimoto. Alpacaeval: An automatic evaluator of instruction-following  
531 models, 2023.

532 Yi-Chen Li, Tian Xu, Yang Yu, Xuqin Zhang, Xiong-Hui Chen, Zhongxiang Ling, Ningjing Chao,  
533 Lei Yuan, and Zhi-Hua Zhou. Generalist reward models: Found inside large language models.  
534 *arXiv preprint arXiv:2506.23235*, 2025.

540 Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao,  
 541 Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint*  
 542 *arXiv:2412.19437*, 2024.

543 Subhabrata Mukherjee, Arindam Mitra, Ganesh Jawahar, Sahaj Agarwal, Hamid Palangi, and Ahmed  
 544 Awadallah. Orca: Progressive learning from complex explanation traces of gpt-4. *arXiv preprint*  
 545 *arXiv:2306.02707*, 2023.

546 Andrew Y Ng and Stuart J Russell. Algorithms for inverse reinforcement learning. In *Proceedings of the*  
 547 *Seventeenth International Conference on Machine Learning*, pp. 663–670, 2000.

548 Andrew Y Ng, Daishi Harada, and Stuart Russell. Policy invariance under reward transformations:  
 549 Theory and application to reward shaping. In *Proceedings of the Sixteenth International Conference*  
 550 *on Machine Learning*, pp. 278–287, 1999.

551 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong  
 552 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow  
 553 instructions with human feedback. *Advances in neural information processing systems*, 35:27730–  
 554 27744, 2022.

555 Xiangyu Qi, Ashwinee Panda, Kaifeng Lyu, Xiao Ma, Subhrajit Roy, Ahmad Beirami, Prateek Mittal,  
 556 and Peter Henderson. Safety alignment should be made more than just a few tokens deep. In *The*  
 557 *Thirteenth International Conference on Learning Representations*.

558 Chongli Qin and Jost Tobias Springenberg. Supervised fine tuning on curated data is reinforcement  
 559 learning (and can be improved). *arXiv preprint arXiv:2507.12856*, 2025.

560 Rafael Rafailov, Joey Hejna, Ryan Park, and Chelsea Finn. From  $r$  to  $q^*$ : Your language model is  
 561 secretly a  $q$ -function. In *First Conference on Language Modeling*.

562 Rafael Rafailov, Kshitij Sharma, Eric Mitchell, and Chelsea Finn. Direct preference optimization:  
 563 Your language model is secretly a reward model. *arXiv preprint arXiv:2305.18290*, 2023.

564 Omar Shaikh, Michelle S Lam, Joey Hejna, Yijia Shao, Hyundong Justin Cho, Michael S Bernstein,  
 565 and Diyi Yang. Aligning language models with demonstrated feedback. In *The Thirteenth*  
 566 *International Conference on Learning Representations*.

567 Zhihong Shao, Yucheng Guo, Yiduo Zhao, et al. Deepseekmath: Pushing the limits of mathematical  
 568 reasoning in open-source models. *arXiv preprint arXiv:2402.03300*, 2024. Introduces Group  
 569 Relative Policy Optimization (GRPO).

570 Hao Sun. Supervised fine-tuning as inverse reinforcement learning. *arXiv preprint arXiv:2403.12017*,  
 571 2024.

572 Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej,  
 573 Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. Gemma 3 technical  
 574 report. *arXiv preprint arXiv:2503.19786*, 2025.

575 Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement  
 576 learning. *Machine learning*, 8(3):229–256, 1992.

577 Teng Xiao, Yige Yuan, Mingxiao Li, Zhengyu Chen, and Vasant G Honavar. On a connection  
 578 between imitation learning and rlhf. In *The Thirteenth International Conference on Learning*  
 579 *Representations*.

580 Teng Xiao, Mingxiao Li, Yige Yuan, Huaisheng Zhu, Chao Cui, and Vasant G Honavar. How to  
 581 leverage demonstration data in alignment for large language model? a self-imitation learning  
 582 perspective. In *2024 Conference on Empirical Methods in Natural Language Processing, EMNLP*  
 583 2024, pp. 13413–13426. Association for Computational Linguistics (ACL), 2024.

584 An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li,  
 585 Chengyuan Li, Dayiheng Liu, Fei Huang, et al. Qwen2 technical report. *CoRR*, 2024.

594 Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang,  
595 Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and  
596 chatbot arena. *Advances in neural information processing systems*, 36:46595–46623, 2023.  
597

598 Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia  
599 Efrat, Ping Yu, Lili Yu, et al. Lima: Less is more for alignment. *Advances in Neural Information  
600 Processing Systems*, 36:55006–55021, 2023.

601 Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. Maximum entropy inverse  
602 reinforcement learning. In *AAAI*, 2008.

603

604

605

606

607

608

609

610

611

612

613

614

615

616

617

618

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

# 648 649 Appendix

## 650 651 A FULL PROOFS AND TECHNICAL DETAILS

### 652 653 A.1 NOTATION, BASIC ASSUMPTIONS, AND IDENTITIES

654  
655 We work on the finite-horizon token MDP  $(\mathcal{S}, \mathcal{A}, f, \rho_0)$  with deterministic transition  $f(s, a) = s|a$   
656 and horizon  $H$ . A trajectory is  $\tau = (s_0, a_0, \dots, s_H)$  with  $s_{t+1} = f(s_t, a_t)$  and  $s_H$  terminal (EOS or  
657 max length). For any policy  $\pi$ , the *occupancy measure* is

$$658 \quad \rho_\pi(s, a) = \sum_{t=0}^{H-1} \Pr_\pi(s_t = s, a_t = a), \quad \langle \rho_\pi, r \rangle = \sum_{s, a} \rho_\pi(s, a) r(s, a).$$

659 For a function  $Q : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ , define the log-partition (soft value) and Boltzmann policy

$$660 \quad V(s) = \beta \log \sum_{a \in \mathcal{A}} e^{Q(s, a) / \beta}, \quad \pi_Q(a | s) = \exp\left(\frac{1}{\beta}(Q(s, a) - V(s))\right),$$

661 with fixed temperature  $\beta > 0$  (we use  $\beta = 1$  when not stated). We frequently use the identity

$$662 \quad \log \pi_Q(a | s) = \frac{1}{\beta}(Q(s, a) - V(s)). \quad (7)$$

### 663 A.2 DERIVATION OF EQ. (1) (OPTIMAL SOFT POLICY AND VALUE)

664 **Setup.** Fix a state  $s$ . Consider the convex optimization problem

$$665 \quad \max_{\pi(\cdot | s) \in \Delta(\mathcal{A})} \sum_a \pi(a | s) Q^*(s, a) - \beta \sum_a \pi(a | s) \log \pi(a | s),$$

666 subject to (i)  $\sum_a \pi(a | s) = 1$ , (ii)  $\pi(a | s) \geq 0$  for all  $a$ . The objective is strictly concave in  $\pi(\cdot | s)$   
667 because the negative entropy  $-\sum \pi \log \pi$  is strictly convex and we *maximize* its negation; hence the  
668 maximizer is unique.

669 **KKT conditions.** Form the Lagrangian

$$670 \quad \mathcal{L}(\pi, \lambda, \{\nu_a\}) = \sum_a \pi(a | s) Q^*(s, a) - \beta \sum_a \pi(a | s) \log \pi(a | s) + \lambda \left( \sum_a \pi(a | s) - 1 \right) + \sum_a \nu_a \pi(a | s),$$

671 with multipliers  $\lambda \in \mathbb{R}$  for the simplex constraint and  $\nu_a \geq 0$  for non-negativity. Stationarity for  
672 every  $a$  gives

$$673 \quad \frac{\partial \mathcal{L}}{\partial \pi(a | s)} = Q^*(s, a) - \beta(1 + \log \pi(a | s)) + \lambda + \nu_a = 0.$$

674 *Complementary slackness:* if  $\pi^*(a | s) > 0$ , then  $\nu_a = 0$ . Since the optimum has full support under  
675 finite  $\beta > 0$  (the entropy term forces interior optimum), we set  $\nu_a = 0$  for all  $a$  and obtain

$$676 \quad \log \pi^*(a | s) = \frac{1}{\beta}(Q^*(s, a) + \lambda - \beta).$$

677 Exponentiating and normalizing by the constraint yields

$$678 \quad \pi^*(a | s) = \frac{\exp(Q^*(s, a) / \beta)}{\sum_{a'} \exp(Q^*(s, a') / \beta)}.$$

679 Defining  $V^*(s) := \beta \log \sum_{a'} \exp(Q^*(s, a') / \beta)$  gives the stated softmax policy and the value expres-  
680 sion in Eq. (1). This completes the derivation.

### 681 A.3 FROM MAXENT IRL TO THE IQ-LEARN REDUCED OBJECTIVE $J^*(Q)$

682 In this section, we give a minimal proof modified from IQ-Learn (Garg et al., 2021). We recall the  
683 MaxEnt IRL saddle objective

$$684 \quad L(\pi, r) = \langle \rho_E - \rho_\pi, r \rangle - H(\pi) - \psi(r), \quad (8)$$

702 with a convex reward regularizer  $\psi$  for identifiability/stability (Ziebart et al., 2008; Ho & Ermon,  
703 2016). For a fixed  $Q$ , minimizing  $L$  over  $\pi$  with the soft entropy yields the Boltzmann policy  $\pi_Q$  in  
704 (7); the corresponding *reduced* objective over  $Q$  (IQ-Learn) is

$$705 \quad 706 \quad J^*(Q) = \mathbb{E}_{(s,a) \sim \rho_E} [Q(s,a) - V(f(s,a))] - \mathbb{E}_{s_0 \sim \rho_0} [V(s_0)], \quad (9)$$

707 where  $V$  is the log-partition induced by  $Q$  and  $f$  is the deterministic environment transition. For  
708 completeness, we expand all steps below.

710 **Detailed derivation.** Write the inner minimization over  $\pi$  at each state  $s$ :

$$711 \quad \min_{\pi(\cdot|s) \in \Delta} \left\{ -\sum_a \pi(a|s) Q(s,a) + \beta \sum_a \pi(a|s) \log \pi(a|s) \right\} = -\max_{\pi(\cdot|s)} \left\{ \sum_a \pi(a|s) Q(s,a) - \beta H(\pi(\cdot|s)) \right\}.$$

714 By Sec. A.2, the maximizer is  $\pi_Q(\cdot|s)$  and the maximized value equals the log-partition  $V(s)$ :

$$715 \quad \max_{\pi(\cdot|s)} \left\{ \sum_a \pi(a|s) Q(s,a) - \beta H(\pi(\cdot|s)) \right\} = V(s).$$

718 Plugging back into (8) and unrolling the entropy term over time yields

$$719 \quad \min_{\pi} L(\pi, r) = \langle \rho_E - \rho_{\pi_Q}, r \rangle - \sum_{t=0}^{H-1} \mathbb{E}_{s_t} [V(s_t)] - \psi(r).$$

722 In IQ-Learn we eliminate  $r$  in favor of  $Q$  using the soft Bellman identity (see next subsection): for  
723  $\gamma = 1$  and deterministic  $f$ ,  $Q(s_t, a_t) = r(s_t, a_t) + V(s_{t+1})$  and hence

$$725 \quad \langle \rho_E, r \rangle = \mathbb{E}_{(s,a) \sim \rho_E} [Q(s,a) - V(f(s,a))].$$

726 The  $\rho_{\pi_Q}$ -term cancels at the saddle (where  $\rho_{\pi^*} = \rho_E$ ), and the initial-state entropy contributes  
727  $-\mathbb{E}_{s_0 \sim \rho_0} [V(s_0)]$ , leading exactly to (9).

#### 729 A.4 PROOF OF PROP. 1: SFT IS EQUIVALENT TO MAXIMIZING $J^*(Q)$

731 We now show that, on the LLM environment with  $\gamma = 1$  and linear conjugate (no extra reward  
732 regularization beyond convexity), maximizing  $J^*(Q)$  equals maximizing the SFT log-likelihood.  
733 Starting from (9),

$$734 \quad \sum_{t=0}^{H-1} (Q(s_t, a_t) - V(f(s_t, a_t))) = \sum_{t=0}^{H-1} (Q(s_t, a_t) - V(s_{t+1})).$$

737 Add and subtract  $V(s_t)$  termwise, then regroup:

$$739 \quad \sum_{t=0}^{H-1} (Q(s_t, a_t) - V(s_t)) + \sum_{t=0}^{H-1} (V(s_t) - V(s_{t+1})) = \sum_{t=0}^{H-1} \log \pi_Q(a_t | s_t) + V(s_0) - V(s_H),$$

742 where we used (7). With the terminal state  $V(s_H) = 0$ , take expectation over expert trajectories and  
743 subtract  $\mathbb{E}[V(s_0)]$  (the last term of (9)) to obtain

$$744 \quad 745 \quad J^*(Q) = \mathbb{E}_{\tau \sim \rho_E} \left[ \sum_{t=0}^{H-1} \log \pi_Q(a_t | s_t) \right].$$

747 Maximizing the objective of IQ-Learn is exactly maximizing the teacher-forced log-likelihood of  
748 expert tokens, i.e., minimizing the token-level SFT cross-entropy. This proves the proposition.

#### 750 A.5 DERIVATION OF EQ. (4): SFT LOGITS AS A SHAPED REWARD

752 We derive the identity used in §3 (Eq. (4)):

$$754 \quad \log \pi_{\text{SFT}}(a_t | s_t) = r(s_t, a_t) + V(s_{t+1}) - V(s_t)$$

755 under the soft-control model with  $\gamma = 1$  and deterministic transition  $s_{t+1} = f(s_t, a_t)$ .

756 **Soft Bellman equations (finite horizon).** For any  $(s_t, a_t)$ ,

$$758 \quad Q(s_t, a_t) = r(s_t, a_t) + V(s_{t+1}), \quad V(s_t) = \beta \log \sum_a \exp\left(\frac{1}{\beta} Q(s_t, a)\right).$$

759 Subtract  $V(s_t)$  from both sides of the first equation and divide by  $\beta$ :

$$761 \quad \frac{1}{\beta}(Q(s_t, a_t) - V(s_t)) = \frac{1}{\beta}r(s_t, a_t) + \frac{1}{\beta}(V(s_{t+1}) - V(s_t)).$$

762 Using (7) on the left gives exactly Eq. (4) (with  $\beta = 1$ ). No approximation is used.

764 **Telescoping of returns and why we remove  $V$ .** Summing Eq. (4) from  $t$  to  $H - 1$  (with  $V(s_H) = 0$ ),

$$766 \quad 767 \quad \sum_{k=t}^{H-1} \log \pi_{\text{SFT}}(a_k | s_k) = \sum_{k=t}^{H-1} r(s_k, a_k) - V(s_t).$$

769 Thus under  $\gamma = 1$ , log-prob returns differ from true returns by a *state-dependent constant*  $-V(s_t)$ .  
770 This constant shift (i) proves that using  $\log \pi_{\text{SFT}}$  as reward yields the *same* policy gradient as using  $r$   
771 (Sec. A.8), and (ii) motivates *eliminating  $V$*  via potential-based shaping to reduce variance and length  
772 bias (Sec. A.9).

#### 773 A.6 DUAL CONTRACTION: REWARD ERROR IS BOUNDED BY POLICY (OCCUPANCY) ERROR

775 We restate the IRL objective (8) and define the reward best response  $\hat{r}(\pi) = \arg \max_r L(\pi, r)$ . Let  
776  $r^*$  be any reward at the IRL saddle (unique up to shaping). Assume  $\psi$  is  $\mu$ -strongly convex in norm  
777  $\|\cdot\|$ . We prove

$$778 \quad \|\hat{r}(\pi) - r^*\| \leq \frac{1}{\mu} \|\rho_\pi - \rho_E\|_*$$

780 where  $\|\cdot\|_*$  is the dual norm to  $\|\cdot\|$ .

782 **First-order conditions and strong monotonicity.** Optimality of the reward player yields

$$783 \quad \nabla \psi(\hat{r}(\pi)) = \rho_E - \rho_\pi, \quad \nabla \psi(r^*) = \rho_E - \rho_{\pi^*}.$$

784 At the saddle  $\rho_{\pi^*} = \rho_E$ , so  $\nabla \psi(r^*) = 0$  and hence

$$785 \quad \nabla \psi(\hat{r}(\pi)) - \nabla \psi(r^*) = \rho_E - \rho_\pi.$$

786 By strong convexity,  $\nabla \psi$  is  $\mu$ -strongly *monotone*:

$$787 \quad \langle \hat{r}(\pi) - r^*, \nabla \psi(\hat{r}(\pi)) - \nabla \psi(r^*) \rangle \geq \mu \|\hat{r}(\pi) - r^*\|^2.$$

788 Combine the last two displays and apply Hölder's inequality in dual norms:

$$789 \quad \mu \|\hat{r}(\pi) - r^*\|^2 \leq \langle \hat{r}(\pi) - r^*, \rho_E - \rho_\pi \rangle \leq \|\hat{r}(\pi) - r^*\| \|\rho_\pi - \rho_E\|_*.$$

790 If  $\hat{r}(\pi) \neq r^*$ , divide both sides by  $\|\hat{r}(\pi) - r^*\|$ ; otherwise the bound is trivial. This proves the claim.

#### 793 A.7 SAFE IMPROVEMENT UNDER A PROXY REWARD (FULL PROOF OF EQ. (6))

795 Let  $J_r(\pi) := \langle \rho_\pi, r \rangle$  be the return under reward  $r$ . For any rewards  $r, \hat{r}$  and policies  $\pi, \pi'$ ,

$$796 \quad J_r(\pi') - J_r(\pi) = \langle \rho_{\pi'} - \rho_\pi, r \rangle = \langle \rho_{\pi'} - \rho_\pi, \hat{r} \rangle + \langle \rho_{\pi'} - \rho_\pi, r - \hat{r} \rangle.$$

797 The first term equals  $J_{\hat{r}}(\pi') - J_{\hat{r}}(\pi)$ . For the second term, apply Hölder with  $\ell_1/\ell_\infty$  duality:

$$798 \quad |\langle \rho_{\pi'} - \rho_\pi, r - \hat{r} \rangle| \leq \|\rho_{\pi'} - \rho_\pi\|_1 \|r - \hat{r}\|_\infty.$$

799 It remains to upper bound  $\|\rho_{\pi'} - \rho_\pi\|_1$ . Writing  $p_t^\pi(s, a) = \Pr_\pi(s_t = s, a_t = a)$ ,

$$800 \quad \begin{aligned} 801 \quad \|\rho_{\pi'} - \rho_\pi\|_1 &= \sum_{s,a} \left| \sum_{t=0}^{H-1} (p_t^{\pi'}(s, a) - p_t^\pi(s, a)) \right| \leq \sum_{t=0}^{H-1} \sum_{s,a} |p_t^{\pi'}(s, a) - p_t^\pi(s, a)| \\ 802 \quad &= \sum_{t=0}^{H-1} \|p_t^{\pi'} - p_t^\pi\|_{\text{TV}} \cdot 2 \leq 2H, \end{aligned}$$

803 since each  $p_t^\pi$  is a probability distribution over  $(s, a)$  (total variation  $\leq 2$ ). Therefore

$$804 \quad J_r(\pi') - J_r(\pi) \geq (J_{\hat{r}}(\pi') - J_{\hat{r}}(\pi)) - 2H \|r - \hat{r}\|_\infty.$$

805 Setting  $m := J_{\hat{r}}(\pi') - J_{\hat{r}}(\pi)$  gives Eq. (6).

810 A.8 POLICY-GRADIENT EQUIVALENCE UNDER  $\gamma = 1$  (REINFORCE BASELINE IDENTITY)  
811812 Let  $r_t := r(s_t, a_t)$  and define the shaped reward  $\tilde{r}_t := \log \pi_{\text{SFT}}(a_t | s_t) = r_t + (V_{t+1} - V_t)$  with  
813  $V_H = 0$ . Define returns from step  $t$ :

814  
815 
$$G_t = \sum_{k=t}^{H-1} r_k, \quad \tilde{G}_t = \sum_{k=t}^{H-1} \tilde{r}_k = G_t - V_t.$$
  
816

817 The REINFORCE gradients are  
818

819  
820 
$$\nabla J_r(\pi) = \mathbb{E} \left[ \sum_{t=0}^{H-1} \nabla \log \pi(a_t | s_t) G_t \right], \quad \nabla J_{\tilde{r}}(\pi) = \mathbb{E} \left[ \sum_{t=0}^{H-1} \nabla \log \pi(a_t | s_t) \tilde{G}_t \right].$$
  
821

822 For any function  $b_t(s_t)$ , using the law of iterated expectations and the identity  $\mathbb{E}_{a \sim \pi(\cdot | s)} [\nabla \log \pi(a | s)] = \nabla \sum_a \pi(a | s) = 0$ , we have  
823

824  
825 
$$\mathbb{E} [\nabla \log \pi(a_t | s_t) b_t(s_t)] = 0.$$

826 Choosing  $b_t = V_t$  yields  $\nabla J_{\tilde{r}}(\pi) = \nabla J_r(\pi)$ . Thus the policy gradient under  $\log \pi_{\text{SFT}}$  equals that  
827 under  $r$ , up to a *state-only* baseline that does not require fitting a critic.  
828829 A.9 CHECKPOINT BASELINE TIGHTNESS AND DYNAMIC RANGE REDUCTION  
830831 Consider the baseline-relative reward  
832

833 
$$\hat{r}(s, a) = \log \pi_{\text{SFT}}(a | s) - \log \pi_{\text{ref}}(a | s), \quad \hat{V}(s) := V_{\text{SFT}}(s) - V_{\text{ref}}(s).$$
  
834

By Sec. A.5, the corresponding return from step  $t$  differs by  $-\hat{V}(s_t)$ . Hence for any trajectory and  $t$ ,

835  
836 
$$|\tilde{G}_t^{\text{SFT}} - \tilde{G}_t^{\text{ref}}| = |\hat{V}(s_t)| \leq \|\hat{V}\|_{\infty}.$$

If  $\pi_{\text{ref}}$  is an SFT checkpoint, empirically  $\|\hat{V}\|_{\infty}$  is small because the two values remain close along the  
837 training path. The dynamic range of token returns is thus reduced by at least  $\|V_{\text{SFT}}\|_{\infty} - \|\hat{V}\|_{\infty}$ ,  
838 stabilizing updates and mitigating EOS/length bias (see also the toy pathology in App. A.11).  
839840 A.10 POTENTIAL-BASED SHAPING INVARIANCE (FINITE-HORIZON, DETERMINISTIC  
841 ENVIRONMENT)  
842843 Define a shaped reward  $r^F(s, a) = r(s, a) + F(s') - F(s)$  with  $s' = f(s, a)$  and any  $F : \mathcal{S} \rightarrow \mathbb{R}$ .  
844 Consider the soft  $Q$ -values for  $\gamma = 1$ :

845  
846 
$$Q^F(s, a) = r^F(s, a) + V^F(s') = r(s, a) + \underbrace{F(s') - F(s)}_{\text{shaping}} + V^F(s').$$
  
847

848 Define  $\tilde{V}(s) := V^F(s) + F(s)$ . Then  
849

850  
851 
$$Q^F(s, a) - \tilde{V}(s) = r(s, a) + V^F(s') - V^F(s) = Q(s, a) - V(s),$$

852 where the last equality follows because the soft Bellman backup  $V(\cdot) = \log \sum_a e^{Q(\cdot, a)}$  is invariant  
853 to adding the same  $F$  to all action-logs at a state. Therefore, by (7),  
854

855 
$$\pi_{Q^F}(a | s) = \pi_Q(a | s) \quad \text{for all } (s, a).$$

856 Thus potential-based shaping preserves the optimal policy and all on-policy distributions (Ng et al.,  
857 1999).  
858A.11 EOS/LENGTH PATHOLOGY WITHOUT A BASELINE: A TOY PROOF  
859860 Assume at each nonterminal state  $s$  there are actions  $\{\text{EOS}\} \cup \mathcal{A}_{\text{cont}}$  and consider the proxy objective  
861 without baseline:  
862

863 
$$J_{\text{naive}}(\pi) = \mathbb{E}_{\pi} \left[ \sum_{t=0}^{T-1} \log \pi_{\text{SFT}}(a_t | s_t) \right], \quad T = (\text{random stopping time at EOS}).$$

864 Suppose (mild) that  $\log \pi_{\text{SFT}}(\text{EOS} \mid s) \geq \max_{a \in \mathcal{A}_{\text{cont}}} \log \pi_{\text{SFT}}(a \mid s)$  for all  $s$  in a subset of  
 865 high measure under  $\pi$ . Then any deviation that delays EOS will, in expectation, *decrease* the sum  
 866 of log-probs (since each additional token contributes a non-positive term no larger than the EOS  
 867 log-prob). Therefore maximizing  $J_{\text{naive}}$  prefers immediate EOS whenever it is locally the highest-  
 868 probability token; this formalizes the “short-output bias” and motivates the baseline subtraction  
 869  $\log \pi_{\text{SFT}} - \log \pi_{\text{ref}}$ .  
 870

## 871 B ADDITIONAL EXPERIMENTAL DETAILS

873 Table 3: Hyperparameters used across all backbones for SFT.  
 874

875 Component	876 Value	877 Component	878 Value
877 Learning rate	878 5e-6	879 Global Batch size	256
878 Max prompt length	1024	879 Max gen length	1024
879 Warmup ratio	0.03	880 Optimizer	Adam

881 Table 4: Hyperparameters used across all backbones for DPR.  
 882

883 Component	884 Value	885 Component	886 Value
885 Learning rate	886 5e-7	887 Global Batch size	128
886 Max prompt length	1024	888 Max gen length	1024
887 KL weight	1e-5	889 Warmup ratio	0.03
888 Reward discount rate	1	890 rollout temperature	1
889 Rollout Batch Size	1024	891 Value clip	0.2
890 Samples per prompt	1	892 Optimizer	Adam

## 893 C THE USE OF LARGE LANGUAGE MODELS

894 We employed LLM to assist with paper writing, primarily for vocabulary and grammar checks, while  
 895 utilizing Copilot for code completion in writing research code. All text or code generated by LLM or  
 896 Copilot undergoes secondary verification or unit testing by authors to ensure accuracy. We affirm that  
 897 the LLM did not participate in any research sections beyond writing and coding assistance.  
 898