

Controllable Generation

Recap: KL-reg RL objective (traj-wise)

$$J(\pi) = \mathbb{E}_{x \sim \nu} \left[\mathbb{E}_{\tau \sim \pi(\cdot | x)} r(x, \tau) - \beta \text{KL} \left(\pi(\cdot | x) \middle| \pi_{ref}(\cdot | x) \right) \right]$$

$$\hat{\pi}(\tau | x) \propto \pi_{ref}(\tau | x) \cdot \exp \left(\frac{r(x, \tau)}{\beta} \right)$$

Stay close to π_{ref}

Optimize reward

Recap: DPO and REBEL

DPO:

$$\arg \max_{\theta} \sum_{x, \tau, \tau', z} \ln \frac{1}{1 + \exp \left(-z \cdot \beta \left(\ln \frac{\pi_{\theta}(\tau|x)}{\pi_{ref}(\tau|x)} - \ln \frac{\pi_{\theta}(\tau'|x)}{\pi_{ref}(\tau'|x)} \right) \right)}$$

Recap: DPO and REBEL

DPO:

$$\arg \max_{\theta} \sum_{x, \tau, \tau', z} \ln \frac{1}{1 + \exp \left(-z \cdot \beta \left(\ln \frac{\pi_{\theta}(\tau | x)}{\pi_{ref}(\tau | x)} - \ln \frac{\pi_{\theta}(\tau' | x)}{\pi_{ref}(\tau' | x)} \right) \right)}$$

REBEL:

$$\theta_{t+1} = \arg \min_{\theta} \mathbb{E}_{x, (\tau, \tau') \sim \pi_{\theta_t}(\cdot | x)} \left(\beta \left(\ln \frac{\pi_{\theta}(\tau | x)}{\pi_{\theta_t}(\tau | x)} - \ln \frac{\pi_{\theta}(\tau' | x)}{\pi_{\theta_t}(\tau' | x)} \right) - \underbrace{(r(x, \tau) - r(x, \tau'))}_{2} \right)$$

One more RL algorithm: GRPO (Deepseek-R1)

Basically some combination PPO clipping with RLoO (Reinforce w/ leave-one-out)

Given π_t , it updates policy to π_{t+1} as follows:

One more RL algorithm: GRPO (Deepseek-R1)

Basically some combination PPO clipping with RLoO (Reinforce w/ leave-one-out)

Given π_t , it updates policy to π_{t+1} as follows:

$$\tau^i \sim \pi_x(\cdot | x)$$

1. Sample a bunch of prompts, for each x , generate k i.i.d responses $\tau^1, \tau^2, \dots, \tau^k$

One more RL algorithm: GRPO (Deepseek-R1)

Basically some combination PPO clipping with RLoO (Reinforce w/ leave-one-out)

Given π_t , it updates policy to π_{t+1} as follows:

1. Sample a bunch of prompts, for each x , generate k i.i.d responses $\tau^1, \tau^2, \dots, \tau^k$
2. Form the following clipping-based objective

One more RL algorithm: GRPO (Deepseek-R1)

Basically some combination PPO clipping with RLoO (Reinforce w/ leave-one-out)

Given π_t , it updates policy to π_{t+1} as follows:

1. Sample a bunch of prompts, for each x , generate k i.i.d responses $\tau^1, \tau^2, \dots, \tau^k$
2. Form the following clipping-based objective

$$\max_{\pi} \underbrace{\sum_{\{x, \tau^1, \tau^2, \dots, \tau^k\}}}_{\Delta} \underbrace{\sum_{i=1}^k}_{\Delta} \min \left\{ \underbrace{\frac{\pi(\tau^i | x)}{\pi_t(\tau^i | x)} A(x, \tau^i)}_{\Delta}, \underbrace{\text{clip} \left(\frac{\pi(\tau^i | x)}{\pi_t(\tau^i | x)}, 1 - \epsilon, 1 + \epsilon \right) A(x, \tau^i)}_{\Delta} \right\}$$

One more RL algorithm: GRPO (Deepseek-R1)

Basically some combination PPO clipping with RLoO (Reinforce w/ leave-one-out)

Given π_t , it updates policy to π_{t+1} as follows:

1. Sample a bunch of prompts, for each x , generate k i.i.d responses $\tau^1, \tau^2, \dots, \tau^k$
2. Form the following clipping-based objective

$$\max_{\pi} \sum_{\{x, \tau^1, \tau^2, \dots, \tau^k\}} \sum_{i=1}^k \min \left\{ \frac{\pi(\tau^i | x)}{\pi_t(\tau^i | x)} A(x, \tau^i), \text{clip} \left(\frac{\pi(\tau^i | x)}{\pi_t(\tau^i | x)}, 1 - \epsilon, 1 + \epsilon \right) A(x, \tau^i) \right\}$$

where:

$$A(x, \tau^i) = \frac{r(\tau^i) - \bar{r}}{\text{std}(r(\tau^1), r(\tau^2), \dots, r(\tau^k))}$$

$$\bar{r} = \frac{1}{k} \sum_{i=1}^k r(x, \tau^i)$$

One more RL algorithm: GRPO (Deepseek-R1)

Basically some combination PPO clipping with RLoO (Reinforce w/ leave-one-out)

Given π_t , it updates policy to π_{t+1} as follows:

1. Sample a bunch of prompts, for each x , generate k i.i.d responses $\tau^1, \tau^2, \dots, \tau^k$
2. Form the following clipping-based objective

$$\max_{\pi} \sum_{\{x, \tau^1, \tau^2, \dots, \tau^k\}} \sum_{i=1}^k \min \left\{ \frac{\pi(\tau^i | x)}{\pi_t(\tau^i | x)} A(x, \tau^i), \text{clip} \left(\frac{\pi(\tau^i | x)}{\pi_t(\tau^i | x)}, 1 - \epsilon, 1 + \epsilon \right) A(x, \tau^i) \right\}$$

where:

$$A(x, \tau^i) = \frac{r(\tau^i) - \bar{r}}{\text{std}(r(\tau^1), r(\tau^2), \dots, r(\tau^k))}$$

Normalize advantage use group responses τ^1, \dots, τ^k , per prompt;

Today's question

So far, DPO, PPO, REBEL, and GRPO all optimize the entire LLM; when LLM is large (e.g., > 70B), we cannot afford to do full parameter optimization...

Today's question

So far, DPO, PPO, REBEL, and GRPO all optimize the entire LLM; when LLM is large (e.g., > 70B), we cannot afford to do full parameter optimization...

Q: can we train small evaluation model (e.g., 3B) to **guide the generation** of a big large black-box model (e.g., 70B)?

Outline

1. KL regularized RL again, but in token space (i.e., s_h, a_h) not traj space
2. Train value/Q functions
3. Controllable generation via guidance from Q/V functions

Notation

Finite horizon MDP with deterministic transition, i.e., $s_{h+1} = f(s_h, a_h)$

Notation

Finite horizon MDP with deterministic transition, i.e., $s_{h+1} = f(s_h, a_h)$

π_{ref} — a black-box large model, don't want to do full backpropagation on it...

Notation

Finite horizon MDP with deterministic transition, i.e., $s_{h+1} = f(s_h, a_h)$

π_{ref} — a black-box large model, don't want to do full backpropagation on it...

We want to optimize KL-regularized RL objective:

Notation

Finite horizon MDP with deterministic transition, i.e., $s_{h+1} = f(s_h, a_h)$

π_{ref} — a black-box large model, don't want to do full backpropagation on it...

We want to optimize KL-regularized RL objective:

$$\max_{\pi} \mathbb{E}_{\tau \sim \pi} \left[\sum_{h=0}^{H-1} r(s_h, a_h) - \beta \text{KL} \left(\pi(\cdot | s_h) \parallel \pi_{ref}(\cdot | s_h) \right) \right]$$

Notation

Finite horizon MDP with deterministic transition, i.e., $s_{h+1} = f(s_h, a_h)$

π_{ref} — a black-box large model, don't want to do full backpropagation on it...

We want to optimize KL-regularized RL objective:

$$\max_{\pi} \mathbb{E}_{\tau \sim \pi} \left[\sum_{h=0}^{H-1} r(s_h, a_h) - \beta \text{KL} \left(\pi(\cdot | s_h) | \pi_{ref}(\cdot | s_h) \right) \right]$$

Let's solve this via Dynamic Programming (backward in time)

DP for solving the KL-regularized RL

$$\max_{\pi} \mathbb{E}_{\tau \sim \pi} \left[\sum_{h=0}^{H-1} r(s_h, a_h) - \beta \text{KL} \left(\pi(\cdot | s_h) | \pi_{ref}(\cdot | s_h) \right) \right]$$

DP for solving the KL-regularized RL

H
 $h=0 \dots \underline{H-1}$

$$\max_{\pi} \mathbb{E}_{\tau \sim \pi} \left[\sum_{h=0}^{H-1} r(s_h, a_h) - \beta \text{KL} \left(\pi(\cdot | s_h) | \pi_{ref}(\cdot | s_h) \right) \right]$$

Base case: $V^*(s_H) = 0$, for the fictitious step H

DP for solving the KL-regularized RL

$$\max_{\pi} \mathbb{E}_{\tau \sim \pi} \left[\sum_{h=0}^{H-1} r(s_h, a_h) - \beta \text{KL} \left(\pi(\cdot | s_h) | \pi_{ref}(\cdot | s_h) \right) \right]$$

Base case: $V^*(s_H) = 0$, for the fictitious step H

Induction step: given $V^*(s_{h+1})$, want to compute $V^*(s_h)$

DP for solving the KL-regularized RL

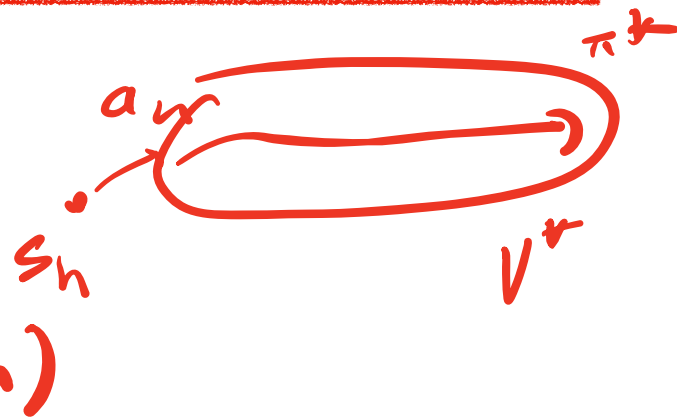
$$\max_{\pi} \mathbb{E}_{\tau \sim \pi} \left[\sum_{h=0}^{H-1} r(s_h, a_h) - \beta \text{KL} \left(\pi(\cdot | s_h) | \pi_{\text{ref}}(\cdot | s_h) \right) \right]$$

Base case: $V^*(s_H) = 0$, for the fictitious step H

Induction step: given $V^*(s_{h+1})$, want to compute $V^*(s_h)$

$$Q^*(s_h, a_h) = r(s_h, a_h) + \mathbb{E}_{s_{h+1} \sim P(\cdot | s_h, a_h)} V^*(s_{h+1})$$

$\rightsquigarrow V^*(s_h)$



DP for solving the KL-regularized RL

$$\max_{\pi} \mathbb{E}_{\tau \sim \pi} \left[\sum_{h=0}^{H-1} r(s_h, a_h) - \beta \text{KL} \left(\pi(\cdot | s_h) | \pi_{\text{ref}}(\cdot | s_h) \right) \right]$$

Base case: $V^*(s_H) = 0$, for the fictitious step H

$$V^*(s) = \max_a Q^*(s, a)$$

Induction step: given $V^*(s_{h+1})$, want to compute $V^*(s_h)$

$$Q^*(s_h, a_h) = r(s_h, a_h) + \mathbb{E}_{s_{h+1} \sim P(\cdot | s_h, a_h)} V^*(s_{h+1})$$

$$\pi^*(a|s_h) \propto \pi_{\text{ref}}(a|s_h) \exp\left(\frac{Q^*(s_h, a)}{\beta}\right)$$

$$\propto \pi_{\text{ref}}(a|s_h) \exp\left(\frac{Q^*(s_h, a)}{\beta}\right)$$

$$\beta \rightarrow 0$$

$$V^*(s_h) = \max_{\pi(\cdot | s_h) \in \Delta(A)} \left[\mathbb{E}_{a \sim \pi(\cdot | s_h)} Q^*(s_h, a) - \beta \text{KL} \left(\pi(\cdot | s_h) | \pi_{\text{ref}}(\cdot | s_h) \right) \right]$$

DP for solving the KL-regularized RL

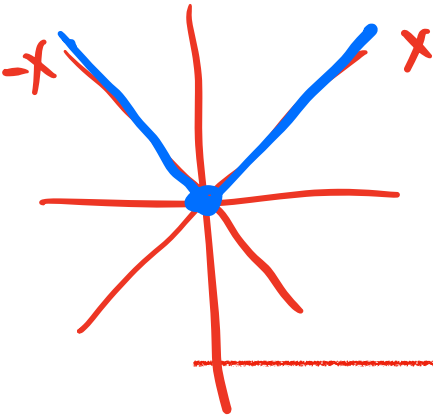
$$\max_{\pi} \mathbb{E}_{\tau \sim \pi} \left[\sum_{h=0}^{H-1} r(s_h, a_h) - \beta \text{KL} \left(\pi(\cdot | s_h) | \pi_{ref}(\cdot | s_h) \right) \right]$$

Base case: $V^*(s_H) = 0$, for the fictitious step H

Induction step: given $V^*(s_{h+1})$, want to compute $V^*(s_h)$

$$Q^*(s_h, a_h) = r(s_h, a_h) + \mathbb{E}_{s_{h+1} \sim P(\cdot | s_h, a_h)} V^*(s_{h+1})$$

$$V^*(s_h) = \max_{\pi(\cdot | s_h) \in \Delta(A)} \mathbb{E}_{a \sim \pi(\cdot | s_h)} Q^*(s_h, a) - \beta \text{KL} \left(\pi(\cdot | s_h) | \pi_{ref}(\cdot | s_h) \right) \Rightarrow \pi^*(a | s_h) \propto \pi_{ref}(a | s_h) \exp(Q^*(s_h, a) / \beta)$$

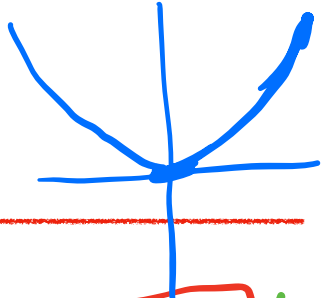


DP for solving the KL-regularized RL

$\max\{x, -x\}$

$$\max_{\pi} \mathbb{E}_{\tau \sim \pi} \left[\sum_{h=0}^{H-1} r(s_h, a_h) - \beta \text{KL} \left(\pi(\cdot | s_h) | \pi_{\text{ref}}(\cdot | s_h) \right) \right]$$

$$\ln(\exp(-x) + \exp(x))$$



Base case: $V^*(s_H) = 0$, for the fictitious step H

Induction step: given $V^*(s_{h+1})$, want to compute $V^*(s_h)$

$$Q^*(s_h, a_h) = r(s_h, a_h) + \mathbb{E}_{s_{h+1} \sim P(\cdot | s_h, a_h)} V^*(s_{h+1})$$

$$V^*(s_h) = \max_{\pi(\cdot | s_h) \in \Delta(A)} \mathbb{E}_{a \sim \pi(\cdot | s_h)} Q^*(s_h, a) - \beta \text{KL} \left(\pi(\cdot | s_h) | \pi_{\text{ref}}(\cdot | s_h) \right)$$

$$V^*(s_h) = \beta \ln \mathbb{E}_{a_h \sim \pi_{\text{ref}}(\cdot | s_h)} \left[\exp(Q^*(s_h, a_h) / \beta) \right]$$

$$V^*(s_h) = \beta \ln \mathbb{E}_{a \sim \pi_{\text{ref}}(\cdot | s_h)} \left(\frac{\exp(Q^*(s_h, a) / \beta)}{\beta} \right)$$

softmax $\approx \max_a$

$\Rightarrow \pi^*(a | s_h) \propto \pi_{\text{ref}}(a | s_h) \exp(Q^*(s_h, a) / \beta)$
 $\beta \rightarrow 0 \quad V^*(s_h) \rightarrow \max_a Q^*(s_h, a)$

DP for solving the KL-regularized RL

$$\max_{\pi} \mathbb{E}_{\tau \sim \pi} \left[\sum_{h=0}^{H-1} r(s_h, a_h) - \beta \text{KL} \left(\pi(\cdot | s_h) | \pi_{ref}(\cdot | s_h) \right) \right]$$

Base case: $V^*(s_H) = 0$, for the fictitious step H

Induction step: given $V^*(s_{h+1})$, want to compute $V^*(s_h)$

$$Q^*(s_h, a_h) = r(s_h, a_h) + \mathbb{E}_{s_{h+1} \sim P(\cdot | s_h, a_h)} V^*(s_{h+1})$$

$$V^*(s_h) = \max_{\pi(\cdot | s_h) \in \Delta(A)} \mathbb{E}_{a \sim \pi(\cdot | s_h)} Q^*(s_h, a) - \beta \text{KL} \left(\pi(\cdot | s_h) | \pi_{ref}(\cdot | s_h) \right) \Rightarrow \pi^*(a | s_h) \propto \pi_{ref}(a | s_h) \exp(Q^*(s_h, a) / \beta)$$

$$V^*(s_h) = \beta \ln \mathbb{E}_{a_h \sim \pi_{ref}(\cdot | s_h)} \left[\exp(Q^*(s_h, a_h) / \beta) \right]$$

(Exercise: show $V^*(s) \rightarrow \max_a Q^*(s, a)$, when $\beta \rightarrow 0$, assuming $\pi_{ref}(a | s) > 0, \forall a$)

DP for solving the KL-regularized RL

$$V^*(s_h) = \beta \ln \mathbb{E}_{a_h \sim \pi_{\text{ref}}(\cdot | s_h)} [\exp(Q^*(s_h, a_h) / \beta)]$$

Now let's assume transition is deterministic, i.e., $s_{h+1} = f(s_h, a_h)$, and see if we can further simplify V^*

$$Q^*(s_n, a_n) = r_n + V^*(s_{n+1}) \quad \text{where } s_{n+1} = f(s_n, a_n)$$

DP for solving the KL-regularized RL

$$V^*(s_h) = \beta \ln \mathbb{E}_{a_h \sim \pi_{ref}(\cdot | s_h)} [\exp(Q^*(s_h, a_h)/\beta)]$$

Now let's assume transition is deterministic, i.e., $s_{h+1} = f(s_h, a_h)$, and see if we can further simplify V^*

$$\exp(V^*(s_h)/\beta) = \mathbb{E}_{a_h \sim \pi_{ref}(\cdot | s_h)} [\exp(r_h/\beta + \underline{V^*(s_{h+1})/\beta})], \text{ where } s_{h+1} = f(s_h, a_h)$$

$$\exp\left(\frac{r}{\beta} + \frac{V^*}{\beta}\right) = \exp\left(\frac{r}{\beta}\right) \cdot \exp\left(\frac{V^*}{\beta}\right)$$

DP for solving the KL-regularized RL

$$V^*(s_h) = \beta \ln \mathbb{E}_{a_h \sim \pi_{ref}(\cdot | s_h)} [\exp(Q^*(s_h, a_h)/\beta)]$$

Now let's assume transition is deterministic, i.e., $s_{h+1} = f(s_h, a_h)$, and see if we can further simplify V^*

$$\exp(V^*(s_h)/\beta) = \mathbb{E}_{a_h \sim \pi_{ref}(\cdot | s_h)} [\exp(r_h/\beta + V^*(s_{h+1})/\beta)], \text{ where } s_{h+1} = f(s_h, a_h)$$

$$= \mathbb{E}_{a_h \sim \pi_{ref}(\cdot | s_h)} \underbrace{\exp(r_h/\beta)} \underbrace{\exp(V^*(s_{h+1})/\beta)}$$

DP for solving the KL-regularized RL

$$V^*(s_h) = \beta \ln \mathbb{E}_{a_h \sim \pi_{\text{ref}}(\cdot | s_h)} [\exp(Q^*(s_h, a_h)/\beta)]$$

Now let's assume transition is deterministic, i.e., $s_{h+1} = f(s_h, a_h)$, and see if we can further simplify V^*

$$\exp(V^*(s_h)/\beta) = \mathbb{E}_{a_h \sim \pi_{\text{ref}}(\cdot | s_h)} [\exp(r_h/\beta + V^*(s_{h+1})/\beta)], \text{ where } s_{h+1} = f(s_h, a_h)$$

$$= \mathbb{E}_{a_h \sim \pi_{\text{ref}}(\cdot | s_h)} \exp(r_h/\beta) \exp(V^*(s_{h+1})/\beta) \quad \text{Recursion again}$$

$$\mathbb{E}_{a_{h+1} \sim \pi_{\text{ref}}(\cdot | s_{h+1})} \exp\left(\frac{r_{h+1}}{\beta}\right) \exp\left(\frac{V^*(s_{h+2})}{\beta}\right)$$

DP for solving the KL-regularized RL

$$V^*(s_h) = \beta \ln \mathbb{E}_{a_h \sim \pi_{ref}(\cdot | s_h)} [\exp(Q^*(s_h, a_h) / \beta)]$$

Now let's assume transition is deterministic, i.e., $s_{h+1} = f(s_h, a_h)$, and see if we can further simplify V^*

$$\exp(V^*(s_h) / \beta) = \mathbb{E}_{a_h \sim \pi_{ref}(\cdot | s_h)} [\exp(r_h / \beta + V^*(s_{h+1}) / \beta)], \text{ where } s_{h+1} = f(s_h, a_h)$$

$$= \mathbb{E}_{a_h \sim \pi_{ref}(\cdot | s_h)} \exp(r_h / \beta) \exp(V^*(s_{h+1}) / \beta)$$

Recursion again

Repeat

$$= \mathbb{E}_{a_h \sim \pi_{ref}(\cdot | s_h)} \exp(r_h / \beta) \mathbb{E}_{a_{h+1} \sim \pi_{ref}(\cdot | s_{h+1})} \exp(r_{h+1} / \beta) \exp(V^*(s_{h+2}))$$

DP for solving the KL-regularized RL

$$V^*(s_h) = \beta \ln \mathbb{E}_{a_h \sim \pi_{ref}(\cdot | s_h)} [\exp(Q^*(s_h, a_h)/\beta)]$$

Now let's assume transition is deterministic, i.e., $s_{h+1} = f(s_h, a_h)$, and see if we can further simplify V^*

$$\exp(V^*(s_h)/\beta) = \mathbb{E}_{a_h \sim \pi_{ref}(\cdot | s_h)} [\exp(r_h/\beta + V^*(s_{h+1})/\beta)], \text{ where } s_{h+1} = f(s_h, a_h)$$

$$= \mathbb{E}_{a_h \sim \pi_{ref}(\cdot | s_h)} \exp(r_h/\beta) \exp(V^*(s_{h+1})/\beta) \quad \text{Recursion again}$$

$$= \mathbb{E}_{a_h \sim \pi_{ref}(\cdot | s_h)} \exp(r_h/\beta) \mathbb{E}_{a_{h+1} \sim \pi_{ref}(\cdot | s_{h+1})} \exp(r_{h+1}/\beta) \exp(V^*(s_{h+2}))$$

$$= \mathbb{E}_{\tau \sim \pi_{ref}(\cdot | s_h)} \exp\left(\sum_{\tau=h}^{H-1} r_\tau / \beta\right)$$

DP for solving the KL-regularized RL

$$V^*(s_h) = \beta \ln \mathbb{E}_{a_h \sim \pi_{ref}(\cdot | s_h)} [\exp(Q^*(s_h, a_h)/\beta)]$$

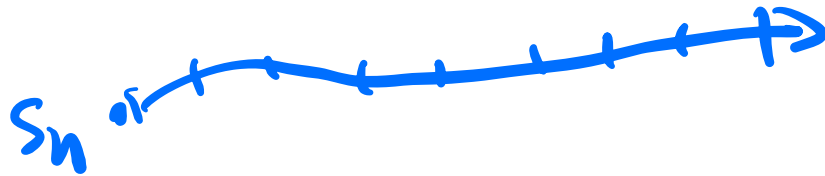
Now let's assume transition is deterministic, i.e., $s_{h+1} = f(s_h, a_h)$, and see if we can further simplify V^*

$$\exp(V^*(s_h)/\beta) = \mathbb{E}_{a_h \sim \pi_{ref}(\cdot | s_h)} [\exp(r_h/\beta + V^*(s_{h+1})/\beta)], \text{ where } s_{h+1} = f(s_h, a_h)$$

$$= \mathbb{E}_{a_h \sim \pi_{ref}(\cdot | s_h)} \underbrace{\exp(r_h/\beta) \exp(V^*(s_{h+1})/\beta)}_{\text{Recursion again}}$$

$$= \mathbb{E}_{a_h \sim \pi_{ref}(\cdot | s_h)} \exp(r_h/\beta) \mathbb{E}_{a_{h+1} \sim \pi_{ref}(\cdot | s_{h+1})} \exp(r_{h+1}/\beta) \exp(V^*(s_{h+2}))$$

$$= \mathbb{E}_{\tau \sim \pi_{ref}(\cdot | s_h)} \exp\left(\sum_{\tau=h}^{H-1} r_\tau/\beta\right)$$



$\tau \sim \pi_{ref}(\cdot | s_h)$: Denotes generating a future trajectory using π_{ref} from state s_h

In summary, when transition is deterministic, we have

$$\forall h, s : \exp(V^*(s_h)/\beta) = \mathbb{E}_{\pi_{ref}} \left[\exp \left(\sum_{\tau=h}^{H-1} r_{\tau}/\beta \right) \mid s_h \right]$$

$$\forall h, s, a : \exp(Q^*(s_h, a_h)/\beta) = \mathbb{E}_{\pi_{ref}} \left[\exp \left(\sum_{\tau=h}^{H-1} r_{\tau}/\beta \right) \mid s_h, a_h \right]$$

In summary, when transition is deterministic, we have

$$\forall h, s : \exp(V^*(s_h)/\beta) = \mathbb{E}_{\pi_{ref}} \left[\exp \left(\sum_{\tau=h}^{H-1} r_{\tau}/\beta \right) \mid s_h \right]$$

s_h → π_{ref}

$$\forall h, s, a : \exp(Q^*(s_h, a_h)/\beta) = \mathbb{E}_{\pi_{ref}} \left[\exp \left(\sum_{\tau=h}^{H-1} r_{\tau}/\beta \right) \mid s_h, a_h \right]$$

s_h, a_h → π_{ref}

Note the expectation is always wrt to the future generated from π_{ref}

In summary, when transition is deterministic, we have

Also recall the optimal policy format:

$$\pi^*(a_h | s_h) \propto \underbrace{\pi_{ref}(a_h | s_h)} \exp \left(\underbrace{Q^*(s_h, a_h) / \beta} \right)$$

In summary, when transition is deterministic, we have

Also recall the optimal policy format:

$$\pi^*(a_h | s_h) \propto \pi_{ref}(a_h | s_h) \exp(Q^*(s_h, a_h) / \beta)$$


As long as we can learn $\exp(Q^*(s, a) / \beta)$, then we can use it to guide π_{ref}

In summary, when transition is deterministic, we have

Also recall the optimal policy format:

$$\pi^*(a_h | s_h) \propto \pi_{ref}(a_h | s_h) \exp(Q^*(s_h, a_h) / \beta)$$

As long as we can learn $\exp(Q^*(s, a) / \beta)$, then we can use it to guide π_{ref}

Q: why this format of π^* is tractable and implementable?

$$\pi_{ref}(\cdot | s_h) \in \mathcal{R}^{|A|} \quad \exp(Q^*(s_h, a) / \beta) \quad \forall a \in \mathcal{A}$$

Outline

1. KL regularized RL again, but in token space (i.e., s_h, a_h) not traj space

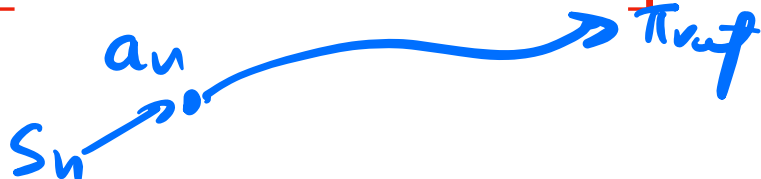
2. Train value/Q functions

$$\pi_{\text{ref}} \cdot \exp\left(\frac{Q^*}{\beta}\right)$$

3. Controllable generation via guidance from Q/V functions

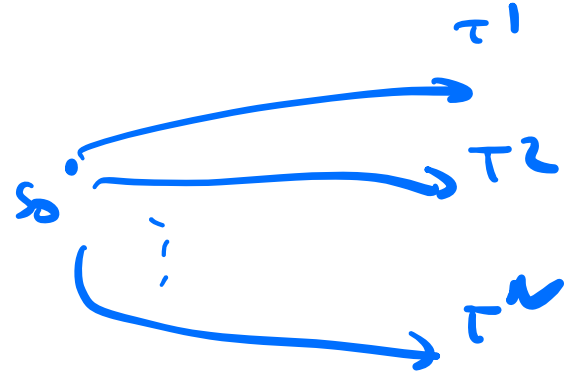
Recall the format of Q^*/V^*

($S_{h+1} \leftarrow f(S_h, a_h)$)

$$\forall h, s : \exp(Q^*(s_h, a_h)/\beta) = \mathbb{E}_{\pi_{ref}} \left[\exp \left(\sum_{\tau=h}^{H-1} r_{\tau}/\beta \right) \mid s_h, a_h \right]$$


$$\text{Learn}(\exp(Q^*/\beta)) = E_{\tau \sim \pi_{ref}} \left(\exp \left(\sum_{t=0}^{H-1} \frac{r_t}{\beta} \right) \right)$$

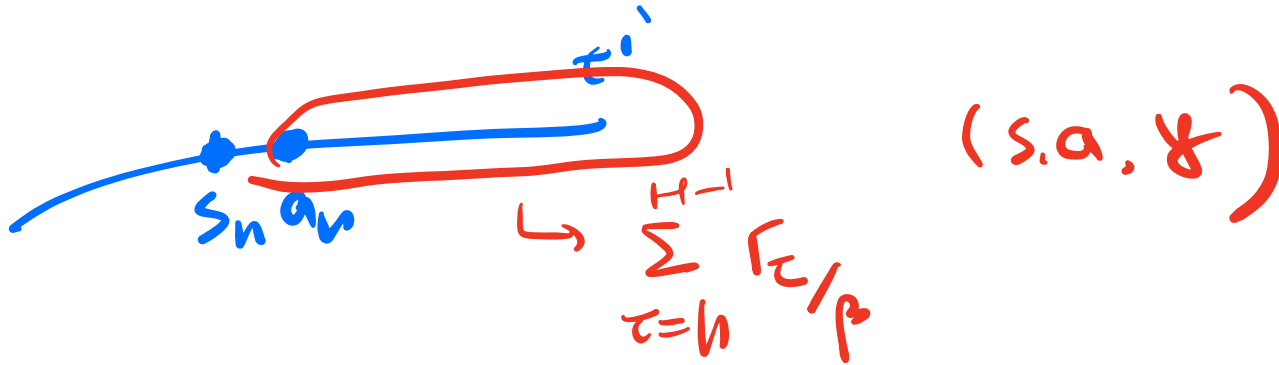
1. Data collection: generate N i.i.d trajectories from π_{ref} , $\tau^1, \dots, \tau^N \sim \pi_{ref}$



Learn $\exp(Q^*/\beta)$

1. Data collection: generate N i.i.d trajectories from π_{ref} , $\tau^1, \dots, \tau^N \sim \pi_{ref}$

2. For each $s_h, a_h \in \tau^i$, compute reward-to-go $y = \sum_{\tau=h}^{H-1} r_\tau$ on τ^i



Learn $\exp(Q^*/\beta)$

1. Data collection: generate N i.i.d trajectories from π_{ref} , $\tau^1, \dots, \tau^N \sim \pi_{ref}$

2. For each $s_h, a_h \in \tau^i$, compute reward-to-go $y = \sum_{\tau=h}^{H-1} r_\tau$ on τ^i

3 Given the data $\{(s, a), y\}$, train g via least square regression:

Learn $\exp(Q^*/\beta)$

1. Data collection: generate N i.i.d trajectories from π_{ref} , $\tau^1, \dots, \tau^N \sim \pi_{ref}$

2. For each $s_h, a_h \in \tau^i$, compute reward-to-go $y = \sum_{\tau=h}^{H-1} r_\tau$ on τ^i

3 Given the data $\{(s, a), y\}$, train g via least square regression:

$$\exp\left(\frac{Q^*}{\beta}\right) \leftarrow \hat{g} = \min_g \sum_{(s,a), y \in \mathcal{D}} \left(\underline{g(s, a)} - \underline{\exp(y/\beta)} \right)^2$$

Learn $\exp(Q^*/\beta)$

1. Data collection: generate N i.i.d trajectories from $\pi_{ref}, \tau^1, \dots, \tau^N \sim \pi_{ref}$

2. For each $s_h, a_h \in \tau^i$, compute reward-to-go $y = \sum_{\tau=h}^{H-1} r_\tau$ on τ^i

3 Given the data $\{(s, a), y\}$, train g via least square regression:

$$\hat{g} = \min_g \sum_{(s,a), y \in \mathcal{D}} (g(s, a) - \exp(y/\beta))^2$$

$y \leftarrow \pi_{ref}$

$\mathbb{E}_{\pi_{ref}} \exp(\frac{y}{\beta})$

Q: what's the Bayes optimal of this regression problem?

Learn $\exp(V^*/\beta)$

$$\hat{g} = \min_g \sum_{(s,a), y \in \mathcal{D}} (g(s, a) - \exp(y/\beta))^2$$

Bayes opt: $\mathbb{E}_{\pi_{ref}} \left[\exp\left(\sum_{\tau=h}^{H-1} r_{\tau}/\beta \right) \mid s_h = s, a_h = a \right] = \exp(Q^*(s, a)/\beta)$

Learn $\exp(V^*/\beta)$

$$\hat{g} = \min_g \sum_{(s,a), y \in \mathcal{D}} (g(s, a) - \exp(y/\beta))^2$$

$$\text{Bayes opt: } \mathbb{E}_{\pi_{ref}} \left[\exp\left(\sum_{\tau=h}^{H-1} r_\tau / \beta\right) \mid s_h = s, a_h = a \right] = \exp(Q^*(s, a) / \beta)$$

Under reasonable conditions, let's assume $\hat{g} \approx \text{Bayes opt}$

$$\hat{g} \approx \exp\left(\frac{Q^*(s, a)}{\beta}\right)$$

Outline

1. KL regularized RL again, but in token space (i.e., s_h, a_h) not traj space

2. Train value/Q functions



3. Controllable generation via guidance from Q/V functions

Controlling generation via Q^* guidance

Least square regression learns

$$\hat{g}(s_h, a_h) \approx \exp(Q^*(s_h, a_h)/\beta) = \mathbb{E}_{\pi_{ref}} \left[\exp \left(\sum_{\tau=h}^{H-1} r_{\tau}/\beta \right) \mid s_h, a_h \right]$$

Controlling generation via Q^* guidance

During inference time, we can generate a trajectory use the following policy:

$$\begin{aligned} \pi^* &\propto \pi_{ref} \exp\left(\frac{Q^*}{\beta}\right) \\ &\propto \pi_{ref}(a|s) \exp\left(\frac{Q^*(s,a)}{\beta}\right) \end{aligned}$$

\hat{g}

Controlling generation via Q^* guidance

During inference time, we can generate a trajectory use the following policy:

$$\pi(a | s_h) \propto \pi_{ref}(a | s_h) \hat{g}(s_h, a)$$

1. Note that we never do backpropagation on π_{ref}

Controlling generation via Q^* guidance

During inference time, we can generate a trajectory use the following policy:

γ
=

$$\pi(a | s_h) \propto \pi_{ref}(a | s_h) \hat{g}(s_h, a)$$

1. Note that we never do backpropagation on π_{ref}
2. $\hat{g}(s_h, a) \approx \exp(Q^*(s_h, a)/\beta)$ is an evaluation function, not a generator!



Experiment results

Guide black-box models to do math reasoning

Reward: 0/1 – 1 means model's final answer is correct solution, 0 otherwise

Experiment results

Guide black-box models to do math reasoning

Reward: 0/1 — 1 means model's final answer is correct solution, 0 otherwise

For reasoning tasks, we often prompt model to think **step-by-step before outputting a final answer**

Experiment results

Guide black-box models to do math reasoning

Reward: 0/1 — 1 means model's final answer is correct solution, 0 otherwise

For reasoning tasks, we often prompt model to think **step-by-step before outputting a final answer**
(aka chain-of-thought (CoT) reasoning)

Experiment results

Guide black-box models to do math reasoning

Reward: 0/1 — 1 means model's final answer is correct solution, 0 otherwise

For reasoning tasks, we often prompt model to think **step-by-step before outputting a final answer**
(aka chain-of-thought (CoT) reasoning)

Problem from GSM8K: According to its nutritional info, a bag of chips has 250 calories per serving. If a 300g bag has 5 servings, how many grams can you eat if your daily calorie target is 2000 and you have already consumed 1800 calories?

Experiment results

Guide black-box models to do math reasoning

Reward: 0/1 — 1 means model's final answer is correct solution, 0 otherwise

For reasoning tasks, we often prompt model to think **step-by-step before outputting a final answer**
(aka chain-of-thought (CoT) reasoning)

Problem from GSM8K: According to its nutritional info, a bag of chips has 250 calories per serving. If a 300g bag has 5 servings, how many grams can you eat if your daily calorie target is 2000 and you have already consumed 1800 calories?

Incorrect Solution from π^{ref} : Let's break this down step by step!

1. The bag of chips has 5 servings, and each serving has 250 calories. To find the total calories in the bag, multiply the number of servings by the calories per serving:

$$5 \text{ servings} \times 250 \text{ calories/serving} = 1250 \text{ calories}$$

2. You have already consumed 1800 calories, and your daily calorie target is 2000. To find out how many more calories you can consume, subtract the calories you've already had from your daily target:

$$2000 - 1800 = 200 \text{ calories}$$

3. Since each serving of chips has 250 calories, you can consume:

$$200 \text{ calories} \div 250 \text{ calories/serving} = 0.8 \text{ servings}$$

4. **Since there are 5 servings in the bag, you can eat:**

$$0.8 \text{ servings} \times 300\text{g/bag} = 240\text{g}$$

So, you can eat up to 240g of the bag of chips.

Here's the answer:

240 g



Experiment results

Guide black-box models to do math reasoning

Reward: 0/1 — 1 means model's final answer is correct solution, 0 otherwise

For reasoning tasks, we often prompt model to think **step-by-step before outputting a final answer**
(aka chain-of-thought (CoT) reasoning)

Problem from GSM8K: According to its nutritional info, a bag of chips has 250 calories per serving. If a 300g bag has 5 servings, how many grams can you eat if your daily calorie target is 2000 and you have already consumed 1800 calories?

Incorrect Solution from π^{ref} : Let's break this down step by step!

1. The bag of chips has 5 servings, and each serving has 250 calories. To find the total calories in the bag, multiply the number of servings by the calories per serving:

$$5 \text{ servings} \times 250 \text{ calories/serving} = 1250 \text{ calories}$$

2. You have already consumed 1800 calories, and your daily calorie target is 2000. To find out how many more calories you can consume, subtract the calories you've already had from your daily target:

$$2000 - 1800 = 200 \text{ calories}$$

3. Since each serving of chips has 250 calories, you can consume:

$$200 \text{ calories} \div 250 \text{ calories/serving} = 0.8 \text{ servings}$$

4. Since there are 5 servings in the bag, you can eat:

$$0.8 \text{ servings} \times 300\text{g/bag} = 240\text{g}$$

So, you can eat up to 240g of the bag of chips.

Here's the answer:

~~240g~~

Experiment results

Guide black-box models to do math reasoning

Reward: 0/1 — 1 means model's final answer is correct solution, 0 otherwise

For reasoning tasks, we often prompt model to think **step-by-step before outputting a final answer**
(aka chain-of-thought (CoT) reasoning)

Problem from GSM8K: According to its nutritional info, a bag of chips has 250 calories per serving. If a 300g bag has 5 servings, how many grams can you eat if your daily calorie target is 2000 and you have already consumed 1800 calories?

Incorrect Solution from π^{ref} : Let's break this down step by step!

1. The bag of chips has 5 servings, and each serving has 250 calories. To find the total calories in the bag, multiply the number of servings by the calories per serving:

$$5 \text{ servings} \times 250 \text{ calories/serving} = 1250 \text{ calories}$$

2. You have already consumed 1800 calories, and your daily calorie target is 2000. To find out how many more calories you can consume, subtract the calories you've already had from your daily target:

$$2000 - 1800 = 200 \text{ calories}$$

3. Since each serving of chips has 250 calories, you can consume:

$$200 \text{ calories} \div 250 \text{ calories/serving} = 0.8 \text{ servings}$$

4. Since there are 5 servings in the bag, you can eat:

$$0.8 \text{ servings} \times 300\text{g}/\text{bag} = 240\text{g}$$

So, you can eat up to 240g of the bag of chips.

Here's the answer:

240 g

π_{ref} made a mistake

Experiment results

Train a small model (1B) $\hat{g}(s, a)$ to approximate $\exp(Q^*(s, a)/\beta)$

Define the new policy as $\pi_{new} \propto \pi_{ref}(a | s) \cdot \hat{g}(s, a)$

Let's compare the generation from π_{ref} and π_{new}

Tip

Incorrect Solution from π^{ref} : Let's break this down step by step!

1. The bag of chips has 5 servings, and each serving has 250 calories. To find the total calories in the bag, multiply the number of servings by the calories per serving:

$$5 \text{ servings} \times 250 \text{ calories/serving} = 1250 \text{ calories}$$

2. You have already consumed 1800 calories, and your daily calorie target is 2000. To find out how many more calories you can consume, subtract the calories you've already had from your daily target:

$$2000 - 1800 = 200 \text{ calories}$$

3. Since each serving of chips has 250 calories, you can consume:

$$200 \text{ calories} \div 250 \text{ calories/serving} = 0.8 \text{ servings}$$

Q

4. Since there are 5 servings in the bag, you can eat:

$$0.8 \text{ servings} \times 300\text{g/bag} = 240\text{g}$$

So, you can eat up to 240g of the bag of chips.

Here's the answer:

240 g

Almost the same

Tip

Correct Solution from Q#: A delicious problem!

Let's break it down step by step:

1. The bag of chips has 250 calories per serving, and there are 5 servings in the bag.

2. The total calories in the bag are: $250 \text{ calories/serving} \times 5 \text{ servings} = 1250 \text{ calories}$.

3. You have already consumed 1800 calories, so you have $2000 - 1800 = 200 \text{ calories}$ left for the day.

4. Since each serving of chips has 250 calories, you can consume: $200 \text{ calories} / 250 \text{ calories/serving} = 0.8 \text{ servings}$.

5. Since there are 5 servings in the bag, you can eat: $0.8 \text{ servings} \times 300\text{g}/5 \text{ servings} = 0.8 \times 60\text{g} = 48\text{g}$.

So, you can eat up to 48g of the bag of chips.

Enjoy your snacking, but don't overdo it!

π^{ref}

Incorrect Solution from π^{ref} : Let's break this down step by step!

1. The bag of chips has 5 servings, and each serving has 250 calories. To find the total calories in the bag, multiply the number of servings by the calories per serving:

$$5 \text{ servings} \times 250 \text{ calories/serving} = 1250 \text{ calories}$$

2. You have already consumed 1800 calories, and your daily calorie target is 2000. To find out how many more calories you can consume, subtract the calories you've already had from your daily target:

$$2000 - 1800 = 200 \text{ calories}$$

3. Since each serving of chips has 250 calories, you can consume:

$$200 \text{ calories} \div 250 \text{ calories/serving} = 0.8 \text{ servings}$$

4. **Since there are 5 servings in the bag, you can eat:**

$$0.8 \text{ servings} \times 300\text{g/bag} = 240\text{g}$$

So, you can eat up to 240g of the bag of chips.

Here's the answer:

240 g

Almost the same

Corrected the mistake

g^{ref}

Correct Solution from $Q\#$: A delicious problem!

Let's break it down step by step:

1. The bag of chips has 250 calories per serving, and there are 5 servings in the bag.

2. The total calories in the bag are: $250 \text{ calories/serving} \times 5 \text{ servings} = 1250 \text{ calories}$.

3. You have already consumed 1800 calories, so you have $2000 - 1800 = 200 \text{ calories}$ left for the day.

4. Since each serving of chips has 250 calories, you can consume: $200 \text{ calories} / 250 \text{ calories/serving} = 0.8 \text{ servings}$.

5. **Since there are 5 servings in the bag, you can eat:** $0.8 \text{ servings} \times 300\text{g}/5 \text{ servings} = 0.8 \times 60\text{g} = 48\text{g}$.

So, you can eat up to 48g of the bag of chips.

Enjoy your snacking, but don't overdo it!

Summary

1. DP on the KL-regularized RL objective
2. With the KL-reg, now the hard \max_a is replaced by a soft-max operator $\ln \mathbb{E} \exp(\dots)$
3. The V^*/Q^* has very simple form when transition is deterministic
4. Simple regression to learn a approximator of $\exp(Q^*/\beta)$ directly, and use it to guide π_{ref} in generation