

PREFERENCE ELICITATION FOR OFFLINE REINFORCEMENT LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

Applying reinforcement learning (RL) to real-world problems is often made challenging by the inability to interact with the environment and the difficulty of designing reward functions. Offline RL addresses the first challenge by considering access to an offline dataset of environment interactions labeled by the reward function. In contrast, Preference-based RL does not assume access to the reward function and learns it from preferences, but typically requires an online interaction with the environment. We bridge the gap between these frameworks by exploring efficient methods for acquiring preference feedback in a fully offline setup. We propose Sim-OPRL, an offline preference-based reinforcement learning algorithm, which leverages a learned environment model to elicit preference feedback on simulated rollouts. Drawing on insights from both the offline RL and the preference-based RL literature, our algorithm employs a pessimistic approach for out-of-distribution data, and an optimistic approach for acquiring informative preferences about the optimal policy. We provide theoretical guarantees regarding the sample complexity of our approach, dependent on how well the offline data covers the optimal policy. Finally, we demonstrate the empirical performance of Sim-OPRL in various environments.

1 INTRODUCTION

While reinforcement learning (RL) (Sutton and Barto, 2018) achieves excellent performance in various decision-making tasks (Kendall et al., 2019; Mirhoseini et al., 2020; Degraeve et al., 2022), its practical deployment remains limited by the requirement of direct interaction with the environment. This can be impractical or unsafe in real-world scenarios. For example, patient management and treatment in intensive care units involve complex decision-making that has often been framed as a reinforcement learning problem (Raghu et al., 2017; Komorowski et al., 2018). However, the timing, dosage, and combination of treatments required are critical to patient safety, and incorrect decisions can lead to severe complications or death, making the use of traditional RL algorithms unfeasible (Gottesman et al., 2019; Tang and Wiens, 2021). Offline RL emerges as a promising solution, allowing policy learning from entirely observational data (Levine et al., 2020).

Still, a challenge with Offline RL is its requirement for an explicit reward function. Quantifying the numerical value of taking a certain action in a given environment state is challenging in many applications (Yu et al., 2021). Preference-based RL offers a compelling alternative, relying on comparisons between different actions or trajectories (Wirth et al., 2017) and being often easier for humans to provide (Christiano et al., 2017). In medical settings, for instance, clinicians may be queried for feedback on which trajectories lead to favorable outcomes. Unfortunately, most algorithms for preference acquisition require environment interaction (Saha et al., 2023; Chen et al., 2022; Lindner et al., 2021) and are therefore not applicable to the offline setting.

Lack of environment interaction and reward learning are thus two critical challenges for real-world RL deployment that are rarely tackled jointly. In this work, we address the problem of preference elicitation for offline reinforcement learning by asking: *What trajectories should we sample to minimize the number of human queries required to learn the best offline policy?* This presents a challenging problem as it combines learning from offline data and active feedback acquisition, two frameworks that require opposing inductive biases for conservatism and exploration, respectively.

To the best of our knowledge, the only strategy proposed in prior work is to acquire feedback directly over samples within an offline dataset of trajectories (Shin et al., 2022, Offline Preference-based Reward Learning (OPRL)). We propose an alternative solution that queries feedback on *simulated rollouts* by leveraging a learned environment model. Our offline preference-based reinforcement learning algorithm, Sim-OPRL, strikes a balance between conservatism and exploration by combining pessimism when handling states out-of-distribution from the observational data (Jin et al., 2021; Zhan et al., 2023a), and optimism in acquiring informative preferences about the optimal policy (Saha et al., 2023; Chen et al., 2022). We study the efficiency of our approach through both theoretical and empirical analysis, demonstrating the superior performance of Sim-OPRL across various environments.

Our contributions are the following: (1) In Section 3, we first formalize the new problem setting of preference elicitation for offline reinforcement learning, which allows for **complementing offline data with preference feedback**. This framework is crucial for real-world RL applications where direct environment interaction is unsafe or impractical and reward functions are challenging to design manually, yet experts can be queried for their knowledge. (2) In Section 5, we provide theoretical guarantees on eliciting preference feedback over samples from an offline dataset, complementing work of Shin et al. (2022). (3) Next, in Section 6, we propose our own **efficient preference elicitation algorithm** based on simulated rollouts in a learned environment model, and establish its improved theoretical guarantees. (4) Finally, we develop a practical implementation of our algorithm and demonstrate its **empirical efficiency** and scalability across various decision-making environments.

2 RELATED WORK

Our problem setting shares similarities with Offline RL and Preference-based RL, which we summarize below. We position ourselves relative to our closest related works in Table 1 and extend our discussion in Appendix B.

Offline RL. Offline Reinforcement Learning has gained significant traction in recent years, as the practicality of training RL agents without environment interaction makes it relevant to real-world applications (Levine et al., 2020). However, learning from observational data only is a source of bias in the model, as the data may not cover the entire state-action space. Offline RL algorithms therefore output pessimistic policies, which has been shown to minimize suboptimality (Jin et al., 2021). Model-based approaches show particular promise for their sample efficiency (Yu et al., 2020; Kidambi et al., 2020; Rigter et al., 2022; Zhai et al., 2024; Uehara and Sun, 2021). In this work, we study the setting where reward signals are unavailable and must be estimated by actively querying preference feedback.

Preference-based RL. Rather than accessing numerical reward values for each state-action pair as in traditional online RL, preference-based RL learns the reward model through collecting pairwise preferences over trajectories (Wirth et al., 2017). Different preference elicitation strategies have been proposed for this framework, generally based on knowing the transition model exactly or on having access to the environment for rollouts (Christiano et al., 2017; Saha et al., 2023; Chen et al., 2022; Lindner et al., 2021; Zhan et al., 2023b; Sadigh et al., 2018; Brown et al., 2020). Prior theoretical and empirical work (Lindner et al., 2021; Chen et al., 2022) show that, in this setting, the most efficient preference elicitation strategy is to actively reduce the set of candidate optimal policies, rather maximize information gain on the reward function – our theoretical and empirical results reach the same conclusion for the offline setting.

Offline Preference-based RL. The development of preference-based RL algorithms based on offline data only is critical to settings where environment interaction is not feasible for safety and

Table 1: Comparison of related work on preference elicitation.

Framework	Offline	Policy-Based Sampling	Robustness Guarantees	Practical Implementation
PbOP (Chen et al., 2022)	✗	✓	✓	✗
MoP-RL (Liu et al., 2023)	✗	✗	✗	✓
REGIME (Zhan et al., 2023b)	✗	✓	✓	✗
FREEHAND (Zhan et al., 2023a)	✓	✗	✓	✗
OPRL (Shin et al., 2022)	✓	✗	✗	✓
Sim-OPRL (Ours)	✓	✓	✓	✓

108 efficiency reasons. Still, this framework remains largely unexplored in the literature. While Zhu
 109 et al. (2023); Zhan et al. (2023a) demonstrate the value of pessimism in offline preference-based re-
 110 inforcement learning, they do not consider how to query feedback actively. On the other hand, Shin
 111 et al. (2022) propose an empirical comparison of different preference sampling trajectories from an
 112 offline trajectories buffer. In Section 5, we provide a theoretical analysis of their approach, then
 113 propose an alternative sampling strategy based on simulated trajectory rollouts in Section 6, which
 114 benefits from both theoretical motivation and superior empirical performance.

115 3 PROBLEM FORMULATION

116 3.1 PRELIMINARIES

119 **Markov Decision Process.** We consider the episodic Markov Decision Process (MDP), defined by
 120 the tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, H, T, R)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, H is the episode
 121 length, $T : \mathcal{S} \times \mathcal{A} \rightarrow \Delta_{\mathcal{S}}$ is the transition function, $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function. We
 122 assume an initial state s_0 , but our analysis generalizes to a fixed initial state distribution. At time t ,
 123 the environment is at state $s_t \in \mathcal{S}$ and an agent selects an action $a_t \in \mathcal{A}$. The agent then receives a
 124 reward $R(s_t, a_t)$ and the environment transitions to state $s_{t+1} \sim T(\cdot | s_t, a_t)$. We describe an agent’s
 125 behavior through a policy function $\pi : \mathcal{S} \rightarrow \Delta_{\mathcal{A}}$, such that $\pi(a|s)$ is the probability of taking action
 126 a in state s . Let $\tau = (s_0, a_0, \dots, s_H, a_H)$ denote the trajectory of state-action pairs of an interaction
 127 episode with the environment. With an abuse of notation, we also write $R(\tau) = \sum_t R(s_t, a_t)$. Let
 128 d_T^π denote the distribution of trajectories (or state-action pairs, overloading notation) induced by
 129 rolling out policy π in transition model T . We denote the expected return of policy π as $V_{T,R}^\pi =$
 130 $\mathbb{E}_{\tau \sim d_T^\pi} [R(\tau)]$, and $\pi^* = \operatorname{argmax}_\pi V_{T,R}^\pi$ denotes the optimal policy in \mathcal{M} .

131 **Preference-based Reinforcement Learning.** Rather than observing rewards for each state and
 132 action, we receive preference feedback over trajectories. For a pair of trajectories (τ_1, τ_2) , we obtain
 133 binary feedback $o \in \{0, 1\}$ about whether τ_1 is preferred to τ_2 . We define the preference function
 134 P_R and assume that preference labels follow the Bradley-Terry model (Bradley and Terry, 1952):

$$135 P_R(\tau_1 \succ \tau_2) := P(o = 1 | \tau_1, \tau_2) = \frac{\exp(R(\tau_1))}{\exp(R(\tau_1)) + \exp(R(\tau_2))} = \sigma(R(\tau_1) - R(\tau_2)), \quad (1)$$

137 where \succ denotes a preference relationship and σ is a monotonous increasing link function. Within
 138 this framework, *preference elicitation* refers to the process of sampling preferences to obtain infor-
 139 mation about both the preference function and the system dynamics (Wirth et al., 2017).

140 3.2 OFFLINE PREFERENCE ELICITATION

143 We assume access to an observational dataset of trajectories $\mathcal{D}_{\text{offline}} = \{\tau : \tau \sim d_T^{\pi_\beta}\}$, where π_β is
 144 an unknown behavioural policy in \mathcal{M} . As in Offline RL, we do *not* have access to the environment
 145 to observe transition dynamics or rewards under alternative action choices. We assume *not* to have
 146 access to the reward function, but we can query preference feedback from a human to generate a
 147 dataset of preferences $\mathcal{D}_{\text{pref}} = \{(\tau_1, \tau_2, o)\}$.

148 **Optimality Criterion.** Based only on our offline dataset $\mathcal{D}_{\text{offline}}$, our goal is to recover a policy
 149 $\hat{\pi}^*$ that minimizes suboptimality in the true environment with as few human preference queries as
 150 possible. Let π_{offline}^* denote the *optimal offline policy* estimated based on the offline data, with access
 151 to the true reward function R , and let ϵ_T denote its suboptimality. Since preference elicitation only
 152 allows us to estimate the reward function, we do not aim to achieve a suboptimality less than ϵ_T .¹
 153 Our objective is then formalized as follows.

154 **Definition 3.1** (Optimality Criterion of Offline Preference Elicitation). *Let π^* be the optimal policy*
 155 *in \mathcal{M} and $\hat{\pi}^*$ be the estimated optimal policy based on an offline dataset $\mathcal{D}_{\text{offline}}$ and $N_p > 0$*
 156 *preference queries. Let ϵ_T be the inherent suboptimality assuming access to the true reward function.*
 157 *We say that a sampling strategy is (ϵ, δ, N_p) -correct if for every $\epsilon \geq \epsilon_T$, with probability at least*
 158 *$(1 - \delta)$, it holds that $V_{T,R}^{\pi^*} - V_{T,R}^{\hat{\pi}^*} \leq \epsilon$.*

159 Our work is the first to formalize this important problem, which faces the challenge of balancing
 160 **exploration** when actively acquiring feedback and **bias mitigation** in learning from offline data.

161 ¹However, ϵ_T is not formally a lower bound for our problem, as shown in Appendix A.3.

Function classes. We estimate the reward function and transition kernel with general function approximation; let \mathcal{F}_R and \mathcal{F}_T denote the classes of functions considered respectively. We also assume a policy class Π . Our theoretical analysis also requires the following assumptions and definitions, which are standard in preference-based RL (Chen et al., 2022; Zhan et al., 2023a).

Assumption 3.1 (Realizability). *The true reward function belongs to the reward class: $R \in \mathcal{F}_R$. The true transition function belongs to the transition class: $T \in \mathcal{F}_T$. The optimal policy belongs to the policy class: $\pi^* \in \Pi$.*

Assumption 3.2 (Boundedness). *The reward function is bounded: $0 \leq \tilde{R}(\tau) \leq R_{\max}$ for all $\tilde{R} \in \mathcal{F}_R$ and all trajectories τ .*

Definition 3.2 (ϵ -bracketing number). *Let \mathcal{F} be a class of real functions $f : \mathcal{X} \rightarrow \mathbb{R}$. We say (l, u) is an ϵ -bracket if $l(x) \leq u(x)$ and $|u(x) - l(x)|_1 \leq \epsilon$ for all $x \in \mathcal{X}$. The ϵ -bracketing number of \mathcal{F} , denoted $\mathcal{N}_{\mathcal{F}}(\epsilon)$, is the minimal number of ϵ -brackets $(l^n, u^n)_{n=1}^N$ needed so that for any $f \in \mathcal{F}$, there is a bracket $i \in [N]$ containing it, meaning $l^i(x) \leq f(x) \leq u^i(x)$ for all $x \in \mathcal{X}$.*

Let $\mathcal{N}_{\mathcal{F}_R}(\epsilon)$ and $\mathcal{N}_{\mathcal{F}_T}(\epsilon)$ denote the ϵ -bracketing numbers of \mathcal{F}_R and \mathcal{F}_T respectively. This measures the complexity of the function classes (Geer, 2000). For instance, with linear rewards $\mathcal{F}_R := \{R : \tau \rightarrow \theta^T \phi(\tau)\}$, the ϵ -bracket number is bounded by $\log \mathcal{N}_{\mathcal{F}_R}(\epsilon) \leq \mathcal{O}(d \log \frac{BG}{\epsilon})$, where $\|\theta\|_2 \leq B$ and $\|\phi(\tau)\|_2 \leq G \forall \tau \in \mathcal{T}$, and d is the dimension of the feature space (Zhan et al., 2023a).

Definition 3.3 (Transition concentrability coefficient, Zhan et al. (2023a)). *The concentrability coefficient w.r.t. transition classes \mathcal{F}_T and the optimal policy π^* is defined as:*

$$C_T(\mathcal{F}_T, \pi^*) = \sup_{\tilde{T} \in \mathcal{F}_T} \left[\frac{\mathbb{E}_{(s,a) \sim d_{\tilde{T}}^{\pi^*}} [|T(\cdot|s,a) - \tilde{T}(\cdot|s,a)|_1]}{\sqrt{\mathbb{E}_{(s,a) \sim \mathcal{D}_{\text{offline}}} [|T(\cdot|s,a) - \tilde{T}(\cdot|s,a)|_1^2]}} \right],$$

The concentrability coefficient measures the coverage of the optimal policy in the offline dataset. Note that C_T is upper-bounded by the density-ratio coefficient: $C_T(\mathcal{F}_T, \pi^*) \leq \sup_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_{\tilde{T}}^{\pi^*}(s,a) / d_{\tilde{T}}^{\pi_\beta}(s,a)$, where π_β is the behavioural policy underlying $\mathcal{D}_{\text{offline}}$.

4 OFFLINE PREFERENCE-BASED RL WITH PREFERENCE ELICITATION

In this section, we propose a general framework for offline preference-based reinforcement learning. The next two sections propose two different preference elicitation strategies. As learning must be carried out in two distinct stages, with environment dynamics based on $\mathcal{D}_{\text{offline}}$ and reward learning on $\mathcal{D}_{\text{pref}}$, we adopt a model-based approach which we summarize in Algorithm 1.

Model Learning. We first leverage the offline data to learn a model of the environment dynamics, fitting a transition model \hat{T} and an uncertainty function u_T through maximum likelihood:

$$\begin{aligned} \hat{T} &= \operatorname{argmax}_{\tilde{T} \in \mathcal{F}_T} \mathbb{E}_{(s,a,s') \sim \mathcal{D}_{\text{offline}}} \left[\log \tilde{T}(s'|s,a) \right], \\ u_T(s,a) &= \max_{\tilde{T}_1, \tilde{T}_2 \in \mathcal{T}} |\tilde{T}_1(\cdot|s,a) - \tilde{T}_2(\cdot|s,a)|_1 \cdot R_{\max}, \end{aligned}$$

where $\mathcal{T} = \{\tilde{T} \in \mathcal{F}_T \mid \mathbb{E}_{(s,a,s') \sim \mathcal{D}_{\text{offline}}} \left[\log \left(\frac{\hat{T}(s'|s,a)}{\tilde{T}(s'|s,a)} \right) \right] \leq \beta_T\}$, defining a confidence set over the maximum likelihood estimator (MLE), and β_T is a margin hyperparameter.

Iterative Preference Elicitation and Reward Learning. As with the transition model, our algorithm estimates the reward function \hat{R} and its uncertainty function through maximum likelihood over iteratively collected preference data $\mathcal{D}_{\text{pref}}$:

$$\begin{aligned} \hat{R} &= \operatorname{argmax}_{\tilde{R} \in \mathcal{F}_R} \mathbb{E}_{(\tau_1, \tau_2, o) \sim \mathcal{D}_{\text{pref}}} [o \log P_{\tilde{R}}(\tau_1 \succ \tau_2) + (1-o) \log P_{\tilde{R}}(\tau_2 \succ \tau_1)], \\ u_R(\tau) &= \max_{\tilde{R}_1, \tilde{R}_2 \in \mathcal{R}} |\tilde{R}_1(\tau) - \tilde{R}_2(\tau)|_1, \end{aligned}$$

where $\mathcal{R} = \{\tilde{R} \in \mathcal{F}_R \mid \mathbb{E}_{(\tau_1, \tau_2, o) \sim \mathcal{D}_{\text{pref}}} \left[\log \left(\frac{P_{\tilde{R}}(\tau_1 \succ \tau_2)}{P_{\tilde{R}}(\tau_1 \succ \tau_2)} \right) \right] \leq \beta_R\}$ defines the confidence set and β_R is a hyperparameter. We also define preference uncertainty between two trajectories τ_1, τ_2 :

$$u_{P_R}(\tau_1, \tau_2) = \max_{\tilde{R}_1, \tilde{R}_2 \in \mathcal{R}} |P_{\tilde{R}_1}(\tau_1 \succ \tau_2) - P_{\tilde{R}_2}(\tau_1 \succ \tau_2)|_1. \quad (2)$$

Algorithm 1 Offline Preference-based Reinforcement Learning with Preference Elicitation**Input:** Observational trajectories dataset $\mathcal{D}_{\text{offline}}$. Significance $\delta \in (0, 1)$, preference budget N_p .**Output:** $\hat{\pi}^*$

- 1: Estimate \hat{T} and u_T via maximum likelihood over the observational data $\mathcal{D}_{\text{offline}}$.
- 2: $\mathcal{D}_{\text{pref}} \leftarrow \emptyset$.
- 3: **for** $k = 1, \dots, N_p$ **do**
- 4: Generate trajectory pairs (τ_1, τ_2) . ▷ **Preference Elicitation:** Sections 5 and 6
- 5: Collect preference label o for (τ_1, τ_2) .
- 6: $\mathcal{D}_{\text{pref}} \leftarrow \mathcal{D}_{\text{pref}} \cup \{(\tau_1, \tau_2, o)\}$.
- 7: Estimate \hat{R} and u_R via maximum likelihood over the preference data $\mathcal{D}_{\text{pref}}$.
- 8: **end for**
- 9: $\hat{\pi}^* \leftarrow \operatorname{argmax}_{\pi \in \Pi} \mathbb{E}_{\tau \sim d_{\hat{T}}^{\pi}} [\hat{R}(\tau) - u_R(\tau) - u_T(\tau)]$

The choice of trajectory sampling strategy for preference elicitation in line 4 (Algorithm 1) is critical to efficiently obtaining an ϵ -optimal policy. We present two possible strategies in Sections 5 and 6.

Pessimistic Policy Optimization. Finally, our algorithm outputs a policy $\hat{\pi}^*$ that is optimal while ensuring robustness to modeling error. This means optimizing for the worst-case value function over the remaining transition and reward uncertainties (Levine et al., 2020):

$$\hat{\pi}^* = \operatorname{argmax}_{\pi \in \Pi} \min_{\hat{T} \in \mathcal{T}, \hat{R} \in \mathcal{R}} V_{\hat{T}, \hat{R}}^{\pi} \quad (3)$$

Based on this objective, we define the pessimistic transition and reward models as follows: $\hat{T}_{\text{inf}}, \hat{R}_{\text{inf}} = \operatorname{argmin}_{\hat{T} \in \mathcal{T}, \hat{R} \in \mathcal{R}} \max_{\pi \in \Pi} V_{\hat{T}, \hat{R}}^{\pi}$. Our analysis provides a worst-case robustness guarantee when considering well-calibrated confidence intervals, as detailed in Sections 5.1 and 6.1. In other words, following prior work (Chen et al., 2022; Zhan et al., 2023a), our theoretical analysis assumes that modeling elements can be identified with no optimization error. We then complement this algorithmic framework with a flexible practical implementation in Section 6.3.

5 PREFERENCE ELICITATION FROM OFFLINE TRAJECTORIES

A first strategy for preference elicitation without environment interaction is to sample trajectories directly from the offline dataset. Shin et al. (2022) propose this approach as Offline Preference-based Reward Learning (OPRL), and design a uniform and uncertainty-sampling variant:

OPRL Uniform Sampling: $\tau_1, \tau_2 \sim \mathcal{D}_{\text{offline}}$

OPRL Uncertainty Sampling: $\tau_1, \tau_2 = \operatorname{argmax}_{\tau_1, \tau_2 \in \mathcal{D}_{\text{offline}}} u_{PR}(\tau_1, \tau_2)$

We provide a theoretical analysis of the performance of OPRL.

5.1 THEORETICAL GUARANTEES.

We obtain the following result, proven in Appendix A.4. The suboptimality of the estimated policy $\hat{\pi}^*$ is bounded by the policy evaluation error for the optimal policy π^* . This error decomposes into a term depending on transition model estimation, and one on reward model estimation.

Theorem 5.1. *For any $\delta \in (0, 1]$, let $\beta_T = c_T^{\text{MLE}} \log(H \mathcal{N}_{\mathcal{F}_T}(1/N_o)/\delta)/N_o$ and $\beta_R = c_R^{\text{MLE}} \log(\mathcal{N}_{\mathcal{F}_R}(1/N_p)/\delta)/N_p$, where $N_o = H|\mathcal{D}_{\text{offline}}|$ is the number of observed transitions in the observational dataset and $c_T^{\text{MLE}}, c_R^{\text{MLE}}$ are universal constants. The policy $\hat{\pi}^*$ estimated by Algorithm 1, with preference elicitation based on offline trajectories, achieves the following suboptimality with probability $1 - \delta$:*

$$V^{\pi^*} - V^{\hat{\pi}^*} \leq \underbrace{H R_{\max} c_T(\mathcal{F}_T, \pi^*) \sqrt{\frac{c_T}{N_o} \log\left(\frac{H}{\delta} \mathcal{N}_{\mathcal{F}_T}\left(\frac{1}{N_o}\right)\right)}}_{\text{transition term } \epsilon_T} + 2\alpha \kappa c_R(\mathcal{F}_R, \pi^*) \underbrace{\sqrt{\frac{c_R}{N_p} \log\left(\frac{1}{\delta} \mathcal{N}_{\mathcal{F}_R}\left(\frac{1}{N_p}\right)\right)}}_{\text{reward term}},$$

where $\alpha = 1$ for uniform sampling or $\alpha \leq 1$ for uncertainty sampling, C_R is a concentrability measure for the reward function, $\kappa = \sup_{r \in [-R_{\max}, R_{\max}]} \frac{1}{\sigma'(r)}$ measures the degree of non-linearity of the link function, and c_T, c_R are universal constants.

In the special case where both the transition and reward functions are learned on a fixed initial preference dataset (no preference elicitation; $|\mathcal{D}_{\text{offline}}| = 2N_p$), we recover Theorem 1 from Zhan et al. (2023a). Importantly, the coefficient α allows us to motivate the superior efficiency of uncertainty sampling over uniform sampling, observed empirically in Shin et al. (2022) and in our own experiments (Section 7). Uncertainty sampling learns accurate reward models with fewer preference queries when $\alpha < 1$, but can perform like uniform sampling in harder problems ($\alpha = 1$).

6 PREFERENCE ELICITATION FROM SIMULATED TRAJECTORIES

We now propose our alternative strategy for generating trajectories for offline preference elicitation: **Simulated Offline Preference-based Reward Learning (Sim-OPRL)**. This method simulates trajectories (τ_1, τ_2) by leveraging the *learned environment model*. This overcomes a limitation of OPRL, which is only designed to reduce uncertainty about the reward functions in \mathcal{R} , by instead reducing uncertainty about which policies are plausibly optimal. Our approach is inspired by efficient online preference elicitation algorithms (Saha et al., 2023; Chen et al., 2022), which we modify for practical implementation. We account for the offline nature of our problem by avoiding regions that are out of the distribution of the data: the sampling strategy is *optimistic* with respect to uncertainty in rewards, but *pessimistic* with respect to uncertainty in transitions.

We summarize our approach to generating simulated trajectories for preference elicitation in Algorithm 2. First, we construct a set of candidate optimal policies Π_{offline} , containing policy π_{offline}^* (optimal policy under the pessimistic model and the true reward function) with high probability – as demonstrated in Appendix A.5.2. Next, within this set of candidate policies, we identify the two most exploratory policies π_1, π_2 , chosen to maximize preference uncertainty u_{P_R} . Finally, we roll out these policies within our learned transition model to generate a trajectory pair (τ_1, τ_2) for preference feedback.

Algorithm 2 Preference Elicitation through Simulated Trajectory Sampling.

Input: Pessimistic transition model \hat{T}_{inf} . Reward confidence set \mathcal{R} and preference uncertainty function u_{P_R} .

Output: (τ_1, τ_2)

- 1: Estimate optimal offline policy set: $\Pi_{\text{offline}} = \{\pi \mid \exists \tilde{R} \in \mathcal{R} : \pi = \operatorname{argmax}_{\pi \in \Pi} \mathbb{E}_{\tau \sim d_{\hat{T}_{\text{inf}}}^{\pi}} [\tilde{R}(\tau)]\}$
 - 2: Identify exploratory policies: $\pi_1, \pi_2 = \operatorname{argmax}_{\pi_1, \pi_2 \in \Pi_{\text{offline}}} \mathbb{E}_{\tau_1 \sim d_{\hat{T}_{\text{inf}}}^{\pi_1}, \tau_2 \sim d_{\hat{T}_{\text{inf}}}^{\pi_2}} [u_{P_R}(\tau_1, \tau_2)]$
 - 3: Rollouts in model: $\tau_1 \sim d_{\hat{T}_{\text{inf}}}^{\pi_1}, \tau_2 \sim d_{\hat{T}_{\text{inf}}}^{\pi_2}$.
-

We first provide a theoretical analysis of the performance of Sim-OPRL, before proposing a practical implementation of our entire preference elicitation and policy optimization algorithm.

6.1 THEORETICAL GUARANTEES

We decompose suboptimality in a similar way to Section 5.1, but obtain a reward suboptimality term that depends on the learned dynamics model instead of the true one, and on π_{offline}^* instead of π^* :

$$V^{\pi^*} - V^{\hat{\pi}^*} \leq \underbrace{(V_{T,R}^{\pi^*} - V_{\hat{T}_{\text{inf}},R}^{\pi^*})}_{\text{transition term } \epsilon_T} + \underbrace{(V_{\hat{T}_{\text{inf}},R}^{\pi_{\text{offline}}^*} - V_{\hat{T}_{\text{inf}},R_{\text{inf}}}^{\pi_{\text{offline}}^*})}_{\text{reward term}}. \quad (4)$$

Analysis of the suboptimality due to transition error is identical to above, but the reward term is thus significantly different. By design, our sampling strategy ensures good coverage of preferences over π_{offline}^* within the learned environment model, which **eliminates the concentrability term for the reward** C_R . We refer the reader to Appendix A.5 for the proof of Theorem 6.1.

Theorem 6.1. *For any $\delta \in (0, 1]$, let $\beta_T = c_T^{\text{MLE}} \log(H\mathcal{N}_{\mathcal{F}_T}(1/N_o)/\delta)/N_o$ and $\beta_R = c_R^{\text{MLE}} \log(\mathcal{N}_{\mathcal{F}_R}(1/N_p)/\delta)/N_p$, where $N_o = H|\mathcal{D}_{\text{offline}}|$ is the number of observed transitions in the observational dataset and $c_T^{\text{MLE}}, c_R^{\text{MLE}}$ are universal constants. The policy $\hat{\pi}^*$ estimated by Algorithm 1, with a preference sampling strategy based on rollouts in the learned transition model,*

324 achieves the following suboptimality with probability $1 - \delta$:

$$325 \quad V^{\pi^*} - V^{\hat{\pi}^*} \leq \underbrace{HR_{\max} C_T(\mathcal{F}_T, \pi^*) \sqrt{\frac{c_T}{N_o} \log\left(\frac{H}{\delta} \mathcal{N}_{\mathcal{F}_T}\left(\frac{1}{N_o}\right)\right)}}_{\text{transition term } \epsilon_T} + \underbrace{2\kappa \sqrt{\frac{c_R}{N_p} \log\left(\frac{1}{\delta} \mathcal{N}_{\mathcal{F}_R}\left(\frac{1}{N_p}\right)\right)}}_{\text{reward term}}. \quad 326 \quad 327 \quad 328$$

329 where $\kappa = \sup_{r \in [-R_{\max}, R_{\max}]} \frac{1}{\sigma'(r)}$ measures the degree of non-linearity of the link function, and
 330 c_T, c_R are universal constants. 331

332 6.2 DISCUSSION

333 Our theoretical results demonstrate that the learned policy can achieve performance comparable to
 334 the optimal policy, and thus satisfy our optimality criterion in Definition 3.1, provided it is covered
 335 by the offline data ($C_T(\mathcal{F}_T, \pi^*), C_R(\mathcal{F}_R, \pi^*) < \infty$). Following our analysis, a suboptimal dataset
 336 requires more preferences to achieve a certain policy performance, as the concentrability terms C_T
 337 or C_R are large. Empirical results in Section 7 confirm that sample efficiency is worse when the
 338 behavioral policy is more suboptimal. 339

340 **Offline Trajectories vs. Simulated Rollouts.** While both OPRL and Sim-OPRL depend on the
 341 offline dataset for estimating environment dynamics, they induce different suboptimality in model-
 342 ing preference feedback. Simulated rollouts are designed to achieve good coverage of the optimal
 343 offline policy π_{offline}^* , which avoids wasting preference budget on trajectories with low rewards or
 344 high transition uncertainty. In contrast, as shown in Zhan et al. (2023a), due to the dependence
 345 of preferences on full trajectories, the reward concentrability term C_R in Theorem 5.1 can be very
 346 large. While sampling from the offline buffer is not sensitive to the quality of the transition model,
 347 good coverage of the optimal policy is needed from both transition and preference data to achieve
 348 low suboptimality. 349

350 **Transition vs. Preference Model Quality.** Our theoretical analysis also suggests an interesting
 351 trade-off in the sample efficiency of our approach, depending on the accuracy of the transition model.
 352 The width of the confidence interval reduces as significance parameter δ or dataset size increase,
 353 or as function class complexity $\mathcal{N}_{\mathcal{F}_T}$ decreases. For a target suboptimality gap ϵ , provided the
 354 optimal offline policy π_{offline}^* has a gap $\epsilon_T < \epsilon$, then the number of preferences required is of the
 355 order of $\mathcal{O}(\log(1/\delta)/(\epsilon - \epsilon_T)^2)$. A more accurate transition model should therefore require fewer
 356 preference samples to achieve a given suboptimality, which we again confirm empirically.

357 6.3 PRACTICAL IMPLEMENTATION

358 We now complete the general algorithmic framework discussed above with a possible implementa-
 359 tion strategy, allowing for empirical validation. In fact, with minor changes to the following frame-
 360 work, our paper also proposes a feasible implementation of related theoretical algorithms (Chen
 361 et al., 2022; Zhan et al., 2023a). We refer the reader to Appendix C for further detail. 362

363 **Model Learning and Policy Optimization.** Following prior work in offline reinforcement learn-
 364 ing (Yu et al., 2020), we train ensembles of N_T and N_R neural network models for the transition
 365 and reward functions on different bootstraps of the data (Lakshminarayanan et al., 2017), denoted
 366 $\{\hat{T}_1, \dots, \hat{T}_{N_T}\}$ and $\{\hat{R}_1, \dots, \hat{R}_{N_R}\}$. We estimate MLE and uncertainty functions as follows:
 367

$$368 \quad \hat{T}(\cdot|s, a) = \frac{1}{N_T} \sum_{i=1}^{N_T} \hat{T}_i(\cdot|s, a); \quad u_T(s, a) = \max_{i, j \in [1, N_T]} |\hat{T}_i(\cdot|s, a) - \hat{T}_j(\cdot|s, a)|_1 \cdot R_{\max} \quad 369 \quad 370$$

$$371 \quad \hat{R}(s, a) = \frac{1}{N_R} \sum_{i=1}^{N_R} \hat{R}_i(s, a); \quad u_R(s, a) = \max_{i, j \in [1, N_R]} |\hat{R}_i(s, a) - \hat{R}_j(s, a)|_1 \quad 372 \quad 373$$

374 Each \hat{R}_i in the ensemble has an associated preference function defined by the Bradley-Terry model,
 375 with σ as the sigmoid function. We obtain preference uncertainty through variation over the en-
 376 semble as in Equation (2). Recall that transition and reward models are trained on $\mathcal{D}_{\text{offline}}$ and $\mathcal{D}_{\text{pref}}$
 377 respectively; for computational efficiency, we sample preferences in batches of to reduce the number
 of reward model updates needed.

We approximate the pessimistic objective in Equation (3) by penalizing the reward function with the uncertainty, as in Lagrangian formulations of model-based offline RL (Yu et al., 2020; Rigter et al., 2022). We solve for the following objective with a traditional reinforcement learning algorithm:

$$\hat{\pi}^* = \operatorname{argmax}_{\pi \in \Pi} \mathbb{E}_{(s,a) \sim d_{\hat{T}}^{\pi}} [\hat{R}(s,a) - \lambda_R u_R(s,a) - \lambda_T u_T(s,a)], \quad (5)$$

where hyperparameters λ_T, λ_R control the degree of conservatism. Note that in our theoretical analysis, this was achieved through parameters β_T, β_R which affect the width of the confidence intervals u_R and u_T , but their exact value cannot be estimated. We show in Appendix A.2 that Equation (5) indeed lower bounds the true value function, under well-calibrated uncertainty estimates.

Near-Optimal Policy Set and Exploratory Policies. Sim-OPRL requires constructing Π_{offline} , a set of near-optimal policies within a pessimistic model of the environment. Following Lindner et al. (2021), we obtain a policy model for each element \hat{R}_i of the reward ensemble. Policy models are optimized to maximize returns under the transition model \hat{T} and the reward function $\hat{R}_i - \lambda_T u_T$, ensuring pessimism w.r.t transitions. Next, the most exploratory policies are identified by generating rollouts of each candidate policy within the learned model \hat{T} . The trajectories (τ_1, τ_2) induced by different policies and maximizing the preference uncertainty function $u_{PR}(\tau_1, \tau_2)$ are used for preference feedback. We refer the reader to Appendix C for further detail.

7 EXPERIMENTAL RESULTS

In this section, we demonstrate the effectiveness of our preference elicitation strategy, Sim-OPRL, across a range of offline reinforcement learning environments and datasets. We demonstrate its **superior performance over OPRL, as expected from our theoretical analysis.**

Since our closest related works do not propose any experimental validation (Chen et al., 2022; Zhan et al., 2023a), we propose a practical implementation of Preference-based Optimistic Planning (PbOP) in Appendix C; this uncertainty-based preference elicitation queries feedback over trajectory rollouts in the *true environment* (Chen et al., 2022). We also compare against OPRL (Shin et al., 2022) with uniform and uncertainty-sampling. Finally, we report the performance of π_{offline}^* and π^* as upper bounds for the performance of our algorithm: the former is trained in the learned transition model with access to the true reward, and the latter has full knowledge of both transition and reward function.

We compare different preference elicitation strategies on a range of environments detailed in Appendix D. Among others, we explore environments from the D4RL benchmark (Fu et al., 2020) identified as particularly challenging offline preference-based reinforcement learning tasks (Shin et al., 2022), as well as a medical simulation designed to model the evolution of patients with sepsis (Oberst and Sontag, 2019). As detailed in Appendix D, these environments consist of high-dimensional state spaces with continuous or discrete action spaces, follow complex transition dynamics, and have sparse or non-linear rewards and termination conditions. This makes them representative of the challenge of learning a reward function and learning offline in a real-world application. In particular, the sepsis simulation environment is commonly used in medically-motivated offline RL work (Tang and Wiens, 2021; Pace et al., 2023), and highlights another advantage of Sim-OPRL over OPRL: it does not require feedback on real trajectories from the observational dataset. In a sensitive setting such as healthcare where data access is carefully controlled, it may be attractive to query experts about *synthetic* trajectories rather than real samples.

Performance against State-of-the-Art. Performance and sample complexity results with different preference elicitation methods are given in Figure 1 and Table 2. Within the offline approaches, Sim-OPRL consistently achieves better environment returns than OPRL with much fewer preference queries. In line with our theoretical analysis, our empirical results therefore demonstrate that policy-based sampling in Sim-OPRL is more efficient than maximizing information gain on the reward function (uncertainty-based OPRL), which echoes similar conclusions reached in prior work on online preference elicitation (Lindner et al., 2021; Chen et al., 2022).

As an upper bound for the performance of our algorithm, we include baselines that have access to the environment: we report the performance of the optimal policy π^* , as well as that of an algorithm querying feedback over optimistic rollouts in the *real* environment (Chen et al., 2022,

Table 2: **Sample complexity N_p under different preference elicitation strategies**, to reach a suboptimality gap of $\epsilon = 20$ over normalized returns. Mean and 95% confidence interval over 6 experiments. The best-performing offline method is highlighted in bold, \times marks when the target suboptimality could not be achieved. Note that PbOP has an advantage by having access to direct interaction with the environment.

Environment	OPRL Uniform	OPRL Uncertainty	Sim-OPRL (Ours)	PbOP (Online)
Star MDP	32 ± 4	30 ± 4	4 ± 2	4 ± 2
Gridworld	105 ± 11	66 ± 7	49 ± 7	32 ± 4
MiniGrid-FourRooms	92 ± 7	53 ± 5	41 ± 5	25 ± 3
HalfCheetah-Random	108 ± 9	71 ± 8	50 ± 10	36 ± 3
Sepsis Simulation	\times	642 ± 72	225 ± 46	75 ± 11

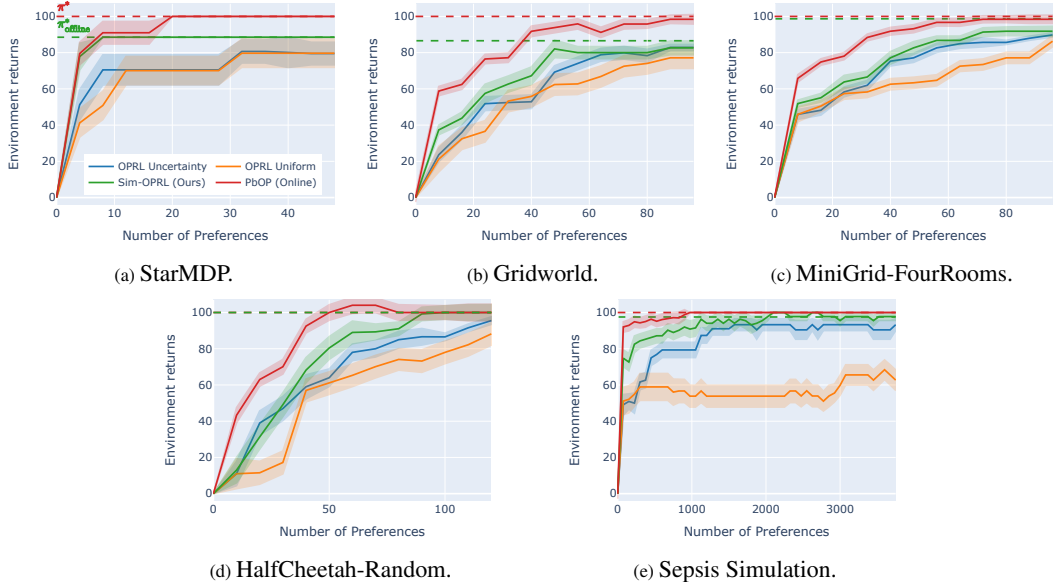


Figure 1: **Environment returns under different preference elicitation strategies**. Mean and 95% confidence interval over 6 experiments. Environment returns are normalized between 0 and 100. Only OPRL and Sim-OPRL are fully offline.

PbOP). In Figure 1, the PbOP method naturally reaches a superior policy with fewer samples as it allows environment interaction and can thus improve its estimate of the transition model in parallel to learning the preference function. As supported by our theoretical analysis, this result stresses the importance of having a high-quality transition model to make our method effective. We explore this in more detail in our following ablations.

Algorithm Ablations. We conduct ablations for our algorithm on a simple tabular MDP, with results in Figure 2. This example (transition and reward details deferred to Appendix D) illustrates the importance of pessimism with respect to the transition model. Even with access to true rewards, π_{offline}^* is pessimistic to avoid the out-of-distribution state, as it is unclear how to reach it. Thus, in Figure 2, we see a drop in performance if pessimism is not applied to the output policy (purple lines). This confirms the theoretical insights from Zhu et al. (2023); Zhan et al. (2023a), who demonstrate the importance of pessimism in offline preference-based RL problems. Pessimism is also crucial in simulated rollouts, to avoid wasting preference budget on regions of low confidence — as value estimates are inaccurate in any case. This is reflected in lower performance without pessimism w.r.t the transition model in Figure 2 (brown line), and which **could be**

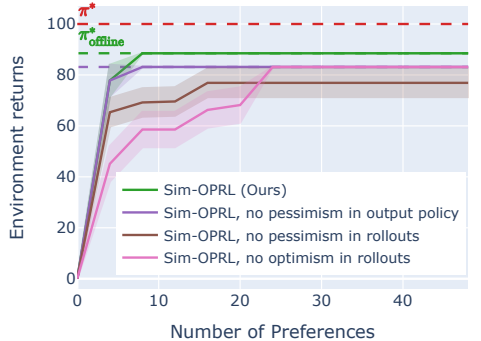


Figure 2: **Algorithm ablations** on the StarMDP.

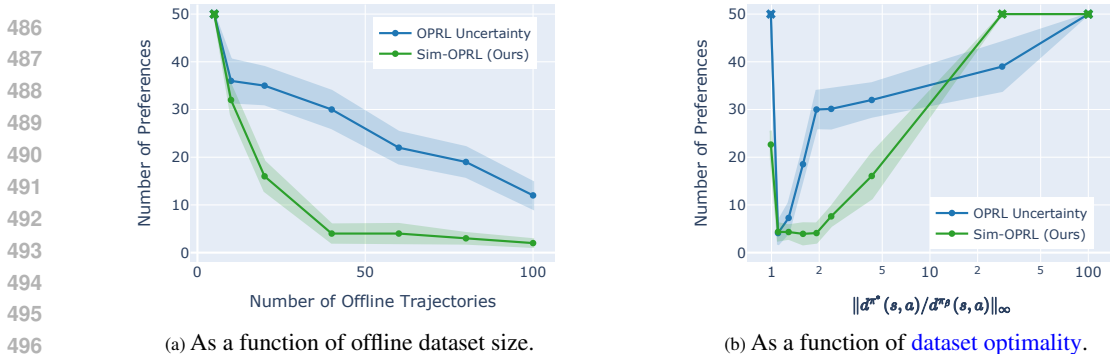


Figure 3: **Preference sample complexity N_p as function of the properties of the observational data**, to reach a suboptimality gap of $\epsilon = 20$ over normalized environment returns (Star MDP). Mean and 95% confidence intervals over 6 experiments. \times marks when the target suboptimality could not be achieved.

seen as the naive adaptation of online preference elicitation methods to our setting (Chen et al., 2022; Lindner et al., 2022). We also note the importance of optimism with respect to the reward uncertainty, both in OPRL in Figure 1 and in our own model-based rollouts in Figure 2.

Transition vs. Preference Model Quality. Next, we empirically study the trade-off between transition and preference model performance in our problem setting. Still in the Star MDP, in the low-data regime, the error ϵ_T incurred in estimating the value function due to the misspecification of the transition model is large. As dictated by our theoretical analysis and as visualized in Figure 3a, this significantly increases the number of preference samples N_p required to achieve good final performance. At the other end of the spectrum, if the offline dataset is large and allows modeling the transition model accurately, then ϵ_T is small and the number of preference samples N_p needed shrinks. We observe a similar trend for both Sim-OPRL and our OPRL uncertainty-sampling baseline.

We also measure how the coverage of the **optimality of the dataset** affects performance in our setting. In Figure 3b, we vary the behavioral policy π_{β} underlying the offline data, ranging from optimal (density ratio coefficient = 1) to highly suboptimal (large density ratio coefficient). The concentrability terms C_T and C_R are challenging to measure as they require considering entire function classes, but we report the accuracy of the maximum likelihood estimate for both models in Appendix E. We observe that preference elicitation methods perform best when the data is close to optimal (with the exception of a fully optimal, non-diverse dataset making reward learning from preferences challenging). More preference samples are required if the observational dataset has poor coverage of the optimal policy (large $C_T(\mathcal{F}_T, \pi^*)$), as the transition and reward models become less accurate for the trajectory distribution of interest. [We also validate this conclusion on HalfCheetah datasets of varying optimality in Appendix E.](#)

8 CONCLUSION

Our work shows the potential of integrating human feedback within the framework of offline RL. We address the challenges of preference elicitation in a fully offline setup by exploring two key methods: sampling from the offline dataset (Shin et al., 2022, OPRL) and generating model rollouts (Sim-OPRL). By employing a pessimistic approach to handle out-of-distribution data and an optimistic strategy to acquire informative preferences, Sim-OPRL balances the need for robustness and informativeness in learning an optimal policy.

We provide theoretical guarantees on the sample complexity of both approaches, demonstrating that performance depends on how well the offline data covers the optimal policy. Empirical evaluations on various environments confirm the practical effectiveness of our algorithm, as Sim-OPRL consistently outperforms OPRL baselines in all settings.

Overall, our approach not only advances the state-of-the-art in offline preference-based RL but also takes a significant step toward improving the practical utility of offline RL. This opens up new avenues for real-world applications of RL in healthcare, robotics, and manufacturing, where interaction with the environment is challenging but domain experts can be queried for feedback. Looking forward, a natural extension will be to explore alternative sources of information from experts, still without direct environment interaction.

REFERENCES

- 540
541
542 Erdem Biyik and Dorsa Sadigh. Batch active preference-based learning of reward functions. In Aude
543 Billard, Anca Dragan, Jan Peters, and Jun Morimoto, editors, *Proceedings of The 2nd Conference*
544 *on Robot Learning*, volume 87 of *Proceedings of Machine Learning Research*, pages 519–528.
545 PMLR, 29–31 Oct 2018. URL [https://proceedings.mlr.press/v87/biyik18a.](https://proceedings.mlr.press/v87/biyik18a.html)
546 [html](https://proceedings.mlr.press/v87/biyik18a.html).
- 547 Erdem Biyik, Malayandi Palan, Nicholas C Landolfi, Dylan P Losey, and Dorsa Sadigh. Asking easy
548 questions: A user-friendly approach to active reward learning. *arXiv preprint arXiv:1910.04365*,
549 2019.
- 550 Ralph Allan Bradley and Milton E Terry. Rank analysis of incomplete block designs: I. the method
551 of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.
- 552
553 Daniel Brown, Russell Coleman, Ravi Srinivasan, and Scott Niekum. Safe imitation learning
554 via fast Bayesian reward inference from preferences. In Hal Daumé III and Aarti Singh, ed-
555 itors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of
556 *Proceedings of Machine Learning Research*, pages 1165–1177. PMLR, 13–18 Jul 2020. URL
557 <https://proceedings.mlr.press/v119/brown20a.html>.
- 558 Xiaoyu Chen, Han Zhong, Zhuoran Yang, Zhaoran Wang, and Liwei Wang. Human-in-the-loop:
559 Provably efficient preference-based reinforcement learning with general function approximation.
560 In *International Conference on Machine Learning*, pages 3773–3793. PMLR, 2022.
- 561 Yifang Chen, Shuohang Wang, Ziyi Yang, Hiteshi Sharma, Nikos Karampatziakis, Donghan Yu,
562 Kevin Jamieson, Simon Shaolei Du, and Yelong Shen. Cost-effective proxy reward model con-
563 struction with on-policy and active learning. *arXiv preprint arXiv:2407.02119*, 2024.
- 564
565 Maxime Chevalier-Boisvert, Bolun Dai, Mark Towers, Rodrigo Perez-Vicente, Lucas Willems,
566 Salem Lahlou, Suman Pal, Pablo Samuel Castro, and Jordan Terry. Minigrid & miniworld: Mod-
567 ular & customizable reinforcement learning environments for goal-oriented tasks. In *Advances in*
568 *Neural Information Processing Systems 36, New Orleans, LA, USA*, December 2023.
- 569 Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep
570 reinforcement learning from human preferences. *Advances in neural information processing sys-*
571 *tems*, 30, 2017.
- 572
573 Jonas Degraeve, Federico Felici, Jonas Buchli, Michael Neunert, Brendan Tracey, Francesco
574 Carpanese, Timo Ewalds, Roland Hafner, Abbas Abdolmaleki, Diego de Las Casas, et al. Mag-
575 netic control of tokamak plasmas through deep reinforcement learning. *Nature*, 602(7897):414–
576 419, 2022.
- 577 Steven Diamond and Stephen Boyd. Cvxpy: A python-embedded modeling language for convex
578 optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.
- 579 Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. D4rl: Datasets for deep
580 data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*, 2020.
- 581
582 Sara A Geer. *Empirical Processes in M-estimation*, volume 6. Cambridge university press, 2000.
- 583
584 Omer Gottesman, Fredrik Johansson, Matthieu Komorowski, Aldo Faisal, David Sontag, Finale
585 Doshi-Velez, and Leo Anthony Celi. Guidelines for reinforcement learning in healthcare. *Nature*
586 *medicine*, 25(1):16–18, 2019.
- 587 Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy
588 maximum entropy deep reinforcement learning with a stochastic actor. In *International confer-*
589 *ence on machine learning*, pages 1861–1870. PMLR, 2018.
- 590
591 Joey Hejna and Dorsa Sadigh. Inverse preference learning: Preference-based rl without a reward
592 function. *Advances in Neural Information Processing Systems*, 36, 2024.
- 593
594 Ying Jin, Zhuoran Yang, and Zhaoran Wang. Is pessimism provably efficient for offline rl? In
595 *International Conference on Machine Learning*, pages 5084–5096. PMLR, 2021.

- 594 Alex Kendall, Jeffrey Hawke, David Janz, Przemyslaw Mazur, Daniele Reda, John-Mark Allen,
595 Vinh-Dieu Lam, Alex Bewley, and Amar Shah. Learning to drive in a day. In *2019 international
596 conference on robotics and automation (ICRA)*, pages 8248–8254. IEEE, 2019.
597
- 598 Rahul Kidambi, Aravind Rajeswaran, Praneeth Netrapalli, and Thorsten Joachims. Morel: Model-
599 based offline reinforcement learning. *Advances in neural information processing systems*, 33:
600 21810–21823, 2020.
601
- 602 Changyeon Kim, Jongjin Park, Jinwoo Shin, Honglak Lee, Pieter Abbeel, and Kimin Lee. Pref-
603 erence transformer: Modeling human preferences using transformers for rl. *arXiv preprint
604 arXiv:2303.00957*, 2023.
605
- 606 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint
607 arXiv:1412.6980*, 2014.
608
- 609 Matthieu Komorowski, Leo A Celi, Omar Badawi, Anthony C Gordon, and A Aldo Faisal. The arti-
610 ficial intelligence clinician learns optimal treatment strategies for sepsis in intensive care. *Nature
611 medicine*, 24(11):1716–1720, 2018.
612
- 613 Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive
614 uncertainty estimation using deep ensembles. *Advances in neural information processing systems*,
30, 2017.
615
- 616 Kimin Lee, Laura Smith, and Pieter Abbeel. Pebble: Feedback-efficient interactive rein-
617 forcement learning via relabeling experience and unsupervised pre-training. *arXiv preprint
618 arXiv:2106.05091*, 2021.
619
- 620 Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tuto-
621 rial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.
622
- 623 David Lindner, Matteo Turchetta, Sebastian Tschitschek, Kamil Ciosek, and Andreas Krause. In-
624 formation directed reward learning for reinforcement learning. *Advances in Neural Information
625 Processing Systems*, 34:3850–3862, 2021.
626
- 627 David Lindner, Andreas Krause, and Giorgia Ramponi. Active exploration for inverse reinforcement
628 learning. *Advances in Neural Information Processing Systems*, 35:5843–5853, 2022.
629
- 630 Qinghua Liu, Alan Chung, Csaba Szepesvári, and Chi Jin. When is partially observable reinforce-
631 ment learning not scary? In *Conference on Learning Theory*, pages 5175–5220. PMLR, 2022a.
632
- 633 Runze Liu, Fengshuo Bai, Yali Du, and Yaodong Yang. Meta-reward-net: Implicitly differentiable
634 reward learning for preference-based reinforcement learning. *Advances in Neural Information
635 Processing Systems*, 35:22270–22284, 2022b.
636
- 637 Yi Liu, Gaurav Datta, Ellen Novoseller, and Daniel S Brown. Efficient preference-based rein-
638 forcement learning using learned dynamics models. In *2023 IEEE International Conference on
639 Robotics and Automation (ICRA)*, pages 2921–2928. IEEE, 2023.
640
- 641 Luckeciano C Melo, Panagiotis Tigas, Alessandro Abate, and Yarin Gal. Deep bayesian active
642 learning for preference modeling in large language models. *arXiv preprint arXiv:2406.10023*,
643 2024.
644
- 645 Azalia Mirhoseini, Anna Goldie, Mustafa Yazgan, Joe Jiang, Ebrahim Songhori, Shen Wang,
646 Young-Joon Lee, Eric Johnson, Omkar Pathak, Sungmin Bae, et al. Chip placement with deep
647 reinforcement learning. *arXiv preprint arXiv:2004.10746*, 2020.
- Michael Oberst and David Sontag. Counterfactual off-policy evaluation with gumbel-max structural
causal models. In *International Conference on Machine Learning*, pages 4881–4890. PMLR,
2019.

- 648 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
649 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kel-
650 ton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike,
651 and Ryan Lowe. Training language models to follow instructions with human feedback. In
652 S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neu-
653 ral Information Processing Systems*, volume 35, pages 27730–27744. Curran Associates, Inc.,
654 2022. URL [https://proceedings.neurips.cc/paper_files/paper/2022/
655 file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf).
- 656 Alizée Pace, Hugo Yèche, Bernhard Schölkopf, Gunnar Ratsch, and Guy Tennenholtz. Delphic
657 offline reinforcement learning under nonidentifiable hidden confounding. In *The Twelfth Interna-
658 tional Conference on Learning Representations*, 2023.
- 659 Jongjin Park, Younggyo Seo, Jinwoo Shin, Honglak Lee, Pieter Abbeel, and Kimin Lee. Surf:
660 Semi-supervised reward learning with data augmentation for feedback-efficient preference-based
661 reinforcement learning. *arXiv preprint arXiv:2203.10050*, 2022.
- 662 Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah
663 Dormann. Stable-baselines3: Reliable reinforcement learning implementations. *Journal of
664 Machine Learning Research*, 22(268):1–8, 2021. URL [http://jmlr.org/papers/v22/
665 20-1364.html](http://jmlr.org/papers/v22/20-1364.html).
- 666 Aniruddh Raghu, Matthieu Komorowski, Imran Ahmed, Leo Celi, Peter Szolovits, and Marzyeh
667 Ghassemi. Deep reinforcement learning for sepsis treatment. *arXiv preprint arXiv:1711.09602*,
668 2017.
- 669 Marc Rigter, Bruno Lacerda, and Nick Hawes. Rambo-rl: Robust adversarial model-based offline
670 reinforcement learning. *Advances in neural information processing systems*, 35:16082–16097,
671 2022.
- 672 Dorsa Sadigh, Anca D Dragan, Shankar Sastry, and Sanjit A Seshia. Active preference-based learn-
673 ing of reward functions. 2018.
- 674 Aadirupa Saha, Aldo Pacchiano, and Jonathan Lee. Dueling rl: Reinforcement learning with tra-
675 jectory preferences. In *International Conference on Artificial Intelligence and Statistics*, pages
676 6263–6289. PMLR, 2023.
- 677 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy
678 optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- 679 Daniel Shin, Anca Dragan, and Daniel S Brown. Benchmarks and algorithms for offline preference-
680 based reward learning. *Transactions on Machine Learning Research*, 2022.
- 681 Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford,
682 Dario Amodei, and Paul F Christiano. Learning to summarize with human feedback. *Advances
683 in Neural Information Processing Systems*, 33:3008–3021, 2020.
- 684 Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- 685 Shengpu Tang and Jenna Wiens. Model selection for offline reinforcement learning: Practical con-
686 siderations for healthcare settings. In *Machine Learning for Healthcare Conference*, pages 2–35.
687 PMLR, 2021.
- 688 Masatoshi Uehara and Wen Sun. Pessimistic model-based offline reinforcement learning under
689 partial coverage. In *International Conference on Learning Representations*, 2021.
- 690 Nils Wilde, Dana Kulić, and Stephen L Smith. Active preference learning using maximum regret.
691 In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages
692 10952–10959. IEEE, 2020.
- 693 Christian Wirth, Riad Akrou, Gerhard Neumann, and Johannes Fürnkranz. A survey of preference-
694 based reinforcement learning methods. *Journal of Machine Learning Research*, 18(136):1–46,
695 2017.

702 Chao Yu, Jiming Liu, Shamim Nemati, and Guosheng Yin. Reinforcement learning in healthcare:
703 A survey. *ACM Computing Surveys (CSUR)*, 55(1):1–36, 2021.
704

705 Tianhe Yu, Garrett Thomas, Lantao Yu, Stefano Ermon, James Y Zou, Sergey Levine, Chelsea Finn,
706 and Tengyu Ma. Mopo: Model-based offline policy optimization. *Advances in Neural Information*
707 *Processing Systems*, 33:14129–14142, 2020.

708 Yuanzhao Zhai, Yiyang Li, Zijian Gao, Xudong Gong, Kele Xu, Dawei Feng, Ding Bo, and Huaimin
709 Wang. Optimistic model rollouts for pessimistic offline policy optimization. *arXiv preprint*
710 *arXiv:2401.05899*, 2024.
711

712 Wenhao Zhan, Masatoshi Uehara, Nathan Kallus, Jason D Lee, and Wen Sun. Provable offline
713 reinforcement learning with human feedback. In *ICML 2023 Workshop The Many Facets of*
714 *Preference-Based Learning*, 2023a.

715 Wenhao Zhan, Masatoshi Uehara, Wen Sun, and Jason D Lee. How to query human feedback
716 efficiently in rl? In *ICML 2023 Workshop The Many Facets of Preference-Based Learning*,
717 2023b.

718 Banghua Zhu, Michael Jordan, and Jiantao Jiao. Principled reinforcement learning with human feed-
719 back from pairwise or k-wise comparisons. In *International Conference on Machine Learning*,
720 pages 43037–43067. PMLR, 2023.
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

A THEORETICAL DETAILS

This appendix provides proofs for the presented theorems and lemmas. In subsection A.1, we provide details on how we define the maximum likelihood estimators and confidence intervals of the preference and transition models. In subsection A.2 we provide the proof that our uncertainty-penalized objective in Equation (5) lower bounds the true value function and thus forms a valid pessimistic framework. In Appendix A.3, we show that the suboptimality of our offline preference elicitation framework is not lower-bounded by the performance of the optimal offline policy. In Appendix A.4, we provide our proof of theorem 5.1, analyzing the suboptimality of preferences sampled from an offline dataset. Finally, in Appendix A.5, we prove Theorem 6.1, which analyzes the suboptimality of preference sampling over simulated rollouts.

A.1 MAXIMUM LIKELIHOOD AND CONFIDENCE INTERVALS

Let \mathcal{F}_g denote a function class over $\mathcal{X} \rightarrow \Delta_{\mathcal{Y}}$, where \mathcal{X}, \mathcal{Y} are measurable sets, and $g \in \mathcal{F}_g$ denotes a function to be estimated.

Let \hat{g} denote the maximum likelihood estimator (MLE) of g based on a dataset $\mathcal{D} = \{(x^n, y^n)\}_{n=1}^N$: $\hat{g} = \operatorname{argmax}_{\tilde{g} \in \mathcal{F}_g} \mathbb{E}_{(x,y) \sim \mathcal{D}} \log(\tilde{g}(y|x))$. We construct the confidence set around the MLE as follows:

$$\mathcal{G} = \{\tilde{g} \in \mathcal{F}_g \mid \mathbb{E}_{(x,y) \sim \mathcal{D}} \left[\log \frac{\hat{g}(y|x)}{\tilde{g}(y|x)} \right] \leq \beta\}$$

Lemma A.1 (MLE Guarantee, Lemma 1 in Zhan et al. (2023a)). *Let $\delta \in (0, 1]$ and define the event \mathcal{E} that $g \in \mathcal{G}$. If*

$$\beta = \frac{c_{MLE}}{N} \log \left(\frac{1}{\delta} \mathcal{N}_{\mathcal{F}_g} \left(\frac{1}{N} \right) \right),$$

where $c_{MLE} > 0$ is a universal constant, then $P(\mathcal{E}) \geq 1 - \delta/2$.

Proof. The proof follows that of Lemma 1 in Zhan et al. (2023a) and uses Cramér-Chernoff’s method.

Let $\bar{\mathcal{B}}$ be a $1/N$ -bracket of \mathcal{F}_g with $|\bar{\mathcal{B}}|_1 = \mathcal{N}_{\mathcal{F}_g}(1/N)$. Denote the set of all right brackets in $\bar{\mathcal{B}}$ by $\tilde{\mathcal{B}} = \{b : \exists b' \text{ s.t. } [b', b] \in \bar{\mathcal{B}}\}$. For $b \in \tilde{\mathcal{B}}$, we have:

$$\begin{aligned} \mathbb{E} \left[\exp \left(\sum_{n=1}^N \log \frac{b(y^n|x^n)}{g(y^n|x^n)} \right) \right] &= \prod_{n=1}^N \mathbb{E} \left[\exp \left(\log \frac{b(y^n|x^n)}{g(y^n|x^n)} \right) \right] \\ &= \prod_{n=1}^N \mathbb{E} \left[\frac{b(y^n|x^n)}{g(y^n|x^n)} \right] \\ &= \prod_{n=1}^N \mathbb{E} \left[\sum_y b(y|x^n) \right] \\ &\leq (1 + 1/N)^N \leq e. \end{aligned}$$

as samples in \mathcal{D} as i.i.d. We use the Tower property in the third step and the fact that b is a $1/N$ -bracket for \mathcal{F}_g in the fourth: there exists $g' \in \mathcal{F}_g$ such that $\|g(\cdot|x) - b(\cdot|x)\|_1 \leq 1/N$ and thus $\|b(\cdot|x)\|_1 \leq 1 + 1/N$, for all $x \in \mathcal{X}$.

Then by Markov’s inequality, for any $\delta \in (0, 1]$, we have:

$$\begin{aligned} P \left(\sum_{n=1}^N \log \frac{b(y^n|x^n)}{g(y^n|x^n)} > \log(\delta) \right) &\leq \mathbb{E} \left[\exp \left(\sum_{n=1}^N \log \frac{b(y^n|x^n)}{g(y^n|x^n)} \right) \right] \cdot \exp(-\log(1/\delta)) \\ &\leq e\delta. \end{aligned}$$

By union bound, we have for all $b \in \tilde{\mathcal{B}}$,

$$P \left(\sum_{n=1}^N \log \frac{b(y^n|x^n)}{g(y^n|x^n)} > c_{MLE} \log \left(\frac{1}{\delta} \mathcal{N}_{\mathcal{F}_g} \left(\frac{1}{N} \right) \right) \right) \leq \delta/2,$$

where $c_{MLE} > 0$ is a universal constant.

Finally, for all $\tilde{g} \in \mathcal{F}_g$, there exists $b \in \tilde{\mathcal{B}}$ such that $g(\cdot|x) \leq \tilde{g}(\cdot|x)$ for all $x \in \mathcal{X}$. As a result, for all $\tilde{g} \in \mathcal{F}_g$, we have:

$$P \left(\sum_{n=1}^N \log \frac{\tilde{g}(y^n|x^n)}{g(y^n|x^n)} > c_{MLE} \log \left(\frac{1}{\delta} \mathcal{N}_{\mathcal{F}_g} \left(\frac{1}{N} \right) \right) \right) \leq \delta/2.$$

□

Under this event \mathcal{E} , we have $g \in \mathcal{G}$ with probability $1 - \delta/2$. A confidence interval constructed via loglikelihood also incurs a bound on the total variation (TV) distance between g and $\tilde{g} \in \mathcal{G}$:

Lemma A.2 (TV-distance to MLE). *Under the event \mathcal{E} , we have, with probability $1 - \delta$, for all $\tilde{g} \in \mathcal{G}$:*

$$\mathbb{E}_{x \sim \mathcal{D}} [\|g(\cdot|x) - \tilde{g}(\cdot|x)\|_1^2] \leq \frac{c}{N} \log \left(\frac{1}{\delta} \mathcal{N}_{\mathcal{F}_g} \left(\frac{1}{N} \right) \right), \quad (6)$$

where $c > 0$ is a universal constant.

Proof. The proof follows that of Liu et al. (2022a), Proposition 14. □

This guarantees that the true function is within an interval around the MLE estimate with high probability.

We apply these lemmas to our MLE estimates of transition and reward functions in Algorithm 1 to obtain the following guarantees.

Let \mathcal{E}_R denote the event $R \in \mathcal{R}$ and \mathcal{E}_T denote the event $T \in \mathcal{T}$, \mathcal{R} and \mathcal{T} denote the respective confidence sets around the MLE. By Lemma A.1, events \mathcal{E}_R and \mathcal{E}_T have probability $1 - \delta/2$ if we choose $\beta_R = c_R^{\text{MLE}} \log(\mathcal{N}_{\mathcal{F}_R}(1/N_p)/\delta)/N_p$ and $\beta_T = c_T^{\text{MLE}} \log(H\mathcal{N}_{\mathcal{F}_T}(1/N_o)/\delta)/N_o$, where $c_R^{\text{MLE}}, c_T^{\text{MLE}}$ are universal constants.

A.2 MODEL-BASED PESSIMISM AND UNCERTAINTY PENALTIES

Lemma A.3 (Telescoping Lemma). *For any reward model $R \in \mathcal{F}_R$, and any two transition models $T, \hat{T} \in \mathcal{F}_T$:*

$$V_{T,R}^\pi - V_{\hat{T},R}^\pi \leq \mathbb{E}_{\tau \sim d_{\hat{T}}^\pi} \left[\sum_{s_j, a_j \in \tau} \|T(\cdot|s_j, a_j) - \hat{T}(\cdot|s_j, a_j)\|_1 \right] \cdot R_{\max}$$

Proof. The proof follows that of Lemma 4.1 in Yu et al. (2020) or Lemma 4 in Zhan et al. (2023a).

Let W_j be the expected return under policy π , with transition model \hat{T} for the first j steps, then transition model T for the rest of the episode. We have:

$$V_{T,R}^\pi - V_{\hat{T},R}^\pi = \sum_{j=0}^{H-1} W_j - W_{j+1}.$$

Now,

$$\begin{aligned} W_j &= R_j + \mathbb{E}_{s_j, a_j \sim \pi, \hat{T}} \left[\mathbb{E}_{s_{j+1} \sim T(\cdot|s_j, a_j)} [V_{T,R}^\pi(s_{j+1})] \right] \\ W_{j+1} &= R_j + \mathbb{E}_{s_j, a_j \sim \pi, \hat{T}} \left[\mathbb{E}_{s_{j+1} \sim \hat{T}(s_j, a_j)} [V_{T,R}^\pi(s_{j+1})] \right] \end{aligned}$$

where R_j is the expected return of the first j steps taken in \hat{T} . Therefore,

$$\begin{aligned} W_j - W_{j+1} &= \mathbb{E}_{s_j, a_j \sim \pi, \hat{T}} \left[\mathbb{E}_{s_{j+1} \sim T(\cdot|s_j, a_j)} [V_{T,R}^\pi(s_{j+1})] - \mathbb{E}_{s_{j+1} \sim \hat{T}(s_j, a_j)} [V_{T,R}^\pi(s_{j+1})] \right] \\ &\leq \mathbb{E}_{s_j, a_j \sim \pi, \hat{T}} \left[\|T(\cdot|s_j, a_j) - \hat{T}(\cdot|s_j, a_j)\|_1 \cdot R_{\max} \right] \end{aligned}$$

under the boundedness assumption for R . Finally, we have:

$$\begin{aligned}
V_{T,R}^\pi - V_{\hat{T},R}^\pi &= \sum_{j=0}^{H-1} W_j - W_{j+1} \\
&= \sum_{j=0}^{H-1} \mathbb{E}_{s_j, a_j \sim \pi, \hat{T}} \left[\mathbb{E}_{s_{j+1} \sim T(\cdot | s_j, a_j)} [V_{T,R}^\pi(s_{j+1})] - \mathbb{E}_{s_{j+1} \sim \hat{T}(s_j, a_j)} [V_{T,R}^\pi(s_{j+1})] \right] \\
&\leq \sum_{j=0}^{H-1} \mathbb{E}_{s_j, a_j \sim \pi, \hat{T}} \left[\|T(\cdot | s_j, a_j) - \hat{T}(\cdot | s_j, a_j)\|_1 \cdot R_{\max} \right] \\
&= \mathbb{E}_{\tau \sim d_{\hat{T}}^\pi} \left[\sum_{s_j, a_j \in \tau} \|T(\cdot | s_j, a_j) - \hat{T}(\cdot | s_j, a_j)\|_1 \cdot R_{\max} \right]
\end{aligned}$$

□

Lemma A.4 (Pessimistic Transition Model). *Under event \mathcal{E}_T , for all $\pi \in \Pi$, $\tilde{R} \in \mathcal{F}_R$:*

$$V_{\tilde{T}, \tilde{R} - u_T}^\pi \leq V_{T, \tilde{R}}^\pi.$$

Proof.

$$\begin{aligned}
V_{\tilde{T}, \tilde{R}}^\pi &= V_{\tilde{T}, \tilde{R}}^\pi - (V_{\tilde{T}, \tilde{R}}^\pi - V_{T, \tilde{R}}^\pi) \\
&\geq \mathbb{E}_{\tau \sim d_{\tilde{T}}^\pi} [\tilde{R}(\tau)] - \mathbb{E}_{\tau \sim d_{\tilde{T}}^\pi} [u_T(\tau)] \\
&= \mathbb{E}_{\tau \sim d_{\tilde{T}}^\pi} [\tilde{R}(\tau) - u_T(\tau)]
\end{aligned}$$

where we have used the telescoping lemma (Lemma A.3), and where $u_T(\tau) = \sum_{(s,a) \in \tau} u_T(s, a) \geq \sum_{(s,a) \in \tau} \|\hat{T}(\cdot | s, a) - T(\cdot | s, a)\|_1 \cdot R_{\max}$ under event \mathcal{E}_T . □

Lemma A.5 (Pessimistic Reward Model). *Under event \mathcal{E}_R , for all $\pi \in \Pi$, $\tilde{T} \in \mathcal{F}_T$:*

$$V_{\tilde{T}, \hat{R} - u_R}^\pi \leq V_{\tilde{T}, R}^\pi.$$

Proof.

$$\begin{aligned}
V_{\tilde{T}, R}^\pi &= V_{\tilde{T}, \hat{R}}^\pi - (V_{\tilde{T}, \hat{R}}^\pi - V_{\tilde{T}, R}^\pi) \\
&= \mathbb{E}_{\tau \sim d_{\tilde{T}}^\pi} [\hat{R}(\tau)] - \mathbb{E}_{\tau \sim d_{\tilde{T}}^\pi} [\hat{R}(\tau) - R(\tau)] \\
&\geq \mathbb{E}_{\tau \sim d_{\tilde{T}}^\pi} [\hat{R}(\tau) - u_R(\tau)]
\end{aligned}$$

where we have used the fact that $|\hat{R}(\tau) - R(\tau)|_1 \leq \sum_{s,a \in \tau} |\hat{R}(s, a) - R(s, a)|_1 = \sum_{(s,a) \in \tau} u_R(s, a) = u_R(\tau)$ under event \mathcal{E}_R . □

Combining the above two lemmas gives the following result:

Corollary A.1. *Under events \mathcal{E}_T and \mathcal{E}_R , for all $\pi \in \Pi$:*

$$V_{\tilde{T}, \hat{R} - u_T - u_R}^\pi \leq V_{\tilde{T}, R}^\pi.$$

This justifies the overall objective considered in our pessimistic policy optimization procedure in Section 6.3.

A.3 SUBOPTIMALITY LOWER BOUND: A COUNTEREXAMPLE

Let $\pi_{\text{offline}}^* = \operatorname{argmax}_{\pi \in \Pi} \min_{\tilde{T} \in \mathcal{T}} V_{\tilde{T}, R}^\pi$ denote the optimal offline policy, which has access to the ground-truth reward function. In this section, we ask whether its suboptimality $\epsilon_T = V_{T, R}^{\pi^*} - V_{T, R}^{\pi_{\text{offline}}^*}$ is a lower bound for the suboptimality of our learned policy $\hat{\pi}^*$ after preference elicitation.




Figure 4: **Tabular MDP.** The environment starts in state s_0 and has horizon $H = 1$. Transition probabilities from state s_0 are given for the two binary actions a_0, a_1 (which send the agent to the other state with complementary probability).

Counterexample. Consider the MDP illustrated in Figure 4. The value function for the optimal policy π^* , which consists of taking action a_0 , is: $V_{T,R}^{\pi^*} = 0.8 \cdot 1 + 0.2 \cdot 0 = 0.8$.

Now assume the following MLE estimate and uncertainty function for both the transition and reward models:

$$\begin{aligned}
 \hat{T}(s_1|s_0, a_0) &= 0.5; & u_T(s_0, a_0) &= 0.4 \\
 \hat{T}(s_1|s_0, a_1) &= 0.5; & u_T(s_0, a_1) &= 0.1 \\
 \hat{r}(s_1) = \hat{r}(s_2) &= 0.5; & u_R(s_1) = u_R(s_2) &= 0.5
 \end{aligned}$$

Assuming access to the learned transition model and the *true* reward function, we pessimistically estimate the value of both actions:

$$\begin{aligned}
 V_{\hat{T}_{\text{inf}}, R}^{a_0} &= 0.1 \cdot 1 + 0.9 \cdot 0 = 0.1 \\
 V_{\hat{T}_{\text{inf}}, R}^{a_1} &= 0.6 \cdot 0 + 0.4 \cdot 1 = 0.4
 \end{aligned}$$

Thus, we have: $\Pi_{\text{offline}}^*(s_0) = \operatorname{argmax}_a V_{\hat{T}_{\text{inf}}, R}^a = a_1$. The offline policy picks the suboptimal action since the worst-case returns of this action are lower than those estimated for a_0 . Evaluating this policy in the real environment, we get $V_{T,R}^{\pi_{\text{offline}}^*} = 0.6 \cdot 0 + 0.4 \cdot 1 = 0.4$.

We now estimate the optimal policy in the learned transition and reward model. Applying pessimism with respect to both models, we get an equal estimated value of 0 for both actions a_0 and a_1 . If policy optimization converges to $\hat{\pi}^* = a_0$, we achieve optimal performance with $V_{T,R}^{\hat{\pi}^*} = 0.8 > V_{T,R}^{\pi_{\text{offline}}^*}$.

This example demonstrates that $\epsilon_T = V_{T,R}^{\pi^*} - V_{T,R}^{\pi_{\text{offline}}^*} = 0.4$ is not a lower bound for the suboptimality of $\hat{\pi}^*$, as policy $\hat{\pi}^*$ can achieve better performance (even optimal performance in this case) if errors in transition and reward model estimation compensate each other.

A.4 SUBOPTIMALITY OF OPRL: PROOF OF THEOREM 5.1

A.4.1 SUBOPTIMALITY DECOMPOSITION

Recall that $\hat{T}_{\text{inf}}, \hat{R}_{\text{inf}} = \operatorname{argmin}_{\tilde{T} \in \mathcal{T}, \tilde{R} \in \mathcal{R}} V_{\tilde{T}, \tilde{R}}^{\pi}$ denote the pessimistic transition and reward models, such that $\hat{\pi}^* = \operatorname{argmax}_{\pi \in \Pi} V_{\hat{T}_{\text{inf}}, \hat{R}_{\text{inf}}}^{\pi}$. We have:

$$\begin{aligned}
 V^{\pi^*} - V^{\hat{\pi}^*} &= V_{T,R}^{\pi^*} - V_{T,R}^{\hat{\pi}^*} \\
 &= (V_{T,R}^{\pi^*} - V_{\hat{T}_{\text{inf}}, \hat{R}_{\text{inf}}}^{\pi^*}) - (V_{T,R}^{\hat{\pi}^*} - V_{\hat{T}_{\text{inf}}, \hat{R}_{\text{inf}}}^{\hat{\pi}^*}) \\
 &\leq (V_{T,R}^{\pi^*} - V_{\hat{T}_{\text{inf}}, \hat{R}_{\text{inf}}}^{\pi^*}) - (V_{T,R}^{\hat{\pi}^*} - V_{\hat{T}_{\text{inf}}, \hat{R}_{\text{inf}}}^{\hat{\pi}^*}) \\
 &\leq V_{T,R}^{\pi^*} - V_{\hat{T}_{\text{inf}}, \hat{R}_{\text{inf}}}^{\hat{\pi}^*}, \tag{7}
 \end{aligned}$$

where we have first used the optimality of $\hat{\pi}^*$ (stating that $V_{\hat{T}_{\text{inf}}, \hat{R}_{\text{inf}}}^{\hat{\pi}^*} \geq V_{\tilde{T}_{\text{inf}}, \hat{R}_{\text{inf}}}^{\pi}$, for all π) and then the pessimism principle (stating that $V_{\hat{T}_{\text{inf}}, \hat{R}_{\text{inf}}}^{\hat{\pi}^*} \leq V_{\tilde{T}, R}^{\hat{\pi}^*}$).

Finally, we decompose the last term above as follows:

$$V^{\pi^*} - V^{\hat{\pi}^*} \leq \underbrace{(V_{\tilde{T}, \hat{R}_{\text{inf}}}^{\pi^*} - V_{\hat{T}_{\text{inf}}, \hat{R}_{\text{inf}}}^{\pi^*})}_{\text{transition term}} + \underbrace{(V_{\tilde{T}, R}^{\pi^*} - V_{\tilde{T}, \hat{R}_{\text{inf}}}^{\pi^*})}_{\text{reward term}} \quad (8)$$

We further analyze each term in the following sections.

A.4.2 ANALYSIS OF THE TRANSITION TERM

In this section, we now upper bound the transition term defined in Equation (8).

Lemma A.6 (Lemma 4, Zhan et al. (2023a)). *Under the event \mathcal{E}_T , with probability $1 - \delta$, we have for all $\tilde{T} \in \mathcal{T}$, for all $\tilde{R} \in \mathcal{G}_R$, for all π :*

$$\mathbb{E}_{d_T^\pi}[\tilde{R}(\tau)] - \mathbb{E}_{d_T^\pi}[\tilde{R}(\tau)] \leq HR_{\max} C_T(\mathcal{F}_T, \pi) \sqrt{\frac{c_T}{N_o} \log \left(\frac{H}{\delta} \mathcal{N}_{\mathcal{F}_T} \left(\frac{1}{N_o} \right) \right)},$$

where $c_T > 0$ is a constant.

Proof. From the telescoping lemma (Lemma A.3), we have:

$$\begin{aligned} V_{\tilde{T}, \tilde{R}}^\pi - V_{\tilde{T}, \tilde{R}}^\pi &\leq R_{\max} \mathbb{E}_{\tau \sim d_T^\pi} \left[\sum_{s_j, a_j \in \tau} \|T(\cdot | s_j, a_j) - \tilde{T}(\cdot | s_j, a_j)\|_1 \right] \\ &\leq HR_{\max} \mathbb{E}_{(s, a) \sim d_T^\pi} [\|T(\cdot | s, a) - \tilde{T}(\cdot | s, a)\|_1] \\ &\leq HR_{\max} C_T(\mathcal{F}_T, \pi) \sqrt{\mathbb{E}_{(s, a) \sim D_{\text{offline}}} [\|T(\cdot | s, a) - \tilde{T}(\cdot | s, a)\|_1^2]} \end{aligned}$$

Under event \mathcal{E}_T , by Lemma A.2, we have, with probability $1 - \delta$, for all $\tilde{T} \in \mathcal{T}$:

$$\mathbb{E}_{(s, a) \sim D_{\text{offline}}} [\|T(\cdot | s, a) - \tilde{T}(\cdot | s, a)\|_1^2] \leq \frac{1}{N_o} c_T \log \left(\frac{H}{\delta} \mathcal{N}_{\mathcal{F}_T} \left(\frac{1}{N_o} \right) \right)$$

This concludes our proof. \square

A.4.3 ANALYSIS OF THE REWARD TERM

Next, we upper bound the reward term defined in Equation (8).

As in Zhan et al. (2023a), we consider the following value function: $V_{\tilde{T}, R}^\pi = \mathbb{E}_{\tau \sim d_T^\pi} [R(\tau)] - \mathbb{E}_{\tau \sim d_{\text{pref}}} [R(\tau)]$, where d_{pref} is a fixed reference trajectory distribution. This baseline subtraction, which doesn't affect either the optimal policy or the analysis of the transition term, is needed as the approximated confidence set is based on the uncertainty in *preference* between two trajectories, not in the reward of a single one.

Definition A.1 (Preference concentrability coefficient). *The concentrability coefficient w.r.t. reward classes \mathcal{F}_R , a target policy π^* and a reference trajectory distribution d_{pref} is defined as:*

$$C_R(\mathcal{F}_R, \pi^*) = \frac{\mathbb{E}_{\tau_1 \sim d_T^{\pi^*}, \tau_2 \sim d_{\text{pref}}} [u_{P_R}(\tau_1, \tau_2)]}{\mathbb{E}_{\tau_1, \tau_2 \sim D_{\text{offline}}} [u_{P_R}(\tau_1, \tau_2)]}$$

Note that, for the purpose of our analysis, our definition differs from that of Zhan et al. (2023a) who instead consider the max ratio of difference in rewards term: $|R(\tau_1) - R(\tau_2) - \tilde{R}(\tau_1) + \tilde{R}(\tau_2)|_1$ over the entire function class \mathcal{F}_R .

Lemma A.7. *Let trajectories for preference elicitation be sampled uniformly from the offline dataset. Under the event \mathcal{E}_R , with probability $1 - \delta$, we have for all $\tilde{T} \in \mathcal{G}_T$, for all $\tilde{R} \in \mathcal{R}$, for all π :*

$$V_{\tilde{T}, R}^{\pi^*} - V_{\tilde{T}, \hat{R}_{\text{inf}}}^{\pi^*} \leq 2\alpha\kappa C_R(\mathcal{F}_R, \pi) \sqrt{\frac{c_R}{N_p} \log \left(\frac{1}{\delta} \mathcal{N}_{\mathcal{F}_R} \left(\frac{1}{N_p} \right) \right)},$$

where $c_R > 0$ is a constant and $\kappa = \sup_{r \in [-R_{\max}, R_{\max}]} \frac{1}{\sigma'(r)}$ measures the degree of non-linearity of the sigmoid function.

Proof.

$$\begin{aligned} V_{\tilde{T}, R}^{\pi^*} - V_{\tilde{T}, \hat{R}_{\text{inf}}}^{\pi^*} &= \mathbb{E}_{\tau \sim d_{\tilde{T}}^{\pi^*}} [R(\tau)] - \mathbb{E}_{\tau \sim d_{\text{pref}}} [R(\tau)] - \mathbb{E}_{\tau \sim d_{\tilde{T}}^{\pi^*}} [\hat{R}_{\text{inf}}(\tau)] + \mathbb{E}_{\tau \sim d_{\text{pref}}} [\hat{R}_{\text{inf}}(\tau)] \\ &= \mathbb{E}_{\tau_1 \sim d_{\tilde{T}}^{\pi^*}, \tau_2 \sim d_{\text{pref}}} [R(\tau_1) - R(\tau_2)] - (\hat{R}_{\text{inf}}(\tau_1) - \hat{R}_{\text{inf}}(\tau_2)) \\ &\leq \kappa \mathbb{E}_{\tau_1 \sim d_{\tilde{T}}^{\pi^*}, \tau_2 \sim d_{\text{pref}}} [|P_R(\tau_1 \succ \tau_2) - P_{\hat{R}_{\text{inf}}}(\tau_1 \succ \tau_2)|] \\ &\leq \kappa \mathbb{E}_{\tau_1 \sim d_{\tilde{T}}^{\pi^*}, \tau_2 \sim d_{\text{pref}}} [u_{P_R}(\tau_1, \tau_2)] \\ &= \kappa C_R(\mathcal{F}_R, \pi^*) \mathbb{E}_{\tau_1, \tau_2 \sim \mathcal{D}_{\text{offline}}} [u_{P_R}(\tau_1, \tau_2)] \end{aligned} \quad (9)$$

where $\kappa = \sup_{r \in [-R_{\max}, R_{\max}]} \frac{1}{\sigma'(r)}$ measures the degree of non-linearity of the sigmoid function. In the first inequality, we have applied the mean value theorem, under Assumption 3.2. In the second inequality, we have used the definition of uncertainty function u_{P_R} as we know $\hat{R}_{\text{inf}} \in \mathcal{R}$.

Now, under event \mathcal{E}_R , by Lemma A.2, we have, with probability $1 - \delta$ for all $\tilde{R} \in \mathcal{R}$:

$$\mathbb{E}_{(\tau_1, \tau_2) \sim \mathcal{D}_{\text{pref}}} [\|P_R(\tau_1 \succ \tau_2) - P_{\tilde{R}}(\tau_1 \succ \tau_2)\|_1^2] \leq \sqrt{\frac{c_R}{N_p} \log \left(\frac{1}{\delta} \mathcal{N}_{\mathcal{F}_R} \left(\frac{1}{N_p} \right) \right)}, \quad (10)$$

where $c_R > 0$ is a constant. This implies the following upper bound for the preference uncertainty function:

$$\mathbb{E}_{(\tau_1, \tau_2) \sim \mathcal{D}_{\text{pref}}} [u_{P_R}(\tau_1, \tau_2)] \leq 2 \sqrt{\frac{c_R}{N_p} \log \left(\frac{1}{\delta} \mathcal{N}_{\mathcal{F}_R} \left(\frac{1}{N_p} \right) \right)} \quad (11)$$

Under uniform sampling, the distribution of preferences in $\mathcal{D}_{\text{pref}}$ is that of the offline dataset:

$$\mathbb{E}_{(\tau_1, \tau_2) \sim \mathcal{D}_{\text{offline}}} [u_{P_R}(\tau_1, \tau_2)] = \mathbb{E}_{(\tau_1, \tau_2) \sim \mathcal{D}_{\text{pref}}} [u_{P_R}(\tau_1, \tau_2)]$$

Thus,

$$V_{\tilde{T}, R}^{\pi^*} - V_{\tilde{T}, \hat{R}_{\text{inf}}}^{\pi^*} \leq 2\kappa C_R(\mathcal{F}_T, \pi) \sqrt{\frac{c_R}{N_p} \log \left(\frac{1}{\delta} \mathcal{N}_{\mathcal{F}_R} \left(\frac{1}{N_p} \right) \right)}.$$

□

Lemma A.8. *Let trajectories for preference elicitation be sampled through uncertainty sampling from the offline dataset. Under the event \mathcal{E}_R , with probability $1 - \delta$, we have for all $\tilde{T} \in \mathcal{G}_T$, for all $\tilde{R} \in \mathcal{R}$, for all π :*

$$V_{\tilde{T}, R}^{\pi^*} - V_{\tilde{T}, \hat{R}_{\text{inf}}}^{\pi^*} \leq 2\alpha\kappa C_R(\mathcal{F}_T, \pi) \sqrt{\frac{c_R}{N_p} \log \left(\frac{1}{\delta} \mathcal{N}_{\mathcal{F}_R} \left(\frac{1}{N_p} \right) \right)},$$

where $c_R > 0$ is a constant and $\alpha \leq 1$.

Proof. The proof follows closely that of Lemma A.7. We introduce the preference concentrability coefficient defined for a general preference dataset:

$$C'_R(\mathcal{F}_T, \pi^*) = \frac{\mathbb{E}_{\tau_1 \sim d_{\tilde{T}}^{\pi^*}, \tau_2 \sim d_{\text{pref}}} [u_{P_R}(\tau_1, \tau_2)]}{\mathbb{E}_{\tau_1, \tau_2 \sim \mathcal{D}_{\text{pref}}} [u_{P_R}(\tau_1, \tau_2)]}$$

We start from Equation (9):

$$\begin{aligned}
V_{T,R}^{\pi^*} - V_{T,\hat{R}_{\text{inf}}}^{\pi^*} &\leq \kappa \mathbb{E}_{\tau_1 \sim d_T^{\pi^*}, \tau_2 \sim d_{\text{pref}}} [u_{P_R}(\tau_1, \tau_2)] \\
&= \kappa C'_R(\mathcal{F}_R, \pi^*) \mathbb{E}_{\tau_1, \tau_2 \sim \mathcal{D}_{\text{pref}}} [u_{P_R}(\tau_1, \tau_2)] \\
&\leq 2\kappa C'_R(\mathcal{F}_R, \pi^*) \sqrt{\frac{c_R}{N_p} \log \left(\frac{1}{\delta} \mathcal{N}_{\mathcal{F}_R} \left(\frac{1}{N_p} \right) \right)}
\end{aligned}$$

where we have used Equation (11).

Now consider the dataset of uncertainty-sampled preferences $\mathcal{D}_{\text{pref}}$. By definition, we have:

$$\mathbb{E}_{\tau_1, \tau_2 \sim \mathcal{D}_{\text{pref}}} [u_{P_R}(\tau_1, \tau_2)] \geq \mathbb{E}_{\tau_1, \tau_2 \sim \mathcal{D}_{\text{offline}}} [u_{P_R}(\tau_1, \tau_2)]$$

Thus, we have: $C'_R(\mathcal{F}_R, \pi^*) \leq C_R(\mathcal{F}_R, \pi^*)$. In other words, we can write: $C'_R(\mathcal{F}_R, \pi^*) = \alpha C_R(\mathcal{F}_R, \pi^*)$, where $\alpha \leq 1$. This concludes our proof. \square

We now conclude the proof of Theorem 5.1 under events \mathcal{E}_R and \mathcal{E}_T .

From Lemma A.6, we upper bound the transition term:

$$V_{T,\hat{R}_{\text{inf}}}^{\pi^*} - V_{\hat{T}_{\text{inf}},\hat{R}_{\text{inf}}}^{\pi^*} \leq HR_{\max} C_T(\mathcal{F}_T, \pi^*) \sqrt{\frac{c_T}{N_o} \log \left(\frac{H}{\delta} \mathcal{N}_{\mathcal{F}_T} \left(\frac{1}{N_o} \right) \right)}$$

From Lemmas A.7 and A.8, we upper bound the reward term:

$$V_{T,R}^{\pi^*} - V_{T,\hat{R}_{\text{inf}}}^{\pi^*} \leq 2\alpha\kappa C_R(\mathcal{F}_R, \pi^*) \sqrt{\frac{c_R}{N_p} \log \left(\frac{1}{\delta} \mathcal{N}_{\mathcal{F}_R} \left(\frac{1}{N_p} \right) \right)},$$

where $\alpha = 1$ for uniform sampling or $\alpha \leq 1$ for uncertainty sampling.

Combining with Equation (8), we obtain Theorem 5.1.

A.5 SUBOPTIMALITY OF SIM-OPRL: PROOF OF THEOREM 6.1

A.5.1 SUBOPTIMALITY DECOMPOSITION

We decompose the suboptimality slightly differently to Equation (7), introducing the optimal offline policy (optimal in the pessimistic model under the *true* reward function): $\pi_{\text{offline}}^* = \operatorname{argmax}_{\pi \in \Pi} V_{\hat{T}_{\text{inf}},R}^{\pi}$.

$$\begin{aligned}
V^{\pi^*} - V^{\hat{\pi}^*} &= V_{T,R}^{\pi^*} - V_{T,R}^{\hat{\pi}^*} \\
&= (V_{T,R}^{\pi^*} - V_{\hat{T}_{\text{inf}},\hat{R}_{\text{inf}}}^{\pi_{\text{offline}}^*}) - (V_{T,R}^{\hat{\pi}^*} - V_{\hat{T}_{\text{inf}},\hat{R}_{\text{inf}}}^{\pi_{\text{offline}}^*}) \\
&\leq (V_{T,R}^{\pi^*} - V_{\hat{T}_{\text{inf}},\hat{R}_{\text{inf}}}^{\pi_{\text{offline}}^*}) - (V_{T,R}^{\hat{\pi}^*} - V_{\hat{T}_{\text{inf}},\hat{R}_{\text{inf}}}^{\hat{\pi}^*}) \\
&\leq V_{T,R}^{\pi^*} - V_{\hat{T}_{\text{inf}},\hat{R}_{\text{inf}}}^{\pi_{\text{offline}}^*} \\
&= (V_{T,R}^{\pi^*} - V_{\hat{T}_{\text{inf}},R}^{\pi^*}) + (V_{\hat{T}_{\text{inf}},R}^{\pi^*} - V_{\hat{T}_{\text{inf}},\hat{R}_{\text{inf}}}^{\pi_{\text{offline}}^*}) \\
&\leq \underbrace{(V_{T,R}^{\pi^*} - V_{\hat{T}_{\text{inf}},R}^{\pi^*})}_{\text{transition term}} + \underbrace{(V_{\hat{T}_{\text{inf}},R}^{\pi^*} - V_{\hat{T}_{\text{inf}},\hat{R}_{\text{inf}}}^{\pi_{\text{offline}}^*})}_{\text{reward term}} \tag{12}
\end{aligned}$$

where we have followed the same analysis as in Appendix A.4.1 and used the optimality of π_{offline}^* in the last inequality.

The analysis of the transition term is identical to the above (Appendix A.4.2). We analyze the reward term next.

A.5.2 ANALYSIS OF THE REWARD TERM

Lemma A.9 (Optimal Offline Policy In Set). *Let Π_{offline} denote the following set of near-optimal pessimistic policies, under the pessimistic transition model \hat{T}_{inf} and the reward confidence set \mathcal{R} :*

$$\Pi_{\text{offline}} = \{\pi \mid \pi = \operatorname{argmax}_{\pi \in \Pi} \mathbb{E}_{\tau \sim d_{\hat{T}_{\text{inf}}}^{\pi}} [\tilde{R}(\tau)] \forall \tilde{R} \in \mathcal{R}\}$$

Under event \mathcal{E}_R , we have $\pi_{\text{offline}}^* \in \Pi_{\text{offline}}$.

Proof. Recall the definition of π_{offline}^* : $\pi_{\text{offline}}^* = \operatorname{argmax}_{\pi \in \Pi} V_{\hat{T}_{\text{inf}}, R}^{\pi}$. Note that there is no need to consider the preference baseline term in V^{π} when building Π_{offline} since it is independent of the policy. Under event \mathcal{E}_R , we have $R \in \mathcal{R}$. Thus, $\pi_{\text{offline}}^* \in \Pi_{\text{offline}}$. \square

Lemma A.10. *Under event \mathcal{E}_R , we have, with probability $1 - \delta$:*

$$V_{\hat{T}_{\text{inf}}, R}^{\pi_{\text{offline}}^*} - V_{\hat{T}_{\text{inf}}, \hat{R}_{\text{inf}}}^{\pi_{\text{offline}}^*} \leq 2\kappa \sqrt{\frac{c_R}{N_p} \log \left(\frac{1}{\delta} \mathcal{N}_{\mathcal{F}_R} \left(\frac{1}{N_p} \right) \right)}$$

Proof.

$$\begin{aligned} & V_{\hat{T}_{\text{inf}}, R}^{\pi_{\text{offline}}^*} - V_{\hat{T}_{\text{inf}}, \hat{R}_{\text{inf}}}^{\pi_{\text{offline}}^*} \\ &= (V_{\hat{T}_{\text{inf}}, R}^{\pi_{\text{offline}}^*} - V_{\hat{T}_{\text{inf}}, \hat{R}}^{\pi_{\text{offline}}^*}) + (V_{\hat{T}_{\text{inf}}, \hat{R}}^{\pi_{\text{offline}}^*} - V_{\hat{T}_{\text{inf}}, \hat{R}_{\text{inf}}}^{\pi_{\text{offline}}^*}) \\ &= \mathbb{E}_{\tau \sim d_{\hat{T}_{\text{inf}}}^{\pi_{\text{offline}}^*}} [R(\tau)] - \mathbb{E}_{\tau \sim d_{\text{pref}}} [R(\tau)] - \mathbb{E}_{\tau \sim d_{\hat{T}_{\text{inf}}}^{\pi_{\text{offline}}^*}} [\hat{R}_{\text{inf}}(\tau)] + \mathbb{E}_{\tau \sim d_{\text{pref}}} [\hat{R}_{\text{inf}}(\tau)] \\ &= \mathbb{E}_{\tau_1 \sim d_{\hat{T}_{\text{inf}}}^{\pi_{\text{offline}}^*}, \tau_2 \sim d_{\text{pref}}} [R(\tau_1) - R(\tau_2)] - \mathbb{E}_{\tau_1 \sim d_{\hat{T}_{\text{inf}}}^{\pi_{\text{offline}}^*}, \tau_2 \sim d_{\text{pref}}} [\hat{R}_{\text{inf}}(\tau_1) - \hat{R}_{\text{inf}}(\tau_2)] \\ &\leq \kappa \mathbb{E}_{\tau_1 \sim d_{\hat{T}_{\text{inf}}}^{\pi_{\text{offline}}^*}, \tau_2 \sim d_{\text{pref}}} [P_R(\tau_1 \succ \tau_2) - P_{\hat{R}_{\text{inf}}}(\tau_1 \succ \tau_2)], \end{aligned}$$

where $\kappa = \sup_{r \in [-R_{\max}, R_{\max}]} \frac{1}{\sigma'(r)}$ measures the degree of non-linearity of the sigmoid function. We have applied the mean value theorem, under Assumption 3.2.

As $R_{\text{inf}} \in \mathcal{R}$, we have: $P_R(\tau_1 \succ \tau_2) - P_{\hat{R}_{\text{inf}}}(\tau_1 \succ \tau_2) \leq u_{P_R}(\tau_1, \tau_2)$.

Let d_{pref} correspond to the distribution of the preference data, which consists of rollouts from exploratory policies within the learned environment model: $d_{\text{pref}} = d_{\hat{T}_{\text{inf}}}^{\pi_1} / 2 + d_{\hat{T}_{\text{inf}}}^{\pi_2} / 2$. Recall that the near-optimal policy set Π_{offline} includes policy π_{offline}^* (Lemma A.9) and that π_1, π_2 are the two more exploratory policies within this set:

$$\mathbb{E}_{\tau_1 \sim d_{\hat{T}_{\text{inf}}}^{\pi_{\text{offline}}^*}, \tau_2 \sim d_{\text{pref}}} [u_{P_R}(\tau_1, \tau_2)] \leq \max_{\pi_1, \pi_2 \in \Pi_{\text{offline}}} \mathbb{E}_{\tau_1 \sim d_{\hat{T}_{\text{inf}}}^{\pi_1}, \tau_2 \sim d_{\hat{T}_{\text{inf}}}^{\pi_2}} [u_{P_R}(\tau_1, \tau_2)].$$

Now, under event \mathcal{E}_R , by Lemma A.2, we have, with probability $1 - \delta$ for all $\tilde{R} \in \mathcal{R}$:

$$\mathbb{E}_{(\tau_1, \tau_2) \sim \mathcal{D}_{\text{pref}}} [\|P_R(\tau_1 \succ \tau_2) - P_{\tilde{R}}(\tau_1 \succ \tau_2)\|_1^2] \leq \frac{c_R}{N_p} \log \left(\frac{1}{\delta} \mathcal{N}_{\mathcal{F}_R} \left(\frac{1}{N_p} \right) \right),$$

where $c_R > 0$ is a constant. This implies the following upper bound for the preference uncertainty function:

$$\mathbb{E}_{(\tau_1, \tau_2) \sim \mathcal{D}_{\text{pref}}} [u_{P_R}(\tau_1, \tau_2)] \leq 2 \sqrt{\frac{c_R}{N_p} \log \left(\frac{1}{\delta} \mathcal{N}_{\mathcal{F}_R} \left(\frac{1}{N_p} \right) \right)}.$$

Thus, we obtain:

$$V_{\hat{T}_{\text{inf}}, R}^{\pi_{\text{offline}}^*} - V_{\hat{T}_{\text{inf}}, \hat{R}_{\text{inf}}}^{\pi_{\text{offline}}^*} \leq 2\kappa \sqrt{\frac{c_R}{N_p} \log \left(\frac{1}{\delta} \mathcal{N}_{\mathcal{F}_R} \left(\frac{1}{N_p} \right) \right)}.$$

The resulting sample complexity of $\mathcal{O}(\frac{\kappa^2 d}{\epsilon^2})$ matches that of active preference learning within a known environment (Saha et al., 2023; Chen et al., 2022). \square

We now conclude the proof of Theorem 6.1 under events \mathcal{E}_R and \mathcal{E}_T .

From Lemma A.6, we upper bound the transition term:

$$V_{T,R}^{\pi^*} - V_{\hat{T}_{\text{inf}},R}^{\pi^*} \leq HR_{\max} C_T(\mathcal{F}_T, \pi^*) \sqrt{\frac{c_T}{N_o} \log \left(\frac{H}{\delta} \mathcal{N}_{\mathcal{F}_T} \left(\frac{1}{N_o} \right) \right)}.$$

From Lemma A.10, we upper bound the reward term:

$$V_{\hat{T}_{\text{inf}},R}^{\pi^*_{\text{offline}}} - V_{\hat{T}_{\text{inf}},\hat{R}_{\text{inf}}}^{\pi^*_{\text{offline}}} \leq 2\kappa \sqrt{\frac{c_R}{N_p} \log \left(\frac{1}{\delta} \mathcal{N}_{\mathcal{F}_R} \left(\frac{1}{N_p} \right) \right)}.$$

Combining with Equation (12), we obtain Theorem 6.1.

B ADDITIONAL RELATED WORK

Table 3: Comparison of broader related work on preference elicitation and preference-based RL.

Framework	Offline	Policy-based Sampling	Robustness Guarantees	Practical Implementation
PREFERENCE-BASED DEEP RL				
D-REX (Brown et al., 2020)	✓	×	×	✓
PEBBLE (Lee et al., 2021)	×	×	×	✓
MRN (Liu et al., 2022b)	×	×	×	✓
SURF (Park et al., 2022)	×	×	×	✓
PT (Kim et al., 2023)	×	×	×	✓
IPL (Hejna and Sadigh, 2024)	✓	×	×	✓
PREFERENCE ELICITATION				
PbOP (Chen et al., 2022)	×	✓	✓	×
REGIME (Zhan et al., 2023b)	×	✓	✓	×
FREEHAND (Zhan et al., 2023a)	✓	×	✓	×
OPRL (Shin et al., 2022)	✓	×	×	✓
MoP-RL (Liu et al., 2023)	×	×	×	✓
Max. Regret (Wilde et al., 2020)	×	✓	×	✓
Biyik and Sadigh (2018)	×	×	✓	✓
Biyik et al. (2019)	×	×	✓	✓
Sim-OPRL (Ours)	✓	✓	✓	✓

In this section, we position this work against additional related literature. We summarize our analysis in Table 3.

First, we note that significant prior work on preference-based reinforcement learning is centered on improving preference modeling or policy optimization for a fixed (uncertainty-based or uniform) preference elicitation strategy (all methods in the top half of Table 3). This focus is orthogonal to our work, and could certainly be combined with our preference elicitation method for optimal performance.

Next, comparing with other preference elicitation methods in Table 3 (bottom half), we find that almost all works consider a **setting with online environment interaction**, which makes them incompatible with our problem setting. Online environment exploration eliminates the need for pessimism against transition dynamics (Levine et al., 2020), whereas we show that it is essential for both theoretical and empirical performance in the offline setting.

The uncertainty-based preference elicitation strategy adopted in most practical PbRL algorithms is to **minimize uncertainty about preferences**, or to maximize information gain about the reward function. From prior theoretical and empirical work (Chen et al., 2022; Lindner et al., 2021), we know that this strategy is often suboptimal compared to actively reducing the set of candidate optimal policies, and both our theoretical and experimental results confirm this.

Finally, note that ours is the first work to provide both a **theoretical and empirical analysis** of the sample complexity of preference elicitation in the fully offline setting.

Beyond Offline Reinforcement Learning. With our community’s recent surge in interest in preference-based reinforcement learning for its application to language model alignment (Stienon et al., 2020; Ouyang et al., 2022), we foresee many possible applications and extensions of our work beyond offline RL. Sample efficiency in preference collection is of paramount importance in Reinforcement Learning from Human Feedback (RLHF): recent works consider active learning pipelines based on uncertainty sampling to improve efficiency (Chen et al., 2024; Melo et al., 2024), similar to OPRL. The insights from our work suggest that a model-based elicitation strategy like Sim-OPRL, designed to simulate responses from near-optimal policies, might perform even better. We believe this line of future work could prove critical to the language model alignment community.

C IMPLEMENTATION DETAILS

We trained all models on two 64-core AMD processors or a single NVIDIA RTX2080Ti GPU. The total wall-clock time for running all experiments presented in this paper amounted to less than 72 hours.

Transition and Reward Function Training. For all baselines, transition and reward models were implemented as linear classifiers (for the Star MDP), as two-layer perceptions with ReLU activation and hidden layer dimension 32 (Gridworld, Sepsis, MiniGrid environments), or as 5-layer MLPs with ReLU activations and hidden sizes [512, 256, 128, 64, 32] for the HalfCheetah environments. Training was carried out for two or one epochs for the transition and reward models respectively, with the Adam optimizer (Kingma and Ba, 2014) and a learning rate of 10^{-3} .

We provide a more detailed practical algorithm for Sim-OPRL in Algorithm 3. For both our method and baselines relying on uncertainty sets (OPRL and PbOP), we estimated uncertainty sets by training models initialized with different random seeds on different bootstraps of the data (sampling 90% of the data with replacement). We consider ensembles of size $|\mathcal{T}| = |\mathcal{R}| = 5$ for both transition and reward models. Hyperparameters λ_T, λ_R control the degree of pessimism in practice and could be considered equivalent to adjusting margin parameters β_T, β_R in our conceptual algorithm proposed in Section 4. Since the exact values prescribed by our theoretical analysis cannot be estimated, the user must set these parameters themselves. Hyperparameter optimization in offline RL is a challenging problem (Levine et al., 2020); for our experiments, we simply set $\lambda_T = 0.5, \lambda_R = 0.1$ (StarMDP, Gridworld) and $\lambda_T = \lambda_R = 1$ for the Sepsis environment.

Near-Optimal Policy Set and Exploratory Policies. Both Sim-OPRL and PbOP require constructing a set of near-optimal policies within a learned model of the environment. Note that the PbOP algorithm in Chen et al. (2022) proposes to construct the near-optimal policy set by considering all policies that have a preference greater than $1/2$ over *all other policies* in Π , under a transition and preference uncertainty bonus. This is infeasible to estimate in practice; we modified the algorithm to allow for practical implementation. The motivation in building the set of plausibly optimal policies remains the same, but the theoretical guarantees may not hold.

We build Π_{offline} by maintaining a policy model for all $\tilde{R} \in \mathcal{R}$, i.e., each element of the reward ensemble. Policy models are optimized to maximize returns under the transition model \hat{T} and the reward function $\tilde{R} - \lambda_T u_T$ (Sim-OPRL) or $\tilde{R} + \lambda_T u_T$ (PbOP). Next, the most exploratory policies are identified by generating 10 rollouts of each of the candidate policies within the learned (SimOPRL) or true (PbOP) model. The trajectories (τ_1, τ_2) maximizing the preference uncertainty function $u_{P_R}(\tau_1, \tau_2)$ are used for preference feedback. In PbOP, the trajectories are then added to the trajectories buffer and the transition model is retrained for 20 (Star MDP, Gridworld) or 200 steps (Sepsis).

Preference Feedback Collection. Preference labels are provided through the ground-truth reward function associated with every environment. As stated in Section 4, for computational efficiency, we sample preferences in batches of 4 (Star MDP, Gridworld) or 100 (Sepsis) to reduce the number of model updates needed.

Policy Optimization. Policy optimization stages, both in estimating optimal policy sets in Sim-OPRL and PbOP and in outputting final policies, are carried out exactly through linear programming

Algorithm 3 Sim-OPRL: Practical Algorithm**Input:** Observational trajectories dataset $\mathcal{D}_{\text{offline}}$. Hyperparameters λ_T, λ_R .**Output:** $\hat{\pi}^*$

- 1: Train an ensemble
- \mathcal{T}
- of transition models via bootstrapping on the observational data
- $\mathcal{D}_{\text{offline}}$
- :

$$\hat{T}(\cdot|s, a) = \frac{1}{|\mathcal{T}|} \sum_{\tilde{T} \in \mathcal{T}} \tilde{T}(\cdot|s, a); \quad u_T(s, a) = \max_{T_1, T_2 \in \mathcal{T}} |T_1(\cdot|s, a) - T_2(\cdot|s, a)|_1 \cdot R_{\max}$$

- 2:
- $\mathcal{D}_{\text{pref}} \leftarrow \emptyset$
- .

- 3:
- for**
- $k = 1, \dots, N_p$
- do**

- 4: Estimate optimal offline policy set:

$$\Pi_{\text{offline}} = \{\pi \mid \pi = \operatorname{argmax}_{\pi \in \Pi} \mathbb{E}_{(s,a) \sim d_{\pi}^{\pi}} [\tilde{R}(s, a) - \lambda_T u_T(s, a)] \forall \tilde{R} \in \mathcal{R}\}$$

- 5: Identify exploratory policies:
- $\pi_1, \pi_2 = \operatorname{argmax}_{\pi_1, \pi_2 \in \Pi_{\text{offline}}} \mathbb{E}_{\tau_1 \sim d_{\pi_1}^{\pi_1}, \tau_2 \sim d_{\pi_2}^{\pi_2}} [u_{P_R}(\tau_1, \tau_2)]$

- 6: Rollouts in model:
- $\tau_1 \sim d_{\pi_1}^{\pi_1}, \tau_2 \sim d_{\pi_2}^{\pi_2}$
- .

- 7: Collect preference label
- o
- for
- (τ_1, τ_2)
- .

- 8:
- $\mathcal{D}_{\text{pref}} \leftarrow \mathcal{D}_{\text{pref}} \cup \{(\tau_1, \tau_2, o)\}$
- .

- 9: Train an ensemble
- \mathcal{R}
- of reward models via bootstrapping of the preference data
- $\mathcal{D}_{\text{pref}}$
- :

$$\hat{R}(s, a) = \frac{1}{|\mathcal{R}|} \sum_{\tilde{R} \in \mathcal{R}} \tilde{R}(s, a); \quad u_R(s, a) = \max_{R_1, R_2 \in \mathcal{R}} |R_1(s, a) - R_2(s, a)|_1$$

10: **end for**

- 11:
- $\hat{\pi}^* \leftarrow \operatorname{argmax}_{\pi \in \Pi} \mathbb{E}_{(s,a) \sim d_{\pi}^{\pi}} [\hat{R}(s, a) - \lambda_R u_R(s, a) - \lambda_T u_T(s, a)]$

for the Star MDP and Gridworld using `cvxopt` (Diamond and Boyd, 2016), based on code from Lindner et al. (2021), using Proximal Policy Optimization (Schulman et al., 2017) implemented in `stable-baselines3` (Raffin et al., 2021) for the Sepsis and MiniGrid environments, and Soft Actor-Critic (Haarnoja et al., 2018) for HalfCheetah. In the latter case, after every preference collection episode, reward and policy models were trained from the checkpoint of the previous iteration, for only 20 steps to minimize computation.

Baselines and Ablations. We implement both OPRL baselines within our model-based offline preference-based algorithm described in Section 4. Uncertainty sampling is taking the pair with maximum preference uncertainty over 45 pairs for every sample, to reduce the load of computing preference uncertainty over the entire trajectory buffer.

Our ablation study for Figure 2 is conducted as follows. For Sim-OPRL without pessimism in the output policy, we output the policy that maximizes the value function under the MLE estimate of the transition and reward function, \hat{T} and \hat{R} , after preference acquisition. For Sim-OPRL without pessimism in the simulated rollouts, we estimate the optimal policy set Π_{offline} in the MLE estimate of the transition model instead of its pessimistic counterpart. Finally, for Sim-OPRL without optimism in the simulated rollouts, we generate rollouts from any two policies in Π_{offline} instead of the most explorative ones.

D ENVIRONMENT DETAILS

Star MDP. We illustrate the transition dynamics underlying the Star MDP in Figure 5. Transition probabilities are 0.9 for all depicted solid arrows, and leave the state unchanged otherwise. Other actions also keep the state unchanged with probability 1. Episodes have length $H = 3$ and start from s_0 . Unless specified otherwise, the offline dataset $\mathcal{D}_{\text{offline}}$ consists of 40 trajectories which only cover states (s_0, s_1, s_3) and (s_3, s_1, s_2) .

Preferences collected over samples from the offline dataset learn slowly about the negative reward in the bottom state, as it is always included in the sampled trajectories. Instead, simulated rollouts can query a direct comparison between the optimal path and one that includes it. This example illustrates

1350
 1351
 1352
 1353
 1354
 1355
 1356
 1357
 1358
 1359
 1360
 1361
 1362
 1363
 1364
 1365
 1366
 1367
 1368
 1369
 1370
 1371
 1372
 1373
 1374
 1375
 1376
 1377
 1378
 1379
 1380
 1381
 1382
 1383
 1384
 1385
 1386
 1387
 1388
 1389
 1390
 1391
 1392
 1393
 1394
 1395
 1396
 1397
 1398
 1399
 1400
 1401
 1402
 1403

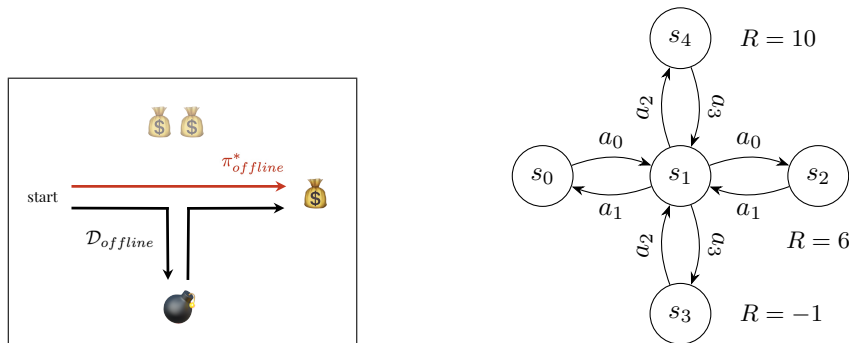


Figure 5: **Star MDP**. Transition probabilities are 0.9 for all solid arrows. Omitted actions or complementary transitions keep the state unchanged.

Start		-1	-1
		-1	-1
	20		
			10

Figure 6: **Gridworld environment**. Rewards at every state are indicated if non-zero. Transition probabilities are 0.9. Thick lines indicate an obstacle, through which state transitions have probability zero.

clearly why querying feedback over simulated rollouts achieves better environment returns than over samples from the offline buffer.

Gridworld. We illustrate the gridworld environment in Figure 6. The environment consists of a 4×4 grid with states associated with different rewards, including a negative-reward region in the top-right corner, a high-reward but unreachable state, and a moderate-reward state at the bottom right corner. Each episode starts in the top-left corner. Transition probabilities for each of the four actions (top, left, bottom, right) are 0.9 for the intended direction, and 0.1 for the others; and action stay remains in the current state with probability 1. Transitions beyond the grid limits or through obstacles have probability zero, with the remainder of the probability mass for each action being distributed amongst other directions equally. The offline dataset contains 150 episodes and the behavioral policy is ϵ -optimal with noise $\epsilon = 0.1$. Episodes have length $H = 10$.

MiniGrid-FourRooms. We also conduct experiments on the minigrid-fourrooms-v0 D4RL dataset (Fu et al., 2020), ignoring all reward information in the offline dataset. In this environment, the agent must navigate a maze consisting of four interconnected rooms and reach a green goal square (Chevalier-Boisvert et al., 2023). Agent and goal squares are randomly placed at the beginning of every episode.

HalfCheetah-Random. The halfcheetah-random-v2 dataset is also part of the D4RL benchmark (Fu et al., 2020). Our choice is motivated by Shin et al. (2022) who identify it as a particularly challenging offline preference-based reinforcement learning task. The dataset consists of 1 million transitions induced by a random policy in the MuJoCo environment Halfcheetah-v2, which rewards agents if they move forward. The observation space is 17-dimensional.

Table 4: **Transition dynamics of the sepsis simulator environment (Oberst and Sontag, 2019).** Table adapted from (Tang and Wiens, 2021).

Step	Variable	Current	New	Change	Variable Affected	Effect
1	Antibiotics	-	on	-	Heart Rate	high \rightarrow normal w.p. 0.5
		-	on	-	Systolic Blood Pressure	high \rightarrow normal w.p. 0.5
		on	off	withdrawn	Heart Rate	normal \rightarrow high w.p. 0.1
		on	off	withdrawn	Systolic Blood Pressure	normal \rightarrow high w.p. 0.5
2	Ventilation	-	on	-	Oxygen Saturation	low \rightarrow normal w.p. 0.7
		on	off	withdrawn	Oxygen Saturation	normal \rightarrow low w.p. 0.1
3	Vasodilators	-	on	-	Systolic Blood Pressure	low \rightarrow normal w.p. 0.7 (non-diabetic)
		-	on	-		normal \rightarrow high w.p. 0.7 (non-diabetic)
		-	on	-	low \rightarrow normal w.p. 0.5 (diabetic)	
		-	on	-	low \rightarrow high w.p. 0.4 (diabetic)	
					normal \rightarrow high w.p. 0.9 (diabetic)	
				Blood Glucose	very low, low \rightarrow normal, normal \rightarrow high, high \rightarrow very high w.p. 0.5 (diabetic)	
		on	off	withdrawn	Systolic Blood Pressure	normal \rightarrow low w.p. 0.1 (non-diabetic)
		on	off	withdrawn		high \rightarrow normal w.p. 0.1 (non-diabetic)
						normal \rightarrow low w.p. 0.05 (diabetic)
						high \rightarrow normal w.p. 0.05 (diabetic)
4	Heart Rate				Vitals spontaneously fluctuate when not affected by treatment (either enabled or withdrawn), the level fluctuates ± 1 w.p. 0.1, except: glucose fluctuates ± 1 w.p. 0.3 (diabetic).	
5	Systolic Blood Pressure					
6	Oxygen Saturation		fluctuate			
7	Blood Glucose		fluctuate			

Sepsis Simulation. The sepsis simulator (Oberst and Sontag, 2019) is a commonly used environment for medically-motivated RL work (Tang and Wiens, 2021). We use the original authors’ publicly available code: <https://github.com/clinicalml/gumbel-max-scm/tree/sim-v2/sepsisSimDiabetes> (MIT license). The state space consists of 1,440 discrete states based on different observational variables (heart rate, blood pressure, oxygen concentration, glucose, diabetes status). The action space consists of three binary treatment options (antibiotic administration, vasopressor administration, and mechanical ventilation). The complex transition dynamics of the environment determine how each treatment affects the value of each vital sign, and are summarized in Table 4; these were designed to reflect patients’ physiology (Oberst and Sontag, 2019). The ground truth reward function is sparse and only assigns a positive reward of +1 to surviving patients and a negative reward of -1 if death occurs (3 or more abnormal vitals) during their stay. The offline trajectories dataset includes 10,000 episodes following an ϵ -optimal policy with noise $\epsilon = 0.1$ and the episode length is $H = 20$.

E ADDITIONAL RESULTS

We include additional results in this section.

In Figure 7, we report the accuracy of the transition and preference model achieved for the Star MDP as we vary the size of optimality of the offline dataset. Accuracy is measured against all possible state transitions and over 100 pairs of random trajectories (random combinations of the 5 states and 4 actions in a sequence of $H = 3$). This complements our analysis in Section 7 and fig. 3. We see a steady improvement in both transition and reward model quality as we increase the amount of observational data in Figure 7a, which explains the observed dependence of N_p on N_o in Figure 3a.

In Figure 7b, we notice low model performance at both extremes of the x-axis. When the dataset is fully optimal, we find that all trajectories involve the same sequence of actions and states, so learning a transition or reward model from this data is challenging. We reach a similar conclusion at the other end of the spectrum at high density ratios, where the coverage the optimal states reduces. We reach highest performance for both models at intermediate values, when diversity of the observational data is high.

Still, it is important to stress that the highest accuracy of both models does not necessarily translate to the best-performing policy: good performance on the distribution induced by the optimal policy is more important, as formalized by the concentrability coefficients.

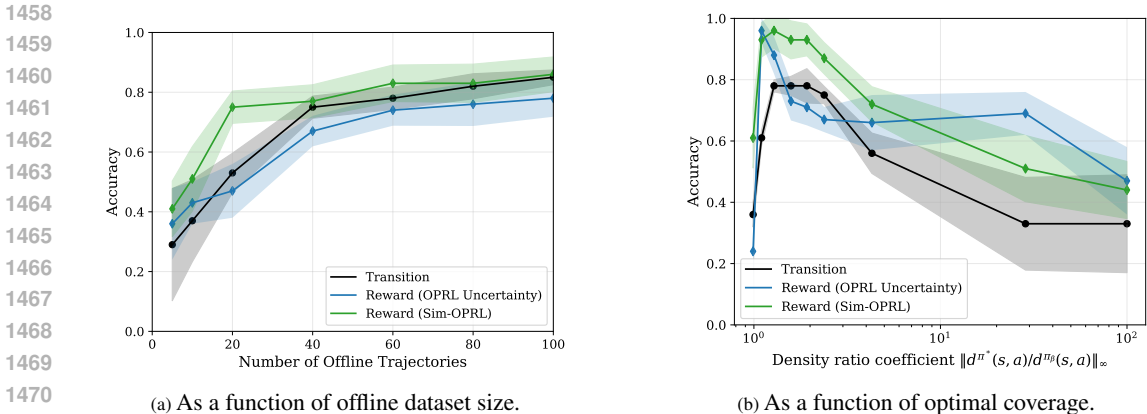


Figure 7: **Transition and preference model accuracy as function of the properties of the observational data (Star MDP).** Preference elicitation is carried out until 10 preferences are queried. Mean and 95% confidence intervals over 20 experiments. Note that the transition model is the same for the two methods, as they have access to the same dataset.

Table 5: **Sample complexity N_p with Sim-OPRL as a function of dataset optimality,** to reach a suboptimality gap of $\epsilon = 20$ over normalized returns for the D4RL HalfCheetah environment. Mean and 95% confidence interval over 6 experiments.

Offline Dataset	OPRL Uniform	OPRL Uncertainty	Sim-OPRL
HalfCheetah-Random	108 ± 9	71 ± 8	50 ± 10
HalfCheetah-Medium	63 ± 8	49 ± 6	36 ± 8
HalfCheetah-Medium-Expert	47 ± 8	45 ± 6	30 ± 8

Table 6: **Policy performance for a fixed preference budget ($N_p = 30$) for different D4RL HalfCheetah datasets.** Normalized environment returns, mean and 95% confidence interval over 6 experiments.

Offline Dataset	OPRL Uniform	OPRL Uncertainty	Sim-OPRL
HalfCheetah-Random	17 ± 7	47 ± 10	49 ± 9
Halfcheetah-Medium	43 ± 6	63 ± 6	69 ± 7
Halfcheetah-Medium-Expert	57 ± 6	72 ± 6	80 ± 6

To complement the analysis with a more complex environment, we also ran different preference elicitation algorithms on D4RL HalfCheetah datasets of varying optimality. In Table 5, we report the number of preference samples needed to achieve a target suboptimality with these different datasets. We reach the same conclusion as above: fewer preferences are needed as the dataset becomes more optimal.

In Table 6, we report policy performance for a fixed preference budget. With a fixed preference budget, policies learned from suboptimal observational datasets achieve lower returns, due to the worse quality of the transition model. Shin et al. (2022) reach similar conclusions when evaluating policy performance under different OPRL methods (in their Table 2).

F BROADER IMPACT

Better preference elicitation strategies for offline reinforcement learning have the potential to facilitate and improve decision-making in real-world safety-critical domains like healthcare or economics, by reducing reliance on direct environment interaction and reducing human effort in providing feedback. Potential downsides could include the amplification of biases in the offline data, potentially leading to suboptimal or unfair policies. Thorough evaluation is therefore crucial to mitigate this before deploying models in such real-world applications. In addition, human preferences may not

1512 be fully captured by binary comparisons. As noted in our conclusion, we hope that future work will
1513 explore richer feedback mechanisms to better model complex decision-making objectives.
1514

1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565