Random Policy Enables In-Context Reinforcement Learning within Trust Horizon

Weiqin Chen[♠], **Santiago Paternain**[♠]

[♠]Department of Electrical, Computer, and Systems Engineering, Rensselaer Polytechnic Institute

Abstract

Pretrained foundation models (FMs) have exhibited extraordinary in-context learning performance, allowing zero-shot (or few-shot) generalization to new environments/tasks not encountered during the pretraining. In the case of reinforcement learning (RL), in-context RL (ICRL) emerges when pretraining FMs on decision-making problems in an autoregressive-supervised manner. Nevertheless, the current state-of-the-art ICRL algorithms, such as Algorithm Distillation, Decision Pretrained Transformer and Decision Importance Transformer, impose stringent requirements on the pretraining dataset concerning the behavior (source) policies, context information, and action labels, etc. Notably, these algorithms either demand optimal policies or require varying degrees of well-trained behavior policies for all pretraining environments. This significantly hinders the application of ICRL to real-world scenarios, where acquiring optimal or well-trained policies for a substantial volume of real-world training environments can be prohibitively expensive or even intractable. To overcome this challenge, we introduce a novel approach, termed State-Action Distillation (SAD), that allows to generate an effective pretraining dataset guided solely by random policies. In particular, SAD selects query states and corresponding action labels by distilling the outstanding state-action pairs from the entire state and action spaces by using random policies within a trust horizon, and then inherits the classical autoregressive-supervised mechanism during the pretraining. To the best of our knowledge, this is the first work that enables effective ICRL under (e.g., uniform) random policies and random contexts. We also establish the quantitative analysis of the trustworthiness as well as the performance guarantees of our SAD approach. Moreover, our empirical results across multiple popular ICRL benchmark environments demonstrate that, on average, SAD outperforms the best baseline by 236.3% in the offline evaluation and by 135.2% in the online evaluation.

Introduction

Pretrained foundation models (FMs) have demonstrated promising performance across a wide variety of domains in artificial intelligence including natural language processing (NLP) (Devlin 2018; Radford 2018; Radford et al. 2019; Brown 2020), computer vision (CV) (Yuan et al. 2021; Sammani, Mukherjee, and Deligiannis 2022; Ma et al. 2023; Chen et al. 2024), and sequential decision-making (Chen et al. 2021; Janner, Li, and Levine 2021; Xu et al. 2022b; Yang et al. 2023; Light et al. 2024a,b). This success is attributed to FMs' impressive capability of in-context learning (Dong et al. 2022; Li et al. 2023; Wei et al. 2023; Wies, Levine, and Shashua 2024) which refers to the ability to infer and understand the new tasks provided with the context information (or prompt) and without model parameters updates. Recently, in-context reinforcement learning (ICRL) (Laskin et al. 2022; Grigsby, Fan, and Zhu 2023; Lin, Bai, and Mei 2023; Sinii et al. 2023; Zisman et al. 2023; Lee et al. 2024; Lu et al. 2024; Wang et al. 2024; Dong et al. 2024) has emerged when FMs are pretrained on sequential decision-making problems. Whereas FMs use texts as the context/prompt in NLP, ICRL treats the state-action-reward tuples as the contextual information for decision-making.

However, the current state-of-the-art (SOTA) ICRL algorithms impose strict requirements on the pretraining datasets. More specifically, Algorithm Distillation (AD) (Laskin et al. 2022) requires the context to contain the complete learning history (from the initial policy to the final-trained policy) of the source (or behavior) RL algorithm for all pretraining environments. In addition, AD requires environments to have short episodes, allowing the context to capture cross-episodic information. This enables AD to learn the improvement operator of the source RL algorithm. Conversely, Decision Pretrained Transformer (DPT) (Lee et al. 2024) partially relaxes the requirement on the context, permitting it to be gathered by random policies and without needing to adhere to the transition dynamics. Nevertheless, DPT necessitates the optimal policy to label an optimal action for any randomly sampled query state across all pretraining environments. To explore the feasibility of ICRL in the absence of optimal policies, Decision Importance Transformer (DIT) (Dong et al. 2024) proposes to leverage the observed state-action pairs in the context data as query states and corresponding action labels. Each state-action pair within the context is assigned a weight in the training process. This weight is proportional to the pseudo-return of the pair. Thus, DIT prioritizes the training on high-pseudo-return pairs. Despite not demanding optimal policies, DIT still requires a substantial context dataset to comprehensively cover all state-action pairs from the state and action spaces. Furthermore, DIT mandates that more than 30% of the transition data in the context be well-trained, and the context should originate from a complete episode.

Notably, acquiring either optimal policies or well-trained policies across a multitude of pretraining environments in real-world scenarios can be prohibitively expensive or even intractable. On the other hand, the transition data available in real-world problems like healthcare (Fatemi et al. 2022; Tang and Wiens 2021)—collected as the context—may not originate from a complete episode. These stringent requirements on the pretraining dataset of the SOTA ICRL algorithms severely limit their practical applications to the real world, especially for those where the transition data exhibits high variance to train an effective policy, and a naive random policy such as the uniform policy becomes a reasonable choice at hand. Consequently, this paper centers on the ICRL that operates without the need for optimal (or any degree of well-trained) policies or episodic context, placing its emphasis on the scenarios under (e.g., uniform) random policies and random contexts only. More background knowledge can be found in the Related Work section (see Appendix A).

Main Contributions

The main contributions of this work are summarized as follows.

- We propose a novel approach termed State-Action Distillation (SAD) to generate the pretraining dataset of ICRL under random policies. Notably, SAD distills the outstanding state-action pairs over the entire state and action spaces for the query states and corresponding action labels (refer to Figure 1), by executing all possible actions under the random policies within a trust horizon.
- To the best of our knowledge, SAD stands as the first method that enables effective ICRL under (e.g., uniform) random policies and random contexts.
- We establish the quantitative analysis of the trustworthiness as well as the performance guarantees of our SAD approach. We substantiate the efficacy of SAD by empirical results on several popular ICRL benchmark environments. On average, SAD significantly outperforms all existing SOTA ICRL algorithms. More concretely, SAD surpasses the best baseline by 236.3% in the offline evaluation and by 135.2% in the online evaluation.

In-Context Reinforcement Learning

This section introduces the background of ICRL mechanisms and three SOTA ICRL algorithms. We start by presenting the preliminaries of ICRL.

Preliminaries

RL problems are generally formulated as Markov Decision Processes (MDPs) (Sutton 2018). An MDP can be represented by a tuple $\tau = (S, A, R, P, \rho)$, where S and A denote finite state and action spaces, $R : S \times A \to \mathbb{R}$ denotes the reward function that evaluates the quality of the decision (action), $P : S \times A \times S \to [0, 1]$ denotes the transition probability that describes the dynamics of the system, and $\rho : S \to [0, 1]$ denotes the initial state distribution.

A policy π defines a mapping from states to probability distributions over actions, providing a strategy that guides the agent in the decision making. The agent interacts with the environment following the policy π and the transition dynamics of the system, and then generates an episode of the transition data $(s_0, a_0, r_0, \cdots)$. The performance measure $J(\pi)$ is defined by the expected discounted cumulative reward under the policy π

$$J(\pi) = \mathbb{E}_{s_0 \sim \rho, a_t \sim \pi(\cdot|s_t), s_{t+1} \sim P(\cdot|s_t, a_t)} \left[\sum_{t=0}^{\infty} \gamma^t r_t \right]. \quad (1)$$

The goal of RL is to find an optimal policy π^* that maximizes $J(\pi)$. It is crucial to recognize that π^* often varies across different MDPs (environments). Thus, the optimal policy for standard RL must be re-learned each time a new environment is encountered. Under this circumstance, ICRL proposes to pretrain a FM on a wide variety of pretraining environments, and then deploy it in the *unseen* test environments without updating parameters in the pretrained model, i.e., zero-shot generalization (Sohn, Oh, and Lee 2018; Mazoure et al. 2022; Zisselman et al. 2023; Kirk et al. 2023).

Supervised Pretraining Mechanism

In this subsection, we introduce the methodology behind ICRL–a supervised pretraining mechanism. Consider two distributions over environments $\mathcal{T}_{\text{train}}$ and $\mathcal{T}_{\text{test}}$ for pretraining and test (evaluation), respectively. Each environment, along with its corresponding MDP τ, can be regarded as an instance drawn from the environment distributions, where each environment may exhibit distinct reward functions and transition dynamics. Given an environment τ, a context/prompt $\mathcal{C} = \{s_i, a_i, r_i, s_i'\}_{i \in [n]}$ refers to a collection of interactions between the agent and the environment τ, sampled from a pretraining context distribution $D_{\text{train}}(\cdot \mid \tau)$, i.e., $\mathcal{C} \sim D_{\text{train}}(\cdot \mid \tau)$. Notably, $D_{\text{train}}(\cdot \mid \tau)$ contains the contextual information regarding the environment τ. We next consider a query state distribution D_q^τ and a label policy that maps the query state to the action label, i.e., $\pi_l : S \to \Delta(a_l)$. Then, the joint distribution over the environment τ, context $\mathcal{C}$, query state s_q, and action label a_l is given by

$$P_{\text{train}}(\tau, \mathcal{C}, s_q, a_l) = \mathcal{T}_{\text{train}}(\tau) \cdot D_{\text{train}}(\mathcal{C}|\tau) \cdot D_q^\tau \cdot \pi_l(a_l|s_q). \quad (2)$$

ICRL follows a supervised pretraining mechanism. More concretely, a FM with parameter θ (denoted by $\mathcal{F}_\theta : \mathbb{C} \times S \to \Delta(a_l)$) is pretrained to predict the action label a_l given the context $\mathcal{C}$ and query state s_q. To do so, the current literature (Laskin et al. 2022; Lee et al. 2024; Dong et al. 2024) often considers the following objective function

$$\theta^* = \operatorname*{argmin}_\theta \mathbb{E}_{P_{\text{train}}} \left[l\left(\mathcal{F}_\theta(\cdot \mid \mathcal{C}, s_q), a_l\right) \right], \quad (3)$$

where $l(\cdot, \cdot)$ represents the loss function, for instance, negative log-likelihood (NLL) for discrete-action problems and mean square error (MSE) for continuous-action scenarios. We note that while the current SOTA ICRL algorithms (AD, DPT, and DIT) adhere to a common objective function (3), they differ significantly in constructing the context, query state, and action label.

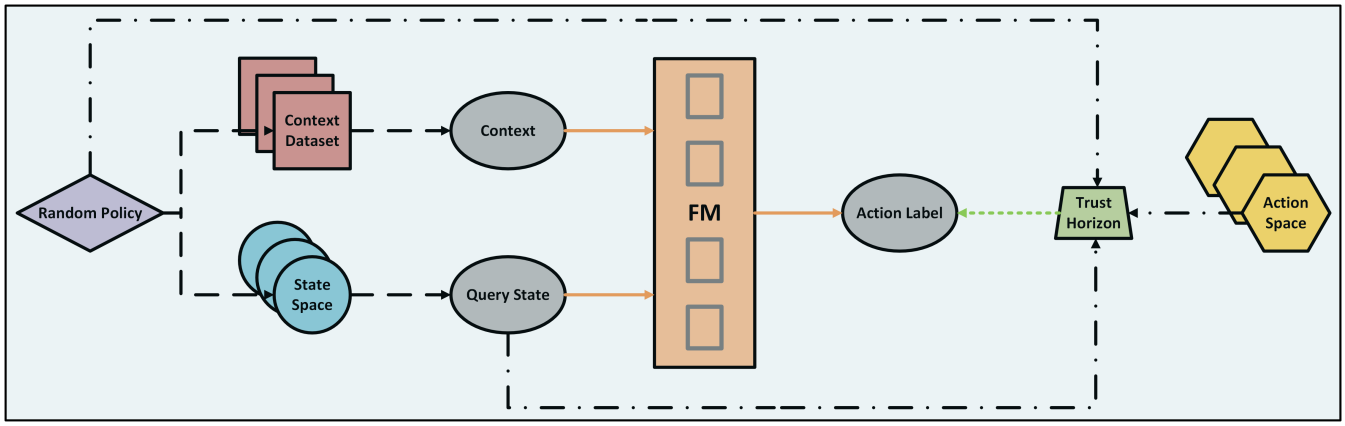

Figure 1: Schematic of the State-Action Distillation (SAD) approach. SAD first collects the context and query state using a random policy, then distills the action label by exploring the action space using the random policy and query state within a trust horizon. The collected contexts, query states, and action labels form the pretraining dataset, which can be employed to pretrain a foundation model for in-context reinforcement learning in a supervised pretraining mechanism.

Algorithm Distillation. Instead of learning an optimal policy for a specific environment, AD proposes to learn a RL algorithm itself across a wide range of environments. This is, the FM in AD is pretrained to imitate the source (or behavior) algorithm over the pretraining environment distribution $\mathcal{T}_{\text{train}}$. In general, AD demands a well-trained source algorithm with its complete learning history (from the initial policy to the final-trained policy). Additionally, AD is restricted to the environments with short episode length by which the context must capture cross-episode information of the source algorithm. In terms of the objective function (3), AD takes the state at the time step t as the query state, a_t from the source algorithm as the action label, and the episodic history data $(s_0, a_0, r_0, \cdots, s_{t-1}, a_{t-1}, r_{t-1})$ to be the context.

Decision Pretrained Transformer. Instead of being stringent on the pretraining context dataset and the environment itself, DPT handles the context and query state in a more general manner. Specifically, DPT considers an entirely random collection of transitions as the context, a random query state s_q drawn from D_q, and an optimal action label corresponding to s_q. Despite less requirements on the context, DPT necessitates access to optimal policies for the optimal action labels in all pretraining environments, which may not be available in real-world applications.

Decision Importance Transformer. DIT proposes to learn ICRL without optimal action labels, following the same supervised pretraining mechanism as DPT. To that end, DIT leverages every possible state and action in the context as the query state and corresponding action label. It is important to point out that DIT requires a (partially) complete episode to form the context, enabling the computation of a pseudo-return for each state-action pair. By mapping the pseudo-return to a weight that reflects the quality of each state-action pair, DIT can pretrain the FM using the DPT structure, augmented by the weight assigned to each query state and action label. In other words, DIT prioritizes the training on high-pseudo-return pairs. Last but not least, DIT still mandates that more than 30% of the context data comes from well-trained policies.

To summarize, it is worth highlighting that all these SOTA ICRL algorithms necessitate varying degrees of well-trained, or even optimal policies during the pretraining phase. However, obtaining optimal or well-trained policies for real-world applications is often prohibitive, as it demands extensive training across a vast number of real-world environments. This challenge becomes even more pronounced in the domains like healthcare (Fatemi et al. 2022), where the transition data exhibits high variance to train an effective policy, and the random policy becomes a reasonable choice at hand. Therefore, executing ICRL under (e.g., uniform) random policies and random contexts is crucial for enabling the practical application of ICRL in the real-world.

State-Action Distillation

In this section, we propose the State-Action Distillation (SAD), an approach for generating the pretraining dataset for ICRL under random policies and random contexts (see Figure 1). We summarize the implementation details of SAD in Algorithms 1-6.

As indicated in (2), the pretraining data consists of the context, the query state and the action label. We start by introducing the generation of the context under random policies in SAD (refer to Algorithm 1). It is important to highlight that the context is derived through interactions with the environment under any given random policy (e.g., uniform random policy), and notably, the context not necessarily originates from a complete episode. These benefits make SAD potentially well-suited for ICRL's real-world applications with random transition data only.

Having collected the random context, we are now in the stage of collecting query states and corresponding action labels for the pretraining of FM under the random policy.

Algorithm 1: Collecting Contexts under Random Policy

Require: Random policy π, context horizon length T, state space S, environment τ, empty context $\mathcal{C} = \emptyset$
for t in $[T]$ **do**
  Sample a state $s \sim S$ and an action $a \sim \pi(\cdot|s)$
  Collect (r, s') by executing action a in the environment τ
  Add (s, a, r, s') to $\mathcal{C}$
end for
Return $\mathcal{C}$

We proceed by recalling that DIT prioritizes the training on high-pseudo-return pairs from the context data collected. To that end, DIT assigns a weight w to the loss function during the pretraining phase that is proportional to the pseudo-return, i.e., $w(s_t, a_t) \propto \sum_{t'=t}^{T} \gamma^{t'-t} r_{t'}$. Nonetheless, we acknowledge that DIT may not explore to train on good state-action pairs under the random policy for two reasons: (i) DIT solely considers to train on the state-action pairs that are observed in the context, which contains limited transition data and is derived by the random policy. (ii) Even for the state-action pairs in the collected context, the pseudo-return does not necessarily prioritize the optimal pair but rather promotes the pair with high immediate reward, as the discount factor applies starting from the current time step with a horizon of $(T - t + 1)$ only, instead of T.

Under this circumstance, our SAD approach advocates for distilling the outstanding query states and action labels by searching across the entire state and action spaces under the random policy. Before proceeding, we first recall the definition of the optimal action in the problems of multi-armed bandit (MAB) and MDP. For any query state s_q, the optimal action in the action space A for s_q corresponds to the action that maximizes the optimal Q-function, i.e.,

$$a^*_{\text{MAB}}(s_q) \triangleq \operatorname*{argmax}_{a \in A} \underbrace{\mathbb{E}\left[r(s_q, a)\right]}_{Q^*_{\text{MAB}}(s_q, a)}, \tag{4}$$

$$a^*_{\text{MDP}}(s_q) \triangleq \operatorname*{argmax}_{a \in A} \underbrace{\mathbb{E}_{\pi^*}\left[\sum_{t=0}^{\infty} \gamma^t r_t | s_0 = s_q, a_0 = a\right]}_{Q^*_{\text{MDP}}(s_q, a)}, \tag{5}$$

where s_q in (4) refers to the singleton state in the MAB problem, and π^* in (5) denotes the optimal policy in the MDP.

However, both a^*_{MAB} and a^*_{MDP} are intractable to obtain in our problem of interest for two reasons: (i) computing the expectation in the Q-function demands to sample infinite episodes; (ii) one can only have access to the random policy, instead of π^*. Therefore, we instead consider (i) stochastic approximation that uses the average as the unbiased estimate of the expectation due to the law of large numbers; (ii) maximizing the return under the random policy.

Trustworthiness of the Random Policy

Subsequently, the crucial question arises: **when can we trust the random policy?** We claim: *The random policy is probabilistically trustworthy for the MAB and MDP problems within a trust horizon.* We formalize this claim for the MAB and MDP respectively in this subsection, which relies on the following assumption.

Assumption 1. *The absolute value of the reward $r(s, a)$ is bounded by a positive constant B for all state-action pairs in the MAB and MDP, i.e., $|r(s, a)| \leq B, \forall (s, a) \in S \times A$.*

Note that Assumption 1 is common in the literature (Azar, Osband, and Munos 2017; Wei et al. 2020; Zhang, Du, and Ji 2021). In particular, in the case of finite state-action spaces, it is always possible to design the reward to avoid the possibility of being unbounded. With Assumption 1 established, we first formalize the trustworthiness of the random policy for the MAB problem in the following theorem.

Theorem 1 (MAB)**.** *Let Assumption 1 hold. Denote by s_q and a^* the singleton state and optimal arm in the MAB problem. Consider a random policy π. Suppose that each arm has been selected N times ($N \in \mathbb{N}_+$, trust horizon) under π. With probability at least $1 - \delta$, it holds that*

$$\frac{1}{N} \sum_{i=1}^{N} r(s_q, a^* \mid \pi) \geq \max_{a \in A \setminus \{a^*\}} \frac{1}{N} \sum_{i=1}^{N} r(s_q, a \mid \pi), \tag{6}$$

when the trust horizon N is greater than

$$\frac{8B^2}{\left(\mathbb{E}[r(s_q, a^*)] - \max\limits_{a \in A \setminus \{a^*\}} \mathbb{E}[r(s_q, a)]\right)^2} \log\left(\frac{1 + \sqrt{1 - \delta}}{\delta}\right). \tag{7}$$

Proof. See Appendix B. □

Theorem 1 implies that the trust horizon N quantifies the trustworthiness of the decision making under the random policy π for MAB problems. Indeed, a larger N implies a higher probability (smaller δ) that the average reward of the optimal arm under π exceeds that of the next-best arm, therefore, making a more reliable decision. We substantiate this claim by empirical evidence (depicted in Figure 5).

In the practical implementation, we simply execute the MAB under the random policy π until every action in the action space A selected exactly N times, discarding any actions that exceed the trust horizon N. Subsequently, we select the action with the maximal average reward as the action label for the singleton state. The detailed procedure for collecting such action labels in MAB is outlined in Algorithm 2.

While the trustworthiness of the random policy in the MAB problem has been established and discussed, the scenario of the MDP presents a distinct challenge. To proceed, we rely on the following assumption.

Assumption 2. *Given a random policy π, assume that*

$$\operatorname*{argmax}_{a \in A} Q^{\pi}_{MDP}(s_q, a) = \operatorname*{argmax}_{a \in A} Q^{*}_{MDP}(s_q, a), \ \forall s_q \in S. \tag{8}$$

It is worth noting that Assumption 2 holds in the MDP problems like grid world navigation (Laskin et al. 2022), which we consider for empirical evaluation in the Experiments section of this paper. Specifically, the action derived

Algorithm 2: Collecting Query States and Action Labels under Random Policy (MAB)

Require: Random policy π, singleton query state s_q, action space A, environment τ, empty average reward list L_r, trust horizon N
Execute the MAB in τ under the random policy π until every action in A selected exactly N times, discarding any actions that exceed the trust horizon N.
for a in $[A]$ **do**
    Record the average reward associated with the action a in the history, and add it to L_r
end for
Obtain $a_l = A(\text{argmax}(L_r))$
Return (s_q, a_l)

from maximizing the return under the random policy becomes (nearly) equivalent to that guided by maximizing the return under the optimal policy, as the maximal return induced by both policies corresponds to navigating to the goal as quickly as possible. We formalize this in Proposition 1 and provide a theoretical proof and empirical validation (refer to Appendix B). Besides, we also acknowledge that Assumption 2 may not hold universally for all MDP problems, and we leave its exploration for future research.

Having introduced Assumption 2, we are in conditions of establishing the trustworthiness of the random policy for MDP, which we formally state in the next theorem. To proceed, let us define

$$Q^{\pi,N}_{\text{MDP}}(s_q, a) = \mathbb{E}_\pi\left[\sum_{t=0}^{N} \gamma^t r(s_t, a_t) | s_0 = s_q, a_0 = a\right], \quad (9)$$

$$\hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a) = \frac{1}{N_{\text{ep}}} \sum_{i=1}^{N_{\text{ep}}} \sum_{t=0}^{N} \left(\gamma^t r(s_t, a_t) | s_0 = s_q, a_0 = a, \pi\right). \quad (10)$$

where $\hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a)$ is an unbiased estimate of $Q^{\pi,N}_{\text{MDP}}(s_q, a)$ and N_{ep} denotes the number of episodes.

Theorem 2 (MDP)**.** *Let Assumptions 1 and 2 hold. Denote by a^* the optimal action given s_q. Consider a random policy π as well as its Q-function Q^π_{MDP}. Consider the trust horizon $N > \log_\gamma(\kappa(1-\gamma)/(2B)) - 1$. Define $\kappa = \min_{s_q \in S}\left(Q^\pi_{MDP}(s_q, a^*) - \max_{a \in A\setminus\{a^*\}} Q^\pi_{MDP}(s_q, a)\right)$. With probability at least $1-\delta$, it holds that*

$$\hat{Q}^{\pi,N}_{MDP}(s_q, a^*) \geq \max_{a \in A\setminus\{a^*\}} \hat{Q}^{\pi,N}_{MDP}(s_q, a), \forall s_q \in S, \quad (11)$$

when the number of episodes N_{ep} satisfies

$$N_{ep} \geq \underbrace{\frac{2\left(1-\gamma^{N+1}\right)^2}{\left(\kappa(1-\gamma)/(2B) - \gamma^{N+1}\right)^2}}_{G_1} \log\left(\frac{1+\sqrt{1-\delta}}{\delta}\right). \quad (12)$$

Proof. See Appendix B. □

Theorem 2 implies that the trust horizon N quantifies the trustworthiness of the decision making under the random policy π for MDP problems. Notice that G_1 in (12) is monotonically decreasing with respect to the trust horizon N when $N > \log_\gamma(\kappa(1-\gamma)/(2B)) - 1$ (see Lemma 2 in Appendix B). Thus, with a fixed number of episodes, a larger N corresponds to a higher probability (smaller δ) that the average reward of the optimal action under π exceeds that of the next-best action, indicating a more reliable decision. This aligns with the intuition that a larger N corresponds to a closer approximation of the infinite-horizon MDP, where the optimal action emerges under the random policy π (by Assumption 2).

Algorithm 3: Collecting Query States and Action Labels under Random Policy (Sparse-Reward MDP)

Require: Random policy π, state space S, action space A, environment τ, trust horizon N
Set `min_step` $= N + 1$
while `min_step` $> N$ **do**
    Sample a query state $s_q \sim S$
    Empty a step list L_s
    for a in $[A]$ **do**
        Initialize the state and action as $s_0 = s_q$, $a_0 = a$
        Run an episode of N steps in τ under the random policy π, and terminate the episode early upon receiving a reward
        Add consumed steps to L_s (add "$N+1$" if no reward is received)
    end for
    `min_step` $= \min(L_s)$
end while
Obtain $a_l = A(\text{argmin}(L_s))$
Return (s_q, a_l)

In the practical implementation, we present two versions of pseudo-codes for the MDP with dense and sparse rewards. In the case of dense rewards, we randomly select a query state from the state space and execute an episode of N steps. Subsequently, we choose the action that maximizes $\hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a)$ across the entire action space A. The implementation details are summarized in Algorithm 5 (see Appendix C). Furthermore, the MDP with sparse rewards are generally more challenging to solve, as the agent does not receive feedback from the environment at each step. For any query state s_q, our goal remains to select the action that maximizes $\hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a)$. However, this approach may prove ineffective if N is too small, as the agent may never reach the goal, resulting in a return of 0 for all actions. To enhance the practicality in the implementation of the sparse-reward MDP, we adopt a different strategy. Specifically, for any query state s_q, we prioritize actions that can achieve the goal within N steps, with the actions consuming fewer steps being preferred. If no action can accomplish the goal within N steps, we sample another query state until a qualified action is identified. This action is then designated as the action label associated with the query state s_q. Details of this implementation is outlined in Algorithm 3.

Having introduced the processes for collecting the context, query state, and action label under the random policy, we can now generate the pretraining dataset by integrating the aforementioned procedures (refer to Algorithm 4). Given the pretraining dataset, the model pretraining procedure as well as the offline and online deployment for SAD are summarized in Algorithm 6 (see Appendix C).

Algorithm 4: State-Action Distillation (SAD) under Random Policy

Require: Empty pretraining dataset $\mathcal{D}$ with size $|\mathcal{D}|$, pretraining environment distribution $\mathcal{T}_{\text{train}}$, random policy π, context horizon length T, state space S, action space A, trust horizon N
for i in $[|\mathcal{D}|]$ **do**
  Sample an environment $\tau \sim \mathcal{T}_{\text{train}}$
  Collect the context $\mathcal{C}$ under the environment τ and the random policy π through Algorithm 1
  Collect the query state s_q and the action label a_l under the environment τ and the random policy π through Algorithm 2 for MAB (Algorithm 3 for sparse-reward MDP, Algorithm 5 for dense-reward MDP)
  Add $(\mathcal{C}, s_q, a_l)$ to the pretraining dataset $\mathcal{D}$
end for
Return $\mathcal{D}$

Performance Guarantees

In this subsection, we provide theoretical guarantees of our SAD approach. Theorems 1 and 2 imply the probability of SAD selecting the optimal action within a trust horizon. Then inspired by DPT, the trajectories generated by SAD take the same distribution as those produced by a well-specified posterior sampling (Osband, Russo, and Van Roy 2013) with a high probability. We formalize this claim in the following corollary.

Corollary 1. *Let hypotheses of Theorems 1 and 2 hold. Denote by l the length of the trajectory. For any environment τ and history data H, SAD and the well-specified posterior sampling follow the same trajectory distribution with probability* $(1-\delta)^l$

$$P_{\mathcal{F}_\theta}(\textit{trajectory}|\tau, H) = P_{ps}(\textit{trajectory}|\tau, H), \forall \textit{trajectory}. \tag{13}$$

Proof. See Appendix B. □

Having established the corollary above, we next investigate the regret bound of SAD in the finite MDP setting (see details in Appendix B). Consider the online cumulative regret of SAD over K episodes in the environment τ as $\text{Regret}_\tau(\mathcal{F}_\theta) = \sum_{k=0}^{K} V_\tau(\pi_\tau^*) - V_\tau(\pi_k)$, where $\pi_k(\cdot \mid s_t) = \mathcal{F}_\theta(\cdot \mid \mathcal{C}_{k-1}, s_t)$. Then, the regret bound of SAD is formally stated as follows.

Corollary 2. *Let hypotheses of Theorems 1 and 2 hold. Given the environment τ and a constant $B' > 0$, suppose that $\sup_\tau \mathcal{T}_{test}(\tau)/\mathcal{T}_{train}(\tau) \leq B'$. In the finite MDP with horizon T, it holds with probability* $(1-\delta)^{KT}$ *that*

$$\mathbb{E}_{\mathcal{T}_{test}}[\textit{Regret}_\tau(\mathcal{F}_\theta)] \leq \widetilde{\mathcal{O}}(B'|S|T^{3/2}\sqrt{K|A|}). \tag{14}$$

Proof. See Appendix B. □

Experiments

In this section, our empirical results on five ICRL benchmark environments (*Gaussian Bandits, Bernoulli Bandits, Darkroom, Darkroom-Large, Miniworld*) substantiate the efficacy of our proposed SAD method. The setup of these environments are deferred to Appendix D.

We compare our proposed SAD approach with three SOTA ICRL algorithms: AD, DPT, and DIT in the aforementioned benchmark environments (see details in Appendix D), and consider the DPT with optimal action labels (DPT*) as the oracle upper bound of SAD. Since all these methods are FM-based, we employ the same transformer architecture (causal GPT2 model (Radford et al. 2019)) and hyperparameters (number of attention layers, number of attention heads, embedding dimensions, etc) across all experiments to ensure a fair comparison. The main hyperparameters employed in this work are summarized in Tables 1-2 (refer to Appendix D).

In all experiments, we employ a uniform random policy to collect context, query states, and action labels, as indicated in Algorithms 1-5. Then, we pretrain and (offline/online) deploy the FM as presented in Algorithm 6 (see Appendix C).

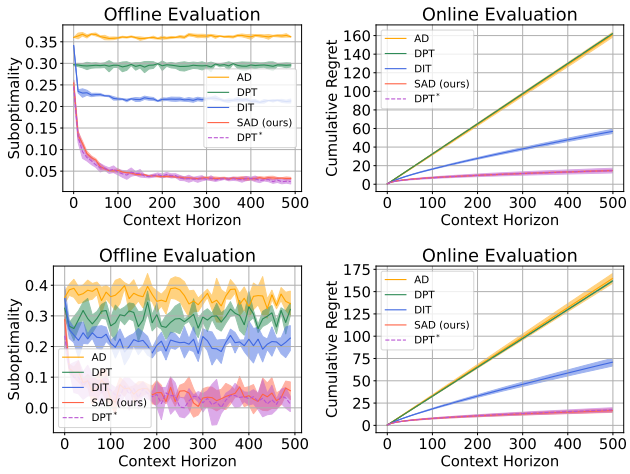

Figure 2: Offline and online evaluations of ICRL algorithms trained under a uniform random policy: AD, DPT, DIT, DPT*, and SAD (ours). Each algorithm contains four independent runs with mean and std. Environment: *Gaussian Bandits* (top row), *Bernoulli Bandits* (bottom row).

Bandits. We adhere to offline and online evaluation metrics for *Bandits* established in (Lee et al. 2024). In the offline evaluation, we utilize the *suboptimality* over different context horizon, defined by $\mu_{a^*} - \mu_a$, where μ_{a^*} and μ_a represent the mean rewards over 200 test environments of the optimal arm and the selected arm, respectively. In online evaluation, we employ *cumulative regret*, defined by $\sum_{t=0}^{T}(\mu_{a^*} - \mu_{a_t})$, where a_t denotes the selected arm at time t. The top row of Figure 2 demonstrates that our SAD approach significantly outperforms three SOTA baselines under uniform random policy, by achieving much lower *suboptimality* and *cumulative regret*. More specifically, let us

define the performance improvement of SAD over baselines in the offline evaluation by $(suboptimality_{\text{baseline}} - suboptimality_{\text{SAD}})/suboptimality_{\text{SAD}}$. Likewise, the performance improvement in the online evaluation is to simply replace the *suboptimality* by *cumulative regret*. Then, SAD surpasses the best baseline, DIT, by achieving 354.0% performance improvements in the offline evaluation and 273.9% in the online evaluation (refer to the first row of Tables 3 and 4 in Appendix D). To evaluate the out-of-distribution performance of SAD compared to other baselines, we test the models pretrained by all methods on the *Gaussian Bandits* and assess their performance on the *Bernoulli Bandits* with no further fine-tuning. The bottom row of Figure 2 illustrates that SAD still achieves lower *suboptimality* and *cumulative regret* than all other baselines, demonstrating a more robust performance in handling out-of-distribution scenarios. More specifically, SAD surpasses the best baseline, DIT, by 289.5% in the offline evaluation and 313.9% in the online evaluation (refer to the second row of Tables 3 and 4 in Appendix D).

In the environments of *Darkroom* and *Miniworld*, we use return as the evaluation metric. Moreover, we define the performance improvement of SAD over the baseline method in both offline and online evaluations by $(Return_{\text{SAD}} - Return_{\text{baseline}})/Return_{\text{baseline}}$.

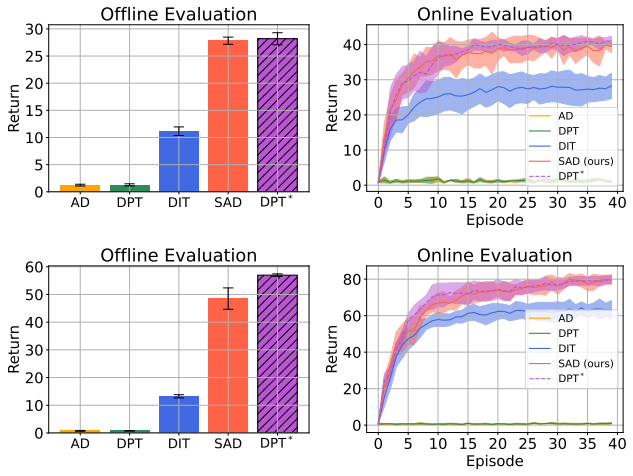

Figure 3: Offline and online evaluations of ICRL algorithms trained under a uniform random policy: AD, DPT, DIT, DPT*, and SAD (ours). Each algorithm contains four independent runs with mean and std. Environment: *DarkRoom* (top row), *DarkRoom-Large* (bottom row).

Darkrooms. Figure 3 demonstrates that our SAD approach significantly outperforms three SOTA baselines in the *Darkroom* and *Darkroom-Large* under uniform random policy, by achieving much higher return. In the *Darkroom*, SAD surpasses the best baseline, DIT, by 149.3% in the offline evaluation and 41.7% in the online evaluation (refer to the third row of Tables 3 and 4). Likewise, in the *Darkroom-Large*, SAD outperforms the best baseline, DIT, by 266.8% in the offline evaluation and 24.7% in the online evaluation (refer to the fourth row of Tables 3 and 4 in Appendix D).

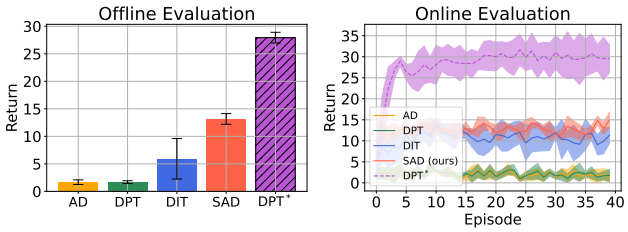

Figure 4: Offline and online evaluations of ICRL algorithms trained under a uniform random policy: AD, DPT, DIT, DPT*, and SAD (ours). Each algorithm contains four independent runs with mean and std. Environment: *Miniworld*.

Miniworld. Figure 4 demonstrates that our SAD approach outperforms the three SOTA baselines in the *Miniworld* under uniform random policy, by achieving a higher return. More specifically, SAD surpasses the best baseline, DIT, by 122.1% in the offline evaluation and 21.7% in the online evaluation (refer to the fifth row of Tables 3 and 4 in Appendix D).

Tables 3 and 4 also imply that SAD significantly outperforms all baselines on average across the five ICRL benchmark environments. In the offline evaluation, SAD exceeds the best baseline DIT by 236.3% on average, the second-best DPT by 2015.9%, and the third-best AD by 2075.2%. In the online evaluation, SAD surpasses DIT by 135.2%, DPT by 3093.8%, and AD by 3208.8% on average. In addition to comparing SAD with the three SOTA ICRL algorithms under the uniform random policy, we also include the empirical performance of the DPT with optimal action labels (DPT*) as the oracle upper bound of SAD. We observe that SAD demonstrates performance comparable to DPT* in tasks such as *Gaussian Bandits*, *Bernoulli Bandits*, *DarkRoom*, and *DarkRoom-Large*. Although *Miniworld* introduces challenges due to its pixel-based inputs and complex environments, SAD under the random policy still achieves approximately 50% of the performance of DPT*. Overall, SAD is within 18.6% of the performance of DPT* in the offline evaluation across the five ICRL tasks, and within 12.3% in the online evaluation (see details in Tables 3 and 4).

Ablation Studies

Trust Horizon. Theorems 1 implies that the uniform random policy is probabilistically trustworthy within a horizon N, with monotonically increasing probability of selecting the optimal action with N in the MAB problem. We substantiate this observation from the theorem by empirical evidence, as presented in Figure 5 (left). Furthermore, we conduct empirical investigations into the influence of the trust horizon N on the performance of the MAB problem, which considers the environments of *Gaussian Bandits*. As expected, a larger N in the MAB problem leads to a better performance (see top row of Figure 6).

We then shift to the MDP problem with the environment of *Darkroom*. Notice that in our practical algorithm for the

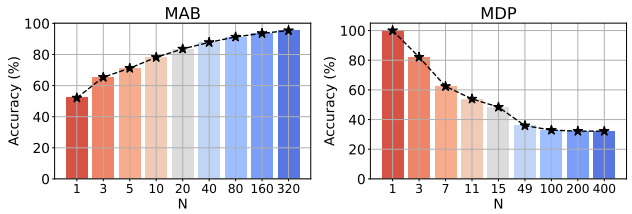

Figure 5: The accuracy (probability) of selecting the optimal action in the MAB and MDP with varying trust horizon N.

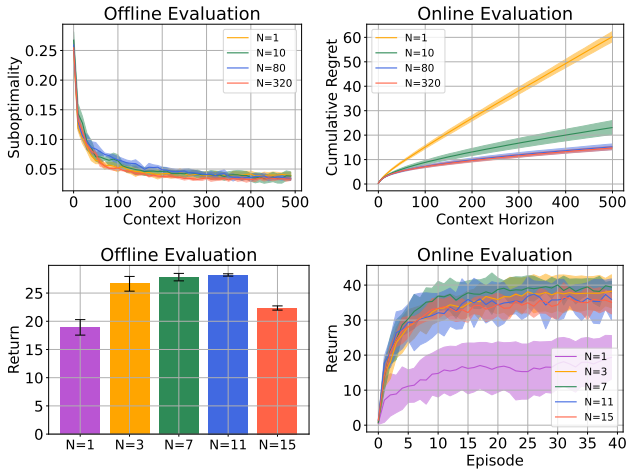

Figure 6: Offline and online evaluation of SAD with varying trust horizon N for MAB (top row) and MDP (bottom row). Each N contains four independent runs with mean and std.

sparse-reward MDP like *Darkroom* (Algorithm 3), we only utilize the state-action pairs that can reach the goal within a trust horizon. Therefore, we solely record the probability/accuracy of selecting the optimal actions on those states, as presented in Figure 5 (right). It shows that the accuracy monotonically decreases with respect to N, which ideally should lead to monotonically decreasing performance as well. However, we note that this is not the case. In particular, large trust horizon N in the MDP leads to the low accuracy of selecting the optimal action, whereas small N may induce the partially short-sighted training of FM, as Algorithm 3 solely trains on the states at most N steps from the goal, instead of all states (refer to Figure 7). Our numerical results in the bottom row of Figure 6 validate this with the fact $N = 7$ performing best, and provide the empirical evidence that either an excessively large or small trust horizon can lead to suboptimality.

Transformer Hyperparameters. We aim to validate the robustness of our proposed SAD approach with respect to the hyperparameters in the transformer block. Concretely, we focus on the number of attention heads and attention layers, as they have large impacts on the model size of the transformer. As depicted in Figure 8, our empirical results in *Darkroom* demonstrate a robust performance across varying numbers of attention heads and attention layers.

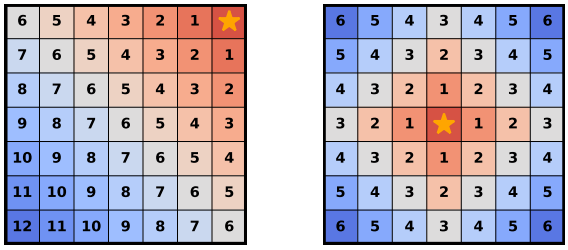

Figure 7: The minimal number of steps required for a query state to reach the goal (the golden star). Left: goal in the upper-right corner. Right: goal in the middle.

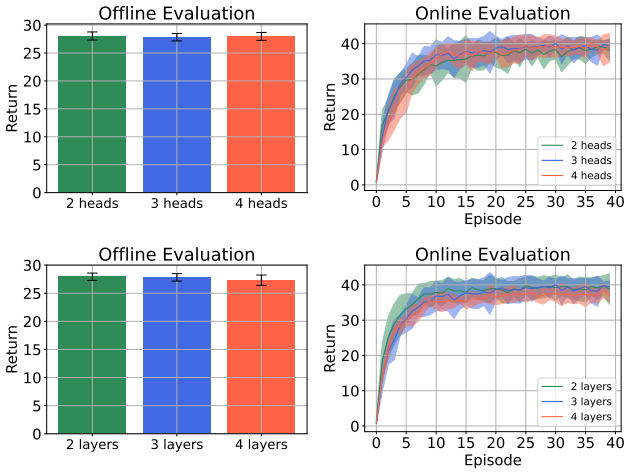

Figure 8: Offline and online evaluations of SAD with different transformer hyperparameters: the number of attention heads (top row), the number of attention layers (bottom row). Each hyperparameter contains four independent runs with mean and std.

Conclusion

In this work, we propose State-Action Distillation (SAD), a novel approach for generating the pretraining dataset for ICRL, which is designed to overcome the limitations imposed by the existing ICRL algorithms like AD, DPT, and DIT in terms of relying on well-trained or even optimal policies to collect the pretraining dataset. SAD leverages solely random policies to construct the pretraining data, significantly promoting the practical application of ICRL in real-world scenarios. We also provide the quantitative analysis of the trustworthiness as well as the performance guarantees of SAD. Moreover, our empirical results on multiple popular ICRL benchmark environments demonstrate significant improvements over the existing baselines in terms of both performance and robustness. Nevertheless, we note that SAD is currently limited to the discrete action space. Extending SAD to handle the continuous action space as well as more complex environments presents a promising direction for the future research.

References

Azar, M. G.; Osband, I.; and Munos, R. 2017. Minimax regret bounds for reinforcement learning. In *International conference on machine learning*, 263–272. PMLR.

Brown, T. B. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.

Chen, L.; Lu, K.; Rajeswaran, A.; Lee, K.; Grover, A.; Laskin, M.; Abbeel, P.; Srinivas, A.; and Mordatch, I. 2021. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34: 15084–15097.

Chen, W.; Mishra, S.; and Paternain, S. 2024. Domain Adaptation for Offline Reinforcement Learning with Limited Samples. *arXiv preprint arXiv:2408.12136*.

Chen, Z.; Wu, J.; Wang, W.; Su, W.; Chen, G.; Xing, S.; Zhong, M.; Zhang, Q.; Zhu, X.; Lu, L.; et al. 2024. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24185–24198.

Devlin, J. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Dong, J.; Guo, M.; Fang, E. X.; Yang, Z.; and Tarokh, V. 2024. In-Context Reinforcement Learning Without Optimal Action Labels. In *ICML 2024 Workshop on In-Context Learning*.

Dong, Q.; Li, L.; Dai, D.; Zheng, C.; Wu, Z.; Chang, B.; Sun, X.; Xu, J.; and Sui, Z. 2022. A survey on in-context learning. *arXiv preprint arXiv:2301.00234*.

Fatemi, M.; Wu, M.; Petch, J.; Nelson, W.; Connolly, S. J.; Benz, A.; Carnicelli, A.; and Ghassemi, M. 2022. Semi-markov offline reinforcement learning for healthcare. In *Conference on Health, Inference, and Learning*, 119–137. PMLR.

Fujimoto, S.; Conti, E.; Ghavamzadeh, M.; and Pineau, J. 2019. Benchmarking batch deep reinforcement learning algorithms. *arXiv preprint arXiv:1910.01708*.

Grigsby, J.; Fan, L.; and Zhu, Y. 2023. Amago: Scalable in-context reinforcement learning for adaptive agents. *arXiv preprint arXiv:2310.09971*.

Hoeffding, W. 1994. Probability inequalities for sums of bounded random variables. *The collected works of Wassily Hoeffding*, 409–426.

Janner, M.; Li, Q.; and Levine, S. 2021. Offline reinforcement learning as one big sequence modeling problem. *Advances in neural information processing systems*, 34: 1273–1286.

Kirk, R.; Zhang, A.; Grefenstette, E.; and Rocktäschel, T. 2023. A survey of zero-shot generalisation in deep reinforcement learning. *Journal of Artificial Intelligence Research*, 76: 201–264.

Kostrikov, I.; Nair, A.; and Levine, S. 2021. Offline reinforcement learning with implicit q-learning. *arXiv preprint arXiv:2110.06169*.

Kumar, A.; Zhou, A.; Tucker, G.; and Levine, S. 2020. Conservative q-learning for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 33: 1179–1191.

Laskin, M.; Wang, L.; Oh, J.; Parisotto, E.; Spencer, S.; Steigerwald, R.; Strouse, D.; Hansen, S.; Filos, A.; Brooks, E.; et al. 2022. In-context reinforcement learning with algorithm distillation. *arXiv preprint arXiv:2210.14215*.

Lee, J.; Xie, A.; Pacchiano, A.; Chandak, Y.; Finn, C.; Nachum, O.; and Brunskill, E. 2024. Supervised pretraining can learn in-context reinforcement learning. *Advances in Neural Information Processing Systems*, 36.

Lee, K.-H.; Nachum, O.; Yang, M. S.; Lee, L.; Freeman, D.; Guadarrama, S.; Fischer, I.; Xu, W.; Jang, E.; Michalewski, H.; et al. 2022. Multi-game decision transformers. *Advances in Neural Information Processing Systems*, 35: 27921–27936.

Levine, S.; Kumar, A.; Tucker, G.; and Fu, J. 2020. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*.

Li, Y.; Ildiz, M. E.; Papailiopoulos, D.; and Oymak, S. 2023. Transformers as algorithms: Generalization and stability in in-context learning. In *International Conference on Machine Learning*, 19565–19594. PMLR.

Light, J.; Cai, M.; Chen, W.; Wang, G.; Chen, X.; Cheng, W.; Yue, Y.; and Hu, Z. 2024a. Strategist: Learning Strategic Skills by LLMs via Bi-Level Tree Search. *arXiv preprint arXiv:2408.10635*.

Light, J.; Xing, S.; Liu, Y.; Chen, W.; Cai, M.; Chen, X.; Wang, G.; Cheng, W.; Yue, Y.; and Hu, Z. 2024b. PIANIST: Learning Partially Observable World Models with LLMs for Multi-Agent Decision Making. *arXiv preprint arXiv:2411.15998*.

Lin, L.; Bai, Y.; and Mei, S. 2023. Transformers as decision makers: Provable in-context reinforcement learning via supervised pretraining. *arXiv preprint arXiv:2310.08566*.

Liu, J.; Zhang, Z.; Wei, Z.; Zhuang, Z.; Kang, Y.; Gai, S.; and Wang, D. 2024. Beyond ood state actions: Supported cross-domain offline reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 13945–13953.

Lu, C.; Schroecker, Y.; Gu, A.; Parisotto, E.; Foerster, J.; Singh, S.; and Behbahani, F. 2024. Structured state space models for in-context reinforcement learning. *Advances in Neural Information Processing Systems*, 36.

Ma, Z.; Hong, J.; Gul, M. O.; Gandhi, M.; Gao, I.; and Krishna, R. 2023. Crepe: Can vision-language foundation models reason compositionally? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10910–10921.

Mazoure, B.; Kostrikov, I.; Nachum, O.; and Tompson, J. J. 2022. Improving zero-shot generalization in offline reinforcement learning using generalized similarity functions. *Advances in Neural Information Processing Systems*, 35: 25088–25101.

Osband, I.; Russo, D.; and Van Roy, B. 2013. (More) efficient reinforcement learning via posterior sampling. *Advances in Neural Information Processing Systems*, 26.

Radford, A. 2018. Improving language understanding by generative pre-training. *OpenAI blog*.

Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.

Reed, S.; Zolna, K.; Parisotto, E.; Colmenarejo, S. G.; Novikov, A.; Barth-Maron, G.; Gimenez, M.; Sulsky, Y.; Kay, J.; Springenberg, J. T.; et al. 2022. A generalist agent. *arXiv preprint arXiv:2205.06175*.

Sammani, F.; Mukherjee, T.; and Deligiannis, N. 2022. Nlx-gpt: A model for natural language explanations in vision and vision-language tasks. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8322–8332.

Sinii, V.; Nikulin, A.; Kurenkov, V.; Zisman, I.; and Kolesnikov, S. 2023. In-context reinforcement learning for variable action spaces. *arXiv preprint arXiv:2312.13327*.

Sohn, S.; Oh, J.; and Lee, H. 2018. Hierarchical reinforcement learning for zero-shot generalization with subtask dependencies. *Advances in neural information processing systems*, 31.

Sutton, R. S. 2018. Reinforcement learning: An introduction. *A Bradford Book*.

Tang, S.; and Wiens, J. 2021. Model selection for offline reinforcement learning: Practical considerations for healthcare settings. In *Machine Learning for Healthcare Conference*, 2–35. PMLR.

Vaswani, A. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*.

Wang, J.; Blaser, E.; Daneshmand, H.; and Zhang, S. 2024. Transformers Learn Temporal Difference Methods for In-Context Reinforcement Learning. *arXiv preprint arXiv:2405.13861*.

Wei, C.-Y.; Jahromi, M. J.; Luo, H.; Sharma, H.; and Jain, R. 2020. Model-free reinforcement learning in infinite-horizon average-reward markov decision processes. In *International conference on machine learning*, 10170–10180. PMLR.

Wei, J.; Wei, J.; Tay, Y.; Tran, D.; Webson, A.; Lu, Y.; Chen, X.; Liu, H.; Huang, D.; Zhou, D.; et al. 2023. Larger language models do in-context learning differently. *arXiv preprint arXiv:2303.03846*.

Wies, N.; Levine, Y.; and Shashua, A. 2024. The learnability of in-context learning. *Advances in Neural Information Processing Systems*, 36.

Xie, S. M.; Raghunathan, A.; Liang, P.; and Ma, T. 2021. An explanation of in-context learning as implicit bayesian inference. *arXiv preprint arXiv:2111.02080*.

Xu, H.; Jiang, L.; Jianxiong, L.; and Zhan, X. 2022a. A policy-guided imitation approach for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 35: 4085–4098.

Xu, M.; Shen, Y.; Zhang, S.; Lu, Y.; Zhao, D.; Tenenbaum, J.; and Gan, C. 2022b. Prompting decision transformer for few-shot policy generalization. In *international conference on machine learning*, 24631–24645. PMLR.

Yang, S.; Nachum, O.; Du, Y.; Wei, J.; Abbeel, P.; and Schuurmans, D. 2023. Foundation models for decision making: Problems, methods, and opportunities. *arXiv preprint arXiv:2303.04129*.

Yuan, L.; Chen, D.; Chen, Y.-L.; Codella, N.; Dai, X.; Gao, J.; Hu, H.; Huang, X.; Li, B.; Li, C.; et al. 2021. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*.

Zhang, Z.; Du, S.; and Ji, X. 2021. Near optimal reward-free reinforcement learning. In *International Conference on Machine Learning*, 12402–12412. PMLR.

Zintgraf, L.; Shiarlis, K.; Igl, M.; Schulze, S.; Gal, Y.; Hofmann, K.; and Whiteson, S. 2019. Varibad: A very good method for bayes-adaptive deep rl via meta-learning. *arXiv preprint arXiv:1910.08348*.

Zisman, I.; Kurenkov, V.; Nikulin, A.; Sinii, V.; and Kolesnikov, S. 2023. Emergence of In-Context Reinforcement Learning from Noise Distillation. *arXiv preprint arXiv:2312.12275*.

Zisselman, E.; Lavie, I.; Soudry, D.; and Tamar, A. 2023. Explore to generalize in zero-shot rl. *Advances in Neural Information Processing Systems*, 36: 63174–63196.

Appendix A: Related Work

Offline Reinforcement Learning

In contrast to the unlimited interactions with the environment in online RL, offline RL seeks to learn optimal policies from a pre-collected and static dataset (Fujimoto et al. 2019; Levine et al. 2020; Kumar et al. 2020; Kostrikov, Nair, and Levine 2021; Chen, Mishra, and Paternain 2024). One of the critical challenges in offline RL is with bootstrapping from out-of-distribution (OOD) actions (Levine et al. 2020; Kumar et al. 2020; Xu et al. 2022a; Liu et al. 2024) due to the mismatch between the behavior policies and the learned policies. To address this issue, the current SOTA offline RL algorithms propose to update pessimistically by either adding a regularization or underestimating the Q-value of OOD actions.

Autoregressive-Supervised Decision Making

In addition to the traditional offline RL methods, autoregressive-supervised mechanisms based on the transformer architecture (Vaswani 2017) have been successfully applied to offline decision making domains by their powerful capability in sequential modeling. The pioneering work in the autoregressive-supervised decision making is the Decision Transformer (DT) (Chen et al. 2021). DT autoregressively models the sequence of actions from the historical offline data conditioned on the sequence of returns in the history. During the inference, the trained model can be queried based on pre-defined target returns, allowing it to generate actions aligned with the target returns. The subsequent works such as Multi-Game Decision Transformer (MGDT) (Lee et al. 2022) and Gato (Reed et al. 2022) have exhibited the success of the autoregressive-supervised mechanisms in learning multi-task policies by fine-tuning or leveraging expert demonstrations in the downstream tasks.

In-Context Reinforcement Learning

However, both traditional offline RL and autoregressive-supervised decision making mechanisms suffer from the poor zero-shot generalization and in-context learning capabilities to new environments, as neither can improve the policy, with a fixed trained model, in context by trial and error. In-context reinforcement learning (ICRL) aims to pretrain a transformer-based FM, such as GPT2 (Radford et al. 2019), across a wide range of pretraining environments. During the evaluation (or inference), the pretrained model can directly infer the unseen environment and learn in-context without the need for updating model parameters. The SOTA ICRL algorithms including AD (Laskin et al. 2022), DPT (Lee et al. 2024) and DIT (Dong et al. 2024) have demonstrated the potential of the ICRL framework. Nevertheless, each of these methods imposes distinct yet strict requirements on the pretraining dataset e.g., requiring well-trained (or even optimal) behavior policies, the context to be episodic and/or substantial, which significantly restrict their practicality in real-world applications. Accordingly, mastering and executing ICRL under (e.g., uniform) random policies and random contexts remains a crucial direction and a critical challenge.

Appendix B: Omitted Proofs

Proof of Theorem 1

Theorem 1 (MAB). *Let Assumption 1 hold. Denote by s_q and a^* the singleton state and optimal arm in the MAB problem. Consider a random policy π. Suppose that each arm has been selected N times ($N \in \mathbb{N}_+$, trust horizon) under π. With probability at least $1-\delta$, it holds that*

$$\frac{1}{N}\sum_{i=1}^{N} r(s_q, a^* \mid \pi) \geq \max_{a \in A \setminus \{a^*\}} \frac{1}{N}\sum_{i=1}^{N} r(s_q, a \mid \pi), \tag{15}$$

when the trust horizon N satisfies

$$N \geq \frac{8B^2}{\left(\mathbb{E}[r(s_q, a^*)] - \max_{a \in A \setminus \{a^*\}} \mathbb{E}[r(s_q, a)]\right)^2} \log\left(\frac{1+\sqrt{1-\delta}}{\delta}\right). \tag{16}$$

Proof. For any action $a \in A \setminus \{a^*\}$, consider two positive constants

$$\epsilon_1 = \alpha\left(\mathbb{E}[r(s_q, a^*)] - \mathbb{E}[r(s_q, a)]\right), \tag{17}$$

$$\epsilon_2 = (1-\alpha)\left(\mathbb{E}[r(s_q, a^*)] - \mathbb{E}[r(s_q, a)]\right), \tag{18}$$

where $\alpha \in [0, 1]$.

Consider the following two inequalities

$$\frac{1}{N}\sum_{i=1}^{N} r(s_q, a^* \mid \pi) \geq \mathbb{E}[r(s_q, a^*)] - \epsilon_1, \tag{19}$$

$$\frac{1}{N}\sum_{i=1}^{N} r(s_q, a \mid \pi) \leq \mathbb{E}[r(s_q, a)] + \epsilon_2. \tag{20}$$

We note that the two inequalities above are the sufficient but not necessary conditions for $\frac{1}{N}\sum_{i=1}^{N} r(s_q, a^* \mid \pi) \geq \frac{1}{N}\sum_{i=1}^{N} r(s_q, a \mid \pi)$ to hold.

Therefore, we simply obtain that

$$P\left(\frac{1}{N}\sum_{i=1}^{N} r(s_q, a^* \mid \pi) \geq \frac{1}{N}\sum_{i=1}^{N} r(s_q, a \mid \pi)\right) \tag{21}$$

$$\geq P\left(\frac{1}{N}\sum_{i=1}^{N} r(s_q, a^* \mid \pi) \geq \mathbb{E}[r(s_q, a^*)] - \epsilon_1, \frac{1}{N}\sum_{i=1}^{N} r(s_q, a \mid \pi) \leq \mathbb{E}[r(s_q, a)] + \epsilon_2\right) \tag{22}$$

$$= P\left(\frac{1}{N}\sum_{i=1}^{N} r(s_q, a^* \mid \pi) \geq \mathbb{E}[r(s_q, a^*)] - \epsilon_1\right) \cdot P\left(\frac{1}{N}\sum_{i=1}^{N} r(s_q, a \mid \pi) \leq \mathbb{E}[r(s_q, a)] + \epsilon_2\right), \tag{23}$$

where the last equation follows from the fact that each arm is independent to other arms. We then lower bound the two probabilities in the previous expression using Hoeffding's inequality (Hoeffding 1994).

Since Assumption 1 implies that $r(\cdot, \cdot) \in [-B, B]$, Hoeffding's inequality yields

$$P\left(\frac{1}{N}\sum_{i=1}^{N} r(s_q, a \mid \pi) - \mathbb{E}[r(s_q, a)] \leq \epsilon_2\right) \geq 1 - \exp\left(-\frac{2N\epsilon_2^2}{(B-(-B))^2}\right), \tag{24}$$

i.e.,

$$P\left(\frac{1}{N}\sum_{i=1}^{N} r(s_q, a \mid \pi) - \mathbb{E}[r(s_q, a)] \leq \epsilon_2\right) \geq 1 - \exp\left(-\frac{N\epsilon_2^2}{2B^2}\right). \tag{25}$$

Likewise, we have

$$P\left(\frac{1}{N}\sum_{i=1}^{N} r(s_q, a^* \mid \pi) - \mathbb{E}[r(s_q, a^*)] \geq -\epsilon_1\right) \geq 1 - \exp\left(-\frac{2N\epsilon_1^2}{(B-(-B))^2}\right), \tag{26}$$

i.e.,

$$P\left(\frac{1}{N}\sum_{i=1}^{N} r(s_q, a^* \mid \pi) - \mathbb{E}[r(s_q, a^*)] \geq -\epsilon_1\right) \geq 1 - \exp\left(-\frac{N\epsilon_1^2}{2B^2}\right). \tag{27}$$

Therefore, it holds for any $\alpha \in [0, 1]$ that

$$P\left(\frac{1}{N}\sum_{i=1}^{N} r(s_q, a^* \mid \pi) \geq \frac{1}{N}\sum_{i=1}^{N} r(s_q, a \mid \pi)\right) \tag{28}$$

$$\geq \left(1 - \exp\left(-\frac{N\epsilon_1^2}{2B^2}\right)\right) \cdot \left(1 - \exp\left(-\frac{N\epsilon_2^2}{2B^2}\right)\right) \tag{29}$$

$$= \left(1 - \exp\left(-\frac{N\alpha^2\left(\mathbb{E}[r(s_q, a^*)] - \mathbb{E}[r(s_q, a)]\right)^2}{2B^2}\right)\right) \cdot \left(1 - \exp\left(-\frac{N(1-\alpha)^2\left(\mathbb{E}[r(s_q, a^*)] - \mathbb{E}[r(s_q, a)]\right)^2}{2B^2}\right)\right). \tag{30}$$

Notice that the maximal value of the previous equation with respect to α reaches at $\alpha = 0.5$. Then it holds that

$$P\left(\frac{1}{N}\sum_{i=1}^{N} r(s_q, a^* \mid \pi) \geq \frac{1}{N}\sum_{i=1}^{N} r(s_q, a \mid \pi)\right) \geq \left(1 - \exp\left(-\frac{N\left(\mathbb{E}[r(s_q, a^*)] - \mathbb{E}[r(s_q, a)]\right)^2}{8B^2}\right)\right)^2, \forall a \in A \setminus \{a^*\}. \tag{31}$$

Since the previous inequality holds for any $a \in A \setminus \{a^*\}$, let us define

$$\bar{a} = \operatorname*{argmax}_{a \in A \setminus \{a^*\}} \frac{1}{N}\sum_{i=1}^{N} r(s_q, a \mid \pi). \tag{32}$$

Then it also holds that

$$P\left(\frac{1}{N}\sum_{i=1}^{N} r(s_q, a^* \mid \pi) \geq \max_{a \in A \setminus \{a^*\}} \frac{1}{N}\sum_{i=1}^{N} r(s_q, a \mid \pi)\right) \geq \left(1 - \exp\left(-\frac{N\left(\mathbb{E}[r(s_q, a^*)] - \mathbb{E}[r(s_q, \bar{a})]\right)^2}{8B^2}\right)\right)^2 \tag{33}$$

$$\geq \left(1 - \exp\left(-\frac{N\left(\mathbb{E}[r(s_q, a^*)] - \max_{a \in A \setminus \{a^*\}} \mathbb{E}[r(s_q, a)]\right)^2}{8B^2}\right)\right)^2. \tag{34}$$

where the last inequality follows from the monotonicity.

To make the previous probability greater than $1 - \delta$, we require

$$\left(1 - \exp\left(-\frac{N\left(\mathbb{E}[r(s_q, a^*)] - \max_{a \in A \setminus \{a^*\}} \mathbb{E}[r(s_q, a)]\right)^2}{8B^2}\right)\right)^2 \geq 1 - \delta. \tag{35}$$

Hence, we require the trust horizon N to satisfy

$$N \geq \frac{8B^2}{\left(\mathbb{E}[r(s_q, a^*)] - \max_{a \in A \setminus \{a^*\}} \mathbb{E}[r(s_q, a)]\right)^2} \log\left(\frac{1 + \sqrt{1 - \delta}}{\delta}\right). \tag{36}$$

This completes the proof.

□

Validation of Assumption 2 in the grid world MDP

For the sake of simplicity and without loss of generality, we consider a single-dimensional grid world MDP with the understanding that Assumption 2 holds for the two-dimensional grid world MDP as well, which is considered in our numerical experiments. The environmental details of the single-dimensional grid world MDP can be found in Figure 9. To proceed, we rely on the lemma below, and Assumption 2 is then validated by Proposition 1 that follows.

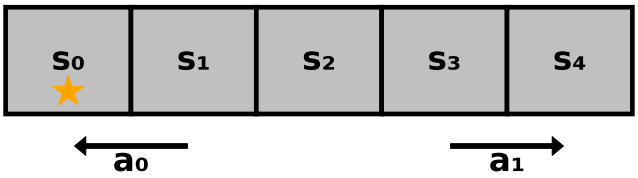

Figure 9: A single-dimensional grid world MDP comprising five states $\{s_0, s_1, s_2, s_3, s_4\}$, where s_0 represents the goal state (golden star). The environment offers two possible actions: a_0 that corresponds to moving left, and a_1 that corresponds to moving right. Crossing the boundaries is strictly prohibited. Any transitions that would result in boundary crossing will be confined to the current position. The reward structure is sparse, with a value of 1 received solely upon reaching the goal state s_0 and a value of 0 otherwise (see Figure 10 (left)). We consider an infinite time horizon with a discounter factor γ.

Lemma 1. *Consider the MDP of a single-dimensional grid world with $S = \{s_0, s_1, s_2, s_3, s_4\}$ and $A = \{a_0, a_1\}$, as depicted in Figure 9. Consider the random policy π in Assumption 2. It holds that*

$$V^{\pi}_{MDP}(s_0) \geq V^{\pi}_{MDP}(s_1) \geq V^{\pi}_{MDP}(s_2) \geq V^{\pi}_{MDP}(s_3) \geq V^{\pi}_{MDP}(s_4). \tag{37}$$

Proof. For simplicity, we consider π to be a uniform random policy, i.e., $P(a_0 \mid s) = P(a_1 \mid s) = 0.5, \forall s \in S$. Recall the Bellman expectation equation

$$V^{\pi}_{\text{MDP}}(s) = \sum_{a \in A} \pi(a \mid s) \left(r(s,a) + \gamma \mathbb{E}_{s'} V^{\pi}_{\text{MDP}}(s') \right). \tag{38}$$

By combining the previous Bellman expectation equation with Figure 9 yields

$$\begin{cases} V^{\pi}_{\text{MDP}}(s_0) = \frac{1}{2}\left(r(s_0,a_0) + \gamma V^{\pi}_{\text{MDP}}(s_0) + r(s_0,a_1) + \gamma V^{\pi}_{\text{MDP}}(s_1) \right), \\ V^{\pi}_{\text{MDP}}(s_1) = \frac{1}{2}\left(r(s_1,a_0) + \gamma V^{\pi}_{\text{MDP}}(s_0) + r(s_1,a_1) + \gamma V^{\pi}_{\text{MDP}}(s_2) \right), \\ V^{\pi}_{\text{MDP}}(s_2) = \frac{1}{2}\left(r(s_2,a_0) + \gamma V^{\pi}_{\text{MDP}}(s_1) + r(s_2,a_1) + \gamma V^{\pi}_{\text{MDP}}(s_3) \right), \\ V^{\pi}_{\text{MDP}}(s_3) = \frac{1}{2}\left(r(s_3,a_0) + \gamma V^{\pi}_{\text{MDP}}(s_2) + r(s_3,a_1) + \gamma V^{\pi}_{\text{MDP}}(s_4) \right), \\ V^{\pi}_{\text{MDP}}(s_4) = \frac{1}{2}\left(r(s_4,a_0) + \gamma V^{\pi}_{\text{MDP}}(s_3) + r(s_4,a_1) + \gamma V^{\pi}_{\text{MDP}}(s_4) \right). \end{cases} \tag{39}$$

Substituting all rewards from Figure 10 (left) into the previous equations yields

$$\begin{cases} V^{\pi}_{\text{MDP}}(s_0) = \frac{1}{2}\left(1 + \gamma V^{\pi}_{\text{MDP}}(s_0) + \gamma V^{\pi}_{\text{MDP}}(s_1) \right), \\ V^{\pi}_{\text{MDP}}(s_1) = \frac{1}{2}\left(1 + \gamma V^{\pi}_{\text{MDP}}(s_0) + \gamma V^{\pi}_{\text{MDP}}(s_2) \right), \\ V^{\pi}_{\text{MDP}}(s_2) = \frac{1}{2}\left(\gamma V^{\pi}_{\text{MDP}}(s_1) + \gamma V^{\pi}_{\text{MDP}}(s_3) \right), \\ V^{\pi}_{\text{MDP}}(s_3) = \frac{1}{2}\left(\gamma V^{\pi}_{\text{MDP}}(s_2) + \gamma V^{\pi}_{\text{MDP}}(s_4) \right), \\ V^{\pi}_{\text{MDP}}(s_4) = \frac{1}{2}\left(\gamma V^{\pi}_{\text{MDP}}(s_3) + \gamma V^{\pi}_{\text{MDP}}(s_4) \right). \end{cases} \tag{40}$$

Given $\gamma \in (0,1)$, the last equation of (40) implies that

$$V^{\pi}_{\text{MDP}}(s_3) = \frac{2-\gamma}{\gamma} V^{\pi}_{\text{MDP}}(s_4) \geq V^{\pi}_{\text{MDP}}(s_4). \tag{41}$$

Then, the fourth equation of (40) yields

$$V^{\pi}_{\text{MDP}}(s_2) = \frac{2}{\gamma} V^{\pi}_{\text{MDP}}(s_3) - V^{\pi}_{\text{MDP}}(s_4) \tag{42}$$

$$\geq \frac{2}{\gamma} V^{\pi}_{\text{MDP}}(s_3) - V^{\pi}_{\text{MDP}}(s_3) \tag{43}$$

$$\geq V^{\pi}_{\text{MDP}}(s_3). \tag{44}$$

Likewise, the third equation of (40) can be rewritten as

$$V^{\pi}_{\text{MDP}}(s_1) = \frac{2}{\gamma} V^{\pi}_{\text{MDP}}(s_2) - V^{\pi}_{\text{MDP}}(s_3) \tag{45}$$

$$\geq \frac{2}{\gamma} V^{\pi}_{\text{MDP}}(s_2) - V^{\pi}_{\text{MDP}}(s_2) \tag{46}$$

$$\geq V^{\pi}_{\text{MDP}}(s_2). \tag{47}$$

Combining the previous inequality with the first two equations of (40) directly yields

$$V^{\pi}_{\text{MDP}}(s_0) \geq V^{\pi}_{\text{MDP}}(s_1). \tag{48}$$

This completes the proof.

□

Proposition 1. *Consider the MDP of a single-dimensional grid world with $S = \{s_0, s_1, s_2, s_3, s_4\}$ and $A = \{a_0, a_1\}$, as depicted in Figure 9. Consider the random policy π in Assumption 2. It holds that*

$$\underset{a \in A}{\operatorname{argmax}}\ Q^{\pi}_{MDP}(s, a) = \underset{a \in A}{\operatorname{argmax}}\ Q^{*}_{MDP}(s, a),\ \forall s \in S. \tag{49}$$

Proof. By the definition of Q^{*}_{MDP} and π^* we obtain

$$Q^{*}_{\text{MDP}}(s, a) = Q^{\pi^*}_{\text{MDP}}(s, a),\ \forall (s, a) \in S \times A. \tag{50}$$

Since the objective of the agent in Figure 9 is to reach the goal state as quickly as possible, and stay still, $\pi^*(s)$ is given by

$$\pi^*(s) = a_0,\ \forall s \in S. \tag{51}$$

Moreover, we have

$$\underset{a \in A}{\operatorname{argmax}}\ Q^{*}_{\text{MDP}}(s, a) = \underset{a \in A}{\operatorname{argmax}}\ Q^{\pi^*}_{\text{MDP}}(s, a) = \pi^*(s) = a_0,\ \forall s \in S. \tag{52}$$

We then turn to consider the learning of Q-function under the random policy π. For simplicity, we consider π to be a uniform random policy, i.e., $P(a_0 \mid s) = P(a_1 \mid s) = 0.5,\ \forall s \in S$.

We next prove that $Q^{\pi}_{\text{MDP}}(s, a_0) \geq Q^{\pi}_{\text{MDP}}(s, a_1),\ \forall s \in S$. We start with the state s_0. Notice that

$$Q^{\pi}_{\text{MDP}}(s_0, a_0) = r(s_0, a_0) + \gamma V^{\pi}_{\text{MDP}}(s_0) = 1 + \gamma V^{\pi}_{\text{MDP}}(s_0), \tag{53}$$

$$Q^{\pi}_{\text{MDP}}(s_0, a_1) = r(s_0, a_1) + \gamma V^{\pi}_{\text{MDP}}(s_1) = \gamma V^{\pi}_{\text{MDP}}(s_1). \tag{54}$$

Lemma 1 implies that $V^{\pi}_{\text{MDP}}(s_0) \geq V^{\pi}_{\text{MDP}}(s_1)$. The previous equations then directly indicate

$$Q^{\pi}_{\text{MDP}}(s_0, a_0) \geq Q^{\pi}_{\text{MDP}}(s_0, a_1). \tag{55}$$

For the state s_1, we have

$$Q^{\pi}_{\text{MDP}}(s_1, a_0) = r(s_1, a_0) + \gamma V^{\pi}_{\text{MDP}}(s_0) = 1 + \gamma V^{\pi}_{\text{MDP}}(s_0), \tag{56}$$

$$Q^{\pi}_{\text{MDP}}(s_1, a_1) = r(s_1, a_1) + \gamma V^{\pi}_{\text{MDP}}(s_2) = \gamma V^{\pi}_{\text{MDP}}(s_2). \tag{57}$$

Lemma 1 implies that $V^{\pi}_{\text{MDP}}(s_0) \geq V^{\pi}_{\text{MDP}}(s_1) \geq V^{\pi}_{\text{MDP}}(s_2)$. Then it holds that

$$Q^{\pi}_{\text{MDP}}(s_1, a_0) \geq Q^{\pi}_{\text{MDP}}(s_1, a_1). \tag{58}$$

For the state s_2, we have

$$Q^{\pi}_{\text{MDP}}(s_2, a_0) = r(s_2, a_0) + \gamma V^{\pi}_{\text{MDP}}(s_1) = \gamma V^{\pi}_{\text{MDP}}(s_1), \tag{59}$$

$$Q^{\pi}_{\text{MDP}}(s_2, a_1) = r(s_2, a_1) + \gamma V^{\pi}_{\text{MDP}}(s_3) = \gamma V^{\pi}_{\text{MDP}}(s_3). \tag{60}$$

Employing Lemma 1 directly yields

$$Q^{\pi}_{\text{MDP}}(s_2, a_0) \geq Q^{\pi}_{\text{MDP}}(s_2, a_1). \tag{61}$$

For the state s_3, we have

$$Q^{\pi}_{\text{MDP}}(s_3, a_0) = r(s_3, a_0) + \gamma V^{\pi}_{\text{MDP}}(s_2) = \gamma V^{\pi}_{\text{MDP}}(s_2), \tag{62}$$

$$Q^{\pi}_{\text{MDP}}(s_3, a_1) = r(s_3, a_1) + \gamma V^{\pi}_{\text{MDP}}(s_4) = \gamma V^{\pi}_{\text{MDP}}(s_4). \tag{63}$$

Employing Lemma 1 directly yields

$$Q^{\pi}_{\text{MDP}}(s_3, a_0) \geq Q^{\pi}_{\text{MDP}}(s_3, a_1). \tag{64}$$

Last but not least, for the state s_4, we have

$$Q^{\pi}_{\text{MDP}}(s_4, a_0) = r(s_4, a_0) + \gamma V^{\pi}_{\text{MDP}}(s_3) = \gamma V^{\pi}_{\text{MDP}}(s_3), \tag{65}$$
$$Q^{\pi}_{\text{MDP}}(s_4, a_1) = r(s_4, a_1) + \gamma V^{\pi}_{\text{MDP}}(s_4) = \gamma V^{\pi}_{\text{MDP}}(s_4). \tag{66}$$

Employing Lemma 1 directly yields

$$Q^{\pi}_{\text{MDP}}(s_4, a_0) \geq Q^{\pi}_{\text{MDP}}(s_4, a_1). \tag{67}$$

Hence we obtain

$$\underset{a \in A}{\text{argmax}}\ Q^{\pi}_{\text{MDP}}(s, a) = a_0,\ \forall s \in S. \tag{68}$$

and then

$$\underset{a \in A}{\text{argmax}}\ Q^{\pi}_{\text{MDP}}(s, a) = a_0 = \underset{a \in A}{\text{argmax}}\ Q^{*}_{\text{MDP}}(s, a),\ \forall s \in S. \tag{69}$$

This completes the proof.

□

Empirical validation of Assumption 2 in the grid world MDP. In addition to the theoretical proof above, we also provide an empirical validation of Assumption 2 in the grid world MDP in Figure 9.

We start by considering the Bellman expectation equation

$$Q^{\pi}_{\text{MDP}}(s, a) = r(s, a) + \gamma \mathbb{E}_{a'} \left[Q^{\pi}_{\text{MDP}}(s', a') \right],\ \forall (s, a) \in S \times A. \tag{70}$$

Let us define the temporal-difference (TD) error as follows

$$\mathcal{E}(s, a) = \left| r(s, a) + \gamma \mathbb{E}_{a'} \left[Q^{\pi}_{\text{MDP}}(s', a') \right] - Q^{\pi}_{\text{MDP}}(s, a) \right|,\ \forall (s, a) \in S \times A. \tag{71}$$

With the understanding that the TD error $\mathcal{E}(s, a)$ of $Q^{\pi}_{\text{MDP}}(s, a)$ is zero, we iteratively learn and update $Q^{\pi}_{\text{MDP}}(s, a)$ by minimizing the TD error. To that end, we consider $\gamma = 0.99$ and a convergence threshold $\epsilon_Q = 10^{-6}$ (can be arbitrarily small). We initialize $Q^{\pi}_{\text{MDP}}(s, a)$ to be full of zeros as in Figure 10 (middle). Subsequently, we consistently update $Q^{\pi}_{\text{MDP}}(s, a)$ under the uniform policy π, until convergence as follows

$$\mathcal{E}(s, a) \leq \epsilon_Q,\ \forall (s, a) \in S \times A. \tag{72}$$

Empirically, we observe that the Q-table $Q^{\pi}_{\text{MDP}}(s, a)$ converges after the 1216^{th} iteration; demonstrated in Figure 11. As depicted in Figure 10 (right), the convergent $Q^{\pi}_{\text{MDP}}(s, a)$ implies that the optimal action for any state under the uniform policy π is a_0 (see the golden stars), i.e.,

$$\underset{a \in A}{\text{argmax}}\ Q^{\pi}_{\text{MDP}}(s, a) = a_0,\ \forall s \in S. \tag{73}$$

Hence we obtain

$$\underset{a \in A}{\text{argmax}}\ Q^{\pi}_{\text{MDP}}(s, a) = a_0 = \underset{a \in A}{\text{argmax}}\ Q^{*}_{\text{MDP}}(s, a),\ \forall s \in S. \tag{74}$$

This validates Assumption 2 empirically.

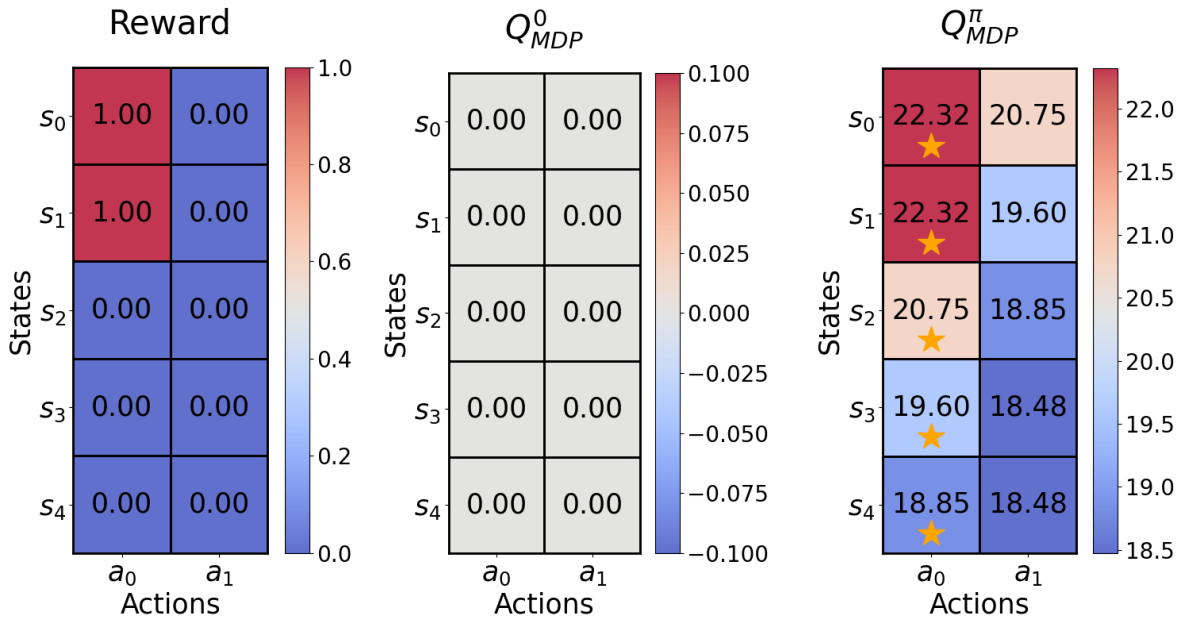

Figure 10: Left: reward table. Middle: initial Q-table. Right: convergent Q-table under the uniform policy π.

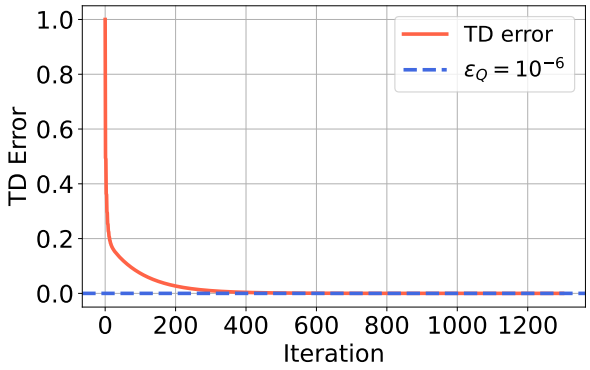

Figure 11: The learning curve of the maximal TD error of $Q^\pi_{\text{MDP}}(s,a)$ over the entire state-action spaces: $\max_{(s,a)\in S\times A} |r(s,a) + \gamma \mathbb{E}_{a'} [Q^\pi_{\text{MDP}}(s',a')] - Q^\pi_{\text{MDP}}(s,a)|$.

Proof of Theorem 2

Theorem 2 (MDP). *Let Assumptions 1 and 2 hold. Denote by a^* the optimal action given s_q. Consider a random policy π as well as its Q-function Q^{π}_{MDP}. Consider the trust horizon $N > \log_\gamma(\kappa(1-\gamma)/(2B)) - 1$. Define $\kappa = \min_{s_q \in S}\left(Q^{\pi}_{MDP}(s_q,a^*) - \max_{a\in A\setminus\{a^*\}} Q^{\pi}_{MDP}(s_q,a)\right)$. With probability at least $1-\delta$, it holds that*

$$\hat{Q}^{\pi,N}_{MDP}(s_q,a^*) \geq \max_{a\in A\setminus\{a^*\}} \hat{Q}^{\pi,N}_{MDP}(s_q,a),\ \forall s_q \in S, \tag{75}$$

when the number of episodes N_{ep} satisfies

$$N_{ep} \geq \frac{2\left(1-\gamma^{N+1}\right)^2}{\left(\kappa(1-\gamma)/(2B) - \gamma^{N+1}\right)^2} \log\left(\frac{1+\sqrt{1-\delta}}{\delta}\right). \tag{76}$$

Proof. For any action $a \in A\setminus\{a^*\}$, let us define

$$Q^{\pi}_{\text{MDP}}(s_q,a) = \underbrace{\mathbb{E}\left[\sum_{t=0}^{N}\gamma^t r(s_t,a_t) \mid s_0=s_q, a_0=a\right]}_{Q^{\pi,N}_{\text{MDP}}(s_q,a)} + \underbrace{\mathbb{E}\left[\sum_{t=N+1}^{\infty}\gamma^t r(s_t,a_t) \mid s_0=s_q, a_0=a\right]}_{\xi_a}, \tag{77}$$

$$Q^{\pi}_{\text{MDP}}(s_q,a^*) = \underbrace{\mathbb{E}\left[\sum_{t=0}^{N}\gamma^t r(s_t,a_t) \mid s_0=s_q, a_0=a^*\right]}_{Q^{\pi,N}_{\text{MDP}}(s_q,a^*)} + \underbrace{\mathbb{E}\left[\sum_{t=N+1}^{\infty}\gamma^t r(s_t,a_t) \mid s_0=s_q, a_0=a^*\right]}_{\xi_{a^*}}. \tag{78}$$

Then $\forall s_q \in S$,

$$\begin{aligned}
&Q^{\pi,N}_{\text{MDP}}(s_q,a^*) - Q^{\pi,N}_{\text{MDP}}(s_q,a) && (79)\\
&= Q^{\pi}_{\text{MDP}}(s_q,a^*) - \xi_{a^*} - (Q^{\pi}_{\text{MDP}}(s_q,a) - \xi_a) && (80)\\
&= Q^{\pi}_{\text{MDP}}(s_q,a^*) - Q^{\pi}_{\text{MDP}}(s_q,a) + \xi_a - \xi_{a^*} && (81)\\
&= Q^{\pi}_{\text{MDP}}(s_q,a^*) - Q^{\pi}_{\text{MDP}}(s_q,a) + \sum_{t=N+1}^{\infty}\gamma^t\left(\mathbb{E}\left[r(s_t,a_t) \mid s_0=s_q,a_0=a\right] - \mathbb{E}\left[r(s_t,a_t) \mid s_0=s_q,a_0=a^*\right]\right) && (82)\\
&\overset{(a)}{\geq} Q^{\pi}_{\text{MDP}}(s_q,a^*) - Q^{\pi}_{\text{MDP}}(s_q,a) - \gamma^{N+1}\frac{2B}{1-\gamma} && (83)\\
&\geq Q^{\pi}_{\text{MDP}}(s_q,a^*) - \max_{a\in A\setminus\{a^*\}} Q^{\pi}_{\text{MDP}}(s_q,a) - \gamma^{N+1}\frac{2B}{1-\gamma} && (84)\\
&\geq \min_{s_q\in S}\left(Q^{\pi}_{\text{MDP}}(s_q,a^*) - \max_{a\in A\setminus\{a^*\}} Q^{\pi}_{\text{MDP}}(s_q,a)\right) - \gamma^{N+1}\frac{2B}{1-\gamma} && (85)\\
&\overset{(b)}{=} \kappa - \gamma^{N+1}\frac{2B}{1-\gamma}, && (86)\\
&\overset{(c)}{>} 0. && (87)
\end{aligned}$$

where (a) follows from Assumption 1 and the properties of geometry series, (b) follows from the definition of κ, (c) is due to $N > \log_\gamma(\kappa(1-\gamma)/(2B)) - 1$. Therefore, let us consider two positive constants as follows

$$\epsilon_1 = \alpha\left(Q^{\pi,N}_{\text{MDP}}(s_q,a^*) - Q^{\pi,N}_{\text{MDP}}(s_q,a)\right), \tag{88}$$

$$\epsilon_2 = (1-\alpha)\left(Q^{\pi,N}_{\text{MDP}}(s_q,a^*) - Q^{\pi,N}_{\text{MDP}}(s_q,a)\right), \tag{89}$$

where $\alpha \in [0,1]$. Consider the following two inequalities

$$\hat{Q}^{\pi,N}_{\text{MDP}}(s_q,a^*) \geq Q^{\pi,N}_{\text{MDP}}(s_q,a^*) - \epsilon_1, \tag{90}$$

$$\hat{Q}^{\pi,N}_{\text{MDP}}(s_q,a) \leq Q^{\pi,N}_{\text{MDP}}(s_q,a) + \epsilon_2. \tag{91}$$

We acknowledge that (90) and (91) are the sufficient but not necessary conditions for $\hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a^*) \geq \hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a)$ to hold. Thus,

$$P\left(\hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a^*) \geq \hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a)\right) \tag{92}$$

$$\geq P\left(\hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a^*) \geq Q^{\pi,N}_{\text{MDP}}(s_q, a^*) - \epsilon_1, \hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a) \leq Q^{\pi,N}_{\text{MDP}}(s_q, a) + \epsilon_2\right) \tag{93}$$

$$= P\left(\hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a^*) \geq Q^{\pi,N}_{\text{MDP}}(s_q, a^*) - \epsilon_1\right) \cdot P\left(\hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a) \leq Q^{\pi,N}_{\text{MDP}}(s_q, a) + \epsilon_2\right) \tag{94}$$

where the last equation follows from the fact that each action is independent to other actions.

Assumption 1 implies that $\sum_{t=0}^{N} (\gamma^t r(s_t, a_t) \mid s_0 = s_q, a_0 = a, \pi) \in [-B\frac{1-\gamma^{N+1}}{1-\gamma}, B\frac{1-\gamma^{N+1}}{1-\gamma}], \forall a \in A$. We therefore lower bound the two probabilities in the previous expression using Hoeffding's inequality (Hoeffding 1994)

$$P\left(\hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a^*) - Q^{\pi,N}_{\text{MDP}}(s_q, a^*) \geq -\epsilon_1\right) \geq 1 - \exp\left(-\frac{N_{\text{ep}}\epsilon_1^2}{2B^2\left(\frac{1-\gamma^{N+1}}{1-\gamma}\right)^2}\right), \tag{95}$$

$$P\left(\hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a) - Q^{\pi,N}_{\text{MDP}}(s_q, a) \leq \epsilon_2\right) \geq 1 - \exp\left(-\frac{N_{\text{ep}}\epsilon_2^2}{2B^2\left(\frac{1-\gamma^{N+1}}{1-\gamma}\right)^2}\right). \tag{96}$$

Then, it holds for any $\alpha \in [0, 1]$ that

$$P\left(\hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a^*) \geq \hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a)\right) \tag{97}$$

$$\geq \left(1 - \exp\left(-\frac{N_{\text{ep}}\epsilon_1^2}{2B^2\left(\frac{1-\gamma^{N+1}}{1-\gamma}\right)^2}\right)\right) \cdot \left(1 - \exp\left(-\frac{N_{\text{ep}}\epsilon_2^2}{2B^2\left(\frac{1-\gamma^{N+1}}{1-\gamma}\right)^2}\right)\right) \tag{98}$$

$$\geq \left(1 - \exp\left(-\frac{N_{\text{ep}}\alpha^2\left(Q^{\pi,N}_{\text{MDP}}(s_q, a^*) - Q^{\pi,N}_{\text{MDP}}(s_q, a)\right)^2}{2B^2\left(\frac{1-\gamma^{N+1}}{1-\gamma}\right)^2}\right)\right) \cdot \left(1 - \exp\left(-\frac{N_{\text{ep}}(1-\alpha)^2\left(Q^{\pi,N}_{\text{MDP}}(s_q, a^*) - Q^{\pi,N}_{\text{MDP}}(s_q, a)\right)^2}{2B^2\left(\frac{1-\gamma^{N+1}}{1-\gamma}\right)^2}\right)\right) \tag{99}$$

Notice that the maximal value of the previous expression with respect to α reaches at $\alpha = 0.5$. Then it holds that

$$P\left(\hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a^*) \geq \hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a)\right) \geq \left(1 - \exp\left(-\frac{N_{\text{ep}}\left(Q^{\pi,N}_{\text{MDP}}(s_q, a^*) - Q^{\pi,N}_{\text{MDP}}(s_q, a)\right)^2}{8B^2\left(\frac{1-\gamma^{N+1}}{1-\gamma}\right)^2}\right)\right)^2, \forall a \in A \setminus \{a^*\}. \tag{100}$$

Since the previous inequality holds for any $a \in A \setminus \{a^*\}$, let us define

$$\hat{a} = \underset{a \in A \setminus \{a^*\}}{\text{argmax}}\ \hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a). \tag{101}$$

Then it also holds that

$$P\left(\hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a^*) \geq \max_{a \in A \setminus \{a^*\}} \hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a)\right) \geq \left(1 - \exp\left(-\frac{N_{\text{ep}}\left(Q^{\pi,N}_{\text{MDP}}(s_q, a^*) - Q^{\pi,N}_{\text{MDP}}(s_q, \hat{a})\right)^2}{8B^2\left(\frac{1-\gamma^{N+1}}{1-\gamma}\right)^2}\right)\right)^2 \tag{102}$$

$$\geq \left(1 - \exp\left(-\frac{N_{\text{ep}}\left(Q^{\pi,N}_{\text{MDP}}(s_q, a^*) - \max_{a \in A \setminus \{a^*\}} Q^{\pi,N}_{\text{MDP}}(s_q, a)\right)^2}{8B^2\left(\frac{1-\gamma^{N+1}}{1-\gamma}\right)^2}\right)\right)^2, \tag{103}$$

where the last inequality follows from the monotonicity and the fact that $Q^{\pi,N}_{\text{MDP}}(s_q, a^*) - Q^{\pi,N}_{\text{MDP}}(s_q, a) \geq 0, \forall a \in A \setminus \{a^*\}$.

Let us define

$$\bar{a} = \underset{a \in A \setminus \{a^*\}}{\operatorname{argmax}} \; Q^{\pi,N}_{\text{MDP}}(s_q, a). \tag{104}$$

Then $\forall s_q \in S$ we obtain

$$\begin{aligned}
&Q^{\pi,N}_{\text{MDP}}(s_q, a^*) - \max_{a \in A \setminus \{a^*\}} Q^{\pi,N}_{\text{MDP}}(s_q, a) && (105)\\
&\overset{(a)}{=} Q^{\pi,N}_{\text{MDP}}(s_q, a^*) - Q^{\pi,N}_{\text{MDP}}(s_q, \bar{a}) && (106)\\
&\overset{(b)}{=} Q^{\pi}_{\text{MDP}}(s_q, a^*) - \xi_{a^*} - (Q^{\pi}_{\text{MDP}}(s_q, \bar{a}) - \xi_{\bar{a}}) && (107)\\
&= Q^{\pi}_{\text{MDP}}(s_q, a^*) - Q^{\pi}_{\text{MDP}}(s_q, \bar{a}) + \xi_{\bar{a}} - \xi_{a^*} && (108)\\
&= Q^{\pi}_{\text{MDP}}(s_q, a^*) - Q^{\pi}_{\text{MDP}}(s_q, \bar{a}) + \sum_{t=N+1}^{\infty} \gamma^t \left(\mathbb{E}\left[r(s_t, a_t) \mid s_0 = s_q, a_0 = \bar{a} \right] - \mathbb{E}\left[r(s_t, a_t) \mid s_0 = s_q, a_0 = a^* \right] \right) && (109)\\
&\overset{(c)}{\geq} Q^{\pi}_{\text{MDP}}(s_q, a^*) - Q^{\pi}_{\text{MDP}}(s_q, \bar{a}) - \gamma^{N+1} \frac{2B}{1-\gamma} && (110)\\
&\overset{(d)}{\geq} Q^{\pi}_{\text{MDP}}(s_q, a^*) - \max_{a \in A \setminus \{a^*\}} Q^{\pi}_{\text{MDP}}(s_q, a) - \gamma^{N+1} \frac{2B}{1-\gamma} && (111)\\
&\geq \min_{s_q \in S} \left(Q^{\pi}_{\text{MDP}}(s_q, a^*) - \max_{a \in A \setminus \{a^*\}} Q^{\pi}_{\text{MDP}}(s_q, a) \right) - \gamma^{N+1} \frac{2B}{1-\gamma} && (112)\\
&\overset{(e)}{=} \kappa - \gamma^{N+1} \frac{2B}{1-\gamma} && (113)\\
&\overset{(f)}{>} 0, && (114)
\end{aligned}$$

where (a) follows from the definition of $\bar{a}$, (b) follows from the definition of Q^{π}_{MDP}, (c) follows from Assumption 1 and the properties of geometry series, (d) follows from the fact that $\bar{a}$ may not be the maximizer of $Q^{\pi}_{\text{MDP}}(s_q, \cdot)$, (e) follows from the definition of κ, (f) is due to $N > \log_\gamma \left(\kappa(1-\gamma)/(2B) \right) - 1$.

Consequently, (103) is monotonically increasing with $Q^{\pi,N}_{\text{MDP}}(s_q, a^*) - \max_{a \in A \setminus \{a^*\}} Q^{\pi,N}_{\text{MDP}}(s_q, a)$. Substituting the previous inequality into (103) and by the monotonicity yields

$$P\left(\hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a^*) \geq \max_{a \in A \setminus \{a^*\}} \hat{Q}^{\pi,N}_{\text{MDP}}(s_q, a) \right) \geq \left(1 - \exp\left(- \frac{N_{\text{ep}} \left(\kappa - \gamma^{N+1} \frac{2B}{1-\gamma} \right)^2}{8B^2 \left(\frac{1-\gamma^{N+1}}{1-\gamma} \right)^2} \right) \right)^2, \tag{115}$$

To make the previous probability greater than $1 - \delta$, we require

$$\left(1 - \exp\left(- \frac{N_{\text{ep}} \left(\kappa - \gamma^{N+1} \frac{2B}{1-\gamma} \right)^2}{8B^2 \left(\frac{1-\gamma^{N+1}}{1-\gamma} \right)^2} \right) \right)^2 \geq 1 - \delta. \tag{116}$$

Thus,

$$\begin{aligned}
N_{\text{ep}} &\geq \log\left(\frac{1 + \sqrt{1-\delta}}{\delta} \right) \frac{8B^2 \left(\frac{1-\gamma^{N+1}}{1-\gamma} \right)^2}{\left(\kappa - \gamma^{N+1} \frac{2B}{1-\gamma} \right)^2} && (117)\\
&= \frac{2\left(1 - \gamma^{N+1}\right)^2}{\left(\kappa\left(1-\gamma\right)/(2B) - \gamma^{N+1} \right)^2} \log\left(\frac{1 + \sqrt{1-\delta}}{\delta} \right). && (118)
\end{aligned}$$

This completes the proof.

□

Technical Lemma

Lemma 2. *Given $N > \log_\gamma(\kappa(1-\gamma)/(2B)) - 1$, G_1 in (12) is monotonically decreasing with respect to the trust horizon N.*

Proof. We proceed by defining $Y = \kappa(1-\gamma)/(2B)$. Since $N > \log_\gamma(\kappa(1-\gamma)/(2B)) - 1$, we can obtain

$$Y = \frac{\kappa(1-\gamma)}{2B} \in (\gamma^{N+1}, 1]. \tag{119}$$

Let $Z = N + 1$, and then we can rewrite G_1 as

$$G_1 = \frac{2\left(1-\gamma^Z\right)^2}{\left(Y-\gamma^Z\right)^2}, \text{ where } Y \in (\gamma^Z, 1]. \tag{120}$$

The chain rule implies that

$$\begin{aligned}
\frac{\partial G_1}{\partial N} &= \frac{\partial G_1}{\partial Z} \cdot \frac{\partial Z}{\partial N} && (121)\\
&= \frac{\partial G_1}{\partial Z} && (122)\\
&= 4\left(\frac{1-\gamma^Z}{Y-\gamma^Z}\right)\frac{-\gamma^Z \log\gamma(Y-\gamma^Z) + (1-\gamma^Z)\gamma^Z\log\gamma}{(Y-\gamma^Z)^2} && (123)\\
&= 4\left(\frac{1-\gamma^Z}{Y-\gamma^Z}\right)\frac{\gamma^Z\log\gamma(1-Y)}{(Y-\gamma^Z)^2} && (124)\\
&\leq 0, && (125)
\end{aligned}$$

where the last inequality follows from $\gamma \in (0, 1]$ and $Y \in (\gamma^Z, 1]$.

This completes the proof.

□

Proof of Corollary 1

Corollary 1. *Let hypotheses of Theorems 1 and 2 hold. Denote by l the length of the trajectory.* For any environment τ and history data H, SAD and the well-specified posterior sampling follow the same trajectory distribution with probability $(1-\delta)^l$

$$P_{\mathcal{F}_\theta}(\text{trajectory} \mid \tau, H) = P_{ps}(\text{trajectory} \mid \tau, H), \forall \text{trajectory}. \tag{126}$$

Proof. To proceed, we rely on the following assumption.

Assumption 3. *Denote by $\mathcal{F}_\theta$ the pretrained FM. $\forall(\mathcal{C}, s_q)$, assume $P_{train}(a \mid \mathcal{C}, s_q) = \mathcal{F}_\theta(a \mid \mathcal{C}, s_q)$ for all $a \in A$.*

Note that Assumption 3 is a common assumption in the in-context learning literature (Xie et al. 2021; Lee et al. 2024), assuming that the pretrained FM fits the pretraining distribution exactly provided with sufficient coverage and data, where the SAD fits with a sufficiently large trust horizon N.

With Assumption 3 established, Theorem 1 of (Lee et al. 2024) implies that (126) holds when the optimal action is selected at each step. In addition, Theorems 1 and 2 indicate that the FM trained by SAD selects the optimal action label with probability $1-\delta$ at each step. Consequently, (126) holds for SAD with probability $(1-\delta)^l$.

This completes the proof. □

Finite MDP setting from (Osband, Russo, and Van Roy 2013)

Let us consider the finite MDP setting as in (Osband, Russo, and Van Roy 2013), where $\mathbb{E}[r(s_t, a_t)] \in [0, 1]$. Denote by S, A, T the state space, action space, and time horizon. Consider the uniform random policy π for sampling the context $\mathcal{C}$ and query state s_q. Denote by $\mathcal{T}_{\text{test}}(\tau)$ and $\mathcal{T}_{\text{train}}(\tau)$ the test and pretraining distribution over the environment τ, respectively. Consider the online cumulative regret of SAD over K episodes in the environment τ as

$$\text{Regret}_\tau(\mathcal{F}_\theta) \triangleq \sum_{k=0}^{K} V_\tau(\pi_\tau^*) - V_\tau(\pi_k), \tag{127}$$

where $\pi_k(\cdot \mid s_t) = \mathcal{F}_\theta(\cdot \mid \mathcal{C}_{k-1}, s_t)$.

Proof of Corollary 2

Corollary 2. *Let hypotheses of Theorems 1 and 2 hold. Given the environment τ and a constant $B' > 0$, suppose that $\sup_\tau \mathcal{T}_{test}(\tau)/\mathcal{T}_{train}(\tau) \leq B'$. In the finite MDP setting above, it holds with probability $(1-\delta)^{KT}$ that*

$$\mathbb{E}_{\mathcal{T}_{test}}\left[Regret_\tau(\mathcal{F}_\theta)\right] \leq \widetilde{\mathcal{O}}(B'|S|T^{3/2}\sqrt{K|A|}). \tag{128}$$

Proof. Theorems 1 and 2 imply that the FM trained by SAD selects the optimal action label with probability $1-\delta$ at each step, while the finite MDP setting above comprises $K \cdot T$ steps. Therefore, it holds with probability $(1-\delta)^{KT}$ that the trained FM $\mathcal{F}_\theta$ is equivalent to the posterior sampling established in Corollary 1. Then, it follows directly from Corollary 6.2 of (Lee et al. 2024) that with probability $(1-\delta)^{KT}$ it holds that

$$\mathbb{E}_{\mathcal{T}_{\text{train}}}\left[\text{Regret}_\tau(\mathcal{F}_\theta)\right] \leq \widetilde{\mathcal{O}}(|S|T^{3/2}\sqrt{K|A|}), \tag{129}$$

where the notation $\widetilde{\mathcal{O}}$ omits the polylogarithmic dependence. Subsequently, by using the bounded likelihood ratio between the test and pretraining distributions yields

$$\begin{aligned}
\mathbb{E}_{\mathcal{T}_{\text{test}}}\left[\text{Regret}_\tau(\mathcal{F}_\theta)\right] &= \int \mathcal{T}_{\text{test}}(\tau)\text{Regret}_\tau(\mathcal{F}_\theta)d(\tau) && (130)\\
&\leq B' \int \mathcal{T}_{\text{train}}(\tau)\text{Regret}_\tau(\mathcal{F}_\theta)d(\tau) && (131)\\
&= B'\mathbb{E}_{\mathcal{T}_{\text{train}}}\left[\text{Regret}_\tau(\mathcal{F}_\theta)\right] && (132)\\
&\leq \widetilde{\mathcal{O}}(B'|S|T^{3/2}\sqrt{K|A|}). && (133)
\end{aligned}$$

This completes the proof.

□

Appendix C: Implementation Details

Pseudo-codes

We provide below the pseudo-codes that are omitted in the main body of the paper.

Algorithm 5: Collecting Query States and Action Labels under Random Policy (Dense-Reward MDP)

Require: Random policy π, state space S, action space A, environment τ, trust horizon N, empty return list L_r
Sample a query state $s_q \sim S$
for a in $[A]$ **do**
    Initialize the state and action as $s_0 = s_q$, $a_0 = a$
    Run an episode of N steps in τ under the random policy π
    Add the discounted episodic return to L_r
end for
Obtain $a_l = A(\mathrm{argmax}(L_r))$
Return (s_q, a_l)

Algorithm 6: Pretraining and Deployment of SAD (Inspired by (Lee et al. 2024))

Require: Pretraining dataset $\mathcal{D}$, initial model parameters θ, test environment distribution $\mathcal{T}_{\text{test}}$, number of episodes N_E
// Model pretraining
while not converged **do**
    Sample $(\mathcal{C}, s_q, a_l)$ from the pretraining dataset $\mathcal{D}$ and predict actions by the model $\mathcal{F}_\theta(\cdot|\mathcal{C}_i, s_q)$ for all $i \in [|\mathcal{C}|]$
    Compute the loss in (3) with respect to the action label a_l and backpropagate to update θ.
end while
// Offline deployment
Sample unseen environments $\tau \sim \mathcal{T}_{\text{test}}$
Sample a context $\mathcal{C} \sim \mathcal{T}_{\text{test}}(\cdot \mid \tau)$
Deploy $\mathcal{F}_\theta$ in τ by selecting $a_t \in \mathrm{argmax}_{a \in A} \mathcal{F}_\theta(a \mid \mathcal{C}, s_t)$ at time step t
// Online deployment
Sample unseen environments $\tau \sim \mathcal{T}_{\text{test}}$ and initialize empty context $\mathcal{C} = \{\}$
for i in $[N_E]$ **do**
    Deploy $\mathcal{F}_\theta$ by sampling $a_t \sim \mathcal{F}_\theta(\cdot \mid \mathcal{C}, s_t)$ at time step t
    Add $(s_0, a_0, r_0, \ldots)$ to $\mathcal{C}$
end for

Appendix D: Experimental Details

Environmental Setup

Gaussian Bandits. We investigate a five-armed bandit problem in which the state space S consists solely of a singleton state s_q. With each arm (action) pulled, the agent receives a reward feedback. The goal is to identify the optimal arm that can maximize the cumulative reward. We consider the reward function for each arm following a Gaussian distribution with mean μ_a and variance σ^2, i.e., $R(\cdot|s_q, a) = \mathcal{N}(\mu_a, \sigma^2)$. Each arm possesses means μ_a drawn from a uniform distribution $U[0, 1]$ and all arms share the same variance $\sigma = 0.3$. We consider the pretraining and test data to have distinct Gaussian distributions with different means.

Bernoulli Bandits. We adopt the same setup as in *Gaussian Bandits*, with the exception that the reward function does not follow a Gaussian distribution. Instead, we model the reward function using a Bernoulli distribution. Specifically, the mean of each arm μ_a is drawn from a Beta distribution $Beta(1, 1)$, and the reward function follows a Bernoulli distribution with probability of success μ_a. To validate the capability of SAD tackling OOD scenarios, we consider the test data drawn from the Bernoulli distribution while the pretraining data drawn from the Gaussian distribution as in the Gaussian bandits.

Darkroom. *Darkroom* (Laskin et al. 2022; Zintgraf et al. 2019) is a two-dimensional navigation task with discrete state and action spaces. The room consists of 7×7 grids ($|S| = 49$), with an unknown goal randomly placed at any of these grids. The agent can select 5 actions: go up, go down, go left, go right, or stay. The horizon length for *Darkroom* is 49, meaning the agent must reach the goal within 49 moves. The challenge of this task arises from its sparse reward structure, i.e., the agent receives a reward of 1 solely upon reaching the goal, and 0 otherwise. Given $7 \times 7 = 49$ available goals, we utilize 39 of these goals ($\sim 80\%$) for pretraining and reserve the remaining 10 ($\sim 20\%$) (unseen during pretraining) for test.

Darkroom-Large. We adopt the same setup as in *Darkroom*, yet with an expanded state space of 10×10 and a longer horizon $T = 100$. Consequently, the agent must explore the environment more extensively due to the sparse reward setting, making this task more challenging than *Darkroom*. We still consider 80% of the 100 available goals for pretraining and the remaining unseen 20% goals for test.

Miniworld. *Miniworld* is a three-dimensional pixel-based navigation task. The agent is situated in a room with four differently colored boxes, one of which is the target (unknown to the agent). The agent must navigate to the target box using $25 \times 25 \times 3$ image observations and by selecting from 4 available actions: turn left, turn right, move forward, or stay. Similar to *Darkroom*, the agent receives a reward of 1 only upon approaching the target box, and 0 otherwise. The high-dimensional pixel inputs clearly render *Miniworld* a much more challenging task than *Darkroom* and *Darkroom-Large*.

Hyperparameters

The main hyperparameters employed in this work are summarized in Tables 1-2.

Table 1: The main hyperparameters of each algorithm

Hyperparameters	AD	DPT	DIT	SAD
Causal transformer	GPT2	GPT2	GPT2	GPT2
Number of attention heads	3	3	3	3
Number of attention layers	3	3	3	3
Embedding size	32	32	32	32
Weight λ	N/A	N/A	500	N/A
Learning rate	0.001	0.001	0.001	0.001
Dropout	0.1	0.1	0.1	0.1

Additional Results

We provide in this subsection additional experimental results that are omitted in the main body of paper.

Figure 12 presents the dataset generation time of our SAD approach, compared with other SOTA baselines (AD, DPT, DIT), where the MAB and MDP problems consider the environments of *Gaussian Bandits* and *Darkroom* respectively. SAD (avg) represents the average of consumed time over different trust horizons (consistent with those in Figure 5). On average, SAD requires the most significant amount of time in the MDP problem. While in the MAB problem, SAD ranks as the second most time-consuming method, with its duration surpassing all except DIT. Notice that the additional computational time of SAD aligns with the prevailing trend of leveraging increased computation to fully harness the advanced reasoning capabilities of FMs (Brown 2020).

Table 2: The main hyperparameters of each environment

Hyperparameters	*Gaussian Bandits*	*Bernoulli Bandits*	*Darkroom*	*Darkroom-Large*	*Miniworld*
Action dimension	5	5	5	5	4
Pixel-based	✗	✗	✗	✗	✓
Trust Horizon	320	320	7	10	3
# of epochs	100	100	100	100	200
Context horizon	500	500	49	100	50
Pretraining/test ratio	0.8/0.2	0.8/0.2	0.8/0.2	0.8/0.2	0.8/0.2
# of environments	100000	100000	24010	100000	40000

Table 3: Performance improvements of SAD compared to baseline algorithms in the offline evaluation.

Environment	SAD vs AD	SAD vs DPT	SAD vs DIT	SAD vs DPT*
Gaussian Bandits	647.0%	508.9%	354.0%	−5.2%
Bernoulli Bandits	553.4%	426.8%	289.5%	−18.7%
Darkroom	2162.2%	2069.6%	149.3%	−1.3%
Darkroom-Large	6325.5%	6389.9%	266.8%	−14.9%
Miniworld	687.7%	684.2%	122.1%	−52.9%
Average	2075.2%	2015.9%	236.3%	−18.6%

Table 4: Performance improvements of SAD compared to baseline algorithms in the online evaluation.

Environment	SAD vs AD	SAD vs DPT	SAD vs DIT	SAD vs DPT*
Gaussian Bandits	933.6%	942.5%	273.9%	−0.4%
Bernoulli Bandits	846.8%	830.9%	313.9%	−0.2%
Darkroom	3053.9%	2893.8%	41.7%	−3.4%
Darkroom-Large	10626.9%	10221.8%	24.7%	−0.1%
Miniworld	582.9%	580.1%	21.7%	−57.3%
Average	3208.8%	3093.8%	135.2%	−12.3%

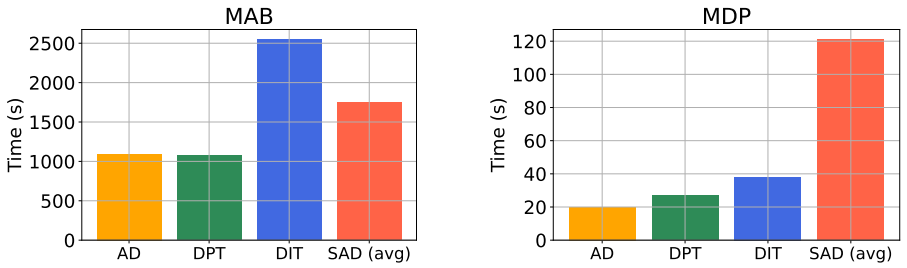

Figure 12: The dataset generation time consumed by SAD averaged over varying trust horizons (as in Figure 5), compared with AD, DPT, and DIT. Left: MAB, Right: MDP.