
Adaptive Multi-Goal Exploration

Jean Tarbouriech
Meta AI & Inria Scool

Omar Darwiche Domingues
Inria Scool

Pierre M enard
OvGU Magdeburg

Matteo Pirodda
Meta AI

Michal Valko
DeepMind

Alessandro Lazaric
Meta AI

Abstract

We introduce a generic strategy for provably efficient multi-goal exploration. It relies on ADAGOAL, a novel goal selection scheme that leverages a measure of uncertainty in reaching states to adaptively target goals that are neither too difficult nor too easy. We show how ADAGOAL can be used to tackle the objective of learning an ε -optimal goal-conditioned policy for the (initially unknown) set of goal states that are reachable within L steps in expectation from a reference state s_0 in a reward-free Markov decision process. In the tabular case with S states and A actions, our algorithm requires $\tilde{O}(L^3 SA\varepsilon^{-2})$ exploration steps, which is nearly minimax optimal. We also readily instantiate ADAGOAL in linear mixture Markov decision processes, yielding the first goal-oriented PAC guarantee with linear function approximation. Beyond its strong theoretical guarantees, we anchor ADAGOAL in goal-conditioned deep reinforcement learning, both conceptually and empirically, by connecting its idea of selecting “uncertain” goals to maximizing value ensemble disagreement.

1 INTRODUCTION

When the extrinsic reward signal is absent or non-informative, a reinforcement learning (RL) agent needs to explore the environment driven by objectives other than reward maximization (see e.g., Chentanez et al., 2005; Oudeyer and Kaplan, 2009; Singh et al., 2010).

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

This general setting is often called *unsupervised* RL, self-supervised RL or intrinsically motivated RL. A noteworthy instance of it is unsupervised *goal-conditioned* RL (GC-RL, see e.g., Colas et al., 2020, for a survey). In this framework, the agent must learn a goal-conditioned policy, which learns a distribution over actions conditioned not only on the current state but also on a goal state that it must reach as quickly as possible (in expectation). For the goal-conditioned policy to be able to reach a variety of goals in the unknown environment, the agent must autonomously set its own goals (via e.g., a curriculum) and learn to effectively reach them.

Deep GC-RL. Recently, GC-RL has been extensively studied in the context of deep RL (e.g., Schaul et al., 2015; Andrychowicz et al., 2017; Florensa et al., 2018; Warde-Farley et al., 2019; Nair et al., 2018; Colas et al., 2019; Zhao et al., 2019; Hartikainen et al., 2020; Ecoffet et al., 2021; Pong et al., 2020; Zhang et al., 2020b; Pitis et al., 2020). GC-RL has notably been shown to be a powerful framework to tackle navigation problems (Florensa et al., 2018), game playing (Ecoffet et al., 2021, on Montezuma’s Revenge) or real-world robotic manipulation tasks (Pong et al., 2020). At a high-level, these deep GC-RL methods follow the same algorithmic structure that alternates between:

- (**GS**) *Goal Selection*: select a candidate goal state;
- (**PE**) *Policy Execution*: deploy an explorative policy conditioned on this goal for a fixed number of steps.

Given this simple and unifying algorithmic structure, the core differences between the methods lie in the goal sampling distribution — e.g., uniform (Kaelbling, 1993; Schaul et al., 2015; Andrychowicz et al., 2017), proportional to rarity (Pong et al., 2020), or targeting goals of increasing complexity (Florensa et al., 2018; Colas et al., 2019; Zhang et al., 2020b) — and in ways to take advantage of each policy execution as much as possible to speed up the learning — e.g., goal relabeling (Andrychowicz et al., 2017) or encoding goal states in

lower-dimensional representations (Pong et al., 2020). The specific choice of **(GS)** and **(PE)** steps directly influence the learning speed as well as the quality of the goal-conditioned policy returned by the algorithm.

While these GC-RL approaches effectively leverage deep RL techniques and are able to achieve impressive results in complex domains, they often lack substantial theoretical understanding and guarantees, even when restricted to the tabular setting.

Theory of GC-RL. On the other hand, GC-RL has been rather sparsely analyzed through a theoretical lens. Notably, Lim and Auer (2012); Tarbouriech et al. (2020b) study the *incremental autonomous exploration* setting, where the agent is restricted to identifying and learning to efficiently reach only the goal states that are *incrementally* attainable from a reference initial state. Meanwhile, Tarbouriech et al. (2021a) study the problem of *goal-free cost-free exploration* (i.e., finding a near-optimal goal-reaching policy for any starting state, goal state and cost function) under the assumption that the Markov decision process (MDP) is *communicating* (i.e., any state is reachable from any other state).

While the existing theoretical GC-RL approaches come with rigorous guarantees on the learning performance, their algorithmic designs are quite involved, far from the interpretable alternation of **(GS)** and **(PE)** steps (see Appendix A for details, where we identify three specific shortcomings). This makes it hard to extract relevant practical insights and to make them amenable to a deep RL implementation.

Objective. In light of these conceptual and algorithmic discrepancies between the two branches of theoretical and deep GC-RL, we seek to:

Design a strategy for unsupervised goal-conditioned RL with both

- 1) *a simple and interpretable algorithmic structure whose conceptual idea can be adapted to deep RL, and*
- 2) *strong learning guarantees under suitable assumptions on the environment (e.g., tabular, linear).*

Contributions. Our main contributions can be summarized as follows:

- We formalize the *multi-goal exploration* (MGE) objective of minimizing the number of exploration steps (i.e., the sample complexity) required to learn a near-optimal goal-conditioned policy for all the goal states that are reachable within a given number of steps in expectation from the initial state. We formally motivate the need of an available reset action to the initial state so as to solve the objective in a reasonable number of exploration steps, and we derive a

lower bound on the sample complexity.

- We introduce ADA_{GOAL}, a novel *goal selection scheme* for unsupervised goal-conditioned RL. It relies on a simple optimization problem that adaptively targets goal states of intermediate difficulty. It also provides an algorithmic *stopping rule* and a *set of candidate goal states* that the agent is confident it can reliably reach. We identify some generic conditions that make ADA_{GOAL} theoretically sound.
- We design ADA_{GOAL}-UCBVI, an instantiation of ADA_{GOAL} in tabular MDPs, and prove that it achieves nearly minimax optimal sample complexity, improving over existing bounds that only hold in special cases.
- Leveraging a similar algorithmic and proof structure, we readily design ADA_{GOAL}-UCRL-VTR for linear-mixture MDPs. This yields the first goal-oriented PAC¹ guarantee with linear function approximation.
- We anchor ADA_{GOAL} in deep GC-RL, both conceptually and empirically, by connecting its idea of selecting “uncertain” goals to a practical approximation of maximizing value ensemble disagreement (Zhang et al., 2020b).

1.1 Additional Related Work

Theory of unsupervised exploration. A recent line of work theoretically analyzes unsupervised exploration in RL (in other settings than GC-RL). Some approaches (e.g., Hazan et al., 2019; Tarbouriech and Lazaric, 2019; Cheung, 2019; Tarbouriech et al., 2020c) learn to induce a desired state(-action) distribution (e.g., maximally entropic), yet they do not analyze how this can be used to solve downstream tasks. In finite-horizon MDPs, Jin et al. (2020) propose a paradigm composed of: an *exploration phase*, where the agent interacts with the environment without the supervision of reward signals; followed by a *planning phase*, where the agent is required to compute a near-optimal policy for some given reward function. If the reward function can be chosen arbitrarily (including adversarially), the problem is called *reward-free exploration* (RFE) (Jin et al., 2020; Kaufmann et al., 2021; Ménard et al., 2021a; Zanette et al., 2020; Wang et al., 2020; Chen et al., 2021b; Zhang et al., 2021a,b). If there is only a finite number of rewards that are fixed a priori yet unknown during exploration, it is called *task-agnostic exploration* (TAE) (Zhang et al., 2020a; Wu et al., 2021b,a). Our MGE problem bears some resemblance to TAE in the sense that we also consider a finite number of unknown tasks to solve (specifically, finding a near-optimal goal-conditioned policy for the unknown

¹Recall that the probably approximately correct (PAC) learning setting provides sample complexity guarantees to find a near-optimal policy at the fixed initial state.

set of reliably reachable goal states). However, TAE (as well as RFE) is limited to the *finite-horizon* setting and does not extend to goal-reaching tasks, which belong to the category of stochastic shortest path problems.

Single-goal exploration (a.k.a. SSP). A special case of multi-goal exploration is when there is a single goal state that is extrinsically fixed throughout learning (i.e., there is no **(GS)**). This is known as the stochastic shortest path (SSP) problem (Bertsekas, 1995). A recent line of work has studied online learning in SSP in the context of regret minimization (Tarbouriech et al., 2020a; Rosenberg et al., 2020; Cohen et al., 2021; Tarbouriech et al., 2021c; Jafarnia-Jahromi et al., 2021; Chen et al., 2021a; Vial et al., 2021; Min et al., 2021). However, as argued by Tarbouriech et al. (2021b), a standard regret-to-PAC conversion does not hold in general in SSP, which implies that these techniques cannot be directly applied to our case. See Section 5 for further discussion.

2 MULTI-GOAL EXPLORATION

The agent interacts with an environment modeled by a Markov decision process (MDP) \mathcal{M} that has no extrinsic reward (i.e., it is *reward-free*), a finite state space \mathcal{S} with $S \triangleq |\mathcal{S}|$ states and a finite action space \mathcal{A} with $A \triangleq |\mathcal{A}|$ actions. Let $s_0 \in \mathcal{S}$ be a designated initial state. We denote by $p(s'|s, a)$ the probability transition from state s to state s' by taking action a .

A deterministic stationary policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ is a mapping between states to actions and we denote by Π the set of all possible policies. We measure the performance of a policy in navigating the MDP and define the shortest-path distance as follows.

Definition 1. For any policy $\pi \in \Pi$ and a pair of states $(s, s') \in \mathcal{S}^2$, let $V^\pi(s \rightarrow s') \in [0, +\infty]$ be the expected number of steps it takes to reach s' starting from s when executing policy π , i.e.,²

$$V^\pi(s \rightarrow s') \triangleq \mathbb{E} \left[\omega_\pi(s \rightarrow s') \right],$$

$$\omega_\pi(s \rightarrow s') \triangleq \inf \left\{ i \geq 0 : s_{i+1} = s' \mid s_1 = s, \pi \right\},$$

where the expectation is w.r.t. the random sequence of states generated by executing π in \mathcal{M} starting from state s . For any state $g \in \mathcal{S}$, let $V^*(s_0 \rightarrow g) \in [0, +\infty]$ be the shortest-path distance from s_0 to g , i.e.,

$$V^*(s_0 \rightarrow g) \triangleq \min_{\pi \in \Pi} V^\pi(s_0 \rightarrow g).$$

²Note that $V^\pi(s \rightarrow s')$ corresponds to the value function (a.k.a. expected cost-to-go) of policy π in a stochastic shortest-path setting (SSP, Bertsekas, 1995) with initial state s , goal state s' and unit cost function. If the policy π reaches s' from s with probability 1, it is said to be *proper*, otherwise it holds that $V^\pi(s \rightarrow s') = +\infty$.

Let $D_0 \in [0, +\infty]$, $D \in [0, +\infty]$ be the MDP's (possibly infinite) s_0 -diameter and diameter, respectively, i.e.,

$$D_0 \triangleq \max_{g \in \mathcal{S}} V^*(s_0 \rightarrow g), \quad D \triangleq \max_{s, g} V^*(s \rightarrow g).$$

We denote by $\mathcal{G} \subseteq \mathcal{S}$ the *goal space*, which corresponds to the set of goal states that the agent may condition its goal-conditioned policy on (in the absence of prior knowledge on the goal space we simply set $\mathcal{G} = \mathcal{S}$). In environments with arbitrary dynamics, there may be some goal states in \mathcal{G} that are *too* difficult for the agent to reliably reach in a reasonable number of exploration steps, or even completely unreachable from s_0 . Consequently, we consider the high-level objective of *learning an accurate goal-conditioned policy for all the goal states that are reliably reachable* from s_0 .

Definition 2 (Reliably L -reachable goal states \mathcal{G}_L). For any threshold $L \geq 1$, we define a goal state $g \in \mathcal{G}$ to be *reliably L -reachable* if $V^*(s_0 \rightarrow g) \leq L$, and we denote by \mathcal{G}_L the set of such goal states, i.e.,

$$\mathcal{G}_L \triangleq \left\{ g \in \mathcal{G} : V^*(s_0 \rightarrow g) \leq L \right\}.$$

We thus seek to learn a goal-conditioned policy that is accurate in reaching the goals in \mathcal{G}_L . A challenge in solving this objective comes from the fact that the set of goals of interest \mathcal{G}_L is initially *unknown* and it has to be discovered online at the same time as learning their corresponding optimal policy. The threshold L can be interpreted as the user's exploration radius of interest around s_0 . In the absence of a pre-specified threshold, the agent can build its own curriculum for L to guide its learning process.

Since in environments with arbitrary dynamics the agent may get stuck in a state without being able to return to s_0 , we introduce the following "reset" assumption.³ In Lemma 5 we will formally motivate its role in solving our learning objective.

Assumption 3. The action space contains a known action $a_{\text{reset}} \in \mathcal{A}$ such that $p(s_0|s, a_{\text{reset}}) = 1$ for any state $s \in \mathcal{S}$.

Consider as input an exploration radius $L \geq 1$, an accuracy level $\varepsilon \in (0, 1]$ and a confidence level $\delta \in (0, 1)$. We now formally define our exploration objective.

Definition 4 (Multi-Goal Exploration — MGE). An algorithm is said to be $(\varepsilon, \delta, L, \mathcal{G})$ -PAC for MGE if

- it stops after some (possibly random) number of exploration steps τ that is less than some polynomial

³This setting should be contrasted with the finite-horizon setting, where each policy resets automatically after H steps, or assumptions on the MDP dynamics such as ergodicity or bounded diameter, which guarantee that it is always possible to find a policy navigating between any two states.

in the relevant quantities $(S, A, L, \varepsilon^{-1}, \log \delta^{-1})$ with probability at least $1 - \delta$,

- it returns a set of goal states \mathcal{X} and a set of policies $\{\hat{\pi}_g\}_{g \in \mathcal{X}}$ such that $\mathbb{P}(\mathcal{C}_1 \cap \mathcal{C}_2) \geq 1 - \delta$, where we define the conditions

$$\mathcal{C}_1 \triangleq \left\{ \forall g \in \mathcal{X}, V^{\hat{\pi}_g}(s_0 \rightarrow g) - V^*(s_0 \rightarrow g) \leq \varepsilon \right\},$$

$$\mathcal{C}_2 \triangleq \left\{ \mathcal{G}_L \subseteq \mathcal{X} \subseteq \mathcal{G}_{L+\varepsilon} \right\}.$$

The objective is to build an $(\varepsilon, \delta, L, \mathcal{G})$ -PAC algorithm for which the MGE sample complexity, that is the number of exploration steps τ , is as small as possible.

Remark 1. Since the goal set \mathcal{G}_L is *unknown*, it may not be possible to exactly *identify* it within a reasonable number of exploration steps. Thus we allow the learner to output a larger set \mathcal{X} of candidate goals and policies. Nonetheless, we constrain an $(\varepsilon, \delta, L, \mathcal{G})$ -PAC algorithm for MGE to return a set \mathcal{X} that is at most contained in the slightly larger set $\mathcal{G}_{L+\varepsilon}$ (i.e., $\mathcal{X} \subseteq \mathcal{G}_{L+\varepsilon}$).

Remark 2. Consider that $\mathcal{G} = \mathcal{S}$. Then the inclusion $\mathcal{G}_L \subseteq \mathcal{S}$ is an equality if \mathcal{M} is communicating (i.e., $D < +\infty$) and if the (unknown) D_0 is lower or equal to L (note that under Assumption 3, $D \leq D_0 + 1$). Moreover, denote by $\mathcal{S}_L^{\rightarrow}$ the set of incrementally L -reachable states defined by Lim and Auer (2012, Definition 5). Then in the special case where $\mathcal{S}_L^{\rightarrow} = \mathcal{G}_L$, the MGE objective of Definition 4 is equivalent to the AX* objective proposed by Tarbouriech et al. (2020b, Definition 5) in the incremental autonomous exploration setting introduced by Lim and Auer (2012).

MGE vs. reset-free MGE. Lemma 5 establishes an exponential separation between MGE and reset-free MGE (i.e., MGE without Assumption 3). This motivates the use of Assumption 3 to solve our learning objective in a reasonable number of exploration steps.

Lemma 5. MGE can be solved in $\text{poly}(S, L, \varepsilon^{-1}, A)$ steps. On the other hand, there exists an MDP and a goal space where any algorithm requires at least $\Omega(D)$ steps to solve reset-free MGE, where D is exponentially larger than $L, S, A, \varepsilon^{-1}$.

MGE lower bound. We now give a worst-case lower bound on the MGE problem (details in Appendix B).

Lemma 6. For any algorithm that is $(\varepsilon, \delta, L, \mathcal{G})$ -PAC for MGE for any MDP and goal space \mathcal{G} , there exists an MDP and a goal space where the algorithm requires, in expectation, at least $\Omega(L^3 S A \varepsilon^{-2})$ exploration steps to stop.

Remark 3. We can relate the dependencies in Lemma 6 with the lower bound of the time steps needed to identify an ε -optimal policy in both γ -discounted MDPs with a generative model — i.e.,

$\Omega((1 - \gamma)^{-3} S A \varepsilon^{-2})$ (Azar et al., 2013) — and online stationary finite-horizon MDPs — i.e., $\Omega(H^3 S A \varepsilon^{-2})$ (Domingues et al., 2021b). This correspondence is not surprising as L captures the “range” of the MGE problem, similar to the effective horizon $1/(1 - \gamma)$ or the horizon H . Also recall that both discounted MDPs and finite-horizon MDPs are subclasses of goal-oriented MDPs, i.e., SSP-MDPs (Bertsekas, 1995).

MGE under linear function approximation.

Lemma 6 shows that the MGE sample complexity must scale with SA in the worst case, which may be prohibitive in the case of large state-action spaces. This motivates us to further analyze MGE under *linear function approximation*. In particular, we focus on the *linear mixture* MDP setting (Ayoub et al., 2020; Zhou et al., 2021), which assumes that the transition probability is a linear mixture of d signed basis measures.

Definition 7 (Linear Mixture MDP, Ayoub et al., 2020; Zhou et al., 2021). *The unknown transition probability p is a linear combination of d signed basis measures $\phi_i(s'|s, a)$, i.e., $p(s'|s, a) \triangleq \sum_{i=1}^d \phi_i(s'|s, a) \theta_i^*$. Meanwhile, for any $V : \mathcal{S} \rightarrow [0, 1]$, $i \in [d]$, $(s, a) \in \mathcal{S} \times \mathcal{A}$, the summation $\sum_{s' \in \mathcal{S}} \phi_i(s'|s, a) V(s')$ is computable. For simplicity, let $\phi \triangleq [\phi_1, \dots, \phi_d]^\top$, $\theta^* \triangleq [\theta_1^*, \dots, \theta_d^*]^\top$ and $\psi_V(s, a) \triangleq \sum_{s' \in \mathcal{S}} \phi(s'|s, a) V(s)$. Without loss of generality, we assume $\|\theta^*\|_2 \leq B$, $\|\psi_V(s, a)\|_2 \leq 1$ for all $V : \mathcal{S} \rightarrow [0, 1]$ and $(s, a) \in \mathcal{S} \times \mathcal{A}$.*

3 OUR ADAGoAL APPROACH

In Algorithm 1, we introduce the common algorithmic structure based on ADAGoAL. We use it to design ADAGoAL-UCBVI that tackles the MGE problem in tabular MDPs, and ADAGoAL-UCRL-VTR that tackles the MGE problem in linear mixture MDPs. Both follow the goal-conditioned structure described in Section 1. The agent sets a horizon of $H = \Omega(L \log L \varepsilon^{-1})$ and splits its learning interaction in algorithmic episodes of length H . At the beginning of each algorithmic episode, it (**GS**) selects a candidate goal state and (**PE**) deploys an explorative (i.e., optimistic) policy conditioned on this goal for H steps before resetting to s_0 . It alternates between these two steps until an adaptive stopping rule is met, at which point the algorithm terminates.⁴

□ (**PE**) **step.** Goal-conditioned finite-horizon Q -functions (Kaelbling, 1993; Schaul et al., 2015) are maintained optimistically. At each episode k and episode step $h \in [H]$, $Q_{k,h}(s, a, g)$ approximates (from below) the number of (expected) steps required to reach any goal $g \in \mathcal{G}$ starting from any state-action

⁴Indeed recall from Definition 4 that the algorithm must adaptively decide when to terminate its learning interaction.

Algorithm 1: ADA_{GOAL}-based algorithmic structure. Blue text denotes ADA_{GOAL}-UCBVI specific steps and purple text denotes ADA_{GOAL}-UCRL-VTR specific steps.

- 1 **Input:** Exploration radius $L \geq 1$, accuracy level $\varepsilon \in (0, 1]$, confidence level $\delta \in (0, 1)$.
- 2 **Input:** Number of states S , number of actions A .
- 3 **Input:** Dimension of feature mapping d , bound B on ℓ_2 -norm of θ^* .
- 4 **Input:** Goal space $\mathcal{G} \subseteq \mathcal{S}$ (otherwise set $\mathcal{G} = \mathcal{S}$).
- 5 Set as horizon $H \triangleq \lceil 5(L+2) \log(10(L+2)/\varepsilon) / \log(2) \rceil$.
- 6 **Initialize:** algorithmic episode index $k = 1$, distance estimates $\mathcal{D}_1(g) = \mathbb{1}[g \neq s_0]$, error estimates $\mathcal{E}_1(g) = H\mathbb{1}[g \neq s_0]$, goal-conditioned finite-horizon Q -values $\mathcal{Q}_{1,h}(s, a, g) = \mathbb{1}[s \neq g]$, for all $(g, s, a, h) \in \mathcal{G} \times \mathcal{S} \times \mathcal{A} \times [H]$.
- 7 **while** stopping rule (1) is not met **do**
- 8 **(i) Goal selection rule:**
- 9 Select as goal state

$$g_k \in \arg \max_{g \in \mathcal{G}} \mathcal{E}_k(g)$$
 subject to: $\mathcal{D}_k(g) \leq L$.
- 10 **(ii) Policy execution rule:**
- 11 For a duration of H steps, run the optimistic goal-conditioned policy $\pi_{g_k}^k$ such that at step h , $\pi_{g_k,h}^k(s) \in \arg \min_{a \in \mathcal{A}} \mathcal{Q}_{k,h}(s, a, g_k)$ (note that $\mathcal{D}_k(g_k) = \min_{a \in \mathcal{A}} \mathcal{Q}_{k,1}(s_0, a, g_k)$).
- 12 Then execute action a_{reset} and increment episode index $k += 1$.
- 13 **(iii) Update and check stopping rule:**
- 14 Update estimates $\mathcal{Q}_k, \mathcal{D}_k, \mathcal{E}_k$ according to (8), (9), (10) using samples collected so far.
- 15 Update estimates $\mathcal{Q}_k, \mathcal{D}_k, \mathcal{E}_k$ according to (29), (30), (31) using samples collected so far.
- 16 Stop the algorithm if

$$\max_{g \in \mathcal{G}: \mathcal{D}_k(g) \leq L} \mathcal{E}_k(g) \leq \varepsilon. \quad (1)$$
- 17 **end**
- 18 **Output:** Goal states $\mathcal{X}_\kappa \triangleq \{g \in \mathcal{G} : \mathcal{D}_\kappa(g) \leq L\}$, and for every $g \in \mathcal{X}_\kappa$, a deterministic, non-stationary policy $\hat{\pi}_g$ that at time step i and state s selects action a according to:

$$\hat{\pi}_g(a|s, i) \triangleq \begin{cases} \arg \min_{a \in \mathcal{A}} \mathcal{Q}_{\kappa,h}(s, a, g) & \text{if } i \equiv h \pmod{H+1} \text{ for } h \in [H], \\ a_{\text{reset}} & \text{if } i \equiv 0 \pmod{H+1}. \end{cases}$$

pair $(s, a) \in \mathcal{S} \times \mathcal{A}$ and executing the optimal goal-reaching policy for $H - h$ steps. Intuitively, the Q -functions will gradually increase, more so for goal states that the agent struggles to reach. This is essentially done by initializing the Q -functions optimistically (i.e., at 0), considering that the cost (i.e., negative reward) is +1 (resp. 0) per time step if the conditioned goal is not reached (resp. reached), and carefully subtracting an exploration bonus to maintain optimism. Given a goal $g_k \in \mathcal{G}$ selected at the beginning of episode k , the **(PE)** step simply amounts to deploying an explorative policy conditioned on g_k , that is, a policy $\pi_{k,h}$ that greedily minimizes the current Q -functions, i.e., $\pi_{k,h}(s) \in \arg \min_{a \in \mathcal{A}} \mathcal{Q}_{k,h}(s, a, g_k)$.

(GS) step. To elect a relevant sequence of candidate goals $(g_k)_{k \geq 1}$, we introduce ADA_{GOAL}, an adaptive goal selection scheme based on a simple constrained optimization problem. It relies on the agent's ability to compute two types of goal-conditioned quantities for any goal $g \in \mathcal{G}$ and episode $k \geq 1$:

- a distance estimate from s_0 to g , denoted by $\mathcal{D}_k(g)$;
- an error of estimating this distance, denoted by $\mathcal{E}_k(g)$.

Conveniently, the distance estimates can be simply instantiated as $\mathcal{D}_k(g) \triangleq \min_{a \in \mathcal{A}} \mathcal{Q}_{k,1}(s_0, a, g)$. Formally, we require the following properties of \mathcal{D} and \mathcal{E} to hold (with high probability):

- **Property 1:** \mathcal{D} is an *optimistic distance estimate*, i.e.,

$$\mathcal{D}_k(g) \leq \mathcal{D}_H^*(g), \quad \forall k \geq 1, \forall g \in \mathcal{G},$$

where $\mathcal{D}_H^*(g) \triangleq \min_{\pi} \mathbb{E}[\omega_{\pi}(s_0 \rightarrow g) \wedge H]$ denotes the shortest-path distance from s_0 to g truncated at H steps. Note that $\mathcal{D}_H^*(g) \in [0, H]$ and that $\lim_{H \rightarrow +\infty} \mathcal{D}_H^*(g) = V^*(s_0 \rightarrow g)$.

- **Property 2:** \mathcal{E} is an *upper bound on the prediction error*, i.e.,

$$|\mathcal{D}_H^*(g) - \mathcal{D}_k(g)| \leq \mathcal{E}_k(g), \quad \forall k \geq 1, \forall g \in \mathcal{G}.$$

Given these two goal-conditioned quantities, ADA_{GOAL} selects at episode k a candidate goal that solves the following constrained optimization problem:

$$g_k \in \arg \max_{g \in \mathcal{G}} \mathcal{E}_k(g) \quad (2a)$$

$$\text{subject to: } \mathcal{D}_k(g) \leq L. \quad (2b)$$

Interpretation. The agent sequentially selects a goal state with highest prediction error \mathcal{E} among those whose distance estimate \mathcal{D} is not too large. If the agent is confident that a goal g is either too easy or too hard to reach, it will assign a low prediction error $\mathcal{E}(g)$. As a




Figure 1: Goal sampling frequency of ADA GOAL-UCBVI over 1000 episodes of length $H = 50$ (with $L = 40$). The grid-world has $S = 52$ states, starting state $s_0 = (0, 0)$ (top left), $A = 5$ actions (4 cardinal ones and a_{reset}). The 4 states of the bottom right room can only be accessed from s_0 by any cardinal action with probability $\eta = 0.001$ (their associated $V^*(s_0 \rightarrow \cdot)$ thus scale with η^{-1}).

result, the objective function in (2a) adaptively samples goal states on the frontier of the learning process. The constraint in (2b), although it is not required for the final sample complexity result, further tightens the goal selection process. Indeed, for any $k \geq 1$, let

$$\mathcal{X}_k \triangleq \{g \in \mathcal{G} : \mathcal{D}_k(g) \leq L\}, \quad \varepsilon_k \triangleq \max_{g \in \mathcal{X}_k} \mathcal{E}_k(g).$$

Then, if as a warm-up we take the limit $H \rightarrow +\infty$, injecting Properties 1 and 2 in (2a-2b) entails that

$$\mathcal{G}_L \subseteq \mathcal{X}_k \subseteq \mathcal{G}_{L+\varepsilon_k}. \quad (3)$$

We thus see that the constraint in (2b) does not remove valid goals in \mathcal{G}_L from the set \mathcal{X}_k of candidate goal states to sample. Second, decreasing ε_k has the dual impact of making the set of candidates goals \mathcal{X}_k closer to \mathcal{G}_L and improving their distance estimates, which motivates the goal selection scheme in (2a). In practice, we consider a *finite* truncation H (line 5 in Algorithm 1), thus we need to account for the bias $\rho_g \triangleq V^*(s_0 \rightarrow g) - \mathcal{D}_H^*(g)$, which can be arbitrarily large for goals g that are hard or impossible to reach. Fortunately, our ADA GOAL strategy will be able to gradually discard such states, hence the final MGE sample complexity will not pay for such terms. In fact, we will later see that the choice of horizon $H = \Omega(L \log L \varepsilon^{-1})$ ensures that $\rho_g = O(\varepsilon)$ for all the (unknown) goal states of interest $g \in \mathcal{G}_L$.

Choice of $\mathcal{Q}, \mathcal{D}, \mathcal{E}$. A key algorithmic design is how to build and update the goal-conditioned \mathcal{Q} -functions and the estimates \mathcal{D} and \mathcal{E} . In Appendices C and D, we will carefully construct them with exact bonus-based estimates for both tabular MDPs and linear mixture MDPs. As suggested by the algorithms’ names, the estimates of the former are inspired from BPI-UCBVI (Ménard et al., 2021a), an algorithm for best policy identification in finite-horizon tabular MDPs, while those of the latter are inspired from UCRL-VTR (Ayoub et al., 2020), an algorithm for regret minimization in finite-horizon linear mixture MDPs. Since all our estimates are optimistic, we see that Algorithm 1 relies on the principle of optimism in the face of uncertainty *both* for the goal selection and the policy execution. Finally, in Section 6, we propose a way to instantiate $\mathcal{Q}, \mathcal{D}, \mathcal{E}$ for a practical implementation in deep RL.

□ **Adaptive stopping rule.** At the end of each algorithmic episode, the estimates are updated using

the samples collected so far, and the algorithm checks whether a stopping rule (1) based on ADA GOAL is triggered, in which case it terminates. This occurs when the prediction errors \mathcal{E} of all the goal states that meet the ADA GOAL constraint (2b) are below the prescribed accuracy level ε . These states then form the set of *candidate goal states* output by Algorithm 1, along with their associated optimistic goal-reaching policies.

Empirical validation. In Figure 1 (see Appendix E for details), we empirically study the sequence of goals selected by ADA GOAL-UCBVI during learning. We design a two-room grid-world with a very small probability of reaching the second room. We see that ADA GOAL is able to discard the states from the second room and target as goals the states in the first room that are “furthest” away from s_0 , which effectively correspond to the fringe of what the agent can reliably reach.

4 SAMPLE COMPLEXITY GUARANTEES

□ **Guarantee for ADA GOAL-UCBVI.** We first bound the MGE sample complexity of ADA GOAL-UCBVI. For simplicity we consider that $\mathcal{G} = \mathcal{S}$, i.e., the goal space spans the entire state space (the results trivially extend to any $\mathcal{G} \subseteq \mathcal{S}$).

Theorem 8. ADA GOAL-UCBVI is $(\varepsilon, \delta, L, \mathcal{S})$ -PAC for MGE and, with probability at least $1 - \delta$, for $\varepsilon \in (0, 1/S]$ its MGE sample complexity is of order⁵ $\tilde{O}(L^3 S A \varepsilon^{-2})$.

Lemma 6 and Theorem 8 imply that the MGE sample complexity of ADA GOAL-UCBVI is nearly minimax optimal for small enough ε and up to logarithmic terms.

In the absence of a pre-specified exploration radius L , the agent can build its own curriculum for L (i.e., design a sequence of increasing L ’s) to guide its learning. In this case, the total sample complexity is (up to a logarithmic factor) the same of ADA GOAL-UCBVI run with the final value of L , as stated below.

Corollary 9. The successive execution of ADA GOAL-UCBVI for an increasing sequence $L \in \{2, 2^2, \dots, 2^f\}$

⁵The notation \tilde{O} in Theorem 8 hides poly-log terms in $\varepsilon^{-1}, S, A, L, \delta^{-1}$. See Lemma 23 in Appendix C.3 for a more detailed bound that includes the poly-log terms.

with $f \in \mathbb{N}^*$ is $(\varepsilon, \delta, L_f, \mathcal{S})$ -PAC for MGE and, with probability at least $1 - \delta$, for $\varepsilon \in (0, 1/S]$ its MGE sample complexity is of order $\tilde{O}(L_f^3 S A \varepsilon^{-2})$, where $L_f = 2^f$.

Finally, we can investigate the special case where \mathcal{M} is communicating and the objective is to learn an ε -optimal goal-conditioned policy for *every* goal state. Since the s_0 -diameter D_0 is unknown, we can use as an initial subroutine the GOSPRL algorithm (Tarbouriech et al., 2021a) to compute an estimate \tilde{D} such that $D_0 \leq \tilde{D} \leq 2D_0$ in $\tilde{O}(D_0^3 S^2 A)$ time steps. Then we can execute ADA_{GOAL}-UCBVI with $L = \tilde{D}$, which leads to the following guarantee.

Corollary 10. *Assume that the MDP \mathcal{M} has a finite and unknown s_0 -diameter D_0 . Then the above strategy is $(\varepsilon, \delta, D_0, \mathcal{S})$ -PAC for MGE and, with probability at least $1 - \delta$, for $\varepsilon \in (0, 1/S]$, its MGE sample complexity is of order $\tilde{O}(D_0^3 S A \varepsilon^{-2})$.*

Comparison with existing bounds. A different although related problem is considered in Tarbouriech et al. (2021a) in *communicating* MDPs, where the agent must find an ε -optimal goal-conditioned policy for any arbitrary starting state, goal state and positive cost function. For small enough ε , their bound in the unit-cost case scales as $\tilde{O}(D^4 \Gamma S A \varepsilon^{-2})$, where Γ is the branching factor which in the worst case is S . We see that Corollary 10 improves over Tarbouriech et al. (2021a) by a factor D_0 as well as a factor $\Gamma \leq S$. Although the latter considers a more demanding cost-free objective (i.e., for any positive cost function), it is unable to avoid its superlinear dependence in S when instantiated in our scenario of unit cost functions, since the algorithm’s design is to estimate uniformly well the transition kernel. Finally, DISCO (Tarbouriech et al., 2020b) designed for incremental autonomous exploration can solve MGE in the special case where $\mathcal{S}_L^\rightarrow = \mathcal{G}_L$, where $\mathcal{S}_L^\rightarrow$ denotes the set of incrementally L -reachable states (see Lim and Auer, 2012, Definition 5). However, the bound of DISCO would translate to $\tilde{O}(L^5 \Gamma S A \varepsilon^{-2})$, which is also worse than Theorem 8.

□ **Guarantee for ADA_{GOAL}-UCRL-VTR.** We now bound the MGE sample complexity of Algorithm 1 in linear mixture MDPs (Definition 7). Since the state space \mathcal{S} may be large, we consider that the known goal space is in all generality a subset of it, i.e., $\mathcal{G} \subseteq \mathcal{S}$, where $G \triangleq |\mathcal{G}|$ denotes the cardinality of the goal space.

Theorem 11. *In linear mixture MDPs, for $\varepsilon \in (0, 1]$, ADA_{GOAL}-UCRL-VTR is $(\varepsilon, \delta, L, \mathcal{G})$ -PAC for MGE and, with probability at least $1 - \delta$, its MGE sample complexity is of order⁶ $\tilde{O}(L^4 d^2 \varepsilon^{-2})$, where d is the dimension of the feature mapping.*

⁶The notation \tilde{O} in Theorem 11 hides poly-log terms in $\varepsilon^{-1}, G, d, L, \delta^{-1}$.

To the best of our knowledge, Theorem 11 yields the first goal-oriented PAC guarantee with linear function approximation. The algorithm’s choice of \mathcal{E} , \mathcal{D} , \mathcal{Q} relies on two regression-based goal-conditioned estimators, one standard “value-targeted” estimator inspired from UCRL-VTR (Ayoub et al., 2020) and one novel “error-targeted” estimator, see Appendix D. We expect that the bound of Theorem 11 can be refined using tighter Bernstein-based estimates, for instance inspired from UCRL-VTR⁺ (Zhou et al., 2021), which we leave as future work. Note that the \tilde{O} notation in Theorem 11 contains a $\log(G)$ factor (which appears when performing a union bound argument over all goals $g \in \mathcal{G}$), and the computational complexity of ADA_{GOAL}-UCRL-VTR scales with G (since the algorithm maintains goal-conditioned estimates). We point out that here we still consider that the MDP has a finite number of states S . This is to be expected given the way a goal is currently modeled (at the granular level of states), independently of how large the state space is, where it may be very hard to visit specific states. Learning in single- or multi-goal RL beyond a finite state space is an interesting direction of future investigation. Note that existing works on single-goal exploration (i.e., SSP regret minimization) with linear function approximation (Vial et al., 2021; Min et al., 2021) also assume that the state space is finite.

5 ANALYSIS OVERVIEW

The proofs of Theorems 8 and 11 (see Appendices C and D) are decomposed in the same following key steps.

▷ **Key step ①: Optimism and gap bounds.** We prove that Properties 1 and 2 hold with high probability, i.e., (i) the quantities \mathcal{D} are optimistic estimates of the optimal goal-conditioned finite-horizon value functions \mathcal{D}_H^* and (ii) the quantities \mathcal{E} are valid upper bounds to the goal-conditioned finite-horizon gaps.

▷ **Key step ②: Bounding the cumulative gap bounds.** We make explicit a function $f_{\mathcal{M}}$ (depending on the MDP \mathcal{M}) that is strictly decreasing in K with $f_{\mathcal{M}}(K) \rightarrow_{K \rightarrow \infty} 0$, such that with high probability,

$$\sum_{k=1}^K \mathcal{D}_H^*(g_k) - \mathcal{D}_k(g_k) \leq \sum_{k=1}^K \mathcal{E}_k(g_k) \leq K \cdot f_{\mathcal{M}}(K), \quad (4)$$

for any number of algorithmic episodes K . Specifically, we establish that $f_{\mathcal{M}}(K) = \tilde{O}(\sqrt{KH^2 SA} + H^2 S^2 A)$ for ADA_{GOAL}-UCBVI, and $f_{\mathcal{M}}(K) = \tilde{O}(\sqrt{KH^3 d^2} + H^2 d^{3/2})$ for ADA_{GOAL}-UCRL-VTR. (4) resembles a *no-regret* property of the exploration algorithm that receives as input the sequence of goals $(g_k)_{k \geq 1}$ prescribed

by ADAGoAL and performs the **(PE)** step. Indeed, intuitively, the aim of the **(PE)** step is to improve the estimation of $\mathcal{D}_k(g_k)$ and make it closer to $\mathcal{D}_H^*(g_k)$, i.e., to decrease the prediction error $\mathcal{E}_k(g_k)$.

▷ **Key step ③: Bounding the sample complexity.**

To bound κ the episode index at which Algorithm 1 terminates, we combine (4) and the termination condition (1) to simultaneously lower and upper bound (with high probability) the cumulative errors \mathcal{E} as

$$\varepsilon \cdot (\kappa - 1) \leq \sum_{k=1}^{\kappa-1} \mathcal{E}_k(g_k) \leq (\kappa - 1) \cdot f_{\mathcal{M}}(\kappa - 1).$$

Inverting this functional inequality in κ yields that κ is finite and bounded as

$$\kappa \leq f_{\mathcal{M}}^{-1}(\varepsilon) + 2. \quad (5)$$

The sample complexity is $(H + 1) \cdot \kappa$ with $H = \tilde{O}(L)$, thus ADAGoAL-UCBVI (resp. ADAGoAL-UCRL-VTR) stops in $\tilde{O}(L^3 S A \varepsilon^{-2} + L^3 S^2 A \varepsilon^{-1})$ (resp. $\tilde{O}(L^4 d^2 \varepsilon^{-2})$) time steps, with high probability.

▷ **Key step ④: Connecting to the original MGE objective.** The key remaining step is to prove that the MGE objective is indeed fulfilled.

Remark 4. Algorithm 1 relies on a finite-horizon construction, with algorithmic episodes of length H . This relates to the reduction of SSP to finite-horizon studied in some SSP regret minimization works (Cohen et al., 2021; Chen and Luo, 2021), which rely on the idea that an SSP problem can be approximated by a finite-horizon problem if the horizon is large enough w.r.t. T_* , the optimal expected hitting time to the goal starting from any state. Two main differences arise in our MGE setting: (i) first, in these works, the goal state is fixed throughout learning and T_* is assumed known, whereas we need to deal with goal selection and find the relevant goals of interest while having to discard those that are poorly reachable or unreachable. (ii) Second, these works ensure that the *empirical* goal-reaching performance of the algorithm’s non-stationary policy *over the whole learning interaction* is good enough (by definition of the regret objective in SSP). As such, they do not show that the *expected* performance of some *candidate* policy is good enough (i.e., the SSP value function is small enough) – in fact, they do not even explicitly prove that the executed policies are proper. The latter property may actually not be possible to obtain since standard regret-to-PAC conversion may not work in SSP (Tarbouriech et al., 2021b). In our MGE objective, the key difference lies in the availability of the reset action (Assumption 3), as we will now see.

Our analysis builds on the following reasoning: given a goal state in $\mathcal{G}_{L+\varepsilon}$, we can find a candidate policy with

near-optimal goal-reaching behavior (i.e., SSP value function) by: (i) first computing a near-optimal policy $\tilde{\pi}$ in the finite-horizon reduction (using the stopping rule (1)), (ii) and then expanding $\tilde{\pi}$ into an infinite-horizon policy via the reset action every H time steps to get our desired candidate policy.

Now, importantly, the above reasoning *only holds* for the goals in $\mathcal{G}_{L+\varepsilon}$, which is an *unknown* set. This is where our ADAGoAL strategy comes into the picture, as it provides a simple and computable *sufficient condition* for a goal to belong to $\mathcal{G}_{L+\varepsilon}$.

Lemma 12. *With probability at least $1 - \delta$, if a goal state $g \in \mathcal{G}$ satisfies $\mathcal{D}_k(g) \leq L$ and $\mathcal{E}_k(g) \leq \varepsilon$ for an episode $k \geq 1$, then $g \in \mathcal{G}_{L+\varepsilon}$.*

We are now ready to put everything together and prove that ADAGoAL-UCBVI and ADAGoAL-UCRL-VTR are (ε, δ, L) -PAC for MGE. The candidate goal states are $\mathcal{X}_\kappa \triangleq \{g \in \mathcal{S} : \mathcal{D}_\kappa(g) \leq L\}$, with candidate policies $\hat{\pi}_g \triangleq (\pi_g^{\kappa+1})^{|H}$. In what follows we reason with high probability. Property 1 ensures that $\mathcal{G}_L \subseteq \mathcal{X}_\kappa$, while Lemma 12 entails that $\mathcal{X}_\kappa \subseteq \mathcal{G}_{L+\varepsilon}$. Finally, for any $g \in \mathcal{X}_\kappa$, combining Property 2 and the termination condition (1) gives that $\pi_g^{\kappa+1}$ is $\varepsilon/9$ -optimal in $\overline{\mathcal{M}}_{g,H}$. As a result, the translation from the finite-horizon to goal-oriented objective (which holds since $g \in \mathcal{G}_{L+\varepsilon}$ and by choice of the horizon $H = \Omega(L \log L \varepsilon^{-1})$, see Lemma 27) yields that $V_g^{\hat{\pi}_g}(s_0) \leq V_g^*(s_0) + \varepsilon$, i.e., $\hat{\pi}_g$ is ε -optimal for the original SSP objective. This concludes the proofs of Theorems 8 and 11.

6 OPERATIONALIZING ADAGoAL IN DEEP RL

In this section, we present a way to operationalize the ADAGoAL idea of targeting goals with high “uncertainty”. We show it can be implemented similar to the deep RL algorithm of Zhang et al. (2020b), and we investigate an aspect that was not considered in the latter paper, which pertains to the capability of ADAGoAL to adapt to an unknown target goal set (\mathcal{G}_L) given a goal set that is possibly misspecified (\mathcal{G}).

First, we notice that \mathcal{Q} and \mathcal{D} in Algorithm 1 can be learned in practice with a goal-conditioned value-based neural network (Schaul et al., 2015). Meanwhile, the prediction errors \mathcal{E} can be approximated by the disagreement between an ensemble of goal-conditioned Q-functions. Interestingly, this approach has already been investigated in the deep GC-RL algorithm VDS (Zhang et al., 2020b), which considers a similar goal-proposal module to prioritize goals that maximize the epistemic uncertainty of the Q-function of the current policy, in order to sample goals at the frontier of the set of goals that the agent is able to reach.

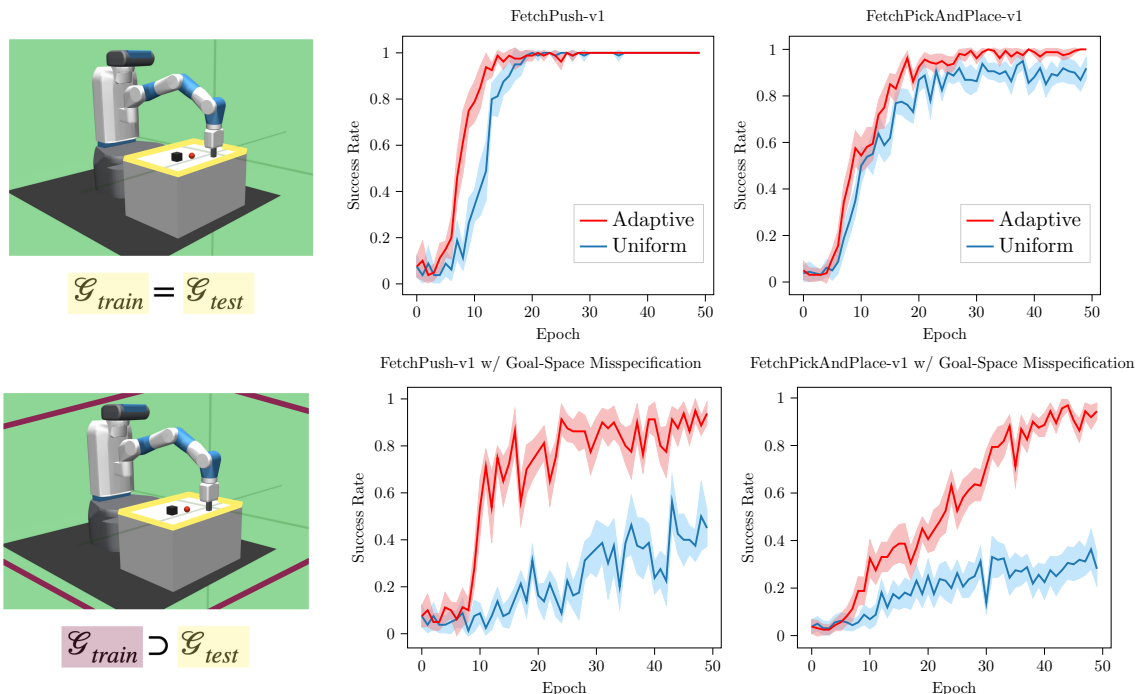


Figure 2: Success rate evaluated on $\mathcal{G}_{\text{test}}$ with the latest policy trained on $\mathcal{G}_{\text{train}}$. The shaded region represents confidence over 5 random seeds. The adaptive goal sampling scheme improves the learning performance over the uniform sampling of HER. This is especially the case in the presence of goal-space misspecification (bottom row), where the training goal space $\mathcal{G}_{\text{train}}$ (delimited in purple) is larger than the test goal space $\mathcal{G}_{\text{test}}$ (delimited in yellow).

Specifically, we take an ensemble $\{Q_j\}_{1 \leq j \leq J}$ of J randomly initialized goal-conditioned Q -functions and we instantiate $\mathcal{E}(g)$ as the standard deviation of the ensemble’s Q_j values conditioned on goal g . Since computing the maximum over $g \in \mathcal{G}$ in (2a) is expensive if the goal space \mathcal{G} is large, the procedure is replaced by first uniformly sampling a set of candidate goals $\{g^{(n)}\}_{n=1}^N \subset \mathcal{G}$, and then selecting a goal $g^{(n)}$ with probability

$$q_n \triangleq \frac{\mathcal{E}(g^{(n)})}{\sum_{n'=1}^N \mathcal{E}(g^{(n')})}, \quad \mathcal{E}(g) \triangleq \text{std}_{1 \leq j \leq J} \left\{ \min_{a \in \mathcal{A}} Q_j(s_0, a, g) \right\}.$$

Moreover, approximating $H \approx L$ renders the constraint in (2b) always valid so it can be omitted. Hence, this approximation of ADAGOAL exactly recovers the goal sampling scheme of Zhang et al. (2020b), which pairs it with Hindsight Experience Replay (HER, Andrychowicz et al., 2017) that performs *uniform* goal sampling.

We consider the multi-goal environments of FetchPush and FetchPickAndPlace, which are sparse-reward simulated robotic manipulation tasks from OpenAI Gym (Plappert et al., 2018). We empirically compare the performance of such an adaptive goal selection that prioritizes “uncertain” goals and the performance of HER’s uniform goal selection (see Appendix F for implementation details). We observe in Figure 2 (top row) that the adaptive goal sampling scheme outperforms the uniform one of HER, which is consistent with the results of Zhang et al. (2020b).

In the above experimental set-up (which is in fact considered by most deep GC-RL works), the goal space $\mathcal{G}_{\text{train}}$ seen at train time is the same as the goal space $\mathcal{G}_{\text{test}}$ on which the agent is evaluated at test time, i.e., the white rectangular table. In the language of the previous sections, by relating $\mathcal{G}_{\text{train}} \leftrightarrow \mathcal{G}$ and $\mathcal{G}_{\text{test}} \leftrightarrow \mathcal{G}_L$, this means that the environment is considered communicating and $\mathcal{G} = \mathcal{G}_L$. However, in some cases, there may be some *misspecification* in the goal space seen during the learning interaction. This may occur if the agent is unaware of the goals of interest, in which case we have that $\mathcal{G}_L \subsetneq \mathcal{G}$, where we recall that \mathcal{G}_L is a priori unknown. We design an experiment to model this scenario by translating the x-y range of $\mathcal{G}_{\text{train}}$ by a factor of $\lambda \geq 1$. Specifically, denoting by (x_0, y_0, z_0) the center of the table and letting $r \triangleq 0.15$, we leave $\mathcal{G}_{\text{test}}$ unchanged yet we expand $\mathcal{G}_{\text{train}} \supset \mathcal{G}_{\text{test}}$ as

$$\begin{aligned} \mathcal{G}_{\text{test}} &\triangleq \{(x_0 + \mathcal{U}(-r, r), y_0 + \mathcal{U}(-r, r), z_0)\}, \\ \mathcal{G}_{\text{train}} &\triangleq \{(x_0 + \mathcal{U}(-\lambda r, \lambda r), y_0 + \mathcal{U}(-\lambda r, \lambda r), z_0)\}, \end{aligned}$$

with $\lambda_{\text{FetchPush}} = 10$, $\lambda_{\text{FetchPickAndPlace}} = 5$. In this scenario, Figure 2 (bottom row) shows that an adaptive goal sampling scheme is particularly pertinent. Intuitively, it enables to discard the set of goals $\mathcal{G}_{\text{train}} \setminus \mathcal{G}_{\text{test}}$ that cannot be reached and thus hinder learning when the agent conditions its behavior on them. This empirically corroborates ADAGOAL’s (theoretically established) ability to adapt to an unknown target goal set (\mathcal{G}_L) given a goal set that is possibly misspecified (\mathcal{G}).

Acknowledgement

We thank Evrard Garcelon, Andrea Tirinzoni and Yann Ollivier for helpful discussion.

References

- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24:2312–2320, 2011.
- Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, Pieter Abbeel, and Wojciech Zaremba. Hindsight experience replay. In *Advances in neural information processing systems*, pages 5048–5058, 2017.
- Alex Ayoub, Zeyu Jia, Csaba Szepesvari, Mengdi Wang, and Lin Yang. Model-based reinforcement learning with value-targeted regression. In *International Conference on Machine Learning*, pages 463–474. PMLR, 2020.
- Mohammad Gheshlaghi Azar, Rémi Munos, and Hilbert J Kappen. Minimax pac bounds on the sample complexity of reinforcement learning with a generative model. *Machine learning*, 91(3):325–349, 2013.
- Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 263–272. JMLR. org, 2017.
- Dimitri Bertsekas. *Dynamic programming and optimal control*, volume 2. 1995.
- Liyu Chen and Haipeng Luo. Finding the stochastic shortest path with low regret: The adversarial cost and unknown transition case. In *International Conference on Machine Learning*, pages 1651–1660. PMLR, 2021.
- Liyu Chen, Mehdi Jafarnia-Jahromi, Rahul Jain, and Haipeng Luo. Implicit finite-horizon approximation and efficient optimal algorithms for stochastic shortest path. *arXiv preprint arXiv:2106.08377*, 2021a.
- Xiaoyu Chen, Jiachen Hu, Lin F Yang, and Liwei Wang. Near-optimal reward-free exploration for linear mixture mdps with plug-in solver. *arXiv preprint arXiv:2110.03244*, 2021b.
- Nuttapong Chentanez, Andrew G Barto, and Sander P Singh. Intrinsically motivated reinforcement learning. In *Advances in neural information processing systems*, pages 1281–1288, 2005.
- Wang Chi Cheung. Regret minimization for reinforcement learning with vectorial feedback and complex objectives. In *Advances in Neural Information Processing Systems*, pages 724–734, 2019.
- Alon Cohen, Yonathan Efroni, Yishay Mansour, and Aviv Rosenberg. Minimax regret for stochastic shortest path. *Advances in Neural Information Processing Systems*, 34, 2021.
- Cédric Colas, Pierre Fournier, Mohamed Chetouani, Olivier Sigaud, and Pierre-Yves Oudeyer. Curious: intrinsically motivated modular multi-goal reinforcement learning. In *International conference on machine learning*, pages 1331–1340. PMLR, 2019.
- Cédric Colas, Tristan Karch, Olivier Sigaud, and Pierre-Yves Oudeyer. Intrinsically motivated goal-conditioned reinforcement learning: a short survey. *arXiv preprint arXiv:2012.09830*, 2020.
- Omar Darwiche Domingues, Yannis Flet-Berliac, Edouard Leurent, Pierre Ménard, Xuedong Shang, and Michal Valko. rlberrry-a reinforcement learning library for research and education, 2021a.
- Omar Darwiche Domingues, Pierre Ménard, Emilie Kaufmann, and Michal Valko. Episodic reinforcement learning in finite mdps: Minimax lower bounds revisited. In *Algorithmic Learning Theory*, pages 578–598. PMLR, 2021b.
- Adrien Ecoffet, Joost Huizinga, Joel Lehman, Kenneth O Stanley, and Jeff Clune. First return, then explore. *Nature*, 590(7847):580–586, 2021.
- Carlos Florensa, David Held, Xinyang Geng, and Pieter Abbeel. Automatic goal generation for reinforcement learning agents. In *International Conference on Machine Learning*, pages 1515–1528, 2018.
- Kristian Hartikainen, Xinyang Geng, Tuomas Haarnoja, and Sergey Levine. Dynamical distance learning for semi-supervised and unsupervised skill discovery. In *International Conference on Learning Representations*, 2020.
- Elad Hazan, Sham Kakade, Karan Singh, and Abby Van Soest. Provably efficient maximum entropy exploration. In *International Conference on Machine Learning*, pages 2681–2691, 2019.
- Mehdi Jafarnia-Jahromi, Liyu Chen, Rahul Jain, and Haipeng Luo. Online learning for stochastic shortest path model via posterior sampling. *arXiv preprint arXiv:2106.05335*, 2021.
- Chi Jin, Akshay Krishnamurthy, Max Simchowitz, and Tiancheng Yu. Reward-free exploration for reinforcement learning. In *International Conference on Machine Learning*, pages 4870–4879. PMLR, 2020.
- Leslie Pack Kaelbling. Learning to achieve goals. In *IJCAI*, pages 1094–1099. Citeseer, 1993.
- Emilie Kaufmann, Pierre Ménard, Omar Darwiche Domingues, Anders Jonsson, Edouard Leurent, and

- Michal Valko. Adaptive reward-free exploration. In *Algorithmic Learning Theory*, pages 865–891. PMLR, 2021.
- Abbas Kazerouni, Mohammad Ghavamzadeh, Yasin Abbasi, and Benjamin Van Roy. Conservative contextual linear bandits. In *Advances in Neural Information Processing Systems*, pages 3910–3919, 2017.
- Tor Lattimore and Marcus Hutter. Pac bounds for discounted mdps. In *International Conference on Algorithmic Learning Theory*, pages 320–334. Springer, 2012.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- Shiau Hong Lim and Peter Auer. Autonomous exploration for navigating in MDPs. In *Conference on Learning Theory*, pages 40–1, 2012.
- Pierre Ménard, Omar Darwiche Domingues, Anders Jonsson, Emilie Kaufmann, Edouard Leurent, and Michal Valko. Fast active learning for pure exploration in reinforcement learning. In *International Conference on Machine Learning*, pages 7599–7608. PMLR, 2021a.
- Pierre Ménard, Omar Darwiche Domingues, Xuedong Shang, and Michal Valko. Ucb momentum q-learning: Correcting the bias without forgetting. In *International Conference on Machine Learning*, pages 7609–7618. PMLR, 2021b.
- Yifei Min, Jiafan He, Tianhao Wang, and Quanquan Gu. Learning stochastic shortest path with linear function approximation. *arXiv preprint arXiv:2110.12727*, 2021.
- Ashvin V Nair, Vitchyr Pong, Murtaza Dalal, Shikhar Bahl, Steven Lin, and Sergey Levine. Visual reinforcement learning with imagined goals. *Advances in Neural Information Processing Systems*, 31:9191–9200, 2018.
- Pierre-Yves Oudeyer and Frederic Kaplan. What is intrinsic motivation? a typology of computational approaches. *Frontiers in neurorobotics*, 1:6, 2009.
- Silviu Pitis, Harris Chan, Stephen Zhao, Bradly Stadie, and Jimmy Ba. Maximum entropy gain exploration for long horizon multi-goal reinforcement learning. In *International Conference on Machine Learning*, pages 7750–7761. PMLR, 2020.
- Matthias Plappert, Marcin Andrychowicz, Alex Ray, Bob McGrew, Bowen Baker, Glenn Powell, Jonas Schneider, Josh Tobin, Maciek Chociej, Peter Welinder, et al. Multi-goal reinforcement learning: Challenging robotics environments and request for research. *arXiv preprint arXiv:1802.09464*, 2018.
- Vitchyr Pong, Murtaza Dalal, Steven Lin, Ashvin Nair, Shikhar Bahl, and Sergey Levine. Skew-fit: State-covering self-supervised reinforcement learning. In *International Conference on Machine Learning*, pages 7783–7792. PMLR, 2020.
- Aviv Rosenberg, Alon Cohen, Yishay Mansour, and Haim Kaplan. Near-optimal regret bounds for stochastic shortest path. In *International Conference on Machine Learning*, pages 8210–8219. PMLR, 2020.
- Tom Schaul, Daniel Horgan, Karol Gregor, and David Silver. Universal value function approximators. In *International conference on machine learning*, pages 1312–1320, 2015.
- Satinder Singh, Richard L Lewis, Andrew G Barto, and Jonathan Sorg. Intrinsically motivated reinforcement learning: An evolutionary perspective. *IEEE Transactions on Autonomous Mental Development*, 2(2):70–82, 2010.
- Jean Tarbouriech and Alessandro Lazaric. Active exploration in markov decision processes. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 974–982, 2019.
- Jean Tarbouriech, Evrard Garcelon, Michal Valko, Matteo Pirotta, and Alessandro Lazaric. No-regret exploration in goal-oriented reinforcement learning. In *International Conference on Machine Learning*, pages 9428–9437. PMLR, 2020a.
- Jean Tarbouriech, Matteo Pirotta, Michal Valko, and Alessandro Lazaric. Improved sample complexity for incremental autonomous exploration in mdps. In *Advances in Neural Information Processing Systems*, volume 33, pages 11273–11284, 2020b.
- Jean Tarbouriech, Shubhanshu Shekhar, Matteo Pirotta, Mohammad Ghavamzadeh, and Alessandro Lazaric. Active model estimation in markov decision processes. In *Conference on Uncertainty in Artificial Intelligence*, pages 1019–1028. PMLR, 2020c.
- Jean Tarbouriech, Matteo Pirotta, Michal Valko, and Alessandro Lazaric. A provably efficient sample collection strategy for reinforcement learning. *Advances in Neural Information Processing Systems*, 34, 2021a.
- Jean Tarbouriech, Matteo Pirotta, Michal Valko, and Alessandro Lazaric. Sample complexity bounds for stochastic shortest path with a generative model. In *Algorithmic Learning Theory*, pages 1157–1178. PMLR, 2021b.

- Jean Tarbouriech, Runlong Zhou, Simon S Du, Matteo Pirotta, Michal Valko, and Alessandro Lazaric. Stochastic shortest path: Minimax, parameter-free and towards horizon-free regret. *Advances in Neural Information Processing Systems*, 34, 2021c.
- Daniel Vial, Advait Parulekar, Sanjay Shakkottai, and R Srikant. Regret bounds for stochastic shortest path problems with linear function approximation. *arXiv preprint arXiv:2105.01593*, 2021.
- Ruosong Wang, Simon S Du, Lin Yang, and Russ R Salakhutdinov. On reward-free reinforcement learning with linear function approximation. In *Advances in Neural Information Processing Systems*, volume 33, pages 17816–17826, 2020.
- David Warde-Farley, Tom Van de Wiele, Tejas Kulkarni, Catalin Ionescu, Steven Hansen, and Volodymyr Mnih. Unsupervised control through non-parametric discriminative rewards. In *International Conference on Learning Representations*, 2019.
- Jingfeng Wu, Vladimir Braverman, and Lin F Yang. Gap-dependent unsupervised exploration for reinforcement learning. *arXiv preprint arXiv:2108.05439*, 2021a.
- Jingfeng Wu, Lin Yang, et al. Accommodating picky customers: Regret bound and exploration complexity for multi-objective reinforcement learning. *Advances in Neural Information Processing Systems*, 34, 2021b.
- Andrea Zanette, Alessandro Lazaric, Mykel J Kochenderfer, and Emma Brunskill. Provably efficient reward-agnostic navigation with linear value iteration. In *Advances in Neural Information Processing Systems*, volume 33, pages 11756–11766, 2020.
- Weitong Zhang, Dongruo Zhou, and Quanquan Gu. Reward-free model-based reinforcement learning with linear function approximation. *Advances in Neural Information Processing Systems*, 34, 2021a.
- Xuezhou Zhang, Yuzhe Ma, and Adish Singla. Task-agnostic exploration in reinforcement learning. *Advances in Neural Information Processing Systems*, 33, 2020a.
- Yunzhi Zhang, Pieter Abbeel, and Lerrel Pinto. Automatic curriculum learning through value disagreement. In *Advances in Neural Information Processing Systems*, pages 7648–7659, 2020b.
- Zihan Zhang, Simon Du, and Xiangyang Ji. Near optimal reward-free reinforcement learning. In *International Conference on Machine Learning*, pages 12402–12412. PMLR, 2021b.
- Rui Zhao, Xudong Sun, and Volker Tresp. Maximum entropy-regularized multi-goal reinforcement learning. In *International Conference on Machine Learning*, pages 7553–7562, 2019.
- Dongruo Zhou, Quanquan Gu, and Csaba Szepesvari. Nearly minimax optimal reinforcement learning for linear mixture markov decision processes. In *Conference on Learning Theory*, pages 4532–4576. PMLR, 2021.

Appendix

Table of Contents

A SHORTCOMINGS OF ALGORITHMIC DESIGNS OF EXISTING THEORETICAL GC-RL METHODS	13
B PROOFS	14
C DETAILS OF ADA_{GOAL}-UCBVI AND ANALYSIS	18
D DETAILS OF ADA_{GOAL}-UCRL-VTR AND ANALYSIS	30
E ABLATION OF THE GOAL SELECTION SCHEME & PROOF OF CONCEPT EXPERIMENT	34
F IMPLEMENTATION DETAILS OF SECTION 6	35

A SHORTCOMINGS OF ALGORITHMIC DESIGNS OF EXISTING THEORETICAL GC-RL METHODS

Here we corroborate our discussion in Section 1 that the existing theoretical GC-RL approaches have quite involved algorithmic designs, far from the interpretable alternation of **(GS)** and **(PE)** steps, thus making them poorly amenable to a more practical implementation. We identify three specific algorithmic shortcomings that we describe below. Specifically, UCBEXPLORE (Lim and Auer, 2012) suffers from Shortcomings 1 and 2, while both DisCo (Tarbouriech et al., 2020b) and GOSPRL (Tarbouriech et al., 2021a) suffer from Shortcoming 3.

- *Shortcoming 1: Cumbersome (PE) step.* Once the method fixes a relevant goal state, it has an elaborate policy execution phase, by running numerous executions of the candidate optimistic goal-conditioned policy and checking after each execution whether an empirical performance check is verified to determine whether to continue the policy execution phase or not.

In contrast, our ADA_{GOAL}-based approach has a simple **(PE)** step, which executes a single rollout of the candidate optimistic goal-conditioned policy for a fixed number of steps before resetting and selecting a (possibly) different, more relevant candidate goal.

- *Shortcoming 2: Possibly eliminatory (GS) step.* The method establishes a strict separation between goal states that have been identified as reliably reachable by a policy and those that have not. In particular, the goal states belonging to the former category can never be selected again as candidate goals, even if the agent could later discover more promising policies leading to it.

In contrast, our ADA_{GOAL}-based approach only implicitly targets goal states of intermediate difficulty (without relying on an explicit distinction between goal states that have already been identified as being reliably reachable and the other goal states). This also means that ADA_{GOAL} does not need to formalize a notion of “fringe” or “border”, which can make it more amenable to implementation.

- *Shortcoming 3: Exhaustive/non-adaptive stopping condition.* The method has a predefined stopping condition, that poorly (or does not) adapts to the samples collected so far. For instance, the method explicitly fixes a minimum number of samples to have collected from *every* state-action pair of interest, or fixes a minimum number of times a candidate goal must be selected.

In contrast, our ADA_{GOAL}-based approach relies on an *adaptive*, data-driven strategy to select goals and to terminate the algorithm. This may also pave the way for problem-dependent sample complexity guarantees.

B PROOFS

Proof of Equation (3). Here we take the limit $H \rightarrow +\infty$. Let $g \in \mathcal{G}_L$, then Property 1 entails that $\mathcal{D}_k(g) \leq V^*(s_0 \rightarrow g) \leq L$, thus $g \in \mathcal{X}_k$, therefore $\mathcal{G}_L \subseteq \mathcal{X}_k$. Now let $g \in \mathcal{X}_k$, then by Property 2 and definition of ε_k , we have that $V^*(s_0 \rightarrow g) \leq \mathcal{D}_k(g) + \mathcal{E}_k(g) \leq L + \mathcal{E}_k(g) \leq L + \varepsilon_k$, therefore $\mathcal{X}_k \subseteq \mathcal{G}_{L+\varepsilon_k}$. \square

Proof of Equation (5). Fix any finite episode index $K < \kappa$, where κ denotes the (possibly unbounded) episode index at which ADA_{GOAL}-UCBVI terminates. By design of ADA_{GOAL}, we have that $\varepsilon \leq \mathcal{E}_k(g_k)$ for every $k \leq K$. We assume that $\kappa > 1$ otherwise the result is trivially true. Define $\kappa' \triangleq \min\{\kappa - 1, K\} \geq 1$. Summing the inequality above yields $\varepsilon \cdot \kappa' \leq \sum_{k=1}^{\kappa'} \mathcal{E}_k(g_k) \leq \kappa' \cdot f_{\mathcal{M}}(\kappa')$. This implies that $\varepsilon \leq f_{\mathcal{M}}(\kappa')$, in which case $f_{\mathcal{M}}^{-1}(\varepsilon) \geq \kappa'$, since $f_{\mathcal{M}}^{-1}$ is strictly decreasing (like $f_{\mathcal{M}}$). Since the last inequality holds for any finite $K < \kappa$, letting $K \rightarrow +\infty$ implies that κ is finite and bounded by $f_{\mathcal{M}}^{-1}(\varepsilon) + 2$. \square

B.1 Proof of Lemma 5

We assume throughout that $L \geq 2$ and $\varepsilon \in (0, 1]$. On the one hand, Theorem 8 proves that ADA_{GOAL}-UCBVI is $(\varepsilon, \delta, L, \mathcal{G})$ -PAC for MGE with sample complexity of order $\tilde{O}(L^3 S A \varepsilon^{-2})$, therefore MGE is solvable in $\text{poly}(S, L, \varepsilon^{-1}, A)$ steps (up to poly-log factors). On the other hand, reset-free MGE is a special case of the cost-free goal-free exploration problem in communicating MDPs studied by Tarbouriech et al. (2021a), whose algorithm GOSPRL can solve it in $\text{poly}(S, D, \varepsilon^{-1}, A)$ steps (up to poly-log factors). Now we prove that there exists an MDP such that any algorithm requires at least $\Omega(D)$ steps to solve the reset-free MGE problem (i.e., without Assumption 3), where the MDP's diameter D can be made exponentially larger than $L, S, A, \varepsilon^{-1}$.

To this end, we design in Figure 3 a communicating MDP \mathcal{M}_{a^\dagger} that does not satisfy Assumption 3, with $A \geq 4$ actions (including a special action $a^\dagger \in \mathcal{A}$) and $S = 5$ states, where $\mathcal{S} \triangleq \{s_0, g, x, \tilde{s}, \bar{s}\}$, and all states apart from \bar{s} are reliably L -reachable from s_0 . We define as goal space $\mathcal{G} \triangleq \{g\}$. We consider the problem of learning an ε -optimal goal-reaching policy with goal state g from the starting state s_0 , i.e., finding a policy π such that $V^\pi(s_0 \rightarrow g) \leq V^*(s_0 \rightarrow g) + \varepsilon$, which is a sub-problem of the MGE objective.

Given $\eta \in (0, 1)$ and an unknown action $a^\dagger \in \mathcal{A}$, we define the following transition probabilities for all $a \in \mathcal{A}$,

$$\begin{aligned} p(\tilde{s}|s_0, a) &\triangleq \eta, & p(x|s_0, a) &\triangleq 1 - \eta, \\ p(g|x, a) &\triangleq \zeta, & p(x|x, a) &\triangleq 1 - \zeta, \\ p(g|\tilde{s}, a) &\triangleq \mathbf{1}[a = a^\dagger], & p(\bar{s}|\tilde{s}, a) &\triangleq \mathbf{1}[a \neq a^\dagger], \\ p(g|\bar{s}, a) &\triangleq \frac{\eta}{2}, & p(\bar{s}|\bar{s}, a) &\triangleq 1 - \frac{\eta}{2}, \\ p(s_0|g, a) &\triangleq 1, \end{aligned}$$

where we can set any $\zeta = O(1/L)$ such that g is reliably L -reachable from s_0 (i.e., $V^*(s_0 \rightarrow g)$). It is easy to see that the MDP's diameter verifies $D = \alpha\eta^{-1}$ for a constant $\alpha > 0$. Finally, we denote by \mathcal{F} the family of MDPs of the form of Figure 3 parameterized by $a^\dagger \in \{1, \dots, A\}$, i.e., $\mathcal{F} \triangleq \{\mathcal{M}_{a^\dagger}\}_{a^\dagger \in \{1, \dots, A\}}$.

We define the event \mathcal{J}_t that state \tilde{s} has never been visited by the agent by time t (recalling that it initially starts at state s_0), i.e., $\mathcal{J}_t \triangleq \{n^t(\tilde{s}) = 0\}$ (note that its probability is the same for all MDPs in \mathcal{F}). We now fix an MDP \mathcal{M}_{a^\dagger} and denote by π^* its optimal policy, i.e., $\pi^* \in \arg \min_\pi V^\pi(\cdot \rightarrow g)$, which in particular selects action a^\dagger at state \tilde{s} . First, we consider any deterministic algorithm \mathfrak{A} whose candidate policy $\hat{\pi}$ does not select action a^\dagger when it is in state \tilde{s} . Then it holds that

$$\begin{aligned} V^{\hat{\pi}}(\tilde{s} \rightarrow g) &= 1 + \sum_{y \in \mathcal{S}} p(y|s_0, \hat{\pi}(\tilde{s})) V^{\hat{\pi}}(y \rightarrow g) \geq 1 + V^{\hat{\pi}}(\bar{s} \rightarrow g) = 1 + \frac{2}{\eta}, \\ V^{\hat{\pi}}(\tilde{s} \rightarrow g) - V^{\pi^*}(\tilde{s} \rightarrow g) &\geq \frac{2}{\eta}, \end{aligned} \tag{6}$$




Figure 3: Illustration of the MDP instance \mathcal{M}_{a^\dagger} .

$$\begin{aligned}
 V^\pi(s_0 \rightarrow g) &= 1 + \eta V^\pi(\tilde{s} \rightarrow g) + (1 - \eta)V^\pi(x \rightarrow g), \quad \forall \pi, \\
 V^{\hat{\pi}}(s_0 \rightarrow g) - V^{\pi^*}(s_0 \rightarrow g) &= (1 - \eta) \underbrace{(V^{\hat{\pi}}(x \rightarrow g) - V^{\pi^*}(x \rightarrow g))}_{\geq 0} + \eta \underbrace{(V^{\hat{\pi}}(\tilde{s} \rightarrow g) - V^{\pi^*}(\tilde{s} \rightarrow g))}_{\geq 2/\eta} \geq 2 > \varepsilon,
 \end{aligned} \tag{7}$$

thus $\hat{\pi}$ has a sub-optimality gap larger than ε . This means that under the event \mathcal{J}_t , where the algorithm cannot know which is the favorable action a^\dagger , it holds that with probability at least $1 - \frac{1}{A} \geq \frac{3}{4}$, any deterministic algorithm's candidate policy $\hat{\pi}$ does not select the action a^\dagger at state \tilde{s} thus its value function is not ε -optimal. Second, we note that we can easily extent to the case where \mathfrak{A} outputs stochastic actions at state \tilde{s} . Given that a^\dagger is unknown, in the best case scenario it can set $\hat{\pi}(\cdot|\tilde{s}) = 1/A$. Then we can retrace the reasoning above and replace (6) with $V^{\hat{\pi}}(\tilde{s}) \geq 1 + \frac{A-1}{A} \frac{2}{\eta}$, and thus (7) with $V^{\hat{\pi}}(s_0) - V^{\pi^*}(s_0) \geq 2 \frac{A-1}{A} > 1 \geq \varepsilon$ since $A > 2$, which leads to the same result that $\hat{\pi}$ is not ε -optimal on at least one of the MDPs in \mathcal{F} .

Now we study how the probability of the event \mathcal{J}_t evolves over time t , i.e. we bound the time required to visit \tilde{s} at least once, which we denote by \tilde{T} . Recall that \tilde{s} can only be reached with probability η from s_0 following any action. The random variable \tilde{T} can be seen as an upper bound of a random variable distributed geometrically with success probability η , thus Chernoff's inequality entails that with probability at least $\frac{1}{2}$ we have $\tilde{T} \geq \frac{1}{9\eta}$, i.e., the event \mathcal{J}_t for $t = \lfloor 1/(9\eta) \rfloor$ holds with probability at least $\frac{1}{2}$.

Putting everything together, there exists an MDP in \mathcal{F} such that with probability at least $\frac{1}{4}$, the number of time steps required to output a candidate policy that is ε -optimal for goal state g is at least $\frac{1}{9\eta} = \Omega(D)$, where η can be made arbitrarily small, so in particular D can be exponentially larger than L . Hence in this MDP instance, no algorithm can solve MGE in $\text{poly}(L)$ steps with overwhelming probability. While here we considered $S = 5$ for simplicity, note that the result easily holds for any $S \geq 5$ by replacing the state x in the construction above with a set of $S - 4$ states; the only property that must be still verified is that g remains reliably L -reachable from s_0 .

Therefore, for any $L \geq 2, S \geq 5, A \geq 4, \varepsilon \in (0, 1]$, there exists an MDP instance and a goal space for which any algorithm requires at least $\Omega(D)$ time steps to solve reset-free MGE, where the diameter D can be exponentially larger than $L, S, A, \varepsilon^{-1}$, which concludes the proof of Lemma 5.

B.2 Proof of Lemma 6

In the following, we will prove in Lemma 14 a lower bound on the expected number of exploration steps to find an ε -optimal SSP policy from s_0 for a specific goal state g that is reliably L -reachable from s_0 . Such BPI-SSP objective corresponds to our MGE objective with goal space $\mathcal{G} = \{g\}$ and it thus induces a lower bound on the MGE problem, which will conclude the proof of Lemma 6.

Notation. We largely follow the notations and definitions of Domingues et al. (2021b). We define an RL algorithm as a history-dependent policy π used to interact with the environment. In the BPI setting, where we eventually stop and recommend a policy, an algorithm is defined as a triple $(\pi, \tau, \hat{\pi}_\tau)$ where τ is a stopping time and $\hat{\pi}_\tau$ is a Markov policy recommended after τ time steps. We now write more formally our objective.

Definition 13 (BPI-SSP). *An algorithm $(\pi, \tau, \hat{\pi}_\tau)$ is (ε, δ, L) -PAC for best-policy identification in a stochastic shortest path problem satisfying Assumption 3 (BPI-SSP) if $V^*(s_0) \leq L$ and if the policy $\hat{\pi}_\tau$ returned after τ time steps satisfies, for the initial state s_0 ,*

$$\mathbb{P}_{\pi, \mathcal{M}} \left[V^{\hat{\pi}_\tau}(s_0) - V^*(s_0) \leq \varepsilon \right] \geq 1 - \delta,$$

where we denote the goal state by g and the SSP value of any policy π at any state s by $V^\pi(s) \triangleq \mathbb{E}_{\pi, \mathcal{M}}[\inf\{i \geq 0 : s_{i+1} = g\} | s_1 = s]$, and $V^*(s) \triangleq \min_\pi V^\pi(s)$.

Lemma 14 (BPI-SSP Lower Bound). *There exist absolute constants $L_0, S_0, A_0, \varepsilon_0, \delta_0$ such that, for any $L \geq L_0, A \geq A_0, \varepsilon \leq \varepsilon_0, \delta \leq \delta_0$ and $S_0 \leq S \leq A^{\frac{5}{3}-2}$, and for any algorithm $(\pi, \tau, \hat{\pi}_\tau)$ that is (ε, δ, L) -PAC for BPI-SSP in any finite MDP with S states and A actions, there exists an MDP \mathcal{M} with a goal state belonging to \mathcal{G}_L and an absolute constant β such that*

$$\mathbb{E}_{\pi, \mathcal{M}}[\tau] \geq \beta \frac{L^3 S A}{\varepsilon^2} \log\left(\frac{1}{\delta}\right).$$

Proof of Lemma 14.

We first define our family of hard MDPs for $S = 4$ states, and the extension to any S states can be done as in Domingues et al. (2021b) as explained later. Consider the hard MDP illustrated in Figure 4, where all states incur a cost of 1 apart from the goal state s_g (a.k.a. “good” state). The agent stays in the initial state s_0 with probability $1 - q$, and goes to a *decision state* s_d with probability q . For any action a taken in s_d , the agent reaches s_g with probability $1/2$ and a “bad” state s_b with probability $1/2$, except if an action a^* is chosen, that increases to $1/2 + \tilde{\varepsilon}$ the probability of reaching s_g . From s_b , the agent returns to the initial state s_0 with probability 1. The goal state s_g is absorbing, and the agent stays there unless the reset action is taken, which brings the agent back to s_0 . Note that the MDP satisfies Assumption 3 (the arrows of the reset action from s_d to s_0 and from s_g to s_0 are not represented in Figure 4 for visual convenience). Moreover, we define the following parameters

$$H \triangleq \lceil \frac{L}{2} - 1 \rceil, \quad q \triangleq 1/H, \quad \tilde{\varepsilon} \triangleq \frac{\varepsilon}{2(H+1)}.$$




Figure 4: Illustration of our hard MDP.

Note that this hard MDP instance is inspired from hard MDPs used in prior lower-bound constructions (see e.g., Lattimore and Hutter, 2012; Domingues et al., 2021b), albeit with slight modifications. Indeed, a key difference with respect to the discounted MDP setting (Lattimore and Hutter, 2012) or the finite-horizon MDP setting (Domingues et al., 2021b) is that in our case, the agent has access to an anytime reset action (Assumption 3). This implies that we cannot do as prior works that rely on absorption properties of states in the MDP (e.g., the “good” and “bad” states s_b and s_g) in order to compound errors and add an extra effective horizon term (either $1/(1 - \gamma)$ or H) in the sample complexity (i.e., to go from quadratic to cubic). The only absorption property we can rely on here is at the initial state s_0 . It turns out that this will be sufficient in our setting to compound error and go from L^2 to L^3 dependence. The intuition for this is that the SSP value function generates a cost of +1 at each time step until the goal state is reached, which compounds errors more than the usual reward-based value function in the sparse-reward MDP constructions of Lattimore and Hutter (2012); Domingues et al. (2021b).

We consider a family of MDPs of the form of Figure 4, parameterized by $a^* \in \{1, \dots, A\}$, where we denote by \mathcal{M}_{a^*} the MDP such that a^* increases the probability by $\tilde{\varepsilon}$ of reaching the goal state s_g from state s_d . For any policy π we denote its SSP value (with goal state s_g) at state s in \mathcal{M}_{a^*} by $V_{a^*}^\pi(s)$.

We also define a reference MDP \mathcal{M}_0 , where $\tilde{\varepsilon} = 0$, that is, there is no special action increasing the probability of reaching the goal state s_g . We denote by $\mathbb{P}_{a^*}[\cdot]$ and $\mathbb{E}_{a^*}[\cdot]$ the probability measure and the expectation in the MDP \mathcal{M}_{a^*} by following the algorithm π and by $\mathbb{P}_0[\cdot]$ and $\mathbb{E}_0[\cdot]$ the corresponding operators in \mathcal{M}_0 .

The optimal value function does not depend on the MDP parameter a^* , and for any MDP \mathcal{M}_{a^*} we get

$$\begin{aligned} V^*(s_0) &= \frac{1}{q} + \left(\frac{1}{2} + \tilde{\varepsilon}\right) + \left(1 - \frac{1}{2} - \tilde{\varepsilon}\right)(1 + V^*(s_0)) \\ \implies V^*(s_0) &= \left(\frac{1}{q} + 1\right) \frac{1}{1/2 + \tilde{\varepsilon}}. \end{aligned}$$

Note that by our choice of q , it importantly holds that $V^*(s_0) \leq L$, i.e., $s_g \in \mathcal{G}_L$.

Meanwhile, the value function of the recommended policy $\hat{\pi}_\tau$ in \mathcal{M}_{a^*} is

$$V_{a^*}^{\hat{\pi}_\tau}(s_0) = \left(\frac{1}{q} + 1\right) \frac{1}{1/2 + \tilde{\varepsilon} \cdot \hat{\pi}_\tau(a^*|s_d)}.$$

As a result,

$$V_{a^*}^{\widehat{\pi}_\tau}(s_0) - V^*(s_0) = \left(\frac{1}{q} + 1\right) \frac{\widetilde{\varepsilon}(1 - \widehat{\pi}_\tau(a^*|s_d))}{(1/2 + \widetilde{\varepsilon})(1/2 + \widetilde{\varepsilon} \cdot \widehat{\pi}_\tau(a^*|s_d))} \leq \underbrace{4\left(\frac{1}{q} + 1\right)\widetilde{\varepsilon}}_{=2\varepsilon} \cdot (1 - \widehat{\pi}_\tau(a^*|s_d)),$$

and thus

$$V_{a^*}^{\widehat{\pi}_\tau}(s_0) - V^*(s_0) < \varepsilon \iff \widehat{\pi}_\tau(a^*|s_d) > \frac{1}{2}.$$

We observe that this analysis of the suboptimality gap of $\widehat{\pi}_\tau$ in terms of SSP value functions can be mapped to the one of Domingues et al. (2021b, Proof of Theorem 7) for their finite-horizon value functions, despite the different MDP constructions. This means that we can retrace the steps of Domingues et al. (2021b, Proof of Theorem 7). In the following, we use the notation $\stackrel{(D)}{\geq}$ to signify that the inequality stems from following the same corresponding steps of Domingues et al. (2021b, Proof of Theorem 7). In particular, we similarly define the event

$$\mathcal{E}_{a^*}^\tau \triangleq \left\{ \widehat{\pi}_\tau(a^*|s_d) > \frac{1}{2} \right\},$$

and since the algorithm is assumed to be (ε, δ, L) -PAC for BPI-SSP for any MDP, we have

$$\mathbb{P}_{a^*} \left[\mathcal{E}_{a^*}^\tau \right] = \mathbb{P}_{a^*} \left[V_{a^*}^{\widehat{\pi}_\tau}(s_0) < V^*(s_0) + \varepsilon \right] \geq 1 - \delta.$$

We now proceed with lower bounding the expectation of the sample complexity τ in the reference MDP \mathcal{M}_0 . We define

$$N_{a^*}^\tau \triangleq \sum_{t=1}^{\tau} \mathbb{1}[S_t = s_d, A_t = a^*], \quad N^\tau \triangleq \sum_{a^*} N_{a^*}^\tau.$$

Following the steps of Domingues et al. (2021b, Proof of Theorem 7), we have that

$$\mathbb{E}_0 \left[N_{a^*}^\tau \right] 4\varepsilon^2 \stackrel{(D)}{\geq} \mathbb{E}_0 \left[N_{a^*}^\tau \right] \text{kl} \left(\frac{1}{2}, \frac{1}{2} + \widetilde{\varepsilon} \right) \stackrel{(D)}{\geq} \left[\left(1 - \mathbb{P}_0 \left[\mathcal{E}_{a^*}^\tau \right] \right) \log \left(\frac{1}{\delta} \right) - \log(2) \right],$$

thus

$$\mathbb{E}_0 \left[N_{a^*}^\tau \right] \geq \frac{1}{4\varepsilon^2} \left[\left(1 - \mathbb{P}_0 \left[\mathcal{E}_{a^*}^\tau \right] \right) \log \left(\frac{1}{\delta} \right) - \log(2) \right].$$

Summing all over MDP instances, we obtain following Domingues et al. (2021b, Proof of Theorem 7) that

$$\mathbb{E}_0 \left[N^\tau \right] = \sum_{a^*} \mathbb{E}_0 \left[N_{a^*}^\tau \right] \stackrel{(D)}{\geq} \frac{A}{8\varepsilon^2} \log \left(\frac{1}{\delta} \right).$$

While the proof of Domingues et al. (2021b, Theorem 7) can stop at this stage, our proof requires an additional step of linking this back to the sample complexity τ , since the latter is not defined in terms of number of episodes but in terms of number of time steps.

For any $i \leq \tau$, denote by $W_i(s_0 \rightarrow s_d)$ the random variable of the number of time steps required to reach s_d starting from s_0 for the i -th time — note importantly that this quantity is independent of the algorithm and of the MDP parameter a^* , and we can write $\mathbb{E}[W_i(s_0 \rightarrow s_d)] = \mathbb{E}[W(s_0 \rightarrow s_d)] = \frac{1}{q}$. Then from Wald's equation we have

$$\mathbb{E}_0[\tau] \geq \mathbb{E}_0 \left[\sum_{i=1}^{N^\tau} W_i(s_0 \rightarrow s_d) \right] = \mathbb{E}[W(s_0 \rightarrow s_d)] \cdot \mathbb{E}_0[N^\tau] \geq \frac{1}{q} \cdot \alpha \frac{1}{\varepsilon^2} A \log \left(\frac{1}{\delta} \right) \geq \alpha' \frac{L^3 A}{\varepsilon^2} \log \left(\frac{1}{\delta} \right),$$

where α and α' are absolute constants.

Finally, as mentioned, the extension to any S to make appear the multiplicative dependence on S can be done by following the steps done in Domingues et al. (2021b, Proof of Theorem 7) (which relies on their Assumption 1, see their Theorem 10 for the relaxed statement). The idea of the construction is to consider not just 1 decision state s_d but $S - 3$ of them, where only one of them possesses the favorable action a^* ; intuitively this generates SA actions instead of A (see e.g., Lattimore and Szepesvári, 2020, Section 38.7), thus leading to the additional S factor in the sample complexity. \square

C DETAILS OF ADA_{GOAL}-UCBVI AND ANALYSIS

In this section, we focus on ADA_{GOAL}-UCBVI. In Section C.1, we introduce useful notation. In Section C.2, we define the exact choice of estimates \mathcal{D} , \mathcal{E} , \mathcal{Q} used by ADA_{GOAL}-UCBVI (line 13 of Algorithm 1). Then in Section C.3, we provide the full proof of Theorem 8, by establishing the key steps followed in Section 4. Throughout the analysis we consider that $\mathcal{G} = \mathcal{S}$, $\varepsilon \in (0, 1]$ and $\delta \in (0, 1)$.

C.1 Notation

Given a goal state $g \in \mathcal{G}$, denote by \mathcal{M}_g the unit-cost SSP-MDP which adds a self-loop at g to the original MDP \mathcal{M} , and denote by p_g its transition function and c_g its cost function. Formally, let

$$c_g(s, a) \triangleq \mathbb{1}[s \neq g], \quad p_g(s'|s, a) \triangleq \begin{cases} p(s'|s, a) & \text{if } s \neq g \\ \mathbb{1}[s' = g] & \text{if } s = g. \end{cases}$$

For any (possibly non-stationary) policy $\pi = (\pi_h)_{h \geq 1}$, let V_g^π be its SSP value function (i.e., expected cost-to-go) in \mathcal{M}_g , i.e.,

$$V_g^\pi(s_0) \triangleq \mathbb{E} \left[\sum_{h=1}^{+\infty} c_g(s_h, a_h) \mid s_1 = s_0, \pi, \mathcal{M}_g \right],$$

where $a_h \triangleq \pi_h(s_h)$ and $s_{h+1} \sim P_g(s_h, a_h)$. Let $\pi_g^* \in \arg \min_{\pi} V_g^\pi$ and $V_g^* \triangleq V_g^{\pi_g^*}$. We now define the set of finite-horizon goal-conditioned models.

Definition 15 (Finite-Horizon Goal-Conditioned Models). *Fix a horizon $H \geq 1$. For any goal state $g \in \mathcal{G}$, denote by $\overline{\mathcal{M}}_{g,H}$ the finite-horizon model that corresponds to starting from state s_0 and interacting for H steps with the original MDP \mathcal{M} in which state g is made absorbing. $\overline{\mathcal{M}}_{g,H}$ admits as cost function $\overline{c}_g \triangleq c_g$ and as transition function $\overline{p}_g \triangleq p_g$.*

Remark 5. Note that for a given goal state $g \in \mathcal{G}$, the cost function \overline{c}_g is known to the agent, while the MDP transitions \overline{p}_g are unknown with some (known) goal-specific changes w.r.t. the original MDP \mathcal{M} (namely, the self-loop at g). An alternative way of framing the problem is that there is one single MDP with state space $\mathcal{S} \times \mathcal{G}$, i.e., with state variable (s, g) .

For a finite-horizon policy $\pi \in \Pi_H$, denote by $\overline{V}_{g,h}^\pi$ its finite-horizon value function at step $1 \leq h \leq H$ in the finite-horizon instance $\overline{\mathcal{M}}_{g,H}$, i.e.,

$$\overline{V}_{g,h}^\pi(s_0) \triangleq \mathbb{E} \left[\sum_{h'=h}^H \overline{c}_g(s_{h'}, a_{h'}) \mid s_1 = s_0, \pi, \overline{\mathcal{M}}_{g,H} \right].$$

We define the corresponding optimal value function as $\overline{V}_{g,h}^* \triangleq \min_{\pi} \overline{V}_{g,h}^\pi$. Observe that $\overline{V}_{g,1}^*(s_0) = \mathcal{D}_H^*(g)$ (notation used in Properties 1 and 2 of Section 3).

Let (s_i, a_i, s_{i+1}) be the state, action and next state observed by an algorithm at time step i . Let $n^k(s, a) \triangleq \sum_{i=1}^{kH} \mathbb{1}[(s_i, a_i) = (s, a)]$ be the number of times state-action pair (s, a) was visited in the first k episodes and $n^k(s, a, s') \triangleq \sum_{i=1}^{kH} \mathbb{1}[(s_i, a_i, s_{i+1}) = (s, a, s')]$. We define the empirical transitions as $\widehat{p}^k(s'|s, a) \triangleq n^k(s, a, s')/n^k(s, a)$ if $n^k(s, a) > 0$, and $\widehat{p}^k(s'|s, a) \triangleq 1/S$ otherwise. Also, $pX(s, a) \triangleq \mathbb{E}_{s' \sim p(\cdot|s,a)}[X(s')]$ denotes the expectation operator w.r.t. the transition probabilities p and $\pi_h Y(s) \triangleq Y(s, \pi_h(s))$ denotes the composition with policy π at step h , so that $p\pi_h Z(s, a) \triangleq \mathbb{E}_{s' \sim p(\cdot|s,a)}[Z(s', \pi_h(s'))]$. Finally, we denote the clip function by $\text{clip}(x, y, z) \triangleq \max(\min(x, z), y)$.

Finally, in the analysis we denote by $p_{g,h}^k(s, a)$ the probability of reaching state-action pair (s, a) at step h under policy π_g^k in the true MDP. We also define the pseudo-counts as $\overline{n}^k(s, a) \triangleq \sum_{h=1}^H \sum_{l=1}^k p_{g_l,h}^l(s, a)$, where $g_l \in \mathcal{S}$ denotes the goal state selected by ADA_{GOAL}-UCBVI at the beginning of algorithmic episode l .

C.2 Algorithmic Choices of \mathcal{D} , \mathcal{E} , \mathcal{Q} for ADAGoAL-UCBVI

We generalize to our goal-conditioned scenario the estimates used by BPI-UCBVI (Ménard et al., 2021a), a recent algorithm designed for Best-Policy Identification (BPI) in finite-horizon non-stationary MDPs. First, we build the optimistic goal-conditioned Q-values and value functions in the finite-horizon models $\overline{\mathcal{M}}_{g,H}$ for $g \in \mathcal{S}$ and $h \leq H$ as follows, $\widetilde{Q}_{g,h}^0(s,a) \triangleq \mathbb{1}[s \neq g]$,

$$\begin{aligned} \widetilde{Q}_{g,h}^k(s,a) &\triangleq \text{clip} \left(\mathbb{1}[s \neq g] - 3 \sqrt{\text{Var}_{\widehat{p}^k}(\widetilde{V}_{g,h+1}^k)(s,a)} \frac{\beta^*(n^k(s,a), \delta)}{n^k(s,a)} - 14H^2 \frac{\beta(n^k(s,a), \delta)}{n^k(s,a)} \right. \\ &\quad \left. - \frac{1}{H} \widehat{p}^k(\overline{V}_{g,h+1}^k - \widetilde{V}_{g,h+1}^k)(s,a) + \widehat{p}^k \widetilde{V}_{g,h+1}^k(s,a), 0, H \right), \\ \widetilde{V}_{g,h}^k(s) &\triangleq \min_{a \in \mathcal{A}} \widetilde{Q}_{g,h}^k(s,a), \quad \widetilde{V}_{g,h}^k(g) \triangleq 0, \quad \widetilde{V}_{g,H+1}^k(s) \triangleq 0, \end{aligned}$$

where we define the variance of $\widetilde{V}_{g,h+1}^k$ with respect to $\widehat{p}^k(\cdot|s,a)$ as $\text{Var}_{\widehat{p}^k}(\widetilde{V}_{g,h+1}^k)(s,a) \triangleq \sum_{s'} \widehat{p}^k(s'|s,a) (\widetilde{V}_{g,h+1}^k(s') - \widehat{p}^k \widetilde{V}_{g,h+1}^k(s,a))^2$, where $\beta(n, \delta) = \widetilde{O}(S \log(n/\delta))$ and $\beta^*(n, \delta) = \widetilde{O}(\log(n/\delta))$ are some exploration thresholds and \overline{V}_g^k is a pessimistic finite-horizon goal-conditioned value function; see Appendix C.3 for the complete definitions. Let $\pi_{g,h}^{k+1}$ be the greedy policy with respect to the lower bounds $\widetilde{Q}_{g,h}^k$. We recursively define the functions \overline{U}_g^k for $g \in \mathcal{S}$ and $h \leq H$ as follows, $\overline{U}_{g,h}^0(s,a) \triangleq H \mathbb{1}[s \neq g]$,

$$\begin{aligned} \overline{U}_{g,h}^k(s,a) &\triangleq \text{clip} \left(6 \sqrt{\text{Var}_{\widehat{p}^k}(\widetilde{V}_{g,h+1}^k)(s,a)} \frac{\beta^*(n^k(s,a), \delta)}{n^k(s,a)} + 36H^2 \frac{\beta(n^k(s,a), \delta)}{n^k(s,a)} \right. \\ &\quad \left. + \left(1 + \frac{3}{H} \right) \widehat{p}^k \pi_{g,h+1}^{k+1} \overline{U}_{g,h+1}^k(s,a), 0, H \right), \\ \overline{U}_{g,h}^k(g,a) &\triangleq 0, \quad \overline{U}_{g,H+1}^k(s,a) \triangleq 0, \end{aligned}$$

We are now ready to define the distance and error estimates of ADAGoAL-UCBVI (Algorithm 1) as follows

$$\mathcal{Q}_{k,h}(s,a,g) \triangleq \widetilde{Q}_{g,h}^{k-1}(s,a), \quad (8)$$

$$\mathcal{D}_k(g) \triangleq \widetilde{V}_{g,1}^{k-1}(s_0), \quad (9)$$

$$\mathcal{E}_k(g) \triangleq \pi_{g,1}^k \overline{U}_{g,1}^{k-1}(s_0) + \frac{8\varepsilon}{9}. \quad (10)$$

C.3 Proof of Theorem 8

In this section, we provide the full proof of Theorem 8, by establishing the key steps followed in Section 4. Appendices C.3.1, C.3.2 and C.3.3, which prove “key steps ①, ②, ③”, focus on more “standard” technical tools (e.g., high-probability events, variance-aware concentration inequalities); in particular building on the analysis of Ménard et al. (2021a) on the sample complexity of BPI in finite-horizon MDPs and extending it to our goal-conditioned scenario. Then Appendix C.3.4 proves “key step ④”, which focuses on the technical novelty of the ADAGoAL goal selection scheme that is specific to the multi-goal exploration setting.

We begin by stating a simple property that we will rely on throughout the analysis.

Lemma 16. *For any state-action pair $(s,a) \in \mathcal{S} \times \mathcal{A}$ and goal state $g \in \mathcal{S}$, consider any vector $Y \in \mathbb{R}^{\mathcal{S}}$ such that $Y(g) = 0$, then $pY(s,a) = p_g Y(s,a)$, where we recall that*

$$p_g(s'|s,a) \triangleq \begin{cases} p(s'|s,a) & \text{if } s \neq g \\ \mathbb{1}[s'=g] & \text{if } s = g. \end{cases}$$

Proof. It is easy to see that

$$pY(s,a) = \sum_{s' \in \mathcal{S} \setminus \{g\}} p(s'|s,a) Y(s') + p(g|s,a) \underbrace{Y(g)}_{=0}$$

$$\begin{aligned}
 &= \sum_{s' \in \mathcal{S} \setminus \{g\}} p_g(s'|s, a) Y(s') + p_g(g|s, a) \underbrace{Y(g)}_{=0} \\
 &= p_g Y(s, a).
 \end{aligned}$$

□

Thanks to the above observation, our analysis will not require to handle goal-conditioned (true or empirical) transition probabilities, and will only need to deal with the (true or empirical) transition probabilities of the original MDP \mathcal{M} .

C.3.1 Proof of “key step ①”

Concentration events. Here we define the high-probability event \mathcal{U} on which we condition our statements. We follow the notation of [Ménard et al. \(2021a, Appendix A\)](#) and define the three following favorable events: \mathcal{E} the event where the empirical transition probabilities are close to the true ones, \mathcal{E}^{cnt} the event where the pseudo-counts are close to their expectation, and \mathcal{E}^* where the empirical means of the optimal goal-conditioned value functions are close to the true ones. Denoting by KL the Kullback-Leibler divergence, we set

$$\begin{aligned}
 \mathcal{E} &\triangleq \left\{ \forall k \in \mathbb{N}, \forall (s, a) \in \mathcal{S} \times \mathcal{A} : \text{KL}(\hat{p}^k(\cdot|s, a), p(\cdot|s, a)) \leq \frac{\beta(n^k(s, a), \delta)}{n^k(s, a)} \right\}, \\
 \mathcal{E}^{\text{cnt}} &\triangleq \left\{ \forall k \in \mathbb{N}, \forall (s, a) \in \mathcal{S} \times \mathcal{A} : n^k(s, a) \geq \frac{1}{2} \bar{n}^k(s, a) - \beta^{\text{cnt}}(\delta) \right\}, \\
 \mathcal{E}^* &\triangleq \left\{ \forall k \in \mathbb{N}, \forall h \in [H], \forall (s, a) \in \mathcal{S} \times \mathcal{A}, \forall g \in \mathcal{S} : \right. \\
 &\quad \left. |(\hat{p}^k - p) \bar{V}_{g, h+1}^*(s, a)| \leq \min \left(H, \sqrt{2 \text{Var}_p(\bar{V}_{g, h+1}^*)(s, a) \frac{\beta^*(n^k(s, a), \delta)}{n^k(s, a)}} + 3H \frac{\beta^*(n^k(s, a), \delta)}{n^k(s, a)} \right) \right\}.
 \end{aligned}$$

We define the intersection of these events as

$$\mathcal{U} \triangleq \mathcal{E} \cap \mathcal{E}^{\text{cnt}} \cap \mathcal{E}^*. \quad (11)$$

We prove that for the right choice of the functions β the above event holds with high probability.

Lemma 17. *For the following choices of functions β ,*

$$\begin{aligned}
 \beta(n, \delta) &\triangleq \log(3S^2 AH/\delta) + S \log(8e(n+1)), \\
 \beta^{\text{cnt}}(\delta) &\triangleq \log(3S^2 AH/\delta), \\
 \beta^*(n, \delta) &\triangleq \log(3S^2 AH/\delta) + \log(8e(n+1)),
 \end{aligned}$$

it holds that $\mathbb{P}(\mathcal{U}) \geq 1 - \delta$.

Proof. The only difference with respect to the concentration inequalities of [Ménard et al. \(2021a, Appendix A\)](#) is that we need to take a union bound over the goal states $g \in \mathcal{S}$ when concentrating our optimal goal-conditioned value functions. We thus set $\delta \leftarrow \delta/S$ in the choices of functions β compared to [Ménard et al. \(2021a\)](#). As a result, by [Ménard et al. \(2021a, Theorem 3 & 4 & 5\)](#) we have that $\mathbb{P}(\mathcal{E}) \geq 1 - \frac{\delta}{3}$, $\mathbb{P}(\mathcal{E}^{\text{cnt}}) \geq 1 - \frac{\delta}{3}$ and $\mathbb{P}(\mathcal{E}^*) \geq 1 - \frac{\delta}{3}$, respectively. Applying a union to the above three inequalities, we conclude that $\mathbb{P}(\mathcal{U}) \geq 1 - \delta$. □

We recall the definitions of the functions \bar{U}_g^k , $\tilde{Q}_{g,h}^k$ and $\tilde{V}_{g,h}^k$ in [Appendix C.2](#). They rely on the pessimistic finite-horizon goal-conditioned values \bar{V}_g^k defined as

$$\bar{Q}_{g,h}^k(s, a) \triangleq \text{clip} \left(\mathbb{1}[s \neq g] + 3 \sqrt{\text{Var}_{\hat{p}^k}(\tilde{V}_{g, h+1}^k)(s, a) \frac{\beta^*(n^k(s, a), \delta)}{n^k(s, a)}} + 14H^2 \frac{\beta^*(n^k(s, a), \delta)}{n^k(s, a)} \right)$$

$$\begin{aligned}
 & + \frac{1}{H} \widehat{p}^k (\overline{V}_{g,h+1}^k - \widetilde{V}_{g,h+1}^k)(s, a) + \widehat{p}^k \overline{V}_{g,h+1}^k(s, a), 0, H \Big), \\
 \overline{V}_{g,h}^k(s) & \triangleq \min_{a \in \mathcal{A}} \overline{Q}_{g,h}^k(s, a), \quad \overline{V}_{g,h}^k(g) \triangleq 0, \quad \overline{V}_{g,H+1}^k(s) \triangleq 0.
 \end{aligned}$$

Finally, we define the following quantities

$$\begin{aligned}
 \overset{\circ}{Q}_{g,h}^k(s, a) & \triangleq \max \left(\mathbf{1}[s \neq g] + p \overset{\circ}{V}_{g,h+1}^k(s, a), \text{clip} \left(\mathbf{1}[s \neq g] + 3 \sqrt{\text{Var}_{\widehat{p}^k}(\widetilde{V}_{g,h+1}^k)(s, a)} \frac{\beta^*(n^k(s, a), \delta)}{n^k(s, a)} \right. \right. \\
 & \quad \left. \left. + 14H^2 \frac{\beta(n^k(s, a), \delta)}{n^k(s, a)} + \frac{1}{H} \widehat{p}^k (\overline{V}_{g,h+1}^k - \widetilde{V}_{g,h+1}^k)(s, a) + \widehat{p}^k \overline{V}_{g,h+1}^k(s, a), 0, H \right) \right), \\
 \overset{\circ}{V}_{g,h}^k(s) & \triangleq \pi_{g,h}^{k+1} \overset{\circ}{Q}_{g,h}^k(s, a), \\
 \overset{\circ}{V}_{g,h}^k(g) & \triangleq 0, \\
 \overset{\circ}{V}_{g,H+1}^k(s) & \triangleq 0.
 \end{aligned}$$

We have the following property, which is the equivalent of [Ménard et al. \(2021a, Lemma 6\)](#) and is proved likewise.

Lemma 18. *On the event \mathcal{U} , for all $(s, a, g, h) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times [H]$ and for every episode k , it holds that*

$$\begin{aligned}
 \overset{\circ}{Q}_{g,h}^k(s, a) & \geq \max \left(\overline{Q}_{g,h}^k(s, a), \overline{Q}_{g,h}^{\pi_g^{k+1}}(s, a) \right), \\
 \overset{\circ}{V}_{g,h}^k(s) & \geq \max \left(\overline{V}_{g,h}^k(s), \overline{V}_{g,h}^{\pi_g^{k+1}}(s) \right).
 \end{aligned}$$

We now derive “key step ①” by establishing that Properties 1 and 2 hold. Specifically, we show that **(i)** the functions $\widetilde{V}_{g,1}$ are optimistic estimates of the optimal goal-conditioned finite-horizon value functions and **(ii)** the functions $\overline{U}_{g,1}$ serve as valid upper bounds to the goal-conditioned finite-horizon gaps, as shown below.

Lemma 19. *On the event \mathcal{U} , it holds that for every episode k and goal $g \in \mathcal{S}$,*

$$\begin{aligned}
 \overline{V}_{g,1}^{\pi_g^{k+1}}(s_0) - \overline{V}_{g,1}^*(s_0) & \leq \overline{V}_{g,1}^{\pi_g^{k+1}}(s_0) - \widetilde{V}_{g,1}^k(s_0) \\
 & \leq \pi_{g,1}^{k+1} \overline{U}_{g,1}^k(s_0).
 \end{aligned}$$

Proof. On the event \mathcal{U} , using [Lemma 18](#), we upper bound the goal-conditioned gap at episode t as

$$\overline{V}_{g,1}^{\pi_g^{k+1}}(s_0) - \overline{V}_{g,1}^*(s_0) \leq \overline{V}_{g,1}^{\pi_g^{k+1}}(s_0) - \widetilde{V}_{g,1}^k(s_0) \leq \overset{\circ}{V}_{g,1}^k(s_0) - \widetilde{V}_{g,1}^k(s_0).$$

Next, following the same reasoning as in [Ménard et al. \(2021a, Proof of Lemma 2\)](#), we obtain by induction on h that for all state-action pairs (s, a) and goal states g ,

$$\overset{\circ}{Q}_{g,h}^k(s, a) - \widetilde{Q}_{g,h}^k(s, a) \leq \overline{U}_{g,h}^k(s, a). \tag{12}$$

In particular for the initial layer $h = 1$ and initial state $s = s_0$, we get that

$$\overset{\circ}{V}_{g,1}^k(s_0) - \widetilde{V}_{g,1}^k(s_0) = \pi_{g,1}^{k+1} (\overset{\circ}{Q}_{g,1}^k - \widetilde{Q}_{g,1}^k)(s_0) \leq \pi_{g,1}^{k+1} \overline{U}_{g,1}^k(s_0).$$

□

C.3.2 Proof of “key step ②”

Lemma 20. *On the event \mathcal{U} , for every goal state $g \in \mathcal{S}$ and episode k , it holds that*

$$\pi_{g,1}^{k+1} \overline{U}_{g,1}^k(s_0) \leq 24e^{13} H \sqrt{\sum_{h=1}^H \sum_{s,a} p_{g,h}^{k+1}(s, a) \frac{\beta^*(\bar{n}^k(s, a), \delta)}{\bar{n}^k(s, a) \vee 1}} + 336e^{13} H^2 \sum_{s,a} \left[\sum_{h=1}^H p_{g,h}^{k+1}(s, a) \frac{\beta(\bar{n}^k(s, a), \delta)}{\bar{n}^k(s, a) \vee 1} \right] \wedge 1,$$

where we recall that $p_{g,h}^{k+1}(s, a)$ denotes the probability of reaching (s, a) at step h under policy π_g^{k+1} .

Proof. Similar to [Ménard et al. \(2021a, Steps 1 & 2 in proof of Theorem 2\)](#), we begin by upper-bounding $\bar{U}_{g,h}^k(s, a)$ for all (s, a, h, g, k) . If $n^k(s, a) > 0$, by definition of $\bar{U}_{g,h}^k$ we have that

$$\bar{U}_{g,h}^k(s, a) \leq 6\sqrt{\text{Var}_{\hat{p}^k}(\tilde{V}_{g,h+1}^k)(s, a) \frac{\beta^*(n^k(s, a), \delta)}{n^k(s, a)}} + 36H^2 \frac{\beta(n^k(s, a), \delta)}{n^k(s, a)} + \left(1 + \frac{3}{H}\right) \hat{p}^k \pi_{g,h+1}^{k+1} \bar{U}_{g,h+1}^k(s, a). \quad (13)$$

We now replace the empirical transition probabilities with the true ones. Using the Bernstein-type technical inequality of [Ménard et al. \(2021a, Lemma 10\)](#) and that $0 \leq \bar{U}_{g,h}^k \leq H$, we get

$$\begin{aligned} (\hat{p}^k - p) \pi_{g,h+1}^{k+1} \bar{U}_{g,h+1}^k(s, a) &\leq \sqrt{2\text{Var}_p(\pi_{g,h+1}^{k+1} \bar{U}_{g,h+1}^k)(s, a) \frac{\beta(n^k(s, a), \delta)}{n^k(s, a)}} + \frac{2}{3} H \frac{\beta(n^k(s, a), \delta)}{n^k(s, a)} \\ &\leq \frac{1}{H} p \pi_{g,h+1}^{k+1} \bar{U}_{g,h+1}^k(s, a) + 3H^2 \frac{\beta(n^k(s, a), \delta)}{n^k(s, a)}, \end{aligned}$$

where in the last line we used $\text{Var}_p(\pi_{g,h+1}^{k+1} \bar{U}_{g,h+1}^k)(s, a) \leq H \pi_{g,h+1}^{k+1} \bar{U}_{g,h+1}^k(s, a)$ and $\sqrt{xy} \leq x + y$ for all $x, y \geq 0$. We then replace the variance of the upper confidence bound under the empirical transition probabilities by the variance of the optimal value function under the true transition probabilities. Using the technical lemmas of [Ménard et al. \(2021a, Lemma 11 & 12\)](#) that control the deviation in variances w.r.t. the choice of transition probabilities, we obtain that

$$\begin{aligned} \text{Var}_{\hat{p}^k}(\tilde{V}_{g,h+1}^k)(s, a) &\leq 2\text{Var}_p(\tilde{V}_{g,h+1}^k)(s, a) + 4H^2 \frac{\beta(n^k(s, a), \delta)}{n^k(s, a)} \\ &\leq 4\text{Var}_p(\bar{V}_{g,h+1}^{\pi_{g,h+1}^{k+1}})(s, a) + 4Hp(\bar{V}_{g,h+1}^{\pi_{g,h+1}^{k+1}} - \bar{V}_{g,h+1}^{\pi_{g,h+1}^{k+1}})(s, a) + 4H^2 \frac{\beta(n^k(s, a), \delta)}{n^k(s, a)} \\ &\leq 4\text{Var}_p(\bar{V}_{g,h+1}^{\pi_{g,h+1}^{k+1}})(s, a) + 4Hp\pi_{g,h+1}^{k+1} \bar{U}_{g,h+1}^k(s, a) + 4H^2 \frac{\beta(n^k(s, a), \delta)}{n^k(s, a)}, \end{aligned}$$

where we used (12) in the last inequality. Next, using $\sqrt{x+y} \leq \sqrt{x} + \sqrt{y}$, $\sqrt{xy} \leq x + y$, and $\beta^*(n, \delta) \leq \beta(n, \delta)$ leads to

$$\begin{aligned} \sqrt{\text{Var}_{\hat{p}^k}(\tilde{V}_{g,h+1}^k)(s, a) \frac{\beta^*(n^k(s, a), \delta)}{n^k(s, a)}} &\leq 2\sqrt{\text{Var}_p(\bar{V}_{g,h+1}^{\pi_{g,h+1}^{k+1}})(s, a) \frac{\beta^*(n^k(s, a), \delta)}{n^k(s, a)}} + (2H + 4H^2) \frac{\beta(n^k(s, a), \delta)}{n^k(s, a)} \\ &\quad + \frac{1}{H} p \pi_{g,h+1}^{k+1} \bar{U}_{g,h+1}^k(s, a) \\ &\leq 2\sqrt{\text{Var}_p(\bar{V}_{g,h+1}^{\pi_{g,h+1}^{k+1}})(s, a) \frac{\beta^*(n^k(s, a), \delta)}{n^k(s, a)}} + 6H^2 \frac{\beta(n^k(s, a), \delta)}{n^k(s, a)} \\ &\quad + \frac{1}{H} p \pi_{g,h+1}^{k+1} \bar{U}_{g,h+1}^k(s, a). \end{aligned}$$

Combining these two inequalities with (13) yields

$$\begin{aligned} \bar{U}_{g,h}^k(s, a) &\leq 12\sqrt{\text{Var}_p(\bar{V}_{g,h+1}^{\pi_{g,h+1}^{k+1}})(s, a) \frac{\beta^*(n^k(s, a), \delta)}{n^k(s, a)}} + 36H^2 \frac{\beta(n^k(s, a), \delta)}{n^k(s, a)} \\ &\quad + \frac{6}{H} p \pi_{g,h+1}^{k+1} \bar{U}_{g,h+1}^k(s, a) + 36H^2 \frac{\beta(n^k(s, a), \delta)}{n^k(s, a)} \\ &\quad + \left(1 + \frac{3}{H}\right) \frac{1}{H} p \pi_{g,h+1}^{k+1} \bar{U}_{g,h+1}^k + \left(1 + \frac{3}{H}\right) 3H^2 \frac{\beta(n^k(s, a), \delta)}{n^k(s, a)} \\ &\quad + \left(1 + \frac{3}{H}\right) p \pi_{g,h+1}^{k+1} \bar{U}_{g,h+1}^k(s, a) \\ &\leq 12\sqrt{\text{Var}_p(\bar{V}_{g,h+1}^{\pi_{g,h+1}^{k+1}})(s, a) \frac{\beta^*(n^k(s, a), \delta)}{n^k(s, a)}} + 84H^2 \frac{\beta(n^k(s, a), \delta)}{n^k(s, a)} \end{aligned}$$

$$+ \left(1 + \frac{13}{H}\right) p \pi_{g,h+1}^{k+1} \bar{U}_{g,h+1}^k(s, a).$$

Since by construction, $\bar{U}_{g,h}^k(s, a) \leq H$, we have that for all $n^k(s, a) \geq 0$,

$$\begin{aligned} \bar{U}_{g,h}^k(s, a) &\leq 12 \sqrt{\text{Var}_p(\bar{V}_{g,h+1}^{\pi_g^{k+1}})(s, a) \left(\frac{\beta^*(n^k(s, a), \delta)}{n^k(s, a)} \wedge 1 \right)} + 84H^2 \left(\frac{\beta(n^k(s, a), \delta)}{n^k(s, a)} \wedge 1 \right) \\ &+ \left(1 + \frac{13}{H}\right) p \pi_{g,h+1}^{k+1} \bar{U}_{g,h+1}^k(s, a). \end{aligned}$$

Unfolding the previous inequality and using $(1 + 13/H)^H \leq e^{13}$ we get

$$\begin{aligned} \pi_{g,1}^{k+1} \bar{U}_{g,1}^k(s_0) &\leq 12e^{13} \sum_{h=1}^H \sum_{s,a} p_{g,h}^{k+1}(s, a) \sqrt{\text{Var}_p(\bar{V}_{g,h+1}^{\pi_g^{k+1}})(s, a) \left(\frac{\beta^*(n^k(s, a), \delta)}{n^k(s, a)} \wedge 1 \right)} \\ &+ 84e^{13} H^2 \sum_{h=1}^H \sum_{s,a} p_{g,h}^{k+1}(s, a) \left(\frac{\beta(n^k(s, a), \delta)}{n^k(s, a)} \wedge 1 \right). \end{aligned}$$

Using that $\pi_{g,1}^{k+1} \bar{U}_{g,1}^k(s_0) \leq H$, we can clip the above bound as follows

$$\begin{aligned} \pi_{g,1}^{k+1} \bar{U}_{g,1}^k(s_0) &\leq 12e^{13} \sum_{h=1}^H \sum_{s,a} p_{g,h}^{k+1}(s, a) \sqrt{\text{Var}_p(\bar{V}_{g,h+1}^{\pi_g^{k+1}})(s, a) \left(\frac{\beta^*(n^k(s, a), \delta)}{n^k(s, a)} \wedge 1 \right)} \\ &+ 84e^{13} H^2 \sum_{s,a} \left[\sum_{h=1}^H p_{g,h}^{k+1}(s, a) \left(\frac{\beta(n^k(s, a), \delta)}{n^k(s, a)} \wedge 1 \right) \right] \wedge 1. \end{aligned} \quad (14)$$

From the technical lemma of [Ménard et al. \(2021a, Lemma 8\)](#) that relates counts to pseudo-counts,

$$\frac{\beta(n^k(s, a), \delta)}{n^k(s, a)} \wedge 1 \leq 4 \frac{\beta(\bar{n}^k(s, a), \delta)}{\bar{n}^k(s, a) \vee 1},$$

thus we can replace the counts by the pseudo-counts in (14) as

$$\begin{aligned} \pi_{g,1}^{k+1} \bar{U}_{g,1}^k(s_0) &\leq 24e^{13} \sum_{h=1}^H \sum_{s,a} p_{g,h}^{k+1}(s, a) \sqrt{\text{Var}_p(\bar{V}_{g,h+1}^{\pi_g^{k+1}})(s, a) \frac{\beta^*(\bar{n}^k(s, a), \delta)}{\bar{n}^k(s, a) \vee 1}} \\ &+ 336e^{13} H^2 \sum_{s,a} \left[\sum_{h=1}^H p_{g,h}^{k+1}(s, a) \frac{\beta(\bar{n}^k(s, a), \delta)}{\bar{n}^k(s, a) \vee 1} \right] \wedge 1. \end{aligned} \quad (15)$$

We now apply the law of total variance (see e.g., [Azar et al., 2017](#) or [Ménard et al., 2021a, Lemma 7](#)) in order to further upper-bound the first sum in (15). In particular, by Cauchy-Schwarz inequality, we obtain

$$\begin{aligned} &\sum_{h=1}^H \sum_{s,a} p_{g,h}^{k+1}(s, a) \sqrt{\text{Var}_p(\bar{V}_{g,h+1}^{\pi_g^{k+1}})(s, a) \frac{\beta^*(\bar{n}^k(s, a), \delta)}{\bar{n}^k(s, a) \vee 1}} \\ &\leq \sqrt{\sum_{h=1}^H \sum_{s,a} p_{g,h}^{k+1}(s, a) \text{Var}_p(\bar{V}_{g,h+1}^{\pi_g^{k+1}})(s, a)} \sqrt{\sum_{h=1}^H \sum_{s,a} p_{g,h}^{k+1}(s, a) \frac{\beta^*(\bar{n}^k(s, a), \delta)}{\bar{n}^k(s, a) \vee 1}} \\ &\leq \sqrt{\mathbb{E}_{\pi_g^{k+1}} \left[\left(\sum_{h=1}^H \mathbb{1}[s_h \neq g] - \bar{V}_{g,1}^{\pi_g^{k+1}}(s_0) \right)^2 \right]} \sqrt{\sum_{h=1}^H \sum_{s,a} p_{g,h}^{k+1}(s, a) \frac{\beta^*(\bar{n}^k(s, a), \delta)}{\bar{n}^k(s, a) \vee 1}} \\ &\leq H \sqrt{\sum_{h=1}^H \sum_{s,a} p_{g,h}^{k+1}(s, a) \frac{\beta^*(\bar{n}^k(s, a), \delta)}{\bar{n}^k(s, a) \vee 1}}. \end{aligned}$$

Plugging this in (15) concludes the proof of [Lemma 20](#). \square

We are now ready to derive “key step ②” which controls the cumulative gap bounds, see Equation 4.

Lemma 21. *On the event \mathcal{U} , for any number of episodes $K \geq 1$, it holds that*

$$\begin{aligned} \sum_{k=0}^{K-1} \pi_{g_{k+1},1}^{k+1} \bar{U}_{g_{k+1},1}^k(s_0) &\leq 48e^{13} \sqrt{K} \sqrt{H^2 SA \log(HK+1) \beta^*(K, \delta)} + 1344e^{13} H^2 SA \beta(K, \delta) \log(HK+1) \\ &\quad + 48e^{13} H^2 SA \sqrt{\beta^*(K, \delta)}. \end{aligned}$$

Proof. Plugging in the bound of Lemma 20 yields

$$\begin{aligned} &\sum_{k=0}^{K-1} \pi_{g_{k+1},1}^{k+1} \bar{U}_{g_{k+1},1}^k(s_0) \\ &\leq 24e^{13} H \sum_{k=0}^{K-1} \sqrt{\sum_{h=1}^H \sum_{s,a} p_{g_{k+1},h}^{k+1}(s,a) \frac{\beta^*(\bar{n}^k(s,a), \delta)}{\bar{n}^k(s,a) \vee 1}} + 336e^{13} H^2 \sum_{k=0}^{K-1} \sum_{s,a} \left[\sum_{h=1}^H p_{g,h}^{k+1}(s,a) \frac{\beta(\bar{n}^k(s,a), \delta)}{\bar{n}^k(s,a) \vee 1} \right] \wedge 1 \\ &\leq 24e^{13} H \sqrt{\beta^*(HK, \delta)} \sum_{k=0}^{K-1} \sqrt{\sum_{s,a} \frac{\bar{n}^{k+1}(s,a) - \bar{n}^k(s,a)}{\bar{n}^k(s,a) \vee 1}} + 336e^{13} H^2 \beta(HK, \delta) \sum_{s,a} \sum_{k=0}^{K-1} \left[\frac{\bar{n}^{k+1}(s,a) - \bar{n}^k(s,a)}{\bar{n}^k(s,a) \vee 1} \right] \wedge 1, \end{aligned}$$

where we used that $\beta(\cdot, \delta)$ and $\beta^*(\cdot, \delta)$ are increasing. We define $\mathcal{J} \triangleq \{k \in [0, K-1] : \bar{n}^k(s,a) < \bar{n}^{k+1}(s,a) - \bar{n}^k(s,a) - 1\}$. Applying Lemma 22 gives that

$$\begin{aligned} \sum_{k \in \mathcal{J}} \sqrt{\sum_{s,a} \frac{\bar{n}^{k+1}(s,a) - \bar{n}^k(s,a)}{\bar{n}^k(s,a) \vee 1}} &\leq \underbrace{\sum_{s,a} \sum_{k \in \mathcal{J}} \sqrt{\frac{\bar{n}^{k+1}(s,a) - \bar{n}^k(s,a)}{\bar{n}^k(s,a) \vee 1}}}_{\leq 2H} \leq 2SAH, \\ \sum_{k \notin \mathcal{J}} \sqrt{\sum_{s,a} \frac{\bar{n}^{k+1}(s,a) - \bar{n}^k(s,a)}{\bar{n}^k(s,a) \vee 1}} &\leq \sqrt{K} \sqrt{\underbrace{\sum_{s,a} \sum_{k \notin \mathcal{J}} \frac{\bar{n}^{k+1}(s,a) - \bar{n}^k(s,a)}{\bar{n}^k(s,a) \vee 1}}_{\leq 4 \log(HK+1)}} \leq 2\sqrt{KSA \log(HK+1)}, \\ \sum_{s,a} \underbrace{\sum_{k=0}^{K-1} \left[\frac{\bar{n}^{k+1}(s,a) - \bar{n}^k(s,a)}{\bar{n}^k(s,a) \vee 1} \right] \wedge 1}_{\leq 4 \log(HK+1)} &\leq 4SA \log(HK+1). \end{aligned}$$

Putting everything together yields Lemma 21. Note that as opposed to [Ménard et al. \(2021a\)](#), we are in the setting of stationary transition probabilities (and cost functions), which is why we are able to shave a factor H in the main order term of the bound of the cumulative gap bounds (also recall that their sample complexity bound is in terms of exploration episodes and not exploration steps as ours). \square

Lemma 22 (Technical lemma). *For $T \in \mathbb{N}^*$ and $(u_t)_{t \in \mathbb{N}^*}$, for any sequence where $u_t \in [0, H]$ for some constant $H > 0$ and $U_t \triangleq \sum_{l=1}^t u_l$, let $\Omega \triangleq \{t \in [0, T] : U_t < u_{t+1} - 1\}$ and $\omega \triangleq \max\{t \in \Omega\}$. Then it holds that*

$$\begin{aligned} \sum_{t \in \Omega} \sqrt{\frac{u_{t+1}}{U_t \vee 1}} &\leq 2H, \\ \sum_{t \notin \Omega} \frac{u_{t+1}}{U_t \vee 1} &\leq 4 \log(U_{T+1} + 1), \\ \sum_{t=0}^T \left[\frac{u_{t+1}}{U_t \vee 1} \right] \wedge 1 &\leq 4 \log(U_{T+1} + 1). \end{aligned}$$

Proof. First, note that for any $t \in \Omega$, $\frac{u_{t+1}}{U_t \vee 1} \geq 1$, therefore

$$\sum_{t \in \Omega} \sqrt{\frac{u_{t+1}}{U_t \vee 1}} \leq \sum_{t \in \Omega} \frac{u_{t+1}}{U_t \vee 1} \leq \sum_{t \in \Omega} u_{t+1} \leq \sum_{t=0}^{\omega} u_{t+1} = U_{\omega+1} = U_{\omega} + u_{\omega} \leq u_{\omega+1} - 1 + u_{\omega} \leq 2H.$$

Second, if $t \notin \Omega$, then $2U_t + 2 \geq U_{t+1} + 1$, therefore

$$\frac{u_{t+1}}{U_t \vee 1} \leq 4 \frac{u_{t+1}}{2U_t + 2} \leq 4 \frac{U_{t+1} - U_t}{U_{t+1} + 1},$$

which yields that

$$\sum_{t \notin \Omega} \frac{u_{t+1}}{U_t \vee 1} \leq 4 \sum_{t \notin \Omega} \frac{U_{t+1} - U_t}{U_{t+1} + 1} \leq 4 \sum_{t=0}^T \int_{U_t}^{U_{t+1}} \frac{1}{x+1} dx \leq 4 \log(U_{T+1} + 1).$$

Third, combining the two cases above and further noticing that $4 \frac{U_{t+1} - U_t}{U_{t+1} + 1} = \frac{4u_{t+1}}{U_t + u_{t+1} + 1} \geq \frac{4u_{t+1}}{2u_{t+1}} \geq 1$ for all $t \in \Omega$, it holds that

$$\left[\frac{u_{t+1}}{U_t \vee 1} \right] \wedge 1 \leq 4 \frac{U_{t+1} - U_t}{U_{t+1} + 1},$$

thus

$$\sum_{t=0}^T \left[\frac{u_{t+1}}{U_t \vee 1} \right] \wedge 1 \leq 4 \sum_{t=0}^T \frac{U_{t+1} - U_t}{U_{t+1} + 1} \leq 4 \sum_{t=0}^T \int_{U_t}^{U_{t+1}} \frac{1}{x+1} dx \leq 4 \log(U_{T+1} + 1).$$

□

C.3.3 Proof of “key step ③”

Lemma 23. *The MGE sample complexity τ of algorithm ADAGoAL-UCBVI can be bounded with probability at least $1 - \delta$ by⁷*

$$\tau = O\left(\frac{L^3 SA}{\varepsilon^2} \cdot \log^3\left(\frac{LSA}{\varepsilon \delta}\right) \cdot \log^3\left(\frac{L}{\varepsilon}\right) + \frac{L^3 S^2 A}{\varepsilon} \cdot \log^3\left(\frac{LSA}{\varepsilon \delta}\right) \cdot \log^3\left(\frac{L}{\varepsilon}\right) \right).$$

Proof. We assume that the event \mathcal{U} holds and fix a (finite) episode $K < \kappa$, where κ denotes the (possibly unbounded) episode index at which ADAGoAL-UCBVI terminates. For any $k \leq K$, denote by g_k the goal selected by ADAGoAL-UCBVI at the beginning of episode k , then by design of the stopping rule (1) and by choice of prediction error \mathcal{E}_k (10), it holds that

$$\varepsilon \leq \pi_{g_k, 1}^k \bar{U}_{g_k, 1}^{k-1}(s_0) + \frac{8\varepsilon}{9}.$$

By summing the previous inequality for all $k \leq K$ and plugging in the bound of Lemma 21, we get that

$$\begin{aligned} \frac{\varepsilon}{9} K &\leq 48e^{13} \sqrt{K} \sqrt{H^2 SA \log(HK + 1) \beta^*(K, \delta)} + 1344e^{13} H^2 SA \beta(K, \delta) \log(HK + 1) \\ &\quad + 48e^{13} H^2 SA \sqrt{\beta^*(K, \delta)}. \end{aligned}$$

We assume that $\kappa > 1$ otherwise the result is trivially true. Defining $\kappa' \triangleq \min\{\kappa - 1, K\}$ and using the definition of β and β^* given in Lemma 17, we get the following functional inequality in κ'

$$\varepsilon \kappa' \leq x_1 \sqrt{\kappa' H^2 SA \log(H\kappa')} (\log(3S^2 AH/\delta) + \log(16e\kappa')) + x_2 H^2 S^2 A \log(3S^2 AH/\delta) \log^2(H\kappa'),$$

for some absolute constants x_1, x_2 . There remains to invert the above inequality to obtain an upper bound on κ' . We use the auxiliary inequality of Lemma 24 instantiated with scalars $B = x_1 \sqrt{H^2 SA \log(3S^2 AH/\delta)}/\varepsilon$, $C = x_2 H^2 S^2 A \log(3S^2 AH/\delta)/\varepsilon$ and $\alpha = 16eH$. This yields that

$$\kappa' \leq O\left(\frac{H^2 SA}{\varepsilon^2} \log^3\left(\frac{HSA}{\varepsilon \delta}\right) + \frac{H^2 S^2 A}{\varepsilon} \log^3\left(\frac{HSA}{\delta \varepsilon}\right) \right). \quad (16)$$

⁷We say that $f = O(g)$ if there exists an absolute constant $\iota > 0$ (independent of the MDP instance) such that $f \leq \iota g$.

Since (16) holds for $\kappa' = \min\{\kappa - 1, K\}$ for any finite $K < \kappa$, letting $K \rightarrow +\infty$ implies that κ is finite and bounded as in (16).

The last step is to relate the above bound on the number of algorithmic episodes κ to the MGE sample complexity of ADAGOAL-UCBVI denoted by τ . Since the algorithmic episodes are of length H and separated by a one-step execution of the reset action a_{reset} , it holds that $\tau \leq (H + 1)\kappa$. We finally plug in the choice of horizon $H \triangleq \lceil 5(L + 2) \log(10(L + 2)/\varepsilon) / \log(2) \rceil$ to conclude the proof of Lemma 23. \square

Lemma 24 (An auxiliary inequality). *For any positive scalars $B, C \geq 1$ and $\alpha \geq e$, it holds for any $X \geq 2$ that*

$$X \leq B\sqrt{X} \log(\alpha X) + C \log^2(\alpha X) \implies X \leq O(B^2 \log^2(\alpha B) + C \log^2(\alpha C)).$$

Proof. On the one hand, assume that $X \leq B\sqrt{X} \log(\alpha X)$, then $\frac{X}{2} \leq -\frac{X}{2} + B\sqrt{X} \log(\alpha X)$. From the technical lemma of [Kazerouni et al. \(2017, Lemma 8\)](#), $-\frac{X}{2} + B\sqrt{X} \log(\alpha X) \leq \frac{32B^2}{9} [\log(4B\sqrt{\alpha e})]^2$, thus $X \leq \frac{64B^2}{9} [\log(4B\sqrt{\alpha e})]^2$. On the other hand, assume that $X \leq C \log^2(\alpha X)$. Using that $\log(x) \leq x^\beta / \beta$ for all $x \geq 0, \beta > 0$, we get $X \leq C(8\alpha^{1/8} X^{1/8})^2 \leq 64C\alpha^{1/4} X^{1/4}$, thus $X \leq (64C)^{4/3} \alpha^{1/3}$, thus $X \leq C \log^2(64\alpha^{4/3} C^{4/3})$. Now, assume that $X \leq B\sqrt{X} \log(\alpha X) + C \log^2(\alpha X)$. Then $X \leq 2 \max\{B\sqrt{X} \log(\alpha X), C \log^2(\alpha X)\}$. From above we can bound each term separately, which concludes the proof. \square

C.3.4 Proof of “key step ④”

We finally establish “key step ④”, which focuses on the technical novelty of the ADAGOAL goal selection scheme that is specific to the multi-goal exploration setting.

Recall for any $H \in \mathbb{N}^*$ the definition of the finite-horizon MDP $\overline{\mathcal{M}}_{g,H} = \{H, S, A, \overline{p}_g, \overline{c}_g\}$, where we recall that $\overline{p}_g = p_g$ and $\overline{c}_g = c_g$. We denote by $\overline{\pi}_{g,H}^*$ the optimal policy in $\overline{\mathcal{M}}_{g,H}$ as well as $\overline{V}_{g,H,h}(s)$ the optimal value function starting from state s at step h . We also define $\overline{p}_{g,H}^{\overline{\pi}_{g,H}^*}(s_H \neq g | s_1 = s)$ the probability of reaching state g starting from state s with the policy $\pi \in \Pi_H$ in the MDP $\overline{\mathcal{M}}_{g,H}$. When it is clear from the context we drop the dependence on the horizon H in the previous notations.

The following lemma controls the probability of not reaching a goal in $\mathcal{G}_{L+\varepsilon}$ with the optimal policy in the finite-horizon reduction MDP.

Lemma 25. *For $g \in \mathcal{G}_{L+\varepsilon}$, for all $H \geq 2(L + 2)$, for all $s \in S$,*

$$\overline{p}_{g,H}^{\overline{\pi}_{g,H}^*}(s_H \neq g | s_1 = s) \leq e^{-\log(2)H / (4(L+2))}.$$

Proof. By induction it holds

$$\begin{aligned} \overline{p}_{g,H}^{\overline{\pi}_{g,H}^*}(s_H \neq g | s_1 = s) &= \sum_{s' \neq g} \overline{p}_{g,H}^{\overline{\pi}_{g,H}^*}(s_{H-M+1} = s' | s_1 = s) \overline{p}_{g,H}^{\overline{\pi}_{g,H}^*}(s_H \neq g | s_{H-M+1} = s') \\ &\leq \overline{p}_{g,H}^{\overline{\pi}_{g,H}^*}(s_{H-M+1} \neq g | s_1 = s) \max_{s'} \overline{p}_{g,H}^{\overline{\pi}_{g,H}^*}(s_H \neq g | s_{H-M+1} = s') \\ &\leq \prod_{j=0}^{\lfloor H/M \rfloor} \max_{s'} \overline{p}_{g,H}^{\overline{\pi}_{g,H}^*}(s_{H-jM} \neq g | s_{H-(j+1)M+1} = s'). \end{aligned}$$

Then thanks to the Markov inequality and the optimal Bellman equations solved by $\overline{\pi}_{g,H}^*$ we obtain

$$\begin{aligned} \overline{p}_{g,H}^{\overline{\pi}_{g,H}^*}(s_{H-jM} \neq g | s_{H-(j+1)M+1} = s') &\leq \overline{p}_{g,H}^{\overline{\pi}_{g,H}^*}(s_H \neq g | s_{H-(j+1)M+1} = s') \\ &\leq \frac{\overline{V}_{g,H,H-(j+1)M+1}^*(s')}{M} \\ &= \frac{\overline{V}_{g,(j+1)M,1}^*(s')}{M} \\ &\leq \frac{\overline{V}_{g,(j+1)M,1}^*(s')}{M} \leq \frac{V_g^*(s')}{M} \leq \frac{L+2}{M}, \end{aligned}$$

where the last inequality uses the existence of a resetting action (Assumption 3) and the fact that $g \in \mathcal{G}_{L+\varepsilon}$ with $\varepsilon \leq 1$. Choosing $M = 2(L+2)$ allows us to conclude

$$\bar{p}_{g,H}^{\pi^*}(s_H \neq g | s_1 = s) \leq e^{-\lfloor H/M \rfloor \log(2)} \leq e^{-\log(2)H/(4(L+2))},$$

where in the last inequality we used $\lfloor x \rfloor \geq x/2$ for $x \geq 1$. \square

We define the class of non-stationary, infinite-horizon policies that perform the reset action whenever the goal state is not reached after H steps.

Definition 26 (Resetting policies). *For any π , we denote by $\pi^{\lfloor H \rfloor}$ the non-stationary policy that, until the goal is reached, successively executes the actions prescribed by π for H steps and takes action a_{reset} , i.e., at time step i and state s it executes the following action:*

$$\pi^{\lfloor H \rfloor}(a|s, i) \triangleq \begin{cases} a_{\text{reset}} & \text{if } i \equiv 0 \pmod{H+1}, \\ \pi(a|s, i) & \text{otherwise.} \end{cases}$$

We denote by $\Pi^{\lfloor H \rfloor}$ the set of such resetting policies.

We now establish two key lemmas. First, we show that, equipped with a near-optimal policy for the finite-horizon model $\bar{\mathcal{M}}_{g,H}$, expanding it into an infinite-horizon policy via the reset provides a near-optimal goal-reaching policy in the original MDP \mathcal{M}_g as long as the goal state g belongs to $\mathcal{G}_{O(L)}$ and the horizon H is large enough.

Lemma 27. *For $g \in \mathcal{G}_{L+\varepsilon}$ and $H \geq 5(L+2) \log(10(L+2)/\varepsilon) / \log(2)$, it holds that*

$$\bar{V}_{g,1}^*(s_0) \leq V_g^*(s_0) \leq \bar{V}_{g,1}^*(s_0) + \frac{\varepsilon}{9},$$

and if a policy $\tilde{\pi}$ is $\varepsilon/9$ -optimal in $\bar{\mathcal{M}}_{g,H}$ then

$$V_g^{\tilde{\pi}^{\lfloor H \rfloor}}(s_0) \leq V_g^*(s_0) + \varepsilon.$$

Proof. We have trivially $\bar{V}_{g,H,1}^*(s_0) \leq V_g^*(s_0)$. Thanks to Lemma 25 it holds

$$\bar{q}_{g,H}^* \triangleq \bar{p}_{g,H}^{\pi^*}(s_{H+1} \neq g | s_1 = s_0) \leq \bar{p}_{g,H}^{\pi^*}(s_H \neq g | s_1 = s_0) \leq \frac{\varepsilon}{10(L+2)} \leq \frac{1}{30}. \quad (17)$$

Thanks to (17) and the definition of a resetting policy, we can conclude that

$$\begin{aligned} V_g^*(s_0) &\leq V_g^{\bar{\pi}_{g,H}^*}{}^{\lfloor H \rfloor}(s_0) = \bar{V}_{g,H,1}^*(s_0) + \bar{q}_{g,H}^*(1 + V_g^{\bar{\pi}_{g,H}^*}{}^{\lfloor H \rfloor}(s_0)) \\ &= \bar{V}_{g,H,1}^*(s_0) + \frac{\bar{q}_{g,H}^*}{1 - \bar{q}_{g,H}^*}(1 + \bar{V}_{g,H,1}^*(s_0)) \\ &\leq \bar{V}_{g,H,1}^*(s_0) + \frac{30}{29} \frac{\varepsilon}{10(L+2)}(1 + L + \varepsilon) \\ &\leq \bar{V}_{g,H,1}^*(s_0) + \frac{\varepsilon}{9}. \end{aligned}$$

Thus it holds that

$$\bar{V}_{g,H,1}^*(s_0) \leq V_g^*(s_0) \leq \bar{V}_{g,H,1}^*(s_0) + \frac{\varepsilon}{9}. \quad (18)$$

For the second part of the lemma, first note that

$$\bar{V}_{g,H,1}^{\tilde{\pi}}(s_0) = \sum_{h=1}^H \bar{p}_{g,H}^{\tilde{\pi}}(s_h \neq g | s_1 = s_0) \geq \bar{V}_{g,H-L,1}^{\tilde{\pi}}(s_0) + L \bar{p}_{g,H}^{\tilde{\pi}}(s_H \neq g | s_1 = s_0),$$

where in the inequality we used that $\bar{p}_{g,H}^{\tilde{\pi}}(s_h \neq g | s_1 = s_0) \geq \bar{p}_{g,H}^{\tilde{\pi}}(s_H \neq g | s_1 = s_0)$. Using successively the fact that $\tilde{\pi}$ is $\frac{\varepsilon}{9}$ -optimal in $\bar{\mathcal{M}}_{g,H}$, the inequality above and (18) we obtain

$$V_g^*(s_0) + \frac{\varepsilon}{9} \geq \bar{V}_{g,H,1}^*(s_0) + \frac{\varepsilon}{9}$$

$$\begin{aligned}
 &\geq \bar{V}_{g,H,1}^{\tilde{\pi}}(s_0) \\
 &\geq \bar{V}_{g,H-L,1}^{\tilde{\pi}}(s_0) + L\bar{p}_{g,H}^{\tilde{\pi}}(s_H \neq g|s_0) \\
 &\geq V_g^*(s_0) - \frac{\varepsilon}{9} + L\bar{p}_{g,H}^{\tilde{\pi}}(s_H \neq g|s_0).
 \end{aligned}$$

The previous sequence of inequalities entails that $\bar{p}_{g,H}^{\tilde{\pi}}(s_H \neq g) \leq (2\varepsilon)/(9L)$. Now we can upper bound the value of the resetting extension of $\tilde{\pi}$. Indeed, for $\tilde{q} \triangleq \bar{p}_{g,H}^{\tilde{\pi}}(s_{H+1} \neq g|s_1 = s_0) \leq \bar{p}_{g,H}^{\tilde{\pi}}(s_H \neq g|s_1 = s_0)$ we have using that $\tilde{\pi}$ is $\frac{\varepsilon}{9}$ -optimal in $\bar{\mathcal{M}}_{g,H}$ with $g \in \mathcal{G}_{L+\varepsilon}$ that

$$\begin{aligned}
 V_g^{\tilde{\pi}|H}(s_0) &= \bar{V}_{g,H,1}^{\tilde{\pi}}(s_0) + \frac{\tilde{q}}{1-\tilde{q}}(1 + \bar{V}_{g,H,1}^{\tilde{\pi}}(s_0)) \\
 &\leq \bar{V}_{g,H,1}^*(s_0) + \frac{\varepsilon}{9} + \frac{2\varepsilon}{9L} \frac{1}{1-2/9} \left(1 + L + \varepsilon + \frac{\varepsilon}{9}\right) \\
 &\leq V_g^*(s_0) + \varepsilon.
 \end{aligned}$$

□

The second key lemma that we prove is that any goal state that meets the constraint (2b) with small enough prediction error (2a) must belong to $\mathcal{G}_{L+\varepsilon}$.

Restatement of Lemma 12. *With probability at least $1 - \delta$, if a goal state $g \in \mathcal{G}$ satisfies $\mathcal{D}_k(g) \leq L$ and $\mathcal{E}_k(g) \leq \varepsilon$ for an episode $k \geq 1$, then $g \in \mathcal{G}_{L+\varepsilon}$.*

Proof. Consider that the event \mathcal{U} defined in (11) holds. Consider a goal state g such that $\mathcal{D}_k(g) \leq L$ and $\mathcal{E}_k(g) \leq \varepsilon$ at an episode $k \geq 1$. Then

$$\begin{aligned}
 \bar{V}_{g,H,1}^*(s_0) &\stackrel{(i)}{\leq} \bar{V}_{g,H,1}^{\tilde{\pi}^k}(s_0) + \pi_{g,1}^{k+1} \bar{U}_{g,1}^k(s_0) \\
 &\stackrel{(ii)}{=} \mathcal{D}_k(g) + \mathcal{E}_k(g) - \frac{8\varepsilon}{9} \\
 &\stackrel{(iii)}{\leq} L + \frac{\varepsilon}{9},
 \end{aligned} \tag{19}$$

where (i) comes from Lemma 19, (ii) stems from the choice of \mathcal{D} and \mathcal{E} estimates and (iii) comes from the conditions on g . Following the steps of the proof of Lemma 25, we have that

$$\bar{p}_{g,H}^{\tilde{\pi}^k}(s_H \neq g|s_1 = s) \leq \prod_{j=0}^{\lfloor H/M \rfloor} \max_{s'} \bar{p}_{g,H}^{\tilde{\pi}^k}(s_{H-jM} \neq g|s_{H-(j+1)M+1} = s'),$$

and that

$$\bar{p}_{g,H}^{\tilde{\pi}^k}(s_{H-jM} \neq g|s_{H-(j+1)M+1} = s') \leq \frac{\bar{V}_{g,(j+1)M,1}^*(s')}{M} \leq \frac{\bar{V}_{g,H,1}^*(s')}{M} \leq \frac{1 + \bar{V}_{g,H,1}^*(s_0)}{M} \leq \frac{1 + L + \varepsilon/9}{M},$$

where the before last inequality uses the existence of a resetting action (Assumption 3) and the last inequality uses (19). Choosing $M = 2(L + 2)$ gives

$$\bar{p}_{g,H}^{\tilde{\pi}^k}(s_H \neq g|s_1 = s) \leq e^{-\log(2)H/(4(L+2))}. \tag{20}$$

We now follow the steps of the proof of Lemma 27. Thanks to (20) and the choice of H it holds

$$\bar{q}_{g,H}^* \triangleq \bar{p}_{g,H}^{\tilde{\pi}^k}(s_{H+1} \neq g|s_1 = s_0) \leq \bar{p}_{g,H}^{\tilde{\pi}^k}(s_H \neq g|s_1 = s_0) \leq \frac{\varepsilon}{10(L+2)} \leq \frac{1}{30}. \tag{21}$$

Thanks to (21) and the definition of a resetting policy, we obtain that

$$V_g^*(s_0) \leq V_g^{\tilde{\pi}^k|H}(s_0) = \bar{V}_{g,H,1}^*(s_0) + \bar{q}_{g,H}^*(1 + V_g^{\tilde{\pi}^k|H}(s_0))$$

$$\begin{aligned}
 &= \bar{V}_{g,H,1}^*(s_0) + \frac{\bar{q}_{g,H}^*}{1 - \bar{q}_{g,H}^*} (1 + \bar{V}_{g,H,1}^*(s_0)) \\
 &\stackrel{(i)}{\leq} L + \frac{\varepsilon}{9} + \frac{30}{29} \frac{\varepsilon}{10(L+2)} \left(1 + L + \frac{\varepsilon}{9}\right) \\
 &\leq L + \frac{2\varepsilon}{9},
 \end{aligned}$$

where (i) uses (19). Therefore we have that $g \in \mathcal{G}_{L+\varepsilon}$, which concludes the proof. \square

We now have all the tools to prove that when ADAGoAL-UCBVI terminates, it fulfills the MGE objective of Definition 4.

Lemma 28. *If the algorithm ADAGoAL-UCBVI stops, it is (ε, δ, L) -PAC for MGE.*

Proof. Consider that the event \mathcal{U} defined in (11) holds, and that the algorithm ADAGoAL-UCBVI has stopped at episode κ . Recall that we define $\mathcal{D}_\kappa(g) \triangleq \bar{V}_{g,1}^{\kappa}(s_0)$ and $\mathcal{E}_\kappa(g) \triangleq \pi_{g,1}^{\kappa+1} \bar{U}_{g,1}^\kappa(s_0) + \frac{8\varepsilon}{9}$. Denoting $\mathcal{X}_\kappa \triangleq \{g \in \mathcal{S} : \mathcal{D}_\kappa(g) \leq L\}$, the stopping rule (Equation 1) implies that $\max_{g \in \mathcal{X}_\kappa} \mathcal{E}_\kappa(g) \leq \varepsilon$. We now prove that $\mathcal{G}_L \subseteq \mathcal{X}_\kappa \subseteq \mathcal{G}_{L+\varepsilon}$. On the one hand, it holds that

$$\mathcal{D}_\kappa(g) = \bar{V}_{g,1}^{\kappa}(s_0) \leq \bar{V}_{g,1}^*(s_0) \leq V_g^*(s_0),$$

which ensures that $\mathcal{G}_L \subseteq \mathcal{X}_\kappa$. On the other hand, consider that $g \in \mathcal{X}_\kappa$, then $\mathcal{D}_\kappa(g) \leq L$ and $\mathcal{E}_\kappa(g) \leq \varepsilon$, which implies that $g \in \mathcal{G}_{L+\varepsilon}$ from Lemma 12, therefore $\mathcal{X}_\kappa \subseteq \mathcal{G}_{L+\varepsilon}$.

We now prove that the candidate policies of ADAGoAL-UCBVI are near-optimal goal-reaching policies. Consider any $g \in \mathcal{X}_\kappa$. Combining the result of Lemma 19 and Equation 1, we obtain that

$$\bar{V}_{g,1}^{\pi_g^{\kappa+1}}(s_0) - \bar{V}_{g,1}^*(s_0) \leq \pi_{g,1}^{\kappa+1} \bar{U}_{g,1}^\kappa(s_0) \leq \mathcal{E}_\kappa(g) - \frac{8\varepsilon}{9} \leq \frac{\varepsilon}{9},$$

thus the policy $\pi_g^{\kappa+1}$ is $\frac{\varepsilon}{9}$ -optimal in $\bar{\mathcal{M}}_{g,H}$. As a result, denoting by $\hat{\pi}_g \triangleq (\pi_g^{\kappa+1})^H$ the candidate policy of ADAGoAL-UCBVI, we have from Lemma 27 that

$$V_g^{\hat{\pi}_g}(s_0) \leq V_g^*(s_0) + \varepsilon,$$

i.e., $\hat{\pi}_g$ is ε -optimal for the original SSP objective. Putting everything together, we have that

$$\mathbb{P}\left(\{\mathcal{G}_L \subseteq \mathcal{X}_\kappa \subseteq \mathcal{G}_{L+\varepsilon}\} \cap \left\{\forall g \in \mathcal{X}_\kappa, V_g^{\hat{\pi}_g}(s_0 \rightarrow g) - V_g^*(s_0 \rightarrow g) \leq \varepsilon\right\}\right) \geq \mathbb{P}(\mathcal{U}) \geq 1 - \delta.$$

which ensures that ADAGoAL-UCBVI is (ε, δ, L) -PAC for MGE. \square

C.4 Putting everything together

Restatement of Theorem 8. *ADAGoAL-UCBVI is $(\varepsilon, \delta, L, \mathcal{S})$ -PAC for MGE and, with probability at least $1 - \delta$, for $\varepsilon \in (0, 1/S]$ its MGE sample complexity is of order⁸ $\tilde{O}(L^3 S A \varepsilon^{-2})$.*

Proof. The result comes from combining Lemmas 23 and 28. \square

⁸The notation \tilde{O} in Theorem 8 hides poly-log terms in $\varepsilon^{-1}, S, A, L, \delta^{-1}$. See Lemma 23 in Appendix C.3 for a more detailed bound that includes the poly-log terms.

D DETAILS OF ADAGOAL-UCRL-VTR AND ANALYSIS

In this section, we provide details on the ADAGOAL-UCRL-VTR algorithm and the guarantee of Theorem 11 which bounds its MGE sample complexity in linear mixture MDPs. Recall that since the state space \mathcal{S} may be large, we consider that the known goal space is in all generality a subset of it, i.e., $\mathcal{G} \subseteq \mathcal{S}$, where $G \triangleq |\mathcal{G}|$ denotes the cardinality of the goal space.

First of all, we extend the linear mixture definition (Definition 7) to handle our multi-goal setting. For any goal $g \in \mathcal{G}$, we define

$$p_g(s'|s, a) \triangleq \langle \phi_g(s'|s, a), \theta_g^* \rangle, \quad \text{with} \quad \theta_g^* \triangleq \begin{pmatrix} \theta^* \\ 1 \end{pmatrix} \in \mathbb{R}^{d+1}, \quad \phi_g(s'|s, a) \triangleq \begin{pmatrix} \mathbb{1}[s \neq g] \phi(s'|s, a) \\ \mathbb{1}[s = g] \mathbb{1}[s' = g] \end{pmatrix} \in \mathbb{R}^{d+1}.$$

We see that by construction,

$$p_g(s'|s, a) = \begin{cases} p(s'|s, a) & \text{if } s \neq g \\ \mathbb{1}[s' = g] & \text{if } s = g. \end{cases}$$

D.1 Overview of ADAGOAL-UCRL-VTR and Choice of \mathcal{E} , \mathcal{D} , \mathcal{Q} in line 14 of Algorithm 1

Here we focus on the specificities of ADAGOAL-UCRL-VTR in the linear mixture MDP setting (refer to Section 3 for the description of the algorithmic structure that is common to ADAGOAL-UCBVI), i.e., we explain how to define the estimates \mathcal{D} , \mathcal{E} , \mathcal{Q} in line 14 of Algorithm 1. At a high level, ADAGOAL-UCRL-VTR uses two regression-based goal-conditioned estimators of the unknown parameter vector θ_g^* of each goal $g \in \mathcal{G}$:

- *Value-targeted estimator.* The first estimator minimizes a ridge regression problem with the target being the past value functions. This is similar to the UCRL-VTR algorithm for linear mixture MDPs (Ayoub et al., 2020) and follow-up work (e.g., Zhou et al., 2021; Zhang et al., 2021a). This step is used to compute the distance estimates \mathcal{D} (and \mathcal{Q}) for ADAGOAL.
- *Error-targeted estimator.* The second estimator is novel and minimizes a ridge regression problem with the target being past “error functions”, that are computable upper bounds on the goal-conditioned gaps. This step is used to compute the prediction errors \mathcal{E} for ADAGOAL.

▷ *Value-targeted estimator.* First, ADAGOAL-UCRL-VTR builds a goal-conditioned estimator θ_g for the unknown parameter vector θ_g^* of each goal $g \in \mathcal{G}$, as well as a goal-conditioned covariance matrix Σ_g of the feature mappings, which characterizes the uncertainty of the estimator θ_g . Similar to UCRL-VTR, θ_g is computed as the minimizer to a ridge regression problem with the target being the past value functions, i.e.,

$$\theta_{g,k+1} \leftarrow \arg \min_{\theta \in \mathbb{R}^{d+1}} \lambda \|\theta\|^2 + \sum_{k'=1}^k \sum_{h=1}^H \left(\langle \theta, \phi_{V_{g,k',h}}(s_h^{k'}, a_h^{k'}) \rangle - V_{g,k',h}(s_{h+1}^{k'}) \right),$$

which has a closed-form solution given in (23). Leveraging θ_g and subtracting an exploration bonus term, ADAGOAL-UCRL-VTR builds optimistic goal-conditioned estimators $Q_{g,k,h}(\cdot, \cdot)$ (25) and $V_{g,k,h}(\cdot)$ (27) for the optimal action-value and value functions $\bar{Q}_{g,h}^*(\cdot, \cdot)$ and $\bar{V}_{g,h}^*(\cdot)$. The associated goal-conditioned policy is the greedy policy of the calculated optimistic Q -values (26).

▷ *Error-targeted estimator.* The main addition compared to existing works on linear mixture MDPs is that ADAGOAL-UCRL-VTR also builds goal-conditioned errors denoted by $U_{g,k,h}$ (28) that upper bound the goal-conditioned gaps (defined as the difference between the value function of the current greedy policy and the optimistic value estimates). They rely on an additional estimator $\hat{\theta}_{g,k}$ and covariance matrix $\hat{\Sigma}_{g,k}$ based on the errors $\{U_{g,k',h}\}_{k' \leq k-1, h}$, instead of the values $\{V_{g,k',h}\}_{k' \leq k-1, h}$ as considered before. Specifically, $\hat{\theta}_g$ minimizes the ridge regression problem with contexts $\phi_{U_{g,k',h}}(s_h^{k'}, a_h^{k'})$ and targets $U_{g,k',h}(s_{h+1}^{k'})$, i.e.,

$$\hat{\theta}_{g,k+1} \leftarrow \arg \min_{\theta \in \mathbb{R}^{d+1}} \lambda \|\theta\|^2 + \sum_{k'=1}^k \sum_{h=1}^H \left(\langle \theta, \phi_{U_{g,k',h}}(s_h^{k'}, a_h^{k'}) \rangle - U_{g,k',h}(s_{h+1}^{k'}) \right),$$

which has a closed-form solution given in (24).

▷ *Algorithmic notation and updates.*

Let B be an upper bound of the ℓ_2 -norm of θ^* (see Definition 7) and set as regularization parameter $\lambda \triangleq 1/(B+1)^2$. Also define the confidence radius

$$\beta_k \triangleq H\sqrt{d \log(3(1 + kH^3(B+1)^2)/\delta)} + 1. \quad (22)$$

At the first episode indexed by $k = 1$, we initialize for every goal $g \in \mathcal{G}$ and $h \in [H]$ the following quantities

$$\Sigma_{g,1,h}, \dot{\Sigma}_{g,1,h} \triangleq \lambda \mathbf{I}, \quad \mathbf{b}_{g,1,h}, \dot{\mathbf{b}}_{g,1,h} \triangleq \mathbf{0}, \quad \theta_{g,1}, \dot{\theta}_{g,1,h} \triangleq \mathbf{0}, \quad V_{g,1,H+1}(\cdot) \triangleq 0, \quad U_{g,1,H+1}(\cdot) \triangleq 0.$$

We now explain how the various estimates are updated during an episode k with goal state denoted by g_k . Over the trajectory of episode k , given the current state visited at step h denoted by s_h^k , the executed action is denoted by $a_h^k \triangleq \pi_{g_k,h}^k(s_h^k)$ and the next state is denoted by s_{h+1}^k . Then for every goal $g \in \mathcal{G}$ and for $h = 1, \dots, H$, we set

$$\begin{aligned} \Sigma_{g,k,h+1} &\triangleq \Sigma_{g,k,h} + \phi_{V_{g,k,h}}(s_h^k, a_h^k) \phi_{V_{g,k,h}}(s_h^k, a_h^k)^\top, \\ \mathbf{b}_{g,k,h+1} &\triangleq \mathbf{b}_{g,k,h} + \phi_{V_{g,k,h}}(s_h^k, a_h^k) V_{g,k,h}(s_{h+1}^k), \\ \dot{\Sigma}_{g,k,h+1} &\triangleq \dot{\Sigma}_{g,k,h} + \phi_{U_{g,k,h}}(s_h^k, a_h^k) \phi_{U_{g,k,h}}(s_h^k, a_h^k)^\top, \\ \dot{\mathbf{b}}_{g,k,h+1} &\triangleq \dot{\mathbf{b}}_{g,k,h} + \phi_{U_{g,k,h}}(s_h^k, a_h^k) U_{g,k,h}(s_{h+1}^k), \end{aligned}$$

and for every goal $g \in \mathcal{G}$, we set

$$\Sigma_{g,k+1,1} \triangleq \Sigma_{g,k,H+1}, \quad \mathbf{b}_{g,k+1,1} \triangleq \mathbf{b}_{g,k,H+1}, \quad \theta_{g,k+1} \triangleq \Sigma_{g,k+1,1}^{-1} \mathbf{b}_{g,k+1,1}, \quad (23)$$

$$\dot{\Sigma}_{g,k+1,1} \triangleq \dot{\Sigma}_{g,k,H+1}, \quad \dot{\mathbf{b}}_{g,k+1,1} \triangleq \dot{\mathbf{b}}_{g,k,H+1}, \quad \dot{\theta}_{g,k+1} \triangleq \dot{\Sigma}_{g,k+1,1}^{-1} \dot{\mathbf{b}}_{g,k+1,1}. \quad (24)$$

We proceed by recursively defining for every episode k , goal $g \in \mathcal{G}$ and $h = H, \dots, 1$,

$$Q_{g,k,h}(\cdot, \cdot) \triangleq \text{clip} \left(\mathbf{1}[\cdot \neq g] + \langle \theta_{g,k}, \phi_{V_{g,k,h+1}}(\cdot, \cdot) \rangle - \beta_k \left\| \Sigma_{g,k,1}^{-1/2} \phi_{V_{g,k,h+1}}(\cdot, \cdot) \right\|_2, 0, H \right), \quad (25)$$

$$\pi_{g,h}^k(\cdot) \triangleq \arg \min_{a \in \mathcal{A}} Q_{g,k,h}(\cdot, a), \quad (26)$$

$$V_{g,k,h}(\cdot) \triangleq \min_{a \in \mathcal{A}} Q_{g,k,h}(\cdot, a), \quad (27)$$

$$\begin{aligned} U_{g,k,h}(\cdot) &\triangleq \text{clip} \left(2\beta_k \left\| \Sigma_{g,k,1}^{-1/2} \phi_{V_{g,k,h+1}}(s, \pi_{g,h}^k(s)) \right\|_2 + \langle \phi_{U_{g,k,h+1}}(s, \pi_{g,h}^k(s)), \dot{\theta}_{g,k} \rangle \right. \\ &\quad \left. + \beta_k \left\| \dot{\Sigma}_{g,k,1}^{-1/2} \phi_{U_{g,k,h+1}}(s, \pi_{g,h}^k(s)) \right\|_2, 0, H \right). \end{aligned} \quad (28)$$

▷ *Choice of estimates \mathcal{D} , \mathcal{E} , \mathcal{Q} of ADA GOAL-UCRL-VTR (line 14 of Algorithm 1):*

$$\mathcal{Q}_{k,h}(s, a, g) \triangleq Q_{g,k,h}(s, a), \quad (29)$$

$$\mathcal{D}_k(g) \triangleq V_{g,k,1}(s_0), \quad (30)$$

$$\mathcal{E}_k(g) \triangleq U_{g,k,1}(s_0) + \frac{8\varepsilon}{9}. \quad (31)$$

D.2 Proof sketch of Theorem 11

The analysis of ADA GOAL-UCRL-VTR follows the same key steps considered in Section 4 for the analysis of ADA GOAL-UCBVI. We now sketch the ADA GOAL-UCRL-VTR equivalent of the various key steps.

First, note that similar to Lemma 16 in the tabular case, for any state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, goal state $g \in \mathcal{G}$ and vector $Y \in \mathbb{R}^{\mathcal{S}}$ such that $Y(g) = 0$, it holds that $[pY](s, a) = [p_g Y](s, a)$.

We now build the high-probability events. By using the standard self-normalized concentration inequality for vector-valued martingales of Abbasi-Yadkori et al. (2011, Theorem 2), it holds that with probability at least $1 - \delta/3$, for any $k \geq 1$ and $g \in \mathcal{G}$, θ_g^* lies in the ellipsoid

$$\mathcal{C}_{g,k} \triangleq \left\{ \theta \in \mathbb{R}^{d+1} : \left\| \Sigma_{g,k,1}^{1/2} (\theta_{g,k} - \theta) \right\|_2 \leq \beta_k \right\}.$$

The proof of the above statement follows the steps of e.g., [Zhang et al. \(2021a, Lemma A.2\)](#), the only slight difference being that we take an additional union bound over all goals $g \in \mathcal{G}$, hence the presence of G in the confidence radius (22). Furthermore, following the exact same steps as above and by definition of $\mathring{\Sigma}$ and $\mathring{\theta}$, it holds that with probability at least $1 - \delta/3$, for any $k \geq 1$ and $g \in \mathcal{G}$, θ_g^* lies in the ellipsoid

$$\mathcal{C}'_{g,k} \triangleq \left\{ \theta \in \mathbb{R}^{d+1} : \left\| \mathring{\Sigma}_{g,k,1}^{1/2} (\mathring{\theta}_{g,k} - \theta) \right\|_2 \leq \beta_k \right\}.$$

Here the confidence radius is the same as the one in $\mathcal{C}_{g,k}$ since it is chosen to be proportional to the magnitude of the $U_{g,k,h+1}(\cdot)$ function, which lies in $[0, H]$, as does the value function $V_{g,k,h+1}(\cdot)$. In what follows, we assume that the two high-probability events considered above hold, i.e., that the following event holds (it does so with probability at least $1 - 2\delta/3$)

$$\{\forall k \geq 1, \forall g \in \mathcal{G}, \theta_g^* \in \mathcal{C}_{g,k} \cap \mathcal{C}'_{g,k}\}. \quad (32)$$

▷ **Key step ①: Optimism and gap bounds.** The optimism property is standard: following e.g., [Zhang et al. \(2021a, Lemma A.1\)](#) (see also [Zhou et al., 2021, Lemma C.4](#)), it holds that $Q_{g,k,h}(s, a) \leq \overline{Q}_{g,h}^*(s, a)$ and $V_{g,k,h}(s) \leq \overline{V}_{g,h}^*(s)$ for any $(s, a, g) \in \mathcal{S} \times \mathcal{A} \times \mathcal{G}$, $h \in [H]$, $k \geq 1$.

We now depart from a usual regret minimization analysis and examine our errors U , proving that they upper bound the goal-conditioned gaps, formally defined as

$$W_{g,k,h}(s) \triangleq V_{g,h}^{\pi_g^k}(s) - V_{g,k,h}(s).$$

We now prove by induction that $W_{g,k,h}(s) \leq U_{g,k,h}(s)$. The property holds at $H+1$ since $W_{g,k,h}(s) = 0 = U_{g,k,h}(s)$. Assume that $W_{g,k,h+1}(s) \leq U_{g,k,h+1}(s)$, then we start by noticing, similar to [Zhou et al. \(2021, Equation C.10\)](#); [Zhang et al. \(2021a, Lemma A.1\)](#), that

$$W_{g,k,h}(s) \leq 2\beta_k \left\| \mathring{\Sigma}_{g,k,1}^{-1/2} \phi_{V_{g,k,h+1}}(s, \pi_{g,h}^k(s)) \right\|_2 + \underbrace{[pV_{g,h+1}^{\pi_g^k}](s, \pi_{g,h}^k(s)) - [pV_{g,k,h+1}](s, \pi_{g,h}^k(s))}_{\triangleq X}.$$

We bound X as follows

$$\begin{aligned} X &= [pW_{g,k,h+1}](s, \pi_{g,h}^k(s)) \\ &\stackrel{(i)}{\leq} [pU_{g,k,h+1}](s, \pi_{g,h}^k(s)) \\ &= \langle \phi_{U_{g,k,h+1}}(s, \pi_{g,h}^k(s)), \theta_g^* \rangle \\ &= \langle \phi_{U_{g,k,h+1}}(s, \pi_{g,h}^k(s)), \mathring{\theta}_{g,k} \rangle + \langle \phi_{U_{g,k,h+1}}(s, \pi_{g,h}^k(s)), \theta_g^* - \mathring{\theta}_{g,k} \rangle \\ &\stackrel{(ii)}{\leq} \langle \phi_{U_{g,k,h+1}}(s, \pi_{g,h}^k(s)), \mathring{\theta}_{g,k} \rangle + \left\| \mathring{\Sigma}_{g,k,1}^{1/2} (\mathring{\theta}_{g,k} - \theta_g^*) \right\|_2 \left\| \mathring{\Sigma}_{g,k,1}^{-1/2} \phi_{U_{g,k,h+1}}(s, \pi_{g,h}^k(s)) \right\|_2 \\ &\stackrel{(iii)}{\leq} \langle \phi_{U_{g,k,h+1}}(s, \pi_{g,h}^k(s)), \mathring{\theta}_{g,k} \rangle + \beta_k \left\| \mathring{\Sigma}_{g,k,1}^{-1/2} \phi_{U_{g,k,h+1}}(s, \pi_{g,h}^k(s)) \right\|_2, \end{aligned}$$

where (i) comes from the induction hypothesis and because p is a monotone operator w.r.t. the partial ordering of functions, (ii) is by Cauchy-Schwarz, (iii) holds by event (32). Finally, using that $W_{g,k,h}(s) \in [0, H]$ and by definition of $U_{g,k,h}(s)$, we conclude that $W_{g,k,h}(s) \leq U_{g,k,h}(s)$.

▷ **Key step ②: Bounding the cumulative gap bounds.** We now bound $\sum_{k=1}^K U_{g_k,k,1}(s_0)$. It holds that

$$\begin{aligned} &U_{g,k,h}(s_{k,h}) - U_{g,k,h+1}(s_{k,h+1}) \\ &\leq 2\beta_k \min \left\{ 1, \left\| \mathring{\Sigma}_{g,k,1}^{-1/2} \phi_{V_{g,k,h+1}}(s_{k,h}, \pi_{g,h}^k(s_{k,h})) \right\|_2 \right\} + \beta_k \min \left\{ 1, \left\| \mathring{\Sigma}_{g,k,1}^{-1/2} \phi_{U_{g,k,h+1}}(s_{k,h}, \pi_{g,h}^k(s_{k,h})) \right\|_2 \right\} \\ &\quad + \min \left\{ \underbrace{\langle \phi_{U_{g,k,h+1}}(s_{k,h}, \pi_{g,h}^k(s_{k,h})), \mathring{\theta}_{g,k} \rangle - U_{g_k,k,h+1}(s_{k,h+1}), H}_{\triangleq Y} \right\}, \end{aligned}$$

where

$$\begin{aligned}
 Y &\leq \left| \langle \phi_{U_{g,k,h+1}}(s_{k,h}, \pi_{g,h}^k(s_{k,h})), \theta_g^* - \hat{\theta}_{g,k} \rangle \right| + \langle \phi_{U_{g,k,h+1}}(s_{k,h}, \pi_{g,h}^k(s_{k,h})), \theta_g^* \rangle - U_{g,k,h+1}(s_{k,h+1}) \\
 &\leq \left\| \hat{\Sigma}_{g,k,1}^{-1/2} (\hat{\theta}_{g,k} - \theta_g^*) \right\|_2 \left\| \hat{\Sigma}_{g,k,1}^{-1/2} \phi_{U_{g,k,h+1}}(s_{k,h}, \pi_{g,h}^k(s_{k,h})) \right\|_2 + [p_g U_{g,k,h+1}](s_{k,h}, \pi_{g,h}^k(s_{k,h})) - U_{g,k,h+1}(s_{k,h+1}) \\
 &\leq \beta_k \left\| \hat{\Sigma}_{g,k,1}^{-1/2} \phi_{U_{g,k,h+1}}(s_{k,h}, \pi_{g,h}^k(s_{k,h})) \right\|_2 + [p_g U_{g,k,h+1}](s_{k,h}, \pi_{g,h}^k(s_{k,h})) - U_{g,k,h+1}(s_{k,h+1}).
 \end{aligned}$$

Therefore we get by telescopic sum

$$\begin{aligned}
 \sum_{k=1}^K U_{g_k,k,1}(s_0) &= \sum_{k=1}^K \sum_{h=1}^H (U_{g_k,k,h}(s_{k,h}) - U_{g_k,k,h+1}(s_{k,h+1})) \\
 &\leq 2\beta_K \underbrace{\sum_{k=1}^K \sum_{h=1}^H \min \left\{ 1, \left\| \hat{\Sigma}_{g_k,k,1}^{-1/2} \phi_{V_{g_k,k,h+1}}(s_{k,h}, a_{k,h}) \right\|_2 \right\}}_{\triangleq Z_1} \\
 &\quad + 2\beta_K \underbrace{\sum_{k=1}^K \sum_{h=1}^H \min \left\{ 1, \left\| \hat{\Sigma}_{g_k,k,1}^{-1/2} \phi_{U_{g_k,k,h+1}}(s_{k,h}, a_{k,h}) \right\|_2 \right\}}_{\triangleq Z_2} \\
 &\quad + \underbrace{\sum_{k=1}^K \sum_{h=1}^H [p_{g_k} U_{g_k,k,h+1}](s_{k,h}, a_{k,h}) - U_{g_k,k,h+1}(s_{k,h+1})}_{\triangleq Z_3}.
 \end{aligned}$$

We bound Z_1 and Z_2 using Cauchy-Schwarz and the elliptical potential lemma from linear bandits (Abbasi-Yadkori et al., 2011, Lemma 11), see e.g., Zhang et al. (2021a, Proof of Lemma A.3). This yields

$$\begin{aligned}
 Z_1 &\leq \sqrt{KH} \sqrt{\sum_{k=1}^K \sum_{h=1}^H \min \left\{ 1, \left\| \hat{\Sigma}_{g_k,k,1}^{-1/2} \phi_{V_{g_k,k,h+1}}(s_{k,h}, a_{k,h}) \right\|_2^2 \right\}} \\
 &\leq \sqrt{2} \sqrt{KH} \sqrt{\sum_{k=1}^K \sum_{h=1}^H \min \left\{ 1, \left\| \hat{\Sigma}_{g_k,k,h}^{-1/2} \phi_{V_{g_k,k,h+1}}(s_{k,h}, a_{k,h}) \right\|_2^2 \right\}} + 2Hd \log(1 + kH^3/\lambda) \\
 &\leq \sqrt{2KHd \log(1 + KH^3/(d\lambda))} + 2Hd \log(1 + kH^3/\lambda),
 \end{aligned}$$

and likewise for Z_2 . The term Z_3 can be bounded by the Azuma-Hoeffding inequality since its summands form a martingale difference sequence, thus with probability at least $1 - \delta/3$, it holds that $Z_3 \leq H\sqrt{2HK \log(3/\delta)}$.

Putting everything together and using that $\beta_K = \tilde{O}(H\sqrt{d})$, we obtain

$$\sum_{k=1}^K U_{g_k,k,1}(s_0) = \tilde{O}\left(dH^{3/2}\sqrt{K} + H^2d^{3/2}\right).$$

▷ **Key step ③: Bounding the sample complexity.** We follow the reasoning given in Section 4. By construction of the stopping rule (1), the algorithm terminates at an episode κ that verifies

$$\varepsilon \cdot (\kappa - 1) \leq \sum_{k=1}^{\kappa-1} \mathcal{E}_k(g_k) = \frac{8\varepsilon}{9} \cdot (\kappa - 1) + \tilde{O}\left(dH^{3/2}\sqrt{\kappa} + H^2d^{3/2}\right).$$

Solving this functional inequality in κ yields

$$\kappa = \tilde{O}\left(\frac{H^3d^2}{\varepsilon^2} + \frac{H^2d^{3/2}}{\varepsilon}\right).$$

Using that the sample complexity is bounded by $\kappa(H+1)$ and that $H = \tilde{O}(L)$, we conclude the proof.

▷ **Key step ④: Connecting to the original MGE objective.** The proof of this step is identical to the one of ADA GOAL-UCBVI in Appendix C.3.4.




Figure 5: Goal sampling frequency of UNIGOAL-UCBVI (*top row*), RAREGOAL-UCBVI (*middle row*) and ADAGOAL-UCBVI (*bottom row*) over 1000 episodes, split over episodes 0 – 333 (*left column*), episodes 334 – 666 (*middle column*) and episodes 667 – 999 (*right column*). Episodes are of length $H = 50$, the environment is a grid-world with $S = 52$ states, starting state $s_0 = (0, 0)$ (i.e., the top left state), $A = 5$ actions (the 4 cardinal actions plus a_{reset}). The black walls act as reflectors, i.e., if the action leads against the wall, the agent stays in the current position with probability 1. An action fails with probability $p_f = 0.1$, in which case the agent follows (uniformly) one of the other directions. The 4 states of the bottom right room can only be accessed from s_0 by any cardinal action with probability $\eta = 0.001$, thus they are extremely hard to reliably reach as their associated $V^*(s_0 \rightarrow \cdot)$ is very large (scaling with η^{-1}). We select $L = 40$ for ADAGOAL, and $\alpha = 0.1$ for RAREGOAL. For the three methods we follow the practice of M enard et al. (2021b, Section 4) and use their proposed simplified form for the exploration bonuses. The experiment is based on the rberry framework (Domingues et al., 2021a).

E ABLATION OF THE GOAL SELECTION SCHEME & PROOF OF CONCEPT EXPERIMENT

In this section, we single out the role of the adaptive goal selection scheme of ADAGOAL, i.e., step   in Algorithm 1. For simplicity we focus on the tabular case and consider that $\mathcal{G} = \mathcal{S}$. Keeping the remainder of the ADAGOAL-UCBVI algorithm fixed, we compare it to two other ad-hoc goal sampling alternatives:

- UNIGOAL: the goal state is sampled uniformly in $\mathcal{S} \setminus \{s_0\}$, i.e., with probability

$$p^{\text{UNI}}(g) \triangleq (S - 1)^{-1};$$

- RAREGOAL: the goal state is sampled proportionally to its rarity, i.e., with probability

$$p_{\alpha}^{\text{RARE}}(g) \triangleq \frac{(n_{\alpha}^k(g))^{-1}}{\sum_{s \in \mathcal{S} \setminus \{s_0\}} (n_{\alpha}^k(s))^{-1}},$$

where $n^k(s) \triangleq \sum_{t=1}^{kH} \mathbb{1}[s_t = s]$ denotes the number of times state s was visited in the first k episodes, and $n_{\alpha}^k(s) \triangleq \max\{n^k(s), \alpha\}$ for $\alpha \in (0, 1]$.

We can find equivalents of these two goal selection schemes in existing goal-conditioned deep RL methods. The case of a uniform goal sampling distribution prescribed by the environment (i.e., `UNI_GOAL`) is the most common, see e.g., Schaul et al. (2015); Andrychowicz et al. (2017). Meanwhile, the goal sampling scheme of Skew-Fit (Pong et al., 2020), a recent state-of-the-art algorithm for deep GC-RL, gives higher sampling weight to rarer goal states, where rarity is measured by a learned generative model. In the tabular case, a goal state’s rarity can be characterized by the inverse of its visitation count, which corresponds to `RARE_GOAL`.

On the one hand, it is straightforward to show that `UNI_GOAL` achieves a sample complexity of at most $\tilde{O}(L^3 S^2 A \epsilon^{-2})$. Intuitively, it pays for an extra S since it may sample goals that are too easy or too hard, in either case they are not very useful for the agent to improve its learning (and there may be in the worst case $S - 2$ of such non-informative goal states).

On the other hand, we can see that by design `RARE_GOAL` relies on the communicating assumption and may require $\text{poly}(S, A, D, \epsilon^{-1})$ samples to learn an ϵ -optimal goal-conditioned policy on \mathcal{G}_L . Here the dependence on the diameter D is somewhat problematic. Indeed, imagine there exist a set of states $\mathcal{S}_{\text{hard}}$ such that $1 \ll V^*(s_0 \rightarrow s) \leq D$ for $s \in \mathcal{S}_{\text{hard}}$ (i.e., very hard to reach states, e.g., by chance due to environment stochasticity). Then throughout the learning process, `RARE_GOAL` will strive to reach the states in $\mathcal{S}_{\text{hard}}$ and select them as goals, which leads to unsuccessful episodes and a possible waste of samples. Consequently, when goals have varying reachability (e.g., if the environment is highly stochastic), `RARE_GOAL` suffers from an issue of *goal prioritizing*, i.e., too-hard-to-reach states are given too much goal sampling importance.

Finally, we empirically complement our discussion above on the sequence of goals selected by `UNI_GOAL`, `RARE_GOAL` and `ADA_GOAL`. We design a simple two-room grid-world with a very small probability of reaching the second room, and illustrate in Figure 5 the goal sampling frequency of `UNI_GOAL-UCBVI`, `RARE_GOAL-UCBVI` and `ADA_GOAL-UCBVI`. We see that over the course of the learning interaction, as opposed to the designs of `RARE_GOAL` and `UNI_GOAL`, our `ADA_GOAL` strategy is able to successfully discard the states from the bottom right room, which have a negligible probability of being reached. In addition, `ADA_GOAL` is able to target as goals the states in the first room that are furthest away from s_0 , i.e., those at the center of the “spiral”, which effectively correspond to the fringe of what the agent can reliably reach.

F IMPLEMENTATION DETAILS OF SECTION 6

Here we provide the implementation details of our experiments reported in Section 6. Our implementation and hyperparameters of HER (Andrychowicz et al., 2017) are based on the PyTorch open-source codebase of <https://github.com/TianhongDai/hindsight-experience-replay>, which follows the official implementation of HER. As explained in Section 6, we approximate `ADA_GOAL` by computing the disagreement (i.e., standard deviation) of an ensemble of J goal-conditioned Q-functions and selecting a goal proportionally to it among N uniformly sampled goals. Then the policy is conditioned on this goal and executed for an episode of length $H = 50$. This recovers the VDS algorithm of Zhang et al. (2020b). We thus follow the implementation details given in the latter paper. In particular, the value ensemble (for goal selection) is treated as a separate module from the policy optimization (for goal-conditioned policy execution). In each training epoch, each Q-function in the ensemble performs Bellman updates with independently sampled mini-batches, and the policy is updated with DDPG (Lillicrap et al., 2015). Each Q-function in the ensemble is trained with its target network, with learning rate 1e-3, polyak coefficient 0.95, buffer size 1e6, and batch size 1000. Finally, we set $J = 3$ and $N = 1000$.