

Computational Benefits of Intermediate Rewards for Goal-Reaching Policy Learning

Yuexiang Zhai

Christina Baek

University of California, Berkeley

Department of Electrical Engineering & Computer Sciences

Berkeley, CA, 94720, USA

SIMONZHAI@BERKELEY.EDU

KBAEK@CS.CMU.EDU

Zhengyuan Zhou

New York University

Stern School of Business

New York, NY, 10012, USA

ZZHOU@STERN.NYU.EDU

Jiantao Jiao

University of California, Berkeley

Department of Electrical Engineering & Computer Sciences

Department of Statistics

Berkeley, CA, 94720, USA

JIANTAO@EECS.BERKELEY.EDU

Yi Ma

University of California, Berkeley

Department of Electrical Engineering & Computer Sciences

Berkeley, CA, 94720, USA

YIMA@EECS.BERKELEY.EDU

Abstract

Many goal-reaching reinforcement learning (RL) tasks have empirically verified that rewarding the agent on subgoals improves convergence speed and practical performance. We attempt to provide a theoretical framework to quantify the computational benefits of rewarding the completion of subgoals, in terms of the number of synchronous value iterations. In particular, we consider subgoals as one-way *intermediate states*, which can only be visited once per episode and propose two settings that consider these one-way intermediate states: the one-way single-path (OWSP) and the one-way multi-path (OWMP) settings. In both OWSP and OWMP settings, we demonstrate that adding *intermediate rewards* to subgoals is more computationally efficient than only rewarding the agent once it completes the goal of reaching a terminal state. We also reveal a trade-off between computational complexity and the pursuit of the shortest path in the OWMP setting: adding intermediate rewards significantly reduces the computational complexity of reaching the goal but the agent may not find the shortest path, whereas with sparse terminal rewards, the agent finds the shortest path at a significantly higher computational cost. We also corroborate our theoretical results with extensive experiments on the MiniGrid environments using Q-learning and some popular deep RL algorithms.

1. Introduction

Markov decision processes (MDPs) (Bertsekas, 1995; Bertsekas & Tsitsiklis, 1996; Puterman, 2014; Russell & Norvig, 2009) provide a powerful framework for reinforcement learning (RL) and planning (Bertsekas, 2019; Sutton & Barto, 2018; Szepesvári, 2010; Siciliano

et al., 2010; LaValle, 2006; Russell & Norvig, 2009) tasks. In particular, many practical tasks (Brockman et al., 2016; Vinyals et al., 2017, 2019; Berner et al., 2019; Ye et al., 2020) containing specific goal states rely heavily on reward design, especially by rewarding the agent upon its arrival on some *one-way* subgoals (each subgoal can only be visited once per episode) (Racanière et al., 2017; Popov et al., 2017; Vinyals et al., 2017, 2019; Berner et al., 2019; Ye et al., 2020). Wen et al. (2020) introduced a framework that decomposes the original MDP into “subMDPs”, and demonstrated that decomposing an MDP into subMDPs indeed leads to statistical and computational efficiency, under certain assumptions. Wen et al. (2020) provides an insightful framework for studying subgoals from a hierarchical RL perspective, but it does not clearly address how rewarding subgoals helps the goal-reaching tasks. To better understand the reward design for goal-reaching tasks, we consider the one-way subgoals as one-way intermediate states (the non-terminal states where the agent receives intermediate rewards upon its arrival). Then we show that under some *practically verifiable assumptions* on the intermediate states and simple conditions on the intermediate rewards, rewarding the agent on intermediate states is indeed more computationally efficient for learning a *successful policy* that *reaches the goal states* in terms of synchronous value iteration (SVI).

1.1 Motivating Examples

We first start with two examples (the Maze and the Pacman Game in Figure 1) to reveal the existence of one-way intermediate states in different goal-reaching tasks. For the Maze problem, suppose one episode ends when the agent reaches the goal, then the Maze problem does not have one-way intermediate states, since all previously visited states can be revisited in one episode. In contrast, the Pacman game intrinsically possesses one-way intermediate states, because the Pacman cannot revisit the states where the previously consumed food pellets are available.

We conduct a toy experiment on the Pacman game to demonstrate how the design of *intermediate states* and *intermediate rewards* affects the behavior of an agent under the greedy policy (Sutton & Barto, 2018). In the Pacman game, the agent (Pacman) wins when it consumes all the food before being caught by the ghost. Intuitively, if we want the Pacman to win the game, we need to design the MDP such that the agent receives positive rewards for winning, consuming food, surviving (not being caught by the ghost), and negative rewards for losing (being caught by the ghost). With the aforementioned intuition, we design several different reward functions for the Pacman game. Comparing reward settings (1), (2), and (3) in Table 1, the Pacman performs better when the MDP contains positive intermediate rewards for consuming food (visiting one-way intermediate states). This observation matches the common belief that integrating prior knowledge of the task into the reward design is generally helpful. However, designing rewards based on prior knowledge could sometimes negatively affect the performance. In setting (4), where the Pacman receives a positive reward for survival (visiting non one-way intermediate states), the Pacman focuses more on dodging the ghost rather than consuming food, and eventually gets caught by the ghost. At this point, a natural question is:

How does the design of intermediate rewards affect the behavior of a greedy policy for goal-reaching tasks?

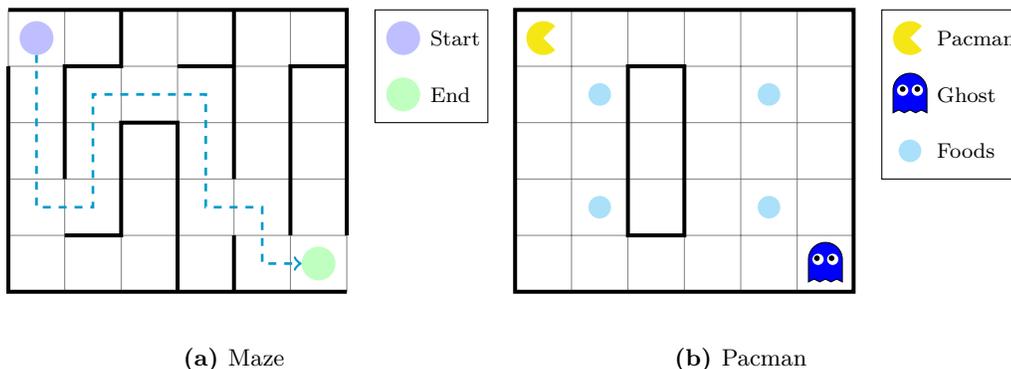


Figure 1: Examples of two practical tasks. The thick curves represent walls. The agent can take actions from {left, right, up, down} and move to the corresponding adjacent locations. The agent stays at the original location if it hits a wall after taking an action. (a): In the maze problem, each state in the MDP is the agent’s position. The goal is to find the end from the start. The shortest path is marked with blue dashed curves. (b): In the Pacman game, each state contains the position of the Pacman, the ghost, and all remaining food pellets. The agent (Pacman) needs to consume all the pellets while avoiding the ghost, and each food pellet can only be eaten *once*. The ghost moves simultaneously with the Pacman. The Pacman wins once all pellets are consumed and loses if it is caught by the ghost. We use the TikZ code from <https://gist.github.com/neic/9546556> to draw the ghost.

	r_f	r_s	0 ep	100 eps	500 eps	1000 eps	1500 eps	2000 eps	2500 eps
(1)	0	0	0.0%	0.1%	0.1%	0.6%	89.5%	93.7%	96.0%
(2)	1	0	0.0%	0.5%	33.0%	78.0%	92.6%	97.8%	98.4%
(3)	10	0	0.0%	0.1%	57.8%	93.6%	95.7%	97.0%	97.9%
(4)	1	1	0.0%	0.0%	1.9%	2.5%	2.2%	7.7%	12.9%

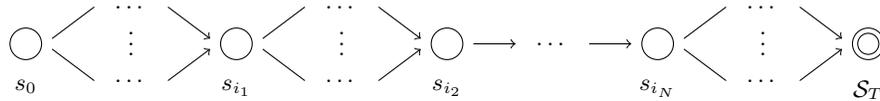
Table 1: The win rate (averaged among 1000 trials) of the Pacman game shown in Figure 1b after {0, 100, 500, 1000, 1500, 2000, 2500} Q-learning training episodes under a greedy policy. We compare the performance under 4 different reward function designs. In all settings, the winning reward is 10 and the losing reward is -10. r_f is the reward of consuming a food pellet. The Pacman receives a survival reward r_s if it is not caught by the ghost after taking an action. Our implementation of the Pacman is based on the code from Berkeley CS188 (<https://inst.eecs.berkeley.edu/~cs188/fa18/project3.html>). Details of the experiments are provided in Appendix C.1.

1.2 Our Formulation and Main Results

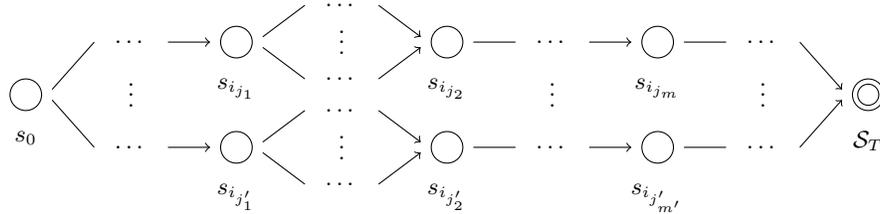
1.2.1 ONE-WAY INTERMEDIATE STATES

To study the conditions under which goal-reaching problems with one-way intermediate states enjoy computational benefits, we consider two one-way intermediate state settings as shown in Figure 2: 1) one-way single-path (OWSP), and 2) one-way multi-path (OWMP). In both settings, we assume the existence of one-way intermediate states (formally defined in Assumption 3.2) that behave like “one-way” checkpoints that *cannot be revisited* in one episode and intermediate rewards are assigned to the arrival at such states. Many practical

RL tasks implicitly adopt this one-way property, as practitioners often identify subgoals and assign a one-time reward for their completion in solving challenging tasks in addition to the terminal rewards which occur once the agent reaches the goal states (Brockman et al., 2016; Vinyals et al., 2017, 2019; Berner et al., 2019; Ye et al., 2020). For the case where intermediate states are non one-way, we provide an example where the agent gets stuck at an intermediate state permanently instead of pursuing the goal, if the intermediate rewards are not “properly” designed. As we have observed in setting (4) of the Pacman game in Table 1, improperly designed intermediate rewards can negatively affect the agent’s ability to find a successful policy that reaches the goal.



(a) In this case, all intermediate states $\mathcal{S}_I = \{s_{i_1}, s_{i_2}, \dots, s_{i_N}\}$ have to be visited in a certain order.



(b) In this case, only a subset of intermediate states $\mathcal{S}_J = \{s_{i_{j_1}}, s_{i_{j_2}}, \dots, s_{i_{j_m}}\}$ (or $\mathcal{S}_{J'} = \{s_{i_{j'_1}}, s_{i_{j'_2}}, \dots, s_{i_{j'_{m'}}}\}$) with at least n states ($m, m' \geq n$) have to be visited in a certain order.

Figure 2: The transition graph of different settings in Assumption 3.3

1.2.2 MAIN RESULTS

We focus on comparing the two following reward settings of *deterministic MDPs*:

- **The sparse reward setting:** the agent only receives terminal rewards B while reaching the goal (terminal states);
- **The intermediate reward settings:** the agent receives *equal magnitude* intermediate rewards B_I upon the arrival of intermediate states, for both OSWP and OWMP intermediate state settings.

Note that the intermediate rewards B_I could be different from the terminal rewards B , and we formally define the reward functions in Table 3 in Section 3.3.

The OWSP Intermediate Reward Setting Intuitively, the OWSP setting, which rewards the one-way intermediate states in addition to terminal rewards, is more computationally efficient than only having terminal rewards alone, because the OWSP setting reduces

the problem from *finding the terminal states* to *finding the closest intermediate states*. For example, considering the environment provided in Figure 3, the number of SVI required to obtain a greedy successful policy using sparse terminal rewards is 8 (the *distance* or minimum required steps from the initial state s_0 to the terminal state \mathcal{S}_T), whereas the OWSP intermediate reward setting (Figure 4) reduces such computational complexity to 3 (the maximum distance between two intermediate states). Intuitively, Figure 4 suggests that the computational complexity of finding a successful policy can be reduced by rewarding the one-way intermediate states. The computational complexity for obtaining a successful policy (in terms of SVI) with sparse terminal rewards and the OWSP intermediate reward settings are formalized in Proposition 4.1 and 4.2, respectively. Moreover, Proposition 4.1 and 4.2 show that a greedy policy follows the shortest path to \mathcal{S}_T in both sparse terminal rewards and OWSP setting.

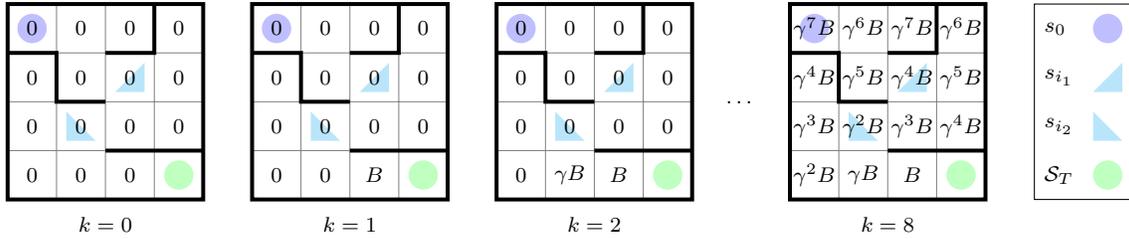


Figure 3: An example of the evolution of a zero-initialized value function via SVI in the one-way single-path (Assumption 3.3 (a)) in sparse reward setting (Table 3). Each block in the 4×4 grid represents a state, the green state is the terminal state, the blue state is the initial state, and the thick curves represent walls that the agent cannot pass through. The cyan triangles are the one-way intermediate states. The orientation of the apex of a given cyan triangle represents the direction of each intermediate state. Namely, s_{i_1} can only be visited from the left, and s_{i_2} can only be visited from the right, and neither s_{i_1} nor s_{i_2} can be revisited. The agent can choose any action from {left, right, up, down}, and the agent will stay in the same state if it takes an action that hits the wall. Note that the value function of the terminal state remains 0 because the MDP stops once the agent reaches \mathcal{S}_T , so for any terminal state $s_t \in \mathcal{S}_T$, $V_k(s_t)$ is never updated.

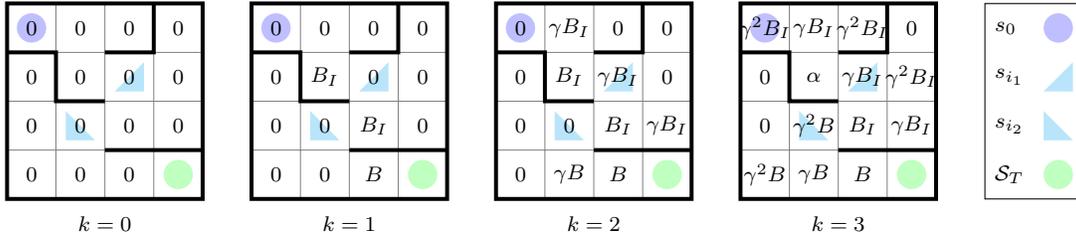


Figure 4: An example of the evolution of a zero-initialized value function via SVI in the one-way single-path (Assumption 3.3 (a)) intermediate reward setting (Table 3). $\alpha = B_I + \gamma^2 B_I$, and the remaining settings are the same as Figure 3.

The OWMP Intermediate Reward Setting Similar to the OWSP setting, the OWMP is also more computationally efficient than using sparse terminal rewards (in terms of SVI), provided the conditions allow a greedy policy to recursively find the closest intermediate

state and eventually reach the terminal state. Theorem 4.3 provides the conditions that enable the agent to find the closest intermediate state when the terminal states are not *directly reachable* (formally defined in Definition 3.4), and when the terminal states are directly reachable, Theorem 4.4 characterizes such conditions and the associated computational complexity of finding a successful policy. Nevertheless, unlike the OWSP setting or the sparse terminal reward setting, such a policy does not necessarily find the shortest path from s_0 to \mathcal{S}_T . Our result in the OWMP setting illustrates a trade-off between the computational complexity and the pursuit of shortest path – adding intermediate rewards on one-way intermediate states generally reduces computational complexity, but it does not necessarily find the shortest path.

Our Contributions In this work, we make the following contributions:

1. We propose a framework that studies the computational complexity for goal-reaching tasks in terms of synchronous value iteration (SVI);
2. We propose two one-way intermediate rewards settings: the one-way single-path (OWSP) and the one-way multi-path intermediate reward settings (OWMP) and connect both settings to practical tasks;
3. In both OWSP and OWMP settings, we demonstrate that assigning intermediate rewards on one-way intermediate states (in addition to terminal rewards) is more computationally efficient than the sparse reward setting where the agent only receives terminal rewards at the terminal states.
4. We reveal a trade-off between the computational complexity and the pursuit of the shortest path in the OWMP intermediate settings versus the sparse reward setting. That is, comparing to the sparse reward setting where the agent only receives terminal rewards and follows the shortest path to the goal, rewarding the agent at one-way intermediate states (in addition to terminal states) generally reduces computational complexity but the agent does not necessarily follow the shortest path.
5. We conduct extensive experiments in both OWSP and OWMP settings to demonstrate the computational benefits of rewarding one-way intermediate states and OWMP’s trade-off between computational complexity and the shortest path. In particular, even though our theoretical analysis adopts SVI, the computational gains and the trade-off appear in both traditional ε -greedy Q-learning and popular deep RL methods such as DQN (Mnih et al., 2015), A2C (Mnih et al., 2016), and PPO (Schulman et al., 2017).

1.3 Organization of this Paper

The rest of this paper is organized as follows. In Section 2, we introduce the basics of MDPs and the definition of successful policies. We then provide the formal definition of intermediate states, the OWSP and OWMP intermediate states in Section 3.1. We then discuss how these assumptions are well reflected in RL tasks observed in practice (Section 3.2). In Section 4, we separately provide the computational complexity of finding a successful trajectory for both the OWSP and OWMP settings. In Section 5, we discuss the connection between the trajectory generated by a greedy policy and the shortest path for each setting.

We then provide experimental results in Section 6 to corroborate our theoretical results in Section 4. In Section 7, we discuss the connection between our work and some prior works, the implication of our results, and potential future directions. The proof and experimental details are deferred to the appendix.

2. Background and Problem Formulation

2.1 Assumptions of the MDP

Basics of MDPs We use a quintuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ to denote an MDP, where \mathcal{S} is a finite state space, \mathcal{A} is a finite action space, $\gamma \in (0, 1)$ is the discount factor, $r : (\mathcal{S} \times \mathcal{A}) \times \mathcal{S} \mapsto [0, \infty)$ is the reward function of each state-action pair (s, a) and its subsequent state s_a , and P is the probability transition kernel of the MDP, where $P(s_a|s, a)$ denotes the probability of the subsequent state s_a of a state-action pair (s, a) . In particular, we focus on *deterministic MDPs*: the transition kernel satisfies $\forall (s, a) \in \mathcal{S} \times \mathcal{A}, \exists s_a \in \mathcal{S}$, such that $P(s_a|s, a) = 1$ and reward function $r(s, a, s_a)$ only depends on the subsequent state s_a . We say a state s' is *reachable* from s if there exists a *path* or a sequence of actions $\{a_1, a_2, \dots, a_n\}$ that takes the agent from s to s' .¹ Note that the definition of reachability *does not* necessarily imply s is reachable from s , since there may not exist a path from s to itself. Our model also considers a fixed *initial state* $s_0 \in \mathcal{S}$ and *terminal states* $\mathcal{S}_T \subset \mathcal{S}$. Without further explanation, we assume the terminal state \mathcal{S}_T is reachable from all $s \in \mathcal{S}$. All in all, MDP begins at state s_0 and stops once the agent reaches any state $s \in \mathcal{S}_T$.

Definition 2.1 (Distance between States) *Given an MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ with initial state $s_0 \in \mathcal{S}$, $\forall s \in \mathcal{S} \setminus \mathcal{S}_T, s' \in \mathcal{S}$, we define the distance $D(s, s') : \mathcal{S} \times \mathcal{S} \mapsto \mathbb{N}$ as the minimum number of required steps (≥ 1) from s to s' . We slightly abuse the notation of $D(\cdot, \cdot)$ by writing $D(s, \mathcal{S}_T) \doteq \min_{s' \in \mathcal{S}_T} D(s, s')$ as the minimum distance from s to \mathcal{S}_T .*

2.2 Successful Q-Functions

In the preceding deterministic MDP formulation, we aim at solving a goal-reaching RL problem (Kaelbling, 1993b; Sutton et al., 2011; Andrychowicz et al., 2017; Andreas et al., 2017; Pong et al., 2018; Ghosh et al., 2019; Eysenbach et al., 2020a, 2020b; Kadian et al., 2020; Fujita et al., 2020; Chebotar et al., 2021; Khazatsky et al., 2021) or a planning problem (Bertsekas & Tsitsiklis, 1996; Boutilier et al., 1999; Sutton et al., 1999; Boutilier et al., 2000; Rintanen & Hoffmann, 2001; LaValle, 2006; Russell & Norvig, 2009; Nasiriany et al., 2019). We say a Q-function is *successful* if its associated *greedy policy* (Sutton & Barto, 2018) leads the agent to the terminal states \mathcal{S}_T from the initial state s_0 .

Definition 2.2 (Successful Q-functions) *Given a deterministic MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$, with initial state s_0 . We say $Q(\cdot, \cdot)$ is a successful Q-function of \mathcal{M} if the greedy policy with respect to Q generates a path $\{s_0, a_0, s_1, a_1, \dots, s_H\}$ such that $\forall i = 0, 1, \dots, H - 1, s_i \notin \mathcal{S}_T$ and $s_H \in \mathcal{S}_T$.*

1. A similar definition of reachability appears in (Forejt et al., 2011).

All theoretical results of this work are based on *synchronous value iteration*:

$$\begin{aligned} V_{k+1}(s) &\doteq \max_{a \in \mathcal{A}} \{Q_k(s, a)\}, \quad \forall s \in \mathcal{S}, \\ Q_k(s, a) &= r(s, a, s_a) + \gamma V_k(s_a), \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}. \end{aligned} \tag{2.1}$$

The convergence of *synchronous* value iteration has been well studied (Puterman, 2014; Sutton & Barto, 2018; Bertsekas, 2019), hence we use it as a starting point for studying the effect of intermediate states and intermediate rewards to be introduced later in Section 3.

3. Intermediate States and Intermediate Rewards

3.1 Intermediate States

We state the formal definition of intermediate states in Definition 3.1, and provide two follow-up assumptions (Assumption 3.2, 3.3) regarding the intermediate states.

Definition 3.1 (Intermediate States) *Given an MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ with initial state $s_0 \in \mathcal{S}$ and terminal states \mathcal{S}_T , we define intermediate states $\mathcal{S}_I = \{s_{i_1}, s_{i_2}, \dots, s_{i_N}\} \subset \mathcal{S} \setminus \mathcal{S}_T$ as the states that satisfy*

$$r(s, a, s_a) > 0, \quad \forall s_a \in \mathcal{S}_I, \tag{3.1}$$

where s_a is the subsequent state of the state-action pair (s, a) .

Assumption 3.2 (One-Way Intermediate States) *Given an MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ with initial state $s_0 \in \mathcal{S}$, terminal states \mathcal{S}_T , and intermediate states \mathcal{S}_I , we assume that each intermediate state $s_{i_j} \in \mathcal{S}_I$ can only be visited at most once in one episode under any policy, namely, $\forall j \in [N], D(s_{i_j}, s_{i_j}) = \infty$.*

Intuitively, Assumption 3.2 characterizes the states that “cannot be revisited” in one episode, upon the agent’s arrival. For example, in the Pacman game (Figure 1b), if the Pacman reaches a location (x, y) that contains a food pellet, then the Pacman cannot go back to previous states where the consumed food pellet is still available at (x, y) . Similarly for the door & key environment (Chevalier-Boisvert et al., 2018) that will be presented in Section 6, once the agent picks up a key at location (x, y) , it possesses the key for the rest of that episode. Assumption 3.2 is widely adopted in practice (Brockman et al., 2016; Vinyals et al., 2017, 2019; Berner et al., 2019; Ye et al., 2020), as many subgoals identified by designers are usually one-way (see more discussion in Section 3.2).

Assumption 3.3 (Different Settings of Intermediate States) *Given an MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ with initial state $s_0 \in \mathcal{S}$, intermediate states set $\mathcal{S}_I = \{s_{i_1}, s_{i_2}, \dots, s_{i_N}\}$ and terminal states \mathcal{S}_T satisfying Assumption 3.2, we study these different settings of \mathcal{S}_I :*

- (a) *Any path from s_0 to \mathcal{S}_T has to visit all states in \mathcal{S}_I in a certain order (i.e., in the order of $s_{i_1}, s_{i_2}, \dots, s_{i_N}$).*
- (b) *Any path from s_0 to \mathcal{S}_T has to visit at least n intermediate states $\{s_{i_j} | s_{i_j} \in \mathcal{S}_I, j \in \mathcal{J} \subset [N], |\mathcal{J}| \geq n\}$ in a certain order.*

Assumptions 3.2 and 3.3 characterize intermediate states as different one-way paths consisting of subgoals that the agent has to complete in a certain order. If we consider all paths from s_0 to \mathcal{S}_T that pass through the same intermediate states as an equivalence class, the goal-reaching problem using setting (a) is a *single-path problem*, while setting (b) is a *multi-path problem*. Additionally, Assumption 3.3 indicates that from a given state s , not every intermediate state $s_i \in \mathcal{S}_I$ can be directly visited (without first visiting any other intermediate states). Thus, it is worth considering the *directly reachable* intermediate (and terminal) states from a given state s and the *minimum distance* between two intermediate states.

Definition 3.4 (Direct Reachability) *Given an MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ with terminal states \mathcal{S}_T and intermediate states $\mathcal{S}_I = \{s_{i_1}, s_{i_2}, \dots, s_{i_N}\}$ satisfying Assumption 3.2, $\forall s \in \mathcal{S} \setminus \mathcal{S}_T$, let $\mathcal{I}_d(s) \subset [N]$ denote the indices of directly reachable intermediate states of s . That is, $\forall j \in \mathcal{I}_d(s)$, there exists a path from s to the intermediate state s_{i_j} in the transition graph of \mathcal{M} that does not visit any other intermediate state.*

Similar to the intermediate states \mathcal{S}_I , we say the terminal states \mathcal{S}_T are *directly reachable* from a state s if there exists a path from s to \mathcal{S}_T that does not contain any intermediate states.

Assumption 3.5 (Minimum Distance) *Given an MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ with terminal states \mathcal{S}_T and intermediate states $\mathcal{S}_I = \{s_{i_1}, s_{i_2}, \dots, s_{i_N}\}$ satisfying Assumption 3.2, $\forall s \in \mathcal{S} \setminus \mathcal{S}_T$, we assume the distance between any two intermediate states is at least $h \in \mathbb{N}^+$, and the distance between any intermediate state and \mathcal{S}_T is also at least h . Namely, $\forall s_{i_j} \in \mathcal{S}_I$, we have*

$$\min_{j' \in \mathcal{I}_d(s_{i_j})} D(s_{i_j}, s_{i_{j'}}) \geq h, \quad \text{and} \quad \min_{j \in [N]} D(s_{i_j}, \mathcal{S}_T) \geq h, \quad (3.2)$$

where $\mathcal{I}_d(s_{i_j})$ is the set of indices of directly reachable intermediate states from s_{i_j} (see Definition 3.4).

In practice, although the minimum distance h between two intermediate states (and an intermediate state to \mathcal{S}_T) is task dependent, it is generally fair to assume that h satisfies $1 < h < D(s_0, \mathcal{S}_T)$.

The aforementioned assumptions (Assumption 3.2, 3.3, and 3.5) are connected to our main results in the following respects:

1. The one-way Assumption 3.2 and 3.3 provide the key theoretical framework for our main result. The key insight of the computational gain in Section 4 is that, if the MDP contains some one-way intermediate states, then one can reduce the computational complexity of finding the terminal states \mathcal{S}_T to the complexity of *finding the closest intermediate states*, which eventually leads to \mathcal{S}_T . The detailed dependence of Assumption 3.2 and 3.3 on our main result is provided in Table 4.
2. We use Assumption 3.5 for quantitatively analyzing the relative scale of the terminal rewards versus the intermediate rewards that *leads the agent to the closet intermediate states* and eventually to \mathcal{S}_T .

Assumption 3.2 and 3.3 are highly task dependent because they characterize the connections of the one-way intermediate states of different tasks. In the next subsection, we will introduce some practical tasks that adopt Assumption 3.2 and 3.3.

3.2 The Interpretation of Intermediate States

We list several practical tasks that adopt Assumption 3.3 in Table 2 to interpret Assumption 3.3. As illustrated in Table 2, the multi-path intermediate state setting (Assumption 3.3 (b)) is more common in practical tasks than the single-path setting (Assumption 3.3 (a)), because the single-path setting is essentially a special case of the multi-path setting.

Assumption 3.3 does not hold for the Maze (See Figure 1a) because the agent can repeatedly visit any state in the maze. As for Go, Assumption 3.3 also does not hold because the existence of subgoals in Go remains ambiguous. Except for the Maze problem and Go, Assumption 3.3 (b) fits the others practical tasks in Table 2 naturally. Moreover, the stronger Assumption 3.3 (a) holds for Montezuma because it

Assumption 3.3	(a)	(b)
Maze (Figure 1a)	✗	✗
Pacman (Figure 1b)	✗	✓
Montezuma ((Brockman et al., 2016))	✓	✓
Go ((Silver et al., 2016, 2017))	✗	✗
Dota2 ((Berner et al., 2019))	✗	✓
StarCraft II ((Vinyals et al., 2017, 2019))	✗	✓
Honor of Kings ((Ye et al., 2020))	✗	✓

Table 2: Examples of some RL applications that adopt different settings in Assumption 3.3.

requires the agent to visit specific states in a certain order, e.g., the agent will need to pick up a key and unlock a door to proceed to the next chapter. The Pacman game satisfies Assumption 3.3 (b) – considering each state in the Pacman game consists of the location of the Pacman, the ghost, and the remaining foods, then the states where “the Pacman consumes a food pellet (the location of the Pacman first coincides with an available food pellet)” can be viewed as a one-way intermediate state. Hence each episode of the Pacman game contains n (the total number of food) intermediate states, and these states appear in the order in which the number of available food is decreasing. As for Dota2, StarCraft II, and Honor of Kings, their winning conditions require the agent to “destroy” the enemy’s base. However, the enemy’s base is not assailable before the agent completes several sub-tasks. For example, in Dota2 or Honor of Kings, 3 towers block each of the roads to the enemy’s base. The agent must first sequentially destroy the 3 towers to reach the enemy’s base.² If one views the state where the agent destroys a tower in Dota2 or Honor of Kings as an intermediate state, Assumption 3.3 (b) is naturally satisfied, because the agent need to visit at least 3 intermediate states to attack the enemy’s base.

Note that even though some of the practical tasks (Pacman, Montezuma, Dota2, StarCraft II, and Honor of Kings) mentioned in Table 2 satisfy Assumption 3.3, our theoretical results to be introduced in Section 4 may not be directly applied to these tasks; this is because they contain other factors that cannot be characterized by our current theoretical assumptions (e.g., the deterministic MDP assumption mentioned in Section 2.1, the equal magnitude assumption to be mentioned in Section 3.3).

2. Information on the rules and gameplay of Dota are available online; i.e., <https://purgegamers.true.io/g/dota-2-guide/>. The rule of Honor of Kings is similar to Dota2.

3.3 Rewards

We assume the reward function $r(s, a, s_a)$ of a state-action pair (s, a) only depends on the subsequent state s_a . In particular, our theoretical results (Section 4) and experiments (Section 6) focus on comparing these two reward settings: the *sparse reward setting* where *there are no intermediate states in the environment* hence the agent only receives positive rewards at \mathcal{S}_T , and the *intermediate reward setting*, meaning that *the environment contains intermediate states* so the agent receives rewards at both \mathcal{S}_T and \mathcal{S}_I . Formally, we write the reward functions of the sparse reward setting and the intermediate reward setting as follows:

	Sparse Rewards		Intermediate Rewards		
$r(s, a, s_a)$	B	0	B	B_I	0
Condition	$s_a \in \mathcal{S}_T$	$s_a \in \mathcal{S} \setminus \mathcal{S}_T$	$s_a \in \mathcal{S}_T$	$s_a \in \mathcal{S}_I$	$s_a \in \mathcal{S} \setminus (\mathcal{S}_T \cup \mathcal{S}_I)$

Table 3: The sparse reward setting and the intermediate reward setting.

As shown in Table 3, we assume the agent receives terminal rewards B once it reaches $s_t \in \mathcal{S}_T$, and receives intermediate rewards B_I once it reaches an intermediate state $s_i \in \mathcal{S}_I$. Note that we are *not claiming* that the applications provided in Table 2 satisfies the intermediate reward setting in Table 3 where all intermediate rewards have the same magnitude. In fact, the reward setting in Table 3 only matches the Pacman game in Table 2, where we assign all food with equal intermediate rewards. Still, our theoretical results in Section 4 relying on the equal magnitude of intermediate rewards can be further generalized to the case where intermediate rewards are different. We defer the discussion of such generalization to unequal magnitude of intermediate rewards in Section 7 and leave the formal studies of the generalization for future.

4. Main Results

We study the conditions for both the sparse reward and intermediate reward setting, under which the Q-function $Q_k(\cdot, \cdot)$ is a successful Q-function (Definition (2.2)), where $Q_k(\cdot, \cdot)$ is the Q-function after k update of SVI (2.1) from zero initialization:

$$Q_0(s, a) = 0, V_0(s) = 0, \forall (s, a) \in \mathcal{S} \times \mathcal{A}. \tag{4.1}$$

And if $Q_k(\cdot, \cdot)$ is indeed a successful Q-function, we further discuss the computational complexity k (the minimum number of SVI) of obtaining a successful $Q_k(\cdot, \cdot)$.

For the sparse reward setting, we first show that for a large enough k , $Q_k(\cdot, \cdot)$ is a successful Q-function and provide its computational complexity in Section 4.1. Next, we discuss the following MDP settings as shown in Table 3.

	Assumption 3.2	Assumption 3.3
The One-Way Single-Path (OWSP) setting	✓	(a)
The One-Way Multi-Path (OWMP) setting	✓	(b)
The Non One-Way (NOW) setting	✗	✗

Table 4: Requirements of different assumptions for the main result.

For the OWSP setting, we show that $Q_k(\cdot, \cdot)$ is a successful Q-function for sufficiently large k and provide its computational complexity in Section 4.2. For the OWMP setting, we introduce a sufficient condition under which $Q_k(\cdot, \cdot)$ is successful for a large enough k , and the corresponding computational complexity in Section 4.3. As for the Non One-Way (NOW) setting, we provide an example where $Q_k(\cdot, \cdot)$ is not successful for any k in Section 4.4. The required assumptions for the OWSP, OWMP, and NOW are provided in Table 4.

4.1 Sparse Rewards

Given zero-initialized value functions (4.1), Figure 3 demonstrates the evolution of $V_k(s)$ as the number of SVI (k) increases. A direct implication of Figure 3 is that, at iteration k , $\forall s \in \mathcal{S} \setminus \mathcal{S}_T$, given $d = D(s, \mathcal{S}_T)$, the value function $V_k(\cdot)$ satisfies:

$$V_k(s) = \gamma^{d-1} B \cdot \mathbb{1}\{d \leq k\}. \tag{4.2}$$

The derivation of (4.2) is provided in Lemma B.1. With (4.2), we know that when $k \geq D(s_0, \mathcal{S}_T)$, the value function at the initial state $V_k(s_0)$ would be positive. Hence, a greedy policy that recursively finds the next state with the largest value from s_0 , will generate a path which eventually reaches \mathcal{S}_T , whenever $k \geq D(s_0, \mathcal{S}_T)$. More precisely, we have:

Proposition 4.1 (Sparse Rewards) *Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ be a deterministic MDP with initial state s_0 and terminal states \mathcal{S}_T . If the reward function $r(\cdot)$ follows the sparse reward setting (Table 3) and the value function and Q-function are zero-initialized (4.1), then after any $k \geq D(s_0, \mathcal{S}_T)$ synchronous value iteration updates (2.1), the Q-function Q_k is a successful Q-function, and a greedy policy follows the shortest path from s_0 to \mathcal{S}_T .*

We provide a sketch proof here and leave the details in Appendix A.2. With (4.2) from Lemma B.1, we can write the value function in this setting as $V_k(s) = \gamma^{d-1} B \cdot \mathbb{1}\{d \leq k\}$, $d = D(s, \mathcal{S}_T)$. Hence a greedy policy taking the agent to the subsequent state with maximum value function will lead the agent one step closer to \mathcal{S}_T . Recursively applying the same argument, we conclude that after $k \geq D(s_0, \mathcal{S}_T)$ SVI, an agent following the greedy policy finds \mathcal{S}_T from s_0 .

Note that our result in the sparse reward setting (Proposition 4.1) might seem similar to the classic convergence result for SVI via dynamical programming (Bertsekas, 1995; Sutton & Barto, 2018). However, our setting only focuses on the conditions that allow a greedy to find \mathcal{S}_T from the initial state s_0 , rather than all state $s \in \mathcal{S}$. Therefore, the computational complexity in the sparse reward setting relies on the distance $D(s_0, \mathcal{S}_T)$. Furthermore, in the OWSP and OWMP settings to be introduced later, our proofs contain share the similar method as the sparse reward setting, hence Proposition 4.1 serves as a good preliminary result for the later parts.

4.2 One-Way Single-Path Intermediate Rewards

Similar to the sparse reward setting, we first illustrate the evolution of the one-way single-path (OWSP) intermediate rewards in Figure 4. As shown in Figure 4, after k synchronous iterations, the value function at each state s equals to the sum of discounted rewards from all future intermediate states and the terminal states. The value function $V_k(s)$ of state s at iteration k is provided in Lemma B.2. In this case, when k satisfies

$$k \geq \max\{D(s_0, s_{i_1}), D(s_{i_1}, s_{i_2}), \dots, D(s_{i_{N-1}}, s_{i_N}), D(s_{i_N}, \mathcal{S}_T)\},$$

$Q_k(\cdot, \cdot)$ is a successful Q-function and a greedy policy will recursively find the next intermediate state, eventually reaching \mathcal{S}_T . More precisely, we have:

Proposition 4.2 (Single-path Intermediate States) *Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ be a deterministic MDP with initial state s_0 , intermediate states $\mathcal{S}_I = \{s_{i_1}, s_{i_2}, \dots, s_{i_N}\}$, and terminal states \mathcal{S}_T . Suppose \mathcal{M} satisfies Assumption 3.2 and 3.3 (a). If the reward function $r(\cdot)$ follows the intermediate reward setting (Table 3) and the value function and Q-function are zero-initialized (4.1), then after*

$$k \geq d_{\max} \doteq \max\{D(s_0, s_{i_1}), D(s_{i_1}, s_{i_2}), \dots, D(s_{i_{N-1}}, s_{i_N}), D(s_{i_N}, \mathcal{S}_T)\} \quad (4.3)$$

synchronous value iteration updates (2.1), the Q-function Q_k is a successful Q-function, and a greedy policy follows the shortest path from s_0 to \mathcal{S}_T .

We defer the details to Appendix A.3. To outline, we provide an explicit formulation of $V_k(s)$ for this setting in Lemma B.2, and show that when $k \geq d_{\max}$, a greedy agent will move one step forward to the next closest intermediate state and eventually reach \mathcal{S}_T .

Comparing to Proposition 4.1, we know that, the computational complexity (in terms of obtaining a successful Q-function) of an MDP \mathcal{M} with the sparse reward setting can be reduced from $D(s_0, \mathcal{S}_T)$ to

$$\max\{D(s_0, s_{i_1}), D(s_{i_1}, s_{i_2}), \dots, D(s_{i_{N-1}}, s_{i_N}), D(s_{i_N}, \mathcal{S}_T)\},$$

if \mathcal{M} adopts the OWSP intermediate reward setting.

4.3 One-Way Multi-Path Intermediate Rewards

As shown in Table 2, comparing to the OWSP setting, the OWMP setting applies to more practical tasks and can be considered as a special case of the OWSP setting. Therefore, Proposition 4.2 for the OWSP setting cannot be directly generalized to the OWMP case, since moving towards one intermediate state does not necessarily lead to \mathcal{S}_T in the OWMP case.

Still, the intuition for the computational complexity of obtaining a successful Q-function in the OWMP setting is similar to the OWSP setting. For a state s , under some conditions regarding rewards B , B_I , and intermediate states \mathcal{S}_I , a greedy agent finds the closest directly reachable intermediate state $s_{i_{j_1}}$ of s (Definition 3.4) after a sufficient number SVI. From any intermediate $s_{i_{j_m}} \in \mathcal{S}_I$, the greedy agent recursively finds the closest directly reachable intermediate state $s_{i_{j_{m+1}}}$ of $s_{i_{j_m}}$ and eventually reaches \mathcal{S}_T , since intermediate

states cannot be revisited and the number of total intermediate states is finite. The next theorem characterizes the sufficient conditions for an agent starting from $s \in \mathcal{S} \setminus \mathcal{S}_T$ following a greedy policy to find the closest directly reachable intermediate of s .

Theorem 4.3 (Finding the Closest \mathcal{S}_I) *Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ be a deterministic MDP with initial state s_0 , intermediate states $\mathcal{S}_I = \{s_{i_1}, s_{i_2}, \dots, s_{i_N}\}$, and terminal states \mathcal{S}_T . Suppose \mathcal{M} satisfies Assumption 3.2, 3.3 (b), and 3.5, if the reward function $r(\cdot)$ follows the intermediate reward setting (Table 3) and the value function and Q -function are zero-initialized (4.1), then $\forall s \in \mathcal{S} \setminus \mathcal{S}_T$ that cannot directly reach \mathcal{S}_T , after*

$$k \geq d = D(s, \mathcal{S}_I) \doteq \min_{j \in \mathcal{I}_d(s)} D(s, s_{i_j}) \quad (4.4)$$

synchronous value iteration updates (2.1), an agent following the greedy policy will find the closest directly reachable intermediate state $s_{i_{j^}}$ of s ($D(s, s_{i_{j^*}}) = D(s, \mathcal{S}_I)$), given $\frac{B}{B_I} \in \left(0, \frac{1}{1-\gamma^h}\right)$ if $\gamma + \gamma^h \leq 1$ or $\frac{B}{B_I} \in \left(\frac{1}{1-\gamma^h}, \frac{1-\gamma}{\gamma^{1+h}}\right)$ if $\gamma + \gamma^h > 1$, where h is the minimum distance between any two intermediate states (Assumption 3.5).*

Here we highlight some key insights of the proof and leave the details in Appendix A.4. We first show that, when k is not large ($k < D(s, \mathcal{S}_I) + h$), $V_k(s)$ will only be affected by the closest \mathcal{S}_I . Hence the value function $V_k(s)$ has the same property as the sparse reward setting provided in Proposition 4.1, namely, $V_k(s)$ increases as $D(s, \mathcal{S}_I)$ decreases. Then when $k \geq D(s, \mathcal{S}_I) + h$, we show that the monotonicity between $V_k(s)$ and $D(s, \mathcal{S}_I)$ still holds using the relative magnitude B/B_I . More specifically, we use the bound of B/B_I to show the reward from pursuing the closest \mathcal{S}_I will be larger than the overall rewards of any other trajectories that *do not* pursuit the closest \mathcal{S}_I . With the monotonicity of $V_k(s)$ and $D(s, \mathcal{S}_I)$, we conclude a greedy agent finds the closest intermediate state.

A direct implication of Theorem 4.3 is that, when the magnitude of the terminal reward B is not significantly larger than the intermediate reward B_I , $\forall s \in \mathcal{S} \setminus \mathcal{S}_T$, after $k \geq \max_{s \in \mathcal{S} \setminus \mathcal{S}_T} \min_{j \in \mathcal{I}_d(s)} D(s, s_{i_j})$ iteration of SVI updates, a greedy agent eventually finds to an intermediate state from which \mathcal{S}_T is directly reachable. When \mathcal{S}_T is directly reachable, the next theorem illustrates the condition that enables a greedy agent follow to pursue the shortest path to \mathcal{S}_T .

Theorem 4.4 (Finding the Shortest Path to \mathcal{S}_T) *Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ be a deterministic MDP with initial state s_0 , intermediate states $\mathcal{S}_I = \{s_{i_1}, s_{i_2}, \dots, s_{i_N}\}$, and terminal states \mathcal{S}_T . Suppose \mathcal{M} satisfies Assumption 3.2, 3.3 (b), and 3.5, the reward function $r(\cdot)$ follows the intermediate reward setting (Table 3), and the value function, Q -function are zero-initialized (4.1). $\forall s \in \mathcal{S} \setminus \mathcal{S}_T$, let*

$$d = D(s, \mathcal{S}_T) \quad \text{and} \quad d_I = D(s, \mathcal{S}_I) \doteq \min_{j \in \mathcal{I}_d(s)} D(s, s_{i_j}). \quad (4.5)$$

If \mathcal{S}_T is directly reachable from s , and d and d_I satisfy

$$d < \begin{cases} d_I + \log_{\frac{1}{\gamma}} \left[(1 - \gamma^h) \frac{B}{B_I} \right], & \text{if } \frac{B}{B_I} < \frac{1}{1-\gamma^h}, \\ d_I + \log_{\frac{1}{\gamma}} \left(\frac{B}{B_I + \gamma^h B} \right), & \text{if } \frac{B}{B_I} \geq \frac{1}{1-\gamma^h}, \end{cases} \quad \text{and} \quad d < d_I + h - 1, \quad (4.6)$$

where h is the minimum distance between two intermediate states (Assumption 3.5), then after $k \geq d$ synchronous value iteration updates (2.1), an agent following the greedy policy will pursue the shortest path to \mathcal{S}_T .

The proof of Theorem 4.4 shares the same idea as the proof of Theorem 4.3. We first provide an explicit expression of the value function $V_k(s)$ for $k \leq d_I + h$, in Lemma B.3. In addition, we demonstrate that $V_k(s)$ is monotonic decreasing as $D(s, \mathcal{S}_T)$ increases when $k \leq d_I + h$, since $V_k(s)$ is only affected by the closest \mathcal{S}_I and \mathcal{S}_T when $k \leq d_I + h$. Then we show the value function is monotonically decreasing as the distance $D(s, \mathcal{S}_T)$ increases for all $k \geq d$, when the conditions in Theorem 4.4 are satisfied. Similar to Theorem 4.3, we use the relative magnitude B/B_I to show the the monotonicity of $V_k(s)$ still holds when $k > d_I + h$. More precisely, we use the bound of B/B_I to ensure the rewards from pursuing \mathcal{S}_T will be larger than all future rewards from any other trajectories that *do not* pursuit \mathcal{S}_T . Finally, we use the monotonicity of $V_k(s)$ to show that greedy policies find the shortest path to \mathcal{S}_T . The proof of Theorem 4.4 is provided in Appendix A.5.

Intuitively, Theorem 4.4 indicates that, if a state s is “close enough” to the terminal states \mathcal{S}_T , then the greedy policy will lead the agent from s to \mathcal{S}_T following the shortest path, given sufficient number of value iteration updates. Note that Theorem 4.4 implicitly uses the fact that, in the OWMP setting, there could exist some states from which both \mathcal{S}_I and \mathcal{S}_T are directly reachable (see Figure 2b).

All in all, when B and B_I satisfy the conditions in Theorem 4.3 and when

$$k \geq \max_{s \in \mathcal{S} \setminus \mathcal{S}_T} \min_{j \in \mathcal{I}_d(s)} D(s, s_{i_j}),$$

a greedy agent will recursively find the closest directly reachable intermediate state, until the terminal state \mathcal{S}_T is directly reachable. When \mathcal{S}_T is directly reachable and $D(s, \mathcal{S}_T)$, $D(s, \mathcal{S}_I)$ satisfy the conditions provided in Theorem 4.4, the agent following the greedy policy will pursue the shortest path to \mathcal{S}_T . We shall clarify that our theoretical results (Theorem 4.3 and 4.4) on the OWMP setting only provides *sufficient conditions* of obtaining successful Q-functions, but successful Q-functions can actually be obtained under broader conditions, as will be presented in Section 6.

4.4 Non One-Way Intermediate Rewards

The theoretical results in the non one-way intermediate (NOW) reward setting is out of the scope of this work, since we do not have additional assumptions on the intermediate states. Still, we provide an example where SVI *never* finds a successful Q-function in Example 4.5 to demonstrate the effect of non-ideal reward design.

Example 4.5 (Synchronous Value Iteration Fails) *Suppose we are given an MDP environment as shown in Figure 5, with $B = 10$, $B_I = 100$, and $\gamma = 0.9$. The value function converges when $k \geq 2$. However, $\forall k \geq 1$, a greedy policy will stay in s_i instead of moving to \mathcal{S}_T .*

As shown in Example 4.5, if we set the intermediate rewards B_I on some *non one-way* intermediate states, and when the intermediate rewards are significantly larger than the

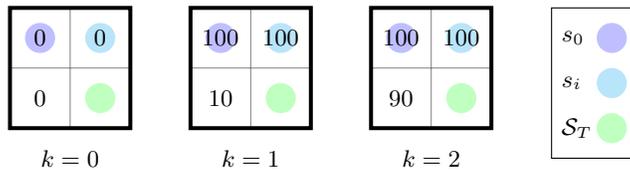


Figure 5: An example of the evolution of a zero-initialized value function during SVI in the non one-way intermediate reward setting, with $B = 10$, $B_I = 100$, and $\gamma = 0.9$. The agent may move left, right, up, or down. The cyan circle is the intermediate state that can be *repeatedly* visited from s_0 and from itself (by taking actions “right” and “up” and hitting a wall). The remaining settings are the same as Figure 3.

terminal rewards, a greedy policy will lead the agent to visit the intermediate states repeatedly rather than pursuing the terminal states. A direct implication from Example 4.5 is that, to prevent an agent from staying at non one-way intermediate states, the intermediate rewards on these non one-way intermediate states should be relatively small.

5. Connections to Finding the Shortest Path

We discuss the connections between the different reward settings (sparse reward, OWSP, OWMP, NOW) and their connections to finding the shortest path in this section.

The Sparse Reward Setting The sparse reward setting *requires no prior knowledge nor assumptions of the subgoals of the tasks*, because it contains no intermediate states. Hence, it can further be applied to general goal-reaching tasks. In addition, as illustrated in Proposition 4.1, the successful policy obtained by SVI pursues the shortest path but at the cost of high computational complexity ($D(s_0, \mathcal{S}_T)$ total number of SVI). The computational complexity in the sparse reward setting motivates the use of intermediate rewards: practical tasks with well-designed intermediate rewards should be more computationally efficient than the same task with only sparse terminal rewards.

The OWSP Intermediate Reward Setting As mentioned in Section 3.2, the OWSP intermediate reward setting only applies to limited number of tasks seen in practice. However, when the subgoals of the goal-reaching tasks indeed *has the one-way single-path structure*, Proposition 4.2 shows that the computational complexity of obtaining a successful Q-function can be significantly reduced, and a greedy policy will find the shortest path from s_0 to \mathcal{S}_T .

The OWMP Intermediate Reward Setting Comparing to the OWSP setting, the OWMP intermediate setting has much broader practical applications. Though the OWMP intermediate setting is more prevalent, the downside is that generally a greedy policy with a successful Q-function will pursue the closest intermediate states instead of the shortest path to \mathcal{S}_T . Figure 6 illustrates the trade-off between computational complexity and achieving the shortest path is provided: A greedy policy reaches \mathcal{S}_T when k (the number of SVI) equals 1 or 2, by recursively pursuing the closest intermediate state (path 1). When $k > 2$ and the terminal rewards B and intermediate rewards B_I satisfy the conditions stated in Theorem 4.4, the greedy policy will find \mathcal{S}_T via the shortest path (path 2).

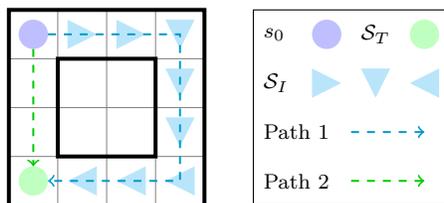


Figure 6: The trade-off between the minimum computational complexity and the pursuit of the shortest path in the OWMP intermediate reward setting. The cyan isosceles triangles are the one-way intermediate states and the orientations of each isosceles triangle represents the direction of each intermediate state – the agent at a given intermediate state can only visit the next state pointed by the orientation of the apex. The remaining settings are the same as Figure 3.

6. Experiment

We experimentally verify in several OpenAI Gym MiniGrid environments that the agent is able to learn a successful trajectory more quickly in OWSP and OWMP intermediate reward settings than the sparse reward setting. We verify our findings on ϵ -greedy asynchronous Q-learning algorithm, and three popular deep RL algorithms: DQN (Mnih et al., 2015), A2C (Mnih et al., 2016), and PPO (Schulman et al., 2017). For all experiments, the agent observes the whole environment. For asynchronous Q-learning, each state is represented as a string encoding of the grid. For deep RL algorithms, each state is an image of the grid. The detailed parameters and additional related experiments are provide in Appendix C.2.2. See <https://github.com/kebaek/minigrid> for the code to run all presented experiments.

The main purpose of our experiments is to justify the following indications from our theoretical results: 1) adding intermediate rewards on one-way intermediate states reduces computational complexity; 2) there is a trade-off between computational complexity and the pursuit of the shortest path in the OWMP setting. Although we have discussed many examples of RL applications in Table 2, we choose the OpenAI Gym MiniGrid environments (Brockman et al., 2016; Chevalier-Boisvert et al., 2018) rather than the Games (Pacman, Montezuma, Dota2, StarCraft II, and Honor of Kings) discussed in Table 2 due to limited computational resources and the intrinsic randomness in these Game.

6.1 Environmental Setting

Single-Path Maze The 7x7 grid maze (Figures 7a and 7b) consists of a single path that the agent must navigate through to reach the terminal state. This environment will be used to study the computational benefit of having well-designed intermediate rewards for the OWSP setting. See Figure 7 for possible actions and reward design. Note that this maze environment is one-way since each intermediate reward may only be obtained once, i.e., the blue circle disappears from the environment once the agent reaches the corresponding square.

3-Door/4-Door These 9x9 grid mazes (Figures 7c and 7d) consist of 3 different paths to to the terminal state each sealed by a series of locked doors. Each door has a corresponding key of the same color. The goal of the agent is to reach the terminal state by picking up the corresponding keys to unlock all the doors along at least one of the three paths. Note that

this environment is one-way since a door cannot lock once unlocked and a key cannot be dropped once picked up. These environments will be used to study the computational benefit of having well-designed intermediate rewards and the trade-off between computational complexity and shortest path for the OWMP setting (Section 4.3).

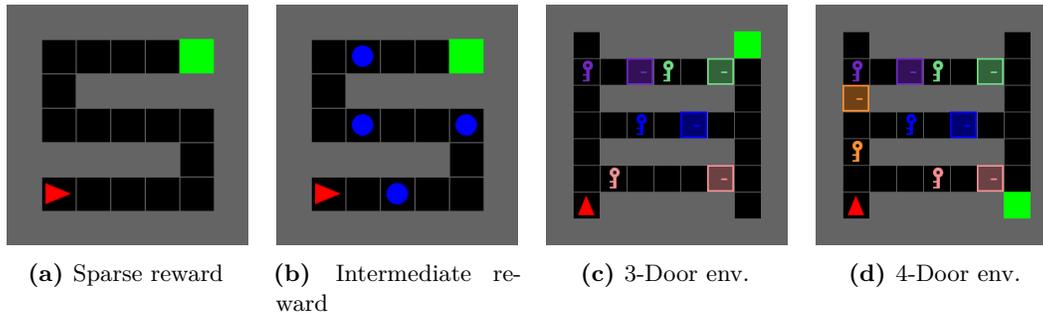


Figure 7: The red triangle is the agent and the green square is the terminal state. (a)-(b): The Single-Path Maze environment. Agent takes actions from {go forward, turn 90° , turn -90° }. For the sparse reward setting (a), the agent receives a terminal reward of +10. For the intermediate reward setting (b), the agent also receives +1 for arriving at any square with a blue circle. (c)-(d): The Door-Key environments. Agent takes actions from {go forward, turn 90° , turn -90° , pick up key, open door}. For the sparse reward setting of (c) and (d), the terminal reward is +10. For the intermediate reward setting of (c) and (d), the agent also receives +2 for picking up a key or opening a door. All rewards in environments (a)-(d) can only be obtained *once* per episode.

6.2 Results

On the Single-Path Maze environment and the 3-Door environment, we compare the number of training episodes required for the agent to find a successful policy in the sparse reward setting and the intermediate reward setting (See Figure 7a and 7c for more details about reward design). For 0.8 ϵ -greedy Q-learning, we train 100 independent models and evaluate each model once, for a total of 100 trials. For the deep RL algorithms, we train 10 independent models, and evaluate each model 10 times with different seeds, for a total of 100 trials. Win rate is computed as the number of trials that reach the terminal state out of 100.

As expected, we observe that it takes ϵ -greedy Q-learning 36 episodes to reach a win rate of 100% whereas in the sparse reward setting, ϵ -greedy Q-learning is only able to reach a win rate of 10% for the same number of episodes (See Table 5). Similarly for the 3-Door setting, if the agent is also rewarded for picking up keys and opening doors, significantly less training episodes are required to obtain a win rate of 100% (See Table 5). We observe a similar phenomena on popular deep RL algorithms: DQN, A2C, and PPO (See Figure 8). Our experimental findings on the computational benefits of using intermediate rewards corroborate our theoretical claim in Section 4.

On the 4-Door environment, we test two intermediate reward settings where the agent is either rewarded 10 or 1000 points for reaching the terminal state. Additionally, the agent is rewarded 2 points for either picking up a key or opening a door. The shortest successful path to the terminal state (12 steps) has the least number of doors to unlock, thus less intermediate rewards, whereas the longest of the three paths to the terminal state contains at

Maze (Figure 7a and 7b)			3-Door (Figure 7c)		
# Episodes	Sparse	Intermediate	# Episodes	Sparse	Intermediate
18	6/100	59/100	40	3/100	90/100
24	7/100	82/100	80	10/100	100/100
30	6/100	95/100	120	43/100	100/100
36	10/100	100/100	160	74/100	100/100

Table 5: Asynchronous Q-learning: Computational Complexity. We report the number of wins out of 100 trials (100 training sessions evaluated once each) after different training episodes for both the sparse reward and intermediate reward settings.

# Episodes	Setting 1: +10			Setting 2: +1000		
	Wins	Rewards	Steps	Wins	Rewards	Steps
50	64/100	11.28 ± 2.08	95.52 ± 82.58	71/100	11.12 ± 2.30	84.39 ± 79.48
150	99/100	11.2 ± 2.03	24.71 ± 18.10	100/100	11.06 ± 1.99	22.53 ± 3.56
350	100/100	11.1 ± 1.93	22.65 ± 3.40	100/100	10.14 ± 2.60	20.88 ± 4.70
750	100/100	11.18 ± 1.80	22.85 ± 3.24	100/100	9.24 ± 2.85	19.39 ± 5.54
1550	100/100	11.64 ± 1.42	23.36 ± 2.62	100/100	5.52 ± 1.10	12.95 ± 1.64
3150	100/100	12.0 ± 0.0	24.0 ± 0.0	100/100	4.46 ± 0.84	12.23 ± 0.42

Table 6: 4-Door (Figure 7d) Trade-Off for Asynchronous Q-Learning. We report the number of wins out of 100 trials (100 training sessions evaluated once each), averaged total intermediate rewards, and averaged number of steps taken after different training episodes. The agent receives a terminal reward of +10 in Setting 1 +1000 in Setting 2. The agent also receives +2 for picking up a key and +2 for opening a door in both settings. The episode maxes out at 324 steps.

least 3 doors to unlock, thus more possible intermediate rewards to collect. We compare the average steps required during evaluation for the agent to reach the terminal state between these two settings. From Section 4.3, we expect that in this OWMP setting, the agent identifies the shortest successful path with Q-learning given a well-designed discount factor and a good ratio between intermediate/terminal rewards. For a discount factor of 0.9, we observe that if the terminal reward is 1000, the path taken by an agent trained by ϵ -greedy Q-learning converges to the shortest path of 12 steps. On the other hand, if the terminal reward is 10, the path taken by the agent converges to that of 24 steps and on average collects more intermediate rewards (See Table 6). Similarly, with a discount factor of 0.8 for DQN and 0.9 for A2C and PPO, we observe that the agent chooses the shortest successful trajectory when rewarded 1000 points for reaching the terminal state during training, and a longer successful trajectory with more intermediate rewards when rewarded 10 points. Our experimental findings on the trade-off between computational efficiency and the shortest path corroborate our theoretical claim in Section 4.3.

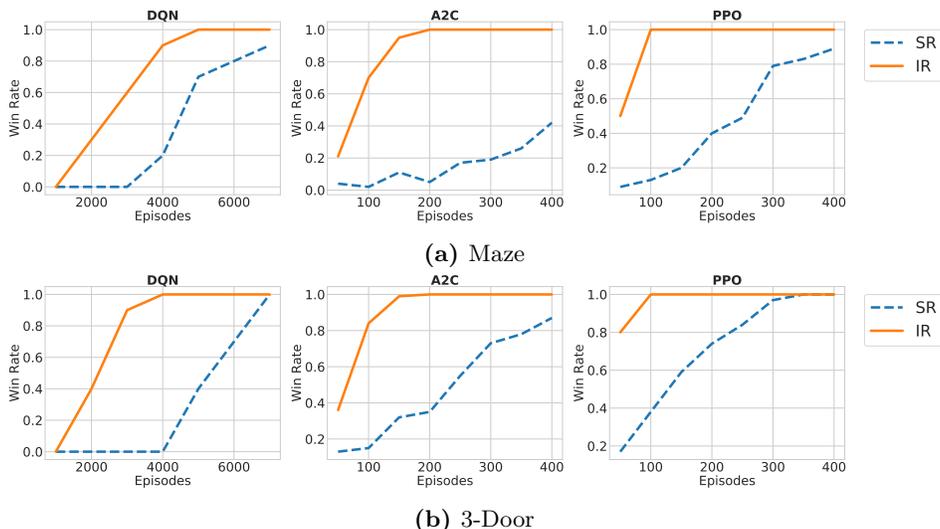


Figure 8: Deep RL: Computational Complexity. We compare the average number of steps an agent takes to reach the terminal state in the Single-Path Maze and 3-Door environments between sparse versus intermediate reward settings. If the agent does not reach the terminal state, the episode maxes out at 324 steps for 3-Door and 196 steps for Single-Path Maze. The results are averaged over 100 trials (10 training sessions evaluated 10 times each)

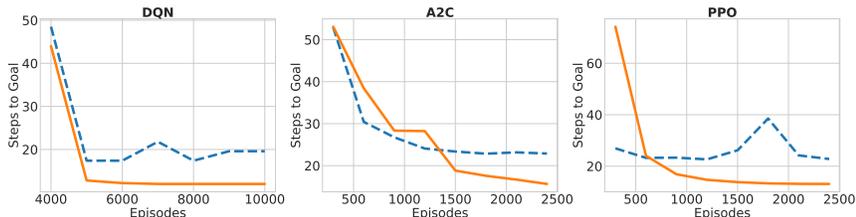


Figure 9: Deep RL: Trade-Off. We compare the average number of steps an agent takes to reach the terminal state in the 4-Door environment if the terminal reward is 10 versus 1000. The agent also receives an intermediate reward of 2 when it picks up a key or opens a door. If the agent does not reach the terminal state, the episode maxes out at 324 steps. The results are averaged over 100 trials (10 training sessions evaluated 10 times each)

7. Related Works and Discussion

7.1 Related Works

Hierarchical RL Hierarchical reinforcement learning and planning are two fundamental problems that have been studied for decades (Dayan & Hinton, 1993; Kaelbling, 1993a; Parr, 1998; Parr & Russell, 1998; McGovern et al., 1998; Sutton et al., 1998; Precup et al., 1998; Sutton et al., 1999; Dietterich, 2000; McGovern & Barto, 2001) (See Barto & Mahadevan, 2003 for an overview of other earlier works on hierarchical RL and Chapter 11.2 in the book Russell & Norvig, 2009 for hierarchical planning). After the success of deep learning (LeCun et al., 2015; Goodfellow et al., 2016), recent works have revisited hierarchical RL under the deep RL framework (Kulkarni et al., 2016; Vezhnevets et al., 2017; Andreas et al., 2017; Le et al., 2018; Xu et al., 2020) (see Chapter 11 in the review paper Li,

2018 for other hierarchical deep RL literature). Prior theoretical attempts on hierarchical RL formulated the problem as MDP decomposition problems (Dean & Lin, 1995; Singh et al., 1998; Meuleau et al., 1998; Wen et al., 2020) or solving subtasks with “bottleneck” states (Sutton et al., 1999; McGovern & Barto, 2001; Stolle & Precup, 2002; Simsek & Barreto, 2008; Solway et al., 2014). Our result is closely related to Wen et al. (2020) in the following two aspects: 1) Our one-way assumption (Assumption 3.2) on the intermediate states is similar to the exit state of subMDPs (Definition 1 in Wen et al., 2020), in the sense that our one-way intermediate states can be viewed as the exit states that separate the MDP into subMDPs; 2) Our quantitative results of different MDP settings suggest that partitioning the large MDP via intermediate states and intermediate rewards generally reduce the computational complexity, which corroborates the computational efficiency of subMDPs in Wen et al. (2020). Though we come to similar theoretical conclusions that intermediate rewards lead to reduced computational complexity, in this work, we simplify down the assumptions and build a theoretical framework that is well connected to practice.

Reward Design With the recent success of deep learning, RL has experienced a renaissance (Krakovsky, 2016) and has demonstrated super-human performance in various applications (Mnih et al., 2013, 2015; Silver et al., 2016, 2017; Vinyals et al., 2017, 2019; Berner et al., 2019; Ye et al., 2020; Fuchs et al., 2021). The reward design varies from task to task. For example, in Go (Silver et al., 2016, 2017), the agent only receives *terminal rewards* at the end of the game (+1 for winning and -1 for losing); for Starcraft II (Vinyals et al., 2019), the reward function is usually a mixture win-loss terminal rewards (+1 on a win, 0 on a draw, and -1 on a loss) and intermediate rewards based on human data; for multiplayer online battle arena (MOBA) games (OpenAI, 2018; Berner et al., 2019; Ye et al., 2020), the reward functions are generally heavily handcrafted (see Table 6 of Berner et al., 2019 and Table 4 of Ye et al., 2020) based on prior knowledge. Besides task-dependent reward design, other works also have studied the reward design for general RL or robotic tasks (Singh et al., 2009, 2010; Sorg et al., 2010; Vezhnevets et al., 2017; Van Seijen et al., 2017; Raileanu & Rocktäschel, 2020; Ratner et al., 2018, 2018; He & Dragan, 2021) (See Guo, 2017; Doroudi et al., 2019 and references therein; also see other literature in Chapter 5 of the review paper Li, 2018).

Goal-Conditioned RL Goal-conditioned RL, the problem of learning a policy that reaches certain goal states, has been empirically studied in many prior works (Kaelbling, 1993b; Sutton et al., 2011; Andrychowicz et al., 2017; Fu et al., 2018; Pong et al., 2018; Ghosh et al., 2019; Eysenbach et al., 2020a, 2020b; Kadian et al., 2020; Fujita et al., 2020; Chebotar et al., 2021; Khazatsky et al., 2021). The goal-conditioned RL is closely related to the sparse reward setting in our framework, where the agent *only* receives terminal rewards at the terminal (goal) states. Moreover, the empirical observations where goal-reaching tasks can be improved via pursuing subgoals (Andreas et al., 2017; Nasiriany et al., 2019) also corroborate with the computational benefits of rewarding intermediate states suggested by our theoretical results in Section 4.

Connection to Reward Shaping (Ng et al., 1999) For a given MDP \mathcal{M} with reward function r , (Ng et al., 1999) proposed a potential-based shaping function F , such that the same MDP \mathcal{M} with shaped reward $F + r$ has the same optimal policy as \mathcal{M} with the original reward function r . Part of our result is related to reward shaping (Ng et al., 1999)

in the sense that, when certain conditions are satisfied (Assumption 3.3 (a)), greedy policies under the sparse reward setting and the intermediate reward setting are the same, as they both follow the shortest path from s_0 to \mathcal{S}_T . Comparing to reward shaping, the advantage of our work is also on the practical side, since the assumptions of the potential function F is generally hard to satisfy. F is usually approximated via neural networks in practice, which requires extra engineering efforts. On the contrary, as we have discussed in Section 3 and Section 6, our assumptions on the one-way intermediate states (Assumption 3.3) and relative magnitude between intermediate rewards B_I and terminal rewards B can easily be implemented in practice.

7.2 Discussion

Practical Implications Our work theoretically verifies the common practice of adding intermediate rewards to speed up training in reinforcement learning. Formally, in order to find a goal-reaching successful policy, adding intermediate rewards based on prior knowledge of the practical tasks is generally more computationally efficient than using sparse terminal rewards alone. However, unless the intermediate rewards are carefully designed (e.g., like the OWSP setting described in Section 4.2), greedy policies usually do not follow the shortest path to the terminal states. To prevent the agent from getting stuck at non one-way intermediate states (e.g., like the case discussed in Example 4.5), we can assign relatively smaller rewards (compared to the terminal rewards) to non one-way intermediate states. Our findings corroborate the reward design of Dota2 in Table 6 of (Berner et al., 2019): one can understand the “Win” (with reward +5) as terminal states; “XP Gained” (with reward +0.002), “Gold Gained” (with reward +0.006), and “Gold Spent” (with reward +0.0006) as *non one-way* intermediate states; “T1 Tower” (with reward +2.25), “T2 Tower” (with reward +3), “T3 Tower” (with reward +4.5), “T4 Tower” (with reward +2.25), “Shrine” (with reward +2.25), and “Barracks” (with reward +6) as *one-way* intermediate states.

Negative Rewards One limitation of this work is that we only consider positive rewards, but in many applications described before, the agent may receive negative rewards upon the arrival of some unfavorable states. We will provide several examples here to provide some intuitions on the effect of negative rewards in Figure 10 and leave formal studies to future. In Figure 10 (a), for an intermediate state with negative reward which is not in the path from s_0 to \mathcal{S}_T , then given enough computational complexity, a greedy agent still finds \mathcal{S}_T . However, if there is an intermediate state with negative reward that is in the path from s_0 to \mathcal{S}_T , the behavior of a greedy agent will depend on the magnitude of the negative rewards. As shown in Figure 10 (b) and (c), when the magnitude of the negative intermediate rewards $r(s, a, s_-)$ is relatively small (as shown in (b), $r(s, a, s_-) = -1$), then a greedy agent will still find \mathcal{S}_T because the overall reward from \mathcal{S}_T and s_- is larger than 0; however, when the magnitude of the negative reward is large (as shown in (c), $r(s, a, s_-) = -10$), then a greedy agent will not find \mathcal{S}_T because the overall reward of s_0 and \mathcal{S}_T is smaller than 0.

Unequal Magnitude Intermediate Rewards Another limitation of this work is the equal magnitude intermediate reward. In Section 3.3, we assume all intermediate rewards are equal magnitude B_I (but can be different from the terminal rewards B). Such equal magnitude assumption helps us provide a bound w.r.t. relative magnitude between termi-

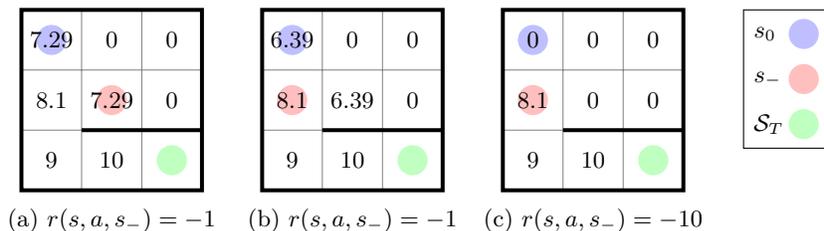


Figure 10: An example of the value functions after $k = 4$ iterations of SVI and with terminal rewards $B = 10, \gamma = 0.9$, with different structures and magnitude of negative intermediate rewards.

nal rewards and intermediate rewards (the bound of B/B_I described in Theorem 4.3 and Theorem 4.4), hence ensuring the monotonicity of the value function $V_k(s)$ w.r.t. $D(s, \mathcal{S}_I)$, $D(s, \mathcal{S}_T)$ in Theorem 4.3, Theorem 4.4, respectively. One way to generalize the equal magnitude intermediate rewards is to assume all intermediate rewards are bounded between $[B_l, B_r]$, and then apply the same techniques as Theorem 4.3 and 4.4 to provide bounds that involve B/B_l and B/B_u .

7.3 Future Works

Our work can be extended in these following directions: 1) Generalize from the deterministic MDPs in our framework to stochastic MDPs; 2) Study the tightness of the bounds for the relative magnitude B/B_I in Theorem 4.3 and 4.4; 3) Generalize to other scenarios that contains negative rewards (e.g., from “bad” terminal states or intermediate states); 4) Generalize to intermediate rewards of unequal magnitude; 5) Study how rewarding “good” non one-way intermediate states (instead of just one-way states) affects the goal-reaching problems in general.

Acknowledgements

We would like to thank anonymous reviewers of JAIR for their insightful suggestions on revising the draft, especially on highlighting the main results and contributions in the introduction. We would also like to thank Professor Mykel Kochenderfer of Stanford for coordinating the revision. We would also like to thank Michael Psenka from UC Berkeley for proofreading the manuscript. YZ would like to thank Haozhi Qi from UC Berkeley for insightful discussion on the implementations of practical deep RL tasks. YZ would also like to thank professor Sergey Levine and Qiyang (Colin) Li from the RAIL lab for insightful suggestions on the connections between this work and goal-conditioned RL. CB would also like to thank Siddharth Verma from UC Berkeley for helpful discussions about deep RL algorithms. ZZ acknowledges the JP Morgan AI Research Grant. JJ acknowledges NSF Grants IIS-1901252 and CCF-1909499. YM acknowledges support from ONR grant N00014-20-1-2002 and the joint Simons Foundation-NSF DMS grant #2031899, as well as support from Berkeley FHL Vive Center for Enhanced Reality and Berkeley Center for Augmented Cognition. YZ and YM acknowledge support from Tsinghua-Berkeley Shenzhen Institute (TBSI) Research Fund.

Appendix A. Extra Definition and Proof of Section 4

A.1 Definition of Correct and Incorrect Actions

We introduce the definition of correct and incorrect actions to facilitate the future proof.

Definition A.1 (Correct and Incorrect Actions) *Let s_a be the subsequent state the state-action pair (s, a) , we define the correct action set $\mathcal{A}^+(s)$ of a state $s \in \mathcal{S} \setminus \mathcal{S}_T$ as*

$$\mathcal{A}^+(s) = \{a | a \in \mathcal{A}, D(s_a, \mathcal{S}_T) = D(s, \mathcal{S}_T) - 1\}. \quad (\text{A.1})$$

Conversely, we can define the incorrect action set of a state $s \in \mathcal{S}$ as

$$\mathcal{A}^-(s) = \{a | a \in \mathcal{A} \setminus \mathcal{A}^+(s)\}. \quad (\text{A.2})$$

A.2 Proof of Proposition 4.1

Proposition A.2 (Sparse Rewards) *Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ be a deterministic MDP with initial state s_0 and terminal states \mathcal{S}_T . If the reward function $r(\cdot)$ follows the sparse reward setting (Table 3) and the value function and Q-function are zero-initialized (4.1), then after any $k \geq D(s_0, \mathcal{S}_T)$ synchronous value iteration updates (2.1), the Q-function Q_k is a successful Q-function, and a greedy policy follows the shortest path from s_0 to \mathcal{S}_T .*

Proof $\forall d \leq k$, let \mathcal{S}^d denote the set of states that is distance d to the desired terminal state \mathcal{S}_T , and let $a^{d+} \in \mathcal{A}^+(s^d)$, $a^{d-} \in \mathcal{A}^-(s^d)$ denote a correct action and an incorrect action of s^d , respectively. From Lemma B.1, we know that

$$V_k(s^d) = \begin{cases} \gamma^{d-1}B, & \forall k, d \in \mathbb{N}^+, d \leq k, \\ 0, & \text{otherwise.} \end{cases}$$

Combine Lemma B.1 with the Q-function update in value iteration (2.1):

$$Q_k(s, a) = r(s, a, s_a) + \gamma V_k(s_a), \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A},$$

we know that

$$Q_k(s^d, a^{d+}) \stackrel{(i)}{=} \begin{cases} \gamma^{d-1}B, & \text{if } d \leq k, \\ 0, & \text{otherwise,} \end{cases} \quad (\text{A.3})$$

$$Q_k(s^d, a^{d-}) \stackrel{(ii)}{\leq} \begin{cases} \gamma^d B, & \text{if } d \leq k, \\ 0, & \text{otherwise,} \end{cases} \quad (\text{A.4})$$

where inequality (ii) holds since Definition A.1 implies that $D(s_{a^{d-}}, \mathcal{S}_T) \geq D(s, \mathcal{S}_T)$. We shall clarify equality (i) for the case when $d = 1$: in this case, $V_k(s) = 0$ always holds $\forall s \in \mathcal{S}_T$ because given a state-action pair (s, a) , the MDP stops once the subsequent state $s_a \in \mathcal{S}_T$, hence we have

$$Q_k(s^1, a^{1+}) = r(s^1, a^{1+}, s_{a^{1+}}^1) + \gamma V_k(s_{a^{1+}}^1) = B. \quad (\text{A.5})$$

Since $\gamma < 1$, we know that

$$Q_k(s^d, a^{d-}) < Q_k(s^d, a^{d+}), \forall k \geq d, \quad (\text{A.6})$$

which implies

$$\arg \max_{a \in \mathcal{A}} Q_k(s^d, a) \in \mathcal{A}^+(s^d), \forall k \geq d. \quad (\text{A.7})$$

Now consider s_0 , let $d_m = D(s_0, \mathcal{S}_T)$, (A.6) implies that any greedy action is a correct action:

$$a_0 = \arg \max_a Q_k(s_0, a) \in \mathcal{A}^+(s_0), \forall k \geq d_m \quad (\text{A.8})$$

By Definition A.1, the subsequent state s_1 of state-action pair (s_0, a_0) has distance at most d_m ($D(s_1, \mathcal{S}_T) \in [d_m - 1, d_m]$) to \mathcal{S}_T , again (A.6) implies that

$$a_1 = \arg \max_a Q_k(s_1, a) \in \mathcal{A}^+(s_1). \quad (\text{A.9})$$

Likewise, we know that the greedy policy generates a trajectory $\{s_0, a_0, s_1, a_1, \dots\}$, such that

$$a_i = \arg \max_a Q_k(s_i, a) \in \mathcal{A}^+(s_i), \quad (\text{A.10})$$

which eventually ends up with a state $s_n \in \mathcal{S}_T$, since Definition A.1 indicates that s_{i+1} is one step closer to \mathcal{S}_T than s_i . Hence, we conclude that Q_k is a successful Q-function. When $k \geq d$, since (A.10) shows that every action taken by the greedy policy will take the agent one step closer to \mathcal{S}_T , we can also conclude that the greedy policy also pursues the shortest path from s_0 to \mathcal{S}_T . \blacksquare

A.3 Proof of Proposition 4.2

Proposition A.3 (Single-path Intermediate States) *Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ be a deterministic MDP with initial state s_0 , intermediate states $\mathcal{S}_I = \{s_{i_1}, s_{i_2}, \dots, s_{i_N}\}$, and terminal states \mathcal{S}_T . Suppose \mathcal{M} satisfies Assumption 3.2 and 3.3 (a). If the reward function $r(\cdot)$ follows the intermediate reward setting (Table 3) and the value function and Q-function are zero-initialized (4.1), then after*

$$k \geq d_{\max} \doteq \max\{D(s_0, s_{i_1}), D(s_{i_1}, s_{i_2}), \dots, D(s_{i_{N-1}}, s_{i_N}), D(s_{i_N}, \mathcal{S}_T)\} \quad (\text{A.11})$$

synchronous value iteration updates (2.1), the Q-function Q_k is a successful Q-function, and a greedy policy follows the shortest path from s_0 to \mathcal{S}_T .

Proof The proof of this Proposition is a direct result of Lemma B.2. Similar to Lemma B.2, $\forall j \in [n + 1]$, let $\mathcal{S}_{[j-1, j]}$ denote the sets

$$\mathcal{S}_{[j-1, j]} = \begin{cases} \{s \in \mathcal{S} | D(s, s_{i_1}) < \infty\} & \text{if } j = 1, \\ \{s \in \mathcal{S} | D(s, s_{i_j}) = \infty, D(s, s_{i_{j+1}}) < \infty\} & \text{if } j = 2, 3, \dots, N, \\ \{s \in \mathcal{S} | D(s, s_{i_n}) = \infty, D(s, \mathcal{S}_T) < \infty\} & \text{if } j = N + 1. \end{cases} \quad (\text{A.12})$$

For a state $s^{\mathbf{d}} \in \mathcal{S}_{[j-1, j]}$, let

$$\mathbf{d} = [d_j, d_{j+1}, \dots, d_{N+1}]^\top \in \mathbb{R}^{n+1}, \quad (\text{A.13})$$

where $d_l = D(s, s_{i_l})$, $\forall l = j, j+1, \dots, N$ and $d_{N+1} = D(s, \mathcal{S}_T)$, denote a vector whose entries are the distance from s to intermediate states s_{i_j} (and the terminal states \mathcal{S}_T). Let $a^{\mathbf{d}^+} \in \mathcal{A}^+(s^{\mathbf{d}})$, $a^{\mathbf{d}^-} \in \mathcal{A}^-(s^{\mathbf{d}})$ denote a correct and an incorrect action of $s^{\mathbf{d}}$. Combine Lemma B.2 with the Q-function update in value iteration (2.1)

$$Q_k(s, a) = r(s, a, s_a) + \gamma V_k(s_a), \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A},$$

we know that

$$Q_k(s^{\mathbf{d}}, a^{\mathbf{d}^+}) = \sum_{l=j}^{N+1} v(k, d_l) B_l = V_k(s), \quad (\text{A.14})$$

$$Q_k(s^{\mathbf{d}}, a^{\mathbf{d}^-}) \stackrel{(i)}{\leq} \gamma \sum_{l=j}^{N+1} v(k, d_l) B_l = \gamma V_k(s), \quad (\text{A.15})$$

where

$$v(k, d_l) = \begin{cases} \gamma^{d_l-1}, & \forall k, d \in \mathbb{N}^+, d_l \leq k+1, \forall l \in [N+1], \\ 0, & \text{otherwise,} \end{cases} \quad \text{and} \quad B_l = \begin{cases} B_I, & \text{if } l \in [N], \\ B, & \text{if } l = N+1, \end{cases} \quad (\text{A.16})$$

and inequality (i) holds because Definition A.1 implies $D(s_{a^{\mathbf{d}^-}}, \mathcal{S}_T) \geq D(s^{\mathbf{d}}, \mathcal{S}_T)$ and Assumption 3.2 guarantees that previous intermediate states $s_{i_1}, s_{i_2}, \dots, s_{i_{j-1}}$ cannot be revisited. Since $\gamma < 1$, notice that when $k \geq d_j$, we have

$$Q_k(s^{\mathbf{d}}, a^{\mathbf{d}^+}) > Q_k(s^{\mathbf{d}}, a^{\mathbf{d}^-}), \quad (\text{A.17})$$

which implies

$$\arg \max_{a \in \mathcal{A}} Q_k(s^{\mathbf{d}^+}, a) \in \mathcal{A}^+(s^{\mathbf{d}}). \quad (\text{A.18})$$

Now consider s_0 , since we assume that $k \geq d_{\max}$, therefore (A.17) implies that the greedy action is a correct action:

$$a_0 = \arg \max_a Q_k(s_0, a) \in \mathcal{A}^+(s_0). \quad (\text{A.19})$$

By Definition A.1 and Assumption 3.2, we consider the following two cases:

- The subsequent state s_1 of state-action pair (s_0, a_0) does not reach s_{i_1} , ($s_1 \in \mathcal{S}_{[0,1]}$), then s_1 has distance at most d_1 to states s_{i_1} :

$$D(s_1, s_{i_1}) \in [d_1 - 1, d_1]. \quad (\text{A.20})$$

Since $k \geq d_{\max} \geq d_1$, (A.17) implies that the greedy action is a correct action:

$$a_1 = \arg \max_a Q_k(s_1, a) \in \mathcal{A}^+(s_1). \quad (\text{A.21})$$

- The subsequent state s_1 of state-action pair (s_0, a_0) reaches s_{i_1} , ($s_1 \in \mathcal{S}_{[1,2]}$), then s_1 has distance at most d_2 to states s_{i_2} :

$$D(s_1, s_{i_2}) \in [d_2 - 1, d_2]. \quad (\text{A.22})$$

Since $k \geq d_{\max} \geq d_2$, (A.17) implies that the greedy action is a correct action:

$$a_1 = \arg \max_a Q_k(s_1, a) \in \mathcal{A}^+(s_1). \quad (\text{A.23})$$

Likewise, we know that greedy execution generates a trajectory $\{s_0, a_0, s_1, a_1, \dots\}$, such that

$$a_i = \arg \max_a Q_k(s_i, a) \in \mathcal{A}^+(s_i), \quad (\text{A.24})$$

which eventually ends up with a state $s_n \in \mathcal{S}_T$, since Definition A.1 indicates that s_{i+1} is one step closer to \mathcal{S}_T than s_i . Hence, we conclude that Q_k is a successful Q-function. When $k \geq d_{\max}$, since (A.24) shows that every action taken by the greedy policy will take the agent one step closer to \mathcal{S}_T , we can also conclude that the greedy policy also pursues the shortest path from s_0 to \mathcal{S}_T . \blacksquare

A.4 Proof of Theorem 4.3

Theorem A.4 (Finding the Closest \mathcal{S}_I) *Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ be a deterministic MDP with initial state s_0 , intermediate states $\mathcal{S}_I = \{s_{i_1}, s_{i_2}, \dots, s_{i_N}\}$, and terminal states \mathcal{S}_T . Suppose \mathcal{M} satisfies Assumption 3.2, 3.3 (b), and 3.5, if the reward function $r(\cdot)$ follows the intermediate reward setting (Table 3) and the value function and Q-function are zero-initialized (4.1), then $\forall s \in \mathcal{S} \setminus \mathcal{S}_T$, if \mathcal{S}_T is not directly reachable from s , then after*

$$k \geq d = D(s, \mathcal{S}_I) \doteq \min_{j \in \mathcal{I}_d(s)} D(s, s_{i_j}) \quad (\text{A.25})$$

synchronous value iteration update (2.1), an agent following the greedy policy will find the closest directly reachable intermediate state $s_{i_{j^}}$ of s ($D(s, s_{i_{j^*}}) = D(s, \mathcal{S}_I)$), given that $\frac{B}{B_I} \in \left(0, \frac{1}{1-\gamma^h}\right)$ when $\gamma + \gamma^h \leq 1$ or $\frac{B}{B_I} \in \left(\frac{1}{1-\gamma^h}, \frac{1-\gamma}{\gamma^{1+h}}\right)$ when $\gamma + \gamma^h > 1$, where h is the minimum distance between two intermediate states (Assumption 3.5).*

Proof Recall the definition of directly reachable intermediate states (Definition 3.4) of state s : $\{s_{i_j}, \forall j \in \mathcal{I}_d(s)\}$. $\forall s' \in \mathcal{S}$, such that

$$d' = D(s', \mathcal{S}_I) = \min_{j \in \mathcal{I}_d(s')} D(s', s_{i_j}) > d \quad (\text{A.26})$$

is the minimum distance from s' to its closest directly reachable intermediate state, it suffices to show $V_k(s) > V_k(s')$, since $V_k(s) > V_k(s')$ can directly lead to $Q_k(s, a_{i+}) > Q_k(s, a_{i-})$:

$$\begin{aligned} Q_k(s, a_{i+}) &\stackrel{(i)}{=} r(s, a_{i+}, s_{a_{i+}}) + \gamma V_k(s_{a_{i+}}) \stackrel{(ii)}{>} \gamma V_k(s_{a_{i-}}) \\ &\stackrel{(iii)}{=} r(s, a_{i-}, s_{a_{i-}}) + \gamma V_k(s_{a_{i-}}) \stackrel{(iv)}{=} Q_k(s, a_{i-}), \end{aligned} \quad (\text{A.27})$$

where $a_{i+} \in \mathcal{A}_I^+(s)$ is an action that leads to the closest directly reachable intermediate states of state s and $a_{i-} \in \mathcal{A}_I^-(s)$ is the other actions that do not. Equality (i) and (iv) hold by the synchronous value iteration update (2.1), inequality (ii) holds since $D(s_{a_{i-}}, \mathcal{S}_I) > D(s_{a_{i+}}, \mathcal{S}_I)$ and we have assumed $V_k(s) > V_k(s')$ holds $\forall s, s'$ such that $D(s', \mathcal{S}_I) > D(s, \mathcal{S}_I)$, and equality (iii) holds because $s_{a_{i-}}$ is not an intermediate state. Moreover, if the agent at s takes an action a_{i+} and moves to $s_{a_{i+}}$, we will have

$$k \geq d - 1 = D(s_{a_{i+}}, s_{i_j}) \quad (\text{A.28})$$

if $s_{a_{i+}}$ is not an intermediate state, which implies that the action given by the greedy policy will still lead $s_{a_{i+}}$ one step forward to the closest directly reachable intermediate states. Therefore, we only need to show $V_k(s) > V_k(s'), \forall s, s'$ such that $D(s', \mathcal{S}_I) > D(s, \mathcal{S}_I)$, which will be our main focus in the remaining proof.

When $d < k < d + h$. In this case, all intermediate rewards are the same (all equal to B_I), Lemma B.1 implies that

$$V_k(s) = v(k, d)B_I, \quad \forall k < d + h, \quad (\text{A.29})$$

where $d = D(s, \mathcal{S}_I)$ and

$$v(k, d) = \begin{cases} \gamma^{d-1}, & d \leq k, \\ 0, & \text{otherwise.} \end{cases} \quad (\text{A.30})$$

Proposition 4.1 implies that the agent will find a path to the closest intermediate state $s_{i_{j_0}^*}$:

$$D(s, s_{i_{j_0}^*}) = D(s, \mathcal{S}_I), \quad (\text{A.31})$$

and

$$V_k(s) = \gamma^{d-1}B_I > \gamma^d B_I \geq V_k(s'). \quad (\text{A.32})$$

When $k \geq d + h$. To complete this theorem, we need to show that, when $k \geq d + h$, $V_k(s) > V_k(s')$ still holds. Let π be any deterministic policy and suppose under policy π , the agent starting from s' will visit the sequence of intermediate states before reaching \mathcal{S}_T :

$$\{s_{i_{j_0}^*}, s_{i_{j_1}^*}, \dots, s_{i_{j_{u^\pi}^*}}, s_{i_{j_{u^\pi+1}^*}}\}, \quad (\text{A.33})$$

where we slightly abuse the notation by assuming $s_{i_{j_{u^\pi+1}^*}} \in \mathcal{S}_T$. Hence, the discounted reward $V_k^\pi(s')$ starting from state s' following policy π satisfies

$$V_k^\pi(s') \leq \underbrace{\gamma^{d-1}B_I + \left[\sum_{l=1}^{u^\pi} \gamma^{d_l-1} \right] B_I}_{\Gamma(\pi)} + \gamma^{d_{u^\pi+1}-1}B, \quad (\text{A.34})$$

where

$$d_l = D(s', s_{i_{j_0}^*}) + \sum_{m=0}^{l-1} D(s_{i_{j_m}^*}, s_{i_{j_{m+1}^*}}) \stackrel{(i)}{\geq} d' + lh, \quad \forall l = 1, 2, \dots, u^\pi + 1, \quad (\text{A.35})$$

where inequality (i) hold because of Assumption 3.5. Hence, we know that

$$\Gamma(\pi) = \gamma^{d'-1} B_I + \left[\sum_{l=1}^{u^\pi} \gamma^{d_l-1} \right] B_I + \gamma^{d_{u^\pi+1}-1} B \stackrel{(ii)}{\leq} \underbrace{\gamma^{d'-1} \left[\sum_{l=0}^{u^\pi} \gamma^{lh} \right] B_I + \gamma^{d'+(u^\pi+1)h-1} B}_{f(u^\pi)}, \quad (\text{A.36})$$

where inequality (ii) holds because of (A.35) and $\gamma < 1$. By the synchronous value iteration update (2.1), and (A.34) the value function of any state s, s' satisfy:

$$\gamma^{d'-1} B_I \leq V_k(s), V_k(s') = \max_{\pi} V_k^\pi(s') \leq \max_{\pi} f(u^\pi), \quad (\text{A.37})$$

Next we will show $\gamma^{d'-1} B_I > f(u^\pi)$ holds $\forall \pi$, (which directly leads to $V_k(s) > V_k(s')$, for any policy from (A.37), after taking the max over all policy π) under these two following conditions: 1) $\frac{B}{B_I} \in \left(0, \frac{1}{1-\gamma^h}\right)$ when $\gamma + \gamma^h \leq 1$; 2) $\frac{B}{B_I} \in \left(\frac{1}{1-\gamma^h}, \frac{1-\gamma}{\gamma^{1+h}}\right)$ when $\gamma + \gamma^h > 1$.

- **When** $\frac{B}{B_I} \in \left(0, \frac{1}{1-\gamma^h}\right)$ **and** $\gamma + \gamma^h \leq 1$.

Since $\frac{B}{B_I} \in \left(0, \frac{1}{1-\gamma^h}\right)$, we know that $B_I - (1 - \gamma^h)B > 0$, hence

$$\begin{aligned} f(u^\pi) &< f(u^\pi) + \gamma^{d'+(u^\pi+1)h-1} [B_I - (1 - \gamma^h)B] \\ &= \gamma^{d'-1} \left[\sum_{l=0}^{u^\pi} \gamma^{lh} \right] B_I + \gamma^{d'+(u^\pi+1)h-1} B + \gamma^{d'+(u^\pi+1)h-1} [B_I - (1 - \gamma^h)B] \\ &= \gamma^{d'-1} \left[\sum_{l=0}^{u^\pi+1} \gamma^{lh} \right] B_I + \gamma^{d'+(u^\pi+2)h-1} B = f(u^\pi + 1) < \dots < f(\infty) \\ &= \gamma^{d'-1} \left[\sum_{l=0}^{\infty} \gamma^{lh} \right] B_I = \frac{\gamma^{d'-1}}{1 - \gamma^h} B_I \stackrel{(i)}{\leq} \frac{\gamma^d}{1 - \gamma^h} B_I \stackrel{(ii)}{\leq} \gamma^{d-1} B_I, \end{aligned} \quad (\text{A.38})$$

where inequality (i) holds because (A.26) implies $d' \geq d + 1$, and inequality (ii) holds because $\gamma + \gamma^h \leq 1$ implies that $\frac{\gamma}{1-\gamma^h} \leq 1$.

- **When** $\frac{B}{B_I} \in \left[\frac{1}{1-\gamma^h}, \frac{1-\gamma}{\gamma^{1+h}}\right)$ **and** $\gamma + \gamma^h > 1$.

When $\gamma + \gamma^h > 1$, we know that the interval $\left(\frac{1}{1-\gamma^h}, \frac{1-\gamma}{\gamma^{1+h}}\right)$ is well defined as $\gamma + \gamma^h > 1$ implies $\frac{1}{1-\gamma^h} < \frac{1-\gamma}{\gamma^{1+h}}$. Since $\frac{B}{B_I} > \frac{1}{1-\gamma^h}$, we know that $(1 - \gamma^h)B - B_I > 0$, hence

$$\begin{aligned}
 f(u^\pi) &< f(u^\pi) + \gamma^{d'+u^\pi h-1} \left[(1 - \gamma^h)B - B_I \right] \\
 &= \gamma^{d'-1} \left[\sum_{l=0}^{u^\pi} \gamma^{lh} \right] B_I + \gamma^{d'+(u^\pi+1)h-1} B + \gamma^{d'+u^\pi h-1} \left[(1 - \gamma^h)B - B_I \right] \\
 &= \gamma^{d'-1} \left[\sum_{l=0}^{u^\pi-1} \gamma^{lh} \right] B_I + \gamma^{d'+u^\pi h-1} B = f(u^\pi - 1) < \dots < f(0) \\
 &= \gamma^{d'-1} B_I + \gamma^{d'+h-1} B \stackrel{(i)}{\leq} \gamma^d B_I + \gamma^{d+h} B \stackrel{(ii)}{\leq} \gamma^{d-1} B_I.
 \end{aligned} \tag{A.39}$$

where inequality (i) holds because (A.26) implies $d' \geq d + 1$, and inequality (ii) holds because $\frac{B}{B_I} \leq \frac{1-\gamma}{\gamma^{1+h}}$ implies that $\gamma B_I + \gamma^{1+h} B \leq B_I$.

Hence, we conclude that $f(u^\pi) < \gamma^{d-1} B_I, \forall \pi$, which leads to the result that $V_k(s) > V_k(s'), \forall k \geq d + h$ and $\forall s, s' \in \mathcal{S}$ satisfying (A.26). This result indicates that we still have $V_k(s) > V_k(s')$ in the future update of synchronous value iteration. Therefore, we conclude that $V_k(s) > V_k(s'), \forall s, s' \in \mathcal{S}$ satisfying $D(s', \mathcal{S}_I) > D(s, \mathcal{S}_I)$, which completes the proof. ■

A.5 Proof of Theorem 4.4

Theorem A.5 (Finding the Shortest Path to \mathcal{S}_T) *Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ be a deterministic MDP with initial state s_0 , intermediate states $\mathcal{S}_I = \{s_{i_1}, s_{i_2}, \dots, s_{i_N}\}$, and terminal states \mathcal{S}_T . Suppose \mathcal{M} satisfies Assumption 3.2, 3.3 (b), and 3.5, the reward function $r(\cdot)$ follows the intermediate reward setting (Table 3), and the value function, Q-function are zero-initialized (4.1). $\forall s \in \mathcal{S} \setminus \mathcal{S}_T$, suppose \mathcal{S}_T is directly reachable from s , and let*

$$d = D(s, \mathcal{S}_T) \quad \text{and} \quad d_I = D(s, \mathcal{S}_I) \doteq \min_{j \in \mathcal{I}_d(s)} D(s, s_{i_j}). \tag{A.40}$$

If d and d_I satisfies:

$$d < \begin{cases} d_I + \log_{\frac{1}{\gamma}} \left[(1 - \gamma^h) \frac{B}{B_I} \right], & \text{if } \frac{B}{B_I} < \frac{1}{1-\gamma^h}, \\ d_I + \log_{\frac{1}{\gamma}} \left(\frac{B}{B_I + \gamma^h B} \right), & \text{if } \frac{B}{B_I} \geq \frac{1}{1-\gamma^h}, \end{cases} \quad \text{and } d < d_I + h - 1, \tag{A.41}$$

where h is the minimum distance between two intermediate states (Assumption 3.5), then after $k \geq d$ synchronous value iteration updates (2.1), an agent following the greedy policy will pursue the shortest path to \mathcal{S}_T .

Proof By the update of Q-function in the synchronous value iteration update (2.1), it suffices to show

$$Q_k(s, a^+) = r(s, a^+, s_{a^+}) + \gamma V_k(s_{a^+}) > r(s, a^-, s_{a^-}) + \gamma V_k(s_{a^-}) = Q_k(s, a^-), \tag{A.42}$$

when condition (A.41) is satisfied. Lemma B.3 implies that when $k < d_I + h$, we have

$$V_k(s) = \max\{v(k, d_I)B_I, v(k, d)B\}, \quad (\text{A.43})$$

where

$$v(k, d) = \begin{cases} \gamma^{d-1}, & \forall k, d \in \mathbb{N}^+, d \leq k, \\ 0, & \text{otherwise.} \end{cases} \quad (\text{A.44})$$

Hence we will focus on proving (A.42) in the remaining proof.

When $d \leq k < d_I + h - 1$. In this case, the conditions in (A.41) first guarantee it exists a k such that $d \leq k < d_I + h - 1$. First, we will show both conditions

$$d < d_I + \log_{\frac{1}{\gamma}} \left[(1 - \gamma^h) \frac{B}{B_I} \right] \quad \text{and} \quad d < d_I + \log_{\frac{1}{\gamma}} \left(\frac{B}{B_I + \gamma^h B} \right),$$

indicate that

$$\gamma^{d_I} B_I < \gamma^d B. \quad (\text{A.45})$$

- For $d < d_I + \log_{\frac{1}{\gamma}} \left[(1 - \gamma^h) \frac{B}{B_I} \right]$, we have

$$\begin{aligned} d - d_I < \log_{\frac{1}{\gamma}} \left[(1 - \gamma^h) \frac{B}{B_I} \right] &\iff \left(\frac{1}{\gamma} \right)^{d-d_I} < (1 - \gamma^h) \frac{B}{B_I} \\ &\iff \gamma^{d_I} B_I < (1 - \gamma^h) \gamma^d B \implies \gamma^{d_I} B_I < \gamma^d B. \end{aligned} \quad (\text{A.46})$$

- For $d < d_I + \log_{\frac{1}{\gamma}} \left(\frac{B}{B_I + \gamma^h B} \right)$, we have

$$\begin{aligned} d - d_I < \log_{\frac{1}{\gamma}} \left(\frac{B}{B_I + \gamma^h B} \right) &\iff \left(\frac{1}{\gamma} \right)^{d-d_I} < \frac{B}{B_I + \gamma^h B} \\ &\iff \gamma^{d_I} (B_I + \gamma^h B) < \gamma^d B \implies \gamma^{d_I} B_I < \gamma^d B. \end{aligned} \quad (\text{A.47})$$

Now consider a correct action $a^+ \in \mathcal{A}^+(s)$ and an incorrect action $a^- \in \mathcal{A}^-(s)$, we will next show

$$Q_k(s, a^+) = v(k, d)B, \quad \text{and} \quad Q_k(s, a^-) \leq \max\{\gamma v(k, d_I - 1)B_I, \gamma v(k, d)B\}. \quad (\text{A.48})$$

- For $Q_k(s, a^+)$, when $d = 1$, we have

$$Q_k(s, a^+) = r(s, a, s_{a^+}) + \gamma V_k(s_{a^+}) \stackrel{(i)}{=} B = v(k, 1)B, \quad (\text{A.49})$$

where equality (i) holds because $s_{a^+} \in \mathcal{S}_T$, hence $r(s, a, s_{a^+}) = B$ and $V_k(s_{a^+}) = 0$.

When $d > 1$, we have

$$Q_k(s, a^+) = r(s, a, s_{a^+}) + \gamma V_k(s_{a^+}) \stackrel{(ii)}{=} \gamma v(k, d-1)B \stackrel{(iii)}{=} v(k, d)B. \quad (\text{A.50})$$

Equality (ii) holds because $r(s, a, s_{a^+}) = 0$ and Lemma B.3 implies $V_k(s_{a^+}) = v(k, d-1)B$. Equality (iii) holds because when $k \geq d$, $\gamma v(k, d-1) = \gamma^{d-1} = v(k, d)$. Hence, we conclude that $Q_k(s, a^+) = v(k, d)B$.

- For $Q_k(s, a^-)$, consider these two following subsets of $\mathcal{A}^-(s)$: $a_{i+}^- \in \mathcal{A}_I^+(s) \cap \mathcal{A}^-(s)$ and $a_{i-}^- \in \mathcal{A} \setminus (\mathcal{A}_I^+(s) \cup \mathcal{A}^+(s))$, where $\mathcal{A}_I^+(s)$ is the set of actions that take state s one step closer to \mathcal{S}_I :

$$\mathcal{A}_I^+(s) \doteq \{a | a \in \mathcal{A}, D(s_a, \mathcal{S}_I) = D(s, \mathcal{S}_I) - 1\}, \quad (\text{A.51})$$

and $D(s, \mathcal{S}_I)$ is defined as

$$D(s, \mathcal{S}_I) \doteq \min_{j \in \mathcal{I}_d(s)} D(s, s_{i_j}) = d_I. \quad (\text{A.52})$$

We will next show that

$$\begin{aligned} Q_k(s, a_{i+}^-) &\leq \max\{\gamma v(k, d_I - 1)B_I, \gamma v(k, d)B\}, \quad \forall a_{i+}^- \in \mathcal{A}_I^+(s) \cup \mathcal{A}^-(s), \\ Q_k(s, a_{i-}^-) &\leq \max\{\gamma v(k, d_I)B_I, \gamma v(k, d)B\}, \quad \forall a_{i-}^- \in \mathcal{A} \setminus (\mathcal{A}_I^+(s) \cup \mathcal{A}^+(s)). \end{aligned} \quad (\text{A.53})$$

When $d_I = 1$, we know that $s_{a_{i+}^-} \in \mathcal{S}_I$, hence $r(s, a_{i+}^-, s_{a_{i+}^-}) = B_I$. Moreover, when $k < d_I + h$, we have $V_k(s_{a_{i+}^-}) = 0$, this is because $V_0(s_{a_{i+}^-})$ is initialized as 0 and it takes at least h update of synchronous value iteration for $V_k(s_{a_{i+}^-})$ to be positive but $k < d_I + h - 1 = h$. Hence, we know that

$$Q_k(s, a_{i+}^-) = B_I = \gamma v(k, 0)B_I \leq \max\{\gamma v(k, 0)B_I, \gamma v(k, d)B\}. \quad (\text{A.54})$$

As for $Q_k(s, a_{i-}^-)$, we have $r(s, a_{i-}^-, s_{a_{i-}^-}) = 0$ and $V_k(s_{a_{i-}^-}) = \max\{v(k, d'_I)B_I, v(k, d')B\}$, where

$$d'_I = D(s_{a_{i-}^-}, \mathcal{S}_I) \geq d_I, \quad d' = D(s_{a_{i-}^-}, \mathcal{S}_T) \geq d. \quad (\text{A.55})$$

Hence, we have

$$\begin{aligned} Q_k(s, a_{i-}^-) &= r(s, a_{i-}^-, s_{a_{i-}^-}) + \gamma V_k(s_{a_{i-}^-}) \\ &\stackrel{(i)}{=} \gamma \max\{v(k, d'_I)B_I, v(k, d')B\} \leq \max\{\gamma v(k, 1)B_I, \gamma v(k, d)B\}, \end{aligned} \quad (\text{A.56})$$

where equality (i) holds by Lemma B.3. Therefore, we conclude that (A.53) holds when $d_I = 1$.

When $d_I > 1$, we have $r(s, a_{i+}^-, s_{a_{i+}^-}) = 0$ and $r(s, a_{i-}^-, s_{a_{i-}^-}) = 0$. In this case, let $d' = D(s_{a_{i-}^-}, \mathcal{S}_T)$, then

$$\begin{aligned} Q_k(s, a_{i+}^-) &= r(s, a_{i+}^-, s_{a_{i+}^-}) + \gamma V_k(s_{a_{i+}^-}) \\ &\stackrel{(i)}{=} \gamma \max\{v(k, d_I - 1)B_I, v(k, d')B\} \leq \max\{\gamma v(k, d_I - 1)B_I, v(k, d)B\}, \end{aligned} \quad (\text{A.57})$$

where equality (i) holds by Lemma B.3. For $Q_k(s, a_{i-}^-)$, we have $r(s, a_{i-}^-, s_{a_{i-}^-}) = 0$ and $V_k(s_{a_{i-}^-}) = \max\{v(k, d'_I)B_I, v(k, d'')B\}$, where

$$d'_I = D(s_{a_{i-}^-}, \mathcal{S}_I) \geq d_I, \quad d'' = D(s_{a_{i-}^-}, \mathcal{S}_T) \geq d. \quad (\text{A.58})$$

Hence, we have

$$\begin{aligned} Q_k(s, a_{i-}^-) &= r(s, a_{i-}^-, s_{a_{i-}^-}) + \gamma V_k(s_{a_{i-}^-}) \\ &\stackrel{(i)}{=} \gamma \max\{v(k, d'_I)B_I, v(k, d'')B\} \leq \max\{\gamma v(k, d_I)B_I, \gamma v(k, d)B\}, \end{aligned} \quad (\text{A.59})$$

where equality (i) holds by Lemma B.3. Therefore, we conclude that (A.53) holds when $d_I > 1$.

A direct implication of (A.53) is that

$$\begin{aligned} Q_k(s, a^-) &< \max\{Q_k(s, a_{i+}^-), Q_k(s, a_{i-}^-)\} \\ &\leq \max\{\max\{\gamma v(k, d_I - 1)B_I, \gamma v(k, d)B\}, \max\{\gamma v(k, d_I)B_I, \gamma v(k, d)B\}\} \\ &= \max\{\gamma v(k, d_I - 1)B_I, \gamma v(k, d)B\}. \end{aligned} \quad (\text{A.60})$$

Therefore, we have shown (A.48) for $d \leq k < d_I + h - 1$. Hence, combine (A.48) and (A.45), we have

$$\begin{aligned} Q_k(s, a^-) &\stackrel{(i)}{\leq} \max\{\gamma v(k, d_I - 1)B_I, \gamma v(k, d)B\} \leq \max\{\gamma^{d_I-1}B_I, \gamma^d B\} \\ &\leq \max\{\gamma^{d_I-1}B_I, \gamma^{d-1}B\} \stackrel{(ii)}{=} \gamma^{d-1}B = v(k, d) = Q_k(s, a^+), \end{aligned} \quad (\text{A.61})$$

where inequality (i) follows (A.48) and equality (ii) holds due to (A.45). Hence, know that a greedy action will select $a^+ \in \mathcal{A}^+(s)$. If we let d', d'_I denote the distance between s_{a^+} to \mathcal{S}_T and \mathcal{S}_I , respectively, we have

$$d' = D(s_{a^+}, \mathcal{S}_T) = d - 1, \quad d'_I = D(s_{a^+}, \mathcal{S}_I) \geq d_I - 1, \quad (\text{A.62})$$

which implies d' and d'_I still satisfy (A.41). Recursively applying the above argument, we conclude that a greedy policy will find the shortest path to \mathcal{S}_T , when $d \leq k < d_I + h - 1$.

When $k \geq d_I + h - 1$. To complete the theorem, we need to show when $k \geq d_I + h - 1$, $Q_k(s, a^+) > Q_k(s, a^-)$ still holds. We have already shown that

$$Q_k(s, a^+) = v(k, d)B = \gamma^{d-1}B,$$

when $d \leq k < d_I + h - 1$, and the Q-function from the synchronous value iteration (2.1) update is non-decreasing, hence when $k \geq d_I + h - 1$, we know that $Q_k(s, a^+) \geq v(k, d) = \gamma^{d-1}B$ and it suffices to show that $\gamma^{d-1}B \geq Q_k(s, a^-)$. Let π be a deterministic policy and suppose the agent starting from s will visit the following sequences of intermediate states before reaching \mathcal{S}_T under policy π :

$$\{s_{i_{j_0}^*}, s_{i_{j_1}^*}, \dots, s_{i_{j_{u^\pi}^*}}, s_{i_{j_{u^\pi+1}^*}}\}, \quad (\text{A.63})$$

where we slightly abuse the notation by assuming $s_{i_{j_{u^\pi+1}^*}} \in \mathcal{S}_T$. By the synchronous value iteration update (2.1), the Q-function of state action pair (s, a^-) under policy π satisfies

$$Q_k^\pi(s, a^-) \leq \underbrace{\gamma^{d_I-1}B_I + \left[\sum_{l=1}^{u^\pi-1} \gamma^{d_l-1} \right] B_I}_{\Gamma(\pi)} + \gamma^{d_{u^\pi-1}}B, \quad (\text{A.64})$$

where

$$d_l = D(s, s_{i_{j_0^*}}) + \sum_{m=0}^{l-1} D(s_{i_{j_m^*}}, s_{i_{j_{m+1}^*}}) \stackrel{(i)}{\geq} d_I + lh, \quad \forall l = 1, 2, \dots, u^\pi + 1, \quad (\text{A.65})$$

where inequality (i) holds because of Assumption 3.5. Hence, we know that

$$\Gamma(\pi) = \gamma^{d_I-1} B_I + \left[\sum_{l=1}^{u^\pi-1} \gamma^{d_l-1} \right] B_I + \gamma^{d_{u^\pi-1}} B \stackrel{(ii)}{\leq} \underbrace{\gamma^{d_I-1} \left[\sum_{l=0}^{u^\pi-1} \gamma^{lh} \right] B_I + \gamma^{d_I+u^\pi h-1} B}_{f(u^\pi)}, \quad (\text{A.66})$$

where inequality (ii) holds because of (A.65) and $\gamma < 1$. By the synchronous value iteration update (2.1), the Q-function of state action pair (s, a^+) when $k \geq d_I + h - 1$ satisfies $Q_k(s, a^+) \geq \gamma^{d-1} B$, and the Q-function of state action pair (s, a^-) satisfies

$$Q_k(s, a^-) = \max_{\pi} Q_k^\pi(s, a^-) \leq \max_{\pi} f(u^\pi). \quad (\text{A.67})$$

Next we will show $\gamma^{d-1} B > f(u^\pi)$ holds $\forall \pi$ (which directly leads to $Q_k(s, a^+) > Q_k(s, a^-)$) in these following two conditions: 1) $d < d_I + \log_{\frac{1}{\gamma}} \left[(1 - \gamma^h) \frac{B}{B_I} \right]$ when $\frac{B}{B_I} \in \left(0, \frac{1}{1-\gamma^h} \right)$; 2) $d < d_I + \log_{\frac{1}{\gamma}} \left(\frac{B}{B_I + \gamma^h B} \right)$ when $\frac{B}{B_I} \geq \frac{1}{1-\gamma^h}$.

- **When** $d < d_I + \log_{\frac{1}{\gamma}} \left[(1 - \gamma^h) \frac{B}{B_I} \right]$ **and** $\frac{B}{B_I} \in \left(0, \frac{1}{1-\gamma^h} \right)$.

Since $\frac{B}{B_I} \in \left(0, \frac{1}{1-\gamma^h} \right)$, we know that $B_I - (1 - \gamma^h)B > 0$, hence

$$\begin{aligned} f(u^\pi) &< f(u^\pi) + \gamma^{d_I+u^\pi h-1} \left[B_I - (1 - \gamma^h)B \right] \\ &= \gamma^{d_I-1} \left[\sum_{l=0}^{u^\pi-1} \gamma^{lh} \right] B_I + \gamma^{d_I+u^\pi h-1} B + \gamma^{d_I+u^\pi h-1} \left[B_I - (1 - \gamma^h)B \right] \\ &= \gamma^{d_I-1} \left[\sum_{l=0}^{u^\pi} \gamma^{lh} \right] B_I + \gamma^{d_I+(u^\pi+1)h-1} B = f(u^\pi + 1) < \dots < f(\infty) \\ &= \gamma^{d_I-1} \left[\sum_{l=0}^{\infty} \gamma^{lh} \right] B_I = \frac{\gamma^{d_I-1}}{1 - \gamma^h} B_I \stackrel{(i)}{<} \gamma^{d-1} B, \end{aligned} \quad (\text{A.68})$$

where the last inequality holds because $d < d_I + \log_{\frac{1}{\gamma}} \left[(1 - \gamma^h) \frac{B}{B_I} \right]$.

- **When** $d < d_I + \log_{\frac{1}{\gamma}} \left(\frac{B}{B_I + \gamma^h B} \right)$ **and** $\frac{B}{B_I} \geq \frac{1}{1-\gamma^h}$.

When $\frac{B}{B_I} > \frac{1}{1-\gamma^h}$, we know that $(1-\gamma^h)B - B_I \geq 0$, hence

$$\begin{aligned}
 f(u^\pi) &< f(u^\pi) + \gamma^{d_I+(u^\pi-1)h-1} \left[(1-\gamma^h)B - B_I \right] \\
 &= \gamma^{d_I-1} \left[\sum_{l=0}^{u^\pi-1} \gamma^{lh} \right] B_I + \gamma^{d_I+u^\pi h-1} B + \gamma^{d_I+(u^\pi-1)h-1} \left[(1-\gamma^h)B - B_I \right] \\
 &= \gamma^{d_I-1} \left[\sum_{l=0}^{u^\pi-2} \gamma^{lh} \right] B_I + \gamma^{d_I+(u^\pi-1)h-1} B = f(u^\pi - 1) < \dots < f(1) \\
 &= \gamma^{d_I-1} B_I + \gamma^{d_I+h-1} B \stackrel{(i)}{\leq} \gamma^{d-1} B.
 \end{aligned} \tag{A.69}$$

where the last inequality holds because $d < d_I + \log_{\frac{1}{\gamma}} \left(\frac{B}{B_I + \gamma^h B} \right)$.

Hence, we conclude that $\gamma^{d-1} B > f(u^\pi), \forall \pi$. This result implies that a greedy action will select $a^+ \in \mathcal{A}^+(s)$. If we let d', d'_I denote the distance between s_{a^+} to \mathcal{S}_T and \mathcal{S}_I , respectively, we will have

$$d' = D(s_{a^+}, \mathcal{S}_T) = d - 1, \quad d'_I = D(s_{a^+}, \mathcal{S}_I) \geq d_I - 1, \tag{A.70}$$

which indicates that d' and d'_I still satisfy (A.41). Recursively applying the above argument, we conclude that a greedy policy will find the shortest path to \mathcal{S}_T , when $k \geq d_I + h - 1$. Combine this result with the case where $d \leq k < d_I + h - 1$, we have shown that the greedy policy finds the shortest path to \mathcal{S}_T , after $k(k > d)$ synchronous value iteration update. ■

Appendix B. Auxiliary Lemmas

Lemma B.1 ($V_k(s)$ with Sparse Rewards) *Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ be a deterministic MDP with desired terminal states \mathcal{S}_T and initial state s_0 . If the reward function $r(\cdot)$ satisfies the sparse reward setting (Table 3) and the Q -function is zero-initialized (4.1), $\forall d \in \mathbb{N}^+$, let \mathcal{S}^d denote the set of states whose distance to the terminal state \mathcal{S}_T is d :*

$$\mathcal{S}^d = \{s \in \mathcal{S} \mid D(s, \mathcal{S}_T) = d\}.$$

Then $\forall k \in \mathbb{N}^+, \forall s^d \in \mathcal{S}^d, V_k(s^d)$ satisfies:

$$V_k(s^d) = \begin{cases} \gamma^{d-1} B, & \forall k, d \in \mathbb{N}^+, d \leq k, \\ 0, & \text{otherwise.} \end{cases} \tag{B.1}$$

Proof We will use induction to prove the results. Recall the synchronous value iteration update (2.1):

$$V_{k+1}(s) = \max_{a \in \mathcal{A}} \{r(s, a, s_a) + \gamma V_k(s_a)\}, \quad \forall s \in \mathcal{S}.$$

We first check the induction condition (B.1) for the initial case $k = 1$.

Initial Condition for $k = 1$. Since the agent only receives reward $r(s, a, s_a) = B$ when $s_a \in \mathcal{S}_T$, by the value iteration update, we have:

$$V_1(s) = \begin{cases} B, & \text{if } s \in \mathcal{S}^1, \\ 0, & \text{otherwise,} \end{cases} \quad (\text{B.2})$$

hence, the initial condition is verified.

Induction. Next, suppose the condition (B.1) holds for $1, 2, \dots, k$, we will show it also holds for $k + 1$. In this case, we only need to verify $V_{k+1}(s^d) = \gamma^k B$ for $d = k + 1$, because the induction assumption already implies

$$V_{k+1}(s^d) = \begin{cases} \gamma^{d-1} B, & \forall k, d \in \mathbb{N}^+, d \leq k, \\ 0, & \text{otherwise.} \end{cases}$$

Note that when $d = k + 1 > 1$, we have

$$r(s^{k+1}, a, s_a^{k+1}) = 0, \forall a \in \mathcal{A}, \text{ and } V_k(s_a^{k+1}) \stackrel{(i)}{\leq} V_k(s^k) = \gamma^{k-1} B, \quad (\text{B.3})$$

where s_a^{k+1} is the subsequent state of state-action pair (s^{k+1}, a) , and inequality (i) becomes equality if only if a is a correct action $a \in \mathcal{A}^+(s)$ (see Definition A.1). Hence, we know that

$$V_{k+1}(s^{k+1}) = \max_{a \in \mathcal{A}} \left\{ r(s^{k+1}, a, s_a) + \gamma V_k(s_a^{k+1}) \right\} = \gamma^k B. \quad (\text{B.4})$$

As a result, we conclude that

$$V_{k+1}(s^d) = \begin{cases} \gamma^{d-1} B, & \forall k, d \in \mathbb{N}^+, d \leq k + 1, \\ 0, & \text{otherwise,} \end{cases}$$

which completes the proof. ■

Lemma B.2 ($V_k(s)$ with Intermediate Rewards Setting (a)) Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ be a deterministic MDP with initial state s_0 , intermediate states $\mathcal{S}_I = \{s_{i_1}, s_{i_2}, \dots, s_{i_N}\}$, and terminal states \mathcal{S}_T . Suppose \mathcal{M} satisfies Assumption 3.2, and 3.3 (a), if the reward function $r(\cdot)$ follows the intermediate reward setting (Table 3) and the Q -function is zero-initialized (4.1), $\forall j \in [N + 1]$, let $\mathcal{S}_{[j-1, j]}$ denote the set of states such that

$$\mathcal{S}_{[j-1, j]} = \begin{cases} \{s \in \mathcal{S} | D(s, s_{i_1}) < \infty\} & \text{if } j = 1, \\ \{s \in \mathcal{S} | D(s, s_{i_j}) = \infty, D(s, s_{i_{j+1}}) < \infty\} & \text{if } j = 2, 3, \dots, N, \\ \{s \in \mathcal{S} | D(s, s_{i_N}) = \infty, D(s, \mathcal{S}_T) < \infty\} & \text{if } j = N + 1. \end{cases} \quad (\text{B.5})$$

Given a state $s^{\mathbf{d}} \in \mathcal{S}_{[j-1, j]}^{\mathbf{d}}$, where $\mathbf{d} = [d_j, d_{j+1}, \dots, d_N, d_{N+1}]^\top \in \mathbb{R}^{n-j+2}$ is a vector, such that $d_j < d_{j+1} < \dots < d_n < d_{N+1}$ denote the distance from $s^{\mathbf{d}}$ to $s_{i_j}, s_{i_{j+1}}, \dots, s_{i_N}, \mathcal{S}_T$, respectively:

$$d_l = D(s^{\mathbf{d}}, s_{i_l}), \forall l = j, j + 1, \dots, n, \text{ and } d_{N+1} = D(s^{\mathbf{d}}, \mathcal{S}_T). \quad (\text{B.6})$$

Then $\forall k \in \mathbb{N}^+, \forall s^{\mathbf{d}} \in \mathcal{S}^{\mathbf{d}}, V_k(s^{\mathbf{d}})$ satisfies the following conditions:

$$V_k(s^{\mathbf{d}}) = \sum_{l=j}^{N+1} v(k, d_l) B_l, \quad (\text{B.7})$$

where

$$v(k, d_l) = \begin{cases} \gamma^{d_l-1}, & \forall k, d \in \mathbb{N}^+, d_l \leq k, \forall l \in [N+1], \\ 0, & \text{otherwise,} \end{cases} \quad \text{and } B_l = \begin{cases} B_I, & \text{if } l \in [N], \\ B, & \text{if } l = N+1. \end{cases} \quad (\text{B.8})$$

Proof We will use induction to prove the results. Recall the synchronous value iteration update (2.1):

$$V_{k+1}(s) = \max_{a \in \mathcal{A}} \{r(s, a, s_a) + \gamma V_k(s_a)\}, \quad \forall s \in \mathcal{S}.$$

We first check the induction condition (B.7) for the initial case $k = 1$.

Initial Condition for $k = 1$. When $k = 1$, the value function of states s before reaching s_{i_j} will not be affected by the rewards from $s_{i_{j+1}}, s_{i_{j+2}}, \dots, s_{i_N}, \mathcal{S}_T$ and will only be updated using the intermediate reward from s_{i_j} . As similarly proven for the sparse reward case (Lemma B.1), we have

$$V_1(s) = \begin{cases} B_I, & \text{if } D(s, s_{i_l}) = 1, \forall l \in [N+1], \\ 0, & \text{otherwise,} \end{cases} \quad (\text{B.9})$$

which implies $V_1(s^{\mathbf{d}}) = \sum_{l=j}^{N+1} v(1, d_l) B_l$, hence, the initial condition is verified.

Induction. Next, suppose the conditions (B.7) hold for $1, 2, \dots, k$, we will show they also hold for $k+1$. By the induction assumption, we know that

$$V_k(s^{\mathbf{d}}) = \sum_{l=j}^{N+1} v(k, d_l) B_l,$$

holds for $1, 2, \dots, k$, therefore we only need to show the induction condition (B.7) holds when it exists $l' \in \{j, j+1, \dots, N+1\}$, such that $d_{l'} = k+1$, because $v(k, d_{l'})$ remains the same, $\forall d_{l'} \neq k+1$ when we change k to $k+1$. Suppose $s^{\mathbf{d}} \in \mathcal{S}_{[j-1, j]}$ satisfies $D(s^{\mathbf{d}}, s_{i_{l'}}) = k+1$, since $d_j < d_{j+1} < \dots < d_{l'} = k+1$, we will discuss the following cases.

- **When $j < N+1$ and $l' < N+1$.** For a correct action $a^+ \in \mathcal{A}^+(s^{\mathbf{d}})$, we have

$$\begin{aligned} & r(s^{\mathbf{d}}, a^+, s_{a^+}^{\mathbf{d}}) + \gamma V_k(s_{a^+}^{\mathbf{d}}) \\ &= \begin{cases} B_I + \gamma V_k(s_{a^+}^{\mathbf{d}}) = B_I + \gamma \sum_{l=j+1}^{l'-1} \gamma^{d_l-2} B_I + \gamma \cdot \gamma^{k-1} B, & \text{if } d_j = 1, \\ \gamma V_k(s_{a^+}^{\mathbf{d}}) = \gamma \sum_{l=j}^{l'-1} \gamma^{d_l-2} B_I + \gamma \cdot \gamma^{k-1} B, & \text{if } d_j > 1. \end{cases} \quad (\text{B.10}) \\ &= \sum_{l=j}^{l'-1} \gamma^{d_l-1} B_I + \gamma^k B_I = \sum_{l=j}^{l'} v(k+1, d_l) B_I = \sum_{l=j}^{N+1} v(k+1, d_l) B_l, \end{aligned}$$

where the last equality holds because of (B.8). For an incorrect action $a^- \in \mathcal{A}^-(s^{\mathbf{d}})$ we have

$$r(s^{\mathbf{d}}, a^-, s_{a^-}^{\mathbf{d}}) + \gamma V_k(s_{a^-}^{\mathbf{d}}) \stackrel{(i)}{=} \gamma V_k(s_{a^-}^{\mathbf{d}}) \stackrel{(ii)}{\leq} \sum_{l=j}^{l'-1} \gamma^{d_l} B_l, \quad (\text{B.11})$$

where equality (i) holds because $r(s^{\mathbf{d}}, a^-, s_{a^-}^{\mathbf{d}}) = 0$ and inequality (ii) holds because of $\gamma < 1$. Since $k + 1 > 1$, we know that $r(s^{\mathbf{d}}, a, s_a^{\mathbf{d}}) = 0$ and hence

$$r(s^{\mathbf{d}}, a^-, s_{a^-}^{\mathbf{d}}) + \gamma V_k(s_{a^-}^{\mathbf{d}}) < r(s^{\mathbf{d}}, a^+, s_{a^+}^{\mathbf{d}}) + \gamma V_k(s_{a^+}^{\mathbf{d}}). \quad (\text{B.12})$$

By the value iteration update (2.1), we know that

$$\begin{aligned} V_{k+1}(s) &= \max_{a \in \mathcal{A}} \{r(s, a, s_a) + \gamma V_k(s_a)\} \\ &= r(s^{\mathbf{d}}, a^+, s_{a^+}^{\mathbf{d}}) + \gamma V_k(s_{a^+}^{\mathbf{d}}) = \sum_{l=j}^{N+1} v(k+1, d_l) B_l, \end{aligned} \quad (\text{B.13})$$

which completes the induction for this case.

- **When $j < N + 1$ and $l' = N + 1$.** For a correct action $a^+ \in \mathcal{A}^+(s^{\mathbf{d}})$, we have

$$\begin{aligned} & r(s^{\mathbf{d}}, a^+, s_{a^+}^{\mathbf{d}}) + \gamma V_k(s_{a^+}^{\mathbf{d}}) \\ &= \begin{cases} B_I + \gamma V_k(s_{a^+}^{\mathbf{d}}) = B_I + \gamma \sum_{l=j+1}^N \gamma^{d_l-2} B_l + \gamma \cdot \gamma^{k-1} B, & \text{if } d_j = 1, \\ \gamma V_k(s_{a^+}^{\mathbf{d}}) = \gamma \sum_{l=j}^N \gamma^{d_l-2} B_l + \gamma \cdot \gamma^{k-1} B, & \text{if } d_j > 1. \end{cases} \quad (\text{B.14}) \\ &= \sum_{l=j}^N \gamma^{d_l-1} B_l + \gamma^k B = \sum_{l=j}^{N+1} v(k+1, d_l) B_l. \end{aligned}$$

For an incorrect action $a^- \in \mathcal{A}^-(s^{\mathbf{d}})$ we have

$$r(s^{\mathbf{d}}, a^-, s_{a^-}^{\mathbf{d}}) + \gamma V_k(s_{a^-}^{\mathbf{d}}) = \gamma V_k(s_{a^-}^{\mathbf{d}}) \leq \sum_{l=j}^N \gamma^{d_l} B_l. \quad (\text{B.15})$$

Since $k + 1 > 1$, we know that $r(s^{\mathbf{d}}, a, s_a^{\mathbf{d}}) = 0$ and hence

$$r(s^{\mathbf{d}}, a^-, s_{a^-}^{\mathbf{d}}) + \gamma V_k(s_{a^-}^{\mathbf{d}}) < r(s^{\mathbf{d}}, a^+, s_{a^+}^{\mathbf{d}}) + \gamma V_k(s_{a^+}^{\mathbf{d}}). \quad (\text{B.16})$$

By the value iteration update (2.1), we know that

$$\begin{aligned} V_{k+1}(s) &= \max_{a \in \mathcal{A}} \{r(s, a, s_a) + \gamma V_k(s_a)\} \\ &= r(s^{\mathbf{d}}, a^+, s_{a^+}^{\mathbf{d}}) + \gamma V_k(s_{a^+}^{\mathbf{d}}) = \sum_{l=j}^{N+1} v(k+1, d_l) B_l, \end{aligned} \quad (\text{B.17})$$

which completes the induction for this case.

- **When $j = N + 1$ and $l' = N + 1$.** For a correct action $a^+ \in \mathcal{A}^+(s^d)$, we have

$$r(s^d, a^+, s_{a^+}^d) + \gamma V_k(s_{a^+}^d) = \gamma^k B = v(k + 1, d_l) B_l. \quad (\text{B.18})$$

For an incorrect action $a^- \in \mathcal{A}^-(s^d)$ we have

$$r(s^d, a^-, s_{a^-}^d) + \gamma V_k(s_{a^-}^d) \stackrel{(i)}{=} 0, \quad (\text{B.19})$$

where equality (i) holds because Definition A.1 implies that $D(s_{a^-}^d, \mathcal{S}_T) \geq D(s^d, \mathcal{S}_T) = k + 1$. Since $k + 1 > 1$, we know that $r(s^d, a^-, s_{a^-}^d) = 0$ and hence

$$r(s^d, a^-, s_{a^-}^d) + \gamma V_k(s_{a^-}^d) < r(s^d, a^+, s_{a^+}^d) + \gamma V_k(s_{a^+}^d). \quad (\text{B.20})$$

By the value iteration update (2.1), we know that

$$\begin{aligned} V_{k+1}(s) &= \max_{a \in \mathcal{A}} \{r(s, a, s_a) + \gamma V_k(s_a)\} \\ &= r(s^d, a^+, s_{a^+}^d) + \gamma V_k(s_{a^+}^d) = v(k + 1, d_{N+1}) B, \end{aligned} \quad (\text{B.21})$$

which completes the induction for this case.

Hence, we know the induction condition (B.7) holds for $k + 1$, which completes the proof. \blacksquare

Lemma B.3 ($V_k(s)$ when \mathcal{S}_T is Directly Reachable) *Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ be a deterministic MDP with initial state s_0 , intermediate states $\mathcal{S}_I = \{s_{i_1}, s_{i_2}, \dots, s_{i_N}\}$, and terminal states \mathcal{S}_T . Suppose \mathcal{M} satisfies Assumption 3.2, 3.3 (b), and 3.5, if the reward function $r(\cdot)$ follows the intermediate reward setting (Table 3) and the value function and Q-function are zero-initialized (4.1), $\forall s \in \mathcal{S} \setminus \mathcal{S}_T$, if \mathcal{S}_T is directly reachable from s and suppose*

$$d = D(s, \mathcal{S}_T) \quad \text{and} \quad d_I = D(s, \mathcal{S}_I) \doteq \min_{j \in \mathcal{I}_d(s)} D(s, s_{i_j}), \quad (\text{B.22})$$

then after $k \leq d_I + h$ synchronous value iteration updates (2.1), the value function $V_k(s)$ satisfies

$$V_k(s) = \max\{v(k, d_I) B_I, v(k, d) B\}, \quad (\text{B.23})$$

where

$$v(k, d) = \begin{cases} \gamma^{d-1}, & \forall k, d \in \mathbb{N}, d \leq k, \\ 0, & \text{otherwise.} \end{cases} \quad (\text{B.24})$$

Proof Similar to the definition of $\mathcal{A}^+(s)$, let $\mathcal{A}_I^+(s)$ denotes the set of actions that lead s one step closer to \mathcal{S}_I :

$$\mathcal{A}_I^+(s) = \{a | a \in \mathcal{A}, D(s_a, \mathcal{S}_I) = D(s, \mathcal{S}_I) - 1\}. \quad (\text{B.25})$$

We will use induction to prove the results. Recall the synchronous value iteration update (2.1):

$$V_{k+1}(s) = \max_{a \in \mathcal{A}} \{r(s, a, s_a) + \gamma V_k(s_a)\}, \quad \forall s \in \mathcal{S}.$$

We first check the induction condition (B.23) for the initial case $k = 1$.

Initial Condition for $k = 1$. When $k = 1$, we know that

$$\begin{aligned} V_1(s) &= \max_a \{r(s, a, s_a) + \gamma V_k(s_a)\} \\ &= \begin{cases} \max\{B_I, B\}, & \text{if } d_I = d = 1, \\ B_I, & \text{if } d_I = 1, d > 1, \\ B, & \text{if } d = 1, d_I > 1, \\ 0, & \text{otherwise} \end{cases} \\ &= \max\{v(1, d_I)B_I, v(1, d)B\}, \end{aligned} \tag{B.26}$$

which verifies the initial condition (B.23).

Induction. We will discuss the following cases: (a) $d_I \leq k + 1, d = 1$, (b) $d_I \leq k + 1, d > 1$, and (c) $d_I > k + 1$.

(a) **When $d_I \leq k + 1, d = 1$.** In this case, we consider these following actions:

$$a^+ \in \mathcal{A}^+(s), a_{i_+}^- \in \mathcal{A}_I^+(s) \setminus \mathcal{A}^+(s), a_{i_-}^- \in \mathcal{A}(s) \setminus (\mathcal{A}_I^+(s) \cup \mathcal{A}^+(s)). \tag{B.27}$$

- **For any $a^+ \in \mathcal{A}^+(s)$.** In this case, $d = 1$ implies that $s_{a^+} \in \mathcal{S}_T$. Hence we have

$$r(s, a^+, s_{a^+}) = B, V_k(s_{a^+}) = 0 \implies r(s, a^+, s_{a^+}) + \gamma V_k(s_{a^+}) = B, \tag{B.28}$$

- **For any $a_{i_+}^- \in \mathcal{A}_I^+(s) \setminus \mathcal{A}^+(s)$.** In this case, let

$$d' = D(s_{a_{i_+}^-}, \mathcal{S}_T), \tag{B.29}$$

by the definition of $\mathcal{A}_I^+(s)$ (B.25), we know that

$$d' = D(s_{a_{i_+}^-}, \mathcal{S}_T) \geq d = 1, D(s_{a_{i_+}^-}, \mathcal{S}_I) = D(s, \mathcal{S}_I) - 1 = d_I - 1 \leq k. \tag{B.30}$$

Hence, the induction condition (B.23) implies that

$$V_k(s_{a_{i_+}^-}) = \max\{v(k, d_I - 1)B_I, v(k, d')B\}. \tag{B.31}$$

If $d_I > 1$, we know that $r(s, a, s_{a_{i_+}^-}) = 0$, and hence

$$\begin{aligned} &r(s, a_{i_+}^-, s_{a_{i_+}^-}) + \gamma V_k(s_{a_{i_+}^-}) \\ &= \gamma \max\{v(k, d_I - 1)B_I, v(k, d')B\} \stackrel{(i)}{=} \max\{v(k + 1, d_I)B_I, \gamma v(k, d')B\}, \end{aligned} \tag{B.32}$$

where equality (i) holds because $\gamma v(k, d_I - 1) = \gamma v(k, k) = \gamma^k = v(k + 1, k + 1) = v(k + 1, d_I)$. When $d_I = 1$, we have

$$r(s, a_{i_+}^-, s_{a_{i_+}^-}) + \gamma V_k(s_{a_{i_+}^-}) = B_I. \tag{B.33}$$

- **For any** $a_{i-}^- \in \mathcal{A} \setminus (\mathcal{A}_I^+(s) \cup \mathcal{A}^+(s))$. In this case, let

$$d'' = D(s_{a_{i+}^-}, \mathcal{S}_T) \geq d = 1, \quad d'_I = D(s_{a_{i+}^-}, \mathcal{S}_I) \geq d_I = k + 1, \quad (\text{B.34})$$

the induction condition (B.23) implies that

$$V_k(s_{a_{i+}^-}) = \max\{v(k, d'_I)B_I, v(k, d'')B\}. \quad (\text{B.35})$$

Similar to the previous case for a_{i+}^- , we also have $r(s, a, s_{a_{i+}^-}) = 0$, and hence

$$r(s, a_{i-}^-, s_{a_{i-}^-}) + \gamma V_k(s_{a_{i-}^-}) = \max\{\gamma v(k, d'_I)B_I, \gamma v(k, d'')B\}. \quad (\text{B.36})$$

Combine (B.28), (B.32), (B.36), and the synchronous value iteration update (2.1), when $d_I > 1$ we have

$$\begin{aligned} V_{k+1}(s) &= \max_{a \in \mathcal{A}} \{r(s, a, s_a) + \gamma V_k(s_a)\} \\ &= \max \{B, \max\{v(k+1, d_I)B_I, \gamma v(k, d')B\}, \max\{\gamma v(k, d'_I)B_I, \gamma v(k, d'')B\}\} \\ &\stackrel{(i)}{=} \max\{v(k+1, d_I)B_I, B\} = \max\{v(k+1, d_I)B_I, v(k+1, 1)B\}, \end{aligned} \quad (\text{B.37})$$

where equality (i) holds because $d', d'' \geq 1$ implies that $\gamma v(k, d')B, \gamma v(k, d'')B \leq \gamma B < B$.

Similarly, using (B.28), (B.33), (B.36) when $d_I = 1$, we have

$$\begin{aligned} V_{k+1}(s) &= \max_{a \in \mathcal{A}} \{r(s, a, s_a) + \gamma V_k(s_a)\} \\ &= \max \{B, B_I, \max\{\gamma v(k, d'_I)B_I, \gamma v(k, d'')B\}\} \\ &\stackrel{(ii)}{=} \max\{B_I, B\} = \max\{v(k+1, 1)B_I, v(k+1, 1)B\}, \end{aligned} \quad (\text{B.38})$$

where equality (ii) holds because $d'_I, d'' > 1$. Hence, we have verified the induction condition (B.23) for $d_I \leq k + 1, d = 1$.

(b) When $d_I \leq k + 1, d > 1$. In this case, we also consider these following actions:

$$a^+ \in \mathcal{A}^+(s), \quad a_{i+}^- \in \mathcal{A}_I^+(s) \setminus \mathcal{A}^+(s), \quad a_{i-}^- \in \mathcal{A}(s) \setminus (\mathcal{A}_I^+(s) \cup \mathcal{A}^+(s)). \quad (\text{B.39})$$

- **For any** $a^+ \in \mathcal{A}^+(s)$. Similarly, let

$$d'_I = D(s_{a^+}, \mathcal{S}_T), \quad (\text{B.40})$$

hence we know that

$$D(s_{a^+}, \mathcal{S}_T) = d - 1 \geq 1, \quad d'_I = D(s_{a^+}, \mathcal{S}_I) \geq D(s, \mathcal{S}_I) - 1 \leq k. \quad (\text{B.41})$$

The induction condition (B.23) implies that

$$V_k(s_{a^+}) = \max\{v(k, d'_I)B_I, v(k, d-1)B\}. \quad (\text{B.42})$$

Notice that $d > 1$ implies that $r(s, a^+, s_{a^+}) = 0$, hence

$$r(s, a^+, s_{a^+}) + \gamma V_k(s_{a^+}) = \max\{\gamma v(k, d'_I)B_I, \gamma v(k, d-1)B\}. \quad (\text{B.43})$$

- **For any** $a_{i+}^- \in \mathcal{A}_I^+(s) \setminus \mathcal{A}^+(s)$. In this case, let

$$d' = D(s_{a_{i+}^-}, \mathcal{S}_T), \quad (\text{B.44})$$

by the definition of $\mathcal{A}_I^+(s)$ (B.25), we know that

$$d' = D(s_{a_{i+}^-}, \mathcal{S}_T) \geq d > 1, \quad D(s_{a_{i+}^-}, \mathcal{S}_I) = D(s, \mathcal{S}_I) - 1 = k. \quad (\text{B.45})$$

Hence, the induction condition (B.23) implies that

$$V_k(s_{a_{i+}^-}) = \max\{v(k, d_I - 1)B_I, v(k, d')B\}. \quad (\text{B.46})$$

If $d_I > 1$, so $r(s, a, s_{a_{i+}^-}) = 0$, and hence

$$\begin{aligned} & r(s, a_{i+}^-, s_{a_{i+}^-}) + \gamma V_k(s_{a_{i+}^-}) \\ &= \gamma \max\{v(k, d_I - 1)B_I, v(k, d')B\} = \max\{v(k+1, d_I)B_I, \gamma v(k, d')B\}, \end{aligned} \quad (\text{B.47})$$

the last equality holds for the same reason described in the previous case for $d_I > 1, d > 1$. When $d_I = 1$, we have

$$r(s, a_{i+}^-, s_{a_{i+}^-}) + \gamma V_k(s_{a_{i+}^-}) = B_I. \quad (\text{B.48})$$

- **For any** $a_{i-}^- \in \mathcal{A} \setminus (\mathcal{A}_I^+(s) \cup \mathcal{A}^+(s))$. Similar to the case where $d_I = k+1, d = 1$, let

$$d'' = D(s_{a_{i-}^-}, \mathcal{S}_T) \geq d > 1, \quad d'_I = D(s_{a_{i-}^-}, \mathcal{S}_I) \geq d_I = k+1, \quad (\text{B.49})$$

the induction condition (B.23) implies that

$$V_k(s_{a_{i-}^-}) = \max\{v(k, d'_I)B_I, v(k, d'')B\}. \quad (\text{B.50})$$

Similar to the previous case for a_{i+}^- , we also have $r(s, a, s_{a_{i-}^-}) = 0$, and hence

$$r(s, a_{i-}^-, s_{a_{i-}^-}) + \gamma V_k(s_{a_{i-}^-}) = \max\{\gamma v(k, d'_I)B_I, \gamma v(k, d'')B\}. \quad (\text{B.51})$$

Combine (B.43), (B.47), (B.51), and the synchronous value iteration update (2.1), when $d_I > 1$ we have

$$\begin{aligned} V_{k+1}(s) &= \max_{a \in \mathcal{A}} \{r(s, a, s_a) + \gamma V_k(s_a)\} \\ &= \max\{\max\{\gamma v(k, d'_I)B_I, \gamma v(k, d-1)B\}, \max\{v(k+1, d_I)B_I, \gamma v(k, d')B\}, \\ &\quad \max\{\gamma v(k, d'_I)B_I, \gamma v(k, d'')B\}\} \\ &\stackrel{(i)}{=} \max\{v(k+1, d_I)B_I, \gamma v(k, d-1)B\} \stackrel{(ii)}{=} \max\{v(k+1, d_I)B_I, v(k+1, d)B\}, \end{aligned} \quad (\text{B.52})$$

where equality (i) holds because $d'_I, d''_I \geq d_I$ implies that $\gamma v(k, d'_I), \gamma v(k, d''_I) \leq \gamma^{d_I} = v(k+1, d_I)$ and $d', d'' > d$ implies that $v(k, d'), v(k, d'') \leq v(k, d-1)$. Equality (ii) holds because when $d_I = k+1$, $v(k+1, d_I) = \gamma^{d_I} > 0$, then if $\gamma v(k, d-1)B > v(k+1, d_I)B_I$,

we know that $\gamma v(k, d-1)B > 0$, which implies $k \geq d-1$ (or equivalently $k+1 \geq d$) hence $v(k+1, d) = \gamma^{d-1} = \gamma v(k, d-1)$.

Similarly, when $d_I = 1$, (B.43), (B.48), (B.51), we have

$$\begin{aligned} V_{k+1}(s) &= \max_{a \in \mathcal{A}} \{r(s, a, s_a) + \gamma V_k(s_a)\} \\ &= \max \{ \max\{\gamma v(k, d'_I)B_I, \gamma v(k, d-1)B\}, B_I, \max\{\gamma v(k, d''_I)B_I, \gamma v(k, d''_I)B\} \} \quad (\text{B.53}) \\ &\stackrel{(iii)}{=} \max\{B_I, \gamma v(k, d-1)B\} = \max\{v(k+1, 1)B_I, v(k+1, d)B\}, \end{aligned}$$

where equality (iii) holds because $d'_I, d''_I \geq 1$. Hence, we have verified the induction condition (B.23) for $d_I \leq k+1, d > 1$.

(c) When $d_I > k+1$. When $d_I > k+1$, we have $\forall a \in \mathcal{A}$, we have $d'_I = D(s_a, \mathcal{S}_I) > k$. By the definition of $v(k, d)$ in (B.24), we know $v(k, d'_I) = 0$. By the induction assumption (B.23), we have

$$V_k(s_a) = \max\{v(k, d'_I)B_I, v(k, d')B\} = v(k, d')B, \quad (\text{B.54})$$

where $d' = D(s_a, \mathcal{S}_T)$. Hence, from the synchronous value iteration (2.1), we have

$$V_{k+1}(s) = \max_{a \in \mathcal{A}} \{r(s, a, s_a) + \gamma V_k(s_a)\} \stackrel{(i)}{=} v(k, d')B = \max\{v(k+1, d'_I)B_I, v(k+1, d')B\}, \quad (\text{B.55})$$

where equality (i) holds by Lemma B.1. Hence we conclude the induction condition (B.23) holds when $d > k+1$.

Conclusion. In conclusion, we have verified the induction condition (B.23) for all three cases: (a) $d_I = k+1, d = 1$, (b) $d_I = k+1, d > 1$, (c) $d_I > k+1$. Hence, we conclude that the induction condition (B.23) holds for all $k < d_I + h$. \blacksquare

Appendix C. Experimental Details

C.1 Experimental Details of Section 1.1

To obtain the win rate in 1, for each experiment, we first run asynchronous Q learning for m episodes, where we use the first $m-100$ episodes for training, and the win rate from the last 100 episodes for testing. We repeat the previous experiment 10 times report the average win rate of all 1000 games. In addition to the reward design specified Table 1, we use the following command: `python pacman.py` under the folder `./cs188/cs188/p3_reinforcement`, with these configurations:

- `-p PacmanQAgent`, set the training algorithm to standard asynchronous Q learning;
- `-x m -n m + 100`, set the training and testing episodes, replace m with the number of episode in Table 1;
- `-l mediumGrid`, set the layout of the game, we use the `mediumGrid` layout for our experiments;
- `-g DirectionalGhost`, set the strategy of the ghost, we opt the `DirectionalGhost`, where the ghost will find the shortest path to the Pacman;

- `-k 1`, set the number of ghosts, the default number of ghosts is 2 in the `mediumGrid` layout.

Similar phenomena reported in Table 1 also appear in larger layout with more ghosts. Also note that in order to obtain a similar win rate, experiments with larger layouts/more ghosts generally require more training episodes.

C.2 Experimental Details of Section 6

C.2.1 ASYNCHRONOUS Q-LEARNING IN SECTION 6

All models for asynchronous q-learning experiments are trained with a learning rate of 0.1, a discount factor of 0.9, and use an 0.8 ϵ -greedy exploration strategy during training. Each state is the entire environment encoded as a string available for MiniGrid (Chevalier-Boisvert et al., 2018) environments.

C.2.2 DEEP RL HYPERPARAMETERS IN SECTION 6

Each training session for deep RL algorithms was run using a GeForce RTX 2080 GPU. Shared parameters are listed in Table 8, and parameters specific to each algorithm is provided in Table 9. For DQN, like asynchronous Q-learning, we use a 0.8 ϵ -greedy exploration strategy.

1	Conv2D(inchannels = 3, outchannels = 16, stride = (2,2))
2	ReLU
3	MaxPool2D(2, 2)
4	Conv2D(inchannels = 16, outchannels = 32, stride = (2,2))
5	ReLU
6	Conv2D(inchannels = 32, outchannels = 64, stride = (2,2))
7	ReLU
8	Linear(embedded size, 64)
9	Tanh
10	Linear(64, number of actions)

Table 7: Network architecture

Parameters	Values
Learning Rate	0.001
Network Architecture	See Table 7
Observability of Env	Fully Observable

Table 8: Parameters

Parameters	DQN	PPO	A2C
optimizer	RMSProp	Adam	RMSProp
Discount Factor (γ)	0.90 (Maze, 3-Door) 0.80 (4-Door)	0.90	0.90
batch size	128	128	N/A
buffer size	100000	N/A	N/A
target net. update interval (steps)	100	N/A	N/A
number of actors	N/A	10	10
steps per actor before update	N/A	128	5
entropy coeff.	N/A	0.01	0.01
value loss coeff.	N/A	0.5	0.5
GAE discount (λ)	N/A	0.95	0.95
max norm of gradient	N/A	0.5	0.5
clipping ϵ	N/A	0.2	N/A
PPO epochs per update	N/A	4	N/A

Table 9: Algorithm Specific Parameters

References

- Andreas, J., Klein, D., & Levine, S. (2017). Modular multitask reinforcement learning with policy sketches. In *International Conference on Machine Learning*, pp. 166–175. PMLR.
- Andrychowicz, M., Wolski, F., Ray, A., Schneider, J., Fong, R., Welinder, P., McGrew, B., Tobin, J., Abbeel, P., & Zaremba, W. (2017). Hindsight experience replay. *arXiv preprint arXiv:1707.01495*.
- Barto, A. G., & Mahadevan, S. (2003). Recent advances in hierarchical reinforcement learning. *Discrete event dynamic systems*, 13(1), 41–77.
- Berner, C., Brockman, G., Chan, B., Cheung, V., Debiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., et al. (2019). Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*.
- Bertsekas, D. P. (1995). *Dynamic Programming and Optimal Control*. Athena Scientific.
- Bertsekas, D. P. (2019). *Reinforcement Learning and Optimal Control*. Athena Scientific.
- Bertsekas, D. P., & Tsitsiklis, J. N. (1996). *Neuro-dynamic programming*. Athena Scientific.
- Boutillier, C., Dean, T., & Hanks, S. (1999). Decision-theoretic planning: Structural assumptions and computational leverage. *Journal of Artificial Intelligence Research*, 11, 1–94.

- Boutilier, C., Dearden, R., & Goldszmidt, M. (2000). Stochastic dynamic programming with factored representations. *Artificial intelligence*, 121(1-2), 49–107.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., & Zaremba, W. (2016). Openai gym.
- Chebatar, Y., Hausman, K., Lu, Y., Xiao, T., Kalashnikov, D., Varley, J., Irpan, A., Eysenbach, B., Julian, R., Finn, C., et al. (2021). Actionable models: Unsupervised offline reinforcement learning of robotic skills. *arXiv preprint arXiv:2104.07749*.
- Chevalier-Boisvert, M., Willems, L., & Pal, S. (2018). Minimalistic gridworld environment for openai gym. <https://github.com/maximecb/gym-minigrid>.
- Dayan, P., & Hinton, G. E. (1993). Feudal reinforcement learning. In Hanson, S., Cowan, J., & Giles, C. (Eds.), *Advances in Neural Information Processing Systems*, Vol. 5. Morgan-Kaufmann.
- Dean, T., & Lin, S.-H. (1995). Decomposition techniques for planning in stochastic domains. In *IJCAI*, Vol. 2, p. 3. Citeseer.
- Dietterich, T. G. (2000). Hierarchical reinforcement learning with the maxq value function decomposition. *Journal of artificial intelligence research*, 13, 227–303.
- Doroudi, S., Alevan, V., & Brunskill, E. (2019). Wheres the reward?. *International Journal of Artificial Intelligence in Education*, 29(4), 568–620.
- Eysenbach, B., Geng, X., Levine, S., & Salakhutdinov, R. (2020a). Rewriting history with inverse rl: Hindsight inference for policy improvement. *arXiv preprint arXiv:2002.11089*.
- Eysenbach, B., Salakhutdinov, R., & Levine, S. (2020b). C-learning: Learning to achieve goals via recursive classification. *arXiv preprint arXiv:2011.08909*.
- Forejt, V., Kwiatkowska, M., Norman, G., & Parker, D. (2011). Automated verification techniques for probabilistic systems. In *International School on Formal Methods for the Design of Computer, Communication and Software Systems*, pp. 53–113. Springer.
- Fu, J., Singh, A., Ghosh, D., Yang, L., & Levine, S. (2018). Variational inverse control with events: A general framework for data-driven reward definition. *arXiv preprint arXiv:1805.11686*.
- Fuchs, F., Song, Y., Kaufmann, E., Scaramuzza, D., & Dürr, P. (2021). Super-human performance in gran turismo sport using deep reinforcement learning. *IEEE Robotics and Automation Letters*, 6(3), 4257–4264.
- Fujita, Y., Uenishi, K., Ummadisingu, A., Nagarajan, P., Masuda, S., & Castro, M. Y. (2020). Distributed reinforcement learning of targeted grasping with active vision for mobile manipulators. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 9712–9719. IEEE.
- Ghosh, D., Gupta, A., Reddy, A., Fu, J., Devin, C., Eysenbach, B., & Levine, S. (2019). Learning to reach goals via iterated supervised learning. *arXiv preprint arXiv:1912.06088*.
- Goodfellow, I., Bengio, Y., Courville, A., & Bengio, Y. (2016). *Deep learning*, Vol. 1. MIT press Cambridge.

- Guo, X. (2017). *Deep learning and reward design for reinforcement learning*. Ph.D. thesis.
- He, J. Z.-Y., & Dragan, A. D. (2021). Assisted robust reward design. *arXiv preprint arXiv:2111.09884*.
- Kadian, A., Truong, J., Gokaslan, A., Clegg, A., Wijmans, E., Lee, S., Savva, M., Chernova, S., & Batra, D. (2020). Sim2real predictivity: Does evaluation in simulation predict real-world performance?. *IEEE Robotics and Automation Letters*, 5(4), 6670–6677.
- Kaelbling, L. P. (1993a). Hierarchical learning in stochastic domains: Preliminary results. In *Proceedings of the tenth international conference on machine learning*, Vol. 951, pp. 167–173.
- Kaelbling, L. P. (1993b). Learning to achieve goals. In *IJCAI*, pp. 1094–1099. Citeseer.
- Khazatsky, A., Nair, A., Jing, D., & Levine, S. (2021). What can i do here? learning new skills by imagining visual affordances. *arXiv preprint arXiv:2106.00671*.
- Krakovsky, M. (2016). Reinforcement renaissance. *Commun. ACM*, 59(8), 1214.
- Kulkarni, T. D., Narasimhan, K. R., Saeedi, A., & Tenenbaum, J. B. (2016). Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation. *arXiv preprint arXiv:1604.06057*.
- LaValle, S. M. (2006). *Planning algorithms*. Cambridge university press.
- Le, H., Jiang, N., Agarwal, A., Dudík, M., Yue, Y., & Daumé, H. (2018). Hierarchical imitation and reinforcement learning. In *International Conference on Machine Learning*, pp. 2917–2926. PMLR.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *nature*, 521(7553), 436–444.
- Li, Y. (2018). Deep reinforcement learning. *arXiv preprint arXiv:1810.06339*.
- McGovern, A., & Barto, A. G. (2001). Automatic discovery of subgoals in reinforcement learning using diverse density..
- McGovern, A., Precup, D., Ravindran, B., Singh, S., & Sutton, R. S. (1998). Hierarchical optimal control of mdps. In *Proceedings of the Tenth Yale Workshop on Adaptive and Learning Systems*, pp. 186–191.
- Meuleau, N., Hauskrecht, M., Kim, K.-E., Peshkin, L., Kaelbling, L. P., Dean, T. L., & Boutilier, C. (1998). Solving very large weakly coupled markov decision processes. In *AAAI/IAAI*, pp. 165–172.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., & Kavukcuoglu, K. (2016). Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pp. 1928–1937.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., & Riedmiller, M. (2013). Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. (2015). Human-level control through deep reinforcement learning. *nature*, 518(7540), 529–533.

- Nasiriany, S., Pong, V. H., Lin, S., & Levine, S. (2019). Planning with goal-conditioned policies. *arXiv preprint arXiv:1911.08453*.
- Ng, A. Y., Harada, D., & Russell, S. (1999). Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*, Vol. 99, pp. 278–287.
- OpenAI (2018). Openai five. <https://blog.openai.com/openai-five/>.
- Parr, R., & Russell, S. (1998). Reinforcement learning with hierarchies of machines. *Advances in neural information processing systems*, pp. 1043–1049.
- Parr, R. E. (1998). *Hierarchical control and learning for Markov decision processes*. University of California, Berkeley Berkeley, CA.
- Pong, V., Gu, S., Dalal, M., & Levine, S. (2018). Temporal difference models: Model-free deep rl for model-based control. *arXiv preprint arXiv:1802.09081*.
- Popov, I., Heess, N., Lillicrap, T., Hafner, R., Barth-Maron, G., Vecerik, M., Lampe, T., Tassa, Y., Erez, T., & Riedmiller, M. (2017). Data-efficient deep reinforcement learning for dexterous manipulation. *arXiv preprint arXiv:1704.03073*.
- Precup, D., Sutton, R. S., & Singh, S. (1998). Multi-time models for temporally abstract planning. *Advances in neural information processing systems*, pp. 1050–1056.
- Puterman, M. L. (2014). *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.
- Racanière, S., Weber, T., Reichert, D. P., Buesing, L., Guez, A., Rezende, D., Badia, A. P., Vinyals, O., Heess, N., Li, Y., et al. (2017). Imagination-augmented agents for deep reinforcement learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp. 5694–5705.
- Raileanu, R., & Rocktäschel, T. (2020). Ride: Rewarding impact-driven exploration for procedurally-generated environments. In *International Conference on Learning Representations*.
- Ratner, E., Hadfield-Menell, D., & Dragan, A. D. (2018). Simplifying reward design through divide-and-conquer. *arXiv preprint arXiv:1806.02501*.
- Rintanen, J., & Hoffmann, J. (2001). An overview of recent algorithms for ai planning. *KI*, 15(2), 5–11.
- Russell, S. J., & Norvig, P. (2009). *Artificial Intelligence: a modern approach* (3 edition). Pearson.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Siciliano, B., Sciavicco, L., Villani, L., & Oriolo, G. (2010). *Robotics: modelling, planning and control*. Springer Science & Business Media.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. (2016). Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587), 484–489.

- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. (2017). Mastering the game of go without human knowledge. *nature*, 550(7676), 354–359.
- Simsek, O., & Barreto, A. S. (2008). Skill characterization based on betweenness. In *Advances in neural information processing systems*, pp. 1497–1504.
- Singh, S., Cohn, D., et al. (1998). How to dynamically merge markov decision processes. *Advances in neural information processing systems*, pp. 1057–1063.
- Singh, S., Lewis, R. L., & Barto, A. G. (2009). Where do rewards come from. In *Proceedings of the annual conference of the cognitive science society*, pp. 2601–2606. Cognitive Science Society.
- Singh, S., Lewis, R. L., Barto, A. G., & Sorg, J. (2010). Intrinsically motivated reinforcement learning: An evolutionary perspective. *IEEE Transactions on Autonomous Mental Development*, 2(2), 70–82.
- Solway, A., Diuk, C., Córdova, N., Yee, D., Barto, A. G., Niv, Y., & Botvinick, M. M. (2014). Optimal behavioral hierarchy. *PLOS Comput Biol*, 10(8), e1003779.
- Sorg, J., Lewis, R. L., & Singh, S. (2010). Reward design via online gradient ascent. *Advances in Neural Information Processing Systems*, 23, 2190–2198.
- Stolle, M., & Precup, D. (2002). Learning options in reinforcement learning. In *International Symposium on abstraction, reformulation, and approximation*, pp. 212–223. Springer.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
- Sutton, R. S., Modayil, J., Delp, M., Degris, T., Pilarski, P. M., White, A., & Precup, D. (2011). Horde: A scalable real-time architecture for learning knowledge from unsupervised sensorimotor interaction. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pp. 761–768.
- Sutton, R. S., Precup, D., & Singh, S. (1999). Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2), 181–211.
- Sutton, R. S., Precup, D., & Singh, S. P. (1998). Intra-option learning about temporally abstract actions.. In *ICML*, Vol. 98, pp. 556–564.
- Szepesvári, C. (2010). Algorithms for reinforcement learning. *Synthesis lectures on artificial intelligence and machine learning*, 4(1), 1–103.
- Van Seijen, H., Fatemi, M., Romoff, J., Laroche, R., Barnes, T., & Tsang, J. (2017). Hybrid reward architecture for reinforcement learning. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., & Garnett, R. (Eds.), *Advances in Neural Information Processing Systems*, Vol. 30. Curran Associates, Inc.
- Vezhnevets, A. S., Osindero, S., Schaul, T., Heess, N., Jaderberg, M., Silver, D., & Kavukcuoglu, K. (2017). Feudal networks for hierarchical reinforcement learning. In *International Conference on Machine Learning*, pp. 3540–3549. PMLR.
- Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., Choi, D. H., Powell, R., Ewalds, T., Georgiev, P., et al. (2019). Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782), 350–354.

- Vinyals, O., Ewalds, T., Bartunov, S., Georgiev, P., Vezhnevets, A. S., Yeo, M., Makhzani, A., Küttler, H., Agapiou, J., Schrittwieser, J., et al. (2017). Starcraft ii: A new challenge for reinforcement learning. *arXiv preprint arXiv:1708.04782*.
- Wen, Z., Precup, D., Ibrahimi, M., Barreto, A., Van Roy, B., & Singh, S. (2020). On efficiency in hierarchical reinforcement learning. *Advances in Neural Information Processing Systems*, 33.
- Xu, X., Huang, T., Wei, P., Narayan, A., & Leong, T.-Y. (2020). Hierarchical reinforcement learning in starcraft ii with human expertise in subgoals selection. *arXiv preprint arXiv:2008.03444*.
- Ye, D., Chen, G., Zhang, W., Chen, S., Yuan, B., Liu, B., Chen, J., Liu, Z., Qiu, F., Yu, H., Yin, Y., Shi, B., Wang, L., Shi, T., Fu, Q., Yang, W., Huang, L., & Liu, W. (2020). Towards playing full moba games with deep reinforcement learning. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., & Lin, H. (Eds.), *Advances in Neural Information Processing Systems*, Vol. 33, pp. 621–632. Curran Associates, Inc.