

# Multi-Timescale Ensemble $Q$ -Learning for Markov Decision Process Policy Optimization

Talha Bozkus  and Urbashi Mitra , *Fellow, IEEE*

**Abstract**—Reinforcement learning (RL) is a classical tool to solve network control or policy optimization problems in unknown environments. The original  $Q$ -learning suffers from performance and complexity challenges across very large networks. Herein, a novel model-free ensemble reinforcement learning algorithm which adapts the classical  $Q$ -learning is proposed to handle these challenges for networks which admit Markov decision process (MDP) models. Multiple  $Q$ -learning algorithms are run on multiple, distinct, synthetically created and structurally related Markovian environments in parallel; the outputs are fused using an adaptive weighting mechanism based on the Jensen-Shannon divergence (JSD) to obtain an approximately optimal policy with low complexity. The theoretical justification of the algorithm, including the convergence of key statistics and  $Q$ -functions are provided. Numerical results across several network models show that the proposed algorithm can achieve up to 55% less average policy error with up to 50% less runtime complexity than the state-of-the-art  $Q$ -learning algorithms. Numerical results validate assumptions made in the theoretical analysis.

**Index Terms**—Markov decision process (MDP), network optimization, ensemble learning, reinforcement learning,  $Q$ -learning

## I. INTRODUCTION

MARKOV Decision Processes (MDPs) are natural mathematical tools for modeling sequential decision-making problems in many large real-world networks [1], [2], [3]. When the underlying system dynamics are observable, the optimization problem of MDPs can be solved by dynamic programming [4]. However, these algorithms are not directly applicable to problems where the underlying MDP is unknown (or non-observable), in which case *model-free* Reinforcement Learning (RL) algorithms such as  $Q$ -learning can be employed to simulate the system dynamics and learn the policies and value functions [5].

Manuscript received 19 July 2023; revised 30 November 2023 and 19 January 2024; accepted 26 February 2024. Date of publication 4 March 2024; date of current version 20 March 2024. This work was supported in part by ARO under Grant W911NF1910269; in part by DOE under Grant DE-SC0021417; in part by Swedish Research Council under Grant 2018-04359; in part by NSF under Grant CCF-2008927, Grant CCF-2200221, Grant CIF-2311653, and Grant RINGS-2148313; in part by ONR under Grant 503400-78050 and Grant N00014-15-1-2550; and in part by USC + Amazon Center on Secure and Trusted Machine Learning. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Shaofeng Zou. (Corresponding author: Talha Bozkus.)

The authors are with the Ming Hsieh Department of Electrical and Computer Engineering, University of Southern California, Los Angeles, CA 90007 USA (e-mail: bozkus@usc.edu; ubli@usc.edu).

Digital Object Identifier 10.1109/TSP.2024.3372699

$Q$ -learning can be employed to solve a variety of optimization and control problems in unknown environments [6], [7], [8], [9]. However, it suffers from several performance and complexity challenges across large MDPs, including high estimation bias and estimation variance, training instability, slow convergence, and high sample complexity. To this end, several variants of  $Q$ -learning have been developed to handle these challenges. Estimation bias is considered in [10], [11], and the estimation variance and training stability are examined in [12], [13]. The convergence rate is improved in [14], and training data efficiency is considered in [15]. In [16], [17], strategies specific to wireless networks are considered. As these algorithms have different objectives, their strategies and implementations also differ. Similar to the original  $Q$ -learning, a single  $Q$ -function estimator is employed in [14], [15]. On the other hand, multiple  $Q$ -function estimators are used in [10], [12], [13], [18], [19], in which each estimator is initialized independently, and their outputs are fused into a single estimate via a weighting mechanism. These algorithms directly operate on the original Markovian environment; however,  $Q$ -learning algorithms using multiple  $Q$ -function estimators on **multiple** Markovian environments have not been well-studied. We will see that the performance and complexity of  $Q$ -learning can be further improved by employing multiple  $Q$ -function estimators on multiple structurally related synthetic Markovian environments operating at different time-scales.

Despite extensive prior work [20], [21], achieving efficient and scalable exploration in large Markovian environments remains a major challenge in reinforcement learning. *Too little exploration* may cause the agent to behave greedily. Consequently, some parts of the environment may never be visited, and the agent may keep selecting sub-optimal actions and get stuck in a local optimum. On the other hand, *too much exploration* may yield a high accumulated cost, preventing the utilization of previous experiences, and be computationally expensive for very large environments. While the previous work generally employs a *single* efficient exploration strategy [20], [21], we herein propose a **two-level exploration** strategy, *i.e.* there are two different sources of exploration: at the *algorithm-level*, we use one of the existing exploration techniques (such as epsilon-greedy  $Q$ -learning [5]), and at the *environment-level*, we utilize multiple, distinct, synthetically created and structurally related Markovian environments, which provide different orders of relationships between states. As we consider products of the probability transition matrix of the original

system to construct our synthetic systems, we deem our approach as having **multiple time scales**, corresponding to the  $n$ -hop transition matrices of the Markov chain. Our goal is to improve the exploration capabilities of the agent and accelerate the exploration stage of  $Q$ -learning.

To this end, we propose a novel ensemble  $Q$ -learning algorithm, where *multiple*  $Q$ -learning algorithms are run in parallel on *multiple* distinct, synthetically created and structurally related Markovian environments. Their outputs are fused into a single  $Q$ -function estimate using an adaptive weighting mechanism based on a Jensen–Shannon divergence between the distributions corresponding to the  $Q$ -functions of different environments. In the end, an approximately optimal deterministic policy is obtained with low complexity. In our prior work [22], [23], we introduced similar algorithms and presented preliminary findings. That initial analysis yielded significant insights into the advantages of employing multiple Markovian environments to improve the accuracy and complexity of the original  $Q$ -learning. Herein, we provide significant improvements over that preliminary work: (i) We propose a more interpretable and computationally cheaper way to construct multiple synthetic Markovian environments. (ii) We remove the constraints and approximations on the structure of the network models; hence, our current design is applicable to a wider range of networks. (iii) We provide a more complete theoretical analysis of the proposed algorithm. (iv) We utilize a new distance metric based on the JSD to improve the accuracy of the adaptive weighting mechanism.

The main contributions of the paper are as follows: (i) We systematically construct the multiple synthetic Markovian environments to enable an efficient and scalable exploration in  $Q$ -learning. (ii) We propose a novel  $Q$ -learning algorithm based on an ensemble of the  $Q$ -functions from multiple Markovian environments. (iii) We provide theoretical analyses on the convergence and error variance of the proposed algorithm. (iv) We simulate the algorithm on several large real-world network classes. Numerical results show that the proposed method outperforms the state-of-the-art  $Q$ -learning algorithms on all networks, achieving up to 55% less average policy error with up to 50% less runtime complexity. In addition, simulations confirm the theoretical analyses.

We use the following notation: the vectors are bold lower case ( $\mathbf{x}$ ), matrices and tensors are bold upper case ( $\mathbf{A}$ ), and sets are in calligraphic font ( $\mathcal{S}$ ).

## II. SYSTEM MODEL AND TOOLS

### A. Infinite Horizon Discounted Cost MDP Model

MDPs are characterized by 4-tuples  $\{\mathcal{S}, \mathcal{A}, p, c\}$ , where  $\mathcal{S}$  and  $\mathcal{A}$  denote the finite state and action spaces, respectively. We denote  $s_t$  as the *state* and  $a_t$  as the *action* taken at discrete time period  $t$ . The transition from state  $s$  to  $s'$  under action  $a$  occurs with probability  $p_a(s, s') = p(s' = s_{t+1} | s = s_t, a = a_t)$ , which is stored in the  $(s, s', a)^{th}$  element of the three-dimensional probability transition tensor (PTT)  $\mathbf{P}$ , and a bounded average cost  $c_a(s) = \sum_{s' \in \mathcal{S}} p_a(s, s') \hat{c}_a(s, s')$  is incurred, which is stored in the  $(s, a)^{th}$  element of the cost matrix

$\mathbf{C}$ , where  $\hat{c}_a(s, s')$  is the instantaneous transition cost from state  $s$  to  $s'$  under action  $a$ . We denote the probability transition matrix (PTM) and cost vector under the action  $a$  by  $\mathbf{P}_a$  and  $\mathbf{c}_a$ , respectively. We focus on infinite horizon discounted cost MDPs, where  $t = \mathbb{Z}^+ \cup \{0\}$ . Our goal is to solve *Bellman's optimality* equation:

$$\mathbf{v}^*(s) = \min_{\pi} \mathbf{v}_{\pi}(s) = \min_{\pi} \mathbb{E}_{\pi} \left[ \sum_{t=0}^{\infty} \gamma^t c_{a_t}(s_t) | s_0 = s \right], \quad (1)$$

$$\pi^*(s) = \operatorname{argmin}_{\pi} \mathbf{v}_{\pi}(s), \quad (2)$$

for all  $s \in \mathcal{S}$ , where  $\mathbf{v}_{\pi}$  is the *value function* [4] under the *policy*  $\pi$ ,  $\mathbf{v}^*$  is the *optimal value function*,  $\pi^*$  is the *optimal policy*, and  $\gamma \in (0, 1)$  is the discount factor. The policy  $\pi$  can define either a specific action per state (*deterministic*) or a distribution over the action space per state (*stochastic*) for each time period. If the policy does not change over time, *i.e.*,  $\pi_t = \pi, \forall t$ , then it is deemed *stationary*. There always exists a deterministic stationary policy that is optimal given a finite state and action spaces [4]. Hence, we, herein, consider deterministic and stationary policies.

### B. $Q$ -Learning

When the system dynamics ( $p$  and  $c$ ) are unknown or non-observable,  $Q$ -learning can be used to solve (1) and (2).  $Q$ -learning seeks to find the optimal policy  $\pi^*$  by learning the  $Q$  functions for all  $(s, a)$  pairs using the following update rule:

$$Q(s, a) \leftarrow (1 - \alpha)Q(s, a) + \alpha(c_a(s) + \gamma \min_{a' \in \mathcal{A}} Q(s', a')), \quad (3)$$

where  $\alpha \in (0, 1)$  is the learning rate. In practice,  $\epsilon$ -greedy policies are used to tackle the *exploration-exploitation* trade-off to ensure that sufficient sampling of the system is captured by visiting each state-action pair sufficiently many times [5]. To this end, a random action is taken with probability  $\epsilon$  (exploration), and a greedy action that minimizes the  $Q$ -function of the next state is taken with probability  $1 - \epsilon$  (exploitation). The agent interacts with the environment and collects samples  $\{s, a, s', c\}$  to update  $Q$ -functions using (3). The learning strategy must specify the trajectory length ( $l$ ) (the number of states in a trajectory) and the minimum number of visits to each state-action pair ( $v$ ), which is generally used as a termination condition for the sampling operation.  $Q$ -functions converge to their optimal values with probability one, *i.e.*,  $Q(s, a) \xrightarrow{w.p. 1} Q^*(s, a)$  for all  $(s, a)$  if necessary conditions are satisfied [24]. The optimal policy and value functions can be inferred from the  $Q$ -functions as follows:

$$\pi^*(s) = \operatorname{argmin}_{a \in \mathcal{A}} Q^*(s, a), \quad \mathbf{v}^*(s) = \min_{a \in \mathcal{A}} Q^*(s, a). \quad (4)$$

### C. Prior Work on Ensemble $Q$ -Learning and Model Ensembles

There has been extensive work on *ensemble*  $Q$ -learning algorithms in the RL literature. For instance, [19] extends double  $Q$ -learning [10] to reduce estimation bias through multiple estimators, while [25] enhances training stability by reducing error

accumulation based on bootstrapping error. In [26], a robust  $Q$ -learning method employs random convex combinations of multiple  $Q$ -functions, and [12] reduces approximation error variance through  $Q$ -function averaging for more stable training. A randomized ensemble double  $Q$ -learning is presented in [13] to improve sample efficiency, while [27] employs pessimistically trained offline  $Q$ -functions for the same purpose.

A variety of works have also leveraged *model ensembles* for a variety of RL problems. For example, [28] approximates the real environment using a linear combination of pre-trained models. To improve sample complexity, [29] utilizes an ensemble of deep neural networks, [30] employs an ensemble of bootstrapped models encoding probability distributions, and [31] introduces an ensemble of Bayesian neural network-based dynamics models. Moreover, [32] improves exploration via model disagreement based on uncertainty estimates from ensembles.

Despite this prior work,  $Q$ -learning algorithms using *multiple  $Q$ -function estimators* on *multiple* Markovian environments (i.e., multiple models) have not been well-studied. By combining the power of ensemble  $Q$ -learning and model ensembles, we will see that sampling and training can be accelerated, and more accurate and stable  $Q$ -functions can be produced.

#### D. Sampling and Creating Multiple Markovian Environments

There are several ways to create multiple environments (and the corresponding PTTs) based on the PTT of the original environment  $\mathbf{P}$ . A natural strategy is to employ some function of  $\mathbf{P}$  (or  $\mathbf{P}^T$ ). In particular, the probability  $p_a(s, s')$  should be related across different environments, and the PTTs of different environments should be row-stochastic (*i.e.* each PTM corresponding to a different action is row-stochastic) or can be converted into row-stochastic environments by employing appropriate normalization without changing the original structure. We herein propose to use  $n$ -hop PTTs ( $\mathbf{P}^n$ ) because (i) they describe the  $n$ -step transition probabilities between states; hence, they are interpretable, corresponding to multiple time-scales, (ii) they are easily computable using efficient matrix multiplication methods, (iii) they are inherently row-stochastic, and (iv) they lead to a nice mathematical analysis (as will be shown later).

There are various factors that suggest employing  $n$ -hop Markovian environments could improve the exploration capabilities of the overall system in several ways: (i) They enable the agent to traverse longer trajectories and uncover new state-action pairs beyond its immediate reach, potentially expediting the agent's understanding of the environment with fewer interactions. (ii) They enable the agent to learn from indirect experiences by simulating trajectories that are not directly observed. (iii) They can encourage the agent to consider longer-term rewards and take actions that may not have immediate rewards, leading to better long-term performance, particularly in environments with sparse rewards or long-term dependencies. (iv) They enable the agent to exploit environment patterns by exploring longer trajectories that uncover complex relationships between actions and outcomes, which can be particularly valuable in structured or repetitive environments such as mazes or

---

#### Algorithm 1 Sampling & Creating Multiple Environments

---

**Input:**  $l, v, K, \mathcal{M}^{(1)}$   
**Output:**  $\mathcal{M}^{(n)}$  for  $n \in \{2, 3, \dots, K\}$

- 1: Initialize each element of  $\hat{\mathbf{P}}_a$  with  $\frac{1}{|\mathcal{S}|}$  for  $a \in \mathcal{A}$
- 2: **while** each  $(s, s')$  in  $\mathcal{M}^{(1)}$  not experienced  $v$  times **do**
- 3:     Choose an initial state  $s$  randomly from  $\mathcal{S}$
- 4:     **repeat**  $l$  **times**
- 5:         Sample  $\{s, a, s'\}$  from  $\mathcal{M}^{(1)}$
- 6:          $\hat{\mathbf{P}}_a(s, s') \leftarrow \hat{\mathbf{P}}_a(s, s') + 1$
- 7:     **end**
- 8: **end while**
- 9: Normalize the sum of each row in  $\hat{\mathbf{P}}_a$  to 1 for  $a \in \mathcal{A}$
- 10: **for**  $n \in \{2, 3, \dots, K\}$  **do**
- 11:     Create  $\hat{\mathbf{P}}_a^n$  by taking  $n^{th}$  matrix power of  $\hat{\mathbf{P}}_a$  for  $a \in \mathcal{A}$
- 12:     Denote  $\mathcal{M}^{(n)}$  as the synthetic Markovian environment corresponding to the  $\hat{\mathbf{P}}^n$
- 13: **end for**

---

puzzles, where the agent must learn to identify and leverage patterns to achieve its objectives. (v) They can help the agent to better handle environments with changing dynamics by enabling it to learn from past experiences that may no longer be directly relevant to the current state of the environment.

The process of sampling, estimation, and constructing multiple Markovian environments is explained in Algorithm 1. The inputs are the original Markovian environment, denoted by  $\mathcal{M}^{(1)}$ , from which the sampling process is performed, the trajectory length ( $l$ ), the number of times each different state transition  $(s \rightarrow s')$  in  $\mathcal{M}^{(1)}$  must be experienced ( $v$ ), and the total number of Markovian environments ( $K$ ). The outputs are the  $K - 1$  synthetic Markovian environments (SME), denoted by  $\mathcal{M}^{(n)}$  for  $n = 2, 3, \dots, K$ . The underlying PTT,  $\mathbf{P}$ , is initially unknown as per the model-free assumption and needs to be estimated to create the PTTs of the multiple environments. Let  $\hat{\mathbf{P}}$  denote the estimated PTT, and  $\hat{\mathbf{P}}_a$  denote the PTM for the action  $a$  in  $\hat{\mathbf{P}}$  for each  $a \in \mathcal{A}$ . In line 1, we initialize each element of  $\hat{\mathbf{P}}_a$  with  $\frac{1}{|\mathcal{S}|}$  so that it is a valid PTM for each  $a \in \mathcal{A}$ . We keep sampling from  $\mathcal{M}^{(1)}$  and updating the elements of  $\hat{\mathbf{P}}$  until each one-step transition between different states under different actions in  $\mathcal{M}^{(1)}$  is experienced at least  $v$  times to ensure that  $\hat{\mathbf{P}}$  is a sufficiently accurate estimate of  $\mathbf{P}$  in lines 2-8. This procedure is known as *sample averaging* [33]. We normalize  $\hat{\mathbf{P}}_a$  row-wise such that the sum of each row in  $\hat{\mathbf{P}}_a$  is 1 for each  $a \in \mathcal{A}$  in line 9. We then create the  $n$ -hop PTTs using  $\hat{\mathbf{P}}$  in line 11. Each different  $\hat{\mathbf{P}}^n$  inherently corresponds to a different Markovian environment and is denoted by  $\mathcal{M}^{(n)}$  in line 12. The relationship between the original Markovian environment  $\mathcal{M}^{(1)}$  and the synthetic Markovian environments  $\mathcal{M}^{(n)}$  ( $n > 1$ ) is given in Fig. 1.

Other approaches to estimate  $\mathbf{P}$  include function approximations and approximate maximum likelihood estimation techniques [34], [35], [36]. However, these approaches have drawbacks: (i) Non-linear function approximations such as neural networks lack interpretability. (ii) They do not appear

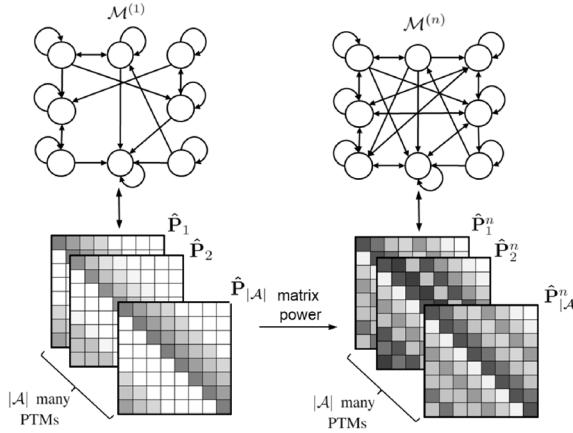


Fig. 1. The relationship between  $\mathcal{M}^{(1)}$  and  $\mathcal{M}^{(n)}$ .

to offer significant computational advantages over sample averaging in sparse Markovian environments. (iii) They generally do not exploit the structural properties of the system. (iv) Training and parameter optimization can be computationally challenging.

The estimation quality affects the accuracy of  $\hat{\mathbf{P}}^n$  (and thus  $\mathcal{M}^{(n)}$ ) differently for each  $n$ . As  $n$  increases, the error from imperfect sampling accumulates due to matrix multiplications. Consequently, higher-order environments may have low accuracy if sampling in the original environment is not done sufficiently. This suggests that  $n$  should not be chosen very large for practical purposes, as we shall see in the numerical results.

### III. nEQL ALGORITHM AND ANALYSIS

In this section, we present the n-hop Ensemble Q-Learning (nEQL) algorithm (Algorithm 2). It is a model-free algorithm since the system dynamics, including transition probabilities and costs, are unknown. Our approach utilizes  $K - 1$  SMEs ( $\mathcal{M}^{(n)}$  for  $n \in 2, \dots, K$ ) in addition to the original Markovian environment ( $\mathcal{M}^{(1)}$ ), resulting in a total of  $K$  Markovian environments. The high-level comparison between the original  $Q$ -learning algorithm, conventional ensemble  $Q$ -learning algorithms and proposed  $Q$ -learning algorithm is shown in Fig. 2, where  $\mathbf{Q}^{(n)}$  represents the  $Q$ -function estimator of the  $Q$ -learning run on  $\mathcal{M}^{(n)}$  for  $n \in [1, K]$ .

The inputs to Algorithm 2 consist of the trajectory length ( $l$ ), the minimum number of visit requirement to each state-action pair ( $v$ ), the update ratio at time  $t$  ( $u_t \in [0, 1]$ ), the total number of Markovian environments ( $K$ ), the empty  $Q$ -tables for  $K$  different environments ( $\mathbf{Q}^{(n)}$  for  $n \in \{1, 2, \dots, K\}$ ), and  $K$  different Markovian environments ( $\mathcal{M}^{(n)}$  for  $n \in \{1, 2, \dots, K\}$ ) since Algorithm 2 requires access to all Markovian environments. Let  $\mathbf{w}_t$  be the weight vector of size  $K$  at time  $t$  (with  $\mathbf{w}_t^{(n)}$  being the  $n^{\text{th}}$  element of  $\mathbf{w}_t$ ). The weight vector at  $t = 0$  ( $\mathbf{w}_0$ ) is initialized randomly to break the symmetry in line 1 (*i.e.* each element is chosen randomly from  $[0, 1]$ , and the vector is softmax-normalized so that  $\sum_{n=1}^K \mathbf{w}_0^{(n)} = 1$ ). The iterations continue until each state-action pair in  $\mathcal{M}^{(1)}$  is visited at least

### Algorithm 2 n-hop Ensemble $Q$ -Learning (nEQL)

---

**Input:**  $l, v, u_t, K, \mathbf{Q}^{(n)}, \mathcal{M}^{(n)}, n \in \{1, 2, \dots, K\}$   
**Output:**  $\mathbf{Q}^{it}, \hat{\pi}$

- 1: Initialize  $\mathbf{w}_0$  randomly,  $\mathbf{Q}_0^{it} \leftarrow \mathbf{0}, t \leftarrow 0$
- 2: **while** each  $(s, a)$  pair in  $\mathcal{M}^{(1)}$  not visited  $v$  times **do**
- 3:   choose common initial state for all  $\mathcal{M}^{(n)}$  randomly
- 4:   **repeat**  $l$  **times**
- 5:     **for** each  $n \in \{1, \dots, K\}$  **do**
- 6:       sample  $\{s, a, s', c\}$  from  $\mathcal{M}^{(n)}$  and update  $\mathbf{Q}_t^{(n)}$  using (3)
- 7:       convert  $\mathbf{Q}_t^{(n)}$  into probabilities  $\hat{\mathbf{Q}}_t^{(n)}$  state-wise using the negative softmax
- 8:        $\mathbf{w}_t^{(n)} \leftarrow 1 - \text{AJSD}(\hat{\mathbf{Q}}_t^{(1)} \parallel \hat{\mathbf{Q}}_t^{(n)})$
- 9:     **end for**
- 10:    $\mathbf{w}_t \leftarrow \text{softmax}(\mathbf{w}_t)$
- 11:    $\mathbf{Q}_{t+1}^{it} \leftarrow u_t \mathbf{Q}_t^{it} + (1 - u_t) \sum_{n=1}^K \mathbf{w}_t^{(n)} \mathbf{Q}_t^{(n)}$
- 12:    $t \leftarrow t + 1$
- 13: **end**
- 14: **end while**
- 15:  $\hat{\pi}(s) \leftarrow \text{argmin}_{a'} \mathbf{Q}^{it}(s, a')$

---

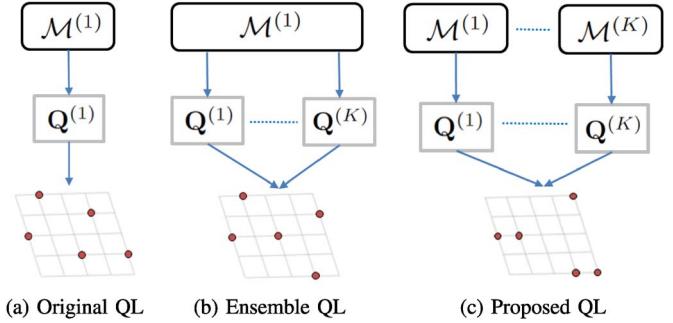


Fig. 2. Classification of Q-learning (QL) algorithms based on their strategies and implementation.

$v$  times (in line 2) to ensure that different state-action dynamics are sufficiently represented.

At the end of each trajectory (*i.e.* every  $l$  time step), all  $K$  Markovian environments are reset, and a common initial state is assigned randomly from  $\{1, 2, \dots, |\mathcal{S}|\}$ , as indicated in line 3. In line 6, independent samples are collected from each different Markovian environment, and corresponding  $Q$ -tables are updated independently. We emphasize that given the common initial state, different actions are taken following the epsilon-greedy policy of each different environment. As a result, different next state and cost pairs are observed for different environments. This procedure is repeated  $l$  times, after which a random but common initial state is set. In line 7, the  $Q$ -functions are converted into probabilities per state using the negative softmax function. For example, if the  $Q$ -functions for a given state are  $[1, 1.4, 0.8, 2]$  (assuming four actions), we input the negative of the  $Q$ -functions ( $[-1, -1.4, -0.8, -2]$ ) to softmax function, and compute the corresponding probabilities as  $[0.31, 0.21, 0.37, 0.11]$ . Recall that we are doing cost *minimization*;

thus, the smaller  $Q$ -function is more likely to correspond to the optimal action. In line 8, we compute the distance between the probability distributions  $\hat{\mathbf{Q}}_t^{(1)}$  and  $\hat{\mathbf{Q}}_t^{(n)}$  using the *averaged Jensen-Shannon divergence* (AJSD) defined as follows:

$$\text{AJSD}(\hat{\mathbf{Q}}_t^{(1)} \| \hat{\mathbf{Q}}_t^{(n)}) = \frac{1}{|\mathcal{A}|} \sum_s \text{JSD}(\hat{\mathbf{Q}}_t^{(1)}(s, :) \| \hat{\mathbf{Q}}_t^{(n)}(s, :)), \quad (5)$$

where  $\hat{\mathbf{Q}}_t^{(n)}(s, :)$  is the probability vector of size  $|\mathcal{A}|$ , and  $\text{JSD}$  between probability distributions  $p$  and  $q$  is defined as [37]:

$$\text{JSD}(p, q) = \frac{1}{2} \left[ \text{KL} \left( p \left\| \frac{p+q}{2} \right. \right) + \text{KL} \left( q \left\| \frac{p+q}{2} \right. \right) \right], \quad (6)$$

where  $\text{KL}$  denotes the Kullback–Leibler divergence using base 2. Although there are several distance measures, including  $l_2$  distance between  $\mathbf{Q}_t^{(1)}$  and  $\mathbf{Q}_t^{(n)}$  or  $\text{KL}$  divergence between  $\mathbf{Q}_t^{(1)}$  and  $\hat{\mathbf{Q}}_t^{(n)}$ , we employ  $\text{JSD}$  because (i)  $\text{JSD}$  is a symmetric measure in contrast to the  $\text{KL}$  divergence. (ii)  $\text{JSD}$  is a smoothed and bounded version of  $\text{KL}$  divergence (bounded to  $[0, 1]$  and hence AJSD is also bounded to  $[0, 1]$ ); thus, it is robust to noise, outliers, or small perturbations in the  $Q$ -functions. (iii) Numerical results show that it provides superior performance to the other measures (see [38]). Herein, a larger  $\mathbf{w}_t^{(n)}$  implies that the two sets of probabilities ( $\hat{\mathbf{Q}}_t^{(1)}$  versus  $\hat{\mathbf{Q}}_t^{(n)}$ ) are closer, so are the corresponding  $Q$ -functions ( $\mathbf{Q}_t^{(1)}$  versus  $\mathbf{Q}_t^{(n)}$ ). The vector  $\mathbf{w}_t$  is softmax-normalized in line 10 and used to update the  $Q$ -function output of Algorithm 2,  $\mathbf{Q}_t^{it}$ , in line 11. When updating  $\mathbf{Q}_t^{it}$ , previous experience are exploited by utilizing fraction  $u_t$  of the  $\mathbf{Q}_t^{it}$  from the previous iteration (*exploitation*), while multiple Markovian environments are sampled based on their weights (*exploration*), and their contributions are weighted by  $1 - u_t$ . In the end, the estimated policy  $\hat{\pi}$  is obtained from  $\mathbf{Q}_t^{it}$  in line 15.

The  $\mathbf{Q}_t^{it}$  (*iterative*) is updated adaptively using the current weight vector  $\mathbf{w}_t$ ; hence, it captures *asymmetric* information between different Markovian environments, *i.e.*, how the utility of samples obtained from  $\mathcal{M}^{(n)}$  may change during iterations. It is likely that  $\mathcal{M}^{(1)}$  provides more useful samples at the beginning as it is the original environment, and there are not enough samples to capture the higher-order relationships versus the first-order relationships. On the other hand,  $\mathcal{M}^{(n)}$  for larger  $n$  contributes more as the iterations increase as the first-order relationships may not aid as much in exploration.

The weights in Algorithm 2 ( $\mathbf{w}_t^{(n)}, n \in \{1, 2, \dots, K\}$ ) converge due to the fact that the distinct  $Q$ -functions converge to their optimal values via  $Q$ -learning [24], and the weights are calculated based on the attendant  $Q$ -functions. We will also verify the convergence of the weights numerically.

We emphasize that there exists a potentially **distinct** optimal policy for each  $K$  different Markovian environment (*i.e.*  $\pi_1^*, \pi_2^*, \dots, \pi_K^*$ ) corresponding to different  $Q$ -functions ( $\mathbf{Q}^{(1)}, \mathbf{Q}^{(2)}, \dots, \mathbf{Q}^{(K)}$ ). Algorithm 2 also yields an ensemble policy ( $\hat{\pi}$ ) corresponding to the  $Q$ -function output  $\mathbf{Q}_t^{it}$ . We will demonstrate that  $\mathbf{Q}_t^{it}$  converges to the optimal  $Q$ -functions of the original environment ( $\mathbf{Q}^{(1)}^*$ ) in the mean-square sense, and thus the ensemble policy ( $\hat{\pi}$ ) also converges to the optimal policy of the original environment  $\pi_1^*$ . This implies that we can

obtain the optimal policy  $\pi_1^*$ , which is the **ultimate** goal, using our proposed algorithm with significantly lower complexity.

Our proposed algorithm combines features of online and offline RL methods. Initially, we construct the PTT of the original environment ( $\hat{\mathbf{P}}$ ) by sampling from the original environment until  $\hat{\mathbf{P}}$  is a sufficiently accurate estimate of  $\mathbf{P}$ . Then, we update the corresponding  $Q$ -functions ( $\mathbf{Q}^{(1)}$ ) by continuously interacting with the original environment in real-time, which is the *online* part of our approach. The PTTs of the multiple SMEs ( $\hat{\mathbf{P}}^n$ ) are constructed using  $\hat{\mathbf{P}}$ , and the corresponding  $Q$ -functions ( $\mathbf{Q}^{(n)}$ ) are updated by collecting synthetic samples from the  $n^{th}$  environment, which is the *offline* part. We emphasize that the estimated PTT ( $\hat{\mathbf{P}}$ ) and PTTs of multiple SMEs ( $\hat{\mathbf{P}}^n$ ) are constructed **only once** and not further updated using the newly collected samples in real-time, which is computationally expensive because of the need to perform matrix multiplications to construct  $\hat{\mathbf{P}}^n$  and normalize  $\hat{\mathbf{P}}$  and  $\hat{\mathbf{P}}^n$  at each iteration. Updating  $\hat{\mathbf{P}}$  beyond a certain point also results in minimal accuracy improvements, as illustrated in Fig. 7. It is also important to note that our approach is **different** from hybrid RL [39], [40], [41], where the agent generally has access to an offline dataset and subsequently collects new data through interacting with the environment. On the other hand, our algorithm *continuously updates* the  $Q$ -functions in real-time while also leveraging pre-collected data to produce more accurate and stable  $Q$ -functions with low complexity.

### A. Theoretical Analysis

In this section, we provide several theoretical results for Algorithm 2. Assume the  $Q$ -function **errors** of the  $n^{th}$  environment follows an arbitrary distribution  $D_n$  with zero mean and finite variance as follows:<sup>1</sup>

$$\mathcal{X}_t^{(n)}(s, a) \stackrel{\text{def}}{=} \mathbf{Q}_t^{(n)}(s, a) - \mathbf{Q}^*(s, a) \sim D_n \left( 0, \frac{\lambda_n^2}{3} \right), \quad (7)$$

for all  $(s, a)$  and  $n$  with  $\lambda_n > 0$  where  $\mathbf{Q}^*$  is the optimal  $Q$ -functions of the original Markovian environment. Prior work has considered the distribution  $D_n$  to be uniform, non-uniform, or normal for the  $n = 1$  case [13], [42], [43], [44]. Herein, we make no assumptions on  $D_n$ . Simulations verify that the true distributions  $D_n$  are, in fact, very close to the normal distributions with zero-mean and finite variance for all  $n$  (see Fig. 9(b) and Fig. 9(e)).

Let  $\mathbb{E}$  and  $\mathbb{V}$  be the expectation and variance operators,  $\lambda = \max_{n \in \{1, 2, \dots, K\}} \lambda_n$ , and  $\mathcal{E}_t(s, a) \stackrel{\text{def}}{=} \mathbf{Q}_t^{it}(s, a) - \mathbf{Q}^*(s, a)$ .

*Proposition 1:* Let  $u_t$  be a constant:  $u_t = u$ . Under Assumption (7), Algorithm 2 produces unbiased  $Q$ -functions in the limit *i.e.*  $\lim_{t \rightarrow \infty} \mathbb{E}[\mathcal{E}_t(s, a)] = 0$ . If the  $Q$ -function errors of a given environment at different times are independent *i.e.*  $\mathcal{X}_{t_1}^{(n)}(s, a) \perp \mathcal{X}_{t_2}^{(n)}(s, a)$  for all  $s, a, n$ ,  $t_1 \neq t_2$ , the error variance in the limit can be upper bounded as:  $\lim_{t \rightarrow \infty} \mathbb{V}[\mathcal{E}_t(s, a)] \leq \frac{(1-u)}{(1+u)} \lambda^2$ . (see Appendix A)

<sup>1</sup>We observe that one can construct **small** state-space examples that do not adhere to this assumption; however, for the large scale examples we consider, numerical results suggest that the assumption is valid (see Fig. 9(b) and Fig. 9(e)).

This proposition shows that under the assumption (7),  $\mathbf{Q}_t^{it}$  is an unbiased estimator of  $\mathbf{Q}^*$  in the limit, and the upper bound on the error variance can be controlled by the parameters  $u$  and  $\lambda$ . Herein, a larger  $\lambda$  implies a higher uncertainty in the  $Q$ -function errors, which makes the upper bound on the variance looser. On the other hand, when the algorithm converges ( $t \rightarrow \infty$ ), a larger  $u$  leads to less reliance on SMEs, reducing the uncertainty arising from multiple environments and yielding a tighter upper bound.

The zero-mean assumption in (7) is only employed to simplify the analysis of Proposition 1. Nevertheless, this assumption can be relaxed as follows:

*Corollary 1:* Proposition 1 is valid under the assumption that the  $Q$ -function errors follow arbitrary distributions, that is:

$$\mathcal{X}_t^{(n)}(s, a) \stackrel{\text{def}}{=} \mathbf{Q}_t^{(n)}(s, a) - \mathbf{Q}^*(s, a) \sim D_n \left( \mu_n, \frac{\lambda_n^2}{3} \right), \quad (8)$$

if the weighted combination of the means at different times is 0, i.e.,  $\sum_{n=1}^K \mathbf{w}_t^{(n)} \mu_n = 0$ , where  $\mu_n$  is the mean of the distribution  $D_n$ . (see Appendix B)

This assumption is less restrictive than (7) since it only requires the weighted convex combination of means to be zero, allowing different means to be non-zero. This result can be numerically validated by Fig. 9(b) and Fig. 9(e).

We emphasize that there is estimation bias between  $\mathbf{Q}^{(1)}$  and  $\mathbf{Q}^{(n)}$  for  $n > 1$  since these  $Q$ -functions are obtained from different environments, and this is reflected by the non-zero means in Equation (8); however, the sign of estimation error (i.e., whether it is overestimation or underestimation) depends on  $n$ . As will be seen later, Algorithm 2 produces unbiased output ( $\mathbf{Q}^{it}$ ). This follows as (i) Algorithm 2 uses a weighted combination of different  $Q$ -functions to update  $\mathbf{Q}^{it}$  and thus the effect of underestimation and overestimation cancel each other, and (ii) a small weight is assigned to the  $Q$ -function with high bias, minimizing the impact of individual biases.

*Corollary 2:* If we remove the independence assumption in Proposition 1, the upper bound on the error variance in the limit can be updated as:  $\lim_{t \rightarrow \infty} \mathbb{V}[\mathcal{E}_t(s, a)] \leq \frac{2\lambda^2}{(1+u)^2} + \frac{(1-u)}{(1+u)} \lambda^2$ . (see [38])

The relaxation of the independence assumption introduces a bias term to the upper bound estimation, resulting in a looser bound than the one in Proposition 1. Nevertheless, a smaller  $\lambda$  or a larger  $u$  tightens the bound as in Proposition 1.

*Corollary 3:* If we use the form  $u_t = 1 - e^{\frac{-t}{c_4}}$  with  $c_4 > 0$  and the independence assumption in Proposition 1, the error variance converges to zero:  $\lim_{t \rightarrow \infty} \mathbb{V}[\mathcal{E}_t(s, a)] = 0$ .

We choose the parameter  $u_t$  such that  $u_t \xrightarrow{t \rightarrow \infty} 1$  as  $u$  should be small initially to explore multiple environments in the beginning (exploration) and should increase to utilize previously obtained samples with time (exploitation *i.e.* less reliance on the synthetic environments as we learn  $\mathbf{Q}_t^{it}$  better). Herein, using multiple Markovian environments accelerates the convergence of  $\mathbf{Q}_t^{it}$  towards  $\mathbf{Q}^*$ , while adjusting  $u_t$  pushes  $\mathbf{Q}_t^{it}$  in the desired direction. Consequently, Algorithm 2 converges and yields the optimal  $Q$ -functions *in the mean-square sense*, which can be shown by combining the results of Proposition 1 and Corollary 3. We will numerically verify that the independence

assumption is **almost always** satisfied. The parameter  $c_4$  adjusts the decay rate of  $u_t$ , which is crucial to tune the amount of exploration. In particular, a larger  $c_4$  implies a slower decay rate, which is needed for larger networks, where more exploration is necessary. This result aligns with Proposition 1 and Corollary 2, as larger values of  $u$  lead to tighter upper bounds on the error variance and  $u = 1$  makes the upper bounds zero.

The structure of  $u_t$  ( $u_t \xrightarrow{t \rightarrow \infty} 1$ ) prevents error accumulation in  $\mathbf{Q}^{it}$  due to the potential lack of knowledge about  $\mathbf{Q}^{(1)}$ , especially in the initial stages. Even if weights at time  $t$  ( $\mathbf{w}_t$ ) are incorrectly assigned, potentially leading to errors in computing  $\mathbf{Q}_{t+1}^{it}$ , the impact of this contribution to the final  $Q$ -function estimate is outweighed by that of  $\mathbf{Q}_{t+2}^{it}$  due to scaling by  $u_t$ , which increases over time ( $u_t < u_{t+1}$ ). Herein,  $\mathbf{Q}_{t+2}^{it}$  is calculated using updated weights ( $\mathbf{w}_{t+1}$ ), which are computed based on a more accurate estimate of  $\mathbf{Q}^{(1)}$ .

The convergence proof of the algorithm in our prior work [23] can be adapted to show the convergence of the Algorithm 2 deterministically without any distribution or independence assumption on the  $Q$ -function errors.

*Proposition 2:* The upper bound on the error variance decreases with the number of Markovian environments  $K$  as  $\mathbb{V}[\mathcal{E}_t(s, a)] \leq \frac{c(\lambda, u)}{K}$  for all  $t$ , where  $c(\lambda, u)$  is a constant of  $K$ . (see Appendix C)

The proposition shows that increasing the number of Markovian environments ( $Q$ -learning algorithms running on different environments simultaneously) reduces the upper bound on variance. This aligns with the primary objective of ensemble algorithms. Unlike the bound in Proposition 1, this upper bound explicitly depends on  $K$ , is valid without any assumptions on independence or structure of  $u_t$ , and holds for all  $t$ . The decay rate of the upper bound is determined by the function  $c(\lambda, u)$ , which incorporates both  $\lambda$  and  $u$ . Furthermore, numerical results will show that increasing  $K$  may not always yield an increasing reduction in the upper bound, *i.e.* there is a diminishing return of increasing  $K$ .

*Proposition 3:* Let the output policy of Algorithm 2 be  $\hat{\pi}$ , and the  $Q$ -functions of the  $n^{th}$  environment under the policy  $\hat{\pi}$  be  $\mathbf{Q}_{\hat{\pi}}^{(n)}$ . Then, Algorithm 2 produces  $Q$ -functions on different environments that satisfy the following upper bound:

$$\|\mathbf{Q}_{\hat{\pi}}^{(1)} - \mathbf{Q}_{\hat{\pi}}^{(n)}\| < \frac{\gamma}{1 - \gamma^n} \frac{1 - \gamma^{n-1}}{1 - \gamma} \|\mathbf{c}_{\hat{\pi}}\|, \quad (9)$$

where  $n > 1$ , the norm  $\|\cdot\|$  is the  $l_2$  norm, and  $\mathbf{c}_{\hat{\pi}}$  is the cost vector under the policy  $\hat{\pi}$ . (see Appendix D)

This proposition shows that the  $Q$ -functions of different environments under the same policy are closely related as a result of the structural relationship between different Markovian environments. As  $n \rightarrow \infty$ , the upper bound primarily depends on the cost function, eliminating the influence of the learning parameter  $\gamma$ . This suggests  $n$  should not be chosen very large for practical purposes. Note that this behavior should not imply the monotonicity of the  $Q$ -functions as a function of  $n$ , as will be seen later, *i.e.*  $\|\mathbf{Q}_{\hat{\pi}}^{(1)} - \mathbf{Q}_{\hat{\pi}}^{(n)}\|$  is a non-monotonic function of  $n$ . It is also worth emphasizing that this result holds without any assumptions on the independence or structure of  $u_t$ .

*Proposition 4:* The  $\mathbf{Q}_{\hat{\pi}}^{(n)}$  satisfy the following (partial) ordering when  $\gamma \rightarrow 1$  (i.e. when the underlying discounted MDP starts to resemble an undiscounted MDP): (see Appendix E)

$$\mathbf{Q}_{\hat{\pi}}^{(n)}(s) \geq \mathbf{Q}_{\hat{\pi}}^{(2n)}(s) \geq \mathbf{Q}_{\hat{\pi}}^{(4n)}(s) \geq \mathbf{Q}_{\hat{\pi}}^{(8n)}(s) \dots, \quad (10)$$

for all  $s$ , where  $\mathbf{Q}_{\hat{\pi}}^{(1)}(s)$  is the largest and  $n$  is an odd number.

This proposition enables assessing the utility of various Markovian environments by ordering the  $Q$ -functions and facilitates determining the most useful (informative) environments to be used in Algorithm 2.

Assume we want to use  $K = 4$  environments. A challenge is to determine which combination of four environments to include in our algorithm, for example:  $\{1^{st}, 2^{nd}, 3^{rd}, 4^{th}\}$  vs  $\{1^{st}, 2^{nd}, 3^{rd}, 5^{th}\}$ , etc. Herein, the notation  $n^{th}$  simply refers to the matrix power used to create Markovian environments. It is important to note that  $n$  can be greater than  $K$ , (i.e.,  $n$  is not limited to  $[1, K]$ ). Thus, in this example, using the  $5^{th}$  environment does not necessarily require using the first five environments as the  $4^{th}$  environment is not used. Herein, we want to collect diverse  $n$ -hop information from different environments while avoiding environments that do not contribute much to achieving the optimal solution. To this end, we assess the similarity of  $\mathbf{Q}^{(n)}$  (for  $n > 1$ ) to  $\mathbf{Q}^{(1)}$  based on the given orderings.  $\mathbf{Q}^{(1)}$  and  $\mathbf{Q}^{(2)}$  are the most similar, making it a logical choice to include both environments. The  $3^{rd}$  environment is added to explore its potential usefulness, given the uncertain similarity between  $\mathbf{Q}^{(3)}$  and  $\mathbf{Q}^{(2)}$ . In contrast,  $\mathbf{Q}^{(4)}$  is less similar to  $\mathbf{Q}^{(1)}$  than  $\mathbf{Q}^{(2)}$ ; thus, we do not include the  $4^{th}$  environment, while we can potentially explore more useful environments. Due to the same reason, we also want to explore the  $5^{th}$  environment. Hence, a reasonable selection is to use  $1^{st}, 2^{nd}, 3^{rd}$  and  $5^{th}$  environments. We emphasize that we use  $K = 4$  environments, but one of the environments is constructed using the  $5^{th}$  matrix power, which is greater than  $K$ . While this approach may not always be optimal, it offers a pragmatic strategy when prior information is lacking.

We emphasize that this result is valid without any independence or structural assumption on  $u_t$ . Moreover, this result is particularly useful for network settings when long-term planning is more important, for example, when the future costs are more important than the immediate costs or taking some actions may not minimize the immediate rewards, but they will be more beneficial in the long run.

#### IV. NUMERICAL RESULTS

In this section, we consider a variety of performance metrics to assess the accuracy and complexity performance of Algorithm 1 and 2 across different network models.

##### A. Network Models

We consider four different network models, which differ in their design, complexity, and implementation. (See [38] for further details)

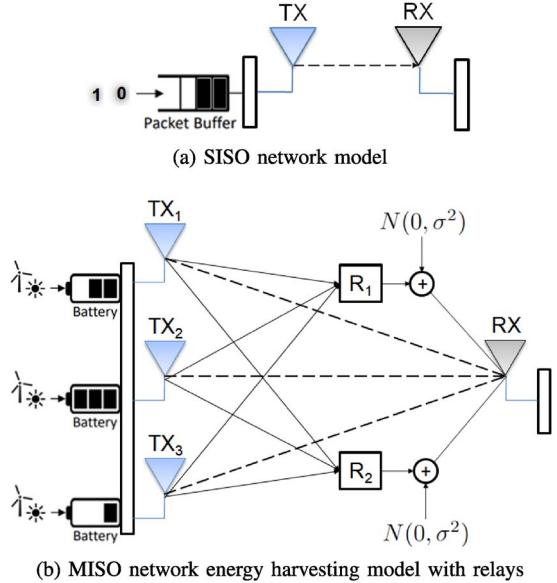


Fig. 3. Examples of wireless network models.

1) *Randomized Graphs:* We consider the Erdős-Rényi (ER) random-directed graph model. The PTM has  $|\mathcal{S}|$  nodes, and each edge is created with a probability 0.2. We create  $|\mathcal{A}|$  many PTMs and concatenate them to obtain the PTT, which is used for sampling and creating the SMEs. The cost function assigns a uniform  $[0, 1]$  random cost to each state-action pair.

2) *Cliff-Walking Environment:* We consider the cliff-walking environment [5]. The number of columns is chosen to be approximately three times the number of rows in the grid so that the number of states is equal to  $|\mathcal{S}|$  (for example, 60 columns and 20 rows  $\rightarrow$  1800 states). If the agent moves to the cliff region, a unit cost is incurred. Moving to the safe grid results in a negative unit cost, while any other movement incurs a cost of 0.01.

3) *SISO Wireless Network Model:* We consider the model of [17], in which there is a single transmitter (TX) and receiver (RX) as shown in Fig. 3(a). The goal is to determine when the transmitter should *transmit data* or *remain silent* to minimize the sum of transmission and packet drop costs.

4) *MISO Energy Harvesting Wireless Network With Gaussian Interference Channels and Multiple Relays:* We consider the model of [45]. An example network with three transmitters ( $\text{TX}_1, \text{TX}_2, \text{TX}_3$ ), a single receiver (RX), and two relays ( $\text{R}_1, \text{R}_2$ ) is shown in Fig. 3(b). The goal is to determine when transmitters should *directly transmit* or *transmit through relays* in order to maximize the overall throughput while minimizing the sum of battery and packet drop costs for each transmitter.

##### B. Average Policy Error Results

Let  $\pi^*$  be the optimal policy from (2), and  $\hat{\pi}$  be the output policy of Algorithm 2. Since our main concern is optimal control, we define the *average policy error* (APE) as follows:

$$APE = \frac{1}{|\mathcal{S}|} \sum_{s=1}^{|\mathcal{S}|} \mathbf{1}(\pi^*(s) \neq \hat{\pi}(s)). \quad (11)$$

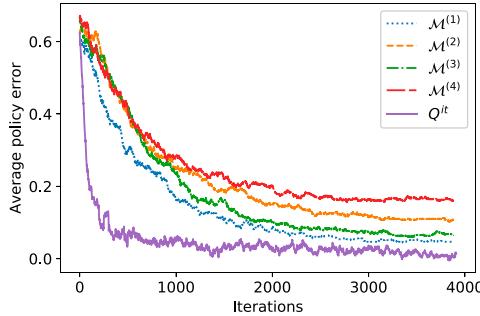


Fig. 4. APE performances across different environments.

We analyze the performance of Algorithm 2 over  $Q$ -learning in Fig. 4. The simulation is carried out using model-4 with network size 5000 with the following parameters:  $\gamma = 0.95$ ,  $\alpha_t = \frac{1}{1 + \frac{t}{100}}$ ,  $u_t = 1 - e^{-\frac{t}{1000}}$ ,  $v = 40$ ,  $l = 10$ ,  $K = 4$ . These parameters are optimized through cross-validation (see Section IV-F for further details). We choose the  $\epsilon$  value for each distinct  $Q$ -learning algorithm as follows:  $\epsilon_t^{(n)} = \max((c_n)^t, 0.01)$  for  $n = 1, 2, 3, 4$  with  $c_1 = 0.95$ ,  $c_2 = c_3 = 0.97$ ,  $c_4 = 0.99$ . We choose these parameters to prioritize learning from useful environments over random exploration. Faster decay rates (i.e., smaller  $c_n$ ) are applied to environments that closely resemble the original environment, informed by Proposition 4. The curves represent APE of  $Q$ -learning algorithms on the original environment  $\mathcal{M}^{(1)}$  and three different SMEs  $\mathcal{M}^{(n)}$  for  $n = 2, 3, 4$ , while  $Q^{it}$  represents the APE of Algorithm 2. Clearly, a near-zero APE (around 0.05) can be achieved with a significantly small number of iterations (around 500). The sharp decline in the  $Q^{it}$  curve at the beginning (up to the 200<sup>th</sup> iteration) corresponds to the *exploration* stage, followed by the *exploitation* stage. Compared to any other algorithm, the exploration stage in Algorithm 2 is fast, which shows the advantages of utilizing multiple Markovian environments to enable a deep and efficient exploration and accelerate the overall training. Notice that individual  $Q$ -learning algorithms can only achieve a slightly higher APE than that of Algorithm 2 if they run for a significantly long time. We observe that the APE results are not monotonic across  $n$  (APE of  $Q$ -learning run on  $\mathcal{M}^{(2)}$  is lower than that of  $\mathcal{M}^{(4)}$ ), which is in line with Proposition 4.

In order to provide a performance comparison, we employ several  $Q$ -learning algorithms, each with different objectives and strategies (number of estimators). Table I provides an overview of these algorithms. Specifically, we focus on value-based model-free algorithms which follow the same strategy as Algorithm 2 to ensure a fair comparison. We also include two algorithms using multiple models (environments): *Ensemble Graph Q-Learning* (EGQL) [23] and *Model-ensemble trust-region policy optimization* (TRPO) [29]. EGQL adopts a different strategy to create multiple SMEs instead of  $n$ -hop systems, employs a weighting mechanism based on policies rather than  $Q$ -functions, and imposes several constraints and assumptions on the system model. TRPO takes a different approach by using an ensemble of deep neural networks for modeling both dynamics and policy (model-based strategy compared to our hybrid

TABLE I  
 $Q$ -LEARNING ALGORITHM AND VARIANTS

Algorithm	Objective	Strategy	
		Est.	Env.
Simple $Q$ (Q) [4]	-	Single	Single
Speedy $Q$ (SQ) [14]	Convergence rate	Single	Single
Double $Q$ (DQ) [10]	Bias	Multi	Single
MaxMin $Q$ (MMQ) [18]	Bias & variance	Multi	Single
Ensemble Bootst. $Q$ (EBQ) [19]	Bias	Multi	Single
Averaged DQN (ADQN) [12]	Stability, Variance	Multi	Single
Model Ensemble (TRPO) [29]	Sample Complexity, Stability	Multi	Multi
Ensemble Graph $Q$ (EGQL) [23]	Variance, Learning speed	Multi	Multi
$n$ -hop Ensemble $Q$ (nEQL)	Variance, Learning speed	Multi	Multi

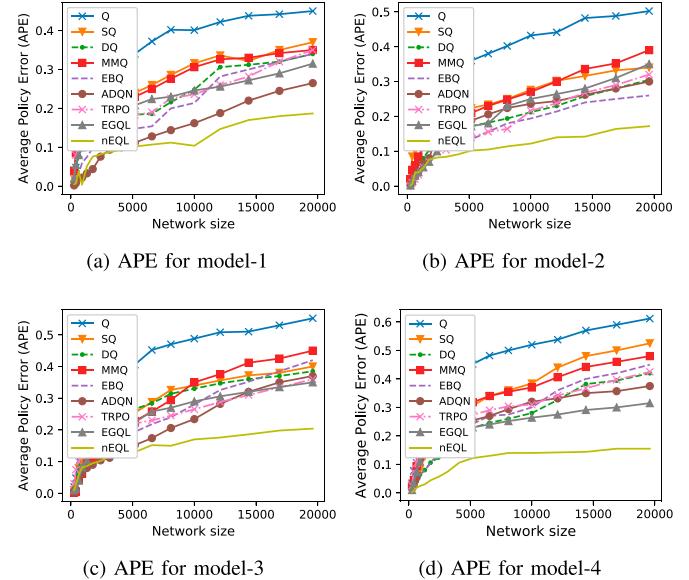


Fig. 5. APE results across different network models.

scheme). It fits this ensemble to a single real-world dataset (in contrast to our use of multiple datasets, some synthetically generated) and employs supervised learning for training, unlike this work. For further details regarding the parameter optimization of each algorithm, refer to [38].

The APE of different algorithms across network size and different models are given in Fig. 5(a)–5(d). Overall, our proposed algorithm consistently achieves a lower APE compared to other algorithms: 30% less for model-1, 35% less for model-2, 45% less for model-3, and 55% less for model-4. The APE gains become more clear for larger networks as using multiple Markovian environments enables deep and efficient exploration by combining the  $n$ -hop relationships between states into a single estimate, and the weighting mechanism based on JSD enables the algorithm fully exploit the most useful environments during training by assigning higher weights.

The proposed algorithm demonstrates significant APE gains, particularly with model-3 and model-4, showing its effectiveness and practicality across real-world networks. In general, ADQN, EBQ, and EGQL produce the lowest APE among other algorithms; however, they have inferior performance compared to nEQL because neural network-based algorithms (ADQN and

TRPO) do not leverage the structural properties of multiple Markovian environments. The performance of EBQ shows the advantages of ensemble algorithms. However, it does not employ a weighting mechanism when calculating the ensemble but uses simple averaging; thus, it performs worse than nEQL, which employs an adaptive weighting mechanism. The performance of EGQL underscores the benefits of using multiple environments (i.e., model ensemble). Nevertheless, it relies on a less accurate and less robust weighting mechanism, does not make use of the simplicity of the  $n$ -hop systems, and suffers from the constraints on the system model. We also observe that as the model complexity increases (from model-1 to model-4), the APE of all algorithms increases. The proposed algorithm, however, yields the least increase in APE, making it the most accurate algorithm. The APE order of the other algorithms remains consistent across different models.

The network models considered herein can be well-modeled with discrete state-spaces with  $|\mathcal{S}| < 20000$  (as in [17], [22], [23], [45]). Numerical simulations show that the performance between the proposed algorithm (tabular-based) and ADQN (neural-network-based) varies as a function of the network size. While there is no clear performance difference for small networks ( $|\mathcal{S}| < 1000$ ), the proposed algorithm offers up to 35% less APE with 20% less runtime for modest-sized networks ( $1000 < |\mathcal{S}| < 10000$ ). For large networks ( $10000 < |\mathcal{S}| < 20000$ ), the advantages of our algorithm become more apparent (55% less APE with 50% less runtime) because of (i) the challenging state-space design for ADQN, (ii) the increased overfitting risk in sparse state spaces for ADQN in data-limited regions, and (iii) the simple averaging mechanism in ADQN that does not capture the changes in the distribution of the  $Q$ -functions over time.

We observe that our algorithm inherits several properties from the traditional  $Q$ -learning algorithm and hence is particularly tailored to finite, but large discrete state-action spaces. For networks with continuous state-spaces, the tabular nature of our algorithm may pose limitations. To this end, replacing tabular  $Q$ -learning with deep  $Q$ -networks is worth exploring.

### C. Computational Complexity Results

The Algorithm 2 can be shown to have the following average time-complexity  $O\left(\frac{|\mathcal{S}||\mathcal{A}|v}{K}f(l, \epsilon)\right)$ , where  $f$  is some non-monotonic function of  $l$  and  $\epsilon$ . The derivation closely follows that presented in [22]. The runtime complexity increases with the network size ( $|\mathcal{S}|$  and  $|\mathcal{A}|$ ) as well as the number of visit requirement to each state-action pair ( $v$ ). On the other hand, the non-monotonicity of  $f$  in  $l, \epsilon$  implies that there are optimal values for  $l, \epsilon$ ; thus, parameter-tuning is required. The complexity is also inversely proportional to the number of Markovian environments ( $K$ ), which may seem counter-intuitive. However, the number of samples that need to be collected from each Markovian environment decreases with  $K$ , leading to an overall reduction in the training runtime complexity.

The runtime of nEQL includes time spent on sampling (visiting each state-action pair at least  $v$  times), constructing multiple SMEs  $K - 1$  times using matrix power operations, and the

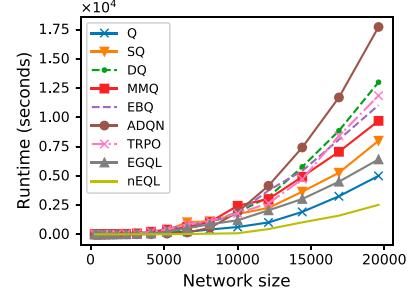


Fig. 6. Runtimes for different algorithms.

time until convergence of  $Q$ -learning algorithms on different Markovian environments. The overall runtime is determined by the slowest  $Q$ -learning algorithm, as they are run in parallel. Other algorithms' runtime results represent the time until each algorithm converges. This runtime measure also serves as a measure of *computational power complexity*. Across fixed network sizes and four different models, the runtime of algorithms is comparable. Therefore, a single runtime result is presented in Fig. 6. The proposed algorithm achieves 50% less runtime than the other algorithms across large networks, which can be attributed to several reasons: (i) Utilizing multiple Markovian environments reduces the need for long trajectories to capture distant node relationships. (ii) The amount of exploration is improved due to the running of multiple Markovian environments simultaneously. (iii) The algorithm can be terminated early with minimal APE as shown in Fig. 4. (iv) The average runtime complexity decreases with  $K$  hyperbolically. (v) Employing different epsilon-greedy policies allows us to capture the distinct characteristics of different Markovian environments. The complexity reduction is independent of the network model or underlying assumptions, making the proposed algorithm an efficient approach for learning various complex environments.

The original  $Q$ -learning has a relatively lower runtime than the other benchmarks as it is a straightforward algorithm. SQ aims to improve the learning (and convergence) rate; therefore, it also produces a relatively lower runtime. EGQL also uses a strategy similar to the current approach using multiple models, which makes it relatively faster than other benchmarks. On the other hand, EBQ has multiple estimators, while ADQN and TRPO suffer from the need for long training periods, which becomes computationally intractable for large networks. It is worth noting that as the model complexity increases (e.g.,  $Q \rightarrow NQ \rightarrow ADQN$ ), the corresponding APE generally decreases, but the runtime complexity increases. However, nEQL achieves a small APE with a small runtime, overcoming the performance-complexity trade-off.

To solve (1) and (2), one could consider using dynamic programming methods such as value iteration instead of  $Q$ -learning, given that we estimate the PTTs of multiple Markovian environments using  $\hat{\mathbf{P}}^n$ . We herein follow a similar iterative procedure to Algorithm 2. While we keep sampling to update  $\hat{\mathbf{P}}$  continuously, we construct the PTTs of multiple SMEs using the current estimate of  $\hat{\mathbf{P}}$  every 10 iterations, and obtain the value functions  $\mathbf{v}_t^{(n)}$  for  $n \in \{1, 2, \dots, K\}$  using the

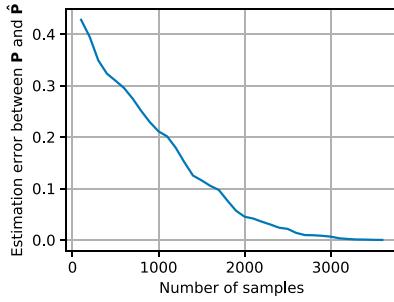


Fig. 7. Estimation error between  $\mathbf{P}$  and  $\hat{\mathbf{P}}$  vs number of samples collected using sample averaging.

value iteration algorithm. We then assign the weights  $\mathbf{w}_t^{(n)}$  as:

$$\mathbf{w}_t^{(n)} \leftarrow -\|\mathbf{v}_t^{(1)} - \mathbf{v}_t^{(n)}\|_2, \quad \forall n \quad (12)$$

$$\mathbf{w}_t \leftarrow \text{softmax}(\mathbf{w}_t), \quad (13)$$

We finally update the value function of the proposed algorithm  $\mathbf{v}_t^{it}$  similar to Algorithm 2 using  $u_t$  (with the same form as in Corollary 3). We carry out these operations iteratively  $k$  times, where  $k$  is sufficiently large to ensure the convergence. This approach has runtime-complexity  $O(k|\mathcal{S}|^2|\mathcal{A}|)$  since we carry out value iteration  $k$  times, which has a complexity of  $O(|\mathcal{S}|^2|\mathcal{A}|)$ . In our simulations, with the settings in Section IV-B, this strategy incurs 50% more runtime complexity and 60% more APE across large networks, making it inferior to our proposed algorithm.

In order to show the computational advantages of Proposition 4 in determining the optimal set of Markovian environments to be employed in Algorithm 2, we choose  $K = 4$  environments over the set of orders  $n = \{1, 2, 3, 4, 5, 6\}$  by (i) using partial ordering from Proposition 4 and (ii) trying every possible distinct combination of four Markovian environments and choosing the one that minimizes APE (i.e. exhaustive search). The same simulation settings in Section IV-B are employed. Numerical results show that Proposition 4 gives the optimal set of Markovian environments with 35% less runtime complexity than exhaustive search. The complexity reduction further increases with the network size and  $K$ , which implies that Proposition 4 is particularly useful for large networks.

We demonstrate the quality of the estimation of the PTT  $\hat{\mathbf{P}}$  as a function of the total number of samples collected from the original Markovian environment  $\mathcal{M}^{(1)}$  in Algorithm 1 in Fig. 7 using sample averaging [33]. We employ the model-4 with a network size of 5000. We measure the quality of the estimation via the  $l_2$  norm as follows:

$$\text{estimation error} = \frac{1}{|\mathcal{A}|} \sum_{a \in \mathcal{A}} \|\mathbf{P}_a - \hat{\mathbf{P}}_a\|_2. \quad (14)$$

Numerical simulations show that an estimation error below 0.05, which requires around 1800 samples, can produce a sufficiently accurate  $\hat{\mathbf{P}}$  in Algorithm 1, leading to less than 10% APE in Algorithm 2. We observe that a relatively small number of samples is sufficient for estimating  $\mathbf{P}$  to achieve a small APE, which can be attributed to the sparsity and several structural properties of the underlying PTT of the network model (as in

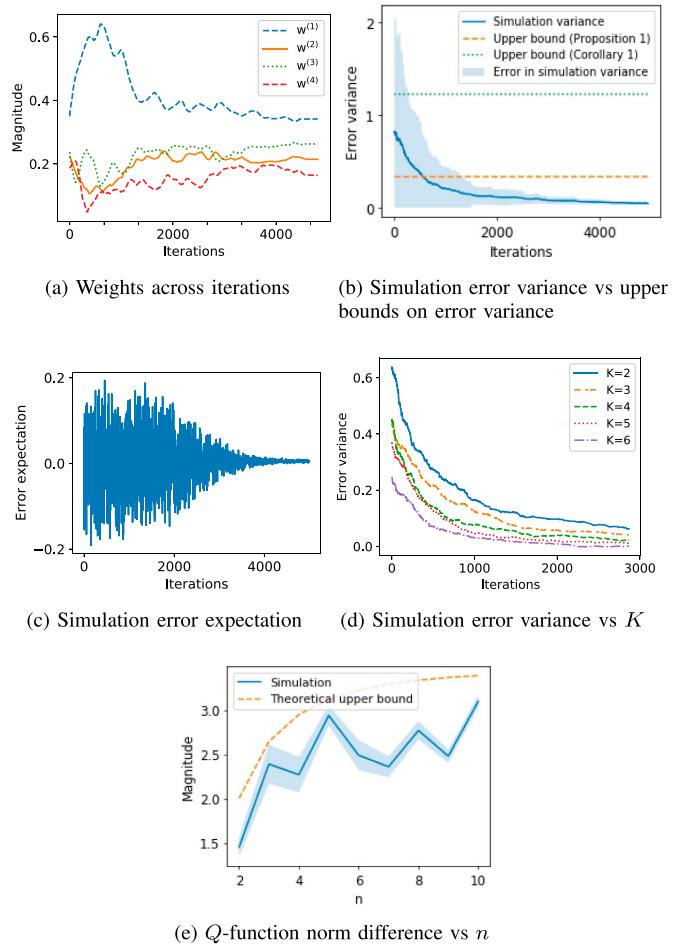


Fig. 8. Numerical validation of assumptions for theoretical results.

[22]). These features also allow us to use shorter trajectories ( $l \approx 10$ ) and select random initial states more frequently in Algorithm 1, thereby further accelerating the exploration of the entire state-space.

#### D. Convergence of Weights

We herein demonstrate the changing weights over time using the simulation settings from Section IV-B in Fig. 8(a). There is a sharp increase in  $\mathbf{w}^{(1)}$  at the beginning because  $\mathcal{M}^{(1)}$  is the original environment, and there are not enough samples to capture the higher-order relationships. Moreover, it is not clear which  $\mathcal{M}^{(n)}$  provides the most useful samples as the weights  $\mathbf{w}^{(n)}$  for  $n > 1$  keep changing. As iterations continue,  $\mathbf{w}^{(1)}$  decreases up to some point, but  $\mathcal{M}^{(1)}$  remains the most useful environment. The weights  $\mathbf{w}^{(n)}$  for  $n > 1$  increase and converge to a fixed value. The final magnitudes of  $\mathbf{w}^{(n)}$  are non-monotonic across  $n$  ( $\mathbf{w}^{(1)} > \mathbf{w}^{(3)} > \mathbf{w}^{(2)} > \mathbf{w}^{(4)}$ ), validating the partial ordering of Proposition 4. We note that similar weight patterns can be also shown across different network parameters and models, although (i) the final values of the weights, (ii) the order of environment utilities, and (iii) the iteration index at which the weights converge may vary.

### E. Numerical Validation of Propositions

In this section, we simulate the results in the propositions and compare theoretical results with simulation results. The same simulation settings in Section IV-B are employed.

We compute the simulation variance of the  $Q$ -function errors numerically as follows:

$$\mathbb{V}[\mathcal{E}_t(s, a)] \approx \frac{1}{2\Delta_t} \sum_{t'=-\Delta_t}^{t+\Delta_t} \mathcal{E}_t(s, a)^2 - \left[ \frac{1}{2\Delta_t} \sum_{t'=-\Delta_t}^{t+\Delta_t} \mathcal{E}_t(s, a) \right]^2, \quad (15)$$

with  $\Delta_t \ll t$ . The upper bounds on the variance from Proposition 1 and Corollary 2 (with  $u_t = 0.5$  and  $\lambda = 1$ ) and the simulation variance (with  $(s, a) = (6, 2)$  and  $\Delta_t = 20$ ) are shown in Fig. 8(b), where the blue-shaded region represents the simulation error. As iterations continue, the simulation variance becomes smaller than the upper bound from Proposition 1; hence, using a time-varying  $u_t$  produces more accurate results. The simulation variance eventually converges to zero, which is in line with Corollary 3. Herein,  $\lambda$  is chosen to be 1, and changing  $\lambda$  may affect the initial error variance and its decay rate. Furthermore, the simulation bias (i.e., the expectation of  $\mathcal{E}_t(6, 2)$ ) is shown in Fig. 8(c), where the blue-shaded area represents the simulation error. The simulation bias eventually converges to zero, which is consistent with Fig. 8(b) and shows the accuracy of the assumption (8).

The simulation error variances across different  $K$  are shown in Fig. 8(d). For each  $K$ , the simulation variance reduces and converges to zero eventually. The convergence may require more than 3000 iterations, particularly with small  $K$ , due to slower convergence as indicated by the theoretical runtime complexity in Section IV-C. Hence, small non-zero error variances are observed at  $t = 3000$ . At a fixed iteration index, the simulation variance also reduces as  $K$  increases, which is in line with Proposition 2. We also observe the diminishing return of increasing  $K$ . In particular, the reduction in error variance from  $K$  systems to  $K + 1$  systems becomes less significant as  $K$  increases. This implies that increasing  $K$  may not always yield substantial improvements because increasing  $K$  results in new environments that exhibit high structural similarity with those of lower; the corresponding PTTs converge to fixed tensors. Consequently, sampling from these new environments may not provide novel samples. Furthermore, a larger  $K$  leads to a reduction in the weights assigned to existing environments, which diminishes the impact of each individual environment on the ensemble estimate.

Both the upper bound on the  $l_2$  norm difference between the  $Q$ -function vectors of the original environment and  $n^{th}$  environment under the output policy of Algorithm 2 as a function of  $n \in [2, 10]$  (from Proposition 3) and the actual norm difference is illustrated in Fig. 8(e), where the blue curve is the mean, and the blue shaded area is the standard deviation of the results over 20 simulations. The norm difference, which represents the utility of the  $n^{th}$  order environment, exhibits a non-monotonic behavior across different  $n$ , and the results align with Proposition 4. We also observe that the upper

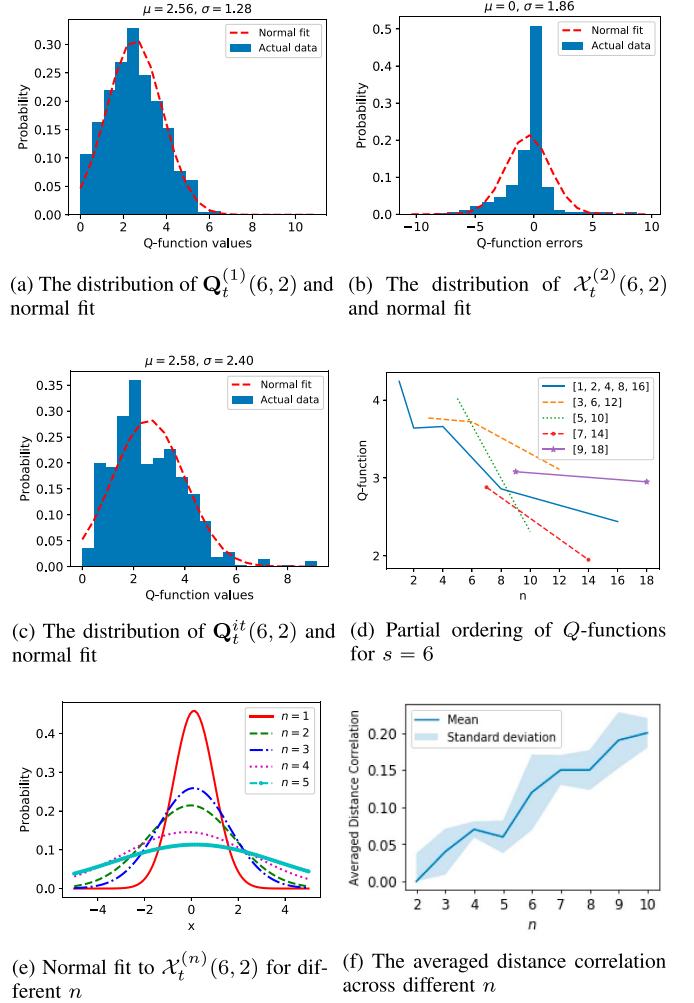


Fig. 9. Numerical validation of assumptions in theoretical results.

bound converges as  $n$  increases, yet the tightness of the bound depends on  $n$ .

The estimated distributions of  $Q_t^{(1)}(s, a)$ ,  $X_t^{(n)}(s, a)$  and  $Q_t^{it}(s, a)$  over time (until convergence) for  $(s, a, n) = (6, 2, 2)$  are shown in Fig. 9(a), Fig. 9(b) and Fig. 9(c), respectively. These distributions can be accurately modeled using normal distributions with corresponding means and variances. Furthermore, we observe that  $X_t^{(2)}(6, 2)$  has zero-mean, which shows the accuracy of the zero-mean assumption (7) and thus also the relaxed assumption (8). Similarly, Fig. 9(e), which shows the normal fitting of Fig. 9(b) duplicated for various values of  $n$ , demonstrates that the  $Q$ -function errors of different environments can be well-modeled by zero-mean normal distributions with different variances. This figure further validates the assumptions (7) and (8). Similar results are also achieved for different  $(s, a)$  pairs.

The  $Q$ -function of  $s = 6$  of the  $n^{th}$  environment under the output policy of Algorithm 2 is shown in Fig. 9(d) for  $n \in [1, 20]$ . The simulation is carried out with  $\gamma = 1 - 10^{-5}$ . We observe that partial orderings given in Proposition 4 hold. For example we have  $Q_{\hat{\pi}}^{(1)}(6) \geq Q_{\hat{\pi}}^{(2)}(6) \geq Q_{\hat{\pi}}^{(4)}(6) \geq Q_{\hat{\pi}}^{(8)}(6)$  and

TABLE II  
OPTIMIZED HYPER-PARAMETER VALUES OF ALGORITHM 2

Params	Small Networks	Modest-Sized Networks	Large Networks
$l$	$l \in [1, 5]$	$l \in [5, 10]$	$l \in [10, 20]$
$K$	$K \in \{2, 3\}$	$K \in \{3, 4, 5\}$	$K \in \{5, 6, 7, 8\}$
$\alpha_t$	$c_1 \in \{10^2, 5 \cdot 10^2\}$	$c_1 \in \{10^2, 10^3\}$	$c_1 \in \{10^3, 10^4\}$
$\epsilon_t$	$c_2 \in \{0.9, 0.95\}$ $c_3 \in \{0.01, 0.1\}$	$c_2 \in \{0.95, 0.99\}$ $c_3 \in \{0.01, 0.05\}$	$c_2 \in \{0.99, 0.999\}$ $c_3 \in \{0.005, 0.01\}$
$u_t$	$c_4 \in \{10^2, 5 \cdot 10^2\}$	$c_4 \in \{10^2, 10^3\}$	$c_4 \in \{5 \cdot 10^3, 10^4\}$

$\mathbf{Q}_{\hat{\pi}}^{(3)}(6) \geq \mathbf{Q}_{\hat{\pi}}^{(6)}(6) \geq \mathbf{Q}_{\hat{\pi}}^{(12)}(6)$  but the relationships between the  $Q$ -functions of different partial ordering groups are non-monotonic across  $n$  such as  $\mathbf{Q}_{\hat{\pi}}^{(6)}(6) \geq \mathbf{Q}_{\hat{\pi}}^{(4)}(6)$  and  $\mathbf{Q}_{\hat{\pi}}^{(7)}(6) \geq \mathbf{Q}_{\hat{\pi}}^{(16)}(6)$ . Extensive simulations demonstrate that similar orderings hold for all states  $s$ .

To assess the practicality of the assumption employed in Proposition 1, we compute the averaged distance correlation (ADC) [46] between  $\mathcal{X}_{t_1}^{(n)}(6, 2)$  and  $\mathcal{X}_{t_2}^{(n)}(6, 2)$  with  $n \in [2, 10]$  and averaging the results over all  $t_1 \neq t_2 \in [0, 5000]$ . ADC is particularly used because it captures both linear and non-linear correlations, is robust to outliers, and does not assume any particular distribution of the variables like Pearson's correlation coefficient. The results are shown in Fig. 9(f). For example, for  $n = 2$ , the ADC is almost 0, indicating that the  $Q$ -function errors of the 2<sup>nd</sup> environment at different times are independent. When  $n$  is modestly large, the ADC is sufficiently small to infer that the correlation is not statistically significant, and the independence assumption holds. However, for very large values of  $n$ , weak correlations may emerge, suggesting that very large values of  $n$  might not provide the best performance.

There are several reasons that explain the independence of the  $Q$ -function errors at different times: (i) Each environment is Markovian; hence, the errors in  $Q$ -function estimates can be considered independent. (ii) Each environment is stationary (with constant transition probabilities and reward distributions over time), and the errors in  $Q$ -function estimates can be assumed independent across different time steps. (iii) The epsilon-greedy policy introduces randomness during the action selection process. (iv) The time-varying learning rate ( $\alpha_t$ ), epsilon probability ( $\epsilon_t$ ) and the update ratio ( $u_t$ ) introduces time-dependent randomness and variability.

#### F. Parameter Tuning

The parameters of Algorithm 2 require fine-tuning across various settings. We discuss how to select  $v, l, K, \alpha_t, \epsilon_t$  for small ( $|\mathcal{S}| \leq 10^3$ ), modest-sized ( $|\mathcal{S}| \in [10^3, 10^4]$ ), and large networks ( $|\mathcal{S}| \geq 10^4$ ). Table II summarizes the parameters that yield near-optimal APE values in our numerical results. However, optimal parameter selection requires cross-validation.

The number of visits ( $v$ ) needed for each state-action pair is independent of network size. The proposed algorithm allows a small  $v$ , around  $v \approx 40$ , to achieve near-optimal performance. Additionally, the proposed algorithm enables us to have short

trajectories ( $l$ ) while ensuring: (i) near-optimal APE performance, (ii) minimized runtime and computations per trajectory, (iii) sufficient capture of samples from neighboring states, (iv) preservation of initial state importance despite discounting in long trajectories, and (v) prevention of following redundant paths and loops that provide no new samples.

Increasing the number of Markovian environments ( $K$ ) reduces the average runtime complexity of the algorithm. However, there are drawbacks to consider: (i) emphasizing high-order node relationships may lead to the loss of low-order node relationships, (ii)  $\hat{\mathbf{P}}^{(n)}$  converges to a fixed tensor as  $n$  increases, rendering samples from corresponding SMEs redundant and potentially degrading performance, (iii) a diminishing return is observed in the error variance (as shown in Fig. 8(d)), and (iv) memory requirements increase. Hence,  $K$  should be small enough to avoid these drawbacks, while also increasing with the network size to prevent a rise in other parameters (specifically  $v, l$ ) and reduce runtime complexity.

The learning rate ( $\alpha_t$ ) must adhere to the convergence conditions of  $Q$ -learning [24] and have a suitable decay to adjust the learning speed. We assume the form  $\alpha_t = \frac{1}{1+t/c_1}$ , where  $c_1 > 0$  determines the decay rate and should increase with  $|\mathcal{S}|$ ,  $v$ ,  $l$ , and  $K$ . On the other hand, the parameter  $\epsilon_t$  is essential for balancing exploration and exploitation. We use the form  $\epsilon_t = \max((c_2)^t, c_3)$ , where  $c_2 > 0$  adjusts the decay rate, and  $0 < c_3 \ll 1$  determines the minimum exploration probability. As the system parameters increase,  $c_2$  should also increase to ensure sufficient exploration. Furthermore,  $c_3$  should be small and positive, allowing for exploration with a low probability when the policy is nearly converged.

The parameter  $u_t$  is time-varying such that  $u_t \xrightarrow{t \rightarrow \infty} 1$  to balance between exploration and exploitation as explained after Corollary 3. To this end, there are several possible choices for the structure of  $u_t$ , including  $u_t = 1 - e^{\frac{-t}{c_4}}$  and  $u_t = 1 - \frac{1}{1 + \frac{t}{c_4}}$  with  $c_4 > 0$  or  $u_t = 1 - (c_4)^t$  with  $c_4 \in (0, 1)$ . Numerical results show that the performance of the algorithm is maximized for the network models we considered with  $u_t = 1 - e^{\frac{-t}{c_4}}$  with optimized  $c_4$ . In particular, this specific choice leads to 15% less APE with 20% less runtime complexity than the other two alternative choices using the simulation settings in Section IV-B. Additionally, this choice of  $u_t$  maintains stable performance even when network parameters change significantly, making it a robust choice.

#### G. Memory Complexity

The memory needs of the proposed algorithm increase with  $K$  as the  $Q$ -functions of  $K$  different Markovian environments need to be stored in tables. To alleviate this challenge, several strategies can be employed, including a linear and non-linear approximation of  $Q$ -functions or state aggregation methods [47], [48], [49], [50]. In our prior work [22], we proposed a special state-action aggregation algorithm to handle the increasing memory complexity of a model-free learning algorithm, similar to Algorithm 2, where  $Q$ -functions of state-action pairs that incur the same cost are grouped into a single  $Q$ -function. The algorithm is applicable to networks that have bounded and

smoothly changing cost functions (as in model-2, 3, and 4) and can significantly reduce the memory needs while negligibly increasing the APE.

We note that more sophisticated function approximation algorithms can also be designed to work with any kind of network; however, they require the design of optimal basis functions, selection of appropriate features, or design of aggregation schemes, which are out of the scope of this paper.

## V. CONCLUSION

In this paper, we presented a novel ensemble  $Q$ -learning algorithm to overcome the performance and complexity challenges of the original  $Q$ -learning across large MDPs. The proposed algorithm employs multiple  $Q$ -function estimators on multiple, distinct, synthetically created, and structurally related Markovian environments that run at different time-scales and fuses the outputs into a single estimate based on an adaptive weighting mechanism using Jensen-Shannon divergence. Extensive simulations across a variety of real-world networks show that the proposed algorithm produces a near-optimal policy with significantly lower complexity and outperforms other  $Q$ -learning algorithms in terms of accuracy and complexity. Several theoretical upper bounds on the error variance are given, and it is shown that the error variance decreases with the number of Markovian environments. An upper bound on the difference between the  $Q$ -functions of different environments is also given. In the end, it is shown that the simulation results closely follow the theoretical results. Our analysis provides initial insights and justifications for the advantages of leveraging multiple synthetic Markovian environments in ensemble reinforcement learning. However, there are areas for further improvement and understanding of the nEQL algorithm. Specifically, our ongoing work focuses on devising computationally efficient sampling methods to ensure good data coverage (as described in [51]) across all Markovian environments. Additionally, we are exploring alternative approaches to construct synthetic systems that go beyond relying on  $n$ -hop environments to further improve the generalizability and practicality of our algorithm. Finally, we are working on the extension of the proposed algorithm to continuous state-spaces.

## APPENDIX

### A. Proof of Proposition 1

The following expressions are valid for all  $(s, a)$  pairs; hence, we drop the  $(s, a)$  notation for simplicity. We first prove the expectation.

$$\lim_{t \rightarrow \infty} \mathcal{E}_t = \lim_{t \rightarrow \infty} \mathbf{Q}_t^{it} - \mathbf{Q}^*. \quad (16)$$

$$= \lim_{t \rightarrow \infty} (1-u) \sum_{i=0}^{t-1} u^{t-i-1} \sum_{n=1}^K \mathbf{w}_i^{(n)} \mathbf{Q}_i^{(n)} - \mathbf{Q}^*. \quad (17)$$

$$= \lim_{t \rightarrow \infty} (1-u) \sum_{i=0}^{t-1} u^{t-i-1} \sum_{n=1}^K \mathbf{w}_i^{(n)} (\mathbf{Q}_i^{(n)} - \mathbf{Q}^*). \quad (18)$$

$$= \lim_{t \rightarrow \infty} (1-u) \sum_{i=0}^{t-1} u^{t-i-1} \sum_{n=1}^K \mathbf{w}_i^{(n)} \mathcal{X}_i^{(n)}, \quad (19)$$

where (17) follows from the explicit expression for  $\mathbf{Q}_t^{it}$ , which can be obtained by repeatedly plugging the expression of  $\mathbf{Q}_{t-1}^{it}$  in  $\mathbf{Q}_t^{it}$  in line 11 in Algorithm 2, (18) follows from the fact that  $\sum_{n=1}^K \mathbf{w}_t^{(n)} = 1$  for all  $t$ , and  $(1-u) \sum_{i=0}^{t-1} u^{t-i-1} = 1$  as  $t \rightarrow \infty$ , and (19) follows from (7). If we take the expectation of both sides:

$$\lim_{t \rightarrow \infty} \mathbb{E}[\mathcal{E}_t] = \lim_{t \rightarrow \infty} (1-u) \sum_{i=0}^{t-1} u^{t-i-1} \sum_{n=1}^K \mathbf{w}_i^{(n)} \mathbb{E}[\mathcal{X}_i^{(n)}] = 0 \quad (20)$$

which follows from the linearity of expectation and (7).

We now prove the upper bound on the variance.

$$\begin{aligned} \lim_{t \rightarrow \infty} \mathbb{V}[\mathcal{E}_t] &= \lim_{t \rightarrow \infty} \mathbb{V} \left[ (1-u) \sum_{i=0}^{t-1} u^{t-i-1} \sum_{n=1}^K \mathbf{w}_i^{(n)} \mathcal{X}_i^{(n)} \right]. \end{aligned} \quad (21)$$

$$\begin{aligned} &= \lim_{t \rightarrow \infty} (1-u)^2 \left[ \sum_{i=0}^{t-1} u^{2(t-i-1)} \left[ \sum_{n=1}^K (\mathbf{w}_i^{(n)})^2 \mathbb{V}[\mathcal{X}_i^{(n)}] \right. \right. \\ &\quad \left. \left. + 2 \sum_{n=1}^K \sum_{m=n+1}^K \mathbf{w}_i^{(n)} \mathbf{w}_i^{(m)} \text{Cov}(\mathcal{X}_i^{(n)}, \mathcal{X}_i^{(m)}) \right] \right]. \end{aligned} \quad (22)$$

$$\begin{aligned} &\leq \lim_{t \rightarrow \infty} (1-u)^2 \left[ \sum_{i=0}^{t-1} u^{2(t-i-1)} \left[ \sum_{n=1}^K (\mathbf{w}_i^{(n)})^2 \mathbb{V}[\mathcal{X}_i^{(n)}] \right. \right. \\ &\quad \left. \left. + 2 \sum_{n=1}^K \sum_{m=n+1}^K \mathbf{w}_i^{(n)} \mathbf{w}_i^{(m)} \sqrt{\mathbb{V}[\mathcal{X}_i^{(n)}] \mathbb{V}[\mathcal{X}_i^{(m)}]} \right] \right]. \end{aligned} \quad (23)$$

$$\begin{aligned} &\leq \lim_{t \rightarrow \infty} (1-u)^2 \left[ \sum_{i=0}^{t-1} u^{2(t-i-1)} \left[ \sum_{n=1}^K \mathbf{w}_i^{(n)} \mathbb{V}[\mathcal{X}_i^{(n)}] \right. \right. \\ &\quad \left. \left. + 2 \sum_{n=1}^K \sum_{m=1}^K \mathbf{w}_i^{(n)} \mathbf{w}_i^{(m)} \sqrt{\mathbb{V}[\mathcal{X}_i^{(n)}] \mathbb{V}[\mathcal{X}_i^{(m)}]} \right] \right]. \end{aligned} \quad (24)$$

$$\begin{aligned} &\leq \lim_{t \rightarrow \infty} (1-u)^2 \left[ \sum_{i=0}^{t-1} u^{2(t-i-1)} \left[ \sum_{n=1}^K \mathbf{w}_i^{(n)} \frac{\lambda^2}{3} \right. \right. \\ &\quad \left. \left. + 2 \sum_{n=1}^K \sum_{m=1}^K \mathbf{w}_i^{(n)} \mathbf{w}_i^{(m)} \frac{\lambda^2}{3} \right] \right]. \end{aligned} \quad (25)$$

$$\leq \lim_{t \rightarrow \infty} (1-u)^2 \left[ \sum_{i=0}^{t-1} u^{2(t-i-1)} \lambda^2 \right]. \quad (26)$$

$$\leq \frac{(1-u)}{(1+u)} \lambda^2, \quad (27)$$

where (21) follows from (19), (22) follows from the properties of the variance operator and the independence assumption, (23) follows from the Cauchy-Schwarz inequality for the variance, (24) follows from the fact that  $\mathbf{w}_t^{(n)} \leq 1$  and dropping the constraint in the second summation, (25) follows from (7) and  $\lambda = \max_n \lambda_n$ , (26) follows from the fact that  $\sum_{n=1}^K \mathbf{w}_t^{(n)} = 1$  for all  $t$ , and (27) follows from the infinite geometric sum formula and  $u \in (0, 1)$ .

### B. Proof of Corollary 1

We can show the expectation result as follows:

$$\lim_{t \rightarrow \infty} \mathbb{E}[\mathcal{E}_t] = \lim_{t \rightarrow \infty} (1-u) \sum_{i=0}^{t-1} u^{t-i-1} \sum_{n=1}^K \mathbf{w}_i^{(n)} \mathbb{E}[\mathcal{X}_i^{(n)}]. \quad (28)$$

$$= \lim_{t \rightarrow \infty} (1-u) \sum_{i=0}^{t-1} u^{t-i-1} \sum_{n=1}^K \mathbf{w}_i^{(n)} \mu_n. \quad (29)$$

$$= 0, \quad (30)$$

where (28) follows from (20), (29) follows from (8), and (30) follows from the condition in Corollary 1. The proof for the result on the upper bound of the variance in Proposition 1 remains unchanged.

### C. Proof of Proposition 2

We firstly bound the weight  $\mathbf{w}^{(n)}$  into a tighter interval than  $[0,1]$ . We first do the calculations for  $n \neq 1$ . The maximum value of  $\mathbf{w}^{(n)}$  is obtained when  $\hat{\mathbf{Q}}^{(1)} = \hat{\mathbf{Q}}^{(n)}$  and  $\hat{\mathbf{Q}}^{(i)}$  is maximally different than  $\hat{\mathbf{Q}}^{(1)}$  for all  $i \neq n$ . In this case,  $\mathbf{w}^{(n)} = 1$ , and  $\mathbf{w}^{(i)} = 0$  for all  $i \notin \{n, 1\}$ . When we apply the softmax operator on the  $\mathbf{w}$ , we obtain the following:

$$\mathbf{w}^{(n)} = \frac{e}{2e + (K-2)} \leq \frac{e}{K}. \quad (31)$$

which follows as there are  $K > 1$  different weights, and  $\mathbf{w}^{(1)} = 1$ . On the other hand, the maximum value of  $\mathbf{w}^{(1)}$  is obtained when  $\hat{\mathbf{Q}}^{(i)}$  is maximally different than  $\hat{\mathbf{Q}}^{(1)}$  for all  $i \neq 1$ . In this case,  $\mathbf{w}^{(1)} = 1$ , and  $\mathbf{w}^{(i)} = 0$  for all  $i \neq 1$ . When we apply the softmax on  $\mathbf{w}$ , we obtain the following:

$$\mathbf{w}^{(1)} = \frac{e}{e + (K-1)} \leq \frac{e}{K}. \quad (32)$$

Combining (31) and (32), the following holds for all  $n$ :

$$\mathbf{w}^{(n)} \leq \frac{e}{K}. \quad (33)$$

We note that this bound is useful for  $K > 2$ . On the other hand, the minimum value of  $\mathbf{w}^{(n)}$  for  $n \neq 1$  is obtained when  $\hat{\mathbf{Q}}^{(1)}$  is maximally different than  $\hat{\mathbf{Q}}^{(n)}$  and  $\hat{\mathbf{Q}}^{(i)} = \hat{\mathbf{Q}}^{(1)}$  for all  $i \neq n$ . In this case,  $\mathbf{w}^{(n)} = 0$ , and  $\mathbf{w}^{(i)} = 1$  for all  $i \neq n$ . When we apply the softmax operator on the  $\mathbf{w}$ , we obtain the following:

$$\mathbf{w}^{(n)} = \frac{1}{1 + e(K-1)} \geq \frac{e^{-1}}{K}. \quad (34)$$

which follows because  $K > 1$ . Similarly, the minimum value of  $\mathbf{w}^{(1)}$  is obtained when  $\hat{\mathbf{Q}}^{(i)} = \hat{\mathbf{Q}}^{(1)}$  for all  $i \neq 1$ . In this case,  $\mathbf{w}^{(i)} = 1$  for all  $i$ . When we apply the softmax operator on the  $\mathbf{w}$ , we obtain the following:

$$\mathbf{w}^{(1)} = \frac{1}{K}. \quad (35)$$

Combining (34) and (35), the following holds for all  $n$ :

$$\mathbf{w}^{(n)} \geq \frac{1}{K}. \quad (36)$$

Combining (33) and (36), the weights  $\mathbf{w}^{(n)}$  can be shown to take values in the following interval:  $\mathbf{w}^{(n)} \in [\frac{1}{K}, \frac{e}{K}]$ . Then, the term  $\sum_{n=1}^K (\mathbf{w}_i^{(n)})^2$  in (22) can be upper bounded as:

$$\sum_{n=1}^K (\mathbf{w}_i^{(n)})^2 \leq \sum_{n=1}^K \left(\frac{e}{K}\right)^2 \leq \frac{e^2}{K}. \quad (37)$$

The following expressions are valid for all  $(s, a)$  pairs; hence, we drop the  $(s, a)$  notation for simplicity. Using (37), (7), and  $\lambda = \max_n \lambda_n$ , we can show the following:

$$\mathbb{V} \left[ \sum_{n=1}^K \mathbf{w}_i^{(n)} \mathcal{X}_i^{(n)} \right] \leq \sum_{n=1}^K \left(\mathbf{w}_i^{(n)}\right)^2 \mathbb{V} \left[ \mathcal{X}_i^{(n)} \right] \leq \frac{e^2}{K} \frac{\lambda^2}{3}. \quad (38)$$

Let  $u_0 = 0$  and  $\bar{u}_i = (1-u_i) \prod_{j=i+1}^{t-1} u_j$ . Using the explicit expression for  $\mathbf{Q}_t^{it}$ , which can be obtained by repeatedly plugging the expression of  $\mathbf{Q}_{t-1}^{it}$  in  $\mathbf{Q}_t^{it}$  in line 11 in Algorithm 2 and the facts that  $\sum_{i=0}^t \bar{u}_i = 1$  and  $u_0 = 0$ , the following can be derived:

$$\mathbf{Q}_t^{it} = \sum_{i=0}^{t-1} \bar{u}_i \sum_{n=1}^K \mathbf{w}_i^{(n)} \mathbf{Q}_i^{(n)}, \quad (39)$$

Then we can show the following:

$$\mathbb{V}[\mathcal{E}_t] = \mathbb{V} \left[ \sum_{i=0}^{t-1} \bar{u}_i \sum_{n=1}^K \mathbf{w}_i^{(n)} \mathcal{X}_i^{(n)} \right]. \quad (40)$$

$$= \sum_{i=0}^{t-1} \bar{u}_i^2 \mathbb{V} \left[ \sum_{n=1}^K \mathbf{w}_i^{(n)} \mathcal{X}_i^{(n)} \right] + 2 \sum_{i=0}^{t-1} \sum_{j=i+1}^{t-1} \bar{u}_i \bar{u}_j \text{Cov} \left( \sum_{n=1}^K \mathbf{w}_i^{(n)} \mathcal{X}_i^{(n)}, \sum_{n=1}^K \mathbf{w}_j^{(n)} \mathcal{X}_j^{(n)} \right). \quad (41)$$

$$\leq \sum_{i=0}^{t-1} \bar{u}_i^2 \mathbb{V} \left[ \sum_{n=1}^K \mathbf{w}_i^{(n)} \mathcal{X}_i^{(n)} \right] + 2 \sum_{i=0}^{t-1} \sum_{j=i+1}^{t-1} \bar{u}_i \bar{u}_j \times \sqrt{\mathbb{V} \left( \sum_{n=1}^K \mathbf{w}_i^{(n)} \mathcal{X}_i^{(n)} \right) \mathbb{V} \left( \sum_{n=1}^K \mathbf{w}_j^{(n)} \mathcal{X}_j^{(n)} \right)}. \quad (42)$$

$$\leq \frac{e^2 \lambda^2}{K} \left[ \sum_{i=0}^{t-1} \bar{u}_i^2 + 2 \sum_{i=0}^{t-1} \sum_{j=i+1}^{t-1} \bar{u}_i \bar{u}_j \right]. \quad (43)$$

$$\leq \frac{c(\lambda, u)}{K}, \quad (44)$$

where (40) follows from (39), (41) follows from the properties of the variance operator, (42) follows from the Cauchy-Schwarz inequality for the variance, (43) follows from (38), and (44) follows as  $c(\lambda, u)$  is a constant of  $K$ .

### D. Proof of Proposition 3

If the spectral norm of a matrix  $\mathbf{A}$  is less than 1 ( $\|\mathbf{A}\| < 1$ ), then the following can be shown [52]:

$$(\mathbf{I} - \mathbf{A})^{-1} = \sum_{i=0}^{\infty} \mathbf{A}^i, \quad (45)$$

where  $\mathbf{I}$  is the identity matrix. The  $Q$ -functions under policy  $\hat{\pi}$  can be obtained using Bellman's equation as follows [4]:

$$\mathbf{Q}_{\hat{\pi}} = (\mathbf{I} - \gamma \mathbf{P}_{\hat{\pi}})^{-1} \mathbf{c}_{\hat{\pi}}, \quad (46)$$

where  $\mathbf{P}_{\hat{\pi}}$  and  $\mathbf{c}_{\hat{\pi}}$  are PTM and cost vectors under policy  $\hat{\pi}$ , respectively. The matrix inversion in (46) can be expressed as an infinite sum using (45) since the spectral norm of  $\mathbf{P}_{\hat{\pi}}$  is 1 (as it is a PTM), and when scaled by  $\gamma$ , its spectral norm becomes strictly less than 1. Then, we can proceed as follows:

$$\mathbf{Q}_{\hat{\pi}}^{(1)} - \mathbf{Q}_{\hat{\pi}}^{(n)} = (\mathbf{I} - \gamma \mathbf{P}_{\hat{\pi}})^{-1} \mathbf{c}_{\hat{\pi}} - (\mathbf{I} - \gamma \mathbf{P}_{\hat{\pi}}^n)^{-1} \mathbf{c}_{\hat{\pi}}. \quad (47)$$

$$\begin{aligned} &= \left[ \gamma \mathbf{P}_{\hat{\pi}} \sum_{i=0}^{\infty} (\gamma \mathbf{P}_{\hat{\pi}})^{ni} \sum_{k=0}^{n-2} (\gamma \mathbf{P}_{\hat{\pi}})^k \right. \\ &\quad \left. + \sum_{j=0}^{\infty} (\gamma^{nj} - \gamma^j) \mathbf{P}_{\hat{\pi}}^{nj} \right] \mathbf{c}_{\hat{\pi}}. \end{aligned} \quad (48)$$

$$\leq \gamma \mathbf{P}_{\hat{\pi}} \sum_{i=0}^{\infty} (\gamma \mathbf{P}_{\hat{\pi}})^{ni} \sum_{k=0}^{n-2} (\gamma \mathbf{P}_{\hat{\pi}})^k \mathbf{c}_{\hat{\pi}} \quad (49)$$

where (47) follows from (46), (48) follows from the fact that  $\gamma \|\mathbf{P}_{\hat{\pi}}\| < 1$ , (45) and grouping the common terms in the expansions, and (49) follows from the facts that  $\gamma < 1$  and the elements of  $\mathbf{P}_{\hat{\pi}}$  and that  $\mathbf{c}_{\hat{\pi}}$  are non-negative. If we take the norm of both sides, we obtain:

$$\|\mathbf{Q}_{\hat{\pi}}^{(1)} - \mathbf{Q}_{\hat{\pi}}^{(n)}\| \leq \|\gamma \mathbf{P}_{\hat{\pi}} \sum_{i=0}^{\infty} (\gamma \mathbf{P}_{\hat{\pi}})^{ni} \sum_{k=0}^{n-2} (\gamma \mathbf{P}_{\hat{\pi}})^k \mathbf{c}_{\hat{\pi}}\|. \quad (50)$$

$$\leq \gamma \|\mathbf{P}_{\hat{\pi}}\| \sum_{i=0}^{\infty} (\gamma \|\mathbf{P}_{\hat{\pi}}\|)^{ni} \sum_{k=0}^{n-2} (\gamma \|\mathbf{P}_{\hat{\pi}}\|)^k \|\mathbf{c}_{\hat{\pi}}\|. \quad (51)$$

$$< \gamma \sum_{i=0}^{\infty} \gamma^{ni} \sum_{k=0}^{n-2} \gamma^k \|\mathbf{c}_{\hat{\pi}}\|. \quad (52)$$

$$< \frac{\gamma}{1 - \gamma^n} \frac{1 - \gamma^{n-1}}{1 - \gamma} \|\mathbf{c}_{\hat{\pi}}\|, \quad (53)$$

where (50) follows from (49), (51) follows from the upper bound on the norm of a matrix-vector product [54], (52) follows from the fact that  $\gamma \|\mathbf{P}_{\hat{\pi}}\| < 1$ , and (53) follows from the finite and infinite geometric sum formulas.

### E. Proof of Proposition 4

We analyze  $\mathbf{Q}_{\hat{\pi}}^{(n)}$  when  $\gamma \rightarrow 1$ .

$$\lim_{\gamma \rightarrow 1} \mathbf{Q}_{\hat{\pi}}^{(n)} = \lim_{\gamma \rightarrow 1} (\mathbf{I} - \gamma \mathbf{P}_{\pi}^n)^{-1} \mathbf{c}_{\pi} \quad (54)$$

$$= (\mathbf{I} - \mathbf{P}_{\pi}^n)^{-1} \mathbf{c}_{\pi} \quad (55)$$

$$= \mathbf{c}_{\pi} + \mathbf{P}_{\pi}^n \mathbf{c}_{\pi} + (\mathbf{P}_{\pi}^n)^2 \mathbf{c}_{\pi} + (\mathbf{P}_{\pi}^n)^3 \mathbf{c}_{\pi} + \dots, \quad (56)$$

which follows from (45). Then, we can write the following:

$$\lim_{\gamma \rightarrow 1} \mathbf{Q}_{\hat{\pi}}^{(1)} = \mathbf{c}_{\pi} + \mathbf{P}_{\pi} \mathbf{c}_{\pi} + \mathbf{P}_{\pi}^2 \mathbf{c}_{\pi} + \mathbf{P}_{\pi}^3 \mathbf{c}_{\pi} + \dots \quad (57)$$

$$\lim_{\gamma \rightarrow 1} \mathbf{Q}_{\hat{\pi}}^{(2)} = \mathbf{c}_{\pi} + \mathbf{P}_{\pi}^2 \mathbf{c}_{\pi} + \mathbf{P}_{\pi}^4 \mathbf{c}_{\pi} + \mathbf{P}_{\pi}^6 \mathbf{c}_{\pi} + \dots \quad (58)$$

$$\lim_{\gamma \rightarrow 1} \mathbf{Q}_{\hat{\pi}}^{(3)} = \mathbf{c}_{\pi} + \mathbf{P}_{\pi}^3 \mathbf{c}_{\pi} + \mathbf{P}_{\pi}^6 \mathbf{c}_{\pi} + \mathbf{P}_{\pi}^9 \mathbf{c}_{\pi} + \dots \quad (59)$$

We observe that  $\mathbf{Q}_{\hat{\pi}}^{(1)}$  is the largest since it has all the terms that  $\mathbf{Q}_{\hat{\pi}}^{(n)}$  contains for all  $n \neq 1$ . We also have the following partial orderings in the limiting case ( $\gamma \rightarrow 1$ ):

$$\mathbf{Q}_{\hat{\pi}}^{(2)} \geq \mathbf{Q}_{\hat{\pi}}^{(4)} \geq \mathbf{Q}_{\hat{\pi}}^{(8)} \geq \mathbf{Q}_{\hat{\pi}}^{(16)} \dots \quad (60)$$

$$\mathbf{Q}_{\hat{\pi}}^{(3)} \geq \mathbf{Q}_{\hat{\pi}}^{(6)} \geq \mathbf{v}_{\pi}^{(12)} \dots \quad (61)$$

$$\mathbf{Q}_{\hat{\pi}}^{(5)} \geq \mathbf{Q}_{\hat{\pi}}^{(10)} \geq \mathbf{Q}_{\hat{\pi}}^{(20)} \dots \quad (62)$$

$$\mathbf{Q}_{\hat{\pi}}^{(7)} \geq \mathbf{Q}_{\hat{\pi}}^{(14)} \geq \mathbf{Q}_{\hat{\pi}}^{(28)} \dots, \quad (63)$$

which follow as the elements of  $\mathbf{P}_{\hat{\pi}}$  and  $\mathbf{c}_{\hat{\pi}}$  are non-negative.

### ACKNOWLEDGMENT

The authors thank Prof. Tara Javidi of UCSD for ongoing discussions regarding this work and for suggesting that we consider multi-hop transition matrices based on our work on co-link approximations [45].

### REFERENCES

- Y. Chen, Q. Zhao, and A. Swami, "Distributed spectrum sensing and access in cognitive radio networks with energy constraint," *IEEE Trans. Signal Process.*, vol. 57, no. 2, pp. 783–797, Feb. 2009.
- F. Zhang and V. K. N. Lau, "Closed-form delay-optimal power control for energy harvesting wireless system with finite energy storage," *IEEE Trans. Signal Process.*, vol. 62, no. 21, pp. 5706–5715, Nov. 2014.
- S. Chai and V. K. N. Lau, "Online trajectory and radio resource optimization of cache-enabled UAV wireless networks with content and energy recharging," *IEEE Trans. Signal Process.*, vol. 68, pp. 1286–1299, 2020.
- D. Bertsekas, *Reinforcement Learning and Optimal Control*. Belmont, MA, USA: Athena Scientific, 2019.
- R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA, USA: MIT Press, 2018.
- D. V. Djordjević and V. Krishnamurthy, "Q-learning algorithms for constrained Markov decision processes with randomized monotone policies: Application to MIMO transmission control," *IEEE Trans. Signal Process.*, vol. 55, no. 5, pp. 2170–2181, May 2007.
- S. Kar, J. M. F. Moura, and H. V. Poor, "QD-learning: A collaborative distributed strategy for multi-agent reinforcement learning through consensus + innovations," *IEEE Trans. Signal Process.*, vol. 61, no. 7, pp. 1848–1862, Apr. 2013.
- M. H. Ngo and V. Krishnamurthy, "Monotonicity of constrained optimal transmission policies in correlated fading channels with ARQ," *IEEE Trans. Signal Process.*, vol. 58, no. 1, pp. 438–451, Jan. 2010.
- Z. Hajikhondi-Meybodi, A. Mohammadi, M. Hou, and K. N. Plataniotis, "DQEL: Deep Q-learning for energy-optimized LoS/NLoS UWB node selection," *IEEE Trans. Signal Process.*, vol. 70, pp. 2532–2547, 2022.
- H. Hasselt, "Double Q-learning," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 23, pp. 2613–2621, 2010.
- C. D'Eramo, M. Restelli, and A. Nuara, "Estimating maximum expected value through Gaussian approximation," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2016, pp. 1032–1040.
- O. Anschel, N. Baram, and N. Shimkin, "Averaged-DQN: Variance reduction and stabilization for deep reinforcement learning," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2017, pp. 176–185.
- X. Chen, C. Wang, Z. Zhou, and K. W. Ross, "Randomized ensembled double Q-learning: Learning fast without a model," 2021. [Online]. Available: <https://arxiv.org/abs/2101.05982>
- M. Ghavamzadeh, H. Kappen, M. Azar, and R. Munos, "Speedy Q-learning," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 24, pp. 2411–2419, 2011.
- M. Riedmiller, "Neural fitted Q iteration–first experiences with a data efficient neural reinforcement learning method," in *Proc. Eur. Conf. Mach. Learn.*, Springer, 2005, pp. 317–328.
- L. Liu and U. Mitra, "On sampled reinforcement learning in wireless networks: Exploitation of policy structures," *IEEE Trans. Commun.*, vol. 68, no. 5, pp. 2823–2837, May 2020.

[17] L. Liu, A. Chattopadhyay, and U. Mitra, "On solving MDPs with large state space: Exploitation of policy structures and spectral properties," *IEEE Trans. Commun.*, vol. 67, no. 6, pp. 4151–4165, Jun. 2019.

[18] Q. Lan, Y. Pan, A. Fyshe, and M. White, "Maxmin Q-learning: Controlling the estimation bias of Q-learning," 2020. [Online]. Available: <https://arxiv.org/abs/2002.06487>

[19] O. Peer, C. Tessler, N. Merlis, and R. Meir, "Ensemble bootstrapping for Q-learning," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2021, pp. 8454–8463.

[20] S. Amin, M. Gomrokchi, H. Satija, H. van Hoof, and D. Precup, "A survey of exploration methods in reinforcement learning," 2021. [Online]. Available: <https://arxiv.org/abs/2109.00157>

[21] R. McFarlane, "A survey of exploration strategies in reinforcement learning," McGill Univ., Montreal, QC, Canada, Rep. 3, pp. 17–18, 2018.

[22] T. Bozkus and U. Mitra, "Ensemble link learning for large state space multiple access communications," in *Proc. 30th Eur. Signal Process. Conf. (EUSIPCO)*, 2022, pp. 747–751.

[23] T. Bozkus and U. Mitra, "Ensemble graph Q-learning for large scale networks," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, 2023, pp. 1–5.

[24] F. S. Melo, "Convergence of Q-learning: A simple proof," Institute of Systems and Robotics, Lisbon, Portugal, Tech. Rep., pp. 1–4, 2001. [Online]. Available: <http://users.isr.ist.utl.pt/~mtjspaan/readingGroup/ProofQlearning.pdf>

[25] A. Kumar, J. Fu, M. Soh, G. Tucker, and S. Levine, "Stabilizing off-policy Q-learning via bootstrapping error reduction," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 32, pp. 11784–11794, 2019.

[26] R. Agarwal, D. Schuurmans, and M. Norouzi, "An optimistic perspective on offline reinforcement learning," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2020, pp. 104–114.

[27] S. Lee, Y. Seo, K. Lee, P. Abbeel, and J. Shin, "Offline-to-online reinforcement learning via balanced replay and pessimistic Q-ensemble," in *Proc. Conf. Robot Learn. (PMLR)*, 2022, pp. 1702–1712.

[28] A. Modi, N. Jiang, A. Tewari, and S. Singh, "Sample complexity of reinforcement learning using linearly combined model ensembles," in *Proc. Int. Conf. Artif. Intell. Statist. (PMLR)*, 2020, pp. 2010–2020.

[29] T. Kurutach, I. Clavera, Y. Duan, A. Tamar, and P. Abbeel, "Model-ensemble trust-region policy optimization," 2018, *arXiv:1802.10592*.

[30] K. Chua, R. Calandra, R. McAllister, and S. Levine, "Deep reinforcement learning in a handful of trials using probabilistic dynamics models," in *Proc. Adv. in Neural Inf. Process. Syst.*, vol. 31, pp. 4759–4770, 2018.

[31] Y. Yao, L. Xiao, Z. An, W. Zhang, and D. Luo, "Sample efficient reinforcement learning via model-ensemble exploration and exploitation," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, Piscataway, NJ, USA: IEEE Press, 2021, pp. 4202–4208.

[32] D. Pathak, D. Gandhi, and A. Gupta, "Self-supervised exploration via disagreement," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2019, pp. 5062–5071.

[33] B. A. Craig and P. P. Sendi, "Estimation of the transition matrix of a discrete-time Markov chain," *Health Econ.*, vol. 11, no. 1, pp. 33–42, 2002.

[34] M. Awiszus and B. Rosenhahn, "Markov chain neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, 2018, pp. 2180–2187.

[35] C. Sherlaw-Johnson, S. Gallivan, and J. Burridge, "Estimating a Markov transition matrix from observational data," *J. Oper. Res. Soc.*, vol. 46, no. 3, pp. 405–410, 1995.

[36] J. Tugnait, "Adaptive estimation and identification for discrete systems with Markov jump parameters," *IEEE Trans. Autom. Control*, vol. AC-27, no. 5, pp. 1054–1065, Oct. 1982.

[37] J. Lin, "Divergence measures based on the Shannon entropy," *IEEE Trans. Inf. Theory*, vol. 37, no. 1, pp. 145–151, Jan. 1991.

[38] T. Bozkus and U. Mitra, "Supplementary appendix," 2023. [Online]. Available: <https://github.com/talhabozkus/Multi-Timescale-Ensemble-QLearning>

[39] Y. Song, Y. Zhou, A. Sekhari, J. A. Bagnell, A. Krishnamurthy, and W. Sun, "Hybrid RL: Using both offline and online data can make RL efficient," 2022, *arXiv:2210.06718*.

[40] A. Nair, M. Dalal, A. Gupta, and S. Levine, "Accelerating online reinforcement learning with offline datasets," 2020. [Online]. Available: <https://arxiv.org/abs/2006.09359>

[41] T. Xie, N. Jiang, H. Wang, C. Xiong, and Y. Bai, "Policy fine-tuning: Bridging sample-efficient offline and online reinforcement learning," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 34, 2021, pp. 27395–27407.

[42] S. Thrun and A. Schwartz, "Issues in using function approximation for reinforcement learning," in *Proc. Connectionist Models Summer School*, Hillsdale, NJ, USA: Lawrence Erlbaum, vol. 6, 1993, pp. 1–9.

[43] H. Van Hasselt, A. Guez, and D. Silver, "Deep reinforcement learning with double Q-learning," in *Proc. AAAI Conf. Artif. Intell.*, vol. 30, no. 1, pp. 2094–2100, 2016.

[44] C. D'Eramo et al., "Gaussian approximation for bias reduction in Q-learning," *J. Mach. Learn. Res.*, vol. 22, no. 1, pp. 12690–12740, 2021.

[45] T. Bozkus and U. Mitra, "Link analysis for solving multiple-access MDPs with large state spaces," *IEEE Trans. Signal Process.*, vol. 71, pp. 947–962, 2023.

[46] G. J. Székely, M. L. Rizzo, and N. K. Bakirov, "Measuring and testing dependence by correlation of distances," *Ann. Statist.*, vol. 35, no. 6, pp. 2769–2794, 2007, doi: 10.1214/009053607000000505.

[47] F. S. Melo and M. I. Ribeiro, "Q-learning with linear function approximation," in *Proc. 20th Annu. Conf. Learn. Theory (COLT)*, San Diego, CA, USA: Springer, Jun. 13–15, 2007, pp. 308–322.

[48] V. Mnih et al., "Playing atari with deep reinforcement learning," 2013, *arXiv:1312.5602*.

[49] D. Russo, "Approximation benefits of policy gradient methods with aggregated states," *Manage. Sci.*, vol. 69, no. 11, pp. 6898–6911, 2023.

[50] L. Li, T. J. Walsh, and M. L. Littman, "Towards a unified theory of state abstraction for MDPs," in *Proc. Int. Symp. Artif. Intell. Math. (AI&Math)*, 2006, pp. 531–539.

[51] P. Rashidinejad, B. Zhu, C. Ma, J. Jiao, and S. Russell, "Bridging offline reinforcement learning and imitation learning: A tale of pessimism," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 34, 2021, pp. 11702–11716.

[52] G. Strang, *Linear Algebra and Its Applications*. Belmont, CA, USA: Brooks/Cole, 2006.

**Talha Bozkus** received the B.Sc. degree in electrical and electronics engineering from Bilkent University, Turkey, in 2020. He is currently working toward the Ph.D. degree with Ming Hsieh Department of Electrical and Computer Engineering, University of Southern California, Los Angeles, CA, USA. His research interests include signal processing, network optimization, and reinforcement learning. He was the recipient of a USC Annenberg Fellowship for graduate study and the Bilkent University Comprehensive Scholarship for his undergraduate studies.

**Urbashi Mitra** (Fellow, IEEE) received the B.S. and M.S. degrees from the University of California at Berkeley, Berkeley, CA, USA, and the Ph.D. degree from Princeton University, Princeton, NJ, USA. She is currently the Gordon S. Marshall Chair of engineering with the University of Southern California. Her research interests include wireless communications, biological communication, underwater acoustic communications, communication and sensor networks, and the detection, estimation, and the interface of communication, sensing, and control. She was a recipient of the 1996 National Science Foundation CAREER Award and the 1997 OSU College of Engineering MacQuigg Award for Teaching, the 2000 OSU College of Engineering Lumley Award for Research, the 2001 Okawa Foundation Award, the Texas Instruments Visiting Professorship (Rice University, Fall 2002), the 2009 DCOS Applications and Systems Best Paper Award, the Best Paper Award from the 2012 GLOBECOM Signal Processing Symposium for Communications, the 2012 U.S. National Academy of Engineering Lillian Gilbreth Lectureship, the 2014 to 2015 IEEE Communications Society Distinguished Lecturer, the Women in Communications Engineering Technical Achievement Award from the IEEE Communications Society in 2017, the 2016 U.K. Royal Academy of Engineering Distinguished Visiting Professorship, the 2016 USA Fulbright Scholar Award, the 2016 to 2017 U.K. Leverhulme Trust Visiting Professorship, and the 2021 USC Viterbi School of Engineering Senior Research Award. She was the inaugural Editor-in-Chief of IEEE TRANSACTIONS ON MOLECULAR, BIOLOGICAL AND MULTI-SCALE COMMUNICATIONS. She was an Associate Editor of IEEE TRANSACTIONS ON SIGNAL PROCESSING, from 2012 to 2015, IEEE TRANSACTIONS ON INFORMATION THEORY, from 2007 to 2011, IEEE JOURNAL OF OCEANIC ENGINEERING, from 2006 to 2011, and IEEE TRANSACTIONS ON COMMUNICATIONS, from 1996 to 2001.