A Related works

Given the large body of relevant literature, even surveying the work on Q-learning in a satisfactory manner is beyond the scope of this paper. Here we only mention the most relevant prior works, and compare them to ours in terms of the assumptions needed, the algorithmic approaches considered, and the performance guarantees provided. Table 1 provides key representative works from the literature and contrasts them with our result.

Q-learning has been studied extensively for finite-state MDPs. [43] and [22] are amongst the first to establish its asymptotic convergence. Both of them cast Q-learning as a stochastic approximation scheme — we utilize this abstraction as well. More recent work studies non-asymptotic performance of Q-learning; see, e.g., [41], [18], and [24]. Many variants of Q-learning have also been proposed and analyzed, including Double Q-learning [20], Speedy Q-learning [3], Phased Q-learning [23] and Delayed Q-learning [40].

A standard approach for continuous-state MDPs with known transition kernels, is to construct a reduced model by discretizing state space and show that the new finite MDP approximates the original one. For example, Chow and Tsitsiklis establish approximation guarantees for a multigrid algorithm when the state space is compact [10, [11]]. This result is recently extended to average-cost problems and to general Borel state and action spaces in [37]. To reduce the computational complexity, Rust proposes a randomized version of the multigrid algorithm and provides a bound on its approximation accuracy [36]. Our approach bears some similarities to this line of work: we also use state space discretization, and impose similar continuity assumptions on the MDP model. However, we do not require the transition kernel to be known, nor do we construct a reduced model; rather, we learn the action-value function of the original MDP directly by observing its sample path.

The closest work to this paper is by Szepesvari and Smart [42], wherein they consider a variant of Q-learning combined with local function approximation methods. The algorithm approximates the optimal Q-values at a given set of sample points and interpolates it for each query point. Follow-up work considers combining Q-learning with linear function approximation [28]. Despite algorithmic similarity, their results are distinct from ours: they establish *asymptotic* convergence of the algorithm, based on the assumption that the data-sampling policy is stochastic stationary. In contrast, we provide finite-sample bounds, and our results apply for arbitrary sample paths (including non-stationary policies). Consequently, our analytical techniques are also different from theirs.

Some other closely related work is by Ormoneit and coauthors on model-free reinforcement learning for continuous state with unknown transition kernels [33] [32]. Their approach, called KBRL, constructs a kernel-based approximation of the conditional expectation that appears in the Bellman operator. Value iteration can then be run using the approximate Bellman operator, and asymptotic consistency is established for the resulting fixed points. A subsequent work demonstrates applicability of KBRL to practical large-scale problems [4]. Unlike our approach, KBRL is an offline, batch algorithm in which data is sampled at once and remains the same throughout the iterations of the algorithm. Moreover, the aforementioned work does not provide convergence rate or finite-sample performance guarantee for KBRL. The idea of approximating the Bellman operator by an empirical estimate, has also been used in the context of discrete state-space problems [19]. The approximate operator is used to develop Empirical Dynamic Programming (EDP) algorithms including value and policy iterations, for which non-asymptotic error bounds are provided. EDP is again an offline batch algorithm; moreover, it requires multiple, independent transitions to be sampled for each state, and hence does not apply to our setting with a single sample path.

In terms of theoretical results, most relevant is the work in [30], who also obtain finite-sample performance guarantees for continuous space problems with unknown transition kernels. Extension to the setting with a single sample path is considered in [1]. The algorithms considered therein, including fitted value iteration and Bellman-residual minimization based fitted policy iteration, are different from ours. In particular, these algorithms perform updates in a batch fashion and require storage of all the data throughout the iterations.

There are other papers that provide finite-sample guarantees, such as [25, 12]; however, their settings (availability of i.i.d. data), algorithms (TD learning) and proof techniques are very different from ours. The work by Bhandari et al. [8] also provides a finite sample analysis of TD learning with linear function approximation, for both the i.i.d. data model and a single trajectory. We also note that the work on PAC-MDP methods [34] explores the impact of exploration policy on

learning performance. The focus of our work is estimation of Q-functions rather than the problem of exploration; nevertheless, we believe it is an interesting future direction to study combining our algorithm with smart exploration strategies.

B Proof of Lemma **1**

Proof. Let \mathcal{D} be the set of all functions $u : \mathcal{X} \times \mathcal{A} \to \mathbb{R}$ such that $||u||_{\infty} \leq V_{\max}$. Let \mathcal{L} be the set of all functions $u : \mathcal{X} \times \mathcal{A} \to \mathbb{R}$ such that $u(\cdot, a) \in \operatorname{Lip}(\mathcal{X}, M_r + \gamma V_{\max}M_p), \forall a \in \mathcal{A}$. Take any $u \in \mathcal{D}$, and fix an arbitrary $a \in \mathcal{A}$. For any $x \in \mathcal{X}$, we have

$$|(Fu)(x,a)| = \left| r(x,a) + \gamma \mathbb{E} \left[\max_{b \in \mathcal{A}} u(Z,b) | x, a \right] \right| \le R_{\max} + \gamma V_{\max} = V_{\max},$$

where the last equality follows from the definition of V_{max} . This means that $Fu \in \mathcal{D}$. Also, for any $x, y \in \mathcal{X}$, we have

$$\begin{aligned} |(Fu)(x,a) - (Fu)(y,a)| &= \left| r(x,a) - r(y,a) + \gamma \mathbb{E} \left[\max_{b \in \mathcal{A}} u(Z,b) | x, a \right] - \gamma \mathbb{E} \left[\max_{b \in \mathcal{A}} u(Z,b) | y, a \right] \\ &\leq |r(x,a) - r(y,a)| + \gamma \left| \int_{\mathcal{X}} \max_{b \in \mathcal{A}} u(z,b) \left(p(z|x,a) - p(z|y,a) \right) \lambda(dz) \right| \\ &\leq M_r \rho(x,y) + \gamma \int_{\mathcal{X}} \left| \max_{b \in \mathcal{A}} u(z,b) \left(p(z|x,a) - p(z|y,a) \right) \right| \lambda(dz) \\ &\leq M_r \rho(x,y) + \gamma \left\| u \right\|_{\infty} \cdot \int_{\mathcal{X}} |p(z|x,a) - p(z|y,a)| \lambda(dz) \\ &\leq [M_r + \gamma V_{\max} M_p] \rho(x,y). \end{aligned}$$

This means that $(Fu)(\cdot, a) \in \operatorname{Lip}(\mathcal{X}, M_r + \gamma V_{\max}M_p)$, so $Fu \in \mathcal{L}$. Putting together, we see that F maps \mathcal{D} to $\mathcal{D} \cap \mathcal{L}$, which in particular implies that F maps $\mathcal{D} \cap \mathcal{L}$ to itself. Since $\mathcal{D} \cap \mathcal{L}$ is closed and F is γ -contraction, both with respect to $\|\cdot\|_{\infty}$, the Banach fixed point theorem guarantees that F has a unique fixed point $Q^* \in \mathcal{D} \cap \mathcal{L}$. This completes the proof of the lemma. \Box

C Examples of Nearest Neighbor Regression Methods

Below we describe three representative nearest neighbor regression methods, each of which corresponds to a certain choice of the kernel function K in the averaging procedure (2).

k-nearest neighbor (k-NN) regression: For each x ∈ X, we find its k nearest neighbors in the subset X_h and average their Q-values, where k ∈ [n] is a pre-specified number. Formally, let c_(i)(x) denote the *i*-th closest data point to x amongst the set X_h. Thus, the distance of each state in X_h to x satisfies ρ(x, c₍₁₎(x)) ≤ ρ(x, c₍₂₎(x)) ≤ ··· ≤ ρ(x, c_(n)(x)). Then the k-NN estimate for the Q-value of (x, a) is given by (Γ_{NN}q)(x, a) = ¹/_k Σ^k_{i=1} q(c_(i)(x), a). This corresponds to using in ⁽²⁾/_k the following weighting function

$$K(x, c_i) = \frac{1}{k} \mathbb{1}\{\rho(x, c_i) \le \rho(x, c_{(k)}(x))\}.$$

Under the definition of \mathcal{X}_h in Section 3.1, the assumption 4 is satisfied if we use k = 1. For other values of k, the assumption holds with a potentially different value of h.

Fixed-radius near neighbor regression: We find all neighbors of x up to a threshold distance h > 0 and average their Q-values. The definition of X_h ensures that at least one point c_i ∈ X_h is within the threshold distance h, i.e., ∀x ∈ X, ∃c_i ∈ X_h such that ρ(x, c_i) ≤ h. We then can define the weighting function function according to

$$K(x, c_i) = \frac{\mathbb{1}\{\rho(x, c_i) \le h\}}{\sum_{i=1}^n \mathbb{1}\{\rho(x, c_i) \le h\}}$$

Kernel regression: Here the Q-values of the neighbors of x are averaged in a weighted fashion according to some kernel function [31, 48]. The kernel function φ : ℝ⁺ → [0, 1]

takes as input a distance (normalized by the bandwidth parameter h) and outputs a similarity score between 0 and 1. Then the weighting function $K(x, c_i)$ is given by

$$K(x,c_i) = \frac{\phi\left(\frac{\rho(x,c_i)}{h}\right)}{\sum_{j=1}^n \phi\left(\frac{\rho(x,c_j)}{h}\right)}.$$

For example, a (truncated) Gaussian kernel corresponds to $\phi(s) = \exp\left(-\frac{s^2}{2}\right) \mathbb{1}\{s \le 1\}$. Choosing $\phi(s) = \mathbb{1}\{s \le 1\}$ reduces to the fixed-radius NN regression described above.

D Bounds on Covering time

D.1 Proof of Proposition 1

Proof. Without loss of generality, we may assume that the balls $\{\mathcal{B}_i, i \in [n]\}$ are disjoint, since the covering time will only become smaller if they overlap with each other. Note that under ε -greedy policy, equation (6) implies that $\forall t \ge 0, \forall x \in \mathcal{X}, \forall a \in \mathcal{A},$

$$\Pr\left(x_{m+t} \in \cdot, a_{m+t} = a | x_t = x\right) \ge \frac{\varepsilon}{|\mathcal{A}|} \psi \nu(\cdot).$$
(10)

First assume that the above assumption holds with m = 1. Let $M \triangleq n |\mathcal{A}|$ be the total number of ball-action pairs. Let $(\mathcal{P}_1, \ldots, \mathcal{P}_M)$ be a fixed ordering of these M pairs. For each integer $t \ge 1$, let K_t be the number of ball-action pairs visited up to time t. Let $T \triangleq \inf \{t \ge 1 : K_t = M\}$ be the first time when all ball-action pairs are visited. For each $k \in \{1, 2, \ldots, M\}$, let $T_k \triangleq \{t \ge 1 : K_t = k\}$ be the first time when k pairs are visited, and let $D_k \triangleq T_k - T_{k-1}$ be the time to visit the k-th pair after k - 1 pairs have been visited. We use the convention that $T_0 = D_0 = 0$. By definition, we have $T = \sum_{k=1}^{M} D_k$.

When k - 1 pairs have been visited, the probability of visiting a *new* pair is at least

$$\min_{I \subseteq [M], |I| = M - k + 1} \Pr\left((x_{T_{k-1}+1}, a_{T_{k-1}+1}) \in \bigcup_{i \in I} \mathcal{P}_i | x_{T_{k-1}} \right) \\
= \min_{I \subseteq [M], |I| = M - k + 1} \sum_{i \in I} \Pr\left((x_{T_{k-1}+1}, a_{T_{k-1}+1}) \in \mathcal{P}_i | x_{T_{k-1}} \right) \\
\ge (M - k + 1) \min_{i \in [M]} \Pr_{\pi} \left((x_{T_{k-1}+1}, a_{T_{k-1}+1}) \in \mathcal{P}_i | x_{T_{k-1}} \right) \\
\ge (M - k + 1) \cdot \varphi \nu_{\min} \cdot \frac{\varepsilon}{|\mathcal{A}|},$$

where the last inequality follows from Eq. (10). Therefore, D_k is stochastically dominated by a geometric random variable with mean at most $\frac{|\mathcal{A}|}{(M-k+1)\varepsilon\varphi\nu_{\min}}$. It follows that

$$\mathbb{E}T = \sum_{k=1}^{M} \mathbb{E}D_k \le \sum_{k=1}^{M} \frac{|\mathcal{A}|}{(M-k+1)\varepsilon\varphi\nu_{\min}} = O\left(\frac{|\mathcal{A}|}{\varepsilon\varphi\nu_{\min}}\log M\right).$$

This prove the proposition for m = 1.

For general values of m, the proposition follows from a similar argument by considering the MDP only at times $t = m, 2m, 3m, \ldots$

D.2 Proof of Proposition 2

Proof. We shall use a line of argument similar to that in the proof of Proposition []. We assume that the balls $\{\mathcal{B}_i, i \in [n]\}$ are disjoint. Note that under ε -greedy policy π , for all $t \ge 0$, for all $x \in \mathcal{X}$, we have

$$\Pr_{\pi}\left(a_{t} = \hat{a}_{1}, \dots, a_{t+m-1} = \hat{a}_{m} | x_{t} = x\right) \ge \left(\frac{\varepsilon}{|\mathcal{A}|}\right)^{m}.$$
(11)

The equation (7) implies that

 $\Pr_{\pi} (x_{t+m} \in \cdot | x_t = x)$ $\geq \Pr(x_{t+m} \in \cdot | x_t = x, a_t = \hat{a}_1, \dots, a_{t+m-1} = \hat{a}_m) \times \Pr_{\pi} (a_t = \hat{a}_1, \dots, a_{t+m-1} = \hat{a}_m | x_t = x)$ $\geq \psi \nu(\cdot) \left(\frac{\varepsilon}{|\mathcal{A}|}\right)^m.$

Thus for each $a \in \mathcal{A}$,

$$\Pr_{\pi}\left(x_{t+m} \in \cdot, a_{t+m} = a | x_t = x\right) \ge \psi \nu(\cdot) \left(\frac{\varepsilon}{|\mathcal{A}|}\right)^{m+1}.$$
(12)

We first consider the case m = 1 and use the same notation as in the proof of Proposition 1. When k - 1 pairs have been visited, the probability of visiting a *new* pair is at least

$$\min_{I \subseteq [M], |I|=M-k+1} \Pr\left((x_{T_{k-1}+1}, a_{T_{k-1}+1}) \in \bigcup_{i \in I} \mathcal{P}_i | x_{T_{k-1}} \right)$$

$$\geq (M-k+1) \min_{i \in [M]} \Pr_{\pi} \left((x_{T_{k-1}+1}, a_{T_{k-1}+1}) \in \mathcal{P}_i | x_{T_{k-1}} \right)$$

$$\geq (M-k+1) \cdot \varphi \nu_{\min} \cdot \left(\frac{\varepsilon}{|\mathcal{A}|} \right)^2,$$

where the last inequality follows from Eq. (12). Therefore, D_k , the time to visit the k-th pair after k-1 pairs have been visited, is stochastically dominated by a geometric random variable with mean at most $\frac{(|\mathcal{A}|/\varepsilon)^2}{(M-k+1)\varphi\nu_{\min}}$. It follows that

$$\mathbb{E}T = \sum_{k=1}^{M} \mathbb{E}D_k \le \sum_{k=1}^{M} \frac{\left(\left|\mathcal{A}\right|/\varepsilon\right)^2}{(M-k+1)\varphi\nu_{\min}} = O\left(\frac{\left(\left|\mathcal{A}\right|/\varepsilon\right)^2}{\varphi\nu_{\min}}\log M\right).$$

This prove the proposition for m = 1.

For general values of m, the proposition follows from a similar argument by considering the MDP only at times $t = m, 2m, 3m, \ldots$

E Proof of the Main Result: Theorem

The proof of Theorem I consists of three key steps summarized as follows.

Step 1. Stochastic Approximation. Since the nearest-neighbor approximation of the Bellman operator induces a biased update for q^k at each step, the key step in our proof is to analyze a Stochastic Approximation (SA) algorithm with *biased* noisy updates. In particular, we establish its finite-sample convergence rate in Theorem 3 which does not follow from available convergence theory. This result itself may be of independent interest.

Step 2. Properties of NNQL. To apply the stochastic approximation result to NNQL, we need to characterize some key properties of NNQL, including (i) the stability of the algorithm (i.e., the sequence q^k stays bounded), as established in Lemma 3; (ii) the contraction property of the joint Bellman-NN operator, as established in Lemma 4; and (iii) the error bound induced by discretization of the state space, as established in Lemma 5.

Step 3. Apply SA to NNQL. We apply the stochastic approximation result to establish the finitesample convergence of NNQL. In particular, step 2 above ensures that NNQL satisfies the assumptions in Theorem 3. Applying this theorem, we prove that NNQL converges to a neighborhood of q_h^* , the fixed point of the Joint Bellman-NN operator G, after a sufficient number of iterations. The proof of Theorem 1 is completed by relating q_h^* to the true optimal Q-function Q^* , and by bounding the number of time steps in terms of the the number of iterations and the covering time.

E.1 Stochastic Approximation

Consider a generic iterative stochastic approximation algorithm, where the iterative update rule is has the following form: let θ^t denote the *state* at time t, then it is updated as

$$\theta^{t+1} = \theta^t + \alpha_t \left(F(\theta^t) - \theta^t + w^{t+1} \right), \tag{13}$$

where $\alpha_t \in [0, 1]$ is a step-size parameter, w^{t+1} is a noise term and F is the functional update of interest.

Theorem 3. Suppose that the mapping $F : \mathbb{R}^d \to \mathbb{R}^d$ has a unique fixed point θ^* with $\|\theta^*\|_{\infty} \leq V$, and is a γ -contraction with respect to the ℓ_{∞} norm in the sense that

$$\|F(\theta) - F(\theta')\|_{\infty} \le \gamma \|\theta - \theta'\|_{\infty}$$

for all $\theta, \theta' \in \mathbb{R}^d$, where $0 < \gamma < 1$. Let $\{\mathcal{F}^t\}$ be an increasing sequence of σ -fields so that α_t and w^t are \mathcal{F}^t -measurable random variables, and θ^t be updated as per (13). Let δ_1, δ_2, M, V be non-negative deterministic constants. Suppose that the following hold with probability 1:

- 1. The bias $\Delta^{t+1} = \mathbb{E}\left[w^{t+1} \mid \mathcal{F}^t\right]$ satisfies $\left\|\Delta^{t+1}\right\|_{\infty} \leq \delta_1 + \delta_2 \left\|\theta^t\right\|_{\infty}$, for all $t \geq 0$;
- 2. $\|w^{t+1} \Delta^{t+1}\|_{\infty} \leq M$, for all $t \geq 0$;
- 3. $\|\theta^t\|_{\infty} \leq V$, for all $t \geq 0$.

Further, we choose

$$\alpha_t = \frac{\beta}{\beta + t},\tag{14}$$

where $\beta = \frac{1}{1-\gamma}$. Then for each $0 < \varepsilon < \min\{2V\beta, 2M\beta^2\}$, after

$$T = \frac{48VM^2\beta^4}{\varepsilon^3} \log\left(\frac{32dM^2\beta^4}{\delta\varepsilon^2}\right) + \frac{6V(\beta-1)}{\varepsilon}$$

iterations of (13), with probability at least $1 - \delta$, we have

$$\left\|\theta^T - \theta^*\right\|_{\infty} \le \beta(\delta_1 + \delta_2 V) + \varepsilon.$$

Proof. We define two auxiliary sequences: for $i \in [d]$, let $u_i^0 = \theta_i^0$, $r_i^0 = 0$ and

$$u_{i}^{t+1} = (1 - \alpha_{t})u_{i}^{t} + \alpha_{t} \underbrace{(w_{i}^{t+1} - \Delta_{i}^{t+1})}_{\bar{w}_{i}^{t+1}},$$
$$r_{i}^{t+1} = \theta_{i}^{t+1} - u_{i}^{t+1}.$$

By construction, $\theta^t = u^t + r^t$ for all t. We first analyze the convergence rate of the (u^t) sequence. One has

$$u_i^{t+1} = (1 - \alpha_t)u_i^t + \alpha_t \bar{w}_i^{t+1} = (1 - \alpha_t)(1 - \alpha_{t-1})u_i^{t-1} + (1 - \alpha_t)\alpha_{t-1}\bar{w}_i^t + \alpha_t \bar{w}_i^{t+1} = \sum_{j=1}^{t+1} \eta^{t+1,j} \bar{w}_i^j,$$

where we define

$$\eta^{t+1,j} := \alpha_{j-1} \cdot \prod_{l=j}^t (1-\alpha_l).$$

Note that the centered noise $\bar{w}_i^{t+1} := w_i^{t+1} - \Delta_i^{t+1}$ satisfies

$$\mathbb{E}\left[\bar{w}_{i}^{t+1}|\mathcal{F}^{t}\right] = \mathbb{E}\left[w_{i}^{t+1}|\mathcal{F}^{t}\right] - \Delta_{i}^{t+1} = 0,$$

$$\mathbb{E}\left[\left|\bar{w}_{i}^{t+1}\right||\mathcal{F}^{t}\right] = \mathbb{E}\left[\left|w_{i}^{t+1} - \Delta_{i}^{t+1}\right||\mathcal{F}^{t}\right] \le M.$$
 (15)

Now

$$\mathbb{E}\left[\eta^{t+1,j}\bar{w}_{i}^{j} \mid \mathcal{F}^{j-1}\right] = \eta^{t+1,j}\mathbb{E}\left[\bar{w}_{i}^{j} \mid \mathcal{F}^{j-1}\right] = 0.$$
(16)

With the linear learning rate defined in Eq. (14), and the fact that $\beta = \frac{1}{1-\gamma} > 1, \forall j \in [1, t+1]$, we have

$$\eta^{t+1,j} = \frac{\beta}{j-1+\beta} \cdot \prod_{l=j}^{t} \frac{l}{l+\beta} < \frac{\beta}{j} \prod_{l=j}^{t} \frac{l}{l+1} = \frac{\beta}{t+1}.$$
(17)

Since the centered noise sequence $\{\bar{w}_i^1, \bar{w}_i^2, \dots, \bar{w}_i^{t+1}\}$ is uniformly bounded by M > 0, it follows that

$$\left|\eta^{t+1,j}\bar{w}_i^j\right| \le \frac{M\beta}{t+1}.\tag{18}$$

Define, for $1 \le s \le t+1$,

$$z_s^{t+1,i} := \sum_{j=1}^s \eta^{t+1,j} \bar{w}_i^j, \tag{19}$$

and $z_0^{t+1,i} = 0$. Then it follows that

$$\mathbb{E}\left[z_{s+1}^{t+1,i}|\mathcal{F}^s\right] = z_s^{t+1,i}.$$
(20)

And from (16)-(18), it follows that

$$|z_{s+1}^{t+1,i} - z_s^{t+1,i}| \le \frac{M\beta}{t+1}.$$
(21)

That is, $z_s^{t+1,i}$ is a Martingale with bounded differences. And $u_i^{t+1} = z_{t+1}^{t+1,i}$. This, using Azuma-Hoeffding's inequality, will provide us desired bound on $|u_i^{t+1}|$. To that end, let us recall the Azuma-Hoeffding's inequality.

Lemma 2 (Azuma-Hoeffding). Let X_j be Martingale with respect to filtration \mathcal{F}_j , i.e. $\mathbb{E}[X_{j+1}|\mathcal{F}_j] = X_j$ for $j \ge 1$ with $X_0 = 0$. Further, let $|X_j - X_{j-1}| \le c_j$ with probability 1 for all $j \ge 1$. Then for all $\varepsilon \ge 0$,

$$\Pr\left[|X_n| \ge \varepsilon\right] \le 2 \exp\left(-\frac{\varepsilon^2}{2\sum_{j=1}^n c_j^2}\right)$$

Applying the lemma to $z_j^{t+1,i}$ for $j \ge 0$ with $z_0^{t+1,i} = 0$, (21) and the fact that $u_i^{t+1} = z_{t+1}^{t+1,i}$, we obtain that

$$\Pr\left(\left|u_{i}^{t+1}\right| > \varepsilon\right) \le 2\exp\left(-\frac{(t+1)\varepsilon^{2}}{2M^{2}\beta^{2}}\right).$$
(22)

Therefore, by union bound we obtain

$$\begin{split} \Pr\left(\exists t \geq T_1 \text{ such that } \left|u_i^t\right| > \varepsilon\right) &\leq \sum_{t=T_1}^{\infty} \Pr\left(\left|u_i^t\right| > \varepsilon\right) \\ &\leq 2\sum_{t=T_1}^{\infty} \exp\left(-\frac{t\varepsilon^2}{2M^2\beta^2}\right) \\ &= \frac{2\exp\left(-\frac{T_1\varepsilon^2}{2M^2\beta^2}\right)}{1 - \exp\left(-\frac{\varepsilon^2}{2M^2\beta^2}\right)} \\ &\leq \frac{8M^2\beta^2}{\varepsilon^2} \exp\left(-\frac{T_1\varepsilon^2}{2M^2\beta^2}\right), \end{split}$$

where the last step follows from the fact that $e^{-x} \le 1 - \frac{x}{2}$ for $0 \le x \le \frac{1}{2}$, and $\varepsilon \le M\beta$. By a union bound over all $i \in [d]$, we deduce that

$$\Pr\left(\exists t \ge T_1 \text{ such that } \left\|u^t\right\|_{\infty} > \varepsilon\right) \le \sum_{i \in [d]} \Pr\left(\exists t \ge T_1 \text{ such that } \left|u^t_i\right| > \varepsilon\right)$$
$$\le \frac{8dM^2\beta^2}{\varepsilon^2} \exp\left(-\frac{T_1\varepsilon^2}{2M^2\beta^2}\right). \tag{23}$$

Next we focus on the residual sequence (r^t) . Assume that $\forall t \geq T_1$, $||u^t||_{\infty} \leq \varepsilon_1$, where $0 < \varepsilon_1 < \min\{V, M\beta\}$. For each $i \in [d]$ and $t \geq T_1$, we get

$$\begin{aligned} \left| r_i^{t+1} - \theta_i^* \right| & \text{by definition} \\ &= \left| \theta_i^{t+1} - u_i^{t+1} - \theta_i^* \right| & \text{by definition} \\ &= \left| \theta_i^t + \alpha_t \left(F_i(u^t + r^t) - u_i^t - r_i^t + w_i^{t+1} \right) - u_i^t - \alpha_t \left(-u_i^t + w_i^{t+1} - \Delta_i^{t+1} \right) - \theta_i^* \right| & \text{by definition} \\ &= \left| r_i^t + \alpha_t \left(F_i(u^t + r^t) - r_i^t \right) - \theta_i^* + \alpha_t \Delta_i^{t+1} \right| & \text{rearranging} \\ &= \left| (1 - \alpha_t) (r_i^t - \theta_i^*) + \alpha_t \left(F_i(u^t + r^t) - \theta_i^* \right) + \alpha_t \Delta_i^{t+1} \right| & \text{rearranging} \\ &\leq (1 - \alpha_t) \left| r_i^t - \theta_i^* \right| + \alpha_t \gamma \left\| u^t + r^t - \theta^* \right\|_{\infty} + \alpha_t \left\| \Delta^{t+1} \right\|_{\infty} & F \text{ is } \gamma \text{-contraction} \\ &\leq (1 - \alpha_t) \left| r_i^t - \theta_i^* \right| + \alpha_t \gamma \left\| r^t - \theta^* \right\|_{\infty} + \alpha_t \gamma \varepsilon_1 + \alpha_t \left(\delta_1 + \delta_2 \left\| \theta^t \right\|_{\infty} \right) & \left\| u^t \right\|_{\infty} \leq \varepsilon_1, \forall t \geq T_1 \\ &\leq (1 - \alpha_t) \left| r_i^t - \theta_i^* \right| + \alpha_t \gamma \left\| r^t - \theta^* \right\|_{\infty} + \alpha_t (\gamma \varepsilon_1 + \delta_1 + \delta_2 V) & \left\| \theta^t \right\|_{\infty} \leq V \end{aligned}$$

Taking the maximum over $i \in [d]$ on both sides, we obtain

$$\begin{aligned} \left\| r^{t+1} - \theta^* \right\|_{\infty} &\leq (1 - \alpha_t) \left\| r^t - \theta^* \right\|_{\infty} + \alpha_t \gamma \left\| r^t - \theta^* \right\|_{\infty} + \alpha_t \left(\gamma \varepsilon_1 + \delta_1 + \delta_2 V \right) \\ &= \left(1 - \underbrace{(1 - \gamma)}_{\frac{1}{\beta}} \alpha_t \right) \underbrace{\left\| r^t - \theta^* \right\|_{\infty}}_{D_t} + \alpha_t \underbrace{(\gamma \varepsilon_1 + \delta_1 + \delta_2 V)}_{H}, \qquad \forall t \geq T_1. \end{aligned}$$

For any $\varepsilon_2 > 0$, we will show that after at most

$$T_2 \triangleq \frac{3V(T_1 + \beta - 1)}{\varepsilon_2}$$

iterations, we have

$$\left\| r^{T_2} - \theta^* \right\|_{\infty} \le H\beta + \varepsilon_2.$$

If for some $T\in [T_1,\infty)$ there holds $D_T\leq H\beta+\varepsilon_2,$ then we have

$$D_{T+1} \le \left(1 - \frac{\alpha_t}{\beta}\right) (H\beta + \varepsilon_2) + \alpha_t H \qquad \alpha_t \le \beta$$
$$\le H\beta + \varepsilon_2$$

Indeed by induction, we have

$$D_t \le H\beta + \varepsilon_2, \quad \forall t \ge T.$$

Let $\widehat{T} \triangleq \sup \{t \ge T_1 : D_t > H\beta + \varepsilon_2\}$ be the last time that D_t exceeds $H\beta + \varepsilon_2$. For each $T_1 \le t \le \widehat{T}$, the above argument implies that we must have $D_t > H\beta + \varepsilon_2$. We can rewrite the iteration for $D_{\widehat{T}}$ as follows:

$$D_{\widehat{T}} - H\beta \leq \left(D_{\widehat{T}-1} - H\beta\right) \left(1 - \frac{\alpha_{\widehat{T}-1}}{\beta}\right)$$

$$\leq \left(D_{T_1} - H\beta\right) \prod_{j=T_1}^{\widehat{T}-1} \left(1 - \frac{\alpha_j}{\beta}\right) \qquad \text{Iteration, } D_t - H\beta > \varepsilon_2 > 0$$

$$= \left(D_{T_1} - H\beta\right) \frac{T_1 + \beta - 1}{\widehat{T} + \beta - 1} \qquad \alpha_j = \frac{\beta}{j + \beta}.$$

But we have the bound

$$D_{T_1} = \|r^{T_1} - \theta^*\|_{\infty}$$

= $\|\theta^{T_1} - u^{T_1} - \theta^*\|_{\infty}$
 $\leq \|\theta^{T_1}\|_{\infty} + \|u^{T_1}\|_{\infty} + \|\theta^*\|_{\infty}$
 $\leq 3V,$

where the last step holds because $\|\theta^{T_1}\|_{\infty} \leq V$, $\|\theta^*\| \leq V$ and $\|u^{T_1}\| \leq \varepsilon_1 \leq V$ by assumption. It follows that

$$D_{\widehat{T}} - H\beta \le \frac{3V(T_1 + \beta - 1)}{\widehat{T}}.$$

Using the fact that $\varepsilon_2 \leq D_{\widehat{T}} - H\beta$, we get that

$$\widehat{T} \leq T_2 = \frac{3V(T_1 + \beta - 1)}{\varepsilon_2}$$

Therefore, for each $\varepsilon_2 > 0$, conditioned on the event

$$\left\{ \forall t \ge T_1, \left| u_i^t \right| \le \varepsilon_1 \right\}$$

after at most T_2 iterations, we have

$$\left\| r^{T_2} - \theta^* \right\|_{\infty} \le H\beta + \varepsilon_2.$$

It then follows from the relationship $\theta^t = u^t + r^t$ that

$$\left\|\theta^{T_2} - \theta^*\right\|_{\infty} \le \left\|r^{T_2} - \theta^*\right\|_{\infty} + \left\|u^{T_2}\right\|_{\infty} \le H\beta + \varepsilon_1 + \varepsilon_2 = \beta(\delta_1 + \delta_2 V) + \beta\varepsilon_1 + \varepsilon_2.$$

By (23), taking

$$\delta = \frac{8dM^2\beta^2}{\varepsilon_1^2} \exp\left(-\frac{T_1\varepsilon_1^2}{2M^2\beta^2}\right),$$

i.e.,

$$T_1 = \frac{2M^2\beta^2}{\varepsilon_1^2} \log\left(\frac{8dM^2\beta^2}{\delta\varepsilon_1^2}\right),$$

we are guaranteed that

$$\Pr\left(\forall t \ge T_1, \left\|u^t\right\|_{\infty} \le \varepsilon_1\right) \ge 1 - \delta.$$

By setting $\varepsilon_1 = \frac{\varepsilon}{2\beta} \leq \min\{V, M\beta\}$, and $T_2 = \frac{3V(T_1 + \beta - 1)}{\varepsilon_2}$, i.e.,

$$T_2 = \frac{48VM^2\beta^4}{\varepsilon^3} \log\left(\frac{32dM^2\beta^4}{\delta\varepsilon^2}\right) + \frac{6V(\beta-1)}{\varepsilon}$$

we obtain that the desire result.

E.2 Properties of NNQL

We first introduce some notations. Let \mathcal{Y}^k be the set of all samples drawn at iteration k of the NNQL algorithms and \mathcal{F}^k be the filtration generated by the sequence $\mathcal{Y}^0, \mathcal{Y}^1, \ldots, \mathcal{Y}^{k-1}$. Thus $\{\mathcal{F}^k\}$ is an increasing sequence of σ -fields. We denote by $\mathcal{Y}_k(c_i, a) = \{Y_t \in \mathcal{Y}_k | Y_t \in \mathcal{B}_i, a_t = a\}$ the set of observations Y_t that fall into the neighborhood \mathcal{B}_i of c_i and with action $a_t = a$ at iteration k. Thus the biased estimator G^k be for the joint Bellman-NN operator at the end of iteration k can be written as

$$(G^k q)(c_i, a) = \frac{1}{|\mathcal{Y}_k(c_i, a)|} \sum_{Y_t \in \mathcal{Y}_k(c_i, a)} \left[R_t + \gamma \max_{b \in \mathcal{A}} (\Gamma_{NN} q^k)(Y_{t+1}, b) \right].$$

The updater rule of NNQL (9) can be written as

$$q^{k+1}(c_i, a) = q^k(c_i, a) + \alpha_k \left[(Gq^k)(c_i, a) - q^k(c_i, a) + w^{k+1}(c_i, a) \right],$$

where

$$w^{k+1}(c_i, a) = (G^k q^k)(c_i, a) - (Gq^k)(c_i, a)$$

=
$$\frac{1}{|\mathcal{Y}_k(c_i, a)|} \sum_{Y_t \in \mathcal{Y}_k(c_i, a)} \left[R_t + \gamma \max_{b \in \mathcal{A}} (\Gamma_{NN} q^k)(Y_{t+1}, b) \right]$$

-
$$r(c_i, a) - \gamma \mathbb{E} \left[\max_{b \in \mathcal{A}} (\Gamma_{NN} q^k)(x', b) \mid c_i, a, \mathcal{F}^k \right].$$

c	_	_	-

E.2.1 Stability of NNQL

We first show the stability of NNQL, which is summarized in the following Lemma.

Lemma 3 (Stability of NNQL). Assume that the immediate reward is uniformly bounded by R_{max} and define $\beta = \frac{1}{1-\gamma}$ and $V_{\text{max}} = \beta R_{\text{max}}$. If the initial action-value function q^0 is uniformly bounded by V_{max} , then we have

$$\left\|q^{k}\right\|_{\infty} \leq V_{\max}, \quad and \quad \left|w^{k+1}(c_{i},a) - \mathbb{E}\left[w^{k+1}(c_{i},a) \mid \mathcal{F}^{k}\right]\right| \leq 2V_{\max}, \qquad \forall k \geq 0.$$

Proof. We first prove that $||q^k||_{\infty} \leq V_{\max}$ by induction. For k = 0, it holds by the assumption. Now assume that for any $0 \leq \tau \leq k$, $||q^{\tau}||_{\infty} \leq V_{\max}$. Thus

$$\begin{split} &|q^{k+1}(c_{i},a)| \\ &= \left| q^{k}(c_{i},a) + \alpha_{k} \left[(G^{k}q^{k})(c_{i},a) - q^{k}(c_{i},a) \right] \right| \\ &= \left| (1 - \alpha_{k}) q^{k}(c_{i},a) + \frac{\alpha_{k}}{|\mathcal{Y}_{k}(c_{i},a)|} \sum_{Y_{t} \in \mathcal{Y}_{k}(c_{i},a)} \left[R_{t} + \gamma \max_{b \in \mathcal{A}} (\Gamma_{\mathrm{NN}}q^{k})(Y_{t+1},b) \right] \right| \\ &\leq \left| (1 - \alpha_{k}) q^{k}(c_{i},a) \right| + \frac{\alpha_{k}}{|\mathcal{Y}_{k}(c_{i},a)|} \sum_{Y_{t} \in \mathcal{Y}_{k}(c_{i},a)} \left(|R_{t}| + \gamma \left| \max_{b \in \mathcal{A}} (\Gamma_{\mathrm{NN}}q^{k})(Y_{t+1},b) \right| \right) \right) \\ &= (1 - \alpha_{k}) \left| q^{k}(c_{i},a) \right| + \frac{\alpha_{k}}{|\mathcal{Y}_{k}(c_{i},a)|} \sum_{Y_{t} \in \mathcal{Y}_{k}(c_{i},a)} \left(|R_{t}| + \gamma \max_{b \in \mathcal{A}} \left| \sum_{j=1}^{n} K(Y_{t+1},c_{j})q^{k}(c_{j},b) \right| \right) \\ &\leq (1 - \alpha_{k}) V_{\mathrm{max}} + \alpha_{k} \left(R_{\mathrm{max}} + \gamma \max_{b \in \mathcal{A}} \left| \sum_{j=1}^{n} K(Y_{t+1},c_{j}) \right| V_{\mathrm{max}} \right) \\ &= V_{\mathrm{max}}, \end{split}$$

where the last equality follows from the fact that $\sum_{j=1}^{n} K(Y_{t+1}, c_j) = 1$. Therefore, for all $k \ge 0$, $||q^k||_{\infty} \le V_{\max}$. The bound on w^{k+1} follows from

$$\begin{split} & \left| w^{k+1}(c_i, a) - \mathbb{E} \left[w^{k+1}(c_i, a) \mid \mathcal{F}^k \right] \right| \\ &= \left| (G^k q^k)(c_i, a) - (Gq^k)(c_i, a) - \mathbb{E} \left[(G^k q^k)(c_i, a) - (Gq^k)(c_i, a) \mid \mathcal{F}^k \right] \right| \\ &= \left| (G^k q^k)(c_i, a) - \mathbb{E} \left[(G^k q^k)(c_i, a) \mid \mathcal{F}^k \right] \right| \\ &= \left| \frac{1}{\left| \mathcal{Y}_k(c_i, a) \right|} \sum_{Y_t \in \mathcal{Y}_k(c_i, a)} \left[R_t + \gamma \max_{b \in \mathcal{A}} (\Gamma_{NN} q^k)(Y_{t+1}, b) \right] \right| \\ &- \mathbb{E} \left[\frac{1}{\left| \mathcal{Y}_k(c_i, a) \right|} \sum_{Y_t \in \mathcal{Y}_k(c_i, a)} \left[R_t + \gamma \max_{b \in \mathcal{A}} (\Gamma_{NN} q^k)(Y_{t+1}, b) \right] \mid \mathcal{F}^k \right] \right| \\ &\leq 2R_{\max} + \frac{1}{\left| \mathcal{Y}_k(c_i, a) \right|} \sum_{Y_t \in \mathcal{Y}_k(c_i, a)} \gamma \left| \max_{b \in \mathcal{A}} \sum_{j=1}^n K(Y_{t+1}, c_j) q^k(c_j, b) \right| \\ &+ \gamma \left| \mathbb{E} \left[\frac{1}{\left| \mathcal{Y}_k(c_i, a) \right|} \sum_{Y_t \in \mathcal{Y}_k(c_i, a)} \max_{b \in \mathcal{A}} \sum_{j=1}^n K(Y_{t+1}, c_j) q^k(c_j, b) \mid c_i, a, \mathcal{F}^k \right] \right| \\ &\leq 2R_{\max} + 2\gamma V_{\max} \\ &= 2V_{\max}. \end{split}$$

E.2.2 A contraction operator

The following Lemma states that the joint Bellman-NN operator G is a contraction with modulus γ , and has a unique fixed point that is bounded.

Lemma 4 (Contraction of the Joint-Bellman-NN operator). For each fixed h > 0, the operator G defined in Eq. (5) is a contraction with modulus γ with the supremum norm. There exists a unique function q_h^* such that

$$(Gq_h^*)(c_i, a) = q_h^*(c_i, a), \qquad \forall (c_i, a) \in \mathcal{Z}_h,$$

where $||q_h^*||_{\infty} \leq V_{\max}$.

Proof. Let \mathcal{D} be the set of all functions $q : \mathcal{X}_h \times \mathcal{A} \to \mathbb{R}$ such that $\|q\|_{\infty} \leq V_{\max}$. We first show that the operator G maps \mathcal{D} into itself. Take any $q \in \mathcal{D}$, and fix an arbitrary $a \in \mathcal{A}$. For any $c_i \in \mathcal{X}_h$, we have

$$\begin{aligned} |(Gq)(c_i, a)| &= \left| r(c_i, a) + \gamma \mathbb{E} \left[\max_{b \in \mathcal{A}} (\Gamma_{NN}q)(x', b) | c_i, a \right] \right| \\ &\leq |r(c_i, a)| + \gamma \left| \int_{\mathcal{X}} \left[\max_{b \in \mathcal{A}} \left(\sum_{j=1}^{N_h} K(y, c_j) q(c_j, b) \right) \right] p(y|c_i a) \lambda(dy) \right| \\ &\leq |r(c_i, a)| + \gamma \int_{\mathcal{X}} \left[\max_{b \in \mathcal{A}} \sum_{j=1}^{N_h} K(y, c_j) |q(c_j, b)| \right] p(y|c_i a) \lambda(dy) \\ &\leq R_{\max} + \gamma V_{\max} \\ &= V_{\max}, \end{aligned}$$

where the last step follows from the definition of V_{max} . This means that $Gq \in \mathcal{D}$, so G maps \mathcal{D} to itself.

Now, by the definition of G in Eq. (5), $\forall q, q' \in \mathcal{D}$, we have

$$\begin{split} \|Gq - Gq'\|_{\infty} &= \max_{i \in [n], a \in \mathcal{A}} |(Gq)(c_{i}, a) - (Gq')(c_{i}, a)| \\ &\leq \gamma \max_{i \in [n], a \in \mathcal{A}} \left| \mathbb{E} \left[\max_{b \in \mathcal{A}} \left(\sum_{j=1}^{n} K(x', c_{j}) \left(q(c_{j}, b) - q'(c_{j}, b) \right) \right) \mid c_{i}, a \right] \right| \\ &\leq \gamma \max_{i \in [n], a \in \mathcal{A}} \mathbb{E} \left[\max_{b \in \mathcal{A}} \left(\sum_{j=1}^{n} K(x', c_{j}) \left| q(c_{j}, b) - q'(c_{j}, b) \right| \right) \mid c_{i}, a \right] \\ &\leq \gamma \max_{i \in [n], a \in \mathcal{A}} \mathbb{E} \left[\max_{b \in \mathcal{A}} \left(\sum_{j=1}^{n} K(x', c_{j}) \left\| q - q' \right\|_{\infty} \right) \mid c_{i}, a \right] \\ &\leq \gamma \|q - q'\|_{\infty} \end{split}$$

Therefore G is indeed a contraction on \mathcal{D} with respect to the supremum norm. The Banach fixed point theorem guarantees that G has a unique fixed point $q_h^* \in \mathcal{D}$. This completes the proof. \Box

E.2.3 Discretization error

For each $q \in C(\mathcal{Z}_h)$, we can obtain an extension to the original continuous state space via the Nearest Neighbor operator. That is, define

$$Q(x,a) = (\Gamma_{NN}q)(x,a), \forall (x,a) \in \mathcal{Z}.$$

The following lemma characterizes the distance between the optimal action-value function Q^* and the extension of the fixed-point of the joint NN-Bellman operator G to the space \mathcal{Z} .

Lemma 5 (Discretization error). Define

$$Q_h^* = \Gamma_{NN} q_h^*$$

Let Q^* be the optimal action-value function for the original MDP. Then we have

$$\|Q_h^* - Q^*\| \le \beta Ch,$$

where $C = M_r + \gamma V_{\max} M_p$ and $\beta = \frac{1}{1-\gamma}$.

Proof. Consider an operator H on $C(\mathcal{Z})$ defined as follows:

$$(HQ)(x,a) = (\Gamma_{NN}(FQ))(x,a)$$

= $\sum_{i=1}^{n} K(x,c_i) \left\{ r(c_i,a) + \gamma \mathbb{E} \left[\max_{b \in \mathcal{A}} Q(x',b) \mid c_i,a \right] \right\}$ (24)

We can show that H is a contraction operator with modulus γ .

$$\begin{aligned} -HQ_2 \|_{\infty} &= \max_{a \in \mathcal{A}} \sup_{x \in \mathcal{X}} |(HQ_1)(x, a) - (HQ_2)(x, a)| \\ &= \gamma \max_{a \in \mathcal{A}} \sup_{x \in \mathcal{X}} \left| \mathbb{E} \left[\max_{b \in \mathcal{A}} \left(\sum_{i=1}^n K(x, c_i) \left(Q_1(x', b) - Q_2(x', b) \right) \right) \mid c_i, a \right] \right| \\ &\leq \gamma \max_{a \in \mathcal{A}} \sup_{x \in \mathcal{X}} \mathbb{E} \left[\max_{b \in \mathcal{A}} \left(\sum_{i=1}^n K(x, c_i) \left| Q_1(x', b) - Q_2(x', b) \right| \right) \mid c_i, a \right] \\ &\leq \gamma \max_{a \in \mathcal{A}} \sup_{x \in \mathcal{X}} \mathbb{E} \left[\max_{b \in \mathcal{A}} \left(\sum_{i=1}^n K(x, c_i) \left\| Q_1 - Q_2 \right\|_{\infty} \right) \mid c_i, a \right] \\ &= \gamma \left\| Q_1 - Q_2 \right\|_{\infty} \end{aligned}$$

We can conclude that H is a contraction operator mapping $C(\mathcal{Z})$ to $C(\mathcal{Z})$. Thus H has a unique fixed point $\tilde{Q} \in C(\mathcal{Z})$. Note that

$$H(\Gamma_{\rm NN}q) = \Gamma_{\rm NN}(F(\Gamma_{\rm NN}q)) = \Gamma_{\rm NN}(Gq),$$

we thus have

 $||HQ_1||$

$$HQ_h^* = H\left(\Gamma_{\rm NN}q_h^*\right) = \Gamma_{\rm NN}(Gq_h^*) = \Gamma_{\rm NN}(q_h^*) = Q_h^*$$

That is, the fixed point of H is exactly the extension of the fixed point of G to \mathcal{Z} . Therefore, we have

$$\begin{split} \|Q_{h}^{*} - Q^{*}\|_{\infty} &= \|HQ_{h}^{*} - HQ^{*} + HQ^{*} - Q^{*}\|_{\infty} \\ &\leq \|HQ_{h}^{*} - HQ^{*}\|_{\infty} + \|HQ^{*} - Q^{*}\|_{\infty} \\ &\leq \gamma \|Q_{h}^{*} - Q^{*}\|_{\infty} + \|HQ^{*} - Q^{*}\|_{\infty} \,. \end{split}$$

It follows that

$$\begin{aligned} \|Q_{h}^{*} - Q^{*}\|_{\infty} &\leq \frac{1}{1 - \gamma} \|HQ^{*} - Q^{*}\|_{\infty} \\ &= \frac{1}{1 - \gamma} \|\Gamma_{NN}(FQ^{*}) - Q^{*}\|_{\infty} \\ &= \frac{1}{1 - \gamma} \|\Gamma_{NN}(Q^{*}) - Q^{*}\|_{\infty} \\ &= \frac{1}{1 - \gamma} \sup_{x \in \mathcal{X}} \max_{a \in \mathcal{A}} \left| \sum_{i=1}^{n} K(x, c_{i})Q^{*}(c_{i}, a) - Q^{*}(x, a) \right| \end{aligned}$$

Recall that $Q^*(\cdot, a)$ is Lipschitz with parameter $C = M_r + \gamma V_{\max} M_p$ (see Lemma 1), i.e., for each $a \in \mathcal{A}$,

$$|Q^*(x,a) - Q^*(y,a)| \le C\rho(x,y).$$

From the state space discretization step, we know that the finite grid $\{c_i\}_{i=1}^{N_h}$ is an *h*-net in \mathcal{X} . Therefore, for each $x \in \mathcal{X}$, there exists $c_i \in \mathcal{X}_h$ such that

$$\rho(x, c_i) \le h.$$

Thus $\sum_{i=1}^{n} K(x, c_i) = 1$. Recall our assumption that the weighting function satisfies K(x, y) = 0 if $\rho(x, y) \ge h$. For each $a \in A$, then we have

$$\sup_{x \in \mathcal{X}} \left| \sum_{i=1}^{n} K(x, c_{i}) Q^{*}(c_{i}, a) - Q^{*}(x, a) \right| = \sup_{x \in \mathcal{X}} \left| \sum_{c_{i} \in \mathcal{B}_{x,h}} K(x, c_{i}) Q^{*}(c_{i}, a) - Q^{*}(x, a) \right|$$
$$\leq \sup_{x \in \mathcal{X}} \sum_{c_{i} \in \mathcal{B}_{x,h}} K(x, c_{i}) \left| Q^{*}(c_{i}, a) - Q^{*}(x, a) \right|$$
$$\leq Ch$$

This completes the proof.

E.3 Applying the Stochastic Approximation Theorem to NNQL

We first apply Theorem 3 to establish that NNQL converges to a neighborhood of q_h^* , the fixed point of the Joint Bellman-NN operator G, after a sufficiently large number of iterations. This is summarized in the following theorem.

Theorem 4. Let Assumptions [] and [2] hold. Then for each $0 < \varepsilon < 2V_{\max}\beta$, after

$$k = \frac{192V_{\max}^3\beta^4}{\varepsilon^3}\log\left(\frac{128dV_{\max}^2\beta^4}{\delta\varepsilon^2}\right) + \frac{4V_{\max}(\beta-1)}{\varepsilon}$$

iterations of Nearest-Neighbor Q-learning, with probability at least $1 - \delta$ *, we have*

$$\left\|q^k - q_h^*\right\|_{\infty} \le \beta(\delta_1 + \delta_2 V_{\max}) + \varepsilon.$$

Proof. We will show that NNQL satisfies the assumptions of Theorem 3. It follows from Lemma 4 that the operator G is a γ -contraction with a unique fixed point $||q_h^*||_{\infty} \leq V_{\max}$. For each $Y_t \in \mathcal{Y}_k(c_i, a)$, we have $\rho(Y_t, c_i) \leq h$, $a_t = a$. Thus

$$\begin{split} & \left| \mathbb{E} \left[w^{k+1}(c_i, a) | \mathcal{F}^k \right] \right| \\ = & \left| \mathbb{E} \left[\frac{1}{|\mathcal{Y}_k(c_i, a)|} \sum_{Y_t \in \mathcal{Y}_k(c_i, a)} \left[R_t + \gamma \max_{b \in \mathcal{A}} (\Gamma_{NN} q^k) (Y_{t+1}, b) \right] | \mathcal{F}^k \right] \\ & - r(c_i, a) - \gamma \mathbb{E} \left[\max_{b \in \mathcal{A}} (\Gamma_{NN} q^k) (x', b) | c_i, a, \mathcal{F}^k \right] \right| \\ \leq & \left| \mathbb{E} \left[\frac{1}{|\mathcal{Y}_k(c_i, a)|} \sum_{Y_t \in \mathcal{Y}_k(c_i, a)} R_t - r(c_i, a) | \mathcal{F}^k \right] \right| \\ & + \gamma \left| \mathbb{E} \left[\frac{1}{|\mathcal{Y}_k(c_i, a)|} \sum_{Y_t \in \mathcal{Y}_k(c_i, a)} \max_{b \in \mathcal{A}} (\Gamma_{NN} q^k) (Y_{t+1}, b) | \mathcal{F}^k \right] - \mathbb{E} \left[\max_{b \in \mathcal{A}} (\Gamma_{NN} q^k) (x', b) | c_i, a, \mathcal{F}^k \right] \right|. \end{split}$$

We can bound the first term on the RHS by using Lipschitz continuity of the reward function:

$$\begin{split} & \left| \mathbb{E} \Big[\frac{1}{|\mathcal{Y}_{k}(c_{i},a)|} \sum_{Y_{t} \in \mathcal{Y}_{k}(c_{i},a)} R_{t} - r(c_{i},a) \mid \mathcal{F}^{k} \Big] \right| \\ &= \left| \mathbb{E} \Big[\mathbb{E} \Big[\frac{1}{|\mathcal{Y}_{k}(c_{i},a)|} \sum_{Y_{t} \in \mathcal{Y}_{k}(c_{i},a)} \left(R_{t} - r(c_{i},a) \right) \mid \mathcal{Y}_{k}, \mathcal{F}^{k} \Big] \mid \mathcal{F}^{k} \Big] \right| \\ &= \left| \mathbb{E} \Big[\frac{1}{|\mathcal{Y}_{k}(c_{i},a)|} \sum_{Y_{t} \in \mathcal{Y}_{k}(c_{i},a)} \left(r(Y_{t},a) - r(c_{i},a) \right) \mid \mathcal{F}^{k} \Big] \right| \\ &\leq \mathbb{E} \Big[\frac{1}{|\mathcal{Y}_{k}(c_{i},a)|} \sum_{Y_{t} \in \mathcal{Y}_{k}(c_{i},a)} \left| r(Y_{t},a) - r(c_{i},a) \right| \mid \mathcal{F}^{k} \Big] \\ &\leq \mathbb{E} \Big[\frac{1}{|\mathcal{Y}_{k}(c_{i},a)|} \sum_{Y_{t} \in \mathcal{Y}_{k}(c_{i},a)} M_{r} \rho(Y_{t},c_{i}) \mid \mathcal{F}^{k} \Big] \\ &\leq M_{r}h \end{split}$$
 Lipschitz continuity of $r(\cdot,a)$

The second term on the right hand side can be bounded as follows:

$$\begin{split} & \left| \mathbb{E} \left[\frac{1}{|\mathcal{Y}_{k}(c_{i},a)|} \sum_{Y_{t} \in \mathcal{Y}_{k}(c_{i},a)} \max_{b \in \mathcal{A}}(\Gamma_{\mathrm{NN}}q^{k})(Y_{t+1},b) \mid \mathcal{F}^{k} \right] - \mathbb{E} \left[\max_{b \in \mathcal{A}}(\Gamma_{\mathrm{NN}}q^{k})(x',b) \mid c_{i},a,\mathcal{F}^{k} \right] \right| \\ \leq \mathbb{E} \left[\left| \frac{1}{|\mathcal{Y}_{k}(c_{i},a)|} \sum_{Y_{t} \in \mathcal{Y}_{k}(c_{i},a)} \int_{\mathcal{X}} \left[\max_{b \in \mathcal{A}}(\Gamma_{\mathrm{NN}}q^{k})(y,b) \right] p(y \mid Y_{t},a)\lambda(dy) \\ & - \int_{\mathcal{X}} \left[\max_{b \in \mathcal{A}}(\Gamma_{\mathrm{NN}}q^{k})(y,b) \right] p(y \mid c_{i},a)\lambda(dy) \mid |\mathcal{F}^{k} \right] \\ = \mathbb{E} \left[\left| \frac{1}{|\mathcal{Y}_{k}(c_{i},a)|} \sum_{Y_{t} \in \mathcal{Y}_{k}(c_{i},a)} \int_{\mathcal{X}} \left[\max_{b \in \mathcal{A}}(\Gamma_{\mathrm{NN}}q^{k})(y,b) \right] (p(y \mid Y_{t},a) - p(y \mid c_{i},a))\lambda(dy) \mid |\mathcal{F}^{k} \right] \\ \leq \mathbb{E} \left[\frac{1}{|\mathcal{Y}_{k}(c_{i},a)|} \sum_{Y_{t} \in \mathcal{Y}_{k}(c_{i},a)} \int_{\mathcal{X}} \left[\max_{b \in \mathcal{A}}(\Gamma_{\mathrm{NN}}q^{k})(y,b) \right] p(y \mid Y_{t},a) - p(y \mid c_{i},a)|\lambda(dy) \mid \mathcal{F}^{k} \right] \\ \leq \mathbb{E} \left[\frac{\sup_{y \in \mathcal{X}} \max_{b \in \mathcal{A}}(\Gamma_{\mathrm{NN}}q^{k})(y,b)}{|\mathcal{Y}_{k}(c_{i},a)|} \sum_{Y_{t} \in \mathcal{Y}_{k}(c_{i},a)} \int_{\mathcal{X}} p(y \mid Y_{t},a) - p(y \mid c_{i},a)|\lambda(dy) \mid \mathcal{F}^{k} \right] \\ \leq \sup_{y \in \mathcal{X}} \max_{b \in \mathcal{A}} \left| \sum_{j=1}^{n} K(y,c_{j})q^{k}(c_{j},b) \right| \times \mathbb{E} \left[\frac{1}{|\mathcal{Y}_{k}(c_{i},a)|} \sum_{Y_{t} \in \mathcal{Y}_{k}(c_{i},a)} \int_{\mathcal{X}} p(y \mid Y_{t},a) - p(y \mid c_{i},a)|\lambda(dy) \mid \mathcal{F}^{k} \right] \\ \leq ||q^{k}||_{\infty} \times \mathbb{E} \left[\frac{1}{|\mathcal{Y}_{k}(c_{i},a)|} \sum_{Y_{t} \in \mathcal{Y}_{k}(c_{i},a)} \int_{\mathcal{X}} W_{p}(y)\rho(Y_{t},c_{i})\lambda(dy) \mid \mathcal{F}^{k} \right] \\ \leq ||q^{k}||_{\infty} \wedge M_{p}. \end{split}$$

Putting together, we have

$$\mathbb{E}\left[w^{k+1}(c_i, a) \mid \mathcal{F}^k\right] \leq h(M_r + \gamma M_p \left\|q^k\right\|_{\infty}), \quad \forall (c_i, a) \in \mathcal{Z}_h.$$

Hence the noise w^{k+1} satisfies Assumption 1 of Theorem 3 with

$$\delta_1 = hM_r, \delta_2 = h\gamma M_p.$$

From Lemma 3, we have

$$|w^{k+1}(c_i, a) - \mathbb{E} \left[w^{k+1}(c_i, a) \mid \mathcal{F}^k \right] | \le 2V_{\max}, \quad \forall (c_i, a) \in \mathcal{Z}_h, \\ ||q^k||_{\infty} \le V_{\max}.$$

Therefore, the remaining Assumptions 2-3 of Theorem 3 are satisfied. And the update algorithm uses the learning rate suggested in Theorem 3. Therefore, we conclude that for each $0 < \varepsilon < 2V_{\max}\beta$ (since $\beta \ge 1$ and hence $2V_{\max}\beta \le \min\{2V_{\max}\beta, 4V_{\max}\beta^2\}$), after

$$k = \frac{192V_{\max}^3\beta^4}{\varepsilon^3}\log\left(\frac{128N_h |\mathcal{A}| V_{\max}\beta^4}{\delta\varepsilon^2}\right) + \frac{6V_{\max}(\beta-1)}{\varepsilon}$$

iterations of (9), with probability at least $1 - \delta$, we have

$$\left\|q^k - q_h^*\right\|_{\infty} \le \beta h(M_r + \gamma M_p V_{\max}) + \varepsilon.$$

To prove Theorem [1], we need the following result which bounds the number of time steps required to visit all ball-actions k times with high probability.

Lemma 6. (Lemma 14 in [2], rephrased) Let Assumption 2 hold. Then for all initial state $x_0 \in \mathcal{X}$, and for each integer $k \ge 4$, after a run of $T = 8kL_h \log \frac{1}{\delta}$ steps, the finite space \mathcal{Z}_h is covered at least k times under the policy π with probability at least $1 - \delta$ for any $\delta \in (0, \frac{1}{e})$.

Now we are ready to prove Theorem 1

Proof. We denote by \tilde{Q}_h^k the extension of q^k to \mathcal{Z} via the nearest neighbor operation, i.e., $\tilde{Q}_h^k = \Gamma_{\text{NN}}q^k$. Recall that Q_h^k is the extension of q_h^k (the fixed point of Gq = q) to \mathcal{Z} . We have

It follows from Theorem 4 that, after

$$k = \frac{192V_{\max}^3\beta^4}{\varepsilon_0^3}\log\left(\frac{128N_h |\mathcal{A}| V_{\max}^2\beta^4}{\delta\varepsilon_0^2}\right) + \frac{6V_{\max}(\beta-1)}{\varepsilon_0}$$

iterations, with probability at least $1 - \delta$, we have

$$\left\|\tilde{Q}_{h}^{k}-Q^{*}\right\|_{\infty}\leq\beta h(M_{r}+\gamma M_{p}V_{\max})+\beta Ch+\varepsilon_{0}=2\beta Ch+\varepsilon_{0}.$$

By setting $\varepsilon_0 = \frac{\varepsilon}{2}$ and $h^*(\varepsilon) = \frac{\varepsilon}{4\beta C}$, we have $\left\| \tilde{Q}_h^k - Q^* \right\|_{\infty} \le \varepsilon$. Let N_{h^*} be the h^* -covering number of the metric space (\mathcal{X}, ρ) . Plugging the result of Lemma 6 concludes the proof of Theorem 1

F Proof of Corollary 1

Proof. Since the probability measure ν is uniform over \mathcal{X} , we have $\nu_{\min} \triangleq \min_{i \in [N_{h^*}]} \nu(\mathcal{B}_i) = O(\frac{1}{N_{h^*}})$. By Proposition 1 the expected covering time of a purely random policy is upper bounded by

$$L_{h^*} = O\left(\frac{mN_{h^*}|\mathcal{A}|}{\psi}\log(N_{h^*}|\mathcal{A}|)\right).$$

By Proposition 4.2.12 in [46], the covering number N_{h^*} of $\mathcal{X} = [0,1]^d$ scales as $O((1/h^*)^d)$, which is $O((\beta/\varepsilon)^d)$ with $h^* = \frac{\varepsilon}{4\beta C}$.

From Theorem 1, with probability at least $1 - \delta$ we have $\|Q_{h^*}^T - Q^*\|_{\infty} \leq \varepsilon$, after at most

$$T = O\left(\frac{|\mathcal{A}|\beta^{d+7}}{\varepsilon^{d+3}}\log\left(\frac{2}{\delta}\right)\log\left(\frac{|\mathcal{A}|\beta^{d}}{\varepsilon^{d}}\right)\log\left(\frac{|\mathcal{A}|\beta^{d+6}}{\delta\varepsilon^{d+2}}\right)\right)$$

steps. Corollary [] follows after absorbing the dependence on $|\mathcal{A}|, d, \beta$ into $\kappa \equiv \kappa(|\mathcal{A}|, d, \beta)$ and doing some algebra.

G Proof of Theorem 2

We prove Theorem 2 by connecting the problem of estimating the value function in MDPs to the problem of non-parametric regression, and then leveraging known minimax lower bound for the latter. In particular, we show that a class of non-parametric regression problem can be embedded in an MDP problem, so any algorithm for the latter can be used to solve the former. Prior work on non-parametric regression[45] [39] establishes that a certain number of observations is *necessary* to achieve a given accuracy using *any* algorithms, hence leading to a corresponding necessary condition for the sample size of estimating the value function in an MDP problem.

We now provide the details.

Step 1. Non-parametric regression

Consider the following non-parametric regression problem: Let $\mathcal{X} := [0, 1]^d$ and assume that we have T independent pairs of random variables $(x_1, y_1), \ldots, (x_T, y_T)$ such that

$$\mathbb{E}\left[y_t|x_t\right] = f(x_t), \qquad x_t \in \mathcal{X}$$
(25)

where $x_t \sim \text{uniform}(\mathcal{X})$ and $f : \mathcal{X} \to \mathbb{R}$ is the unknown regression function. Suppose that the conditional distribution of y_t given $x_t = x$ is a Bernoulli distribution with mean f(x). We also assume that f is 1-Lipschitz continuous with respect to the Euclidean norm, i.e.,

$$|f(x) - f(x_0)| \le |x - x_0|, \quad \forall x, x_0 \in \mathcal{X}.$$

Let \mathcal{F} be the collection of all 1-Lipschitz continuous function on \mathcal{X} , i.e.,

 $\mathcal{F} = \operatorname{Lip}\left(\mathcal{X}, 1\right) = \left\{ h | h \text{ is a 1-Lipschitz function on } \mathcal{X} \right\},\$

where $\text{Lip}(\cdot, \cdot)$ is as defined in Section 2. The goal is to estimate f given the observations $(x_1, y_1), \ldots, (x_T, y_T)$ and the prior knowledge that $f \in \mathcal{F}$.

It is easy to verify that the above problem is a special case of the non-parametric regression problem considered in the work by Stone [39] (in particular, Example 2 therein). Let \hat{f}_T denote an arbitrary (measurable) estimator of f based on the training samples $(x_1, y_1), \ldots, (x_T, y_T)$. By Theorem 1 in [39], we have the following result: there exists a c > 0 such that

$$\lim_{T \to \infty} \inf_{\hat{f}_T} \sup_{f \in \mathcal{F}} \Pr\left(\left\| \hat{f}_T - f \right\|_{\infty} \ge c \left(\frac{\log T}{T} \right)^{\frac{1}{2+d}} \right) = 1,$$
(26)

where infimum is over all possible estimators \hat{f}_T .

Translating this result to the non-asymptotic regime, we obtain the following theorem.

Theorem 5. Under the above assumptions, for any number $\delta \in (0, 1)$, there exits some numbers c > 0 and T_{δ} such that

$$\inf_{\hat{f}_n} \sup_{f \in \mathcal{F}} \Pr\left(\left\|\hat{f}_T - f\right\|_{\infty} \ge c \left(\frac{\log T}{T}\right)^{\frac{1}{2+d}}\right) \ge \delta, \quad \text{for all } T \ge T_{\delta}.$$

Step 2. MDP

Consider a class of (degenerate) discrete-time discounted MDPs $(\mathcal{X}, \mathcal{A}, p, r, \gamma)$ where

$$\begin{aligned} \mathcal{X} &= [0,1]^a, \\ \mathcal{A} \text{ is finite,} \\ p(\cdot|x,a) &= p(\cdot|x) \text{ is uniform on } \mathcal{X} \text{ for all } x, a, \\ r(x,a) &= r(x) \text{ for all } a, \\ \gamma &\in (0,1). \end{aligned}$$

In words, the transition is uniformly at random and independent of the current state and the actions taken, and the expected reward is independent on the action taken but dependent on the current state.

Let R_t be the observed reward at step t. We assume that the distribution of R_t given x_t is Bernoulli $(r(x_t))$, independently of $(x_1, x_2, \ldots, x_{t-1})$. The expected reward function $r(x_t) = \mathbb{E}[R(x_t)|x_t]$ is assumed to be 1-Lipschitz and bounded.

It is easy to see that for all $x \in \mathcal{X}$, $a \in \mathcal{A}$,

$$Q^{*}(x,a) = V^{*}(x) = r(x) + \gamma \mathbb{E}\left[V^{*}(X')|x\right]$$
$$= r(x) + \gamma \int_{\mathcal{X}} V^{*}(y)p(y|x)dy$$
$$= r(x) + \gamma \underbrace{\int_{\mathcal{X}} V^{*}(y)dy}_{C},$$
(27)

where the last step holds because $p(\cdot|x)$ is uniform. Integrating both sides over \mathcal{X} , we obtain

$$C = \int_{\mathcal{X}} r(x) dx + \gamma C,$$

whence

$$C = \frac{1}{1 - \gamma} \int_{\mathcal{X}} r(x) dx.$$

It follows from equation (27) that

$$V^*(x) = r(x) + \frac{\gamma}{1 - \gamma} \int_{\mathcal{X}} r(y) dy, \qquad \forall x \in \mathcal{X},$$
(28)

and

$$r(x) = V^*(x) - \gamma \int_{\mathcal{X}} V^*(y) dy, \qquad \forall x \in \mathcal{X}.$$
(29)

Regardless of the exploration policy used, the sample trajectory (x_1, x_2, \ldots, x_T) is i.i.d. and uniformly random over \mathcal{X} , and the observed rewards (R_1, R_2, \ldots, R_T) are independent.

Step 3. Reduction from regression to MDP

Given a non-parametric regression problem as described in Step 1, we may reduce it to the problem of estimating the value function V^* of the MDP described in Step 2. To do this, we set

$$r(x) = f(x) - \gamma \int_{\mathcal{X}} f(y) dy, \qquad \forall x \in \mathcal{X}$$

and

$$R_t = y_t, \qquad t = 1, 2, \dots, T.$$

In this case, it follows from equations (28) and (29) that the value function is given by $V^* = f$. Moreover, the expected reward function $r(\cdot)$ is 1-Lipschitz as it is just $f(\cdot)$ minus a constant, so the assumptions of the MDP in Step 2 are satisfied. This reduction shows that the MDP problem is at least as hard as the nonparametric regression problem, so a lower bound for the latter is also a lower bound for the former. Applying Theorem [5] yields the result stated in Theorem [2].