
Sample Efficient Reinforcement Learning In Continuous State Spaces: A Perspective Beyond Linearity

Dhruv Malik¹ Aldo Pacchiano² Vishwak Srinivasan¹ Yuanzhi Li¹

Abstract

Reinforcement learning (RL) is empirically successful in complex nonlinear Markov decision processes (MDPs) with continuous state spaces. By contrast, the majority of theoretical RL literature requires the MDP to satisfy some form of linear structure, in order to guarantee sample efficient RL. Such efforts typically assume the transition dynamics or value function of the MDP are described by linear functions of the state features. To resolve this discrepancy between theory and practice, we introduce the Effective Planning Window (EPW) condition, a structural condition on MDPs that makes *no* linearity assumptions. We demonstrate that the EPW condition permits sample efficient RL, by providing an algorithm which provably solves MDPs satisfying this condition. Our algorithm requires minimal assumptions on the policy class, which can include multi-layer neural networks with nonlinear activation functions. Notably, the EPW condition is directly motivated by popular gaming benchmarks, and we show that many classic Atari games satisfy this condition. We additionally show the necessity of conditions like EPW, by demonstrating that simple MDPs with slight nonlinearities cannot be solved sample efficiently.

1. Introduction

Over the past decade, reinforcement learning (RL) has emerged as the dominant paradigm for sequential decision making in modern machine learning. During this time period, video games have served as popular means to benchmark the incremental improvement in state of the art RL. The Arcade Learning Environment (ALE), comprising a suite of classic Atari games, is an archetypical example of

such a benchmark (Bellemare et al., 2013). Agents trained by RL efficiently learn to surpass human level performance in such games (Mnih et al., 2013; 2015; Badia et al., 2020).

Motivated by these empirical accomplishments, there has been a major thrust to theoretically characterize the conditions which permit sample efficient RL. A significant line of work has greatly advanced our understanding of the tabular RL setting, where the number of states is finite and relatively small (Simchowitz & Jamieson, 2019; Pananjady & Wainwright, 2021). Sample efficiency bounds in this setting scale with cardinality of the state space. However, in practice this cardinality is often large or infinite. For instance, many gaming applications of RL, ranging in complexity from Atari to Dota, all have continuous state spaces (Berner et al., 2019). These scenarios are handled in the function approximation setting (Du et al., 2019c; 2020a). Here, each state is associated with a known feature, and one desires a sample efficiency bound that scales with the dimensionality of the features (instead of the cardinality of the state space).

To understand when continuous space RL is sample efficient, theoreticians make certain assumptions on the features or the underlying Markov Decision Process (MDP). A prominent assumption, which has appeared in various forms, is that the problem satisfies some sort of *linear* structure. For instance, in the well studied linear MDP, the transitions and rewards are described by linear functions of the features (Yang & Wang, 2019; Jin et al., 2020; Yang & Wang, 2020; Wang et al., 2021). In particular, the transition probabilities at a state-action pair are defined by linear functions of the feature corresponding to that state-action pair. A weaker, but frequently occurring, form of this assumption is that value function of any policy is nearly linear (Du et al., 2020a; Lattimore et al., 2020), or that the optimal value function is linear (Du et al., 2019c). Such linear structure is amenable to theoretical analysis, since it permits analysts to leverage the vast literature on supervised and online learning.

To obtain a holistic understanding of RL, examining such linear structure is certainly important. Nevertheless, it is unclear whether the aforementioned linearity conditions actually hold in practical scenarios. We illustrate this via a very simple example. Consider an MDP where there is a set of n states with the property that taking any action at one of

¹Machine Learning Department, Carnegie Mellon University

²Department of EECS, UC Berkeley. Correspondence to: Dhruv Malik <dhruvm@andrew.cmu.edu>.




Figure 1. An image of the Atari Pong game. The green paddle must move up and down to hit the ball (the white dot) while playing against the opposing orange paddle.

these states leads to the same state. To cast this MDP in the aforementioned linear MDP setting, the dimensionality of the state features would have to scale linearly with n . This precludes the existence of algorithms that can solve this MDP with sample complexity independent of the cardinality of the state space.

Moreover, it has recently been shown both theoretically and empirically that the optimal value function and optimal policy can be very complex, even in ostensibly elementary continuous state space MDPs (Dong et al., 2020a). Since even powerful linear functions such as the Neural Tangent Kernel are significantly worse in terms of representation power than nonlinear neural networks (Allen-Zhu & Li, 2019; Li et al., 2020), it is unclear whether such weaker linear functions can be used to approximate the value function or the underlying policy well.

Even in simple RL gaming benchmarks, there is no evidence that the aforementioned linearity assumptions hold. Indeed, nonlinear neural networks are predominant means to approximate policies and value functions when solving these games in practice. For instance, consider the Pong game from the ALE benchmark, which is depicted in Figure 1. In this game, the agent must use its paddle to prevent the ball from crossing a boundary, while playing against a pseudorandom opposing paddle. Despite the simplicity of Pong, state of the art methods solve this game using neural networks (Mnih et al., 2013; 2015; Badia et al., 2020), and it is not apparent whether this game is linear in any sense.

This reveals a significant gap between the theory and practice of RL. In theory, one usually employs some sort of linearity assumption to ensure efficient RL. Yet, in practice, RL appears to succeed in domains which do not satisfy such linear structure. In an effort to resolve this discrepancy, we ask the following question:

Which structure is typical of popular RL domains, and does this structure permit sample efficient RL?

This question underlies the analysis of our paper. We make the following contributions:

- We propose the Effective Planning Window (EPW) condition, a structural condition for MDPs which goes beyond linearity. Indeed, this condition is compati-

ble with neural network policies, and MDPs satisfying EPW can have highly nonlinear (stochastic) transitions. Informally, this condition requires the agent to *consistently plan C timesteps ahead*, for a value of C significantly smaller than the horizon length. We show that popular Atari benchmark games satisfy this condition.

- We provide a simple algorithm, which exploits the EPW condition to provably solve MDPs satisfying EPW. We prove the sample efficiency of our algorithm, and show that it requires a number of trajectories that is a lower order polynomial of the horizon length and other relevant problem dependent quantities.
- We argue that one must look beyond linear structure, and further motivate the study and necessity of conditions like EPW, by demonstrating that even slightly nonlinear MDPs cannot be solved sample efficiently.

2. Related Work

Linear Function Approximation. The majority of RL literature in the function approximation setting focuses on MDPs that satisfy some form of linear structure. A notable example is the linear MDP, where the transitions and rewards are described by linear functions of the state features (Yang & Wang, 2019; Jin et al., 2020; Yang & Wang, 2020). A weaker form of this assumption is that the value function for each policy is nearly linear (Du et al., 2020a; Lattimore et al., 2020), or that the optimal value function is linear (Du et al., 2019c; Weisz et al., 2021a;b). We note that the algorithm of Weisz et al. (2021a) requires a generative model, while we work in the standard episodic RL setting. As argued earlier, such linear assumptions are unlikely to hold true in practice. In our work, we eschew any sort of linearity assumption.

Nonlinear Function Approximation. Empirically, it is typical to use nonlinear function approximators such as neural networks (Schulman et al., 2015; Levine et al., 2016). But from a theoretical perspective, the understanding of nonlinear function approximation is limited. There is prior work which studies the sample complexity of RL when using function approximation with nonlinear function classes (Wen & Van Roy, 2013; Jiang et al., 2017; Van Roy & Dong, 2019; Dong et al., 2020b; Du et al., 2020b; Wang et al., 2020; Jin et al., 2021; Wang et al., 2021). However, these works often are restricted to MDPs with deterministic transitions (Wen & Van Roy, 2013; Van Roy & Dong, 2019; Du et al., 2020b). In this deterministic setting, an algorithm can repeatedly visit the same state and simply memorize an optimal path. By contrast, we focus on MDPs with *stochastic* transitions, as is typical in many Atari games. Here, an algorithm generally cannot visit the same state more than once, and must generalize beyond the trajectories it samples to learn something global. Moreover, the aforementioned analyses of nonlin-

ear function approximation typically make some stringent assumption on the complexity of the function class (Jiang et al., 2017; Dong et al., 2020b; Wang et al., 2020; Jin et al., 2021; Wang et al., 2021). Such complexity measures either cannot or are not known to handle neural networks. By contrast, our results can handle nonlinear multilayer neural networks. In a different line of work, Dai et al. (2018) study RL with nonlinear function approximators and provide a convergent algorithm for this setting. However, they do not precisely relate the quality of the solution found by their algorithm to the approximation error of the function class.

Linear Quadratic Regulator. To characterize the sample complexity of RL in continuous state spaces, a different line of work investigates the linear quadratic regulator (LQR) (Fazel et al., 2018; Dean et al., 2019; Malik et al., 2020). Here, the transition dynamics of the MDP are assumed to be noisy linear functions of the state and action, and the rewards are quadratic functions of the state and action. We remark that in this setting, the action space is continuous. By contrast, we exclusively study MDPs with finite action spaces, since these are most typical in the RL video game domains that motivate our paper.

3. Problem Formulation

3.1. Problem Statement

Notation & Preliminaries. We use the notation $[n]$ to denote $\{0, 1 \dots n - 1\}$ for any positive integer n . Recall that an undiscounted, finite horizon MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{T}, R, H)$ is defined by a set of states \mathcal{S} , a set of actions \mathcal{A} , a transition function \mathcal{T} which maps from state-action pairs to a probability density defined over states, a reward function R which maps from state-action pairs to non-negative real numbers, and a finite planning horizon H . Throughout our paper, we assume that $\mathcal{S} \subseteq \mathbb{R}^d$ and \mathcal{A} is a finite set. We assume a single initial state s_0 . Note that this is without loss of generality, since we can always extend the horizon to $H + 1$ where a single dummy initial state connects to an arbitrary initial state distribution. For simplicity, we assume that \mathcal{S} can be partitioned into H different levels. This means that for each $s \in \mathcal{S}$ there exists a unique $h \in [H]$ such that it takes h timesteps to arrive at s from s_0 . We say that such a state s lies on level h , and denote \mathcal{S}_h to be the set of states on level h . Note this assumption is without loss of generality, and our results apply to generic MDPs which cannot be partitioned into levels. This is because we can always make the final coordinate of each state encode the number of timesteps that elapsed to reach the state. Taking any action from level $H - 1$ exits the game. The notation $\|x\|_2$ denotes the Euclidean norm of x .

A policy maps each state to a corresponding distribution over actions. In practice, one typically uses a policy that is parameterized by parameters belonging to some set $\Theta \subseteq \mathbb{R}^k$.

We study such policies, and use $\pi(\theta)$ to denote the policy induced by using parameter $\theta \in \Theta$. When discussing a policy which is not parameterized, we simply use π to denote the policy. We use $\pi_s^a(\theta)$ to denote the probability of taking action a at state s when using the policy $\pi(\theta)$. We use $\pi(\Theta)$ to denote $\{\pi(\theta) \text{ s.t. } \theta \in \Theta\}$, which is the set of feasible policies and defines our policy class. Given a vector $\bar{\theta} \in \Theta^H$, we let $\pi(\bar{\theta})$ denote the policy which executes $\pi(\bar{\theta}_h)$ at for any state lying on level $h \in [H]$, where $\bar{\theta}_h$ denotes the h^{th} entry of $\bar{\theta}$. The value of a policy $\pi(\theta)$ in a (stochastic) MDP \mathcal{M} when initialized at state s is denoted $V_{\mathcal{M}}^s(\pi(\theta))$. It is given by $V_{\mathcal{M}}^s(\pi(\theta)) = \mathbb{E} \left[\sum_{h=0}^{H-1} R(s_h, a_h) \right]$, where the expectation is over the trajectory $\{(s_h, a_h)\}_{h=0}^{H-1}$. Given an accuracy ϵ and failure probability tolerance δ , the goal of RL is to find a policy $\pi(\theta)$ which satisfies $V_{\mathcal{M}}^{s_0}(\pi(\theta)) \geq \max_{\theta' \in \Theta} V_{\mathcal{M}}^{s_0}(\pi(\theta')) - \epsilon$ with probability at least $1 - \delta$.

Query Model. We adopt the standard episodic RL setup. During each episode, an agent is allowed to interact with the MDP by starting from s_0 , taking an action and to observe the next state and reward, and repeating. The episode terminates after the agent takes H actions, and the next episode starts at s_0 . The agent thus takes a single trajectory in each episode, and the total query complexity of the agent is measured by the total number of trajectories. Given a desired solution accuracy ϵ and failure probability tolerance δ , we are interested in algorithms which can successfully solve an MDP using a number of trajectories that is at most polynomial in H , $|\mathcal{A}|$, d , k , $\frac{1}{\epsilon}$ and $\frac{1}{\delta}$. If an algorithm provably accomplishes this, we call such an algorithm *sample efficient* or *tractable*. Notably, such algorithms cannot depend on the (possibly uncountable) number of states.

Without any assumptions on the MDP, approximating an optimal policy is intractable. To permit sample efficient RL, prior theoretical work has often assumed that MDP satisfies some form of linear structure. For instance, the transition or value function might be described by a linear function of the states. However, it is well documented that RL is empirically successful in highly nonlinear domains (Schulman et al., 2015; Levine et al., 2016). We aim to reduce this gap between theory and practice. We now formally state the problem that we consider throughout our paper.

Our goal is to present nonlinear characteristic conditions which permit sample efficient RL, and argue that these conditions are satisfied in practice by popular RL domains.

3.2. Effective Planning Window Condition

We first state basic conditions that are satisfied by most RL problems encountered in practice. We will later refine these to obtain our Effective Planning Window (EPW) condition, and then show that EPW enables sample efficient RL.

Let us begin by observing that in practice, the policy class




Figure 2. An image of the Atari Skiing game. The skier must move through flagged checkpoints.

$\pi(\Theta)$ typically satisfies some mild regularity assumptions. We formalize this in the following condition.

Condition 1 (Regular Policy Class). A policy class $\pi(\Theta)$ is said to be Regular when:

- (a) **Bounded Domain.** There exists $B > 0$ such that each $\theta \in \Theta$ satisfies $\|\theta\|_2 \leq B$.
- (b) **Lipschitz Continuous Policies.** There exists $\phi > 0$ such that for any $\theta, \theta' \in \Theta$ and any $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have $|\pi_s^a(\theta) - \pi_s^a(\theta')| \leq \phi \|\theta - \theta'\|_2$.

We stress that this is a very mild condition, and places minimal restrictions on $\pi(\Theta)$. Indeed, a policy parameterized by a multi-layer neural network with a nonlinear activation function satisfies this condition (Fazlyab et al., 2019). Using this condition on the policy class, we now introduce the following Generic Game condition. As we will discuss in the sequel, many popular gaming RL benchmarks such as Atari games satisfy this condition.

Condition 2 (Generic Game). An MDP and Regular policy class pair $(\mathcal{M}, \pi(\Theta))$ form a Generic Game if:

- (a) **Failure States.** There is a set of failure states $\mathcal{F} \subset \mathcal{S}$, and taking any action from a state in \mathcal{F} exits the game.
- (b) **Complete Policy Class.** There exists some $\theta^* \in \Theta$ such that executing $\pi(\theta^*)$ from s_0 arrives at some state in $\mathcal{S}_{H-1} \setminus \mathcal{F}$ almost surely¹. We define \mathcal{S}^* to be the set of all states $s \in \mathcal{S} \setminus \mathcal{F}$ such that executing $\pi(\theta^*)$ from s reaches $\mathcal{S}_{H-1} \setminus \mathcal{F}$ almost surely. If a state lies in \mathcal{S}^* we call it a safe state.
- (c) **Binary Rewards.** For any state $s \in \mathcal{S}_{H-1} \setminus \mathcal{F}$ and any $a \in \mathcal{A}$, $R(s, a) = 1$. For any other state s and any $a \in \mathcal{A}$, $R(s, a) = 0$.

A few comments are in order. Note that \mathcal{F} is essentially used to describe states where the agent has lost the game. Also, observe that in Generic Games, an optimal policy is one that arrives at a non-failure state in level $H - 1$ almost surely. Hence $\pi(\theta^*)$ is indeed an optimal policy.

¹The Generic Game condition can also be defined in the case when this property of θ^* holds true with probability exponentially large in H , as would occur when using a softmax policy class. Our results hold true when using this notion of a Generic Game. We focus on the almost sure case to avoid complicating notation.

Let us now describe how popular Atari games can be cast as Generic Games. Recall the famous Pong game depicted in Figure 1, which is a part of the ALE benchmark (Bellemare et al., 2013). In this game, an RL agent must learn to move the paddle up and down to hit the ball and prevent it from crossing its boundary. Note that in the context of RL, this is a single player game, since the opposing paddle hits the ball back according to a pre-specified stochastic decision rule (which is not trained). The agent loses the game if the ball crosses its own boundary, and wins the game if it hits the ball past the opposing paddle. Another game in the ALE benchmark is the Skiing game, depicted in Figure 2. Here, the agent must move the skier through a series of randomly appearing flagged checkpoints, which appear frequently over a long time horizon. The skier receives a penalty each time it misses a checkpoint.

We claim that an Atari game like Pong or Skiing, together with a neural network policy class, satisfy the Generic Game condition. The first two conditions of Generic Games are easy to verify. Note that the states in Pong (resp. Skiing) are images, so \mathcal{F} includes any state where the ball has crossed the agent’s boundary (resp. where the skier has missed a prior checkpoint), since this corresponds to the agent failing to complete the game. It is known that Atari can be solved using a neural network policy (Mnih et al., 2015), so a policy class parameterized by neural networks is complete.

To ensure that Pong and Skiing satisfy the third condition, we need to design an appropriate binary reward function. For Pong, this is handled by redefining \mathcal{F} to include any state $s \in \mathcal{S}_{H-1}$ where the ball has not crossed the opposing paddle. Similarly for Skiing, this is done by ensuring \mathcal{F} includes any state where the skier has already missed a checkpoint. Then one can simply assign a reward of 1 to any state in $\mathcal{S}_{H-1} \setminus \mathcal{F}$, and 0 to all other states, as required by the Generic Game condition. Hence, playing optimally in this Generic Game framework ensures the ball has moved past the opposing paddle, or the skier has made all checkpoints, corresponding to winning the game.

The aforementioned reward design is an example of reward shaping, which is unavoidable in RL and ubiquitous in practice (Hadfield-Menell et al., 2017). Nevertheless, we stress that the reward function we described above is very similar to the reward function that practitioners already use. Concretely, in Pong one typically assigns a reward of 1 if the ball has moved past the opposing paddle, a reward of -1 if the ball has moved past the agent’s paddle, and a reward of 0 otherwise (Bellemare et al., 2013). And in Skiing, the skier receives reward at the end of the game, in proportion to the number of checkpoints it has cleared (Badia et al., 2020). Our reward function thus requires no more effort to design than the reward functions already in use, since they require the same information, and these are identical in spirit.




Figure 3. Five states from the Pong game. Here we let $t'' < t - C < t' < t$, and the ball is progressively moving towards the lower right corner. At timesteps t', t , the paddle has not lost the game. However, it does not have enough time to react and reach the ball in time. At timestep $t + 1$ the game is over. At timesteps $t'', t - C$, the paddle has enough time to react and reach the ball.

Beyond Pong and Skiing, other Atari games (and similarly themed video games) can be cast in the Generic Game framework. In Appendix A, we describe this reduction for the Atari games Tennis & Journey Escape, and also for the more complex gaming benchmark CoinRun (Cobbe et al., 2019).

We make one more remark about the binary reward structure of Generic Games. There are many games which naturally have a set of goal states, and these immediately can be described as Generic Games. Examples include Pong, Tennis and CoinRun (latter two are described in Appendix A). More generally, however, EPW applies to many games which do not naturally have a binary reward structure. As we discussed, Skiing can be cast in the EPW framework with binary rewards, by ensuring that \mathcal{F} includes any state where the skier has missed a checkpoint. However, in certain scenarios, one may be satisfied with only collecting a large fraction of the checkpoints in Skiing, or more generally, obtaining a large (but not perfect) score in games where one continually collects small rewards. We emphasize that such scenarios can be cast in our Generic Game framework. For instance, in Skiing if checkpoints arrive roughly every x timesteps, and we desire to give a reward of 1 for each checkpoint collected, then we can design \mathcal{F} to include any state at timestep t where the agent has not collected $\Omega(t/x)$ reward thus far. Similar reductions apply to other games where one continually needs to collect a small amount of reward at regular intervals.

Does the Generic Game condition permit sample efficient RL? Unfortunately, there exist Generic Games where the MDP is only slightly nonlinear, but even approximating an optimal policy sample efficiently is impossible. We later show this formally in Proposition 1. So we must further restrict this class of games. In order to refine our notion of a Generic Game, we first state a useful definition.

Definition 1 (x-Ancestor). Given a Generic Game $(\mathcal{M}, \pi(\Theta))$, consider any $h \in [H]$ and any state $s' \in \mathcal{S}_h$. A state $s \in \mathcal{S}$ is an x -ancestor of s' , if $s \in \mathcal{S}_{\max\{0, h-x\}}$ and there exists some $\theta \in \Theta$ such that following $\pi(\theta)$ from s will reach s' with nonzero probability.

We are now formally state our Effective Planning Window (EPW) condition, which refines our notion of Generic Games. For the statement of the condition, recall our notion of S^* , which was defined in the Generic Game condition.

Condition 3 (Effective Planning Window). A Generic Game $(\mathcal{M}, \pi(\Theta))$ satisfies the Effective Planning Window condition with parameter C if there exists $C \in [H]$ such that the following holds. Consider any $s' \in \mathcal{S} \setminus \mathcal{F}$. If s is a C -ancestor of s' , then $s \in S^*$.

Before examining RL benchmark games in the context of this condition, a few comments about the condition itself are in order. The quantity C ensures that any C -ancestor of a non-failure state is a safe state. So if an agent is at timestep t and the game is not over, then at timestep $t - C$ it was in a state from where it could have achieved the highest reward possible in the MDP (if it took the correct sequence of actions). For the purposes of RL, this effectively means that at each timestep, the agent must consistently plan over the next C timesteps instead of the entire horizon length H . When C is small or a constant, then it is reasonable to believe that sample efficient RL is possible.

Of course, any Generic Game satisfies the EPW condition for a choice of $C = H - 1$. However, many popular RL benchmark games satisfy the EPW property with a value of C that is much smaller than H . Informally, the C quantity is the amount of time required by the agent to successfully react to scenarios in the game (without losing). Let us understand this in the Pong and Skiing games.

In Pong, after the opposing paddle hits the ball, the agent must react to the trajectory of the ball and adjust its position accordingly to hit it. If it takes too long to react before it starts adjusting its position, then it will be unable to reach the ball in time. We depict this in Figure 3. More formally, assume that at timestep t the paddle has not lost the game and the ball is moving towards its boundary. At timestep t , the ball may be too close to the boundary, and so the agent will not have enough time to move its paddle fast enough in order to reach the ball in time. However, at timestep $t - C$




Figure 4. Five states from the Skiing game. Here we let $t'' < t - C < t' < t$, and the skier is progressively moving towards the checkpoint. At timesteps t' , t , the skier has not lost. However, it does not have enough time to react and adjust its position to reach the checkpoint in time. At timestep $t + 1$ the game is over. At timesteps t'' , $t - C$, the skier has enough time to react and reach the checkpoint.

the ball is further away from the boundary, so the agent has enough time to move its paddle appropriately in order to react, reach the ball and hit it back. So at timestep $t - C$ the agent lies in a safe state in \mathcal{S}^* , since it has enough time to adjust its paddle and hit the ball back, and hence play optimally. Notably, if we let C' be the number of timesteps it takes for the ball to traverse from one end of the board to the other, then $C \leq C'$. Hence, when H is large and the agent needs to control the paddle for many rounds, then C is a constant independent of H .

Similarly, in the Skiing game, the skier must react to the location of the oncoming checkpoint, and adjust its position accordingly. Formally, assume at timestep t a checkpoint is oncoming. In such a scenario, as depicted in Figure 4, the skier might be too far from the checkpoint in order to actually clear it (even if it moves directly towards the checkpoint). However, at timestep $t - C$ the skier has enough time to adjust its position in order to clear the checkpoint. Hence at timestep $t - C$, the skier is in a safe state in \mathcal{S}^* since it can play optimally from this state. Again, if we let C' be the number of timesteps it takes for a skier to move from the left edge of the screen to the right edge, then $C \leq C'$. Hence, when H is large and there are many checkpoints to be cleared, then C is a constant independent of H , as previously observed in Pong.

Beyond Pong and Skiing, other Atari games satisfy the EPW condition, with a constant value of C . We demonstrate this for the Atari games Tennis & Journey Escape in Appendix A. In Appendix A we additionally show that more complex games, such as CoinRun (Cobbe et al., 2019), also satisfy EPW with a small value of C . We stress that EPW is orthogonal to linearity. Indeed, there are MDPs satisfying linearity but not EPW, and vice versa. We conclude this section by highlighting two important aspects of the EPW condition.

The Magnitude Of C . We treat C as a constant that is independent of and much smaller than H . This is certainly reasonable given our above discussion. So an algorithm incurring $\mathcal{O}(|\mathcal{A}|^C)$ sample complexity is efficient. Furthermore, as we discuss later, there exist EPW games where $\Omega(|\mathcal{A}|^C)$ sample complexity is necessary to solve the game.

The Challenge Of Solving EPW Games. We note that a deterministic EPW game is straightforward to solve, since an agent can just try each of the $|\mathcal{A}|^C$ trajectories when it is at level h , to discover which trajectories do not lead to \mathcal{F} . In such a case, an agent can simply memorize a good path. However, when transitions are stochastic (as in Atari), the agent cannot simply try each trajectory to memorize one that does not lead to \mathcal{F} . This is because in general stochastic MDPs, a finite sample algorithm might only visit any given state at most once. Instead, the algorithm must *learn* and generalize beyond the trajectories it samples, to learn something global about the MDP. Furthermore, we emphasize that stochastic EPW games *cannot* be solved as simply as just splitting the horizon H into H/C distinct planning windows, and then solving these planning problems independently of each other. Instead, the key difficulty is that the agent must *consistently* plan C timesteps ahead. By this, we mean that just because an agent has arrived at a non-failure state at time t , does not imply that at time $t + 1$ it is guaranteed to avoid \mathcal{F} . Indeed, if we execute a policy and the resulting trajectory ends in a failure state after t timesteps, then it is unclear at which of the prior timesteps $\{t - C \dots t - 1\}$ that we took an incorrect action. And we cannot rollback to timestep $t - C$ and rerun the same trajectory to discover when we made a mistake. This complicates the design of efficient algorithms for this setting.

4. Main Results

We now turn to our main results. Before diving into the details, let us provide a brief overview for the results. Our main contribution is an algorithm which sample efficiently solves games satisfying the EPW condition. We prove the efficiency of this algorithm and discuss the role of the EPW condition in permitting sample efficient RL. We then further motivate the study of conditions like EPW, as opposed to the study of linear structure, by proving a lower bound which shows that Generic Games with even slight nonlinearities cannot be solved sample efficiently. With this outline in mind, let us now formally present our main results.

Our primary contribution is an algorithm which exploits the EPW condition to sample efficiently find an optimal policy.

Algorithm 1

- 1: Inputs: MDP \mathcal{M} , policy class $\pi(\Theta)$, sample size n
- 2: Initialize $\bar{\theta}(0) = [\theta_{\text{rand}}, \theta_{\text{rand}} \dots \theta_{\text{rand}}] \in \Theta^{H-1}$
- 3: **for** $t \in \{0, 1 \dots H - 2\}$ **do**
- 4: Sample n trajectories $\{\tau_i\}_{i=1}^n \sim \pi(\bar{\theta}(t))$, where each trajectory $\tau_i = \{(s_{i,h}, a_{i,h})\}_{h=0}^{H-1}$
- 5: Define the loss function $\widehat{L}_t : \Theta \rightarrow \mathbb{R}$ as

$$\widehat{L}_t(\theta) = |\mathcal{A}|^{C+1} \cdot \left[\frac{1}{n} \sum_{i=1}^n \left(\mathbb{1}_{s_{i,t+C+1} \in \mathcal{F}} \prod_{j=0}^C \pi^{a_{i,t+j}}(\theta) \right) \right] \quad (1)$$

- 6: Minimize \widehat{L}_t over Θ to obtain

$$\widehat{\theta}_t \in \underset{\theta \in \Theta}{\operatorname{argmin}} \widehat{L}_t(\theta) \quad (2)$$

- 7: Define $\bar{\theta}(t+1) = [\widehat{\theta}_0, \widehat{\theta}_1 \dots, \widehat{\theta}_{t-1}, \widehat{\theta}_t, \theta_{\text{rand}}, \dots, \theta_{\text{rand}}]$
- 8: **end for**
- 9: **return** $\bar{\theta}(H-1)$

Our algorithm is defined in Algorithm 1. For this method, recall our notation that for a vector $\bar{\theta} \in \Theta^H$, we let $\pi(\bar{\theta})$ denote the policy which executes $\pi(\bar{\theta}_h)$ at each $h \in [H]$, where $\bar{\theta}_h$ denotes the h^{th} entry of $\bar{\theta}$. We make the basic assumption that there exists $\theta_{\text{rand}} \in \Theta$ such that $\pi(\theta_{\text{rand}})$ maps each state to the uniform distribution over \mathcal{A} .

In our main result Theorem 1, we bound the sample complexity of Algorithm 1 and demonstrate that it efficiently finds a near optimal policy for any EPW game. Before stating the theorem, let us discuss Algorithm 1 and provide some intuition for this method. Recall that in Generic Games, a near optimal policy avoids failure states with high probability. The method initializes $\bar{\theta}(0)$ to be the parameter vector which induces a uniformly random policy regardless of the state (Line 2). It then incrementally updates $\bar{\theta}(t)$, in a fashion that ensures $\pi(\bar{\theta}(t))$ avoids failure states with high probability for each $t' \leq t$ (Line 7). Ultimately it returns the parameter vector $\bar{\theta}(H-1)$ (Line 9). More concretely, at each timestep t in the inner loop, the algorithm samples n trajectories from the policy $\pi(\bar{\theta}(t))$ that it has constructed thus far (Line 4). Via these sampled trajectories it defines the empirical loss \widehat{L}_t , as shown in Eq. (1). Intuitively, \widehat{L}_t penalizes those parameters θ such that executing $\pi(\theta)$ over timesteps $\{t, t+1 \dots t+C\}$ arrives at a failure state with high probability. This intuition suggests that a minimizer $\widehat{\theta}_t$ of \widehat{L}_t , as defined in Eq. (2), will assign low probability to trajectories ending in a failure state when playing $\widehat{\theta}_t$. The quantity $\widehat{\theta}_t$ then defines the $t+1^{\text{th}}$ entry of $\bar{\theta}(t+1)$. Note that $\bar{\theta}(t+1)$ agrees with $\bar{\theta}(t)$ in its first t entries.

The form of the loss \widehat{L}_t suggests that one needs to know the exact value of C to use Algorithm 1. We emphasize that any upper bound $C' \geq C$ can be used in place of C in the definition of \widehat{L}_t . As discussed in Section 3.2, such an upper bound is easy to find in gaming domains where we

expect the EPW condition to hold. For example, in Pong we can play the game manually (in OpenAI Gym), and observe the number of timesteps it takes the paddle to traverse its side. This yields $C \leq 15$, while $200 \leq H$. As to how this affects the sample complexity of the method, one can simply substitute C' for C in the bound provided in our main result Theorem 1.

Before formally presenting Theorem 1, a remark on the computational requirements of Algorithm 1 is imperative. The method requires oracle access to a minimizer $\widehat{\theta}_t$ of the loss \widehat{L}_t , which in turn is defined by the policy class $\pi(\Theta)$. In our paper, we impose minimal assumptions on $\pi(\Theta)$. Our motivation for this choice is that in practice, it is most common to parameterize a policy via a multi-layer neural network with nonlinear activation function. Beyond our extremely mild Regularity condition, is unclear which (if any) desirable properties such a policy class satisfies. Hence, for a worst case Regular policy class $\pi(\Theta)$, obtaining even an approximation of $\widehat{\theta}_t$ could be extremely computationally intractable. Nevertheless, we stress that in both theory and practice, stochastic gradient descent and its variants have been shown to efficiently find global minima of loss functions parameterized by neural networks (Li & Liang, 2018; Allen-Zhu et al., 2019; Du et al., 2019b). Furthermore, as we show in our proofs, the function \widehat{L}_t is Lipschitz continuous with a tolerable Lipschitz constant. Hence by Rademacher's Theorem, \widehat{L}_t is differentiable almost everywhere, and it is reasonable to minimize \widehat{L}_t via the stochastic gradient type methods that are popular for minimizing complex neural network losses. We also remark that it is fairly common in the RL literature to assume access to a computational oracle, when studying sample complexity (Agarwal et al., 2014; Du et al., 2019a; Agarwal et al., 2020a; Misra et al., 2020). We now formally state Theorem 1, our main result.

Theorem 1. Fix error tolerance $\epsilon > 0$ and failure proba-

bility tolerance $\delta > 0$. Given any $(\mathcal{M}, \pi(\Theta))$ satisfying the EPW condition, and sample size

$$n = \frac{4H^2|\mathcal{A}|^{2C+2}}{\epsilon^2} \left(\log \left(\frac{2H}{\delta} \right) + k \log \left(1 + \frac{32H|\mathcal{A}|^{C+1}C\phi B}{\epsilon} \right) \right),$$

Algorithm 1 outputs $\bar{\theta}$ satisfying

$$V_{\mathcal{M}}^{s_0}(\pi(\bar{\theta})) \geq V_{\mathcal{M}}^{s_0}(\pi(\theta^*)) - \epsilon$$

with probability at least $1 - \delta$.

A formal proof for the theorem is presented in Appendix B. A few comments are in order. Note that Algorithm 1 samples n trajectories at each timestep in its inner loop. Recall from the definition of our query model in Section 3.1, that we measure the total sample complexity by the total number of trajectories sampled. Hence, the total sample complexity of Algorithm 1 scales as $\mathcal{O}\left(\frac{H^3|\mathcal{A}|^{2C+2}kC}{\epsilon^2}\right)$, where we have discarded logarithmic factors. Observe that the total sample complexity depends only logarithmically on the failure probability tolerance δ , the bound B on the Euclidean norm of Θ and the Lipschitz constant ϕ of the policy. It is also worth noting that the sample complexity has *no* explicit dependence on the dimension d of states. However, it does depend linearly on the dimension k of the policy parameter space Θ , and of course in general k will scale with d . We note that in practical RL, one typically employs a shallow neural net, with only two or three layers. k is relatively small in this regime, in contrast to NLP or vision tasks where models are much larger.

Observe that the sample complexity bound in Theorem 1 has an exponential dependence on C . Recall that in our framework, as motivated at the end of Section 3.2, C is a constant, so our algorithm is indeed efficient. Nevertheless, we remark that as a direct corollary of the work of Du et al. (2020a), this exponential dependence on C cannot be improved by a better algorithm or sharper analysis.

More generally, in the context of existing literature, we provide some intuition for why the EPW condition permits efficient learning. A major issue that hinders sample efficient RL, in both theory and practice, is that an agent must plan over the entire horizon H . Roughly speaking, at each timestep the agent can choose any of $|\mathcal{A}|$ actions, so the total sample complexity required to plan over H timesteps scales as $\Omega(|\mathcal{A}|^H)$. This is a recurrent theme in the RL literature, and various prior works have shown that even when the MDP has some non-trivial structure, in the worst case such a scaling is unavoidable (Du et al., 2020a; Malik et al., 2021). Assuming that the MDP satisfies significant *linear*

structure, is one way to avoid this difficulty. By contrast, we avoid this difficulty while making *no* linearity assumptions. Instead, the EPW condition guarantees that an agent needs only to plan over C timesteps. Hence we exchange the worst case $\Omega(|\mathcal{A}|^H)$ scaling for the benign $\mathcal{O}(|\mathcal{A}|^{2C+2})$ scaling.

We believe that the EPW condition (or other conditions that are similar in spirit) is the *correct* condition for characterizing when sample efficient RL is possible, at least in RL domains like video games. By contrast, the linearity assumptions which prominently appear in prior literature, in addition to lacking clear empirical justification, are quite brittle. To demonstrate this, we leverage prior work to show the existence of Generic Games which have only slight nonlinearities, yet cannot be solved sample efficiently. Before we state this lower bound, we recall two standard definitions.

Definition 2 (Optimal Value Function). *The optimal value function $V_{\mathcal{M}}^* : \mathcal{S} \rightarrow \mathbb{R}$ of an MDP \mathcal{M} is defined as $V_{\mathcal{M}}^*(s) = V_{\mathcal{M}}^s(\pi^*)$, where π^* is an optimal policy of \mathcal{M} .*

Definition 3 (Softmax Linear Policy). *For an MDP \mathcal{M} , a softmax linear policy $\pi(\theta)$ is parameterized by $\theta \in \mathbb{R}^{|\mathcal{A}| \times d}$. Letting θ_i denote the i^{th} row of θ , the policy $\pi(\theta)$ satisfies*

$$\pi_s^{a_i}(\theta) = \frac{\exp(s^T \theta_i)}{\sum_{a_j \in \mathcal{A}} \exp(s^T \theta_j)}.$$

We now demonstrate the existence of Generic Games, which have only slight nonlinearities, where even approximating an optimal policy in a sample efficient manner is impossible. This result is heavily inspired by the recent work of Du et al. (2020a), and we claim no technical novelty. Rather, the purpose of this result in our setting, is to further motivate the importance of studying conditions such as EPW, instead of assuming that the MDP has linear structure.

Proposition 1. *There exists a Generic Game $(\mathcal{M}, \pi(\Theta))$, where d, k, ϕ, B are all at most polynomial in H and \mathcal{A} , and $\pi(\Theta)$ is the class of softmax linear policies, such that the following holds. There exists an unknown neural network $f : \mathbb{R}^d \rightarrow \mathbb{R}$, where f is a linear combination of two ReLU neurons, such that $V_{\mathcal{M}}^*(s) = f(s)$ for all $s \in \mathcal{S}$. Yet, any algorithm requires $\Omega(\min\{|\mathcal{A}|^H, 2^d\})$ trajectories to find, with probability at least $3/4$, a policy π satisfying*

$$V_{\mathcal{M}}^{s_0}(\pi) \geq V_{\mathcal{M}}^{s_0}(\pi(\theta^*)) - 1/4.$$

We stress that the essence of this result follows from Du et al. (2020a), and we only make small modifications to their proof to fit it in our Generic Game setting. We nevertheless provide a proof sketch in Appendix C. The result shows the existence of MDPs for which a softmax linear policy is optimal, and where the optimal value function can be expressed as a neural network with only two ReLU neurons. Despite only this slight nonlinearity, sample efficient RL is impossible. Notice that in the statement of this result, there is no dependency of π on θ . This is because we do not

restrict the algorithm to only search over policies lying in $\pi(\Theta)$. In particular, the policy π mentioned in the result can be *arbitrary*, and does not have to lie in $\pi(\Theta)$.

Proposition 1 demonstrates that if the Generic Game is even slightly nonlinear, as one would expect in practice, sample efficient RL is impossible. So we must look beyond linearity to obtain a realistic characterization of when sample efficient RL is possible. Our EPW condition, which makes no linearity assumptions, is one example of this.

5. Discussion

In this paper, we studied structural conditions which permit sample efficient RL in continuous state spaces, with a focus on conditions that are typical in popular RL domains such as Atari games. We introduced the EPW condition, which in contrast to prior work, makes no linearity assumptions about the MDP structure. We provided an algorithm which provably solves MDPs satisfying EPW. We analyzed the sample complexity of this algorithm, and showed it requires a number of trajectories that is a lower order polynomial of the horizon length and other relevant problem dependent quantities. We also showed that MDPs which have very slight nonlinearities (but do not satisfy EPW) cannot be solved sample efficiently. Our analysis thus highlights the important need to look beyond linear structure, in order to establish the sample efficiency of RL in popular domains.

A number of open questions remain. First, while our EPW condition is directly motivated by RL gaming domains such as Atari, it is unclear whether EPW is satisfied by other RL application domains such as robotics. A natural direction for future work is to study these domains more closely, and identify structure that permits sample efficient RL in such domains. Second, recall that our algorithm requires access to a particular computational oracle. As discussed, we made this computational abstraction since we placed minimal restrictions on the policy class, so in the worst case obtaining such an oracle could be intractable. Nevertheless, we suspect that when using a neural network policy class with an appropriate architecture, one could approximate this oracle efficiently. It would be interesting to precisely characterize when this is possible. Third, it would be interesting to see whether a variant of our theoretically justified algorithm can be deployed in practice. Using our theoretical insight to design a pragmatic method, with strong empirical performance, is an important direction for future work.

Acknowledgements

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE1745016. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the

views of the National Science Foundation.

References

- Agarwal, A., Hsu, D., Kale, S., Langford, J., Li, L., and Schapire, R. Taming the monster: A fast and simple algorithm for contextual bandits. In *Proceedings of the International Conference on Machine Learning*, 2014.
- Agarwal, A., Kakade, S., Krishnamurthy, A., and Sun, W. Flambe: Structural complexity and representation learning of low rank mdps. In *Advances in Neural Information Processing Systems*, 2020a.
- Agarwal, A., Kakade, S. M., Lee, J. D., and Mahajan, G. Optimality and approximation with policy gradient methods in markov decision processes. In *Proceedings of the Conference on Learning Theory*, 2020b.
- Allen-Zhu, Z. and Li, Y. What can resnet learn efficiently, going beyond kernels? In *Advances in Neural Information Processing Systems*, 2019.
- Allen-Zhu, Z., Li, Y., and Song, Z. A convergence theory for deep learning via over-parameterization. In *Proceedings of the International Conference on Machine Learning*, 2019.
- Badia, A. P., Piot, B., Kapturowski, S., Sprechmann, P., Vitvitskyi, A., Guo, Z. D., and Blundell, C. Agent57: Outperforming the Atari human benchmark. In *Proceedings of the International Conference on Machine Learning*, 2020.
- Bellemare, M. G., Naddaf, Y., Veness, J., and Bowling, M. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.
- Berner, C. et al. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arxiv:1912.06680*, 2019.
- Cobbe, K., Klimov, O., Hesse, C., Kim, T., and Schulman, J. Quantifying generalization in reinforcement learning. In *Proceedings of the International Conference on Machine Learning*, 2019.
- Dai, B., Shaw, A., Li, L., Xiao, L., He, N., Liu, Z., Chen, J., and Song, L. SBEED: Convergent reinforcement learning with nonlinear function approximation. In *Proceedings of the International Conference on Machine Learning*, 2018.
- Dean, S., Mania, H., Matni, N., Recht, B., and Tu, S. On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*, pp. 1–47, 2019.

- Dong, K., Luo, Y., Yu, T., Finn, C., and Ma, T. On the expressivity of neural networks for deep reinforcement learning. In *Proceedings of the International Conference on Machine Learning*, 2020a.
- Dong, K., Peng, J., Wang, Y., and Zhou, Y. Root-n-regret for learning in markov decision processes with function approximation and low bellman rank. In *Proceedings of the Conference on Learning Theory*, 2020b.
- Du, S., Krishnamurthy, A., Jiang, N., Agarwal, A., Dudik, M., and Langford, J. Provably efficient RL with rich observations via latent state decoding. In *Proceedings of the International Conference on Machine Learning*, 2019a.
- Du, S., Lee, J., Li, H., Wang, L., and Zhai, X. Gradient descent finds global minima of deep neural networks. In *Proceedings of the International Conference on Machine Learning*, 2019b.
- Du, S. S., Luo, Y., Wang, R., and Zhang, H. Provably efficient q-learning with function approximation via distribution shift error checking oracle. In *Advances in Neural Information Processing Systems*, 2019c.
- Du, S. S., Kakade, S. M., Wang, R., and Yang, L. F. Is a good representation sufficient for sample efficient reinforcement learning? In *International Conference on Learning Representations*, 2020a.
- Du, S. S., Lee, J. D., Mahajan, G., and Wang, R. Agnostic q-learning with function approximation in deterministic systems: Near-optimal bounds on approximation error and sample complexity. In *Advances in Neural Information Processing Systems*, 2020b.
- Fazel, M., Ge, R., Kakade, S., and Mesbahi, M. Global convergence of policy gradient methods for the linear quadratic regulator. In *Proceedings of the International Conference on Machine Learning*, 2018.
- Fazlyab, M., Robey, A., Hassani, H., Morari, M., and Pappas, G. Efficient and accurate estimation of lipschitz constants for deep neural networks. In *Advances in Neural Information Processing Systems*, 2019.
- Hadfield-Menell, D., Milli, S., Abbeel, P., Russell, S. J., and Dragan, A. Inverse reward design. In *Advances in Neural Information Processing Systems*, 2017.
- Jiang, N., Krishnamurthy, A., Agarwal, A., Langford, J., and Schapire, R. E. Contextual decision processes with low Bellman rank are PAC-learnable. In *Proceedings of the International Conference on Machine Learning*, 2017.
- Jin, C., Yang, Z., Wang, Z., and Jordan, M. I. Provably efficient reinforcement learning with linear function approximation. In *Proceedings of the Conference on Learning Theory*, 2020.
- Jin, C., Liu, Q., and Miryoosefi, S. Bellman eluder dimension: New rich classes of rl problems, and sample-efficient algorithms, 2021.
- Lattimore, T., Szepesvari, C., and Weisz, G. Learning with good feature representations in bandits and in RL with a generative model. In *Proceedings of the International Conference on Machine Learning*, 2020.
- Levine, S., Finn, C., Darrell, T., and Abbeel, P. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research*, 17:1334–1373, 2016.
- Li, Y. and Liang, Y. Learning overparameterized neural networks via stochastic gradient descent on structured data. In *Advances in Neural Information Processing Systems*, 2018.
- Li, Y., Ma, T., and Zhang, H. R. Learning over-parametrized two-layer neural networks beyond ntk. In *Proceedings of the Conference on Learning Theory*, 2020.
- Malik, D., Pananjady, A., Bhatia, K., Khamaru, K., Bartlett, P. L., and Wainwright, M. J. Derivative-free methods for policy optimization: Guarantees for linear quadratic systems. *Journal of Machine Learning Research*, 21: 1–51, 2020.
- Malik, D., Li, Y., and Ravikumar, P. When is generalizable reinforcement learning tractable? *arXiv preprint arxiv:2101.00300*, 2021.
- Misra, D., Henaff, M., Krishnamurthy, A., and Langford, J. Kinematic state abstraction and provably efficient rich-observation reinforcement learning. In *Proceedings of the International Conference on Machine Learning*, 2020.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. Playing atari with deep reinforcement learning. In *NIPS Deep Learning Workshop*. 2013.
- Mnih, V. et al. Human-level control through deep reinforcement learning. *Nature*, 518:529–533, 2015.
- Pananjady, A. and Wainwright, M. J. Instance-dependent ℓ_∞ -bounds for policy evaluation in tabular reinforcement learning. *IEEE Transactions on Information Theory*, 67: 566–585, 2021.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. Trust region policy optimization. In *Proceedings of the International Conference on Machine Learning*, 2015.

- Simchowitz, M. and Jamieson, K. G. Non-asymptotic gap-dependent regret bounds for tabular mdps. In *Advances in Neural Information Processing Systems*. 2019.
- Van Roy, B. and Dong, S. Comments on the dukakade-wang-yang lower bounds. *arXiv preprint arxiv:1911.07910*, 2019.
- Vershynin, R. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- Wang, R., Salakhutdinov, R. R., and Yang, L. Reinforcement learning with general value function approximation: Provably efficient approach via bounded eluder dimension. In *Advances in Neural Information Processing Systems*, 2020.
- Wang, Y., Wang, R., Du, S. S., and Krishnamurthy, A. Optimism in reinforcement learning with generalized linear function approximation. In *International Conference on Learning Representations*, 2021.
- Weisz, G., Amortila, P., Janzer, B., Abbasi-Yadkori, Y., Jiang, N., and Szepesvari, C. On query-efficient planning in mdps under linear realizability of the optimal state-value function. In *Proceedings of the Conference on Learning Theory*, 2021a.
- Weisz, G., Amortila, P., and Szepesvari, C. Exponential lower bounds for planning in mdps with linearly-realizable optimal action-value functions. In *Proceedings of the International Conference on Algorithmic Learning Theory*, 2021b.
- Wen, Z. and Van Roy, B. Efficient exploration and value function generalization in deterministic systems. In *Advances in Neural Information Processing Systems*, 2013.
- Yang, L. and Wang, M. Sample-optimal parametric q-learning using linearly additive features. In *Proceedings of the International Conference on Machine Learning*, 2019.
- Yang, L. and Wang, M. Reinforcement learning in feature space: Matrix bandit, kernels, and regret bound. In *Proceedings of the International Conference on Machine Learning*, 2020.