

# Analysis of Q-learning with Adaptation and Momentum Restart for Gradient Descent

Bowen Weng<sup>\*1</sup>, Huaqing Xiong<sup>\*1</sup>, Yingbin Liang<sup>1</sup> and Wei Zhang<sup>†2</sup>

<sup>1</sup>Electrical and Computer Engineering, The Ohio State University, Columbus, OH, USA.

<sup>2</sup>Mechanical and Energy Engineering, Southern University of Science and Technology, China.

{weng.172, xiong.309, liang.889}@osu.edu, zhangw3@sustech.edu.cn

## Abstract

Existing convergence analyses of Q-learning mostly focus on the vanilla stochastic gradient descent (SGD) type of updates. Despite the Adaptive Moment Estimation (Adam) has been commonly used for practical Q-learning algorithms, there has not been any convergence guarantee provided for Q-learning with such type of updates. In this paper, we first characterize the convergence rate for Q-AMSGrad, which is the Q-learning algorithm with AMSGad update (a commonly adopted alternative of Adam for theoretical analysis). To further improve the performance, we propose to incorporate the momentum restart scheme to Q-AMSGrad, resulting in the so-called Q-AMSGradR algorithm. The convergence rate of Q-AMSGradR is also established. Our experiments on a linear quadratic regulator problem show that the two proposed Q-learning algorithms outperform the vanilla Q-learning with SGD updates. The two algorithms also exhibit significantly better performance than the DQN learning method over a batch of Atari 2600 games.

## 1 Introduction

Q-learning [Watkins and Dayan, 1992], as one of the most important model-free reinforcement learning (RL) algorithms, has received considerable attention in recent years [Bertsekas and Tsitsiklis, 1996; Even-Dar and Mansour, 2003; Lu *et al.*, 2018]. The vanilla Q-learning algorithm runs a step of empirical Bellman operator update of the Q-function and a step of stochastic gradient descent (SGD) in an *alternating* fashion. The convergence guarantee of Q-learning has been studied for the tabular case in [Bertsekas, 1995; Even-Dar and Mansour, 2003], for the case with linear function approximation in [Bertsekas and Tsitsiklis, 1996; Zou *et al.*, 2019; Chen *et al.*, 2019b], and also for neural network parameterization in [Xu and Gu, 2019].

However, all the existing theoretical analyses focus on Q-learning algorithms that take simple SGD iterations. Such theory is not applicable to practical Q-learning algorithms that

implement the Adaptive Moment Estimation (Adam) type of updates. In this paper, we study the Q-learning algorithm with Adam-type updates in terms of its theoretical convergence and the performance in benchmark experiments. It is known in optimization that Adam does not always converge, and instead, and a slightly modified variant AMSGrad proposed in [Reddi *et al.*, 2018] has been widely accepted as an alternative to justify the theoretical performance of Adam-type methods. This motivates the first question that we aim to address.

- *Q1: Can we provide the convergence guarantee for Q-learning under AMSGrad updates (i.e., Q-AMSGrad)?*

In conventional optimization problems, *restart* has been incorporated into the gradient descent algorithm with momentum as a simple yet effective scheme to facilitate the acceleration performance [O’donoghue and Candes, 2015]. Hence, it is natural to incorporate such a restart technique into Q-learning and ask the following question about its performance.

- *Q2: Does Q-AMSGrad with momentum restart (i.e., Q-AMSGradR) still converge?*

As aforementioned, both Q-AMSGrad and Q-AMSGradR update alternatingly between one step of Bellman operator update of the Q-function and one step of adaptive momentum update. This is in contrast to the well-known deep Q-Network (DQN) learning [Mnih *et al.*, 2015], which runs in a nested-loop manner with the outer loop consisting of an one-step update of the Q-function and the inner loop consisting of many iterations of supervised learning to fit a *target Q-function*. It is conventionally known that taking *just one* gradient step toward the target Q-function results in the instability of original Q-learning update. Thus the simple alternating update manner mainly aroused interest in theory instead of practice. We are interested in whether the adaptive momentum updates can improve the stability of Q-learning without using the supervised learning process for target network fitting. Therefore, the third question we want to address is to compare these two types of Q-learning algorithms from the experimental perspective.

- *Q3: Do Q-AMSGrad and Q-AMSGradR perform competitively or even better than DQN in experiment?*

This paper addresses the above theoretical and experimental questions with affirmative answers.

## 1.1 Main Contributions

**Theoretically**, we show that with linear function approximation (which is almost the only structure that the current tools

<sup>\*</sup>Equal contribution

<sup>†</sup>Corresponding author

for analysis of Q-learning can handle), both Q-AMSGrad and Q-AMSGradR converge to the global optimal solution under standard assumptions for Q-learning. To the best of our knowledge, this is the first non-asymptotic convergence guarantee on Q-learning that incorporates Adam-type update. Furthermore, a slight adaptation of our proof provides the convergence rate for the AMSGard for conventional strongly convex optimization which has not been studied before and can be of independent interest.

**Experimentally**, we demonstrate that the practical versions of Q-AMSGrad and Q-AMSGradR (referred to as Q-Adam and Q-AdamR) exhibit appealing experimental performance. In a batch of 23 Atari 2600 games, our experiments show that both Q-Adam and Q-AdamR outperform DQN by 50% on average. Furthermore, Q-AdamR effectively reduces the performance variance and achieves a much more stable learning process. In our experiments for the linear quadratic regulator (LQR) problems, Q-AdamR converges even faster than the model-based value iteration (VI) solution. This is a rather surprising result given that the model-based VI has been treated as the performance upper bound for the Q-learning (including DQN) algorithms with target update [Lewis and Vrabie, 2009; Yang *et al.*, 2019].

## 1.2 Related Work

We briefly review the related work as follows.

**Theoretical analysis of Q-learning:** Since proposed in [Watkins and Dayan, 1992], the convergence of Q-learning has been extensively studied, particularly for the case with linear function approximation such as [Bertsekas and Tsitsiklis, 1996; Zou *et al.*, 2019; Chen *et al.*, 2019b; Du *et al.*, 2019], to name a few, and more recently for the case with neural networks in [Xu and Gu, 2019], where the analysis exploits the approximate linear structure of neural networks in the overparamterized regime. All these existing analysis of Q-learning considers the vanilla SGD update, whereas our study is the first to analyze the more involved case with Adam-type updates.

**Convergence analysis of Adam-type algorithms in conventional optimization:** Adam was proposed in [Kingma and Ba, 2014] for speeding up the training of deep neural networks, and the regret bounds were characterized for Adam/AMSGrad in [Kingma and Ba, 2014; Reddi *et al.*, 2018; Tran and others, 2019] for online convex optimization. Recently, convergence analysis of Adam/AMSGrad was provided for nonconvex optimization in [Zou *et al.*, 2018; Zhou *et al.*, 2018; Chen *et al.*, 2019a] and policy gradient [Xiong *et al.*, 2020], in which such Adam-type algorithms were guaranteed to converge to a stationary point. To the best of our knowledge, our study provides the first convergence analysis of the Adam-type algorithms for Q-learning.

**Empirical performance of Q-learning:** DQN learning and its improved variants of dueling network structure [Wang *et al.*, 2016], double Q-learning [Van Hasselt *et al.*, 2016] and variance exploration and sampling schemes [Schaul *et al.*, 2015] have achieved significant success due to their superb performance in practice. In contrast to such nested-loop algorithms (which involves the fitting of a target Q

function periodically), the Q-learning algorithms that strictly follow the alternating updates are much less explored in practice. [Mnih *et al.*, 2016] proposed the asynchronous alternating Q-learning with competitive performance against DQN. However, the algorithm still relies on a slowly moving target network similar to DQN, and the multi-thread learning also complicates the computational setup. [Lu *et al.*, 2018] studied the problem of value overestimation and proposed the non-delusional Q-learning algorithm that employs the so-called pre-conditioned Q-networks, which is also computationally complex. [Knight and Lerner, 2018] proposed a natural gradient propagation to improve the performance, where the gradient implementation is complex. Our experiments in this paper demonstrate that simple alternating Q-learning algorithms Q-AMSGrad and Q-AMSGradR without the complex designs as in [Mnih *et al.*, 2016; Lu *et al.*, 2018; Knight and Lerner, 2018] have competitive and sometimes better performance than DQN.

**Notations** We use  $\|x\| := \|x\|_2$  to denote the  $\ell_2$  norm of a vector  $x$ , and use  $\|x\|_\infty$  to denote the infinity norm. When  $x, y$  are both vectors,  $x/y, xy, x^2, \sqrt{x}$  are all calculated in the element-wise manner, which will be used in the update of Adam and AMSGard. We denote  $[n] = 1, 2, \dots, n$ , and  $\lfloor x \rfloor \in \mathbb{Z}$  as the integer such that  $\lfloor x \rfloor \leq x < \lfloor x \rfloor + 1$ .

## 2 Preliminaries

We consider a Markov decision process with a considerably large or continuous state space  $\mathcal{S} \subset \mathbb{R}^M$  and action space  $\mathcal{A} \subset \mathbb{R}^N$ , a non-negative bounded reward function  $R : \mathcal{S} \times \mathcal{A} \rightarrow [0, R_{\max}]$ , and a transition kernel  $P(s'|s, a)$  that indicates the probability from a state-action pair  $(s, a)$  to a state  $s'$ . We define  $U(s) \subset \mathcal{A}$  as the admissible set of actions at state  $s$ , and  $\pi : \mathcal{S} \rightarrow \mathcal{A}$  as a feasible stationary policy. We seek to solve a discrete-time sequential decision problem as follows:

$$\begin{aligned} & \underset{\pi}{\text{maximize}} \quad J_{\pi}(s_0) = \mathbb{E}_P \left[ \sum_{t=0}^{\infty} \gamma^t R(s_t, \pi(s_t)) \right], \\ & \text{subject to } s_{t+1} \sim P(\cdot|s_t, a_t), \end{aligned} \quad (1)$$

where  $\gamma \in (0, 1)$  is the discount factor. Let  $J^*(s) := J_{\pi^*}(s)$  be the optimal value function when applying the optimal policy  $\pi^*$ . The corresponding optimal Q-function can be defined as

$$Q^*(s, a) := R(s, a) + \gamma \mathbb{E}_P J^*(s'), \quad (2)$$

where  $s' \sim P(\cdot|s, a)$  and we use the same notation hereafter when no confusion arises. In other words,  $Q^*(s, a)$  represents the reward of an agent who starts from state  $s$  and takes action  $a$  at the first step and then follows the optimal policy  $\pi^*$  thereafter.

### 2.1 Q-learning Algorithm

This paper focuses on the Q-learning algorithm that uses a parametric function  $\hat{Q}(s, a; \theta)$  to approximate the Q-function with a parameter  $\theta$  having finite and relatively small dimen-

sions. The update rule of Q-learning is given by

$$T\hat{Q}(s, a; \theta_t) = R(s, a) + \gamma \max_{a' \in U(s')} \hat{Q}(s', a'; \theta_t); \quad (3)$$

$$\theta_{t+1} = \theta_t - \alpha_t \left( \hat{Q}_t(s, a; \theta_t) - T\hat{Q}(s, a; \theta_t) \right) \hat{g}_t, \quad (4)$$

$$\hat{g}_t := \hat{g}(\theta_t; s, a) = \frac{\partial}{\partial \theta_t} \hat{Q}_t(s, a; \theta_t), \quad (5)$$

where  $\alpha_t$  is the step size at time  $t$ . It is clear that Q-learning performs the update by taking one step of temporal target update and one step of parameter learning in an alternating fashion.

## 2.2 Linear Function Approximation

Like most of the related work, we focus on the convergence analysis under the linear function approximation. A linear approximation of the Q-function  $\hat{Q}(s, a; \theta)$  can be written as

$$\hat{Q}(s, a; \theta) = \phi(s, a)^T \theta, \quad (6)$$

where  $\theta \in \mathbb{R}^d$ , and  $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$  is a vector function of size  $d$ , and the elements of  $\Phi$  represent the nonlinear kernel (feature) functions.

**Remark 1.** We note that recent work [Xu and Gu, 2019] established the convergence rate of Q-learning with neural network approximation, which exploits the approximate linear structure of the neural network in the overparameterized regime. Thus, our analysis under the linear function approximation can be generalized to the function class of overparameterized neural networks by applying the techniques developed in recent work [Xu and Gu, 2019].

## 3 Convergence Analysis of Q-AMSGrad

In this section, we characterize the convergence guarantee for Q-learning under Adam-type updates.

### 3.1 Q-AMSGrad Algorithm

Although Adam has obtained great success as an optimizer in deep learning, it is well known that Adam by nature is non-convergent even for simple convex loss functions [Reddi *et al.*, 2018]. Instead, a slightly modified version called AMSGrad [Reddi *et al.*, 2018] is widely used to study the convergence property of Adam-type algorithms in conventional optimization. Here, we apply the update rule of AMSGrad to the Q-learning algorithm and refer to such an algorithm as Q-AMSGrad. Algorithm 1 describes Q-AMSGrad in detail.

More specifically, the iterations of Q-AMSGrad evolve by updating the exponentially decaying average of historical gradients ( $m_t$ ) and squared historical gradients ( $v_t$ ). The hyperparameters  $\beta_1, \beta_2$  are used to exponentially decrease the rate of the moving averages. The difference between AMSGrad and Adam lies in the fact that AMSGrad makes the sequence  $\hat{v}_{t,i}$  increasing along the time step  $t$  for each entry  $i \in [d]$ , whereas Adam does not guarantee such a property.

### 3.2 Convergence Result

Before stating the main theorem, we first introduce some technical assumptions and lemmas for our analysis.

---

#### Algorithm 1 Q-AMSGrad

---

```

1: Input:  $\alpha, \lambda, \theta_1, \beta_1, \beta_2, m_0 = 0, \hat{v}_0 = 0$ .
2: for  $t = 1, 2, \dots, T$  do
3:    $\alpha_t = \frac{\alpha}{\sqrt{t}}, \beta_{1t} = \beta_1 \lambda^t$ 
4:   Observe data  $(s_t, a_t, s_{t+1})$  from policy  $\pi$  and transition probability  $P$ 
5:    $b_t = R(s_t, a_t) + \gamma \max_{a'} \phi^T(s_{t+1}, a') \theta_t$ 
6:    $g_t = (\phi^T(s_t, a_t) \theta_t - b_t) \phi(s_t, a_t)$ 
7:    $m_t = (1 - \beta_{1t})m_{t-1} + \beta_{1t}g_t$ 
8:    $v_t = (1 - \beta_2)\hat{v}_{t-1} + \beta_2 g_t^2$ 
9:    $\hat{v}_t = \max(\hat{v}_{t-1}, v_t), \hat{V}_t = \text{diag}(\hat{v}_1, \dots, \hat{v}_d)$ 
10:   $\theta_{t+1} = \Pi_{\mathcal{D}, \hat{V}_t^{1/4}}(\theta_t - \alpha_t \hat{V}_t^{-\frac{1}{2}} m_t)$ 
    where  $\Pi_{\mathcal{D}, \hat{V}_t^{1/4}}(\theta') = \min_{\theta \in \mathcal{D}} \|\hat{V}_t^{1/4}(\theta' - \theta)\|$ .
11: end for
12: Output:  $\frac{1}{T} \sum_{t=1}^T \theta_t$ 

```

---

**Assumption 1.** For any state-action pair  $(s, a) \in \mathcal{S} \times \mathcal{A}$ , the kernel function  $\phi$  is uniformly bounded and we have

$$\|\phi(s, a)\| \leq 1, \quad \forall (s, a). \quad (7)$$

This assumption is mild since we can normalize the kernel function if the kernel function is uniformly bounded. It is widely applied in the literature to simplify the analysis of RL algorithms with linear function approximation [Bhandari *et al.*, 2018; Chen *et al.*, 2019b].

**Assumption 2.** [Chen *et al.*, 2019b, Lemma 6.7] At each iteration  $t$ , the noisy gradient is unbiased, i.e.  $g_t = \bar{g}_t + \xi_t$  with  $\mathbb{E}\xi_t = 0$  where  $\bar{g}_t = \mathbb{E}[g_t]$ . The equation  $\bar{g}(\theta) = 0$  has a unique solution  $\theta^*$ , and there exists a  $c > 0$ , such that for any  $\theta \in \mathbb{R}^d$  we have

$$(\theta - \theta^*)^T \bar{g}(\theta) \geq c \|\theta - \theta^*\|^2. \quad (8)$$

Assumption 2 has been proved as a key technical lemma in [Chen *et al.*, 2019b] under certain assumptions, which appears to be the weakest among the existing studies for establishing the convergence guarantee for Q-learning with linear function approximation. It is the standard assumption in the related literature to analyze the convergence Q-learning with linear function approximation [Zou *et al.*, 2019; Chen *et al.*, 2019b; Xu and Gu, 2019].

**Assumption 3.** The domain  $\mathcal{D} \subset \mathbb{R}^d$  of approximation parameters is a ball originating at  $\theta = 0$  with bounded diameter containing  $\theta^*$ . That is, there exists  $D_\infty$ , such that  $\|\theta_m - \theta_n\| < D_\infty, \forall \theta_m, \theta_n \in \mathcal{D}$ , and  $\theta^* \in \mathcal{D}$ .

This Assumption can be easily satisfied when we apply a projected algorithm, and is standard in the theoretical analysis of Adam-type algorithms [Chen *et al.*, 2019a; Zhou *et al.*, 2018].

Based on the above assumptions, we can immediately obtain the bounded property of the gradient, which is stated in the following lemma.

**Lemma 1.** Under Assumptions 1 and 3, at each iteration  $t$ , the gradient estimator  $g_t$  in Q-AMSGrad is uniformly bounded.

That is,

$$\|g_t\|_\infty \leq \|g_t\| \leq R_{\max} + (1 + \gamma)D_\infty, \quad \forall t. \quad (9)$$

In addition, we denote  $G_\infty = R_{\max} + (1 + \gamma)D_\infty$  and let  $\{m_t, \hat{v}_t\}$  for  $t = 1, 2, \dots$  be sequences generated by Algorithm 2. Then we have

$$\|\bar{g}_t\| \leq G_\infty, \|m_t\| \leq G_\infty, \|\hat{v}_t\| \leq G_\infty^2.$$

We next provide the non-asymptotic convergence of Q-AMSGrad in the following theorem.

**Theorem 1.** (Convergence of Q-AMSGrad) Suppose  $\alpha_t = \frac{\alpha}{\sqrt{t}}$ ,  $\beta_{1t} = \beta_1 \lambda^t$  and  $\delta = \beta_1/\beta_2$  with  $\delta, \lambda \in (0, 1)$  for  $t = 1, 2, \dots$  in Algorithm 1. Given Assumptions 1 ~ 3, the output of Q-AMSGrad satisfies:

$$\begin{aligned} & \mathbb{E} \|\theta_{out} - \theta^*\| \\ & \leq \frac{B_1}{T} + \frac{B_2}{\sqrt{T}} + \frac{B_3 \sqrt{1 + \log T}}{T} \sum_{i=1}^d \mathbb{E} \|g_{1:T,i}\|, \end{aligned} \quad (10)$$

where  $B_1 = \frac{G_\infty D_\infty^2}{2\alpha c(1-\beta_1)} + \frac{\beta_1 G_\infty D_\infty^2}{2\alpha c(1-\beta_1)(1-\lambda)^2} + \|\theta_1 - \theta^*\|^2$ ,  $B_2 = \frac{dG_\infty D_\infty^2}{2\alpha c(1-\beta_1)}$ , and  $B_3 = \frac{\alpha(1+\beta_1)}{2c(1-\beta_1)^2(1-\delta)\sqrt{1-\beta_2}}$ .

In Theorem 1,  $B_1, B_2, B_3$  in the bound in Equation (10) are constants and independent of time. Therefore, under the choice of the stepsize and hyper-parameters in Algorithm 1, Q-AMSGrad achieves a convergence rate of  $\mathcal{O}\left(\frac{1}{\sqrt{T}}\right)$  when  $\sum_{i=1}^d \|g_{1:T,i}\| << \sqrt{T}$  [Reddi et al., 2018].

**Remark 2.** Our proof of convergence here has two major differences from that for AMSGrad in [Reddi et al., 2018] in conventional optimization: (a) The two algorithms are quite different. Q-AMSGrad is a Q-learning algorithm alternatively finding the best policy with a moving target, whereas AMSGrad is an optimizer for conventional optimization and does not have alternating nature. (b) Our analysis is on the convergence rate whereas [Reddi et al., 2018] provides regret bound. In fact, a slight modification of our proof also provides the convergence rate of AMSGrad for conventional strongly convex optimization, which can be of independent interest. Moreover, our proof avoids the theoretical error in the proof in [Reddi et al., 2018] as pointed out by [Tran and others, 2019].

## 4 Convergence Analysis of Q-AMSGradR

In this section, we propose to incorporate a momentum restart technique to Q-AMSGrad in order to improve its performance. We first introduce the algorithm and then provide the convergence analysis for such an algorithm. We demonstrate its desired experimental performance in Section 5.

### 4.1 Q-AMSGrad Algorithm with Momentum Restart

We introduce the restart technique to Q-AMSGrad and propose Q-AMSGradR as shown in Algorithm 2. Q-AMSGradR applies the same update rule as Algorithm 1, but periodically resets  $m_t, \hat{v}_t$  with a period of  $r$ , i.e.,  $m_t = 0, \hat{v}_t = 0, \forall t = kr, k = 1, 2, \dots$ . We explain such an idea further as follows.

---

### Algorithm 2 Q-AMSGradR

---

```

1: Input:  $\alpha, \lambda, \theta_1, \beta_1, \beta_2, m_0 = 0, \hat{v}_0 = 0$ .
2: for  $t = 1, 2, \dots, T$  do
3:   if  $\text{mod}(t, r) = 0$  then
4:      $m_t = 0, \hat{v}_t = 0$ 
5:   end if
6:    $\alpha_t = \frac{\alpha}{\sqrt{t}}, \beta_{1t} = \beta_1 \lambda^t$ 
7:   Observe data  $(s_t, a_t, s_{t+1})$  from policy  $\pi$  and transition probability  $P$ 
8:    $b_t = R(s_t, a_t) + \gamma \max_{a'} \phi^T(s_{t+1}, a') \theta_t$ 
9:    $g_t = (\phi^T(s_t, a_t) \theta_t - b_t) \phi(s_t, a_t)$ 
10:   $m_t = (1 - \beta_{1t})m_{t-1} + \beta_{1t}g_t$ 
11:   $v_t = (1 - \beta_2)\hat{v}_{t-1} + \beta_2 g_t^2$ 
12:   $\hat{v}_t = \max(\hat{v}_{t-1}, v_t), \hat{V}_t = \text{diag}(\hat{v}_1, \dots, \hat{v}_d)$ 
13:   $\theta_{t+1} = \Pi_{\mathcal{D}, \hat{V}_t^{1/4}} \left( \theta_t - \alpha_t \hat{V}_t^{-\frac{1}{2}} m_t \right)$ 
14: end for
15: Output:  $\frac{1}{T} \sum_{t=1}^T \theta_t$ 

```

---

Traditional momentum-based algorithms largely depend on the historical gradient direction. When part of the historical information is incorrect, the estimation error tends to accumulate. The restart technique can be employed to deal with such an issue. One way to restart the momentum-based methods is to initialize the momentum at some restart iteration. That is, at restart iteration  $r$ , we reset  $m_r, v_r$ , i.e.,  $m_r = 0, v_r = 0$ , which yields  $\theta_{r+1} = \theta_r$ . It is an intuitive implementation technique to adjust the trajectory from time to time, and can usually help mitigate the aforementioned problem while keeping fast convergence property. For the implementation, we execute the restart periodically with a period  $r$ . It turns out that the restart technique can significantly improve the numerical performance, which can be seen in Section 5.

## 4.2 Convergence Result

In the following theorem, we provide the non-asymptotic convergence for Q-AMSGradR.

**Theorem 2.** (Convergence of Q-AMSGradR) Suppose  $\alpha_t = \frac{\alpha}{\sqrt{t}}$ ,  $\beta_{1t} = \beta_1 \lambda^t$  and  $\delta = \beta_1/\beta_2$  with  $\delta, \lambda \in (0, 1)$  for  $t = 1, 2, \dots$  in Algorithm 2. Given Assumptions 1 ~ 3, the output of Q-AMSGradR satisfies:

$$\begin{aligned} & \mathbb{E} \|\theta_{out} - \theta^*\| \\ & \leq \frac{B_1}{T} + \frac{B_2 \sqrt{1 + \log T}}{T} \sum_{i=1}^d \mathbb{E} \|g_{1:T,i}\| \\ & \quad + \frac{B_3}{T} \left( \sqrt{T} + \sum_{k=1}^{\lfloor T/r \rfloor} \sqrt{kr - 1} \right) \\ & \quad + \frac{1}{T} \sum_{k=0}^{\lfloor T/r \rfloor} \left( \frac{G_\infty D_\infty^2}{\alpha} \sqrt{kr + 2} + B_4 \mathbb{E} \|\theta_{kr} - \theta^*\|^2 \right), \end{aligned} \quad (11)$$

where  $B_1 = \frac{\beta_1 D_\infty^2 G_\infty}{2\alpha c(1-\beta_1)(1-\lambda)^2}$ ,  $B_2 = \frac{\alpha(1+\beta_1)}{2c(1-\beta_1)^2(1-\delta)\sqrt{1-\beta_2}}$ ,  $B_3 = \frac{dG_\infty D_\infty^2}{2\alpha c(1-\beta_1)}$ , and  $B_4 = 4c(1 - \beta_1)$ .

Theorem 2 indicates that for Q-AMSGradR to enjoy a convergence rate of  $\mathcal{O}\left(\frac{1}{\sqrt{T}}\right)$ , the restart period  $r$  needs to be sufficiently large and  $\sum_{i=1}^d \|g_{1:T,i}\| << \sqrt{T}$ . In practice as demonstrated by the experiments in Section 5, Q-AMSGradR typically performs well, not necessarily under the theoretical conditions.

## 5 Experimental Performance

In this section, we empirically evaluate the Q-learning algorithms studied in this paper. We first study the linear quadratic regulator (LQR) problem, which serves as a direct numerical demonstration of the convergence analysis under linear function approximation. We then use the Atari 2600 games [Brockman *et al.*, 2016], a classic benchmark for DQN evaluations, to demonstrate the effectiveness of the Q-learning algorithms for complicated tasks. Note that our experiments test the performance of Q-Adam and Q-AdamR, which serve as practical versions of Q-AMSGrad and Q-AMSGradR by adopting Adam for Q-learning and keeping all major properties of Q-AMSGrad and Q-AMSGradR including the alternating update between Q-function update and parameter fitting.

Our main focus here lies in: (a) comparison between vanilla Q-Adam and that with momentum restart (Q-AdamR), (b) comparison between Q-Adam/Q-AdamR and vanilla Q-learning (through LQR), and (c) comparison between Q-Adam/Q-AdamR and DQN (through Atari 2600 games). Both of our experiments show that Q-AdamR outperforms Q-Adam, vanilla Q-learning and DQN in terms of convergence speed and variance reduction. Compared with DQN in the empirical experiments of Atari games, under the same hyper-parameter settings, Q-Adam and Q-AdamR improve the performance of DQN by 50% on average.

### 5.1 DQN Algorithm

As DQN is also included in this work for performance comparison. We recall the update of DQN in the following as reference. Differently from the vanilla Q-learning, DQN updates the parameters in a nested loop. Within the  $t$ -th inner loop, DQN first obtains the target Q-function as in Equation (12), and then uses a neural network to fit the target Q-function by running  $Y$  steps of a certain optimization algorithm as Equation (13). The update rule of DQN is given as follows.

$$T\hat{Q}(s, a; \theta_t^0) = R(s, a) + \gamma \max_{a' \in U(s')} \hat{Q}(s', a'; \theta_t^0), \quad (12)$$

$$\theta_t^Y = \text{Optimizer}(\theta_t^0, T\hat{Q}(s, a; \theta_t^0)), \quad (13)$$

where *Optimizer* can be SGD or Adam for example, and Equation (13) is thus a supervised learning process with  $T\hat{Q}(s, a; \theta_t^0)$  as the "supervisor". At the  $t$ -th outer loop, DQN performs the so-called target update as

$$\theta_{t+1}^0 = (1 - \tau)\theta_t^0 + \tau\theta_t^Y. \quad (14)$$

In practice, when one of the momentum-based optimizers is adopted for Equation (13), such as Adam, it is only initialized once at the beginning of the first inner loop. The historical gradient terms then accumulate throughout multiple inner loops with different targets.

Table 1: Hyper-parameters for LQR experiments

Step size	$\tilde{\tau}$	Adam $\beta_1$	Adam $\beta_2$
0.0001	0.01	0.9	0.999
Restart period $r$		Stop criterion	$\gamma$
100		$\ K_i - K^*\ _2 \leq 10^{-4}$	1

Generally speaking, the difference between Q-Adam/Q-AdamR and DQN mainly lies in Q-Adam/Q-AdamR takes one-step Q-function update and one-step model parameter fitting alternatively, whereas DQN takes one-step Q-function update followed by a sufficient large number of steps for model parameter fitting (towards a target Q-function).

### 5.2 Linear Quadratic Regulator

We numerically validate the performance of Q-Adam and Q-AdamR through an infinite-horizon discrete-time LQR problem. A typical model-based solution (with known dynamics), known as the discrete-time algebraic Riccati equation (DARE), is adopted to derive the optimal policy  $u_t^* = -K^*x_t$ . The performance of the learning algorithm is then evaluated at each step of iterate  $t$  with the Euclidean norm  $\|K_t - K^*\|$ . Given the problem nature of LQR, we also re-scale the loss term of  $L(\theta_t) := \hat{Q}_t(s, a; \theta_t) - T\hat{Q}(s, a; \theta_t)$  in Equation (4) as  $\tilde{L}(\theta_t) = \tilde{\tau}^2 L(\theta_t)$  with some scaling factor  $\tilde{\tau} \in (0, 1]$ , which is beneficial for stabilizing the learning process. The performance result for each method is averaged over 10 trials with different random seeds. All algorithms share the same set of random seeds and are initialized with the same  $\theta_0$ . The hyper-parameters of the learning settings are also consistent and further details are shown in Table 1. Note that for all the implementations, we also adopt the double Q-update [Van Hasselt *et al.*, 2016] to help prevent over-estimations of the Q-value. The performance results are provided in Figure 1. Here we highlight the main observations from the LQR experiments.

- **Q-AdamR outperforms DARE.** In ideal cases where data sampling perfectly emulates the system dynamics and the target is accurately learned in each inner loop, DARE for LQR would become equivalent to the DQN-like update if the neural network is replaced with a parameterized linear function. In practice, such ideal conditions are difficult to satisfy, and hence the actual Q-learning with target update is usually far slower (in terms of the number of steps of target updates) than DARE. Note that Q-AdamR performs significantly well and even converges faster than DARE, and thus implies it is faster than the most well-performing Q-learning with target update.

- **Q-AdamR outperforms Q-Adam.** Overall, under the same batch sampling scheme and restart period, Q-AdamR achieves a faster convergence and smaller variance than Q-Adam.

### 5.3 Atari Games

We apply the Q-Adam and Q-AdamR algorithms to the more challenging tasks of deep convolutional neural network playing a group of Atari 2600 games. The particular DQN we train to compare against adopts the dueling network structure [Wang

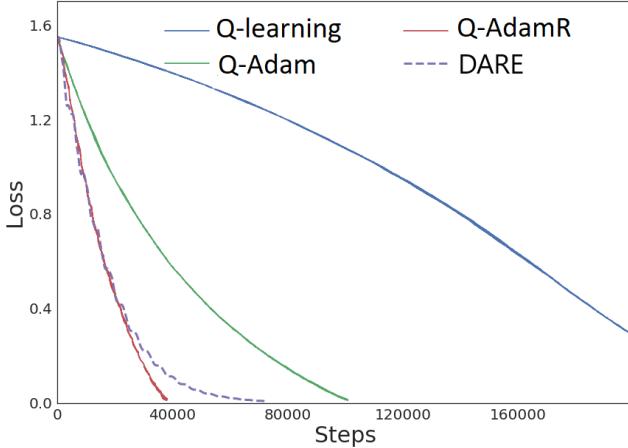


Figure 1: LQR experiments with performance evaluated in terms of policy loss  $\|K_t - K^*\|_2$ .

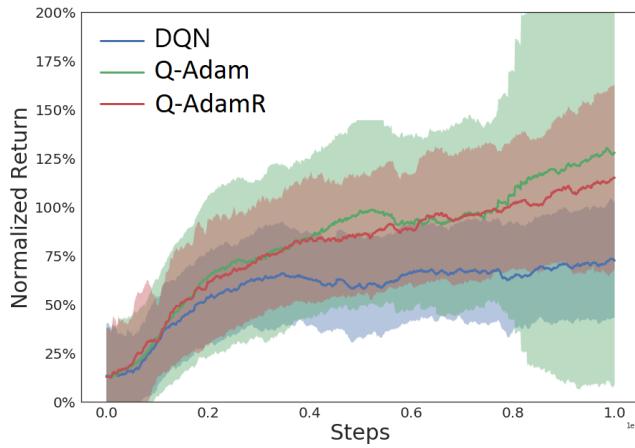


Figure 2: Atari game experiment with performance normalized and averaged over 23 games.

*et al.*, 2016], double Q-learning setup [Van Hasselt *et al.*, 2016],  $\epsilon$ -greedy exploration and experience replay [Mnih *et al.*, 2015]. Adam is also adopted, without momentum restart, as the optimizer for the inner-loop supervised learning process. Q-Adam and Q-AdamR are implemented using the identical setup of network construction, exploration and sampling strategies.

We test all three algorithms with a batch of 23 Atari games. The choice of 10 million steps of iteration is a common setup for benchmark experiments with Atari games. Although this does not guarantee the best performance in comparison with more time-consuming training with 50 million steps or more, it is sufficient to illustrate different performances among the selected methods. The software infrastructure is based on the baseline implementation of OpenAI. Selections of the hyper-parameters are listed in Table 2. We summarize the results in Figure 2. The overall performance is illustrated by first normalizing the return of each method with respect to the results obtained from DQN, and then averaging the performance of all 23 games to obtain the mean return and standard devia-

Table 2: Hyper-parameters for Atari games experiments of DQN, Q-Adam and Q-AdamR

Step size	Scale factor $\tilde{\tau}$	Adam $\beta_1$	Adam $\beta_2$
0.0001	0.0001	0.9	0.999
$r$	Buffer size	$\gamma$	Batch size $B$
$10^4$	$10^5$	0.99	32
Total training steps $K$	Target update frequency		
$10^7$	$10^4$		

tion. Considering we use a smaller buffer size than common practice, DQN is not consistently showing improved return over all tested games. Therefore, the self-normalized average return of DQN in Figure 2 is not strictly increasing from 0 to 100%.

Overall, both Q-Adam and Q-AdamR achieve significant improvement in comparison with the DQN results. In particular, AltQ-Adam increases the performance by over 100% in some of the tasks including Asterix, BeamRider, Enduro, Gopher, etc. However, it also illustrates certain instability with complete failure on Amidar and Assault. This is also captured by the higher variance illustrated on Figure 2. Periodic restart (Q-AdamR) resolves this issue efficiently with an on-par performance on average and far smaller variance. In terms of the maximum average return, Q-Adam and Q-AdamR perform no worse than DQN on 17 and 20 games respectively out of the 23 games being evaluated. Furthermore, if we consider having a final score that is smaller or equal to the start score as a learning failure, DQN fails the learning in 5 out of 23 games (Asteroids, DoubleDunk, Gravitar, Pitfall, Tennis), Q-Adam fails in 3 out of 23 games (Amidar, Assault, Asteroids) and Q-AdamR does not fail in any of the tasks. That is, Q-AdamR not only reduces the variance, but also provides a more consistent performance across the task domain. This implies that momentum restart effectively corrects the accumulated error and stabilizes the training process.

## 6 Conclusion

We study two Q-learning algorithms with Adam-type updates, and demonstrate their superior performance over the vanilla Q-learning and DQN algorithms through a linear quadratic regulator problem and a batch of 23 Atari 2600 games.

It is of considerable future interest to further investigate the potential of the restart scheme. One possible direction is to develop an adaptive restart mechanism with changing period determined by an appropriately defined signal of restart. This will potentially relieve the effort in hyper-parameter tuning of finding a good fixed period.

## Acknowledgements

The work was supported in part by the U.S. National Science Foundation under Grants CCF-1761506, ECCS-1818904, CCF-1909291 and CCF-1900145, and the startup funding of the Southern University of Science and Technology (SUSTech), China.

## References

[Bertsekas and Tsitsiklis, 1996] Dimitri P. Bertsekas and John N Tsitsiklis. *Neuro-Dynamic Programming*, volume 5. Athena Scientific, 1996.

[Bertsekas, 1995] Dimitri P Bertsekas. *Dynamic programming and optimal control*, volume 1. 1995.

[Bhandari *et al.*, 2018] Jalaj Bhandari, Daniel Russo, and Raghav Singal. A finite time analysis of temporal difference learning with linear function approximation. *arXiv preprint arXiv:1806.02450*, 2018.

[Brockman *et al.*, 2016] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.

[Chen *et al.*, 2019a] Xiangyi Chen, Sijia Liu, Ruoyu Sun, and Mingyi Hong. On the convergence of a class of Adam-type algorithms for non-convex optimization. In *International Conference on Learning Representations*, 2019.

[Chen *et al.*, 2019b] Zaiwei Chen, Sheng Zhang, Thinh T. Doan, Siva Theja Maguluri, and John-Paul Clarke. Finite-time analysis of Q-learning with linear function approximation. *arXiv preprint arXiv:1905.11425*, 2019.

[Du *et al.*, 2019] Simon S Du, Yuping Luo, Ruosong Wang, and Hanrui Zhang. Provably efficient Q-learning with function approximation via distribution shift error checking oracle. *arXiv preprint arXiv:1906.06321*, 2019.

[Even-Dar and Mansour, 2003] Eyal Even-Dar and Yishay Mansour. Learning rates for Q-learning. *Journal of Machine Learning Research*, 5:1–25, Dec 2003.

[Kingma and Ba, 2014] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[Knight and Lerner, 2018] Ethan Knight and Osher Lerner. Natural gradient deep Q-learning. *arXiv preprint arXiv:1803.07482*, 2018.

[Lewis and Vrabie, 2009] F. L. Lewis and D. Vrabie. Reinforcement learning and adaptive dynamic programming for feedback control. *IEEE Circuits and Systems Magazine*, 9(3):32–50, Third 2009.

[Lu *et al.*, 2018] Tyler Lu, Dale Schuurmans, and Craig Boutilier. Non-delusional Q-learning and value-iteration. In *Proceedings of the Thirty-second Conference on Neural Information Processing Systems (NeurIPS-18)*, pages 9971–9981, Montreal, QC, 2018.

[Mavrin *et al.*, 2019] Borislav Mavrin, Hengshuai Yao, and Linglong Kong. Deep reinforcement learning with decorrelation. *arXiv preprint arXiv:1903.07765*, 2019.

[Mnih *et al.*, 2013] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.

[Mnih *et al.*, 2015] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.

[Mnih *et al.*, 2016] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48, pages 1928–1937. PMLR, 20–22 Jun 2016.

[O’donoghue and Candes, 2015] Brendan O’donoghue and Emmanuel Candes. Adaptive restart for accelerated gradient schemes. *Foundations of computational mathematics*, 15(3):715–732, 2015.

[Reddi *et al.*, 2018] Sashank J. Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of Adam and beyond. In *International Conference on Learning Representations*, 2018.

[Schaul *et al.*, 2015] Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*, 2015.

[Tran and others, 2019] Phuong Thi Tran et al. On the convergence proof of AMSGrad and a new version. *IEEE Access*, 7:61706–61716, 2019.

[Vamvoudakis, 2017] Kyriakos G Vamvoudakis. Q-learning for continuous-time linear systems: A model-free infinite horizon optimal control approach. *Systems & Control Letters*, 100:14–20, 2017.

[Van Hasselt *et al.*, 2016] Hado Van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double Q-learning. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.

[Vrabie *et al.*, 2009] Draguna Vrabie, O. Pastravanu, Murad Abu-Khalaf, and Frank L. Lewis. Adaptive optimal control for continuous-time linear systems based on policy iteration. *Automatica*, 45(2):477–484, 2009.

[Wang *et al.*, 2016] Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Van Hasselt, Marc Lanctot, and Nando De Freitas. Dueling network architectures for deep reinforcement learning. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ICML’16, pages 1995–2003. JMLR.org, 2016.

[Watkins and Dayan, 1992] Christopher J.C.H. Watkins and Peter Dayan. Q-learning. *Machine Learning*, 8(3-4):279–292, 1992.

[Xiong *et al.*, 2020] Huaqing Xiong, Tengyu Xu, Yingbin Liang, and Wei Zhang. Non-asymptotic convergence of Adam-type reinforcement learning algorithms under markovian sampling. *arXiv preprint arXiv:2002.06286*, 2020.

[Xu and Gu, 2019] Pan Xu and Quanquan Gu. A finite-time analysis of q-learning with neural network function approximation. *arXiv preprint arXiv:1912.04511*, 2019.

[Yang *et al.*, 2019] Zhuora Yang, Yuchen Xie, and Zhaoran Wang. A theoretical analysis of deep Q-learning. *arXiv preprint arXiv:1901.00137*, 2019.

[Zhou *et al.*, 2018] Dongruo Zhou, Yiqi Tang, Ziyan Yang, Yuan Cao, and Quanquan Gu. On the convergence of adaptive gradient methods for nonconvex optimization. *arXiv preprint arXiv:1808.05671*, 2018.

[Zou *et al.*, 2018] Fangyu Zou, Li Shen, Zequn Jie, Weizhong Zhang, and Wei Liu. A sufficient condition for convergences of Adam and RMSProp. *arXiv preprint arXiv:1811.09358*, 2018.

[Zou *et al.*, 2019] Shaofeng Zou, Tengyu Xu, and Yingbin Liang. Finite-sample analysis for SARSA and Q-learning with linear function approximation. *arXiv preprint arXiv:1902.02234*, 2019.

## Supplementary Material

### A Further Details and Results on Experiments

We discuss more details on the experiment setup and provide further results that are not included in Section 5.

#### A.1 Linear Quadratic Regulator

The linear quadratic regulator (LQR) problem is of great interest for control community where Lewis et al. applies PQL to both discrete-time problems [Lewis and Vrabie, 2009] and continuous-time problems [Vamvoudakis, 2017; Vrabie *et al.*, 2009].

We empirically validate the proposed algorithms through an infinite-horizon discrete-time LQR problem defined as

$$\begin{aligned} \underset{\pi}{\text{minimize}} \quad & J = \sum_{t=0}^{\infty} (x_t^T Q x_t + u_t^T R u_t + 2x_t^T N u_t), \\ \text{subject to} \quad & x_{t+1} = Ax_t + Bu_t, \end{aligned}$$

where  $u_t = \pi(x_t)$ .

A typical model-based solution (with known  $A$  and  $B$ ) considers the problem backwards in time and iterates a dynamic equation known as the discrete-time algebraic Riccati equation (DARE):

$$P = A^T P A - (A^T P B + N)(R + B^T P B)^{-1}(B^T P A + N^T) + Q, \quad (15)$$

with the cost-to-go  $P$  being positive definite. The optimal policy satisfies  $u_t^* = -K^* x_t$  with

$$K^* = (R + B^T P B)^{-1}(N^T + B^T P A). \quad (16)$$

For experiments, we parameterize a quadratic Q-function with a matrix parameter  $H$  in the form of

$$Q(x, u; H) = \begin{bmatrix} x \\ u \end{bmatrix}^T \begin{bmatrix} H_{xx} & H_{xu} \\ H_{ux} & H_{uu} \end{bmatrix} \begin{bmatrix} x \\ u \end{bmatrix}. \quad (17)$$

The corresponding linear policy satisfies  $u = -Kx$ , and  $K = H_{uu}^{-1}H_{ux}$ . The performance of the learning algorithm is then evaluated at each step of iterate  $i$  with the Euclidean norm  $\|K_i - K^*\|_2$ .

#### A.2 Atari Games

We list detailed experiments of the 23 Atari games evaluated with the proposed algorithms in Figure 3. All experiments are executed with the same set of two random seeds. Each task takes about 20-hour of wall-clock time on a GPU instance. All three methods being evaluated share similar training time. AltQ-Adam and AltQ-AdamR can be further accelerated in practice with a more memory-efficient implementation considering the target network is not required. We keep our implementation of proposed algorithms consistent with the DQN we are comparing against. Other techniques that are not included in this experiment are also compatible with AltQ-Adam and AltQ-AdamR, such like asynchronous exploration [Mnih *et al.*, 2013] and training with decorrelated loss [Mavrin *et al.*, 2019].

Overall, AltQ-Adam significantly increases the performance by over 100% in some of the tasks including Asterix, BeamRider, Enduro, Gopher, etc. However, it also illustrates certain instability with complete failure on Amidar and Assault. This is mostly caused by the sampling where we are using a relevantly small buffer size with 10% of the common configured size in Atari games with experience replay. Notice that those failures tend to appear when the  $\epsilon$ -greedy exploration has evolved to a certain level where the immediate policy is effectively contributing to the accumulated experience. This potentially amplifies the biased exploration that essentially leads to the observed phenomenon.

Interestingly, AltQ-AdamR that incorporates the restart scheme resolves the problem of high variance of average return brought by AltQ-Adam and provides a more consistent performance across the task domain. This implies that momentum restart effectively corrects the accumulated error and stabilizes the training process.

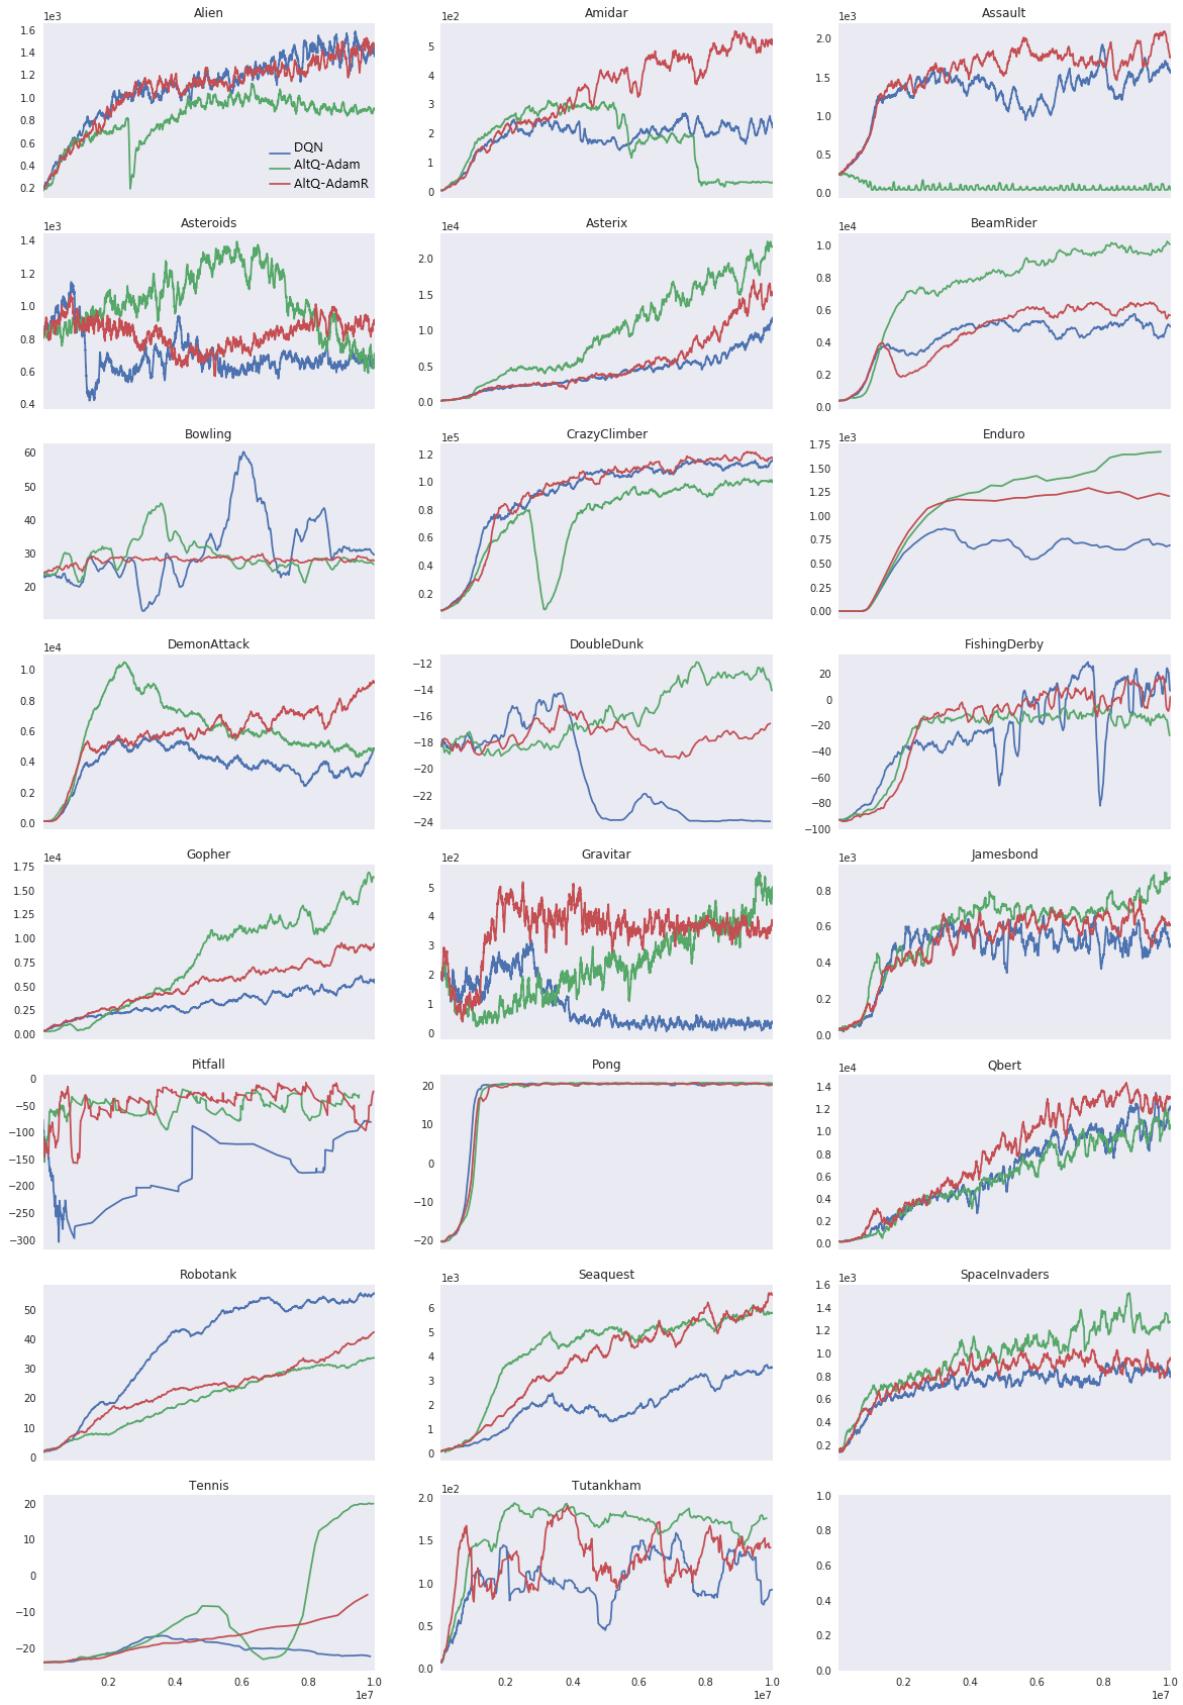


Figure 3: Experiment results of 23 Atari games with DQN, AltQ-Adam and AltQ-AdamR

Task	DQN	AltQ-Adam	AltQ-AdamR
Alien	1529	1125	<b>1587</b>
Amidar	269	313	<b>551</b>
Assault	1925	260	<b>2097</b>
Asteroids	1147	<b>1394</b>	1069
Asterix	11794	<b>22413</b>	17064
BeamRider	5728	<b>10210</b>	6458
Bowling	<b>60</b>	45	30
CrazyClimber	116422	102731	<b>121770</b>
Enduro	866	<b>1671</b>	1291
DemonAttack	5729	<b>10485</b>	9273
DoubleDunk	-14	<b>-12</b>	-15
FishingDerby	<b>29.01</b>	-4	19
Gopher	6066	<b>16863</b>	9508
Gravitar	316	<b>551</b>	518
Jamesbond	663	<b>899</b>	756
Pitfall	-76	-20	<b>-7</b>
Pong	20.68	<b>20.79</b>	20.74
Qbert	13453	12487	<b>14352</b>
Robotank	<b>56</b>	34	42
Seaquest	3652	6121	<b>6624</b>
Spaceinvaders	923	<b>1528</b>	1036
Tennis	-17	<b>20</b>	-5
Tutankham	159	<b>194</b>	191

Table 3: Best empirical return of 23 Atari games with DQN, AltQ-Adam and AltQ-AdamR

## B Proof of Lemma 1

The proof can be proceeded by the assumptions on the bounded kernel function and the bounded domains, which yields

$$\begin{aligned}
g_t &= \left( \phi^T(s_t, a_t) \theta_t - R(s_t, a_t) - \gamma \max_{a'} \phi^T(s_{t+1}, a') \theta_t \right) \phi(s_t, a_t) \\
&\stackrel{(i)}{\leq} \left\| \phi^T(s_t, a_t) \theta_t - R(s_t, a_t) - \gamma \max_{a'} \phi^T(s_{t+1}, a') \theta_t \right\| \|\phi(s_t, a_t)\| \\
&\stackrel{(ii)}{\leq} \|\phi^T(s_t, a_t) \theta_t\| + \|R(s_t, a_t)\| + \gamma \left\| \max_{a'} \phi^T(s_{t+1}, a') \theta_t \right\| \\
&\stackrel{(iii)}{\leq} \|g_t\| \leq R_{\max} + (1 + \gamma) D_{\infty},
\end{aligned}$$

where (i) follows from Cauchy-Schwarz inequality, (ii) follows from the triangle inequality and Assumption 1 and (iii) follows from Assumptions 1 and 3.

Recall that

$$\bar{g}_t = \mathbb{E}_{\mu} \left[ \left( \phi^T(s_t, \pi(s_t)) \theta_t - R(s_t, \pi(s_t)) - \gamma \max_{a'} \phi^T(s_{t+1}, a') \theta_t \right) \phi(s_t, \pi(s_t)) \right],$$

where  $\mu$  is the stationary distribution of the states. Then the bound of  $\bar{g}_t$  can be obtained by using the similar steps and techniques as those of bounding  $g_t$ .

The bounds of  $m_t$  and  $\hat{v}_t$  can be obtained by induction. To this end, we first check that  $m_0 = 0$  and  $\|m_1\| = \|\beta_{11}g_1\| \leq G_{\infty}$ . Assume that  $\|m_{t-1}\| \leq G_{\infty}$ , then we have

$$\begin{aligned}
\|m_t\| &= \|(1 - \beta_{1t})m_{t-1} + \beta_{1t}g_t\| \\
&\leq (1 - \beta_{1t})\|m_{t-1}\| + \beta_{1t}\|g_t\| \\
&\leq (1 - \beta_{1t})G_{\infty} + \beta_{1t}G_{\infty} \\
&= G_{\infty}.
\end{aligned}$$

We next bound  $\hat{v}_t$  similarly. First check  $\hat{v}_0 = 0$  and  $\|\hat{v}_1\| = \|\beta_2 g_1^2\| \leq G_\infty^2$ . Suppose that  $\|\hat{v}_{t-1}\| \leq G_\infty^2$ . Then we have

$$\begin{aligned}\|v_t\| &= \|(1 - \beta_2)v_{t-1} + \beta_2 g_t^2\| \\ &\leq (1 - \beta_2)\|v_{t-1}\| + \beta_2\|g_t^2\| \\ &\leq (1 - \beta_2)\|\hat{v}_{t-1}\| + \beta_2\|g_t^2\| \\ &\leq (1 - \beta_2)G_\infty^2 + \beta_2 G_\infty^2 \\ &= G_\infty^2.\end{aligned}$$

Thus we complete our proof by observing that  $\|\hat{v}_t\| \leq \max\{\|\hat{v}_{t-1}\|, \|v_t\|\} \leq G_\infty^2$ .

## C Proof of Theorem 1

Different from the regret bound for AMSGrad obtained in [Reddi *et al.*, 2018], our analysis is on the convergence rate. In fact, a slight modification of our proof also provides the convergence rate for AMSGrad for conventional strongly convex optimization, which can be of independent interest. Moreover, our proof avoids the theoretical error in the proof in [Reddi *et al.*, 2018] pointed out by [Tran and others, 2019]. Before proving the theorems, we first provide some useful lemmas.

**Lemma 2.** [Reddi *et al.*, 2018, Lemma 2] Let  $\{m_t, \hat{V}_t\}$  for  $t = 1, 2, \dots$  be sequences generated by Algorithm 1. Given  $\alpha_t, \beta_{1t}, \beta_2$  as specified in Theorem 1, we have

$$\begin{aligned}\sum_{t=1}^T \alpha_t \left\| \hat{V}_t^{-\frac{1}{4}} m_t \right\|^2 &\leq \frac{\alpha}{(1 - \beta_1)(1 - \delta)\sqrt{1 - \beta_2}} \sum_{i=1}^d \|g_{1:T,i}\| \sqrt{\sum_{t=1}^T \frac{1}{t}} \\ &\leq \frac{\alpha\sqrt{1 + \log T}}{(1 - \beta_1)(1 - \delta)\sqrt{1 - \beta_2}} \sum_{i=1}^d \|g_{1:T,i}\|.\end{aligned}$$

**Lemma 3.** Let  $\alpha_t = \frac{\alpha}{\sqrt{t}}$  and  $\beta_{1t} = \beta_1 \lambda^t$  for  $t = 1, 2, \dots$ . Then

$$\sum_{t=1}^T \frac{\beta_{1t}}{\alpha_t} \leq \frac{\beta_1}{\alpha(1 - \lambda)^2}. \quad (18)$$

*Proof.* The proof is based on taking the standard sum of geometric sequences.

$$\sum_{t=1}^T \frac{\beta_{1t}}{\alpha_t} = \sum_{t=1}^T \frac{\beta_{1t}\sqrt{t}}{\alpha} \leq \sum_{t=1}^T \frac{\beta_1 \lambda^{t-1} t}{\alpha} = \frac{\beta_1}{\alpha} \left( \frac{1}{(1 - \lambda)} \sum_{t=1}^T \lambda^{t-1} - T\lambda^T \right) \leq \frac{\beta_1}{\alpha(1 - \lambda)^2}. \quad (19)$$

□

With the lemmas above, we are ready to prove Theorem 1. Observe that

$$\theta_{t+1} = \Pi_{\mathcal{D}, \hat{V}_t^{1/4}} \left( \theta_t - \alpha_t \hat{V}_t^{-\frac{1}{2}} m_t \right) = \min_{\theta \in \mathcal{D}} \left\| \hat{V}_t^{1/4} \left( \theta_t - \alpha_t \hat{V}_t^{-\frac{1}{2}} m_t - \theta \right) \right\|.$$

Clearly  $\Pi_{\mathcal{D}, \hat{V}_t^{1/4}}(\theta^*) = \theta^*$  due to Assumption 3. We start from the update of  $\theta_t$  when  $t \geq 2$ .

$$\begin{aligned}\left\| \hat{V}_t^{1/4}(\theta_{t+1} - \theta^*) \right\|^2 &= \left\| \Pi_{\mathcal{D}, \hat{V}_t^{1/4}} \hat{V}_t^{1/4} \left( \theta_t - \theta^* - \alpha_t \hat{V}_t^{-\frac{1}{2}} m_t \right) \right\|^2 \\ &\leq \left\| \hat{V}_t^{1/4} \left( \theta_t - \theta^* - \alpha_t \hat{V}_t^{-\frac{1}{2}} m_t \right) \right\|^2 \\ &= \left\| \hat{V}_t^{1/4}(\theta_t - \theta^*) \right\|^2 + \left\| \alpha_t \hat{V}_t^{-1/4} m_t \right\|^2 - 2\alpha_t(\theta_t - \theta^*)^T m_t \\ &= \left\| \hat{V}_t^{1/4}(\theta_t - \theta^*) \right\|^2 + \left\| \alpha_t \hat{V}_t^{-1/4} m_t \right\|^2 - 2\alpha_t(\theta_t - \theta^*)^T (\beta_{1t} m_{t-1} + (1 - \beta_{1t}) g_t) \\ &\stackrel{(i)}{\leq} \left\| \hat{V}_t^{1/4}(\theta_t - \theta^*) \right\|^2 + \left\| \alpha_t \hat{V}_t^{-1/4} m_t \right\|^2 + \alpha_t \beta_{1t} \left( \frac{1}{\alpha_t} \left\| \hat{V}_t^{1/4}(\theta_t - \theta^*) \right\|^2 + \alpha_t \left\| \hat{V}_t^{-1/4} m_{t-1} \right\|^2 \right) \\ &\quad - 2\alpha_t(1 - \beta_{1t})(\theta_t - \theta^*)^T g_t \\ &\stackrel{(ii)}{\leq} \left\| \hat{V}_t^{1/4}(\theta_t - \theta^*) \right\|^2 + \left\| \alpha_t \hat{V}_t^{-1/4} m_t \right\|^2 + \beta_{1t} \left\| \hat{V}_t^{1/4}(\theta_t - \theta^*) \right\|^2 + \alpha_t^2 \beta_{1t} \left\| \hat{V}_{t-1}^{-1/4} m_{t-1} \right\|^2 \\ &\quad - 2\alpha_t(1 - \beta_{1t})(\theta_t - \theta^*)^T g_t,\end{aligned}$$

where (i) follows from Cauchy-Schwarz inequality, and (ii) holds because  $\hat{v}_{t+1,i} \geq \hat{v}_{t,i}, \forall t, \forall i$ . Next, we take the expectation over all samples used up to time step  $t$  on both sides, which still preserves the inequality. Since we consider i.i.d. sampling case, by letting  $\mathcal{F}_t$  be the filtration of all the sampling up to time  $t$ , we have

$$\mathbb{E} [(\theta_t - \theta^*)^T g_t] = \mathbb{E} [\mathbb{E} [(\theta_t - \theta^*)^T g_t] | \mathcal{F}_{t-1}] = \mathbb{E} [(\theta_t - \theta^*)^T \bar{g}_t]. \quad (20)$$

Thus we have

$$\begin{aligned} \mathbb{E} \left\| \hat{V}_t^{1/4} (\theta_{t+1} - \theta^*) \right\|^2 &\leq \mathbb{E} \left\| \hat{V}_t^{1/4} (\theta_t - \theta^*) \right\|^2 + \alpha_t^2 \mathbb{E} \left\| \hat{V}_t^{-1/4} m_t \right\|^2 + \beta_{1t} \mathbb{E} \left\| \hat{V}_t^{1/4} (\theta_t - \theta^*) \right\|^2 + \alpha_t^2 \beta_{1t} \mathbb{E} \left\| \hat{V}_{t-1}^{-1/4} m_{t-1} \right\|^2 \\ &\quad - 2\alpha_t (1 - \beta_{1t}) \mathbb{E} [(\theta_t - \theta^*)^T g_t] \\ &\stackrel{(i)}{=} \mathbb{E} \left\| \hat{V}_t^{1/4} (\theta_t - \theta^*) \right\|^2 + \alpha_t^2 \mathbb{E} \left\| \hat{V}_t^{-1/4} m_t \right\|^2 + \beta_{1t} \mathbb{E} \left\| \hat{V}_t^{1/4} (\theta_t - \theta^*) \right\|^2 + \alpha_t^2 \beta_{1t} \mathbb{E} \left\| \hat{V}_{t-1}^{-1/4} m_{t-1} \right\|^2 \\ &\quad - 2\alpha_t (1 - \beta_{1t}) \mathbb{E} [(\theta_t - \theta^*)^T \bar{g}_t] \\ &\stackrel{(ii)}{\leq} \mathbb{E} \left\| \hat{V}_t^{1/4} (\theta_t - \theta^*) \right\|^2 + \alpha_t^2 \mathbb{E} \left\| \hat{V}_t^{-1/4} m_t \right\|^2 + \beta_{1t} \mathbb{E} \left\| \hat{V}_t^{1/4} (\theta_t - \theta^*) \right\|^2 + \alpha_t^2 \beta_{1t} \mathbb{E} \left\| \hat{V}_{t-1}^{-1/4} m_{t-1} \right\|^2 \\ &\quad - 2\alpha_t c (1 - \beta_{1t}) \mathbb{E} \|\theta_t - \theta^*\|^2 \\ &\stackrel{(iii)}{\leq} \mathbb{E} \left\| \hat{V}_t^{1/4} (\theta_t - \theta^*) \right\|^2 + \alpha_t^2 \mathbb{E} \left\| \hat{V}_t^{-1/4} m_t \right\|^2 + \beta_{1t} \mathbb{E} \left\| \hat{V}_t^{1/4} (\theta_t - \theta^*) \right\|^2 + \alpha_t^2 \beta_1 \mathbb{E} \left\| \hat{V}_{t-1}^{-1/4} m_{t-1} \right\|^2 \\ &\quad - 2\alpha_t c (1 - \beta_1) \mathbb{E} \|\theta_t - \theta^*\|^2 \\ &\stackrel{(iv)}{\leq} \mathbb{E} \left\| \hat{V}_t^{1/4} (\theta_t - \theta^*) \right\|^2 + \alpha_t^2 \mathbb{E} \left\| \hat{V}_t^{-1/4} m_t \right\|^2 + G_\infty D_\infty^2 \beta_{1t} + \alpha_t^2 \beta_1 \mathbb{E} \left\| \hat{V}_{t-1}^{-1/4} m_{t-1} \right\|^2 \\ &\quad - 2\alpha_t c (1 - \beta_1) \mathbb{E} \|\theta_t - \theta^*\|^2, \end{aligned}$$

where (i) follows from Equation (20), (ii) follows due to Assumption 2 and  $1 - \beta_{1t} > 0$ , (iii) follows from  $\beta_{1t} < \beta_1 < 1$  and  $\mathbb{E} \|\theta_t - \theta^*\|^2 > 0$ , and (iv) follows from  $\left\| \hat{V}_t^{1/4} (\theta_t - \theta^*) \right\|^2 \leq \left\| \hat{V}_t^{1/4} \right\|_2^2 \|\theta_t - \theta^*\|^2 \leq G_\infty D_\infty^2$  by Lemma 1 and Assumption 3. We note that (iii) is the key step to avoid the error in the proof in [Reddi *et al.*, 2018], where we can directly bound  $1 - \beta_{1t}$ , which is impossible in [Reddi *et al.*, 2018]. By rearranging the terms in the above inequality and taking the summation over time steps, we have

$$\begin{aligned} 2c(1 - \beta_1) \sum_{t=2}^T \mathbb{E} \|\theta_t - \theta^*\|^2 &\leq \sum_{t=2}^T \frac{1}{\alpha_t} \left( \mathbb{E} \left\| \hat{V}_t^{1/4} (\theta_t - \theta^*) \right\|^2 - \mathbb{E} \left\| \hat{V}_t^{1/4} (\theta_{t+1} - \theta^*) \right\|^2 \right) + \sum_{t=2}^T \frac{\beta_{1t} G_\infty D_\infty^2}{\alpha_t} \\ &\quad + \sum_{t=2}^T \alpha_t \mathbb{E} \left\| \hat{V}_t^{-1/4} m_t \right\|^2 + \sum_{t=2}^T \alpha_t \beta_1 \mathbb{E} \left\| \hat{V}_{t-1}^{-1/4} m_{t-1} \right\|^2 \\ &\stackrel{(i)}{\leq} \sum_{t=2}^T \frac{1}{\alpha_t} \left( \mathbb{E} \left\| \hat{V}_t^{1/4} (\theta_t - \theta^*) \right\|^2 - \mathbb{E} \left\| \hat{V}_t^{1/4} (\theta_{t+1} - \theta^*) \right\|^2 \right) + \sum_{t=2}^T \frac{\beta_{1t} G_\infty D_\infty^2}{\alpha_t} \\ &\quad + \sum_{t=2}^T \alpha_t \mathbb{E} \left\| \hat{V}_t^{-1/4} m_t \right\|^2 + \sum_{t=2}^T \alpha_{t-1} \beta_1 \mathbb{E} \left\| \hat{V}_{t-1}^{-1/4} m_{t-1} \right\|^2 \\ &\leq \sum_{t=2}^T \frac{1}{\alpha_t} \left( \mathbb{E} \left\| \hat{V}_t^{1/4} (\theta_t - \theta^*) \right\|^2 - \mathbb{E} \left\| \hat{V}_t^{1/4} (\theta_{t+1} - \theta^*) \right\|^2 \right) + \sum_{t=2}^T \frac{\beta_{1t} G_\infty D_\infty^2}{\alpha_t} \\ &\quad + (1 + \beta_1) \sum_{t=1}^T \alpha_t \mathbb{E} \left\| \hat{V}_t^{-1/4} m_t \right\|^2, \end{aligned}$$

where (i) follows from  $\alpha_t < \alpha_{t-1}$ . With further adjustment of the first term in the right hand side of the last inequality, we can

then bound the sum as

$$\begin{aligned}
& 2c(1 - \beta_1) \sum_{t=2}^T \mathbb{E} \|\theta_t - \theta^*\|^2 \\
& \leq \sum_{t=2}^T \frac{1}{\alpha_t} \mathbb{E} \left( \left\| \hat{V}_t^{1/4} (\theta_t - \theta^*) \right\|^2 - \left\| \hat{V}_t^{1/4} (\theta_{t+1} - \theta^*) \right\|^2 \right) + \sum_{t=2}^T \frac{\beta_{1t} G_\infty D_\infty^2}{\alpha_t} \\
& \quad + (1 + \beta_1) \sum_{t=1}^T \alpha_t \mathbb{E} \left\| \hat{V}_t^{-1/4} m_t \right\|^2 \\
& = \frac{\mathbb{E} \left\| \hat{V}_2^{1/4} (\theta_2 - \theta^*) \right\|^2}{\alpha_2} + \sum_{t=3}^T \mathbb{E} \left( \frac{\left\| \hat{V}_t^{1/4} (\theta_t - \theta^*) \right\|^2}{\alpha_t} - \frac{\left\| \hat{V}_{t-1}^{1/4} (\theta_t - \theta^*) \right\|^2}{\alpha_{t-1}} \right) \\
& \quad - \frac{\mathbb{E} \left\| \hat{V}_T^{1/4} (\theta_{T+1} - \theta^*) \right\|^2}{\alpha_T} + \sum_{t=2}^T \frac{\beta_{1t} G_\infty D_\infty^2}{\alpha_t} + (1 + \beta_1) \sum_{t=1}^T \alpha_t \mathbb{E} \left\| \hat{V}_t^{-1/4} m_t \right\|^2 \\
& = \frac{\mathbb{E} \left\| \hat{V}_2^{1/4} (\theta_2 - \theta^*) \right\|^2}{\alpha_2} + \sum_{t=3}^T \mathbb{E} \left( \frac{\sum_{i=1}^d \hat{v}_{t,i}^{1/2} (\theta_{t,i} - \theta_i^*)^2}{\alpha_t} - \frac{\sum_{i=1}^d \hat{v}_{t-1,i}^{1/2} (\theta_{t,i} - \theta_i^*)^2}{\alpha_{t-1}} \right) \\
& \quad - \frac{\mathbb{E} \left\| \hat{V}_T^{1/4} (\theta_{T+1} - \theta^*) \right\|^2}{\alpha_T} + \sum_{t=2}^T \frac{\beta_{1t} G_\infty D_\infty^2}{\alpha_t} + (1 + \beta_1) \sum_{t=1}^T \alpha_t \mathbb{E} \left\| \hat{V}_t^{-1/4} m_t \right\|^2 \\
& = \frac{\mathbb{E} \left\| \hat{V}_2^{1/4} (\theta_2 - \theta^*) \right\|^2}{\alpha_2} + \sum_{t=3}^T \sum_{i=1}^d \mathbb{E} (\theta_{t,i} - \theta_i^*)^2 \left( \frac{\hat{v}_{t,i}^{1/2}}{\alpha_t} - \frac{\hat{v}_{t-1,i}^{1/2}}{\alpha_{t-1}} \right) \\
& \quad - \frac{\mathbb{E} \left\| \hat{V}_T^{1/4} (\theta_{T+1} - \theta^*) \right\|^2}{\alpha_T} + \sum_{t=2}^T \frac{\beta_{1t} G_\infty D_\infty^2}{\alpha_t} + (1 + \beta_1) \sum_{t=1}^T \alpha_t \mathbb{E} \left\| \hat{V}_t^{-1/4} m_t \right\|^2.
\end{aligned}$$

So far we just rearrange the terms in the series sum. Next, we are ready to obtain the upper bound.

$$\begin{aligned}
& 2c(1 - \beta_1) \sum_{t=2}^T \mathbb{E} \|\theta_t - \theta^*\|^2 \\
& \stackrel{(i)}{\leq} \frac{\mathbb{E} \left\| \hat{V}_2^{1/4} (\theta_2 - \theta^*) \right\|^2}{\alpha_2} + D_\infty^2 \sum_{t=3}^T \sum_{i=1}^d \mathbb{E} \left( \frac{\hat{v}_{t,i}^{1/2}}{\alpha_t} - \frac{\hat{v}_{t-1,i}^{1/2}}{\alpha_{t-1}} \right) \\
& \quad - \frac{\mathbb{E} \left\| \hat{V}_T^{1/4} (\theta_{T+1} - \theta^*) \right\|^2}{\alpha_T} + \sum_{t=2}^T \frac{\beta_{1t} G_\infty D_\infty^2}{\alpha_t} + (1 + \beta_1) \sum_{t=1}^T \alpha_t \mathbb{E} \left\| \hat{V}_t^{-1/4} m_t \right\|^2 \\
& \leq \frac{\mathbb{E} \left\| \hat{V}_2^{1/4} (\theta_2 - \theta^*) \right\|^2}{\alpha_2} + D_\infty^2 \sum_{i=1}^d \mathbb{E} \frac{\hat{v}_{T,i}^{1/2}}{\alpha_T} + \sum_{t=2}^T \frac{\beta_{1t} G_\infty D_\infty^2}{\alpha_t} + (1 + \beta_1) \sum_{t=1}^T \alpha_t \mathbb{E} \left\| \hat{V}_t^{-1/4} m_t \right\|^2 \\
& \stackrel{(ii)}{\leq} \frac{G_\infty D_\infty^2}{\alpha_2} + \frac{d G_\infty D_\infty^2 \sqrt{T}}{\alpha} + \frac{\beta_1 G_\infty D_\infty^2}{\alpha(1 - \lambda)^2} + \frac{\alpha(1 + \beta_1) \sqrt{1 + \log T}}{(1 - \beta_1)(1 - \delta) \sqrt{1 - \beta_2}} \sum_{i=1}^d \mathbb{E} \|g_{1:T,i}\|, \tag{21}
\end{aligned}$$

where (i) follows from Assumption 3 and because  $\frac{\hat{v}_{t,i}^{1/2}}{\alpha_t} > \frac{\hat{v}_{t-1,i}^{1/2}}{\alpha_{t-1}}$ , and (ii) follows from Lemmas 1 - 3.

Finally, applying the Jensen's inequality yields

$$\mathbb{E} \|\theta_{out} - \theta^*\|^2 \leq \frac{1}{T} \sum_{t=1}^T \mathbb{E} \|\theta_t - \theta^*\|^2. \tag{22}$$

We conclude our proof by further applying the bound in Equation (21) to Equation (22).

## D Proof of Theorem 2

To prove the convergence for AltQ-AMSGradR, the major technical development beyond the proof of Theorem 1 lies in dealing with the parameter restart. More specifically, the moment approximation terms are reset every  $r$  steps, i.e.,  $m_{kr} = \hat{v}_{kr} = 0$  for  $k = 1, 2, \dots$ , which implies  $\theta_{kr+1} = \theta_{kr}$  for  $k = 1, 2, \dots$ . For technical convenience, we define  $\theta_0 = \theta_1$ . Using the arguments similar to Equation (21), in a time window that does not contain a restart (i.e.  $kr \leq S \leq (k+1)r-1$ ) we have

$$\begin{aligned}
& 2c(1-\beta_1) \sum_{t=kr}^S \mathbb{E} \|\theta_t - \theta^*\|^2 \\
& \stackrel{(i)}{\leq} \frac{G_\infty D_\infty^2}{\alpha_{kr+2}} + \frac{dG_\infty D_\infty^2 \sqrt{S}}{\alpha} + \frac{\alpha(1+\beta_1)}{(1-\beta_1)(1-\delta)\sqrt{1-\beta_2}} \sum_{i=1}^d \mathbb{E} \|g_{kr+1:S,i}\| \sqrt{\sum_{t=kr+1}^S \frac{1}{t}} \\
& \quad + G_\infty D_\infty^2 \sum_{t=kr+2}^S \frac{\beta_{1t}}{\alpha_t} + 2c(1-\beta_1) \left( \mathbb{E} \|\theta_{kr+1} - \theta^*\|^2 + \mathbb{E} \|\theta_{kr} - \theta^*\|^2 \right) \\
& \stackrel{(ii)}{=} \frac{G_\infty D_\infty^2 \sqrt{kr+2}}{\alpha} + \frac{dG_\infty D_\infty^2 \sqrt{S}}{\alpha} + \frac{\alpha(1+\beta_1)}{(1-\beta_1)(1-\delta)\sqrt{1-\beta_2}} \sum_{i=1}^d \mathbb{E} \|g_{kr+1:S,i}\| \sqrt{\sum_{t=kr+1}^S \frac{1}{t}} \\
& \quad + G_\infty D_\infty^2 \sum_{t=kr+2}^S \frac{\beta_{1t}}{\alpha_t} + 4c(1-\beta_1) \mathbb{E} \|\theta_{kr} - \theta^*\|^2,
\end{aligned}$$

where (i) follows from Equation (21) and (ii) follows from  $\theta_{kr+1} = \theta_{kr}$  due to the definition of restart. Then we take the summation over the total time steps and obtain

$$\begin{aligned}
& 2c(1-\beta_1) \sum_{t=1}^T \mathbb{E} \|\theta_t - \theta^*\|^2 \\
& = 2c(1-\beta_1) \left( \sum_{k=1}^{\lfloor T/r \rfloor} \sum_{t=(k-1)r}^{kr-1} \mathbb{E} \|\theta_t - \theta^*\|^2 + \sum_{t=\lfloor T/r \rfloor r}^T \mathbb{E} \|\theta_t - \theta^*\|^2 - \mathbb{E} \|\theta_0 - \theta^*\|^2 \right) \\
& \leq \sum_{k=0}^{\lfloor T/r \rfloor} \left( \frac{G_\infty D_\infty^2}{\alpha} \sqrt{kr+2} + 4c(1-\beta_1) \mathbb{E} \|\theta_{kr} - \theta^*\|^2 \right) + \sum_{k=1}^{\lfloor T/r \rfloor} \frac{dG_\infty D_\infty^2}{\alpha} \sqrt{kr-1} \\
& \quad + \frac{dG_\infty D_\infty^2 \sqrt{T}}{\alpha} + \frac{\alpha(1+\beta_1)}{(1-\beta_1)(1-\delta)\sqrt{1-\beta_2}} \sum_{k=1}^{\lfloor T/r \rfloor} \sum_{i=1}^d \mathbb{E} \|g_{(k-1)r+1:kr-1,i}\| \sqrt{\sum_{t=(k-1)r+1}^{kr-1} \frac{1}{t}} \\
& \quad + \frac{\alpha(1+\beta_1)}{(1-\beta_1)(1-\delta)\sqrt{1-\beta_2}} \sum_{i=1}^d \mathbb{E} \|g_{\lfloor T/r \rfloor r+1:T,i}\| \sqrt{\sum_{t=\lfloor T/r \rfloor r+1}^T \frac{1}{t}} \\
& \quad + G_\infty D_\infty^2 \sum_{k=1}^{\lfloor T/r \rfloor} \sum_{t=(k-1)r+2}^{kr-1} \frac{\beta_{1t}}{\alpha_t} + G_\infty D_\infty^2 \sum_{t=\lfloor T/r \rfloor r+2}^T \frac{\beta_{1t}}{\alpha_t} \\
& \leq \sum_{k=0}^{\lfloor T/r \rfloor} \left( \frac{G_\infty D_\infty^2}{\alpha} \sqrt{kr+2} + 4c(1-\beta_1) \mathbb{E} \|\theta_{kr} - \theta^*\|^2 \right) + \sum_{k=1}^{\lfloor T/r \rfloor} \frac{dG_\infty D_\infty^2}{\alpha} \sqrt{kr-1} \\
& \quad + \frac{dG_\infty D_\infty^2 \sqrt{T}}{\alpha} + \frac{\alpha(1+\beta_1)}{(1-\beta_1)(1-\delta)\sqrt{1-\beta_2}} \sum_{k=1}^{\lfloor T/r \rfloor} \sum_{i=1}^d \mathbb{E} \|g_{(k-1)r+1:kr-1,i}\| \sqrt{\sum_{t=(k-1)r+1}^{kr-1} \frac{1}{t}} \\
& \quad + \frac{\alpha(1+\beta_1)}{(1-\beta_1)(1-\delta)\sqrt{1-\beta_2}} \sum_{i=1}^d \mathbb{E} \|g_{\lfloor T/r \rfloor r+1:T,i}\| \sqrt{\sum_{t=\lfloor T/r \rfloor r+1}^T \frac{1}{t}} + G_\infty D_\infty^2 \sum_{t=1}^T \frac{\beta_{1t}}{\alpha_t}.
\end{aligned}$$

We can bound the term  $G_\infty D_\infty^2 \sum_{t=1}^T \frac{\beta_{1t}}{\alpha_t}$  by Lemma 3. Next, we bound another key term in the above inequality. We first observe that  $\forall k \geq 2, \forall i \in [d]$ ,

$$\begin{aligned} \|g_{(k-1)r+1:kr-1,i}\| \sqrt{\sum_{t=(k-1)r+1}^{kr-1} \frac{1}{t}} &\stackrel{(i)}{\leq} \|g_{(k-1)r+1:kr-1,i}\| \sqrt{\sum_{t=(k-1)r+1}^{kr-1} \frac{1}{t} + |g_{kr,i}| \sqrt{\frac{1}{kr}}} \\ &\stackrel{(ii)}{\leq} \|g_{(k-1)r+1:kr,i}\| \sqrt{\sum_{t=(k-1)r+1}^{kr} \frac{1}{t}}, \end{aligned} \quad (23)$$

where (i) holds due to  $|g_{t,i}| \sqrt{\frac{1}{t}} > 0$  and (ii) follows from Cauchy-Schwarz inequality. Then we have

$$\begin{aligned} &\sum_{k=1}^{\lfloor T/r \rfloor} \sum_{i=1}^d \|g_{(k-1)r+1:kr-1,i}\| \sqrt{\sum_{t=(k-1)r+1}^{kr-1} \frac{1}{t} + \sum_{i=1}^d \|g_{[T/r]r+1:T,i}\| \sqrt{\sum_{t=[T/r]r+1}^T \frac{1}{t}}} \\ &\stackrel{(i)}{\leq} \sum_{k=1}^{\lfloor T/r \rfloor} \sum_{i=1}^d |g_{kr,i}| \sqrt{\frac{1}{kr}} + \sum_{k=1}^{\lfloor T/r \rfloor} \sum_{i=1}^d \|g_{(k-1)r+1:kr-1,i}\| \sqrt{\sum_{t=(k-1)r+1}^{kr-1} \frac{1}{t}} \\ &\quad + \sum_{i=1}^d \|g_{[T/r]r+1:T,i}\| \sqrt{\sum_{t=[T/r]r+1}^T \frac{1}{t}} \\ &= \sum_{k=1}^{\lfloor T/r \rfloor} \sum_{i=1}^d \left( \|g_{(k-1)r+1:kr-1,i}\| \sqrt{\sum_{t=(k-1)r+1}^{kr-1} \frac{1}{t}} + |g_{kr,i}| \sqrt{\frac{1}{kr}} \right) \\ &\quad + \sum_{i=1}^d \|g_{[T/r]r+1:T,i}\| \sqrt{\sum_{t=[T/r]r+1}^T \frac{1}{t}} \\ &\stackrel{(ii)}{\leq} \sum_{k=1}^{\lfloor T/r \rfloor} \sum_{i=1}^d \|g_{(k-1)r+1:kr,i}\| \sqrt{\sum_{t=(k-1)r+1}^{kr} \frac{1}{t}} + \sum_{i=1}^d \|g_{[T/r]r+1:T,i}\| \sqrt{\sum_{t=[T/r]r+1}^T \frac{1}{t}} \\ &= \sum_{i=1}^d \left( \sum_{k=1}^{\lfloor T/r \rfloor} \|g_{(k-1)r+1:kr,i}\| \sqrt{\sum_{t=(k-1)r+1}^{kr} \frac{1}{t}} + \|g_{[T/r]r+1:T,i}\| \sqrt{\sum_{t=[T/r]r+1}^T \frac{1}{t}} \right) \\ &\stackrel{(iii)}{\leq} \sum_{i=1}^d \|g_{1:T,i}\| \sqrt{\sum_{t=1}^T \frac{1}{t}}, \end{aligned}$$

where (i) follows from  $|g_{kr,i}| \sqrt{\frac{1}{kr}}, \forall k \geq 1, \forall i \in [d]$ , (ii) follows from Equation (23) and (iii) holds due to Cauchy-Schwarz inequality. Then we have

$$\begin{aligned} &2c(1-\beta_1) \sum_{t=1}^T \mathbb{E} \|\theta_t - \theta^*\|^2 \\ &\leq \sum_{k=0}^{\lfloor T/r \rfloor} \left( \frac{G_\infty D_\infty^2}{\alpha} \sqrt{kr+2} + 4c(1-\beta_1) \mathbb{E} \|\theta_{kr} - \theta^*\|^2 \right) + \sum_{k=1}^{\lfloor T/r \rfloor} \frac{dG_\infty D_\infty^2}{\alpha} \sqrt{kr-1} \\ &\quad + \frac{dG_\infty D_\infty^2 \sqrt{T}}{\alpha} + \frac{\alpha(1+\beta_1)}{(1-\beta_1)(1-\delta)\sqrt{1-\beta_2}} \sum_{k=1}^{\lfloor T/r \rfloor} \sum_{i=1}^d \mathbb{E} \|g_{(k-1)r:kr-1,i}\| \sqrt{\sum_{t=(k-1)r}^{kr-1} \frac{1}{t}} \\ &\quad + \frac{\alpha(1+\beta_1)}{(1-\beta_1)(1-\delta)\sqrt{1-\beta_2}} \sum_{i=1}^d \mathbb{E} \|g_{[T/r]r:T,i}\| \sqrt{\sum_{t=[T/r]r}^T \frac{1}{t}} + G_\infty D_\infty^2 \sum_{t=1}^T \frac{\beta_{1t}}{\alpha_t} \end{aligned}$$

$$\begin{aligned}
&\leq \sum_{k=0}^{\lfloor T/r \rfloor} \left( \frac{G_\infty D_\infty^2}{\alpha} \sqrt{kr+2} + 4c(1-\beta_1)\mathbb{E} \|\theta_{kr} - \theta^*\|^2 \right) + \sum_{k=1}^{\lfloor T/r \rfloor} \frac{dG_\infty D_\infty^2}{\alpha} \sqrt{kr-1} \\
&\quad + \frac{dG_\infty D_\infty^2 \sqrt{T}}{\alpha} + \frac{\alpha(1+\beta_1)}{(1-\beta_1)(1-\delta)\sqrt{1-\beta_2}} \sum_{i=1}^d \mathbb{E} \|g_{1:T,i}\| \sqrt{\sum_{t=1}^T \frac{1}{t}} + G_\infty D_\infty^2 \sum_{t=1}^T \frac{\beta_{1t}}{\alpha_t} \\
&\stackrel{(i)}{\leq} \sum_{k=0}^{\lfloor T/r \rfloor} \left( \frac{G_\infty D_\infty^2}{\alpha} \sqrt{kr+2} + 4c(1-\beta_1)\mathbb{E} \|\theta_{kr} - \theta^*\|^2 \right) + \sum_{k=1}^{\lfloor T/r \rfloor} \frac{dG_\infty D_\infty^2}{\alpha} \sqrt{kr-1} \\
&\quad + \frac{dG_\infty D_\infty^2 \sqrt{T}}{\alpha} + \frac{\alpha(1+\beta_1)\sqrt{d(1+\log T)}}{(1-\beta_1)(1-\delta)\sqrt{1-\beta_2}} \sum_{i=1}^d \mathbb{E} \|g_{1:T,i}\| + \frac{\beta_1 G_\infty D_\infty^2}{\alpha(1-\lambda)^2},
\end{aligned}$$

where (i) follows from Lemma 2 and Lemma 3.

Finally, applying the Jensen's inequality and the above bound, we obtain

$$\begin{aligned}
\mathbb{E} \|\theta_{out} - \theta^*\|^2 &\leq \frac{1}{T} \sum_{t=1}^T \mathbb{E} \|\theta_t - \theta^*\|^2 \\
&\leq \frac{1}{T} \sum_{k=0}^{\lfloor T/r \rfloor} \left( \frac{G_\infty D_\infty^2}{2c\alpha(1-\beta_1)} \sqrt{kr+2} + 2\mathbb{E} \|\theta_{kr} - \theta^*\|^2 \right) + \frac{1}{T} \sum_{k=1}^{\lfloor T/r \rfloor} \frac{dG_\infty D_\infty^2}{2c\alpha(1-\beta_1)} \sqrt{kr-1} \\
&\quad + \frac{dG_\infty D_\infty^2 \sqrt{T}}{2c\alpha(1-\beta_1)} + \frac{\alpha(1+\beta_1)\sqrt{d(1+\log T)}}{2c(1-\beta_1)^2(1-\delta)\sqrt{1-\beta_2}} \sum_{i=1}^d \mathbb{E} \|g_{1:T,i}\| + \frac{\beta_1 G_\infty D_\infty^2}{2c\alpha(1-\beta_1)(1-\lambda)^2},
\end{aligned}$$

which concludes the proof.