

Accelerating Stochastic Composition Optimization

Mengdi Wang*

MENGDW@PRINCETON.EDU

Department of Operations Research and Financial Engineering
Princeton University
Princeton, NJ 08544, USA

Ji Liu*

JI.LIU.UWISC@GMAIL.COM

Department of Computer Science and Department of Electrical and Computer Engineering
University of Rochester
Rochester, NY 14627, USA

Ethan X. Fang

XXF13@PSU.EDU

Department of Statistics and Department of Industrial and Manufacturing Engineering
Pennsylvania State University
University Park, PA 16802, USA

Editor: Inderjit Dhillon

Abstract

We consider the stochastic nested composition optimization problem where the objective is a composition of two expected-value functions. We propose a new stochastic first-order method, namely the accelerated stochastic compositional proximal gradient (ASC-PG) method. This algorithm updates the solution based on noisy gradient queries using a two-timescale iteration. The ASC-PG is the first proximal gradient method for the stochastic composition problem that can deal with nonsmooth regularization penalty. We show that the ASC-PG exhibits faster convergence than the best known algorithms, and that it achieves the optimal sample-error complexity in several important special cases. We demonstrate the application of ASC-PG to reinforcement learning and conduct numerical experiments.

Keywords: Large-scale optimization, stochastic gradient, composition optimization, sample complexity.

1. Introduction

The popular stochastic gradient methods are well suited for minimizing expected-value objective functions or the sum of a large number of loss functions. Stochastic gradient methods find wide applications in estimation, online learning, and training of deep neural networks. Despite their popularity, they do not apply to the minimization of a nonlinear function involving expected values or a composition between two expected-value functions.

In this paper, we consider the *stochastic composition problem*, given by

$$\min_{\mathbf{x} \in X} H(\mathbf{x}) := \underbrace{\mathbb{E}_v(f_v(\mathbb{E}_w(g_w(\mathbf{x}))))}_{=: F(\mathbf{x})} + R(\mathbf{x}) \quad (1)$$

*. Both authors contributed equally.

where $f(g(\mathbf{x})) = (f \circ g)(\mathbf{x})$ denotes the function composition, $g_w(\cdot) : \mathfrak{R}^n \mapsto \mathfrak{R}^m$ and $f_v(\cdot) : \mathfrak{R}^m \mapsto \mathfrak{R}$ are continuously differentiable functions, v, w are random variables, and $R(\mathbf{x}) : \mathfrak{R}^n \mapsto \mathfrak{R} \cup \{+\infty\}$ is an extended real-valued closed convex function, X is a convex and closed set. We assume throughout that there exists at least one optimal solution $\mathbf{x}^* \in X$ to problem (1). We focus on the case where f_v and g_w are smooth, but we allow R to be a nonsmooth penalty such as the ℓ_1 -norm. We do not require either the outer function f_v or the inner function g_w to be convex or monotone. As a result, the composition problem cannot be reformulated into a saddle point problem in general.

Our algorithmic objective is to develop efficient algorithms for solving problem (1) based on random evaluations of f_v, g_w and their gradients. Our theoretical objective is to analyze the rate of convergence for the stochastic algorithm and to improve it when possible. In the online setting, the iteration complexity of our stochastic methods can be interpreted as a sample-error complexity upper bound for estimating the optimal solution of problem (1).

1.1 Motivating Examples

One motivating example is *reinforcement learning* (Sutton and Barto, 1998). Consider a controllable Markov chain with states $1, \dots, S$. Estimating the value-per-state of a fixed control policy π is known as on-policy learning. It can be casted into an $S \times S$ system of Bellman equations:

$$\gamma P^\pi V^\pi + r^\pi = V^\pi,$$

where $\gamma \in (0, 1)$ is a discount factor, P_{ss}^π is the transition probability from state s to state \tilde{s} , and r_s^π is the expected state transition reward at state s . The solution V^π to the Bellman equation is the value vector, with $V^\pi(s)$ being the total expected reward starting at state s . In the blackbox simulation environment, P^π, r^π are unknown but can be sampled from a simulator. As a result, solving the Bellman equation becomes a special case of the stochastic composition optimization problem:

$$\min_{\mathbf{x} \in X} \|\mathbb{E}[A]\mathbf{x} - \mathbb{E}[\mathbf{b}]\|^2, \tag{2}$$

where A, B, b are random matrices and random vectors such that $\mathbb{E}[A] = I - \gamma P^\pi$ and $\mathbb{E}[b] = r^\pi$. It can be viewed as the composition of the square norm function $f(\cdot) \equiv f_v(\cdot) = \|\cdot\|^2$ and the expected linear function $g(\mathbf{x}) = \mathbb{E}[A]\mathbf{x} - \mathbb{E}[\mathbf{b}]$. We will give more details on the reinforcement learning application in Section 4.

Another motivating example is *risk-averse learning*. For example, consider the mean-variance minimization problem

$$\min_{\mathbf{x} \in X} \mathbb{E}_{a,b}[h(\mathbf{x}; a, b)] + \lambda \text{Var}_{a,b}[h(\mathbf{x}; a, b)],$$

where $h(x; a, b)$ is some loss function parameterized by random variables a and b , and $\lambda > 0$ is a regularization parameter. In particular, the mean-variance minimization takes the composition form where $g(\mathbf{x}) = (\mathbb{E}[\mathbf{x}], \mathbb{E}[h(\mathbf{x}; a, b)])$, and $f(y_1, y_2) = \mathbb{E}[h(y_1; a, b)] + \lambda \mathbb{E}(h(y_1; a, b) - y_2)^2$. Its batch version takes the form

$$\min_{\mathbf{x} \in X} \frac{1}{N} \sum_{i=1}^N h(\mathbf{x}; a_i, b_i) + \frac{\lambda}{N} \sum_{i=1}^N \left(h(\mathbf{x}; a_i, b_i) - \frac{1}{N} \sum_{i=1}^N h(\mathbf{x}; a_i, b_i) \right)^2.$$

Here the variance term is the composition of the mean square function and an expected loss function.

Although the stochastic composition problem (1) was barely studied, it actually finds a broad spectrum of emerging applications in estimation and machine learning (see Wang et al. (2017) for a list of applications). Fast optimization algorithms with theoretical guarantees will lead to new computation tools and online learning methods for a broader problem class, no longer limited to the expectation minimization problem.

1.2 Related Works and Contributions

Contrary to the expectation minimization problem, “unbiased” gradient samples are no longer available for the stochastic composition problem (1). The objective is *nonlinear* in the joint probability distribution of (w, v) , which substantially complicates the problem. In a recent work by Dentcheva et al. (2016), a special case of the stochastic composition problem, i.e., risk-averse optimization, has been studied. A central limit theorem has been established, showing that the K -sample batch problem converges to the true problem at the rate of $O(1/\sqrt{K})$ in a proper sense. For the case where $R(x) = 0$, Wang et al. (2017) has proposed and analyzed a class of *stochastic compositional gradient/subgradient methods* (SCGD). The SCGD involves two iterations of different time scales, one for estimating x^* by a stochastic quasi-gradient iteration, the other for maintaining a running estimate of $g(x^*)$. Wang and Liu (2016) studies the SCGD in the setting where samples are corrupted with Markov noises (instead of i.i.d. zero-mean noises). Both works establish almost sure convergence of the algorithm and several convergence rate results, which are the best-known convergence rate prior to the current paper.

The idea of using two-timescale quasi-gradient traced back to the earlier work Ermoliev (1976). It proposed for the first time a stochastic iteration that updates the solution and an auxiliary variable using two different stepsizes. The incremental treatment of proximal gradient iteration has been studied extensively for the expectation minimization problem, see for examples Nedić and Bertsekas (2001); Nemirovski et al. (2009); Wang et al. (2015); Wang and Bertsekas (2016) and references therein. However, except for Wang et al. (2017) and Wang and Liu (2016), all of these works focus on variants of the expectation minimization problem and do not apply to the stochastic composition problem (1). Another work partially related to this paper is by Dai et al. (2017). They consider a special case of problem (1) arising in kernel estimation, where they assume that all functions f_v 's are convex and their conjugate functions f_v^* 's can be easily obtained/sampled. Under these additional assumptions, they essentially rewrite the problem into a saddle point optimization involving functional variables.

In this paper, we propose a new accelerated stochastic compositional proximal gradient (ASC-PG) method that applies to the penalized problem (1), which is a more general problem than the one considered in Wang et al. (2017). We use a *coupled martingale stochastic analysis* to show that ASC-PG achieves significantly better sample-error complexity in various cases. We also show that ASC-PG exhibits optimal sample-error complexity in two important special cases: the case where the outer function is linear and the case where the inner function is linear. Note that the current work has a preliminary conference version (Wang et al., 2016).

The current journal version provides more complete technical details as well as important theoretical extensions (Theorem 3).

Our major contributions are summarized as follows:

1. We propose the first stochastic *proximal-gradient* method for the stochastic composition problem. This is the first algorithm that is able to address the nonsmooth regularization penalty $R(\cdot)$ without deteriorating the convergence rate.
2. We obtain a convergence rate $O(K^{-4/9})$ for smooth optimization problems that are not necessarily convex, where K is the number of queries to the stochastic first-order oracle. This improves the best known convergence rate and provides a new benchmark for the stochastic composition problem.
3. We provide a comprehensive analysis and results that apply to various special cases. In particular, our results contain as special cases the known optimal rate results for the expectation minimization problem, i.e., $O(1/\sqrt{K})$ for general objectives and $O(1/K)$ for strongly convex objectives.
4. In the special case where the inner function $g(\cdot)$ is a linear mapping, we show that it is sufficient to use one timescale to guarantee convergence. Our result achieves the non-improvable rate of convergence $O(1/K)$ for optimal strongly convex optimization and $O(1/\sqrt{K})$ for convex optimization and nonconvex smooth optimization. It implies that the inner linearity does not bring fundamental difficulty to the stochastic composition problem.
5. We show that the proposed method leads to a new on-policy reinforcement learning algorithm. The new learning algorithm achieves the optimal convergence rate $O(1/\sqrt{K})$ for solving Bellman equations based on K observations of state-to-state transitions.

In comparison with Wang et al. (2017), our analysis is more succinct and leads to stronger results. To the best of our knowledge, Theorems 1,2,3 in this paper provide the best-known rates for the stochastic composition problem.

Paper Organization. Section 2 states the sampling oracle and the accelerated stochastic compositional proximal gradient algorithm (ASC-PG). Section 3 states the convergence rate results in the case of general nonconvex objective and in the case of strongly convex objective, respectively. Section 4 describes an application of ASC-PG to reinforcement learning and gives numerical experiments.

Notations and Definitions. For $\mathbf{x} \in \mathbb{R}^n$, we denote by \mathbf{x}' its transpose, and by $\|\mathbf{x}\|$ its Euclidean norm (i.e., $\|\mathbf{x}\| = \sqrt{\mathbf{x}'\mathbf{x}}$). For two sequences $\{\mathbf{y}_k\}$ and $\{\mathbf{z}_k\}$, we write $\mathbf{y}_k = O(\mathbf{z}_k)$ if there exists a constant $c > 0$ such that $\|\mathbf{y}_k\| \leq c\|\mathbf{z}_k\|$ for each k . We denote by $\mathbf{I}_{\text{condition}}^{\text{value}}$ the indicator function, which returns “value” if the “condition” is satisfied; otherwise 0. We denote by H^* the optimal objective function value of problem (1), denote by X^* the set of optimal solutions, and denote by $\mathcal{P}_S(\mathbf{x})$ the Euclidean projection of \mathbf{x} onto S for any convex set S . We also denote by short that $f(\mathbf{y}) = \mathbb{E}_v[f_v(\mathbf{y})]$ and $g(\mathbf{x}) = \mathbb{E}_w[g_w(\mathbf{x})]$.

2. Algorithm

We focus on the black-box sampling environment. Suppose that we have access to a stochastic first-order oracle, which returns random realizations of first-order information upon queries. This is a typical simulation oracle that is available in both online and batch learning. More specifically, assume that we are given a **Sampling Oracle (SO)** such that

- Given some $\mathbf{x} \in \mathfrak{R}^n$, the **SO** returns a random vector $g_w(x)$ and a noisy subgradient $\nabla g_w(\mathbf{x})$.
- Given some $\mathbf{y} \in \mathfrak{R}^m$, the **SO** returns a noisy gradient $\nabla f_v(\mathbf{y})$.

Now we propose the Accelerated Stochastic Compositional Proximal Gradient (ASC-PG) algorithm, see Algorithm 1. ASC-PG is a generalization of the SCGD proposed by Wang et al. (2017), in which a proximal step is used to replace the projection step.

Algorithm 1 Accelerated Stochastic Compositional Proximal Gradient (ASC-PG)

Require: $\mathbf{x}_1 \in \mathfrak{R}^n$, **SO**, K , stepsize sequences $\{\alpha_k\}_{k=1}^K$, and $\{\beta_k\}_{k=1}^K$.

Ensure: $\{\mathbf{x}_k\}_{k=1}^K$

- 1: Initialize $\mathbf{z}_1 = \mathbf{x}_1$ and $\mathbf{y}_1 = \mathbf{0}$.
- 2: **for** $k = 1, \dots, K$ **do**
- 3: Query the **SO** and obtain gradient samples $\nabla f_{v_k}(\mathbf{y}_k)$, $\nabla g_{w_k}(\mathbf{z}_k)$.
- 4: Update the main iterate by

$$\mathbf{x}_{k+1} = \text{prox}_{\alpha_k R(\cdot)} \left(\mathbf{x}_k - \alpha_k \nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(\mathbf{y}_k) \right).$$

- 5: Update auxillary iterates by an *extrapolation-smoothing* scheme:

$$\begin{aligned} \mathbf{z}_{k+1} &= \left(1 - \frac{1}{\beta_k}\right) \mathbf{x}_k + \frac{1}{\beta_k} \mathbf{x}_{k+1}, \\ \mathbf{y}_{k+1} &= (1 - \beta_k) \mathbf{y}_k + \beta_k g_{w_{k+1}}(\mathbf{z}_{k+1}), \end{aligned}$$

where the sample $g_{w_{k+1}}(\mathbf{z}_{k+1})$ is obtained via querying the **SO**.

- 6: **end for**
-

In Algorithm 1, the *extrapolation-smoothing scheme* (i.e., the (\mathbf{y}, \mathbf{z}) -step) is critical to the acceleration of convergence. The acceleration is due to the fast running estimation of the unknown quantity $g(\mathbf{x}_k) := \mathbb{E}_w[g_w(\mathbf{x}_k)]$. At iteration k , the running estimate \mathbf{y}_k of $g(\mathbf{x}_k)$ is obtained using a weighted smoothing scheme, corresponding to the \mathbf{y} -step; while the new query point \mathbf{z}_{k+1} is obtained through extrapolation, corresponding to the \mathbf{z} -step. The updates are constructed in a way such that \mathbf{y}_k is a nearly unbiased estimate of $g(\mathbf{x}_k)$. To see how the extrapolation-smoothing scheme works, we let the weights be

$$\xi_t^{(k)} = \begin{cases} \beta_t \prod_{i=t+1}^k (1 - \beta_i), & \text{if } k > t \geq 0 \\ \beta_k, & \text{if } k = t \geq 0. \end{cases} \quad (3)$$

Then, we can verify the following important relations:

$$\mathbf{x}_{k+1} = \sum_{t=0}^k \xi_t^{(k)} \mathbf{z}_{t+1}, \quad \mathbf{y}_{k+1} = \sum_{t=0}^k \xi_t^{(k)} g_{w_{t+1}}(\mathbf{z}_{t+1}),$$

which essentially say that \mathbf{x}_{k+1} is a damped weighted average of $\{z_{t+1}\}_0^{k+1}$ and \mathbf{y}_{k+1} is a damped weighted average of $\{g_{w_{t+1}}(z_{t+1})\}_0^{k+1}$.

An Analytical Example of the Extrapolation-Smooth Scheme Now consider the special case where $g_w(\cdot)$ is always a linear mapping $g_w(z) = A_w z + b_w$ and $\beta_k = 1/(k+1)$. We can verify that $\xi_t^{(k)} = 1/(k+1)$ for all t . Then we have

$$g(\mathbf{x}_{k+1}) = \frac{1}{k+1} \sum_{t=0}^k \mathbb{E}[A_w] \mathbf{z}_{t+1} + \mathbb{E}[b_w], \quad \mathbf{y}_{k+1} = \frac{1}{k+1} \sum_{t=0}^k A_{w_{t+1}} \mathbf{z}_{t+1} + \frac{1}{k+1} \sum_{t=0}^k \mathbf{b}_{w_{t+1}}.$$

In this way, we can see that the scaled error

$$k(\mathbf{y}_{k+1} - g(\mathbf{x}_{k+1})) = \sum_{t=0}^k (A_{w_{t+1}} - \mathbb{E}[A_w]) \mathbf{z}_{t+1} + \sum_{t=0}^k (\mathbf{b}_{w_{t+1}} - \mathbb{E}[b_w])$$

is a *zero-mean* and *zero-drift* martingale. Under additional technical assumptions, we have

$$\mathbb{E}[\|\mathbf{y}_{k+1} - g(\mathbf{x}_{k+1})\|^2] \leq O\left(\frac{1}{k}\right).$$

Note that the zero-drift property of the error martingale is the key to the fast convergence rate. The zero-drift property comes from the near-unbiasedness of \mathbf{y}_k , which is due to the special construction of the extrapolation-smoothing scheme. In the more general case where g_w is not necessarily linear, we can use a similar argument to show that \mathbf{y}_k is a nearly unbiased estimate of $g(\mathbf{x}_k)$. As a result, the extrapolation-smoothing (\mathbf{y}, \mathbf{z})-step ensures that \mathbf{y}_k tracks the unknown quantity $g(\mathbf{x}_k)$ efficiently.

3. Main Results

We present our main theoretical results in this section. Note that for ease of presentation, we defer all detailed proofs for the theorems, and technical lemmas to Appendix. Let us begin by stating our assumptions. Note that all assumptions involving random realizations of v, w hold with probability 1.

Assumption 1 *The samples generated by the SO are unbiased in the following sense:*

1. $\mathbb{E}_{\{w_k, v_k\}}(\nabla g_{w_k}^\top(\mathbf{x}) \nabla f_{v_k}(\mathbf{y})) = \nabla g^\top(\mathbf{x}) \nabla f(\mathbf{y}) \quad \forall k = 1, 2, \dots, K, \quad \forall \mathbf{x}, \forall \mathbf{y}.$
2. $\mathbb{E}_{w_k}(g_{w_k}(\mathbf{x})) = g(\mathbf{x}) \quad \forall \mathbf{x}.$

Note that w_k and v_k are not necessarily independent.

Assumption 2 *The sample gradients and values generated by the SO satisfy*

$$\mathbb{E}_w(\|g_w(\mathbf{x}) - g(\mathbf{x})\|^2) \leq \sigma^2 \quad \forall \mathbf{x}.$$

Assumption 3 *The sample gradients generated by the SO are uniformly bounded, and the penalty function R has bounded gradients.*

$$\|\nabla f_v(\mathbf{x})\| \leq \Theta(1), \quad \|\nabla g_w(\mathbf{x})\| \leq \Theta(1), \quad \|\partial R(\mathbf{x})\| \leq \Theta(1) \quad \forall \mathbf{x}, \forall w, \forall v$$

Assumption 4 *There exists $L_F, L_f, L_g > 0$ such that*

1. $F(\mathbf{z}) - F(\mathbf{x}) \leq \langle \nabla F(\mathbf{x}), \mathbf{z} - \mathbf{x} \rangle + \frac{L_F}{2} \|\mathbf{z} - \mathbf{x}\|^2 \quad \forall \mathbf{x} \quad \forall \mathbf{z}.$
2. $\|\nabla f_v(\mathbf{y}) - \nabla f_v(\mathbf{w})\| \leq L_f \|\mathbf{y} - \mathbf{w}\| \quad \forall \mathbf{y} \quad \forall \mathbf{w} \quad \forall v.$
3. $\|g(\mathbf{x}) - g(\mathbf{z}) - \nabla g(\mathbf{z})^\top (\mathbf{x} - \mathbf{z})\| \leq \frac{L_g}{2} \|\mathbf{x} - \mathbf{z}\|^2 \quad \forall \mathbf{x} \quad \forall \mathbf{z}.$

Condition 1 of Assumption 4 requires that the function $F(\cdot)$ is L_F -smooth. Condition 2 requires that f_v has Lipschitz gradient for all v . Condition 3 requires that the function $g(\cdot)$ is L_g smooth. Note that the case where $L_f = 0$ or $L_g = 0$ coincides with the special case where either f_v or g is linear, respectively. The constants L_F, L_f, L_g jointly characterize the smoothness and complexity of stochastic composition optimization. They do not admit any straightforward dependence relation.

Our first main result concerns with general optimization problems which are not necessarily convex.

Theorem 1 (Smooth (Nonconvex) Optimization) *Let Assumptions 1, 2, 3, and 4 hold. Denote by $F(\mathbf{x}) := (\mathbb{E}_v(f_v) \circ \mathbb{E}_w(g_w))(\mathbf{x})$ for short and suppose that $R(\mathbf{x}) = 0$ in (1) and $\mathbb{E}(F(\mathbf{x}_k))$ is bounded from above. Choose $\alpha_k = k^{-a}$ and $\beta_k = 2k^{-b}$ where $a \in (0, 1)$ and $b \in (0, 1)$ in Algorithm 1. Then we have*

$$\frac{\sum_{k=1}^K \mathbb{E}(\|\nabla F(\mathbf{x}_k)\|^2)}{K} \leq O(K^{a-1} + L_f^2 L_g K^{4b-4a} \mathbf{I}_{4a-4b=1}^{\log K} + L_f^2 K^{-b} + K^{-a}). \quad (4)$$

If $L_g \neq 0$ and $L_f \neq 0$, choose $a = 5/9$ and $b = 4/9$, yielding

$$\frac{\sum_{k=1}^K \mathbb{E}(\|\nabla F(\mathbf{x}_k)\|^2)}{K} \leq O(K^{-4/9}). \quad (5)$$

If $L_g = 0$ or $L_f = 0$, choose $a = b = 1/2$, yielding

$$\frac{\sum_{k=1}^K \mathbb{E}(\|\nabla F(\mathbf{x}_k)\|^2)}{K} \leq O(K^{-1/2}). \quad (6)$$

The result of Theorem 1 strictly improves the best-known results given by Wang et al. (2017). First the result of (5) improves the finite-sample error bound from $O(k^{-2/7})$ to $O(k^{-4/9})$ for general convex and nonconvex optimization. This improves the best known convergence rate and provides a new benchmark for the stochastic composition problem. Note that it is possible to relax the condition “ $\mathbb{E}(F(\mathbf{x}_k))$ is bounded from above” in Theorem 1. However, it would make the analysis more cumbersome and yield an additional term $\log K$ in the error bound.

Our second main result concerns strongly convex objective functions. We say that the objective function H is *optimally strongly convex* with parameter $\lambda > 0$ if

$$H(\mathbf{x}) - H(\mathcal{P}_{X^*}(\mathbf{x})) \geq \lambda \|\mathbf{x} - \mathcal{P}_{X^*}(\mathbf{x})\|^2 \quad \forall \mathbf{x}. \quad (7)$$

(see Liu and Wright (2015)). Note that any strongly convex function is optimally strongly convex, but the reverse does not hold. For example, the objective function (2) in on-policy reinforcement learning is always optimally strongly convex (even if $\mathbb{E}(A)$ is a rank deficient matrix), but not necessarily strongly convex.

We point out that there are other class of functions that are not strongly convex, for which first-order algorithms still achieve linear convergence. See Karimi et al. (2016) and Necoara et al. (2015) for examples. For ease of presentation, in this work, we only consider optimally strongly convex functions.

Theorem 2 (*Strongly Convex Optimization*) *Suppose that the objective function $H(\mathbf{x})$ in (1) is optimally strongly convex with parameter $\lambda > 0$ defined in (7). Set $\alpha_k = C_a k^{-a}$ and $\beta_k = C_b k^{-b}$ where $C_a > 4\lambda$, $C_b > 2$, $a \in (0, 1]$, and $b \in (0, 1]$ in Algorithm 1. Under Assumptions 1, 2, 3, and 4, we have*

$$\mathbb{E}(\|\mathbf{x}_K - \mathcal{P}_{X^*}(\mathbf{x}_K)\|^2) \leq O\left(K^{-a} + L_f^2 L_g K^{-4a+4b} + L_f^2 K^{-b}\right). \quad (8)$$

If $L_g \neq 0$ and $L_f \neq 0$, choose $a = 1$ and $b = 4/5$, yielding

$$\mathbb{E}(\|\mathbf{x}_K - \mathcal{P}_{X^*}(\mathbf{x}_K)\|^2) \leq O(K^{-4/5}). \quad (9)$$

If $L_g = 0$ or $L_f = 0$, choose $a = 1$ and $b = 1$, yielding

$$\mathbb{E}(\|\mathbf{x}_K - \mathcal{P}_{X^*}(\mathbf{x}_K)\|^2) \leq O(K^{-1}). \quad (10)$$

Let us discuss the results of Theorem 2. In the general case where $L_f \neq 0$ and $L_g \neq 0$, the convergence rate in (9) is consistent with the result of Wang et al. (2017). Now consider the special case where $L_g = 0$, i.e., the inner mapping is linear. This result finds an immediate application to Bellman error minimization problem (2) which arises from reinforcement learning problem in (and with ℓ_1 norm regularization). The proposed ASC-PG algorithm is able to achieve the optimal rate $O(1/K)$ without any extra assumption on the outer function f_v . To the best of our knowledge, this is the best (also optimal) sample-error complexity for on-policy reinforcement learning.

Theorem 3 (*General Convex Optimization*) *Suppose that both $F(\mathbf{x})$ and $H(\mathbf{x})$ are convex, and the feasible set X is bounded. Set $\alpha_k = \Theta(k^{-a})$ and $\beta_k = \Theta(k^{-b})$ with $a \in (0, 1)$, and $b \in (0, 1]$ in Algorithm 1. Under Assumptions 1, 2, 3, and 4, we have for any $c > 0$:*

$$\begin{aligned} H(\bar{\mathbf{x}}_K) - H^* \leq O\left(& K^{a-1} + K^{a-c} \mathbf{I}_{c=1}^{\log K} + K^{-a} \right. \\ & \left. + L_f^2 L_g K^{-5a+4b+c} \mathbf{I}_{6a-4b-c=1}^{\log K} + L_f^2 K^{-a-b+c} \mathbf{I}_{2a+b-c=1}^{\log K} \right). \end{aligned} \quad (11)$$

where $\bar{\mathbf{x}}_K := \frac{\sum_{k=2}^{K+1} \alpha_k \mathbf{x}_k}{\sum_{k=2}^{K+1} \alpha_k}$. If $L_g \neq 0$ and $L_f \neq 0$, choose $a = 5/7$, $b = 4/7$, $c = 1$, yielding

$$H(\bar{\mathbf{x}}_K) - H^* \leq O(K^{-2/7}). \quad (12)$$

If $L_g = 0$ or $L_f = 0$, choose $a = 1/2$, $b = 1$, and $c = 1$, yielding

$$H(\bar{\mathbf{x}}_K) - H^* \leq O(K^{-1/2} \log K). \quad (13)$$

Theorem 3 gives stronger results than the best-known results of Wang et al. (2017) and requires milder assumptions. In the general case where $L_f \neq 0$ and $L_g \neq 0$, the convergence rate in (12) matches the result of Wang et al. (2017). When either $L_g = 0$ (i.e., the inner mapping is linear) or $L_f = 0$ (i.e., the outer mapping is linear), the proposed ASC-PG algorithm is able to achieve the optimal rate $O(1/\sqrt{K})$ up to a logarithmic factor.

Remarks Theorems 1, 2, and 3 give important implications about the special cases where $L_f = 0$ or $L_g = 0$. In these cases, we argue that our convergence rate (10) is “*optimal*” with respect to the sample size K . To see this, it is worth pointing out the $O(1/K)$ rate of convergence is optimal for strongly convex expectation minimization problem. Because the expectation minimization problem is a special case of the problem (1), the $O(1/K)$ convergence rate must be optimal for the stochastic composition problem too. Similarly, the $O(1/\sqrt{K})$ rate of convergence is also optimal for general convex optimization due to the same argument.

- Consider the case where $L_f = 0$, which means that the outer function $f_v(\cdot)$ is linear with probability 1. Then the stochastic composition problem (1) reduces to an expectation minimization problem since $(\mathbb{E}_v f_v \circ \mathbb{E}_w g_w)(\mathbf{x}) = \mathbb{E}_v(f_v(\mathbb{E}_w g_w(\mathbf{x}))) = \mathbb{E}_v \mathbb{E}_w(f_v \circ g_w)(\mathbf{x})$. Therefore, it makes a perfect sense to obtain the optimal convergence rate.
- Consider the case where $L_g = 0$, which means that the inner function $g(\cdot)$ is a linear mapping. The result is quite surprising. Note that even $g(\cdot)$ is a linear mapping, it does not reduce problem (1) to an expectation minimization problem. However, the ASC-PG still achieves the optimal convergence rate. This suggests that, when inner linearity holds, the stochastic composition problem (1) is not fundamentally more difficult than the expectation minimization problem.

The convergence rate results unveiled in Theorems 1 and 2 are the best known results for the composition problem. We believe that they provide important new result which provides insights into the complexity of the stochastic composition problem.

4. Application to Reinforcement Learning

In this section, we apply the proposed ASC-PG algorithm to conduct policy value evaluation in reinforcement learning through attacking Bellman equations. Suppose that there are in total S states. Let the policy of interest be π . Denote the value function of states by $V^\pi \in \mathbb{R}^S$, where $V^\pi(s)$ denotes the value of being at state s under policy π . The Bellman equation of the problem is

$$V^\pi(s_1) = \mathbb{E}_\pi\{r_{s_1, s_2} + \gamma \cdot V^\pi(s_2) | s_1\} \text{ for all } s_1, s_2 \in \{1, \dots, S\},$$

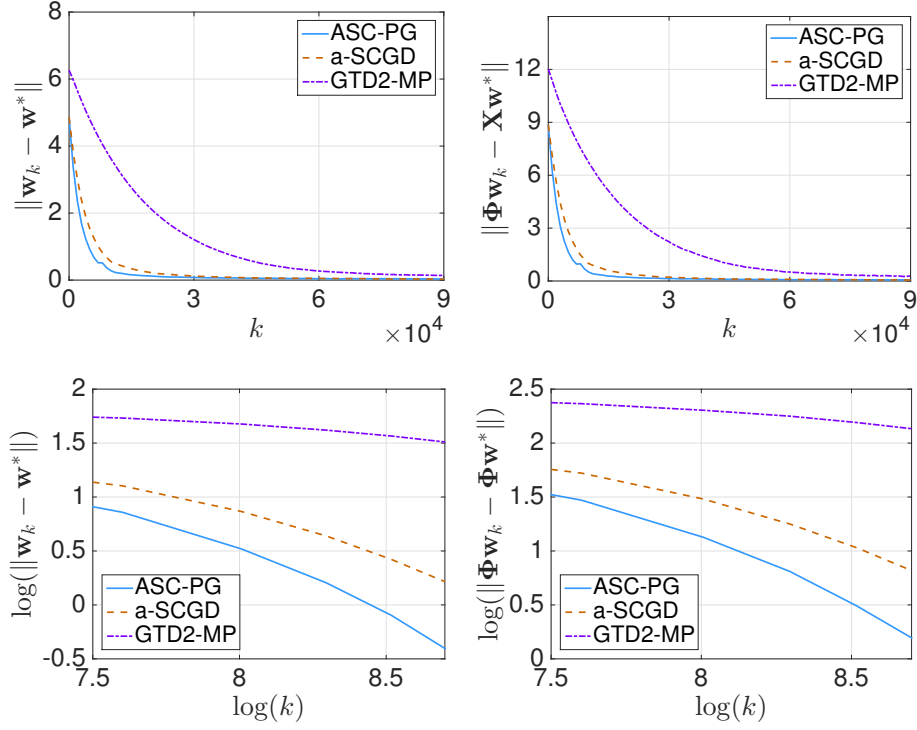


Figure 1: Empirical convergence rate of the ASC-PG algorithm and the GTD2-MP algorithm under Experiment 1 averaged over 100 runs, where \mathbf{w}_k denotes the solution at the k -th iteration.

where r_{s_1, s_2} denotes the reward of moving from state s_1 to s_2 , and the expectation is taken over all possible future state s_2 conditioned on current state s_1 and the policy π . We have that the solution $V^* \in \mathbb{R}^S$ to the above equation satisfies that $V^* = V^\pi$. Here a moderately large S will make solving the Bellman equation directly impractical. To resolve the curse of dimensionality, in many practical applications, we approximate the value of each state by some linear map of its feature $\phi_s \in \mathbb{R}^d$, where $d < S$ to reduce the dimension. In particular, we assume that $V^\pi(s) \approx \phi_s^T \mathbf{w}^*$ for some $\mathbf{w}^* \in \mathbb{R}^m$.

To compute \mathbf{w}^* , we formulate the problem as a Bellman residual minimization problem that

$$\min_{\mathbf{w}} \sum_{s=1}^S (\phi_s^T \mathbf{w} - q_{\pi, s}(\mathbf{w}))^2,$$

where $q_{\pi, s}(\mathbf{w}) = \mathbb{E}_\pi \{r_{s, s'} + \gamma \cdot \phi_{s'}^T \mathbf{w}\} = \sum_{s'} P_{ss'}^\pi (\{r_{s, s'} + \gamma \cdot \phi_{s'}^T \mathbf{w}\})$; $\gamma < 1$ is a discount factor, and $r_{s, s'}$ is the random reward of transition from state s to state s' . It is clearly seen that the proposed ASC-PG algorithm could be directly applied to solve this problem where we take

$$g(\mathbf{w}) = (\phi_1^T \mathbf{w}, q_{\pi, 1}(\mathbf{w}), \dots, \phi_S^T \mathbf{w}, q_{\pi, S}(\mathbf{w})) \in \mathbb{R}^{2S},$$

$$f\left((\phi_1^T \mathbf{w}, q_{\pi, 1}(\mathbf{w}), \dots, \phi_S^T \mathbf{w}, q_{\pi, S}(\mathbf{w}))\right) = \sum_{s=1}^S (\phi_s \mathbf{w} - q_{\pi, s}(\mathbf{w}))^2 \in \mathbb{R}.$$

We point out that the $g(\cdot)$ function here is a linear map. By our theoretical analysis, we expect to achieve a faster $O(1/k)$ rate of convergence, which is justified empirically in our later simulation study.

We consider three experiments, where in the first two experiments, we compare our proposed accelerated ASC-PG algorithm with accelerated SCGD (a-SCGD) algorithm (Wang et al., 2017) and the recently proposed GTD2-MP algorithm (Liu et al., 2015). Also, in the first two experiments, we do not add any regularization term, i.e. $R(\cdot) = 0$. We point out here that in experiments 1 and 2, since there is no penalization term, i.e., $R(\cdot) = 0$, the ASC-PG algorithm essentially matches the a-SCGD algorithm in Wang et al. (2017). However, note that the choice of stepsizes $\{\alpha_k\}$ and $\{\beta_k\}$ of ASC-PG are different from the stepsizes of a-SCGD, where for ASC-PG $\alpha_k = k^{-1}$ and $\beta_k = k^{-1}$, and for a-SCGD $\alpha_k = k^{-1}$ and $\beta_k = k^{-4/5}$. In the third experiment, we add an ℓ_1 -penalization term $\lambda\|\mathbf{w}\|_1$. In all cases, we choose the step sizes via comparison studies as in Dann et al. (2014):

- Experiment 1: We use the Baird’s example (Baird, 1995), which is a well-known example to test the off-policy convergent algorithms. This example contains $S = 6$ states, and two actions at each state. We refer the readers to Baird (1995) for more detailed information of the example.
- Experiment 2: We generate a Markov decision problem (MDP) using similar setup as in White and White (2016). In each instance, we randomly generate an MDP which contains $S = 100$ states, and three actions at each state. The dimension of the Given one state and one action, the agent can move to one of four next possible states. In our simulation, we generate the transition probabilities for each MDP instance uniformly from $[0, 1]$ and normalize the sum of transitions to one, and we generate the reward for each transition also uniformly in $[0, 1]$.
- Experiment 3: We generate the data same as Experiment 2 except that we have a larger $d = 100$ dimensional feature space, where only the first 4 components of \mathbf{w}^* are non-zeros. We add an ℓ_1 -regularization term, $\lambda\|\mathbf{w}\|_1$, to the objective function.

Denote by \mathbf{w}_k the solution at the k -th iteration. For the first two experiments, we report the empirical convergence performance $\|\mathbf{w}_k - \mathbf{w}^*\|$ and $\|\Phi\mathbf{w}_k - \Phi\mathbf{w}^*\|$, where $\Phi = (\phi_1, \dots, \phi_S)^T \in \mathfrak{R}^{S \times d}$ and $\Phi\mathbf{w}^* = V$, and all \mathbf{w}_k ’s are averaged over 100 runs, in the first two subfigures of Figures 1 and 2. It is seen that the ASC-PG algorithm achieves the fastest convergence rate empirically in both experiments. To further evaluate our theoretical results, we plot $\log(t)$ vs. $\log(\|\mathbf{w}_k - \mathbf{w}^*\|)$ (or $\log(\|\Phi\mathbf{w}_k - \Phi\mathbf{w}^*\|)$) averaged over 100 runs for the first two experiments in the second two subfigures of Figures 1 and 2. The empirical results further support our theoretical analysis that $\|\mathbf{w}_k - \mathbf{w}^*\|^2 = O(1/k)$ for the ASC-PG algorithm when $g(\cdot)$ is a linear mapping.

For Experiment 3, as the optimal solution is unknown, we run the ASC-PG algorithm for one million iterations and take the corresponding solution as the optimal solution $\hat{\mathbf{w}}^*$, and we report $\|\mathbf{w}_k - \hat{\mathbf{w}}^*\|$ and $\|\Phi\mathbf{w}_k - \Phi\hat{\mathbf{w}}^*\|$ averaged over 100 runs in Figure 3. It is seen the the ASC-PG algorithm achieves fast empirical convergence rate.




Figure 2: Empirical convergence rate of the ASC-PG algorithm and the GTD2-MP algorithm under Experiment 2 averaged over 100 runs, where \mathbf{w}_k denotes the solution at the k -th iteration.




Figure 3: Empirical convergence rate of the ASC-PG algorithm with the ℓ_1 -regularization term $\lambda \|\mathbf{w}\|_1$ under Experiment 3 averaged over 100 runs, where \mathbf{w}_k denotes the solution at the t -th iteration.

5. Conclusion

We develop the first proximal gradient method for stochastic composition optimization with nonsmooth penalties. The algorithm updates by interacting with a stochastic first-order

oracle. Finite-sample convergence rates are established under a variety of assumptions, which provide new rate benchmarks that improve the best-known results prior to this paper. Application of the ASC-PG to reinforcement learning leads to a new on-policy learning algorithm, which achieves faster convergence than the best known algorithms. For future research, it remains open whether or under what circumstances the current $O(K^{-4/9})$ can be further improved. Another direction is to customize and adapt the algorithm and analysis to more specific problems arising from reinforcement learning and risk-averse optimization, in order to fully exploit the potential of the proposed method.

Acknowledgments

We thank the Editor, the Action Editor and two anonymous reviewers for the careful consideration and invaluable suggestions, which help us further improve the paper. This project is in part supported by NSF grants CCF1718513, CNS-1548078, DMS-10009141, CMMI-1653435, and by the National Institute on Drug Abuse grant P50 DA039838. The content of this manuscript is solely the responsibility of the authors and does not necessarily represent the official views of the National Institute on Drug Abuse or the National Institutes of Health.

Appendix

This appendix includes all detailed proofs for the theorems in the main text. The key idea of deriving the rate convergence of our algorithm is to find a sequence of random vectors associated with the output solutions at each iteration, which converge at fast rates. To better analyze the convergence process, we have constructed two auxiliary sequences of random variables $\{m_k\}$ and $\{n_k\}$, which are defined in the beginning of Lemma 7, that converge jointly to 0 at a fast rate. The convergence of $\{m_k\}$ and $\{n_k\}$ further imply that $\{\mathbf{x}_k\}$ and $\{\mathbf{y}_k\}$ converge rapidly to their limits respectively. Then we are able to improve the convergence rate of \mathbf{y}_k . Intuitively, we show that the improved iterate \mathbf{y}_k tracks the unknown quantity $\mathbb{E}[g(x_k)]$ more closely than in the previous work. This explains why our new algorithm achieves an improved rate of convergence and sample complexity.

A. Proof to Theorem 1

In this section, we provide the detailed proof for Theorem 1. We start by some technical lemmas, The first one provides a convergence result of the generated solutions, where we bound the term $\|\mathbf{x}_k - \mathbf{x}_{k+1}\|^2$ as k increases.

Lemma 4 *Under Assumption 3, two subsequent iterates in Algorithm 1 satisfy*

$$\|\mathbf{x}_k - \mathbf{x}_{k+1}\|^2 \leq \Theta(\alpha_k^2).$$

Proof From the definition of the proximal operation, we have

$$\mathbf{x}_{k+1} = \text{prox}_{\alpha_k R(\cdot)}(\mathbf{x}_k - \alpha_k \nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(\mathbf{y}_k))$$

$$= \arg \min_x \frac{1}{2} \|\mathbf{x} - \mathbf{x}_k + \alpha_k \nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(\mathbf{y}_k)\|^2 + \alpha_k R(\mathbf{x}).$$

The optimality condition suggests the following equality:

$$\mathbf{x}_{k+1} - \mathbf{x}_k = -\alpha_k (\nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(\mathbf{y}_k) + \mathbf{s}_{k+1}) \quad (14)$$

where $\mathbf{s}_{k+1} \in \partial R(\mathbf{x}_{k+1})$ is some vector in the sub-differential set of $R(\cdot)$ at \mathbf{x}_{k+1} . Then apply the boundedness condition in Assumption 3 to yield

$$\begin{aligned} \|\mathbf{x}_{k+1} - \mathbf{x}_k\| &= \alpha_k \|(\nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(\mathbf{y}_k) + \mathbf{s}_{k+1})\| \\ &\leq \alpha_k (\|\nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(\mathbf{y}_k)\| + \|\mathbf{s}_{k+1}\|) \\ &\leq \alpha_k (\|\nabla g_{w_k}^\top(\mathbf{x}_k)\| \|\nabla f_{v_k}(\mathbf{y}_k)\| + \|\mathbf{s}_{k+1}\|) \\ &\stackrel{\text{(Assumption 3)}}{\leq} \Theta(1) \alpha_k, \end{aligned}$$

which implies the claim. \blacksquare

Next, we provide a simple bound which quantifies the difference between the random variable $\|\nabla g_w^\top(\mathbf{x}) \nabla_v f(g(\mathbf{x})) - \nabla g_w^\top(\mathbf{x}) \nabla_v f(\mathbf{y})\|$ and its population version $\|\mathbf{y} - g(\mathbf{x})\|$.

Lemma 5 *Under Assumptions 3 and 4, we have*

$$\|\nabla g_w^\top(\mathbf{x}) \nabla_v f(g(\mathbf{x})) - \nabla g_w^\top(\mathbf{x}) \nabla_v f(\mathbf{y})\| \leq \Theta(L_f \|\mathbf{y} - g(\mathbf{x})\|).$$

Proof We have

$$\begin{aligned} \|\nabla g_w^\top(\mathbf{x}) \nabla_v f(g(\mathbf{x})) - \nabla g_w^\top(\mathbf{x}) \nabla_v f(\mathbf{y})\| &\leq \|\nabla g_w^\top(\mathbf{x})\| \|\nabla f_v(g(\mathbf{x})) - \nabla f_v(\mathbf{y})\| \\ &\stackrel{\text{(Assumption 3)}}{\leq} \Theta(1) \|\nabla f_v(g(\mathbf{x})) - \nabla f_v(\mathbf{y})\| \\ &\stackrel{\text{(Assumption 4)}}{\leq} \Theta(L_f) \|\mathbf{y} - g(\mathbf{x})\|. \end{aligned}$$

It completes the proof. \blacksquare

Then, we prove two lemmas which provides the convergence rate of $\mathbb{E}(\|\mathbf{y}_k - g(\mathbf{x}_k)\|^2)$.

Lemma 6 *Given two sequences of positive scalars $\{s_k\}_{k=1}^\infty$ and $\{\gamma_k\}_{k=1}^\infty$ satisfying*

$$s_{k+1} \leq (1 - \gamma_k + C_1 \gamma_k^2) s_k + C_2 k^{-a} \quad (15)$$

where $C_1 \geq 0$, $C_2 \geq 0$, and $a \geq 0$. Letting $\gamma_k \in \mathfrak{R}$ as $\gamma_k = C_3 k^{-b}$ where $b \in (0, 1]$ and $C_3 > 2$, the sequence can be bounded by

$$v_k \leq C k^{-c}$$

where C and c are defined as

$$C := \max_{k \leq (C_1 C_3^2)^{1/b+1}} s_k k^c + \frac{C_2}{C_3 - 2} \quad \text{and} \quad c := a - b.$$

In other words, we have

$$s_k \leq \Theta(k^{-a+b}).$$

Proof We prove it by induction. First it is easy to verify that the claim holds for $k \leq (C_1 C_3^2)^{1/b}$ from the definition for C . Next we prove from “ k ” to “ $k + 1$ ”, that is, given $s_k \leq Ck^{-c}$ for $k > (C_1 C_3^2)^{1/b}$, we need to prove $s_{k+1} \leq C(k+1)^{-c}$.

$$\begin{aligned} s_{k+1} &\leq (1 - \gamma_k + C_1 \gamma_k^2) s_k + C_2 k^{-a} \\ &\leq (1 - C_3 k^{-b} + C_1 C_3^2 k^{-2b}) C k^{-c} + C_2 k^{-a} \\ &= C k^{-c} - C C_3 k^{-b-c} + C C_1 C_3^2 k^{-2b-c} + C_2 k^{-a}. \end{aligned} \quad (16)$$

To prove that (16) is bounded by $C(k+1)^{-c}$, it suffices to show that

$$\Delta := (k+1)^{-c} - k^{-c} + C_3 k^{-b-c} - C_1 C_3^2 k^{-2b-c} > 0 \quad \text{and} \quad C \geq \frac{C_2 k^{-a}}{\Delta}.$$

From the convexity of function $h(t) = t^{-c}$, we have the inequality $(k+1)^{-c} - k^{-c} \geq (-c)k^{-c-1}$. Therefore we obtain

$$\begin{aligned} \Delta &\geq -ck^{-c-1} + C_3 k^{-b-c} - C_1 C_3^2 k^{-2b-c} \\ &\stackrel{(b \leq 1, k > (C_1 C_3^2)^{1/b})}{\geq} (C_3 - 2)(k^{-b-c}) \\ &\stackrel{(C_3 > 2)}{>} 0. \end{aligned}$$

To verify the second one, we have

$$\frac{C_2 k^{-a}}{\Delta} \leq \frac{C_2}{C_3 - 2} k^{-a+b+c} \stackrel{(c=a+b)}{=} \frac{C_2}{C_3 - 2} \leq C$$

where the last inequality is due to the definition of C . It completes the proof. \blacksquare

Lemma 7 Choose β_k to be $\beta_k = C_b k^{-b}$ where $C_b > 2$, $b \in (0, 1]$, and $\alpha_k = C_a k^{-a}$. Under Assumptions 1 and 2, we have

$$\mathbb{E}(\|\mathbf{y}_k - g(\mathbf{x}_k)\|^2) \leq L_g \Theta(k^{-4a+4b}) + \Theta(k^{-b}). \quad (17)$$

Proof Denote by m_{k+1}

$$m_{k+1} := \sum_{t=0}^k \xi_t^{(k)} \|\mathbf{x}_{k+1} - \mathbf{z}_{t+1}\|^2$$

and n_{k+1}

$$n_{k+1} := \left\| \sum_{t=0}^k \xi_t^{(k)} (g_{w_{t+1}}(\mathbf{z}_{t+1}) - g(\mathbf{z}_{t+1})) \right\|$$

for short.

From Lemma 10 in (Wang et al., 2017), we have

$$\|\mathbf{y}_k - g(\mathbf{x}_k)\|^2 \leq \left(\frac{L_g}{2} m_k + n_k \right)^2 \leq L_g m_k^2 + 2n_k^2. \quad (18)$$

From Lemma 11 in (Wang et al., 2017), m_{k+1} can be bounded by

$$m_{k+1} \leq (1 - \beta_k)m_k + \beta_k q_k + \frac{2}{\beta_k} \|\mathbf{x}_k - \mathbf{x}_{k+1}\|^2 \quad (19)$$

where q_k is bounded by

$$\begin{aligned} q_{k+1} &\leq (1 - \beta_k)q_k + \frac{4}{\beta_k} \|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2 \\ &\stackrel{\text{(Lemma 4)}}{\leq} (1 - \beta_k)q_k + \frac{\Theta(1)\alpha_k^2}{\beta_k} \\ &\leq (1 - \beta_k)q_k + \Theta(k^{-2a+b}). \end{aligned}$$

Taking $s_k = q_k$, and $\gamma_k = \beta_k$, applying Lemma 6 implies the following decay rate

$$q_k \leq \Theta(k^{-2a+b+b}) = \Theta(k^{-2a+2b}).$$

Together with (19), we have

$$\begin{aligned} m_{k+1} &\leq (1 - \beta_k)m_k + \beta_k q_k + \frac{2}{\beta_k} \|\mathbf{x}_k - \mathbf{x}_{k+1}\|^2 \\ &\leq (1 - \beta_k)m_k + \Theta(k^{-2a+b}) + \Theta(k^{-2a+b}) \\ &\leq (1 - \beta_k)m_k + \Theta(k^{-2a+b}), \end{aligned}$$

which leads to

$$m_k \leq \Theta(k^{-2a+2b}) \quad \text{and} \quad m_k^2 \leq \Theta(k^{-4a+4b}). \quad (20)$$

by using Lemma 6 again. Then we estimate the upper bound for $\mathbb{E}(n_k^2)$. From Lemma 11 in (Wang et al., 2017), we know $\mathbb{E}(n_k^2)$ is bounded by

$$\mathbb{E}(n_{k+1}^2) \leq (1 - \beta_k)^2 \mathbb{E}(\|n_k\|^2) + \beta_k^2 \sigma_g^2 = (1 - 2\beta_k + \beta_k^2) \mathbb{E}(\|n_k\|^2) + \beta_k^2 \sigma_g^2.$$

By using Lemma 6 again, we have

$$\mathbb{E}(n_k^2) \leq \Theta(k^{-b}). \quad (21)$$

Now we are ready to estimate the upper bound of $\|\mathbf{y}_{k+1} - g(\mathbf{x}_{k+1})\|^2$ by following (18)

$$\begin{aligned} \mathbb{E}(\|\mathbf{y}_k - g(\mathbf{x}_k)\|^2) &\leq L_g \mathbb{E}(m_k^2) + 2\mathbb{E}(n_k^2) \\ &\stackrel{(20)+(21)}{\leq} L_g \Theta(k^{-4a+4b}) + \Theta(k^{-b}). \end{aligned}$$

It completes the proof. ■

Now we are ready to prove Theorem 1.

Proof to Theorem 1. From the Lipschitzian condition in Assumption 4, we have

$$F(\mathbf{x}_{k+1}) - F(\mathbf{x}_k)$$

$$\begin{aligned}
 &\leq \langle \nabla F(\mathbf{x}_k), \mathbf{x}_{k+1} - \mathbf{x}_k \rangle + \frac{L_F}{2} \|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2 \\
 \text{(Lemma 6)} \quad &\leq -\alpha_k \langle \nabla F(\mathbf{x}_k), \nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(\mathbf{y}_k) \rangle + \Theta(\alpha_k^2) \\
 &= -\alpha_k \|\nabla F(\mathbf{x}_k)\|^2 + \underbrace{\alpha_k \langle \nabla F(\mathbf{x}_k), \nabla F(\mathbf{x}_k) - \nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(\mathbf{y}_k) \rangle}_{=:T} \\
 &\quad + \Theta(\alpha_k^2) \tag{22}
 \end{aligned}$$

Next we estimate the upper bound for $\mathbb{E}(T)$:

$$\begin{aligned}
 \mathbb{E}(T) &= \mathbb{E}(\langle \nabla F(\mathbf{x}_k), \nabla F(\mathbf{x}_k) - \nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(g(\mathbf{x}_k)) \rangle) \\
 &\quad + \mathbb{E}(\langle \nabla F(\mathbf{x}_k), \nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(g(\mathbf{x}_k)) - \nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(\mathbf{y}_k) \rangle) \\
 \text{(Assumption 1)} \quad &\stackrel{=}{=} \mathbb{E}(\langle \nabla F(\mathbf{x}_k), \nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(g(\mathbf{x}_k)) - \nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(\mathbf{y}_k) \rangle) \\
 &\leq \frac{1}{2} \mathbb{E}(\|\nabla F(\mathbf{x}_k)\|^2) + \frac{1}{2} \mathbb{E}(\|\nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(g(\mathbf{x}_k)) - \nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(\mathbf{y}_k)\|^2) \\
 \text{(Lemma 5)} \quad &\leq \frac{1}{2} \mathbb{E}(\|\nabla F(\mathbf{x}_k)\|^2) + \Theta(L_f^2) \mathbb{E}(\|\mathbf{y}_k - g(\mathbf{x}_k)\|^2).
 \end{aligned}$$

Take expectation on both sides of (22) and substitute $\mathbb{E}(T)$ by its upper bound:

$$\begin{aligned}
 &\frac{\alpha_k}{2} \|\nabla F(\mathbf{x}_k)\|^2 \\
 &\leq \mathbb{E}(F(\mathbf{x}_k)) - \mathbb{E}(F(\mathbf{x}_{k+1})) + \Theta(L_f^2 \alpha_k) \mathbb{E}(\|\mathbf{y}_k - g(\mathbf{x}_k)\|^2) + \Theta(\alpha_k^2) \\
 \text{(Lemma 7)} \quad &\leq \mathbb{E}(F(\mathbf{x}_k)) - \mathbb{E}(F(\mathbf{x}_{k+1})) + L_g \Theta(L_f^2 \alpha_k) \Theta(k^{-4a+4b}) + \Theta(L_f^2 \alpha_k k^{-b}) + \Theta(\alpha_k^2) \\
 &\leq \mathbb{E}(F(\mathbf{x}_k)) - \mathbb{E}(F(\mathbf{x}_{k+1})) + L_f^2 L_g \Theta(k^{-5a+4b}) + L_f^2 \Theta(k^{-a-b}) + \Theta(k^{-2a})
 \end{aligned}$$

which suggests that

$$\begin{aligned}
 &\mathbb{E}(\|\nabla F(\mathbf{x}_k)\|^2) \\
 &\leq 2\alpha_k^{-1} \mathbb{E}(F(\mathbf{x}_k)) - 2\alpha_k^{-1} \mathbb{E}(F(\mathbf{x}_{k+1})) + L_f^2 L_g \Theta(k^{-4a+4b}) + L_f^2 \Theta(k^{-b}) + \Theta(k^{-a}) \\
 &\leq 2k^a \mathbb{E}(F(\mathbf{x}_k)) - 2k^a \mathbb{E}(F(\mathbf{x}_{k+1})) + L_f^2 L_g \Theta(k^{-4a+4b}) + L_f^2 \Theta(k^{-b}) + \Theta(k^{-a}). \tag{23}
 \end{aligned}$$

Sum Eq. (23) from $k = 1$ to K and obtain

$$\begin{aligned}
 \frac{\sum_{k=1}^K \mathbb{E}(\|\nabla F(\mathbf{x}_k)\|^2)}{K} &\leq 2K^{-1} \alpha_1^{-1} F(\mathbf{x}_1) + K^{-1} \sum_{k=2}^K ((k+1)^a - k^a) \mathbb{E}(F(\mathbf{x}_k)) \\
 &\quad + K^{-1} \sum_{k=1}^K L_f^2 L_g \Theta(k^{-4a+4b}) + K^{-1} L_f^2 \sum_{k=1}^K \Theta(k^{-b}) + K^{-1} \sum_{k=1}^K \Theta(k^{-a}) \\
 &\leq 2K^{-1} F(\mathbf{x}_0) + K^{-1} \sum_{k=2}^K a k^{a-1} \mathbb{E}(F(\mathbf{x}_k)) \\
 &\quad + K^{-1} \sum_{k=1}^K L_f^2 L_g \Theta(k^{-4a+4b}) + K^{-1} L_f^2 \sum_{k=1}^K \Theta(k^{-b}) + K^{-1} \sum_{k=1}^K \Theta(k^{-a})
 \end{aligned}$$

$$\leq O(K^{a-1} + L_f^2 L_g K^{4b-4a} \mathbf{I}_{4a-4b=1}^{\log K} + L_f^2 K^{-b} + K^{-a}),$$

where the second inequality uses the fact that $h(t) = t^a$ is a concave function suggesting $(k+1)^a \leq k^a + ak^{a-1}$, and the last inequality uses the condition $\mathbb{E}(F(\mathbf{x}_k)) \leq \Theta(1)$.

Letting $a^* = 5/9$ and $b^* = 4/9$, we obtain the desired convergence rate $O(K^{-4/9})$. \blacksquare

B. Proof to Theorem 2

In this section, we provide the detailed proof for Theorem 2. The proof is based on the following key lemma, which provides a ‘‘contraction’’ property of $(\mathbb{E}(H(\mathbf{x}_{k+1})) - H^*)$ and $\mathbb{E}(\|\mathbf{x}_{k+1} - \mathcal{P}_{X^*}(\mathbf{x}_{k+1})\|^2)$.

Lemma 8 *Assume that both $F(\mathbf{x})$ and $R(\mathbf{x})$ are convex. Under Assumptions 1, 2, 3, and 4, the iterates generated by Algorithm 1 satisfies for any sequence of positive scalars $\{\phi_k\}$:*

$$\begin{aligned} & 2\alpha_k(\mathbb{E}(H(\mathbf{x}_{k+1})) - H^*) + \mathbb{E}(\|\mathbf{x}_{k+1} - \mathcal{P}_{X^*}(\mathbf{x}_{k+1})\|^2) \\ & \leq (1 + \phi_k)\mathbb{E}(\|\mathbf{x}_k - \mathcal{P}_{X^*}(\mathbf{x}_k)\|^2) + \Theta(\alpha_k^3) + \Theta(L_f^2 \alpha_k^2 / \phi_k)\mathbb{E}(\|\mathbf{y}_k - g(\mathbf{x}_k)\|^2) + \Theta(\alpha_k^2). \end{aligned} \quad (24)$$

Proof Following the line of the proof to Lemma 4, we have

$$\mathbf{x}_{k+1} - \mathbf{x}_k = -\alpha_k(\nabla g_{w_k}^\top(\mathbf{x}_k)\nabla f_{v_k}(\mathbf{y}_k) + \mathbf{s}_{k+1}) \quad (25)$$

where $\mathbf{s}_{k+1} \in \partial R(\mathbf{x}_{k+1})$ is some vector in the sub-differential set of $R(\cdot)$ at \mathbf{x}_{k+1} . Then we consider $\|\mathbf{x}_{k+1} - \mathcal{P}_{X^*}(\mathbf{x}_{k+1})\|^2$:

$$\begin{aligned} & \|\mathbf{x}_{k+1} - \mathcal{P}_{X^*}(\mathbf{x}_{k+1})\|^2 \\ & \leq \|\mathbf{x}_{k+1} - \mathbf{x}_k + \mathbf{x}_k - \mathcal{P}_{X^*}(\mathbf{x}_k)\|^2 \\ & = \|\mathbf{x}_k - \mathcal{P}_{X^*}(\mathbf{x}_k)\|^2 - \|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2 + 2\langle \mathbf{x}_{k+1} - \mathbf{x}_k, \mathbf{x}_{k+1} - \mathcal{P}_{X^*}(\mathbf{x}_k) \rangle \\ (25) \quad & = \|\mathbf{x}_k - \mathcal{P}_{X^*}(\mathbf{x}_k)\|^2 - \|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2 - 2\alpha_k \langle \nabla g_{w_k}^\top(\mathbf{x}_k)\nabla f_{v_k}(\mathbf{y}_k) + \mathbf{s}_{k+1}, \mathbf{x}_{k+1} - \mathcal{P}_{X^*}(\mathbf{x}_k) \rangle \\ & = \|\mathbf{x}_k - \mathcal{P}_{X^*}(\mathbf{x}_k)\|^2 - \|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2 + 2\alpha_k \langle \nabla g_{w_k}^\top(\mathbf{x}_k)\nabla f_{v_k}(\mathbf{y}_k), \mathcal{P}_{X^*}(\mathbf{x}_k) - \mathbf{x}_{k+1} \rangle \\ & \quad + 2\alpha_k \langle \mathbf{s}_{k+1}, \mathcal{P}_{X^*}(\mathbf{x}_k) - \mathbf{x}_{k+1} \rangle \\ & \leq \|\mathbf{x}_k - \mathcal{P}_{X^*}(\mathbf{x}_k)\|^2 - \|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2 + 2\alpha_k \langle \nabla g_{w_k}^\top(\mathbf{x}_k)\nabla f_{v_k}(\mathbf{y}_k), \mathcal{P}_{X^*}(\mathbf{x}_k) - \mathbf{x}_{k+1} \rangle \\ & \quad + 2\alpha_k (R(\mathcal{P}_{X^*}(\mathbf{x}_k)) - R(\mathbf{x}_{k+1})) \quad (\text{due to the convexity of } R(\cdot)) \\ & \leq \|\mathbf{x}_k - \mathcal{P}_{X^*}(\mathbf{x}_k)\|^2 - \|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2 + 2\alpha_k \underbrace{\langle \nabla F(\mathbf{x}_k), \mathcal{P}_{X^*}(\mathbf{x}_k) - \mathbf{x}_{k+1} \rangle}_{T_1} \\ & \quad + 2\alpha_k \underbrace{\langle \nabla g_{w_k}^\top(\mathbf{x}_k)\nabla f_{v_k}(\mathbf{y}_k) - \nabla F(\mathbf{x}_k), \mathcal{P}_{X^*}(\mathbf{x}_k) - \mathbf{x}_{k+1} \rangle}_{T_2} \\ & \quad + 2\alpha_k (R(\mathcal{P}_{X^*}(\mathbf{x}_k)) - R(\mathbf{x}_{k+1})) \end{aligned} \quad (26)$$

where the second equality follows from $\|a+b\|^2 = \|b\|^2 - \|a\|^2 + 2\langle a, a+b \rangle$ with $a = \mathbf{x}_{k+1} - \mathbf{x}_k$ and $b = \mathbf{x}_k - \mathcal{P}_{X^*}(\mathbf{x}_k)$. We next estimate the upper bound for T_1 and T_2 respectively:

$$T_1 = \langle \nabla F(\mathbf{x}_k), \mathbf{x}_k - \mathbf{x}_{k+1} \rangle + \langle \nabla F(\mathbf{x}_k), -\mathbf{x}_k + \mathcal{P}_{X^*}(\mathbf{x}_k) \rangle$$

$$\begin{aligned}
 &\leq \underbrace{F(\mathbf{x}_k) - F(\mathbf{x}_{k+1}) + \frac{L_F}{2} \|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2}_{\text{due to Assumption 4}} + \underbrace{F(\mathcal{P}_{X^*}(\mathbf{x}_k)) - F(\mathbf{x}_k)}_{\text{due to the convexity of } F(\cdot)} \\
 &= F(\mathcal{P}_{X^*}(\mathbf{x}_k)) - F(\mathbf{x}_{k+1}) + \frac{L_F}{2} \|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2 \\
 &\leq F(\mathcal{P}_{X^*}(\mathbf{x}_k)) - F(\mathbf{x}_{k+1}) + \Theta(\alpha_k^2),
 \end{aligned}$$

where the last inequality uses Lemma 4.

$$\begin{aligned}
 T_2 &= \langle \nabla F(\mathbf{x}_k) - \nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(\mathbf{y}_k), \mathbf{x}_k - \mathcal{P}_{X^*}(\mathbf{x}_k) \rangle \\
 &\quad + \langle \nabla F(\mathbf{x}_k) - \nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(\mathbf{y}_k), \mathbf{x}_{k+1} - \mathbf{x}_k \rangle \\
 &\leq \underbrace{\langle \nabla F(\mathbf{x}_k) - \nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(g(\mathbf{x}_k)), \mathbf{x}_k - \mathcal{P}_{X^*}(\mathbf{x}_k) \rangle}_{T_{2,1}} \\
 &\quad + \underbrace{\langle \nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(g(\mathbf{x}_k)) - \nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(\mathbf{y}_k), \mathbf{x}_k - \mathcal{P}_{X^*}(\mathbf{x}_k) \rangle}_{T_{2,2}} \\
 &\quad + \frac{\alpha_k}{2} \underbrace{\|\nabla F(\mathbf{x}_k) - \nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(\mathbf{y}_k)\|^2}_{T_{2,3}} + \frac{1}{2\alpha_k} \|\mathbf{x}_k - \mathbf{x}_{k+1}\|^2
 \end{aligned}$$

where the last line is due to the inequality $\langle a, b \rangle \leq \frac{1}{2\alpha_k} \|a\|^2 + \frac{\alpha_k}{2} \|b\|^2$. For $T_{2,1}$, we have $\mathbb{E}(T_{2,1}) = 0$ due to Assumption 1. For $T_{2,2}$, we have

$$\begin{aligned}
 T_{2,2} &\leq \frac{\alpha_k}{2\phi_k} \|\nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(g(\mathbf{x}_k)) - \nabla g_{w_k}^\top(\mathbf{x}_k) \nabla f_{v_k}(\mathbf{y}_k)\|^2 + \frac{\phi_k}{2\alpha_k} \|\mathbf{x}_k - \mathcal{P}_{X^*}(\mathbf{x}_k)\|^2 \\
 &\stackrel{\text{(Lemma 5)}}{\leq} \Theta \left(L_f^2 \frac{\alpha_k}{\phi_k} \right) \|\mathbf{y}_k - g(\mathbf{x}_k)\|^2 + \frac{\phi_k}{2\alpha_k} \|\mathbf{x}_k - \mathbf{x}_{k+1}\|^2.
 \end{aligned}$$

$T_{2,3}$ can be bounded by a constant

$$T_{2,3} \leq 2\|\nabla F(\mathbf{x}_k)\|^2 + 2\|\nabla g_{w_k}^\top \nabla f_{v_k}(\mathbf{y}_k)\|^2 \stackrel{\text{(Assumption 3)}}{\leq} \Theta(1).$$

Take expectation on T_2 and put all pieces into it:

$$\mathbb{E}(T_2) \leq \Theta \left(L_f^2 \frac{\alpha_k}{\phi_k} \right) \|\mathbf{y}_k - g(\mathbf{x}_k)\|^2 + \frac{1}{2\alpha_k} (\phi_k \|\mathbf{x}_k - \mathcal{P}_{X^*}(\mathbf{x}_k)\|^2 + \|\mathbf{x}_k - \mathbf{x}_{k+1}\|^2) + \Theta(\alpha_k).$$

Taking expectation on both sides of (26) and plugging the upper bounds of T_1 and T_2 into it, we obtain

$$\begin{aligned}
 &2\alpha_k(\mathbb{E}(H(\mathbf{x}_{k+1})) - H^*) + \mathbb{E}(\|\mathbf{x}_{k+1} - \mathcal{P}_{X^*}(\mathbf{x}_{k+1})\|^2) \\
 &\leq (1 + \phi_k) \mathbb{E}(\|\mathbf{x}_k - \mathcal{P}_{X^*}(\mathbf{x}_k)\|^2) + \Theta(\alpha_k^3) + \Theta(L_f^2 \alpha_k^2 / \phi_k) \mathbb{E}(\|\mathbf{y}_k - g(\mathbf{x}_k)\|^2) + \Theta(\alpha_k^2),
 \end{aligned}$$

which completes the proof. \blacksquare

Then, we prove the Theorem 2 based on the previous lemma.

Proof to Theorem 2 Apply the optimally strong convexity in (7) to Lemma 8, yielding

$$\begin{aligned} & (1 + 2\lambda\alpha_k)\mathbb{E}(\|\mathbf{x}_{k+1} - \mathcal{P}_{X^*}(\mathbf{x}_{k+1})\|^2) \\ & \leq (1 + \phi_k)\mathbb{E}(\|\mathbf{x}_k - \mathcal{P}_{X^*}(\mathbf{x}_k)\|^2) + \Theta(\alpha_k^3) + \Theta(L_f^2\alpha_k^2/\phi_k)\mathbb{E}(\|\mathbf{y}_k - g(\mathbf{x}_k)\|^2) + \Theta(\alpha_k^2). \end{aligned}$$

It follows by dividing $1 + 2\lambda\alpha_k$ on both sides

$$\begin{aligned} & \mathbb{E}(\|\mathbf{x}_{k+1} - \mathcal{P}_{X^*}(\mathbf{x}_{k+1})\|^2) \\ & \leq \frac{1 + \phi_k}{1 + 2\lambda\alpha_k}\mathbb{E}(\|\mathbf{x}_k - \mathcal{P}_{X^*}(\mathbf{x}_k)\|^2) + \Theta(\alpha_k^3) + \Theta(L_f^2\alpha_k^2/\phi_k)\mathbb{E}(\|\mathbf{y}_k - g(\mathbf{x}_k)\|^2) + \Theta(\alpha_k^2). \end{aligned}$$

Choosing $\phi_k = \lambda\alpha_k - 2\lambda^2\alpha_k^2 \geq 0.5\lambda\alpha_k$ yields

$$\begin{aligned} & \mathbb{E}(\|\mathbf{x}_{k+1} - \mathcal{P}_{X^*}(\mathbf{x}_{k+1})\|^2) \\ & \leq (1 - \lambda\alpha_k)\mathbb{E}(\|\mathbf{x}_k - \mathcal{P}_{X^*}(\mathbf{x}_k)\|^2) + \Theta(\alpha_k^2) + \frac{\Theta(L_f^2\alpha_k)}{\lambda}\mathbb{E}(\|g(\mathbf{x}_k) - \mathbf{y}_k\|^2) \\ & \leq (1 - \lambda\alpha_k)\mathbb{E}(\|\mathbf{x}_k - \mathcal{P}_{X^*}(\mathbf{x}_k)\|^2) + \Theta(k^{-2a}) + \Theta(L_gL_f^2k^{-5a+4b} + L_f^2k^{-a-b}). \end{aligned}$$

Apply Lemma 6 and substitute the subscript k by K to obtain the first claim in (8)

$$\mathbb{E}(\|\mathbf{x}_K - \mathcal{P}_{X^*}(\mathbf{x}_K)\|^2) \leq O\left(K^{-a} + L_f^2L_gK^{-4a+4b} + L_f^2K^{-b}\right).$$

The followed specification of a and b can easily verified. ■

C. Proof to Theorem 3

In what follows, we provide the detailed proof to Theorem 3.

Proof to Theorem 3 Let $P_k = 2\alpha_k(\mathbb{E}(H(\mathbf{x}_{k+1})) - H^*) + \mathbb{E}(\|\mathbf{x}_{k+1} - \mathcal{P}_{X^*}(\mathbf{x}_{k+1})\|^2)$, which is the left hand side of inequality (24). Summing P_k up from $k = 1$ to $k = K$ and taking $\phi_k = \Theta(k^{-c})$ yields

$$\begin{aligned} & 2\sum_{k=1}^K \alpha_k(\mathbb{E}(H(\mathbf{x}_{k+1})) - H^*) + \mathbb{E}(\|\mathbf{x}_{K+1} - \mathcal{P}_{X^*}(\mathbf{x}_{K+1})\|^2) \\ & \leq (1 + \phi_1)\mathbb{E}(\|\mathbf{x}_1 - \mathcal{P}_{X^*}(\mathbf{x}_1)\|^2) + \sum_{k=2}^K \phi_k\mathbb{E}(\|\mathbf{x}_k - \mathcal{P}_{X^*}(\mathbf{x}_k)\|^2) \\ & \quad + \sum_{k=2}^K \Theta(L_f^2\alpha_k^2/\phi_k)\mathbb{E}(\|\mathbf{y}_k - g(\mathbf{x}_k)\|^2) + \Theta(\alpha_k^2) \\ & \leq \Theta(1) + \Theta\left(\sum_{k=2}^K \phi_k\right) + \Theta(K^{1-2a}) + \sum_{k=2}^K \Theta(L_f^2\alpha_k^2/\phi_k)\mathbb{E}(\|\mathbf{y}_k - g(\mathbf{x}_k)\|^2) \end{aligned}$$

$$\leq \Theta(1) + \Theta\left(K^{1-c}\mathbf{I}_{c=1}^{\log K}\right) + O(K^{1-2a}) + \sum_{k=2}^K \Theta(L_f^2 \alpha_k^2 / \phi_k) \mathbb{E}(\|\mathbf{y}_k - g(\mathbf{x}_k)\|^2). \quad (27)$$

Note that the second inequality holds by the condition that the feasible set X is bounded, and thus $\|x_k - \mathcal{P}_{X^*}(x_k)\|^2$ is bounded.

We continue to bound the last term on the right hand side of (27):

$$\begin{aligned} & \sum_{k=2}^K \Theta(L_f^2 \alpha_k^2 / \phi_k) \mathbb{E}(\|\mathbf{y}_k - g(\mathbf{x}_k)\|^2) \\ \stackrel{\text{(Lemma 7)}}{\leq} & \sum_{k=2}^K L_f^2 L_g \Theta(k^{-6a+4b-c}) + L_f^2 \Theta(k^{-2a-b-c}) \\ \leq & L_f^2 L_g \Theta(K^{1-6a+4b-c} \mathbf{I}_{6a-4b-c=1}^{\log K}) + L_f^2 \Theta(K^{1-2a-b-c} \mathbf{I}_{2a+b-c=1}^{\log K}). \end{aligned} \quad (28)$$

Plugging (28) into (27) and dividing $2 \sum_{k=1}^K \alpha_k$ on both sides:

$$\begin{aligned} & \frac{2 \sum_{k=1}^K \alpha_k H(\mathbf{x}_{k+1}) - H^*}{2 \sum_{k=1}^K \alpha_k} \\ \leq & \frac{\Theta(1) + \Theta\left(K^{1-c}\mathbf{I}_{c=1}^{\log K}\right) + \Theta(K^{1-2a})}{2 \sum_{k=1}^K \alpha_k} \\ & + \frac{L_f^2 L_g \Theta(K^{1-6a+4b+c} \mathbf{I}_{6a-4b-c=1}^{\log K}) + L_f^2 \Theta(K^{1-2a-b+c} \mathbf{I}_{2a-b+c=1}^{\log K})}{2 \sum_{k=1}^K \alpha_k} \\ \leq & O(K^{a-1}) + O(K^{a-c} \mathbf{I}_{c=1}^{\log K}) + O(K^{-a}) \\ & + O(L_f^2 L_g K^{-5a+4b+c} \mathbf{I}_{6a-4b-c=1}^{\log K}) + O(L_f^2 K^{-a-b+c} \mathbf{I}_{2a+b-c=1}^{\log K}) \\ = & O\left(K^{a-1} + K^{a-c} \mathbf{I}_{c=1}^{\log K} + K^{-a} + L_f^2 L_g K^{-5a+4b+c} \mathbf{I}_{6a-4b-c=1}^{\log K} + L_f^2 K^{-a-b+c} \mathbf{I}_{2a+b-c=1}^{\log K}\right). \end{aligned}$$

where the last inequality uses the fact $\sum_{k=1}^K \alpha_k \geq \Theta(K^{1-a})$. To finish the proof of (11), we use Jensen's inequality

$$\frac{2 \sum_{k=1}^K \alpha_k H(\mathbf{x}_{k+1})}{2 \sum_{k=1}^K \alpha_k} \geq H(\bar{\mathbf{x}}_K).$$

■

References

- L Baird. Residual algorithms: Reinforcement learning with function approximation. In *International Conference on Machine Learning (ICML)*, pages 30–37, 1995.
- B. Dai, N. He, Y. Pan, B. Boots, and L. Song. Learning from conditional distributions via dual embeddings. In *Artificial Intelligence and Statistics (AISTATS)*, pages 1458–1467, 2017.

- C. Dann, G. Neumann, and J. Peters. Policy evaluation with temporal differences: A survey and comparison. *The Journal of Machine Learning Research*, 15(1):809–883, 2014.
- D. Dentcheva, S. Penev, and A. Ruszczyński. Statistical estimation of composite risk functionals and risk optimization problems. *Annals of the Institute of Statistical Mathematics*, pages 1–24, 2016.
- Y. M. Ermoliev. *Methods of Stochastic Programming*. Monographs in Optimization and OR, Nauka, Moscow, 1976.
- Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases (KDD)*, pages 795–811. Springer, 2016.
- B. Liu, J. Liu, M. Ghavamzadeh, S. Mahadevan, and M. Petrik. Finite-sample analysis of proximal gradient td algorithms. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2015.
- J. Liu and S. J. Wright. Asynchronous stochastic coordinate descent: Parallelism and convergence properties. *SIAM Journal on Optimization*, 25(1):351–376, 2015.
- I Necoara, Yu Nesterov, and F Glineur. Linear convergence of first order methods for non-strongly convex optimization. *arXiv preprint arXiv:1504.06298*, 2015.
- A. Nedić and D. P. Bertsekas. Incremental subgradient methods for nondifferentiable optimization. *SIAM Journal on Optimization*, 12:109–138, 2001.
- A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19:1574–1609, 2009.
- R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT press, 1998.
- M. Wang and D. P. Bertsekas. Stochastic first-order methods with random constraint projection. *SIAM Journal on Optimization*, 26(1):681–717, 2016.
- M. Wang and J. Liu. A stochastic compositional subgradient method using Markov samples. *Winter Simulation Conference*, 2016.
- M. Wang, Y. Chen, J. Liu, and Y. Gu. Random multi-constraint projection: Stochastic gradient methods for convex optimization with many constraints. *arXiv preprint arXiv:1511.03760*, 2015.
- M. Wang, J. Liu, and E. X. Fang. Accelerating stochastic composition optimization. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1714–1722, 2016.
- M. Wang, E. X. Fang, and H. Liu. Stochastic compositional gradient descent: algorithms for minimizing compositions of expected-value functions. *Mathematical Programming*, 161(1-2):419–449, 2017.

- A. White and M. White. Investigating practical linear temporal difference learning. In *International Conference on Autonomous Agents & Multiagent Systems (AAMAS)*, pages 494–502, 2016.