# Fast Stochastic Variance Reduced ADMM for Stochastic Composition Optimization 

Yue Yu and Longbo Huang<br>Institute for Interdisciplinary Information Sciences, Tsinghua University<br>yu-y14@mails.tsinghua.edu.cn, longbohuang@tsinghua.edu.cn


#### Abstract

We consider the stochastic composition optimization problem proposed in [Wang et al., 2016a], which has applications ranging from estimation to statistical and machine learning. We propose the first ADMM-based algorithm named com-SVRADMM, and show that com-SVR-ADMM converges linearly for strongly convex and Lipschitz smooth objectives, and has a convergence rate of $O(\log S / S)$, which improves upon the $O\left(S^{-4 / 9}\right)$ rate in [Wang et al., 2016b] when the objective is convex and Lipschitz smooth. Moreover, com-SVR-ADMM possesses a rate of $O(1 / \sqrt{S})$ when the objective is convex but without Lipschitz smoothness. We also conduct experiments and show that it outperforms existing algorithms.


## 1 Introduction

Recently, [Wang et al., 2016a] proposed the stochastic composition optimization of the following form:

$$
\begin{equation*}
\min _{x}\left(\mathbf{E}_{i} f_{i} \circ \mathbf{E}_{j} g_{j}\right)(x) . \tag{1}
\end{equation*}
$$

Here $x \in \mathbb{R}^{q}, \mathbf{E}_{i} f_{i}=\mathbf{E}_{i}\left\{f_{i}(x)\right\},(f \circ g)(x) \triangleq f(g(x))$ denotes the composite function, and $i, j$ are random variables. Problem (1) has been shown in [Wang et al., 2016a] to include several important applications in estimation and machine learning.
In this paper, we focus on extending the formulation to include linear constraints, and consider the following variant of Problem (1):

$$
\begin{array}{cl}
\min _{x, \omega} & F(x)+R(\omega) \\
\text { s.t. } & A x+B \omega=0 . \tag{3}
\end{array}
$$

Here $F(x) \triangleq \frac{1}{n} \sum_{i=1}^{n} f_{i}\left(\frac{1}{m} \sum_{j=1}^{m} g_{j}(x)\right), x \in \mathbb{R}^{q}, \omega \in \mathbb{R}^{l}$, $A \in \mathbb{R}^{p \times q}, B \in \mathbb{R}^{p \times l}, g_{j}: \mathbb{R}^{q} \mapsto \mathbb{R}^{r}$ and $f_{i}: \mathbb{R}^{r} \mapsto \mathbb{R}$ are continuous functions, and $R(\omega): \mathbb{R}^{l} \mapsto \mathbb{R}$ is a closed convex function. The reason to consider the specific form of Problem (2) is as follows. (i) In practice, random variables such as $i$ and $j$ are obtained from problem-dependent data sets. Thus,
they often only take values in a finite set with certain frequencies (captured by the first term in the objective (2)). (ii) Such problems often require the solutions to satisfy certain regularizing conditions (imposed by the term $R(\omega)$ and constraint (3)). Note here that the uniform distribution of $i$ and $j$ (the $\frac{1}{n}$ and $\frac{1}{m}$ ) in (2) is not critical. In Section 4, we show that our algorithm is also applicable under other distributions.

### 1.1 Motivating Examples

We first present a few motivating examples of formulation (2). The first example is a risk-averse learning problem discussed in [Wang et al., 2016b], which can be formulated into the following mean-variance minimization problems, i.e.,

$$
\begin{array}{cl}
\min _{x} & \mathbf{E}_{\epsilon} h_{\epsilon}(x)+\lambda \mathbf{V a r}_{\epsilon} h_{\epsilon}(x) \\
\text { s.t. } & A x=0 . \tag{5}
\end{array}
$$

Here $h_{\epsilon}(x)$ is the loss function w.r.t variable $x$ and is parameterized by random variable $\epsilon$, and $\operatorname{Var}_{\epsilon}(x) \triangleq \mathbf{E}_{\epsilon}\left\{\left[h_{\epsilon}(x)-\right.\right.$ $\left.\left.\mathbf{E}_{\epsilon} h_{\epsilon}(x)\right]^{2}\right\}$ denotes its variance. We see that this example is of the form (2), where $\mathbf{E}_{\epsilon} h_{\epsilon}(x)$ plays the role of the regularizer and the variance term is the composition functions. There are many other problems that can be formulated into the mean-variance optimization as in (4), e.g., portfolio management [Alexander and Baptista, 2004].
The second motivating example is dynamic programming [Sutton and Barto, 1998; Dai et al., 2016]. In this case, one can often approximate the value of each state by an inner product of a state feature $\phi_{s}$ and a target variable $w$. Then, the policy learning problem can be formulated into minimizing the Bellman residual as follows:

$$
\begin{equation*}
\min _{w} \sum_{s=1}^{S}\left(\left\langle\phi_{s}, w\right\rangle-\sum_{s^{\prime}} P_{s, s^{\prime}}^{\pi}\left(r_{s, s^{\prime}}+\gamma\left\langle\phi_{s^{\prime}}, w\right\rangle\right)\right)^{2}+R(w), \tag{6}
\end{equation*}
$$

where $P_{s, s^{\prime}}^{\pi}$ denotes the transition probabilities under a policy $\pi$, and $\gamma$ denotes the discounting factor. Note that this problem also has the form of Problem (2).

The third example is multi-stage stochastic programming [Shapiro et al., 2014]. For example, a two-stage optimization scenario often requires solving the following problem:

$$
\min _{x} \mathbf{E}_{v}\left(\min _{y} \mathbf{E}_{u \mid v}(U(x, v, y, u))\right) .
$$

Here $x, y$ are decision variables, $v, u$ are the corresponding random variables, and the function $U$ is the utility function.

In this case, the expectation $\mathbf{E}_{u \mid v}(\cdot)$ is the inner function and $\min _{y}(\cdot)$ is the outer function in Problem (2).

From these examples, we see that formulation (2) is general and includes important applications. Thus, it is important to develop fast and robust algorithms for solving (2).

### 1.2 Related Works

The stochastic composition optimization problem was first proposed in [Wang et al., 2016a], where two solution algorithms, Basic SCGD and accelerated SCGD, were proposed. The algorithms were shown to achieve a sublinear convergence rate for convex and strongly convex cases, and were also shown to converge to a stationary point in the nonconvex case. Later, [Wang et al., 2016b] proposed a proximal gradient algorithm called ASC-PG to improve the convergence rate when both inner and outer functions are smooth. However, the convergence rate is sublinear and their results do not include the regularizer when the objective functions are not strongly convex. In [Lian et al., 2016], the authors solved the finite sample case of stochastic composition optimization and obtained two linear-convergent algorithms based on the stochastic variance reduction gradient technique (SVRG) proposed in [Johnson and Zhang, 2013]. However, the algorithms do not handle the regularizer either.

The ADMM algorithm, on the other hand, was first proposed in [Glowinski and Marroco, 1975; Gabay and Mercier, 1976] and later reviewed in [Boyd et al., 2011]. Since then, several ADMM-based stochastic algorithms have been proposed, e.g., [Ouyang et al., 2013; Suzuki and others, 2013; Wang and Banerjee, 2013]. However, these algorithms all possess sublinear convergence rates. Therefore, several recent works tried to combine the variance reduction scheme and ADMM to accelerate convergence. For instance, SVRGADMM was proposed in [Zheng and Kwok, 2016a]. It was shown that SVRG-ADMM converges linearly when the objective is strongly convex, and has a sublinear convergence rate in the general convex case. Another recent work [Zheng and Kwok, 2016b] further proved that SVRG-ADMM converges to a stationary point with a rate $O\left(\frac{1}{T}\right)$ when the objective is nonconvex. In [Qian and Zhou, 2016], the authors used acceleration technique in [Allen-Zhu, 2016; Hien et al., 2016] to further improve the convergence rate of SVRG-ADMM. However, all aforementioned variancereduced ADMM algorithms cannot be directly applied to solving the stochastic composition optimization problem.

### 1.3 Contribution

In this paper, we propose an efficient algorithm called com-SVR-ADMM, which combines ideas of SVRG and ADMM, to solve stochastic composition optimization. Our algorithm is based on the SVRG-ADMM algorithm proposed in [Zheng and Kwok, 2016a], which does not apply to composite optimization problems. We consider three different objective functions in Problem (2), and show that our algorithm achieves the following performance.

- When $F(x)$ is strongly convex and Lipschitz smooth, and $R(\omega)$ is convex, our algorithm converges linearly. This convergence rate is comparable with those of com-

SVRG-1 and com-SVRG-2 in [Lian et al., 2016]. However, com-SVRG-1 and com-SVRG-2 do not take the commonly used regularization penalty into consideration. Experimental results also show that com-SVRADMM converges faster than com-SVRG-1 and com-SVRG-2.

- When $F(x)$ is convex and Lipschitz smooth, and $R(\omega)$ is convex, com-SVR-ADMM has a sublinear rate of $O\left(\frac{\log (S+1)}{S}\right)$, where $S$ is the number of outer iterations. ${ }^{1}$ This result outperforms the $O\left(S^{-4 / 9}\right)$ convergence rate of ASC-PG in [Wang et al., 2016b]. ${ }^{2}$
- When $F(x)$ and $R(\omega)$ are general convex functions (not necessarily Lipschitz smooth), com-SVR-ADMM achieves a rate of $O\left(\frac{1}{\sqrt{S}}\right)$. To the best of our knowledge, this is the first convergence result for stochastic composite optimization with general convex problems without Lipschitz smoothness.


### 1.4 Notation

For vector $x$ and a positive definite matrix $G$, the $G$-norm of vector $x$ is defined as $\|x\|_{G}=\sqrt{x^{T} G x}$. For a matrix $X$, $\|X\|$ denotes its spectral norm, $\sigma_{\max }(X), \sigma_{\min }(X)$ denote its largest and smallest eigenvalue, respectively. $X^{\dagger}$ denotes the pseudoinverse of $X . \quad \tilde{\nabla} R(\omega)$ denotes a noisy subgradient of nonsmooth $R(\omega)$. For a function $g(x): \mathbb{R}^{q} \mapsto \mathbb{R}^{r}$, $\partial g(x) \in \mathbb{R}^{r \times q}$ denotes its Jacobian matrix. $\nabla f_{i_{k}}(g(x))$ denotes the gradient of $f_{i_{k}}(\cdot)$ at point $y=g(x)$.

## 2 Algorithm

Recall that the stochastic composition problem we want to solve has the following form:

$$
\begin{array}{cl}
\min _{x, \omega} & F(x)+R(\omega) \\
\text { s.t. } & A x+B \omega=0
\end{array}
$$

where $F(x) \triangleq \frac{1}{n} \sum_{i=1}^{n} f_{i}\left(\frac{1}{m} \sum_{j=1}^{m} g_{j}(x)\right)$. For clarity, we denote $F(x)=\frac{1}{n} \sum_{i=1}^{n} F_{i}(x), F_{i}(x)=f_{i}(g(x)), g(x)=\mathbf{E} g(x)=$ $\frac{1}{m} \sum_{j=1}^{m} g_{j}(x)$. Therefore, $\nabla F_{i}(x)=(\partial g(x))^{T} \nabla f_{i}(g(x))$.

Our proposed procedure adopts the ADMM scheme. At every iteration the primal variables $(x, \omega)$ are obtained by minimizing the following augmented Lagrangian equation parameterized with $\rho>0$, i.e.,

$$
\begin{aligned}
& L_{\rho}(x, \omega, \lambda) \\
& =F(x)+R(\omega)+\langle\lambda, A x+B \omega\rangle+\frac{\rho}{2}\|A x+B \omega\|_{2}^{2}
\end{aligned}
$$

The update of dual variable $\lambda$ is similar to that under gradient descent with the stepsize equaling $\rho$. We also based our algorithm on a sampling oracle as in [Wang et al., 2016b]. Specifically, we assume a stochastic first-order oracle, which,

[^0]```
Algorithm 1 com-SVR-ADMM for strongly convex stochas-
tic composition optimization
    Input: \(K, M, N, \eta, \rho, \tilde{x}^{0}, \tilde{\omega}^{0}, \tilde{\lambda}^{0}=-\left(A^{T}\right)^{\dagger} \nabla F\left(\tilde{x}^{0}\right) ;\)
    for \(s=1,2, \ldots\) do
        \(\tilde{x}=\tilde{x}^{s-1}, x^{0}=\tilde{x}^{s-1}, \omega^{0}=\tilde{\omega}^{s-1}, \lambda^{0}=\tilde{\lambda}^{s-1} ;\)
        \(g(\tilde{x})=\frac{1}{m} \sum_{j=1}^{m} g_{j}(\tilde{x}) ; \quad \quad(m\) queries)
        evaluate \(\nabla F(\tilde{x}) ; \quad(m+n\) queries \()\)
        for \(k=0\) to \(K-1\) do
            \(\omega^{k+1}=\arg \min _{\omega} R(\omega)+\left\langle\lambda^{k}, B \omega\right\rangle+\frac{\rho}{2} \| A x^{k}+\)
            \(B \omega \|_{2}^{2}\);
            uniformly sample \(N_{k}\) and calculate \(\hat{g}\left(x^{k}\right)\) using (9);
            ( \(2 N\) queries)
            uniformly sample \(i_{k}, j_{k}\) and calculate \(\nabla \hat{F}_{i_{k}}\left(x^{k}\right)\) us-
            ing (8);
                                    (4 queries)
            \(x^{k+1}=\arg \min _{x}\left\langle\nabla \hat{F}_{i_{k}}\left(x^{k}\right), x-x^{k}\right\rangle+\left\langle\lambda^{k}, A x\right\rangle+\)
            \(\frac{\rho}{2}\left\|A x+B \omega^{k+1}\right\|_{2}^{2}+\frac{1}{2 \eta}\left\|x-x^{k}\right\|_{2}^{2} ;\)
            \(\lambda^{k+1}=\lambda^{k}+\rho\left(A x^{k+1}+B \omega^{k+1}\right) ;\)
        end for
        \(\tilde{x}^{s}=\frac{1}{K} \sum_{k=1}^{K} x^{k}, \quad \tilde{w}^{s}=\frac{1}{K} \sum_{k=1}^{K} w^{k}, \quad \tilde{\lambda}^{s}=\)
        \(-\left(A^{T}\right)^{\dagger} \nabla F\left(\tilde{x}^{s}\right) ;\)
end for
Output: \(\tilde{x}^{s}, \tilde{w}^{s}\).
```

if queried, returns a noisy gradient/subgradient or function value of $f_{i}(\cdot)$ and $g_{j}(\cdot)$ at a given point.

In the following sections, we introduce the stochastic variance reduced ADMM algorithm for solving stochastic compositional optimization (com-SVR-ADMM). We present com-SVR-ADMM in three different scenarios, i.e., strongly convex and Lipschitz smooth, general convex and Lipschitz smooth, and general convex. Algorithm 1 shows how com-SVR-ADMM is used in the strongly convex case, while Algorithm 2 is for the second and third cases.

### 2.1 Compositional Stochastic Variance Reduced ADMM for Strongly Convex Functions

As in SVRG, com-SVR-ADMM has $K$ inner loops inside each outer iteration. At every outer iteration, we need to keep track of a reference point $\tilde{x}$ (Step 3 in Algorithm 1) for computing $g(\tilde{x})$ defined as

$$
\begin{equation*}
g(\tilde{x})=\frac{1}{m} \sum_{j=1}^{m} g_{j}(\tilde{x}) \tag{7}
\end{equation*}
$$

and evaluate $\partial g(\tilde{x})$, which is the Jacobian matrix of $g(x)$ at point $\tilde{x}$. The updates of $\omega^{k+1}$ and $\lambda^{k+1}$ are the same as those in batch ADMM [Boyd et al., 2011]. The main difference lies in the update for $x^{k+1}$, in that we use a stochastic sample $i_{k}$ and replace $F_{i_{k}}(x)$ with its first-order approximation, and then approximate $\nabla F_{i_{k}}\left(x^{k}\right)$ by

$$
\begin{align*}
\nabla \hat{F}_{i_{k}}\left(x^{k}\right)= & \left(\partial g_{j_{k}}\left(x^{k}\right)\right)^{T} \nabla f_{i_{k}}\left(\hat{g}\left(x^{k}\right)\right)  \tag{8}\\
& -\left(\partial g_{j_{k}}(\tilde{x})\right)^{T} \nabla f_{i_{k}}(g(\tilde{x}))+\nabla F(\tilde{x})
\end{align*}
$$

Here $i_{k}, j_{k}$ are uniformly sampled from $\{1,2, \ldots, n\}$ and $\{1, \ldots, m\}$, respectively. $\hat{g}\left(x^{k}\right)$ is an estimation of $g\left(x^{k}\right)$ defined as follows:

$$
\begin{equation*}
\hat{g}\left(x^{k}\right)=g(\tilde{x})-\frac{1}{N} \sum_{1 \leq j \leq N}\left(g_{\mathbb{N}_{k}[j]}(\tilde{x})-g_{\mathbb{N}_{k}[j]}\left(x^{k}\right)\right), \tag{9}
\end{equation*}
$$

where $\mathbb{N}_{k}$ is a mini-batch and is obtained by uniformly and randomly sampling from $\{1, \ldots, m\}$ for $N$ times (with replacement) and $\mathbb{N}_{k}[j]$ is the $j$ th element of $\mathbb{N}_{k}$. In step 10 of Algorithm 1, we add a proximal term $\frac{1}{2 \eta}\left\|x-x^{k}\right\|_{2}^{2}$ to control the distance between the next point $x^{k+1}$ and the current point $x^{k}$. The parameter $\eta$ is a constant and plays the role of stepsize as in [Ouyang et al., 2013].

Note that our estimated gradient $\nabla \hat{F}_{i_{k}}\left(x^{k}\right)$ is biased due to the composition objective function and its form is the same as com-SVRG-1 in [Lian et al., 2016]. However, we remark that our algorithm is not a trivial extension of com-SVRG1 due to the existence of linear constraint and Lagrangian dual variable. Moreover, com-SVR-ADMM can handle regularization penalty while com-SVRG-1 cannot. Also, the update of $\tilde{\lambda}^{s}$ uses the pseudoinverse of $A$. In the common case when $A$ is sparse, one can use the efficient Lanczos algorithm [Golub and Van Loan, 2012] to compute $A^{\dagger}$. Note that step 10 in Algorithm 1 often involves computing $A^{T} A$. The memory complexity for this step can be alleviated by algorithms proposed in recent works, e.g., [Zheng and Kwok, 2016a; Zhang et al., 2011].

### 2.2 Compositional Stochastic Variance Reduced ADMM for General Convex Functions

In this section, we describe how com-SVR-ADMM handles general convex composition problems with Lipschitz smoothness. Without strong convexity, we need to make subtle changes. As shown in Algorithm 2, besides changes in variable initialization and output, another key difference is the approximation of $\nabla F_{i_{k}}(x)$, where we use $g\left(x^{k}\right)$ instead of $\hat{g}\left(x^{k}\right)$, i.e.,

$$
\begin{align*}
\nabla \hat{F}_{i_{k}}\left(x^{k}\right)= & \left(\partial g_{j_{k}}\left(x^{k}\right)\right)^{T} \nabla f_{i_{k}}\left(g\left(x^{k}\right)\right)  \tag{10}\\
& -\left(\partial g_{j_{k}}(\tilde{x})\right)^{T} \nabla f_{i_{k}}(g(\tilde{x}))+\nabla F(\tilde{x})
\end{align*}
$$

Note that in the cases of interest (see Assumption 1 below), the approximated gradient $\nabla \hat{F}_{i_{k}}\left(x^{k}\right)$ is unbiased. The next change is the stepsize for updating $x$. In step 10 of Algorithm 2, we use a positive definite matrix $G_{k}$ in the proximal term. ${ }^{3}$ Therefore, the stepsize depends on two parameters: $\eta_{s}$ and $G_{k}$, as shown in (11), where $s$ and $k$ are the iteration counters for outer and inner iteration, respectively. Here $L_{F}$ is a parameter of Lipschitz smoothness and will be specified in our assumptions in next section.

$$
\begin{align*}
& \eta_{s}=\frac{1}{(s+1) L_{F}}, \quad G_{0} \succeq G_{1} \succeq G_{2} \succeq \ldots \succeq G_{K-1}  \tag{11}\\
& G_{0}=\frac{1}{s} I, \quad G_{K-1}=\frac{1}{s+1} I, \quad G_{K}=\frac{1}{s+1} I
\end{align*}
$$

[^1]```
Algorithm 2 com-SVR-ADMM for general convex stochas-
tic composition optimization
    Input: \(S, K, N, \eta_{s}, \rho, \tilde{x}^{0}=\hat{x}^{0}, \hat{\omega}^{0}, \hat{\lambda}^{0}, \hat{G}^{0}=I ;\)
    for \(s=1,2, \ldots, S\) do
        \(\tilde{x}=\tilde{x}^{s-1}, x^{0}=\hat{x}^{s-1}, \omega^{0}=\hat{\omega}^{s-1}, \lambda^{0}=\hat{\lambda}^{s-1}\),
        \(G_{0}=\hat{G}^{s-1}\);
        \(g(\tilde{x})=\frac{1}{m} \sum_{j=1}^{m} g_{j}(\tilde{x}) ; \quad\) ( \(m\) queries)
        evaluate \(\nabla F(\tilde{x}) ; \quad(m+n\) queries \()\)
    for \(k=0\) to \(K-1\) do
        \(\omega^{k+1}=\arg \min _{\omega} R(\omega)+\left\langle\lambda^{k}, B \omega\right\rangle+\frac{\rho}{2} \| A x^{k}+\)
        \(B \omega \|_{2}^{2}\);
        calculate \(g\left(x^{k}\right)=\frac{1}{m} \sum_{j=1}^{m} g_{j}\left(x^{k}\right) ; \quad(m\) queries)
        uniformly sample \(i_{k}, j_{k}\) and calculate \(\nabla \hat{F}_{i_{k}}\left(x^{k}\right)\) us-
        ing (10);
                            (4 queries)
        \(x^{k+1}=\arg \min _{x}\left\langle\nabla \hat{F}_{i_{k}}\left(x^{k}\right), x-x^{k}\right\rangle+\left\langle\lambda^{k}, A x\right\rangle+\)
        \(\frac{\rho}{2}\left\|A x+B \omega^{k+1}\right\|_{2}^{2}+\frac{1}{2 \eta_{s}}\left\|x-x^{k}\right\|_{G_{k}}^{2} ;\)
        \(\lambda^{k+1}=\lambda^{k}+\rho\left(A x^{k+1}+B \omega^{k+1}\right) ;\)
        end for
        \(\tilde{x}^{s}=\frac{1}{K} \sum_{k=1}^{K} x^{k}, \quad \tilde{\omega}^{s}=\frac{1}{K} \sum_{k=1}^{K} \omega^{k}, \quad \tilde{\lambda}^{s}=\frac{1}{K} \sum_{k=1}^{K} \lambda^{k}\),
        \(\hat{x}^{s}=x^{K}, \hat{\omega}^{s}=\omega^{K}, \hat{\lambda}^{s}=\lambda^{K}, \hat{G}^{s}=G_{K} ;\)
end for
Output: \(\bar{x}=\frac{1}{S} \sum_{s=1}^{S} \tilde{x}^{s}, \bar{\omega}=\frac{1}{S} \sum_{s=1}^{S} \tilde{\omega}^{s}\).
```

That is, $G_{k}$ is nonincreasing for $k=0,1, \ldots, K$. Then, according to the definition of $G$-norm and (11), we have:

$$
\begin{equation*}
\frac{1}{2 \eta_{s}}\left\|x-x^{k}\right\|_{G_{k}}^{2}=\frac{1}{2 \eta_{s, k}}\left\|x-x^{k}\right\|_{2}^{2} \tag{12}
\end{equation*}
$$

where $\eta_{s, k}=\frac{\eta_{s}}{\cos \left(G_{k}\right)}$ and $\operatorname{co}\left(G_{k}\right)=a$ if $G_{k}=a I$, and $a$ is a scalar. Therefore, $\eta_{s, k}$ serves as the stepsize [Ouyang et al., 2013], and it can be verified that $\eta_{s, k}$ satisfies the following properties:

$$
\begin{gather*}
\eta_{s, 0}=\frac{s}{(s+1) L_{F}}, \eta_{s, K-1}=\frac{1}{L_{F}}, \eta_{s, K}=\frac{1}{L_{F}}  \tag{13}\\
\eta_{s, 0} \leq \eta_{s, 1} \leq \ldots \leq \eta_{s, K-1}
\end{gather*}
$$

That is, $\eta_{s, k}$ changes from $\frac{s}{(s+1) L_{F}}$ to $\frac{1}{L_{F}}$ in stage $s$. Note that even though $\eta_{s, k}$ is not a constant, it still has a reasonable value and does not need to vanish. This feature is helpful for convergence acceleration.

### 2.3 General Convex Functions without Lipschitz Smoothness

In the previous two sections, we present the procedures of com-SVR-ADMM for the strongly convex and general convex settings, both under the Lipschitz smooth assumption of $F(x)$. In this section, we further investigate the case when the smooth condition is relaxed. We still use Algorithm 2, except
that the values $\eta_{s}$ and $G_{k}$ are changed to

$$
\begin{gather*}
\eta_{s}=\frac{1}{s+1}, \quad G_{0} \succeq G_{1} \succeq G_{2} \succeq \ldots \succeq G_{K-1} \\
G_{0}=\frac{1}{\sqrt{s}} I, \quad G_{K-1}=\frac{1}{\sqrt{s+1}} I, \quad G_{K}=\frac{1}{\sqrt{s+1}} I . \tag{14}
\end{gather*}
$$

Therefore, using the same technique in (12), it can be verified that $\eta_{s, k}$ in this setting changes from $\frac{\sqrt{s}}{s+1}$ to $\frac{1}{\sqrt{s+1}}$ in stage $s$ and decreases to zero. Note that in this case, the number of oracle calls at each step is the same as that in section 2.2.

Although the algorithm we proposed appears similar to the SVRG-ADMM algorithm in [Zheng and Kwok, 2016a], it is very different due to the composition nature of the objective function (which is not considered in SVRG-ADMM) and the stochastic variance reduced gradients in (8) and (10). These differences make it impossible to directly apply SVRG-ADMM and require a very different analysis for the new algorithm. Readers interested in the full proofs can refer to our technical report [Yu and Huang, 2017].

## 3 Theoretical Results

In this section, we analyze the convergence performance of com-SVR-ADMM under the three cases described in section 2. Below, we first state our assumptions. Note that the assumptions are not restrictive and are commonly made in the literature, e.g., [Wang et al., 2016b; Ouyang et al., 2013; Wang et al., 2016a; Zheng and Kwok, 2016b].
Assumption 1. (i) For each $i \in\{1, \ldots, n\}, F_{i}$ is convex and continuously differentiable, $R(\omega)$ is convex (can be nonsmooth). Moreover, there exists an optimal primal-dual solution ( $x^{*}, \omega^{*}, \lambda^{*}$ ) for Problem (2).
(ii) The feasible set $\mathbb{X}$ for $x$ is bounded and denote $D=$ $\max _{x, y \in \mathbb{X}}\|x-y\|$.
(iii) For randomly sampled $i_{k} \in\{1, \ldots, n\}, j_{k} \in\{1, \ldots, m\}$ and $\forall x$, we assume the following unbiased properties:

$$
\begin{gather*}
E\left(\left(\partial g_{j_{k}}(x)\right)^{T} \nabla f_{i_{k}}(g(x))\right)=\nabla F(x)  \tag{15}\\
E\left(\partial g_{j_{k}}(x)\right)=\partial g(x), \quad E\left(\nabla F_{i_{k}}(x)\right)=\nabla F(x)
\end{gather*}
$$

Assumption 2. $F$ is strongly convex with parameter $\mu_{F}>0$, i.e., $\forall x$,

$$
\begin{equation*}
F(x)-F\left(x^{*}\right) \geq\left\langle\nabla F\left(x^{*}\right), x-x^{*}\right\rangle+\frac{\mu_{F}}{2}\left\|x-x^{*}\right\|_{2}^{2} \tag{16}
\end{equation*}
$$

Assumption 3. Matrix A has full row rank.
Assumption 4. There exists a positive constant $L_{F}$, such that $\forall i \in\{1, \ldots, n\}, \forall j \in\{1, \ldots, m\}$ and $\forall x$, $y$, we have
$\left\|\left(\partial g_{j}(x)\right)^{T} \nabla f_{i}(g(x))-\left(\partial g_{j}(y)\right)^{T} \nabla f_{i}(g(y))\right\| \leq L_{F}\|x-y\|$.
Assumption 5. For each $i \in\{1, \ldots, n\}, f_{i}$ is Lipschitz smooth with positive parameter $L_{f}$, that is, $\forall x, y$, we have

$$
\begin{equation*}
\left\|\nabla f_{i}(y)-\nabla f_{i}(x)\right\| \leq L_{f}\|y-x\| . \tag{17}
\end{equation*}
$$

Assumption 6. For every $j \in\{1, \ldots, m\}, \partial g_{j}(x)$ is bounded, and for all $x, y, \exists C_{G}, L_{G}>0$ that satisfy

$$
\begin{gather*}
\left\|g_{j}(x)-g_{j}(y)\right\| \leq C_{G}\|x-y\|, \quad\left\|\partial g_{j}(x)\right\| \leq C_{G} \\
\left\|g_{j}(x)-g_{j}(y)\right\| \leq L_{G}\|x-y\|^{2} \tag{18}
\end{gather*}
$$

For clarity, we also use the following notations used in the theorems:

$$
\begin{align*}
& u=\left[\begin{array}{l}
x \\
\omega
\end{array}\right], u^{k}=\left[\begin{array}{l}
x^{k} \\
\omega^{k}
\end{array}\right], \tilde{u}^{s}=\left[\begin{array}{c}
\tilde{x}^{s} \\
\tilde{\omega}^{s}
\end{array}\right], \bar{u}=\left[\begin{array}{l}
\bar{x} \\
\bar{\omega}
\end{array}\right], \\
& G(u)= F(x)-F\left(x^{*}\right)-\left\langle\nabla F\left(x^{*}\right), x-x^{*}\right\rangle  \tag{19}\\
& \quad+R(\omega)-R\left(\omega^{*}\right)-\left\langle\tilde{\nabla} R\left(\omega^{*}\right), \omega-\omega^{*}\right\rangle .
\end{align*}
$$

It can be verified that $G(u)$ is always non-negative due to the convexity of $F(x)$ and $R(\omega)$. The following theorem and corollary show that Algorithm 1 has a linear convergence rate.

Proposition 1. Under Assumption 4, we have $\forall i \in\{1, \ldots, n\}$ and $\forall x, y$ :

$$
\begin{equation*}
\left\|\nabla F_{i}(x)-\nabla F_{i}(y)\right\| \leq L_{F}\|x-y\|, \tag{20}
\end{equation*}
$$

i.e., each $F_{i}$ is Lipschitz smooth. Moreover, it implies $\|\nabla F(x)-\nabla F(y)\| \leq L_{F}\|x-y\|$.
Theorem 1. Under Assumptions 1, 2, 3, 4, 5 and 6, if $0<$ $\eta \leq 1 / L_{F}$, then under Algorithm 1,

$$
\begin{equation*}
\gamma_{1} \mathbf{E}\left[G\left(\tilde{u}^{s}\right)\right] \leq \gamma_{2} G\left(\tilde{u}^{s-1}\right) \tag{21}
\end{equation*}
$$

where (denote $\sigma(N)=\sqrt{1 / N}$ )

$$
\begin{aligned}
\gamma_{1}= & \left(2 \eta-\frac{32 \eta^{2} C_{G}^{4} L_{f}^{2}}{\mu_{F} N}-\frac{48 \eta^{2} L_{F}^{2}+8 \eta D C_{G} L_{f} L_{G} \sigma(N)}{\mu_{F}}\right) K \\
\gamma_{2}= & (K+1)\left(\frac{32 \eta^{2} C_{G}^{4} L_{f}^{2}}{\mu_{F} N}+\frac{48 \eta^{2} L_{F}^{2}+8 \eta D C_{G} L_{f} L_{G} \sigma(N)}{\mu_{F}}\right) \\
& +\frac{2}{\mu_{F}}+\frac{2 \eta \rho\left\|A^{T} A\right\|}{\mu_{F}}+\frac{2 L_{F} \eta}{\rho \sigma_{\min }\left(A A^{T}\right)}
\end{aligned}
$$

Corollary 1. Suppose the conditions in Theorem 1 hold. Then, there exist positive $\Theta(1)$ constants $K$ (number of inner iterations) and $N$ (mini-batch size) such that $\gamma_{1}, \gamma_{2}>$ $0, \gamma=\gamma_{2} / \gamma_{1}<1$. Thus, Algorithm 1 converges linearly.

From Corollary 1, if we want to achieve $\mathbf{E}\left[G\left(\tilde{u}^{s}\right)\right] \leq \epsilon$, $\forall \epsilon>0$, the number of steps we need to take is roughly $s \geq \log \left(\frac{G\left(\tilde{u}^{0}\right)}{\epsilon}\right) / \log \left(\frac{1}{\gamma}\right)$. In each iteration, we need $2 m+$ $n+K(2 N+4)$ oracle calls. Therefore, the overall query complexity is $O\left((m+n+K N) \log \frac{1}{\epsilon}\right)$. For comparison, the query complexity is $O\left(\left(m+n+\kappa^{4}\right) \log (1 / \epsilon)\right)$ for com-SVRG-1 and $O\left(\left(m+n+\kappa^{3}\right) \log (1 / \epsilon)\right)$ for com-SVRG-2 [Lian et al., 2016], where $\kappa$ is a parameter related to condition number. We will see in simulations in section 4 that the overall query complexity of com-SVR-ADMM is lower than com-SVRG-1 and com-SVRG-2.

Now we prove the convergence property of com-SVRADMM under Assumptions 1 and 4.
Theorem 2. Under Assumptions 1 and 4 , if $\eta_{s}$ and $G_{k}$ are chosen as in (11), under Algorithm 2,

$$
\begin{align*}
& \mathbf{E}(G(\bar{u})+\Lambda\|A \bar{x}+B \bar{\omega}\|) \\
& \leq  \tag{22}\\
& \leq \frac{4 L_{F} D^{2} \log (S+1)}{S}+\frac{L_{F} D^{2} \log S}{2 K S} \\
& \\
& \quad+\frac{L_{F} D^{2}+\rho D^{2}\left\|A^{T} A\right\|+\frac{2}{\rho}\left\|\hat{\lambda}^{0}-\lambda^{*}\right\|_{2}^{2}+\frac{2}{\rho} \Lambda^{2}}{2 K S}
\end{align*}
$$

where $\Lambda>0$.

From Theorem 2, we see that com-SVR-ADMM has an $O\left(\frac{\log (S+1)}{S}\right)$ convergence rate under the general convex and Lipschitz smooth condition. It improves upon the convergence rate $O\left(S^{-4 / 9}\right)$ in the recent work [Wang et al., 2016b]. In Theorem 2, we consider both the convergence property of function value and feasibility violation. Since $G(u)$ and $\|A \bar{x}+B \bar{\omega}\|$ are both non-negative, each term has an $O\left(\frac{\log (S+1)}{S}\right)$ convergence rate.

In the following theorem, we show that our algorithm exhibits $O\left(\frac{1}{\sqrt{S}}\right)$ convergence rate for both the objective value and feasibility violation, when the objective is a general convex function.
Assumption 7. The gradients/subgradients of all $f_{i}, F_{i}, g_{j}$ and $R(\omega)$ are bounded and $\left\|\nabla F_{i}(x)\right\| \leq C_{F}, \quad C_{F}>0$. Moreover, $B$ is invertible and $A, B$ are bounded.
Theorem 3. Under Assumptions 1 and 7 , denote
$\dot{x}^{s}=\frac{1}{K} \sum_{k=0}^{K-1} x^{k}, \tilde{z}^{s}=\left[\begin{array}{c}\dot{x}^{s} \\ \tilde{w}^{s}\end{array}\right], \bar{z}=\frac{1}{S} \sum_{s=1}^{S} \tilde{z}^{s}$. If $\eta_{s}$ and $G_{k}$ are chosen as in (14), there exists a positive $\Theta(1)$ constant $\rho$ such that, under Algorithm 2,

$$
\begin{align*}
& \mathbf{E}(G(\bar{z})+\Lambda\|A \bar{x}+B \bar{\omega}\|) \\
& \leq \frac{C_{1}\left(C_{4}+C_{F}\right)}{\sqrt{S}}+\frac{D^{2}}{K \sqrt{S}}+\frac{C_{3} \log (S+1)}{S}  \tag{23}\\
& \\
& +\frac{D^{2}+\rho\left\|A^{T} A\right\| D^{2}+\frac{2}{\rho}\left\|\hat{\lambda}^{0}-\lambda^{*}\right\|_{2}^{2}+\frac{2}{\rho} \Lambda^{2}}{2 K S}
\end{align*}
$$

where $\Lambda, C_{1}, C_{3}, C_{4}$ are positive constants.
The reason for the introduction of $\bar{z}$ is similar to the step taken in [Ouyang et al., 2013], and is due to the lack of Lipschitz smooth property. This result implies an $O\left(\frac{1}{\sqrt{S}}\right)$ convergence rate for both objective value and feasibility violation.

## 4 Experiments

In this section, we conduct experiments and compare com-SVR-ADMM to existing algorithms. We consider two experiment scenarios, i.e., the portfolio management scenario from [Lian et al., 2016] and the reinforcement learning scenario from [Wang et al., 2016b]. Since the objective functions in both scenarios are strongly convex and Lipshitz smooth, we only provide results for Algorithm 1.

### 4.1 Portfolio Management

Portfolio management is usually formulated as meanvariance minimization of the following form:
$\min _{x}-\frac{1}{n} \sum_{i=1}^{n}\left\langle r_{i}, x\right\rangle+\frac{1}{n} \sum_{i=1}^{n}\left(\left\langle r_{i}, x\right\rangle-\frac{1}{n} \sum_{j=1}^{n}\left\langle r_{j}, x\right\rangle\right)^{2}+R(x)$,
where $r_{i} \in \mathbb{R}^{N}$ for $i \in\{1, \ldots, n\}, N$ is the number of assets, and $n$ is the number of observed time slots. Thus, $r_{i}$ is the observed reward in time slot $i$. We compare our proposed com-SVR-ADMM with three benchmarks: com-SVRG-1, com-SVRG-2 from [Lian et al., 2016], and SGD. In order to compute the unbiased stochastic gradient of SGD, we first enumerate all samples in the data set of $g$ to calculate $g(x)$ and


Figure 1: Portfolio Management with $\operatorname{cov}=2$.


Figure 2: Portfolio Management with $\operatorname{cov}=10$. The other parameters have the same value as Figure 1.


Figure 3: On-policy learning experiment with $S=200, d=100$.
$\partial g(x)$, then evaluate $\partial g(x) \nabla f_{i}(g(x))$ for a random sample $i$. Using the same definition of $g_{j}(x)$ and $f_{i}(y)$ and the same parameters generation method as [Lian et al., 2016], we set the regularization to $\left.R(x)=\frac{\mu}{2} \right\rvert\,\|x\|_{2}^{2}$, where $\mu>0$.

The experimental results are shown in Figure 1 and Figure 2. Here the $y$-axis represents the objective value minus optimal value and the $x$-axis is the number of oracle calls or CPU time. We set $N=200, n=2000$. cov is the parameter used for reward covariance matrix generation [Lian et al., 2016]. In Figure 1, $\operatorname{cov}=2$, and $\operatorname{cov}=10$ in Figure 2. All shared parameters in the four algorithms, e.g., stepsize, have the same values. We can see that all SVRG based algorithms perform much better than SGD, and com-SVR-ADMM outperforms two other linear convergent algorithms.

### 4.2 Reinforcement Learning

Here we consider the problem (6), which can be used for onpolicy learning [Wang et al., 2016b]. In our experiment, we assume there are finite states and the number of states is $S$. $\pi$ is the policy in consideration. $P_{s, s^{\prime}}^{\pi}$ is the transition probability from state $s$ to $s^{\prime}$ given policy $\pi, \gamma$ is a discount factor, $\phi_{s} \in \mathbb{R}^{d}$ is the feature of state $s$. Here we use a linear product $\left\langle\phi_{s}, w\right\rangle$ to approximate the value of state $s$. Our goal is


Figure 4: On-policy learning experiment with $S=200, d=200$. The other parameter values is the same as Figure 3.
to find the optimal $w \in \mathbb{R}^{d}$.
We use the following specifications for oracles $g_{s^{\prime}}(w)$ and $f_{s}(y)$ :

$$
\begin{aligned}
& g_{s^{\prime}}(w)=\left(\phi_{1}^{T} w, r_{1, s^{\prime}}+\gamma \phi_{s^{\prime}}^{T} w, \ldots, \phi_{S}^{T} w, r_{S, s^{\prime}}+\gamma \phi_{s^{\prime}}^{T} w\right)^{T}, \\
& f_{s}(y)=(y[2 s-1]-y[2 s])^{2}
\end{aligned}
$$

Note here $g_{s^{\prime}}(w) \in \mathbb{R}^{2 S}$, and $y[i]$ denote the $i$-th element of vector $y$. All shared parameters in four algorithms have the same values. Note here that the calculation of $\mathbf{E}[g(w)]$ is no longer under uniform distribution. We use the given transition probability. In this experiment, the transition probability is randomly generated and then regularized. The reward is also randomly generated. In addition, we include a regularization term $R(w)=\frac{\mu}{2}\|w\|_{2}^{2}$ with $\mu>0$. The results are shown in Figure 3 and Figure 4. It can be seen that our proposed com-SVR-ADMM achieves faster convergence compared to the benchmark algorithms.

## 5 Conclusion

In this paper, we propose an ADMM-based algorithm, called com-SVR-ADMM, for stochastic composition optimization. We show that when the objective function is strongly convex and Lipschitz smooth, com-SVR-ADMM converges linearly. In the case when the objective function is convex (not necessarily strongly convex) and Lipschitz smooth, com-SVR-ADMM improves the theoretical convergence rate from $O\left(S^{-4 / 9}\right)$ in [Wang et al., 2016b] to $O\left(\frac{\log S}{S}\right)$. When the objective is only assumed to be convex, com-SVR-ADMM has a convergence rate of $O\left(\frac{1}{\sqrt{S}}\right)$. Experimental results show that com-SVR-ADMM outperforms existing algorithms.

## Acknowledgements

This work was supported in part by the National Natural Science Foundation of China Grants 61672316, 61303195, the Tsinghua Initiative Research Grant, and the China youth 1000-talent grant.

## References

[Alexander and Baptista, 2004] Gordon J Alexander and Alexandre M Baptista. A comparison of var and cvar constraints on portfolio selection with the mean-variance model. Management science, 50(9):1261-1273, 2004.
[Allen-Zhu, 2016] Zeyuan Allen-Zhu. Katyusha: Accelerated variance reduction for faster sgd. ArXiv e-prints, abs/1603.05953, 2016.
[Boyd et al., 2011] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. Foundations and Trends $\mathbb{R}^{\text {® }}$ in Machine Learning, 3(1):1-122, 2011.
[Dai et al., 2016] Bo Dai, Niao He, Yunpeng Pan, Byron Boots, and Le Song. Learning from conditional distributions via dual kernel embeddings. arXiv preprint arXiv:1607.04579, 2016.
[Gabay and Mercier, 1976] Daniel Gabay and Bertrand Mercier. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. Computers \& Mathematics with Applications, 2(1):17-40, 1976.
[Glowinski and Marroco, 1975] Roland Glowinski and A Marroco. Sur l'approximation, par éléments finis d'ordre un, et la résolution, par pénalisation-dualité d'une classe de problèmes de dirichlet non linéaires. Revue française d'automatique, informatique, recherche opérationnelle. Analyse numérique, 9(2):41-76, 1975.
[Golub and Van Loan, 2012] Gene H Golub and Charles F Van Loan. Matrix computations, volume 3. JHU Press, 2012.
[Hien et al., 2016] Le Thi Khanh Hien, Canyi Lu, Huan Xu, and Jiashi Feng. Accelerated stochastic mirror descent algorithms for composite non-strongly convex optimization. arXiv preprint arXiv:1605.06892, 2016.
[Johnson and Zhang, 2013] Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In Advances in Neural Information Processing Systems, pages 315-323, 2013.
[Lian et al., 2016] Xiangru Lian, Mengdi Wang, and Ji Liu. Finite-sum composition optimization via variance reduced gradient descent. arXiv preprint arXiv:1610.04674, 2016.
[Ouyang et al., 2013] Hua Ouyang, Niao He, Long Tran, and Alexander G Gray. Stochastic alternating direction method of multipliers. ICML (1), 28:80-88, 2013.
[Qian and Zhou, 2016] Chao Zhang Hui Qian and Zebang Shen Tengfei Zhou. Accelerated stochastic admm for empirical risk minimization. arXiv preprint arXiv:1611.04074, 2016.
[Shapiro et al., 2014] Alexander Shapiro, Darinka Dentcheva, et al. Lectures on stochastic programming: modeling and theory, volume 16. Siam, 2014.
[Sutton and Barto, 1998] Richard S Sutton and Andrew G Barto. Reinforcement learning: An introduction, volume 1. MIT press Cambridge, 1998.
[Suzuki and others, 2013] Taiji Suzuki et al. Dual averaging and proximal gradient descent for online alternating direction multiplier method. In ICML (1), pages 392-400, 2013.
[Wang and Banerjee, 2013] Huahua Wang and Arindam Banerjee. Online alternating direction method (longer version). arXiv preprint arXiv:1306.3721, 2013.
[Wang et al., 2016a] Mengdi Wang, Ethan X Fang, and Han Liu. Stochastic compositional gradient descent: Algorithms for minimizing compositions of expected-value functions. Mathematical Programming, 161(1-2):419449, 2016.
[Wang et al., 2016b] Mengdi Wang, Ji Liu, and Ethan X. Fang. Accelerating stochastic composition optimization. In Advances In Neural Information Processing Systems, pages 1714-1722, 2016.
[Yu and Huang, 2017] Yue Yu and Longbo Huang. Fast stochastic variance reduced admm for stochastic composition optimization. arXiv preprint arXiv:1705.04138, 2017.
[Zhang et al., 2011] Xiaoqun Zhang, Martin Burger, and Stanley Osher. A unified primal-dual algorithm framework based on bregman iteration. Journal of Scientific Computing, 46(1):20-46, 2011.
[Zheng and Kwok, 2016a] Zheng and Kwok. Fast-and-light stochastic admm. arXiv preprint arXiv:1604.07070, 2016.
[Zheng and Kwok, 2016b] Shuai Zheng and James T Kwok. Stochastic variance-reduced admm. arXiv preprint arXiv:1604.07070, 2016.


[^0]:    ${ }^{1}$ The number of inner iterations is constant.
    ${ }^{2}$ Note that ASC-PG is not based on SVRG and does not have inner loops.

[^1]:    ${ }^{3}$ The corresponding proximal term of Algorithm 1 can be viewed to have $G_{k}=I$.

