

Stochastic Non-convex Optimization with Strong High Probability Second-order Convergence

Mingrui Liu

Department of Computer Science, The University of Iowa, Iowa City, IA 52242

mingrui-liu@uiowa.edu

Tianbao Yang

Department of Computer Science, The University of Iowa, Iowa City, IA 52242

tianbao-yang@uiowa.edu

Abstract

In this paper, we study stochastic non-convex optimization with non-convex random functions. Recent studies on non-convex optimization revolve around establishing second-order convergence, i.e., converging to a nearly second-order optimal stationary points. However, existing results on stochastic non-convex optimization are limited, especially with a high probability second-order convergence. We propose a novel updating step (named NCG-S) by leveraging a stochastic gradient and a noisy negative curvature of a stochastic Hessian, where the stochastic gradient and Hessian are based on a proper mini-batch of random functions. Building on this step, we develop two algorithms and establish their high probability second-order convergence. To the best of our knowledge, the proposed stochastic algorithms are the first with a second-order convergence in *high probability* and a time complexity that is *almost linear* in the problem’s dimensionality.

1 Introduction

In this paper, we consider the following stochastic optimization problem:

$$\min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x}) = \mathbb{E}_{\xi} [f(\mathbf{x}; \xi)], \quad (1)$$

where $f(\mathbf{x}; \xi)$ is a random function but not necessarily convex. The above formulation plays an important role for solving many machine learning problems, e.g., deep learning [12].

A prevalent algorithm for solving the problem is stochastic gradient descent (SGD) [10]. However, SGD can only guarantee convergence to a first-order stationary point (i.e., $\|\nabla f(\mathbf{x})\| \leq \epsilon_1$, where $\|\cdot\|$ denotes the Euclidean norm) for non-convex optimization, which could be a saddle point. A potential solution to address this issue is to find a nearly second-order stationary point \mathbf{x} such that $\|\nabla f(\mathbf{x})\| \leq \epsilon_1 \ll 1$, and $-\lambda_{\min}(\nabla^2 f(\mathbf{x})) \leq \epsilon_2 \ll 1$, where $\lambda_{\min}(\cdot)$ denotes the smallest eigenvalue. When the objective function is non-degenerate (e.g., strict saddle [9] or whose Hessian at all saddle points has a negative eigenvalue), an approximate second-order stationary point is close to a local minimum.

Although there emerged a number of algorithms for finding a nearly second-order stationary point for non-convex optimization with a deterministic function [13, 7, 5, 6, 1, 4, 16], results for stochastic non-convex optimization are still limited. There are three closely related works [9, 18, 2]. A summary of algorithms in these works and their convergence results is presented in Table 1. It is notable that Natasha2, which involves switch between several sub-routines including SGD, a degenerate version of Natasha1.5 for finding a first-order stationary point, and an online power method (i.e., the Oja’s algorithm [14]) for computing the negative curvature (i.e., the eigen-vector corresponding to the minimum eigen-value) of the Hessian matrix, is more complex than noisy SGD and SGLD.

In this paper, we propose new stochastic optimization algorithms for solving (1). Similar to several existing algorithms, we also use the negative curvature to escape from saddle points. The key difference is that we compute a noisy negative curvature based on a proper mini-batch of sampled random functions. A novel updating step is proposed that follows a stochastic gradient or the noisy negative curvature depending on which decreases the objective value most. Building on this step, we present two algorithms that have different time complexities. A summary of our results and comparison with previous similar results are presented in Table 1. To the best of our knowledge, the proposed algorithms are the first for stochastic non-convex optimization with a second-order convergence in *high probability* and a time complexity that is *almost linear* in the problem’s dimensionality. It is also notable that our result is much stronger than the mini-batch SGD analyzed in [11] for stochastic non-convex optimization in that (i) we use the same number of IFO as in [11] but achieve the

Table 1: Comparison with existing stochastic algorithms for achieving an (ϵ_1, ϵ_2) -second-order stationary solution to (1), where p is a number at least 4, IFO (incremental first-order oracle) and ISO (incremental second-order oracle) are terminologies borrowed from [15], representing $\nabla f(\mathbf{x}; \xi)$ and $\nabla^2 f(\mathbf{x}; \mathbf{v})$ respectively, T_h denotes the runtime of ISO and T_g denotes the runtime of IFO. The proposed algorithms SNCG have two variants with different time complexities, where the result marked with $*$ has a practical improvement detailed later.

algo.	oracle	second-order guarantee in expectation or high probability	time complexity
Noisy SGD [9]	IFO	$(\epsilon, \epsilon^{1/4})$, expectation	$\tilde{O}(T_g d^p \epsilon^{-4})$
SGLD [18]	IFO	$(\epsilon, \epsilon^{1/2})$, high probability	$\tilde{O}(T_g d^p \epsilon^{-4})$
Natasha2 [2]	IFO + ISO	$(\epsilon, \epsilon^{1/2})$, expectation	$\tilde{O}(T_g \epsilon^{-3.5} + T_h \epsilon^{-2.5})$
SNCG	IFO + ISO	$(\epsilon, \epsilon^{1/2})$, high probability	$\tilde{O}(T_g \epsilon^{-4} + T_h \epsilon^{-3})^*$ $\tilde{O}(T_g \epsilon^{-4} + T_h \epsilon^{-2.5})$

second-order convergence using a marginal number of ISO; (ii) our high probability convergence is for a solution from a single run of the proposed algorithms instead of from multiple runs and using a boosting technique as in [11].

Before moving to the next section, we would like to remark that stochastic algorithms with second-order convergence result are recently proposed for solving a finite-sum problem [15], which alternates between a first-order sub-routine (e.g., stochastic variance reduced gradient) and a second-order sub-routine (e.g., Hessian descent). Since full gradients are computed occasionally, they are not applicable to the general stochastic non-convex optimization problem (1) and hence are excluded from comparison. Nevertheless, our idea of the proposed NCG-S step that lets negative curvature descent competes with the gradient descent can be borrowed to reduce the number of stochastic Hessian-vector products in their Hessian descent. We will elaborate this point later.

2 Preliminaries and Building Blocks

Our goal is to find an (ϵ_1, ϵ_2) -second order stationary point \mathbf{x} such that $\|\nabla f(\mathbf{x})\| \leq \epsilon_1$, and $\lambda_{\min}(\nabla^2 f(\mathbf{x})) \geq -\epsilon_2$. To this end, we make the following assumptions regarding (1).

Assumption 1. (i) Every random function $f(\mathbf{x}; \xi)$ is twice differentiable, and it has Lipschitz continuous gradient, i.e., there exists $L_1 > 0$ such that $\|\nabla f(\mathbf{x}; \xi) - \nabla f(\mathbf{y}; \xi)\| \leq L_1 \|\mathbf{x} - \mathbf{y}\|$, (ii) $f(\mathbf{x})$ has Lipschitz continuous Hessian, i.e., there exists $L_2 > 0$ such that $\|\nabla^2 f(\mathbf{x}) - \nabla^2 f(\mathbf{y})\|_2 \leq L_2 \|\mathbf{x} - \mathbf{y}\|$, (iii) given an initial point \mathbf{x}_0 , there exists $\Delta < \infty$ such that $f(\mathbf{x}_0) - f(\mathbf{x}_*) \leq \Delta$, where \mathbf{x}_* denotes the global minimum of $f(\mathbf{x})$; (iv) there exists $G > 0$ such that $\mathbb{E}[\exp(\|\nabla f(\mathbf{x}; \xi) - \nabla f(\mathbf{x})\|/G)] \leq \exp(1)$ holds.

Remark: The first three assumptions are standard assumptions for non-convex optimization in order to establish second-order convergence. The last assumption is standard for stochastic optimization necessary for high probability analysis.

The proposed algorithms require noisy first-order information at each iteration and maybe noisy second-order information. We first discuss approaches to compute these information, which will lead us to the updating step NCG-S. To compute noisy first-order information, we use incremental first-order oracle (IFO) that takes \mathbf{x} as input and returns $\nabla f(\mathbf{x}; \xi)$. In particular, at a point \mathbf{x} we sample a set of random variables $\mathcal{S}_1 = \{\xi_1, \xi_2, \dots, \}$ and compute a stochastic gradient $\mathbf{g}(\mathbf{x}) = \frac{1}{|\mathcal{S}_1|} \sum_{\xi_i \in \mathcal{S}_1} \nabla f(\mathbf{x}; \xi_i)$ such that $\|\mathbf{g}(\mathbf{x}) - \nabla f(\mathbf{x})\| \leq \epsilon_4 \leq \min(\frac{1}{2\sqrt{2}}\epsilon_1, \epsilon_2^2/(24L_2))$ holds with high probability. This can be guaranteed by the following lemma.

Lemma 1. Suppose Assumption 1 (iv) holds. Let $\mathbf{g}(\mathbf{x}) = \frac{1}{|\mathcal{S}_1|} \sum_{\xi_i \in \mathcal{S}_1} \nabla f(\mathbf{x}; \xi_i)$. For any $\epsilon_4, \delta \in (0, 1)$, $\mathbf{x} \in \mathbb{R}^d$, when $|\mathcal{S}_1| \geq \frac{4G^2(1+3\log^2(1/\delta))}{\epsilon_4^2}$, we have $\Pr(\|\mathbf{g}(\mathbf{x}) - \nabla f(\mathbf{x})\| \leq \epsilon_4) \geq 1 - \delta$.

The lemma can be proved by using large deviation theorem of vector-valued martingales (e.g., see [11][Lemma 4]).

To compute noisy second-order information, we calculate a noisy negative curvature of a stochastic Hessian that is sufficiently close to the true Hessian. In particular, at a point \mathbf{x} we sample a set of random variables $\mathcal{S}_2 = \{\xi'_1, \xi'_2, \dots, \}$ and compute a noisy negative curvature \mathbf{v} of the stochastic Hessian $H(\mathbf{x}) = \frac{1}{|\mathcal{S}_2|} \sum_{\xi'_i \in \mathcal{S}_2} \nabla^2 f(\mathbf{x}; \xi'_i)$, where $|\mathcal{S}_2|$ is sufficiently large such that $\|H(\mathbf{x}) - \nabla^2 f(\mathbf{x})\|_2 \leq \epsilon_3 \leq \epsilon_2/24$

Algorithm 1: The stochastic NCG step: $(\mathbf{x}^+, \mathbf{v}^\top H(\mathbf{x})\mathbf{v}) = \text{NCG-S}(\mathbf{x}, \varepsilon, \delta, \epsilon_1, \epsilon_2)$

- 1 **Input:** $\mathbf{x}, \varepsilon, \delta, \epsilon_1, \epsilon_2$;
 - 2 let $\mathbf{g}(\mathbf{x})$ and $H(\mathbf{x})$ be a stochastic gradient and Hessian according to Lemma 1 and 2;
 - 3 Find a unit vector \mathbf{v} such that $\lambda_{\min}(H(\mathbf{x})) \geq \mathbf{v}^\top H(\mathbf{x})\mathbf{v} - \varepsilon$ according to Lemma 3;
 - 4 **if** $-\frac{\epsilon_2^2}{2L_2^2} \mathbf{v}^\top H(\mathbf{x})\mathbf{v} - \frac{11\epsilon_3^3}{48L_2^2} > \frac{\|\mathbf{g}(\mathbf{x})\|^2}{4L_1} - \frac{\epsilon_1^2}{8L_1}$ **then**
 - 5 | Compute $\mathbf{x}^+ = \mathbf{x} - \frac{\epsilon_2}{L_2} \text{sign}(\mathbf{v}^\top \mathbf{g}(\mathbf{x}))\mathbf{v}$;
 - 6 **else**
 - 7 | Compute $\mathbf{x}^+ = \mathbf{x} - \frac{1}{L_1} \mathbf{g}(\mathbf{x})$;
 - 8 **return** $\mathbf{x}^+, \mathbf{v}^\top H(\mathbf{x})\mathbf{v}$
-

holds with high probability, where $\|\cdot\|_2$ denotes the spectral norm of a matrix. This can be guaranteed according to the following lemma.

Lemma 2. *Suppose Assumption 1 (i) holds. Let $H(\mathbf{x}) = \frac{1}{|\mathcal{S}_2|} \sum_{\xi_i \in \mathcal{S}_2} \nabla^2 f(\mathbf{x}; \xi_i)$. For any $\epsilon_3, \delta \in (0, 1)$, $\mathbf{x} \in \mathbb{R}^d$, when $|\mathcal{S}_2| \geq \frac{16L_2^2}{\epsilon_3^2} \log(\frac{2d}{\delta})$, we have $\Pr(\|H(\mathbf{x}) - \nabla^2 f(\mathbf{x})\|_2 \leq \epsilon_3) \geq 1 - \delta'$.*

The above lemma can be proved by using matrix concentration inequalities. Please see [17][Lemma 4] for a proof. To compute a noisy negative curvature of $H(\mathbf{x})$, we can leverage approximate PCA algorithms [3, 8] using the incremental second-order oracle (ISO) that can compute $\nabla^2 f(\mathbf{x}; \xi)\mathbf{v}$.

Lemma 3. *Let $H = \frac{1}{m} \sum_{i=1}^m H_i$ where $\|H_i\|_2 \leq L_1$. There exists a randomized algorithm \mathcal{A} such that with probability at least $1 - \delta$, \mathcal{A} produces a unit vector \mathbf{v} satisfying $\lambda_{\min}(H) \geq \mathbf{v}^\top H\mathbf{v} - \varepsilon$ with a time complexity of $\tilde{O}(T_h^1 \max\{m, m^{3/4} \sqrt{L_1/\varepsilon}\})$, where T_h denotes the time of computing $H_i\mathbf{v}$ and \tilde{O} suppresses a logarithmic term in $d, 1/\delta, 1/\varepsilon$.*

NCG-S: the updating step. With the approaches for computing noisy first-order and second-order information, we present a novel updating step called NCG-S in Algorithm 1, which uses a competing idea that takes a step along the noisy negative gradient direction or the noisy negative curvature direction depending on which decreases the objective value more. One striking feature of NCG-S is that the noise level in computing a noisy negative curvature of $H(\mathbf{x})$ is set to a free parameter ε instead of the target accuracy level ϵ_2 as in many previous works [1, 4, 17], which allows us to design an algorithm with a much reduced number of ISO calls in practice. The following lemma justifies the fact of sufficient decrease in terms of the objective value of each NCG-S step.

Lemma 4. *Suppose Assumption 1 holds. Conditioned on the event $\mathcal{A} = \{\|H(\mathbf{x}_j) - \nabla^2 f(\mathbf{x}_j)\|_2 \leq \epsilon_3\} \cap \{\|\mathbf{g}(\mathbf{x}_j) - \nabla f(\mathbf{x}_j)\| \leq \epsilon_4\}$ where $\epsilon_3 \leq \epsilon_2/24$ and $\epsilon_4 \leq \min(\frac{1}{2\sqrt{2}}\epsilon_1, \epsilon_2^2/(24L_2))$, the update $\mathbf{x}_{j+1} = \text{NCG-S}(\mathbf{x}_j, \varepsilon, \delta, \epsilon_2)$ satisfies $f(\mathbf{x}_j) - f(\mathbf{x}_{j+1}) \geq \max\left(\frac{1}{4L_1} \|\mathbf{g}(\mathbf{x}_j)\|^2 - \frac{\epsilon_1^2}{8L_1}, \frac{-\epsilon_2^2 \mathbf{v}_j^\top H(\mathbf{x}_j)\mathbf{v}_j}{2L_2^2} - \frac{11\epsilon_3^3}{48L_2^2}\right)$.*

3 The Proposed Algorithms: SNCG

In this section, we present two variants of the proposed algorithms based on the NCG-S step shown in Algorithm 2 and Algorithm 3. The differences of these two variants are (i) SNCG-1 uses NCG-S at every iteration to update the solution, while SNCG-2 only uses NCG-S when the approximate gradient's norm is small; (ii) the noise level ε for computing the noisy negative curvature (as in Lemma 3) in SNCG-1 is set to $\max(\epsilon_2, \|\mathbf{g}(\mathbf{x}_j)\|^\alpha)/2$ adaptive to the magnitude of the stochastic gradient, where $\alpha \in (0, 1]$ is a parameter that characterizes $\epsilon_2 = \epsilon_1^\alpha$. In contrast, the noise level ε in SNCG-2 is simply set to $\epsilon_2/2$. These differences lead to different time complexities of the two algorithms.

Theorem 1. *Suppose Assumption 1 holds, $\epsilon_3 \leq \epsilon_2/24$ and $\epsilon_4 \leq \min(\frac{1}{2\sqrt{2}}\epsilon_1, \epsilon_2^2/(24L_2))$. With probability $1 - \delta$, SNCG-1 terminates with at most $[1 + \max(\frac{48L_2^2}{\epsilon_3^2}, \frac{8L_1}{\epsilon_1^2}) \Delta]$ NCG-S steps, and furthermore, each NCG-S step requires time in the order of $\tilde{O}\left(T_h|\mathcal{S}_2| + T_h|\mathcal{S}_2|^{3/4} \frac{\sqrt{L_1}}{\max(\epsilon_2, \|\mathbf{g}(\mathbf{x}_j)\|^\alpha)^{1/2}} + |\mathcal{S}_1|T_g\right)$; SNCG-2 terminates with at most $\frac{8L_1}{\epsilon_1^2} \Delta$ SG steps and at most $(1 + \frac{48L_2^2}{\epsilon_3^2}) \Delta$ NSG-S steps, each NCG-S step requires time in the order of $\tilde{O}\left(T_h|\mathcal{S}_2| + T_h|\mathcal{S}_2|^{3/4} \frac{\sqrt{L_1}}{\epsilon_2^{1/2}} + |\mathcal{S}_1|T_g\right)$. Upon termination, with probability $1 - 3\delta$, both algorithms return a solution \mathbf{x}_{j^*} such that $\|\nabla f(\mathbf{x}_{j^*})\| \leq 2\epsilon_1$ and $\lambda_{\min}(\nabla^2 f(\mathbf{x}_{j^*})) \geq -2\epsilon_2$.*

Algorithm 2: SNCG-1: $(\mathbf{x}_0, \epsilon_1, \alpha, \delta)$

1 **Input:** $\mathbf{x}_0, \epsilon_1, \alpha, \delta$
2 Set $\mathbf{x}_1 = \mathbf{x}_0, \epsilon_2 = \epsilon_1^\alpha, \delta' = \delta / (1 + \max(\frac{48L_2^2}{\epsilon_2^3}, \frac{8L_1}{\epsilon_1^2}) \Delta)$
3 **for** $j = 1, 2, \dots$, **do**
4 $(\mathbf{x}_{j+1}, \mathbf{v}_j^\top H(\mathbf{x}_j) \mathbf{v}_j) = \text{NCG-S}(\mathbf{x}_j, \max(\epsilon_2, \|\mathbf{g}(\mathbf{x}_j)\|^\alpha) / 2, \delta', \epsilon_1, \epsilon_2)$
5 **if** $\mathbf{v}_j^\top H(\mathbf{x}_j) \mathbf{v}_j > -\epsilon_2 / 2$ **and** $\|\mathbf{g}(\mathbf{x}_j)\| \leq \epsilon_1$ **then**
6 return \mathbf{x}_j

Algorithm 3: SNCG-2: $(\mathbf{x}_0, \epsilon_1, \delta)$

1 **Input:** $\mathbf{x}_0, \epsilon_1, \delta$
2 Set $\mathbf{x}_1 = \mathbf{x}_0, \delta' = \delta / (1 + \max(\frac{48L_2^2}{\epsilon_2^3}, \frac{8L_1}{\epsilon_1^2}) \Delta)$
3 **for** $j = 1, 2, \dots$, **do**
4 Compute $\mathbf{g}(\mathbf{x}_j)$ according to Lemma 1
5 **if** $\|\mathbf{g}(\mathbf{x}_j)\| \geq \epsilon_1$ **then**
6 compute $\mathbf{x}_{j+1} = \mathbf{x}_j - \frac{1}{L_1} \mathbf{g}(\mathbf{x}_j)$ // SG step
7 **else**
8 compute $(\mathbf{x}_{j+1}, \mathbf{v}_j^\top H(\mathbf{x}_j) \mathbf{v}_j) = \text{NCG-S}(\mathbf{x}_j, \epsilon_2 / 2, \delta', \epsilon_1, \epsilon_2)$
9 **if** $\mathbf{v}_j^\top H(\mathbf{x}_j) \mathbf{v}_j > -\epsilon_2 / 2$ **then**
10 return \mathbf{x}_j

Remark: To analyze the time complexity, we can plug in the order of $|\mathcal{S}_1|$ and $|\mathcal{S}_2|$ as in Lemma 1 and Lemma 2. It is not difficult to show that when $\epsilon_2 = \sqrt{\epsilon_1}$, the worst-case time complexities of these two algorithms are given in Table 1, where the result marked by * corresponds to SNCG-1. However, this worst-case result is computed by simply bounding $T_h / \sqrt{\max(\epsilon_2, \|\mathbf{g}(\mathbf{x})\|^\alpha)}$ by $T_h / \sqrt{\epsilon_2}$. In practice, before reaching a saddle point (i.e., $\|\mathbf{g}(\mathbf{x}_j)\| \geq \epsilon_1$), the number of ISO calls for each NCG-S step in SNCG-1 can be less than that of each NCG-S step in SNCG-2. In addition, the NCG-S step in SNCG-1 can be faster than the SG step in SNCG-2 before reaching a saddle point. More importantly, the idea of competing between gradient descent and negative curvature descent and the adaptive noise parameter ϵ for computing the noisy negative curvature can be also useful in other algorithms. For example, in [15] the Hessian descent (also known as negative curvature descent) can take the competing idea and uses adaptive noise level for computing a noisy negative curvature.

4 Conclusion

In this paper, we have proposed new algorithms for stochastic non-convex optimization with strong high probability second-order convergence guarantee. To the best of our knowledge, the proposed stochastic algorithms are the first with a second-order convergence in high probability and a time complexity that is almost linear in the problem's dimensionality.

References

- [1] N. Agarwal, Z. Allen Zhu, B. Bullins, E. Hazan, and T. Ma. Finding approximate local minima faster than gradient descent. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 1195–1199, 2017.
- [2] Z. Allen-Zhu. Natasha 2: Faster non-convex optimization than sgd. *CoRR*, /abs/1708.08694, 2017.
- [3] Z. Allen Zhu and Y. Li. Even faster SVD decomposition yet without agonizing pain. In *Advances in Neural Information Processing Systems 29 (NIPS)*, pages 974–982, 2016.
- [4] Y. Carmon, J. C. Duchi, O. Hinder, and A. Sidford. Accelerated methods for non-convex optimization. *CoRR*, abs/1611.00756, 2016.

- [5] C. Cartis, N. I. M. Gould, and P. L. Toint. Adaptive cubic regularisation methods for unconstrained optimization. part i: motivation, convergence and numerical results. *Mathematical Programming*, 127(2):245–295, Apr 2011.
- [6] C. Cartis, N. I. M. Gould, and P. L. Toint. Adaptive cubic regularisation methods for unconstrained optimization. part ii: worst-case function- and derivative-evaluation complexity. *Mathematical Programming*, 130(2):295–319, Dec 2011.
- [7] A. Conn, N. Gould, and P. Toint. *Trust Region Methods*. MPS-SIAM Series on Optimization. Society for Industrial and Applied Mathematics, 2000.
- [8] D. Garber, E. Hazan, C. Jin, S. M. Kakade, C. Musco, P. Netrapalli, and A. Sidford. Faster eigenvector computation via shift-and-invert preconditioning. In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, pages 2626–2634, 2016.
- [9] R. Ge, F. Huang, C. Jin, and Y. Yuan. Escaping from saddle points — online stochastic gradient for tensor decomposition. In P. Grünwald, E. Hazan, and S. Kale, editors, *Proceedings of The 28th Conference on Learning Theory (COLT)*, volume 40, pages 797–842. PMLR, 03–06 Jul 2015.
- [10] S. Ghadimi and G. Lan. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- [11] S. Ghadimi, G. Lan, and H. Zhang. Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization. *Math. Program.*, 155(1-2):267–305, Jan. 2016.
- [12] I. Goodfellow, Y. Bengio, and A. Courville. *Deep learning*. MIT press, 2016.
- [13] Y. Nesterov and B. T. Polyak. Cubic regularization of newton method and its global performance. *Mathematical Programming*, 108(1):177–205, 2006.
- [14] E. Oja. Simplified neuron model as a principal component analyzer. *Journal of mathematical biology*, 15(3):267–273, 1982.
- [15] S. J. Reddi, M. Zaheer, S. Sra, B. Póczos, F. Bach, R. Salakhutdinov, and A. J. Smola. A generic approach for escaping saddle points. *arXiv preprint arXiv:1709.01434*, 2017.
- [16] C. W. Royer and S. J. Wright. Complexity analysis of second-order line-search algorithms for smooth nonconvex optimization. *arXiv preprint arXiv:1706.03131*, 2017.
- [17] P. Xu, F. Roosta-Khorasani, and M. W. Mahoney. Newton-type methods for non-convex optimization under inexact hessian information. *CoRR*, abs/1708.07164, 2017.
- [18] Y. Zhang, P. Liang, and M. Charikar. A hitting time analysis of stochastic gradient langevin dynamics. In *Proceedings of the 30th Conference on Learning Theory (COLT)*, pages 1980–2022, 2017.

A Proof of Lemma 3

We first introduce a proposition, which is the Theorem 2.5 in [1].

Proposition 1. *Let $M \in \mathbb{R}^{d \times d}$ be a symmetric matrix with eigenvalues $1 \geq \lambda_1 \dots \geq \lambda_d \geq 0$. Then with probability at least $1 - p$, the Algorithm AppxPCA produces a unit vector \mathbf{v} such that $\mathbf{v}^\top M \mathbf{v} \geq (1 - \delta_+)(1 - \epsilon)\lambda_{\max}(M)$. The total running time is $\tilde{O}\left(T_h^1 \max\{m, \frac{m^{3/4}}{\sqrt{\epsilon}}\} \log^2\left(\frac{1}{\epsilon^2 \delta_+}\right)\right)$.*

Proof of Lemma 3. Define $M = I - \frac{H}{L_1}$, then M satisfies the condition in the Proposition 1. Then we know that with probability at least $1 - p$, the Algorithm AppxPCA produces a vector \mathbf{v} satisfying

$$\mathbf{v}^\top \left(I - \frac{H}{L_1} \right) \mathbf{v} \geq (1 - \delta_+)(1 - \epsilon) \left(1 - \frac{\lambda_{\min}(H)}{L_1} \right),$$

which implies that

$$L_1 - \mathbf{v}^\top H \mathbf{v} \geq (1 - \delta_+ - \epsilon + \delta_+ \epsilon)(L_1 - \lambda_{\min}(H)) \geq (1 - \delta_+ - \epsilon)(L_1 - \lambda_{\min}(H)).$$

By simple algebra, we have

$$\begin{aligned} \lambda_{\min}(H) &\geq \mathbf{v}^\top H \mathbf{v} - (\delta_+ + \epsilon)(L_1 - \lambda_{\min}(H)) \\ &\geq \mathbf{v}^\top H \mathbf{v} - 2L_1(\delta_+ + \epsilon). \end{aligned}$$

By setting $\epsilon = \delta_+ = \frac{\epsilon}{4L_1}$, we can finish the proof. \square

B Proof of Lemma 4

Proof. Define $\eta_j = \frac{\epsilon_2}{L_2} \text{sign}(\mathbf{v}_j^\top \mathbf{g}(\mathbf{x}_j))$. Next, we analyze the objective decrease for j -th NCG-S step conditioned on the event $\mathcal{A} = \{\|H(\mathbf{x}_j) - \nabla^2 f(\mathbf{x}_j)\|_2 \leq \epsilon_3 \cap \|\mathbf{g}(\mathbf{x}_j) - \nabla f(\mathbf{x}_j)\| \leq \epsilon_4\}$ and let $\Pr(\mathcal{A}) = 1 - \delta'$.

By L_2 -Lipschitz continuity of Hessian, we know that

$$f(\mathbf{x}_{j+1}^1) - f(\mathbf{x}_j) \leq -\eta_j \nabla f(\mathbf{x}_j)^\top \mathbf{v}_j + \frac{\eta_j^2}{2} \mathbf{v}_j^\top (\nabla^2 f(\mathbf{x}_j) - H(\mathbf{x}_j)) \mathbf{v}_j + \frac{\eta_j^2}{2} \mathbf{v}_j^\top H(\mathbf{x}_j) \mathbf{v}_j + \frac{L_2}{6} |\eta_j|^3.$$

where \mathbf{x}_{j+1}^1 is an update of \mathbf{x}_j following \mathbf{v}_j in NCG-S. Note that $\epsilon_4 \leq \epsilon_2^2/24L_2$, and then we have

$$-\eta_j \nabla f(\mathbf{x}_j)^\top \mathbf{v}_j = -\eta_j \mathbf{g}(\mathbf{x}_j)^\top \mathbf{v}_j + \eta_j (\mathbf{g}(\mathbf{x}_j) - \nabla f(\mathbf{x}_j))^\top \mathbf{v}_j \leq |\eta_j \epsilon_4| \leq \frac{\epsilon_2^3}{24L_2^2} \quad (2)$$

$$\mathbf{v}_j^\top (\nabla^2 f(\mathbf{x}_j) - H(\mathbf{x}_j)) \mathbf{v}_j \leq \epsilon_3 \leq \epsilon_2/24 \quad (3)$$

Then it follows that

$$f(\mathbf{x}_{j+1}^1) - f(\mathbf{x}_j) \leq \frac{\epsilon_2^3}{24L_2^2} + \frac{\epsilon_2^3}{48L_2^2} + \frac{\epsilon_2^2 \mathbf{v}_j^\top H(\mathbf{x}_j) \mathbf{v}_j}{2L_2^2} + \frac{\epsilon_2^3}{6L_2^2} = - \underbrace{\left(\frac{-\epsilon_2^2 \mathbf{v}_j^\top H(\mathbf{x}_j) \mathbf{v}_j}{2L_2^2} - \frac{11\epsilon_2^3}{48L_2^2} \right)}_{\Delta_1}. \quad (4)$$

Similarly, let \mathbf{x}_{j+1}^2 denote an update of \mathbf{x}_j following $\mathbf{g}(\mathbf{x}_j)$ in NCG-S, we have

$$\begin{aligned} f(\mathbf{x}_{j+1}^2) - f(\mathbf{x}_j) &\leq (\mathbf{x}_{j+1}^2 - \mathbf{x}_j)^\top \nabla f(\mathbf{x}_j) + \frac{L_1}{2} \|\mathbf{x}_{j+1}^2 - \mathbf{x}_j\|^2 \\ &= -\frac{1}{L_1} \mathbf{g}(\mathbf{x}_j)^\top \nabla f(\mathbf{x}_j) + \frac{\|\mathbf{g}(\mathbf{x}_j)\|^2}{2L_1} \\ &= -\frac{1}{L_1} \mathbf{g}(\mathbf{x}_j)^\top \mathbf{g}(\mathbf{x}_j) + \frac{1}{L_1} \mathbf{g}(\mathbf{x}_j)^\top (\mathbf{g}(\mathbf{x}_j) - \nabla f(\mathbf{x}_j)) + \frac{\|\mathbf{g}(\mathbf{x}_j)\|^2}{2L_1} \\ &\leq -\frac{1}{2L_1} \|\mathbf{g}(\mathbf{x}_j)\|^2 + \frac{1}{4L_1} \|\mathbf{g}(\mathbf{x}_j)\|^2 + \frac{1}{L_1} \|\mathbf{g}(\mathbf{x}_j) - \nabla f(\mathbf{x}_j)\|^2 \\ &= -\frac{1}{4L_1} \|\mathbf{g}(\mathbf{x}_j)\|^2 + \frac{1}{L_1} \epsilon_4^2 \leq - \underbrace{\frac{1}{4L_1} \|\mathbf{g}(\mathbf{x}_j)\|^2 + \frac{\epsilon_1^2}{8L_1}}_{-\Delta_2} \end{aligned}$$

where we use $\epsilon_4 \leq \frac{1}{2\sqrt{2}}\epsilon_1$. According to the update of NCG-S, if $\Delta_1 > \Delta_2$, we have $\mathbf{x}_{j+1} = \mathbf{x}_{j+1}^1$ and then $f(\mathbf{x}_j) - f(\mathbf{x}_{j+1}) \geq \Delta_1 = \max(\Delta_1, \Delta_2)$. If $\Delta_2 \geq \Delta_1$, we have $\mathbf{x}_{j+1} = \mathbf{x}_{j+1}^2$ and then $f(\mathbf{x}_j) - f(\mathbf{x}_{j+1}) \geq \Delta_2 = \max(\Delta_1, \Delta_2)$. Therefore, with probability $1 - \delta'$ we have,

$$f(\mathbf{x}_j) - f(\mathbf{x}_{j+1}) \geq \max\left(\frac{1}{4L_1}\|\mathbf{g}(\mathbf{x}_j)\|^2 - \frac{\epsilon_1^2}{8L_1}, \frac{-\epsilon_2^2\mathbf{v}_j^\top H(\mathbf{x}_j)\mathbf{v}_j}{2L_2^2} - \frac{11\epsilon_2^3}{48L_2^2}\right)$$

□

C Proof of Theorem 1

Proof. • We first prove the result of SNCG-1. For the j -th NCG-S step, define the event $\mathcal{A} = \{\|H(\mathbf{x}_j) - \nabla^2 f(\mathbf{x}_j)\|_2 \leq \epsilon_3\} \cap \{\|\mathbf{g}(\mathbf{x}_j) - \nabla f(\mathbf{x}_j)\| \leq \epsilon_4\}$ and let $\Pr(\mathcal{A}) = 1 - \delta'$ (we can choose ϵ_3 and ϵ_4 to make it hold). Since the Algorithm SCNG-1 calls NCG-S as a subroutine, then by Lemma 4, we know that with probability at least $1 - \delta'$,

$$f(\mathbf{x}_j) - f(\mathbf{x}_{j+1}) \geq \max\left(\frac{1}{4L_1}\|\mathbf{g}(\mathbf{x}_j)\|^2 - \frac{\epsilon_1^2}{8L_1}, \frac{-\epsilon_2^2\mathbf{v}_j^\top H(\mathbf{x}_j)\mathbf{v}_j}{2L_2^2} - \frac{11\epsilon_2^3}{48L_2^2}\right).$$

If $\mathbf{v}_j^\top H(\mathbf{x}_j)\mathbf{v}_j \leq -\epsilon_2/2$, we have $\Delta_1 \geq \frac{\epsilon_2^3}{48L_2^2}$ and

$$f(\mathbf{x}_j) - f(\mathbf{x}_{j+1}) \geq \frac{\epsilon_2^3}{48L_2^2}$$

If $\|\mathbf{g}(\mathbf{x}_j)\| > \epsilon_1$, we have $\Delta_2 \geq \frac{\epsilon_1^2}{8L_1}$ and

$$f(\mathbf{x}_j) - f(\mathbf{x}_{j+1}) \geq \frac{\epsilon_1^2}{8L_1}$$

Therefore, before the algorithm terminates, i.e., for all iterations $j \leq j_* - 1$, we have either $\mathbf{v}_j^\top H(\mathbf{x}_j)\mathbf{v}_j \leq -\epsilon_2/2$ or $\|\mathbf{g}(\mathbf{x}_j)\| > \epsilon_1$. In either case, the following holds with probability $1 - \delta'$

$$f(\mathbf{x}_j) - f(\mathbf{x}_{j+1}) \geq \min\left(\frac{\epsilon_1^2}{8L_1}, \frac{\epsilon_2^3}{48L_2^2}\right),$$

from which we can derive the upper bound of j_* , which is $j_* \leq [1 + \max\left(\frac{48L_2^2}{\epsilon_2^3}, \frac{8L_1}{\epsilon_1^2}\right)\Delta]$. Next, we show that upon termination, we achieve an $(2\epsilon_1, 2\epsilon_2)$ -second order stationary point with high probability. In particular, with probability $1 - \delta'$ we have

$$\|\nabla f(\mathbf{x}_{j_*})\| \leq \|\nabla f(\mathbf{x}_{j_*}) - \mathbf{g}(\mathbf{x}_{j_*})\| + \|\mathbf{g}(\mathbf{x}_{j_*})\| \leq \epsilon_4 + \epsilon_1 \leq 2\epsilon_1.$$

and with probability $1 - \delta'$

$$\lambda_{\min}(H(\mathbf{x}_{j_*})) \geq \mathbf{v}_{j_*}^\top H(\mathbf{x}_{j_*})\mathbf{v}_{j_*} - \max(\epsilon_2, \|g(\mathbf{x}_{j_*})\|^\alpha)/2 \geq -\epsilon_2$$

In addition, with probability $1 - \delta'$, we have

$$\lambda_{\min}(\nabla^2 f(\mathbf{x}_{j_*})) \geq \lambda_{\min}(H(\mathbf{x}_{j_*})) - \epsilon_3 \geq -2\epsilon_2$$

As a result, by using union bound, we have with probability $1 - 3j_*\delta' = 1 - 3\delta$, we have

$$\|\nabla f(\mathbf{x}_{j_*})\| \leq 2\epsilon_1, \quad \lambda_{\min}(\nabla^2 f(\mathbf{x}_{j_*})) \geq -2\epsilon_2$$

Finally, the time complexity of each iteration follows Lemma 3.

- The proof of the result of SNCG-2 is similar. For simplicity, we use the same notation unless specified. According to the Algorithm SNCG-2, we know that when $\|\mathbf{g}(\mathbf{x}_j)\| \geq \epsilon_1$, the SG step guarantees that

$$f(\mathbf{x}_{j+1}) - f(\mathbf{x}_j) \leq -\frac{1}{4L_1}\|\mathbf{g}(\mathbf{x}_j)\|^2 + \frac{\epsilon_1^2}{8L_1} \leq -\frac{\epsilon_1^2}{8L_1}$$

When $\mathbf{v}_j^\top H(\mathbf{x})\mathbf{v}_j \leq -\epsilon_2/2$, then the NCG-S step guarantees that

$$f(\mathbf{x}_{j+1}) - f(\mathbf{x}_j) \leq -\max\left(\frac{\epsilon_2^2}{48L_2^2}, \frac{1}{4L_1}\|\mathbf{g}(\mathbf{x}_j)\|^2 - \frac{\epsilon_1^2}{8L_1}\right) \leq -\frac{\epsilon_2^3}{48L_2^2}$$

According to the update rule, it is easy to see that conditioned on the event \mathcal{A} , the SG step always decrease the objective value (with high probability) and the NCG-S step decreases the objective value at all steps except for the very last iteration (with high probability). Denote j_* by the number of iterations in the Algorithm SNCG-2. By the sufficient decrease argument, we know that with probability at least $1 - j_*\delta'$, the algorithm terminates, where

$$j_* \leq 1 + \max\left(\frac{48L_2^2}{\epsilon_2^3}, \frac{8L_1}{\epsilon_1^2}\right).$$

By the relationship between δ' and δ , we know that the algorithm terminates with probability at least $1 - \delta$.

Note that $f(\mathbf{x}_0) - f(\mathbf{x}_*) \leq \Delta$, and hence with probability $1 - j_*\delta'$, we have at most $\frac{8L_1}{\epsilon_1^2}\Delta$ stochastic gradient evaluations and at most $(1 + \frac{48L_2^3}{\epsilon_2^3})\Delta$ stochastic Hessian-vector product evaluations before termination.

Next, we show that upon termination, we achieve an $(2\epsilon_1, 2\epsilon_2)$ -second order stationary point with high probability. In particular, with probability $1 - \delta'$ we have

$$\|\nabla f(\mathbf{x}_{j_*})\| \leq \|\nabla f(\mathbf{x}_{j_*}) - \mathbf{g}(\mathbf{x}_{j_*})\| + \|\mathbf{g}(\mathbf{x}_{j_*})\| \leq \epsilon_4 + \epsilon_1 \leq 2\epsilon_1.$$

and with probability $1 - \delta'$

$$\lambda_{\min}(H(\mathbf{x}_{j_*})) \geq \mathbf{v}_{j_*}^\top H(\mathbf{x}_{j_*})\mathbf{v}_{j_*} - \epsilon_2/2 \geq -\epsilon_2$$

In addition, with probability $1 - \delta'$, we have

$$\lambda_{\min}(\nabla^2 f(\mathbf{x}_{j_*})) \geq \lambda_{\min}(H(\mathbf{x}_{j_*})) - \epsilon_3 \geq -2\epsilon_2$$

As a result, by using union bound, we have with probability $1 - 3j_*\delta' = 1 - 3\delta$, we have

$$\|\nabla f(\mathbf{x}_{j_*})\| \leq 2\epsilon_1, \quad \lambda_{\min}(\nabla^2 f(\mathbf{x}_{j_*})) \geq -2\epsilon_2$$

Finally, the time complexity of each iteration follows Lemma 3.

□