

Explicit Convergence Rates of Greedy and Random Quasi-Newton Methods

Dachao Lin

*Academy for Advanced Interdisciplinary Studies
Peking University
Beijing, China*

LINDACHAO@PKU.EDU.CN

Haishan Ye*

*School of Management
Xi'an Jiaotong University
Xi'an, China*

YEHAISHAN@XJTU.EDU.CN

Zhihua Zhang

*School of Mathematical Sciences
Peking University
Beijing, China*

ZHZHANG@MATH.PKU.EDU.CN

Editor: Philipp Hennig

Abstract

Optimization is important in machine learning problems, and quasi-Newton methods have a reputation as the most efficient numerical methods for smooth unconstrained optimization. In this paper, we study the explicit superlinear convergence rates of quasi-Newton methods and address two open problems mentioned by Rodomanov and Nesterov (2021b). First, we extend Rodomanov and Nesterov (2021b)'s results to random quasi-Newton methods, which include common DFP, BFGS, SR1 methods. Such random methods employ a random direction for updating the approximate Hessian matrix in each iteration. Second, we focus on the specific quasi-Newton methods: SR1 and BFGS methods. We provide improved versions of greedy and random methods with provable better explicit (local) superlinear convergence rates. Our analysis is closely related to the approximation of a given Hessian matrix, unconstrained quadratic objective, as well as the general strongly convex, smooth, and strongly self-concordant functions.

Keywords: quasi-Newton methods, superlinear convergence, local convergence, rate of convergence, Broyden family, SR1, BFGS, DFP

1. Introduction

Many machine learning problems can be formulated to the minimization of an objective defined as the expectation over a set of random functions (Liu and Nocedal, 1989; Bottou and Le Cun, 2005; Shalev-Shwartz and Srebro, 2008; Mokhtari and Ribeiro, 2014, 2015). Specifically, given the training sample $\mathbf{z} \sim \mathcal{D}$, where \mathcal{D} is the data distribution, we consider

*. Corresponding Author.

an optimization function $f : \mathbb{R}^d \rightarrow \mathbb{R}$:

$$\min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x}) = \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \ell(\mathbf{x}, \mathbf{z}) + R(\mathbf{x}),$$

where $\ell(\mathbf{x}, \mathbf{z})$ is the loss with respect to the training sample \mathbf{z} , and $R(\mathbf{x})$ is some regularization function, such as $\|\mathbf{x}\|_2^2$. When \mathcal{D} is the empirical distribution of training samples $\{\mathbf{z}_i\}_{i=1}^n$, we could recover the classical finite-sum empirical risk minimization:

$$\min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{x}, \mathbf{z}_i) + R(\mathbf{x}).$$

Such a finite-sum formulation encapsulates a wide variety of machine learning problems including least squares regression, support vector machines (SVM), logistic regression, neural networks, and graphical models.

Previous methods mainly use first-order methods by evaluating objective function gradients $\nabla f(\mathbf{x})$, such as gradient descent, stochastic gradient descent, accelerated gradient descent (Nesterov, 2003), Adagrad (Duchi et al., 2011), Adam (Kingma and Ba, 2015), etc. These methods dominate the current optimization methods of machine learning problems, and have affordable computation complexity in each iteration. However, these first-order methods generally only have a linear or sublinear convergence rate even if the objective has nice properties.

Recently, second-order methods have also received great attention due to their fast convergence rates compared to first-order methods. But second-order methods, such as Newton’s method, are impractical because the exact Hessian matrix $\nabla^2 f(\mathbf{x})$ needs high computation cost in general cases. Common wisdom proposes quasi-Newton methods by replacing Hessian matrices with some reasonable approximations. The approximation is updated in iterations based on some special formulas from the previous variation.

Quasi-Newton methods have a broad application in machine learning problems (Bordes et al., 2009; Yu et al., 2010; Mokhtari and Ribeiro, 2015; Yuan and Li, 2020; Ye et al., 2020; Liu and Owen, 2021). There exist various quasi-Newton algorithms with different Hessian approximations. The three most popular versions are the *Davidon-Fletcher-Powell (DFP) method* (Fletcher and Powell, 1963; Davidon, 1991), the *Broyden-Fletcher-Goldfarb-Shanno (BFGS) method* (Broyden, 1970a,b; Fletcher, 1970; Goldfarb, 1970; Shanno, 1970), and the *Symmetric Rank 1 (SR1) method* (Broyden, 1967; Davidon, 1991), all of which belong to the Broyden family (Broyden, 1967) of quasi-Newton algorithms. The most attractive property of quasi-Newton methods compared to the classical first-order methods, is their superlinear convergence, which can trace back to the 1970s (Powell, 1971; Broyden et al., 1973; Dennis and Moré, 1974). However, the superlinear convergence rates provided in prior work are asymptotic (Stachurski, 1981; Griewank and Toint, 1982; Byrd et al., 1987; Yabe and Yamaki, 1996; Kovalev et al., 2020). The results only show that the ratio of successive residuals tends to zero as the running iterations approach to infinity, i.e.,

$$\lim_{k \rightarrow +\infty} \frac{\|\mathbf{x}_{k+1} - \mathbf{x}_*\|}{\|\mathbf{x}_k - \mathbf{x}_*\|} = 0, \text{ or } \|\mathbf{x}_{k+1} - \mathbf{x}_*\| = o(\|\mathbf{x}_k - \mathbf{x}_*\|),$$

where $\{\mathbf{x}_k\}$ is iterative update sequence, k is the iteration counter, and \mathbf{x}_* is the optimal solution. It is unknown whether the residuals converge like $O(c^{k^2})$, $O(k^{-k})$, where $c \in (0, 1)$

is some constant. Hence, the theory is inadequate and there still lacks of a specific superlinear convergence rate. Additionally, machine learning problems have requirement of the explicit convergence rates to compare the performance and design better algorithms for applications. Therefore, *to give a better guidance of quasi-Newton methods in machine learning problems, we are still interested in the explicit rates of quasi-Newton methods.*

Recently, Rodomanov and Nesterov (2021b) gave the first explicit local superlinear convergence for their proposed new quasi-Newton methods. They introduced greedy quasi-Newton updates by greedily selecting from basis vectors to maximize a certain measure of progress, and established an explicit non-asymptotic bound on the local superlinear convergence rate correspondingly. However, as Rodomanov and Nesterov (2021b) stated, “greedy methods require additional information beyond just the gradient of the objective function.” A natural idea might be to replace the greedy strategy with a randomized one. Indeed, the strategy of randomness has almost the same performance as the greedy one, which has been observed in Rodomanov and Nesterov (2021b)’s experiments. Therefore, one can expect that it should be possible to establish similar theoretical results about its superlinear convergence, but they did not provide theoretical guarantees. This raises the issue: *can we give explicit superlinear rates for random quasi-Newton methods theoretically?* In addition, Rodomanov and Nesterov (2021b)’s proofs are mainly applicable to the DFP methods because they reduced all possible Broyden family to the DFP update based on the monotonicity property (see Lemma 5). However, the SR1 and BFGS updates are more popular and faster than the DFP update in practice, which also has been verified in their experiments. Thus, it is natural to ask *can we obtain separate superlinear rates for different quasi-Newton methods?*

In this work, we solve the above two problems rigorously. We extend Rodomanov and Nesterov (2021b)’s results into random quasi-Newton methods, and improve the local superlinear convergence rates by our revised greedy or random SR1 and BFGS methods. We present our contribution in detail as follows:

- First, we extend Rodomanov and Nesterov (2021b)’s results to random quasi-Newton methods, which use a random direction for updating the approximate Hessian matrix. Our superlinear convergence rate is of the form $(1 - \frac{1}{d\kappa+1})^{k(k-1)/2}$ with high probability, which is similar as the greedy-type methods proposed by Rodomanov and Nesterov (2021b, Theorem 4.9). Here, κ is the condition number of the objective function, k is the current iteration, and d is the dimension of parameters.
- Second, for specific quasi-Newton methods, including SR1 and BFGS methods, we provide improved versions of greedy and random methods. We show that for approximating a fixed Hessian matrix, both the methods share a faster condition-number-free convergence. Particularly, we can obtain the superlinear convergence rate $O((1 - \frac{\kappa}{d})_+)$ for the SR1 update, and the linear convergence rate $O((1 - \frac{1}{d})^k)$ for the BFGS update, where $(x)_+ = \max\{x, 0\}$. Both the findings improve the original convergence rate $O((1 - \frac{1}{d\kappa})^k)$ by Rodomanov and Nesterov (2021b, Theorem 2.5).
- Third, we extend our analysis to a practical scheme, showing (local) superlinear convergence under our proposed greedy/random SR1 and BFGS update, when applied to unconstrained quadratic objective or strongly self-concordant functions. We list our

| Quasi-Newton Methods | Superlinear Rates | k_0 |
|---|--|---------------------------------------|
| Greedy Broyden (Rodomanov and Nesterov, 2021b) | $(1 - \frac{1}{d\kappa})^{k(k-1)/2} (\frac{1}{2})^k (1 - \frac{1}{2\kappa})^{k_0}$ | $O(d\kappa \ln(d\kappa))$ |
| Random Broyden (Corollary 12) | $(1 - \frac{1}{d\kappa+1})^{k(k-1)/2} (\frac{1}{2})^k (1 - \frac{1}{2\kappa})^{k_0}$ | $O(d\kappa \ln(d\kappa/\delta))$ |
| Greedy BFGS*/SR1 (Corollary 20) | $(1 - \frac{1}{d})^{k(k-1)/2} (\frac{1}{2})^k (1 - \frac{1}{2\kappa})^{k_0}$ | $O((d + \kappa) \ln(d\kappa))$ |
| Random BFGS/SR1 (Corollary 20) | $(1 - \frac{1}{d+1})^{k(k-1)/2} (\frac{1}{2})^k (1 - \frac{1}{2\kappa})^{k_0}$ | $O((d + \kappa) \ln(d\kappa/\delta))$ |

Table 1: Comparison of the existing specific superlinear convergence rates of the random or greedy quasi-Newton methods in the view of $\lambda_f(\cdot)$ (shown in Eq. (3)) under strongly self-concordant objective, where d is the dimension of parameters, κ is the condition number of the objective function, k_0 is the iteration number last for the first phase and k is the iteration number of the subsequent second phase. For the randomized methods, the presented rates hold with probability at least $1 - \delta$. (*): Our greedy BFGS method is not practical.

results in Table 1 with the same formulation as the work of Rodomanov and Nesterov (2021b). Note that in general, the convergence goes through two phases. The first phase lasts for k_0 iterations, and only has a linear convergence rate $O((1 - \frac{1}{2\kappa})^{k_0})$. The second phase has a superlinear convergence rate $O((1 - \frac{1}{d})^{k(k-1)/2})$. Our revised bound takes fewer first-phase iterations k_0 as well as a faster (condition-number-free) superlinear convergence rate in the second phase compared to Rodomanov and Nesterov (2021b)'s results.

1.1 Other Related Work

In addition to the work of Rodomanov and Nesterov (2021b), there are other results of explicit local superlinear convergence analysis along this line of research. Rodomanov and Nesterov (2021c) analyzed the well-known DFP and BFGS methods, which are based on a standard Hessian update direction through the previous variation. They demonstrated faster initial convergence rates, while slower final rates compared to Rodomanov and Nesterov (2021b)'s results. Rodomanov and Nesterov (2021a) improved Rodomanov and Nesterov (2021c)'s results by reducing the dependence of the condition number κ to $\ln \kappa$, though having similar worse long-history behavior. Jin and Mokhtari (2020) provided a non-asymptotic dimension-free superlinear convergence rate of the original Broyden family when the initial Hessian approximation is also good enough. However, the two issues mentioned earlier remain open.

The remainder of this paper is organized as follows. We present preliminaries in Section 2, and discuss the rates of random quasi-Newton methods in Section 3. In Section 4, we show faster superlinear convergence rates of our revised greedy/random SR1 and BFGS methods. Then in Section 5, we show comparison with the work of Rodomanov and Nesterov (2021b)

in detail. We give some empirical results in Section 6. Finally, we conclude our results in Section 7.

2. Preliminaries

First of all, we present some notation. We denote vectors by lowercase bold letters (e.g., \mathbf{u}, \mathbf{x}), and matrices by capital bold letters (e.g., $\mathbf{W} = [w_{ij}]$). We use $\mathbf{e}_1, \dots, \mathbf{e}_d$ for the d -dimensional standard coordinate directions, and $(x)_+ = \max\{x, 0\}$ for $x \in \mathbb{R}$. Let $\lambda_{\max}(\mathbf{A}) = \lambda_1(\mathbf{A}) \geq \dots \geq \lambda_d(\mathbf{A})$ be the eigenvalues of a real symmetric matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$, and $\|\cdot\|$ denotes the standard Euclidean norm (ℓ_2 -norm) for vectors, or induced ℓ_2 -norm (spectral norm) for a given matrix: $\|\mathbf{A}\| = \sup_{\|\mathbf{u}\|=1, \mathbf{u} \in \mathbb{R}^d} \|\mathbf{A}\mathbf{u}\|$. We denote $\mathcal{S}^{d-1} := \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\| = 1\}$ as the standard Euclidean sphere in \mathbb{R}^d , and $\text{Unif}(\mathcal{S}^{d-1})$ as the uniform distribution from \mathcal{S}^{d-1} . We use $\mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ as the standard Gaussian distribution, where $\mathbf{I}_d \in \mathbb{R}^{d \times d}$ is the identity matrix.

For two symmetric matrices \mathbf{A} and $\mathbf{B} \in \mathbb{R}^{d \times d}$, we denote $\mathbf{A} \succeq \mathbf{B}$ (or $\mathbf{B} \preceq \mathbf{A}$) if $\mathbf{A} - \mathbf{B}$ is a positive semi-definite matrix, and $\mathbf{A} \succ \mathbf{B}$ (or $\mathbf{B} \prec \mathbf{A}$) if $\mathbf{A} - \mathbf{B}$ is a positive definite matrix. Following Rodomanov and Nesterov (2021b)'s notation, for a given positive definite matrix \mathbf{A} (i.e., $\mathbf{A} \succ 0$), we induce a pair of conjugate Euclidean norms: $\|\mathbf{x}\|_{\mathbf{A}} := \sqrt{\mathbf{x}^\top \mathbf{A} \mathbf{x}}$ and $\|\mathbf{x}\|_{\mathbf{A}}^* := \sqrt{\mathbf{x}^\top \mathbf{A}^{-1} \mathbf{x}}$. When $\mathbf{A} = \nabla^2 f(\mathbf{x}) \succ 0$ for some $\mathbf{x} \in \mathbb{R}^d$, we prefer to use notation $\|\cdot\|_{\mathbf{x}}$ and $\|\cdot\|_{\mathbf{x}}^*$, provided that there is no ambiguity with the reference function f .

Next, we introduce some common definitions used in this paper below.

Definition 1 (Strongly convex and smooth) *A twice differentiable function $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is μ -strongly convex and L -smooth ($\mu, L > 0$), if*

$$\mu \mathbf{I}_d \preceq \nabla^2 f(\mathbf{x}) \preceq L \mathbf{I}_d, \forall \mathbf{x} \in \mathbb{R}^d.$$

Additionally, the condition number of a μ -strongly convex and L -smooth function is $\kappa := L/\mu$.

We also need the same assumption of *strongly self-concordancy* followed by Rodomanov and Nesterov (2021b). And Rodomanov and Nesterov (2021b, Section 4) have already mentioned several properties and examples of strongly self-concordant functions, such as a strongly convex function with Lipschitz continuous Hessians.

Definition 2 (Strongly self-concordant) *A twice differentiable function $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is M -strongly self-concordant ($M > 0$), if the Hessians are close to each other in the sense that*

$$\nabla^2 f(\mathbf{y}) - \nabla^2 f(\mathbf{x}) \preceq M \|\mathbf{y} - \mathbf{x}\|_{\mathbf{z}} \nabla^2 f(\mathbf{w}), \forall \mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{w} \in \mathbb{R}^d.$$

Finally, we recall the rate of convergence used in this paper.

Definition 3 (R-Linear/Superlinear convergence) *Suppose a scalar sequence $\{x_k\}$ converges to 0 with*

$$\lim_{k \rightarrow +\infty} \frac{|x_{k+1}|}{|x_k|} = q \in [0, 1).$$

Now suppose another sequence $\{y_k\}$ converges to y^ and satisfies that $|y_k - y^*| \leq |x_k|, \forall k \geq 0$. We say $\{y_k\}$ converges superlinearly if and only if $q = 0$, linearly if and only if $q \in (0, 1)$.*

2.1 Notation for Convergence Analysis

For convergence analysis, we introduce two measures which describe the approximation precision of the positive definite matrices:

$$\sigma_{\mathbf{A}}(\mathbf{G}) := \text{tr}[(\mathbf{G} - \mathbf{A})\mathbf{A}^{-1}] = \text{tr}(\mathbf{G}\mathbf{A}^{-1}) - d, \text{ where } \mathbf{G} \succeq \mathbf{A} \succ \mathbf{0}, \quad (1)$$

and

$$\tau_{\mathbf{A}}(\mathbf{G}) := \text{tr}(\mathbf{G} - \mathbf{A}), \text{ where } \mathbf{G} \succeq \mathbf{A} \succ \mathbf{0}. \quad (2)$$

Moreover, we estimate the convergence rate of a strongly convex objective $f(\mathbf{x})$ by the local norm of the gradient:

$$\lambda_f(\mathbf{x}) := \|\nabla f(\mathbf{x})\|_{\mathbf{x}}^* = \sqrt{\nabla f(\mathbf{x})^\top [\nabla^2 f(\mathbf{x})]^{-1} \nabla f(\mathbf{x})}, \quad \mathbf{x} \in \mathbb{R}^d. \quad (3)$$

Note that $\sigma_{\mathbf{A}}(\mathbf{G})$ and $\lambda_f(\mathbf{x})$ are also introduced in the work of Rodomanov and Nesterov (2021b). When applied to the update sequences $\{\mathbf{x}_k\}$ and $\{\mathbf{G}_k\}$ from a specific algorithm, we also denote the following notation for brevity:

$$\lambda_k := \lambda_f(\mathbf{x}_k), \quad \sigma_k := \sigma_{\nabla^2 f(\mathbf{x}_k)}(\mathbf{G}_k) \quad \text{and} \quad \tau_k := \tau_{\nabla^2 f(\mathbf{x}_k)}(\mathbf{G}_k). \quad (4)$$

2.2 Quasi-Newton Updates

Before starting our theoretical results, we briefly review a class of quasi-Newton updating rules for approximating a positive definite matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$. We follow the definition by Rodomanov and Nesterov (2021b), employing the following family of updates which describes the Broyden family (Nocedal and Wright, 2006, Section 6.3) of quasi-Newton updates, parameterized by a scalar $\tau \in \mathbb{R}$.

Definition 4 *Let $\mathbf{G} \succeq \mathbf{A} \succ \mathbf{0}$. For any $\mathbf{u} \in \mathbb{R}^d$, if $\mathbf{G}\mathbf{u} = \mathbf{A}\mathbf{u}$, we define $\text{Broyd}_\tau(\mathbf{G}, \mathbf{A}, \mathbf{u}) := \mathbf{G}$. Otherwise, i.e., $\mathbf{G}\mathbf{u} \neq \mathbf{A}\mathbf{u}$, we define*

$$\begin{aligned} \text{Broyd}_\tau(\mathbf{G}, \mathbf{A}, \mathbf{u}) := & \tau \left[\mathbf{G} - \frac{\mathbf{A}\mathbf{u}\mathbf{u}^\top \mathbf{G} + \mathbf{G}\mathbf{u}\mathbf{u}^\top \mathbf{A}}{\mathbf{u}^\top \mathbf{A}\mathbf{u}} + \left(\frac{\mathbf{u}^\top \mathbf{G}\mathbf{u}}{\mathbf{u}^\top \mathbf{A}\mathbf{u}} + 1 \right) \frac{\mathbf{A}\mathbf{u}\mathbf{u}^\top \mathbf{A}}{\mathbf{u}^\top \mathbf{A}\mathbf{u}} \right] \\ & + (1 - \tau) \left[\mathbf{G} - \frac{(\mathbf{G} - \mathbf{A})\mathbf{u}\mathbf{u}^\top (\mathbf{G} - \mathbf{A})}{\mathbf{u}^\top (\mathbf{G} - \mathbf{A})\mathbf{u}} \right]. \end{aligned} \quad (5)$$

As mentioned in the work of Rodomanov and Nesterov (2021b), we can recover several well-known quasi-Newton methods for several choices of τ .

For $\tau = 0$, Eq. (5) corresponds to the well-known SR1 update:

$$\text{SR1}(\mathbf{G}, \mathbf{A}, \mathbf{u}) := \mathbf{G} - \frac{(\mathbf{G} - \mathbf{A})\mathbf{u}\mathbf{u}^\top (\mathbf{G} - \mathbf{A})}{\mathbf{u}^\top (\mathbf{G} - \mathbf{A})\mathbf{u}}, \text{ if } \mathbf{G}\mathbf{u} \neq \mathbf{A}\mathbf{u}, \quad (6)$$

and for $\tau = 1$, it corresponds to the well-known DFP update:

$$\text{DFP}(\mathbf{G}, \mathbf{A}, \mathbf{u}) := \mathbf{G} - \frac{\mathbf{A}\mathbf{u}\mathbf{u}^\top \mathbf{G} + \mathbf{G}\mathbf{u}\mathbf{u}^\top \mathbf{A}}{\mathbf{u}^\top \mathbf{A}\mathbf{u}} + \left(\frac{\mathbf{u}^\top \mathbf{G}\mathbf{u}}{\mathbf{u}^\top \mathbf{A}\mathbf{u}} + 1 \right) \frac{\mathbf{A}\mathbf{u}\mathbf{u}^\top \mathbf{A}}{\mathbf{u}^\top \mathbf{A}\mathbf{u}}, \text{ if } \mathbf{u} \neq \mathbf{0}. \quad (7)$$

Finally, when $\tau = \frac{\mathbf{u}^\top \mathbf{A} \mathbf{u}}{\mathbf{u}^\top \mathbf{G} \mathbf{u}} \in [0, 1]$, we recover the famous BFGS update¹:

$$\text{BFGS}(\mathbf{G}, \mathbf{A}, \mathbf{u}) := \mathbf{G} - \frac{\mathbf{G} \mathbf{u} \mathbf{u}^\top \mathbf{G}}{\mathbf{u}^\top \mathbf{G} \mathbf{u}} + \frac{\mathbf{A} \mathbf{u} \mathbf{u}^\top \mathbf{A}}{\mathbf{u}^\top \mathbf{A} \mathbf{u}}, \text{ if } \mathbf{u} \neq \mathbf{0}. \quad (8)$$

The Broyden family has matrix monotonicity below, showing the relationship among these quasi-Newton methods.

Lemma 5 (Rodomanov and Nesterov, 2021b, Lemmas 2.1 and 2.2) *If $\mathbf{0} \prec \mathbf{A} \preceq \mathbf{G} \preceq \eta \mathbf{A}$ for some $\eta \geq 1$, then we have for any $\mathbf{u} \in \mathbb{R}^d$, and $\tau_1, \tau_2 \in \mathbb{R}$ with $\tau_1 \leq \tau_2$ such that*

$$\text{Broyd}_{\tau_1}(\mathbf{G}, \mathbf{A}, \mathbf{u}) \preceq \text{Broyd}_{\tau_2}(\mathbf{G}, \mathbf{A}, \mathbf{u}).$$

And for any $\tau \in [0, 1]$, we have $\mathbf{A} \preceq \text{Broyd}_\tau(\mathbf{G}, \mathbf{A}, \mathbf{u}) \preceq \eta \mathbf{A}$.

2.3 Greedy and Random Quasi-Newton Updates

Rodomanov and Nesterov (2021b) proposed a greedy version for selecting the direction \mathbf{u} :

$$\text{(Greedy Broyden)} \quad \hat{\mathbf{u}}_{\mathbf{A}}(\mathbf{G}) := \arg \max_{\mathbf{u} \in \{\mathbf{e}_1, \dots, \mathbf{e}_d\}} \frac{\mathbf{u}^\top \mathbf{G} \mathbf{u}}{\mathbf{u}^\top \mathbf{A} \mathbf{u}}, \quad (9)$$

which provides superlinear convergence of the form $(1 - \frac{1}{d\lambda})^{k(k-1)/2}$. They also conducted experiments to verify the performance of their greedy methods, which actually is competitive with the standard versions. Moreover, they gave random quasi-Newton updates, that is,

$$\text{(Random Broyden)} \quad \mathbf{u} \sim \mathcal{D}$$

for some predefined distribution \mathcal{D} . They observed that choosing a random direction uniformly from the standard Euclidean sphere, i.e., $\mathbf{u} \sim \text{Unif}(\mathcal{S}^{d-1})$, does not make superlinear convergence looser, and is only slightly slower than the greedy versions. However, they did not provide the theory to support their experimental findings. We describe the distribution \mathcal{D} explicitly, and give a rigorous proof of the superlinear rates of such random methods in this paper.

3. Rates of Random Quasi-Newton Methods

We follow the same roadmap as the work of Rodomanov and Nesterov (2021b). We begin with the analysis of quasi-Newton methods for approximating a target matrix. Then we extend the scheme to unconstrained quadratic minimization. Finally, we move to general strongly self-concordant functions.

3.1 Matrix Approximation

We first consider approximating a positive definite matrix \mathbf{A} which satisfies

$$\mu \mathbf{I}_d \preceq \mathbf{A} \preceq L \mathbf{I}_d, \quad (10)$$

1. See Eq. (2.6) in the work of Rodomanov and Nesterov (2021b) for derivation.

Algorithm 1 Random quasi-Newton updates

 Initialization: Choose $\mathbf{G}_0 \succeq \mathbf{A}$.

for $k \geq 0$ **do**

 Choose $\tau_k \in [0, 1]$ and \mathbf{u}_k from distribution \mathcal{D} which satisfies Eq. (12).

 Compute $\mathbf{G}_{k+1} = \text{Broyd}_{\tau_k}(\mathbf{G}_k, \mathbf{A}, \mathbf{u}_k)$.

end for

where $L \geq \mu > 0$, and $\varkappa := L/\mu$ is the condition number of \mathbf{A} . We use the measure $\sigma_{\mathbf{A}}(\mathbf{G})$ to describe the closeness between matrix \mathbf{A} and the current approximate matrix \mathbf{G} . When $\mathbf{G}\mathbf{u} \neq \mathbf{A}\mathbf{u}$, one iteration update of Broyden family leads to

$$\begin{aligned} \sigma_{\mathbf{A}}(\mathbf{G}_+) &\stackrel{(1)(5)}{=} \tau \left[\sigma_{\mathbf{A}}(\mathbf{G}) - 2 \cdot \frac{\mathbf{u}^\top \mathbf{G}\mathbf{u}}{\mathbf{u}^\top \mathbf{A}\mathbf{u}} + \left(\frac{\mathbf{u}^\top \mathbf{G}\mathbf{u}}{\mathbf{u}^\top \mathbf{A}\mathbf{u}} + 1 \right) \right] \\ &\quad + (1 - \tau) \left[\sigma_{\mathbf{A}}(\mathbf{G}) - \frac{\mathbf{u}^\top (\mathbf{G} - \mathbf{A})\mathbf{A}^{-1}(\mathbf{G} - \mathbf{A})\mathbf{u}}{\mathbf{u}^\top (\mathbf{G} - \mathbf{A})\mathbf{u}} \right] \\ &= \sigma_{\mathbf{A}}(\mathbf{G}) - \left[\tau \cdot \frac{\mathbf{u}^\top (\mathbf{G} - \mathbf{A})\mathbf{u}}{\mathbf{u}^\top \mathbf{A}\mathbf{u}} + (1 - \tau) \cdot \frac{\mathbf{u}^\top (\mathbf{G} - \mathbf{A})\mathbf{A}^{-1}(\mathbf{G} - \mathbf{A})\mathbf{u}}{\mathbf{u}^\top (\mathbf{G} - \mathbf{A})\mathbf{u}} \right], \end{aligned}$$

where $\mathbf{G}_+ = \text{Broyd}_\tau(\mathbf{G}, \mathbf{A}, \mathbf{u})$. Note that we always have $\mathbf{G}_+ \succeq \mathbf{A}$ for $\tau \in [0, 1]$ if $\mathbf{G} \succeq \mathbf{A}$ from Lemma 5. Thus, by the Cauchy-Schwarz inequality and $\mathbf{G} \succeq \mathbf{A}$, we have

$$\frac{\mathbf{u}^\top (\mathbf{G} - \mathbf{A})\mathbf{A}^{-1}(\mathbf{G} - \mathbf{A})\mathbf{u}}{\mathbf{u}^\top (\mathbf{G} - \mathbf{A})\mathbf{u}} \geq \frac{\mathbf{u}^\top (\mathbf{G} - \mathbf{A})\mathbf{u}}{\mathbf{u}^\top \mathbf{A}\mathbf{u}}.$$

Hence, we obtain when $\mathbf{G}\mathbf{u} \neq \mathbf{A}\mathbf{u}$,

$$\begin{aligned} \sigma_{\mathbf{A}}(\mathbf{G}_+) &\leq \sigma_{\mathbf{A}}(\mathbf{G}) - \frac{\mathbf{u}^\top (\mathbf{G} - \mathbf{A})\mathbf{u}}{\mathbf{u}^\top \mathbf{A}\mathbf{u}} \stackrel{(10)}{\leq} \sigma_{\mathbf{A}}(\mathbf{G}) - \frac{1}{L} \cdot \frac{\mathbf{u}^\top (\mathbf{G} - \mathbf{A})\mathbf{u}}{\mathbf{u}^\top \mathbf{u}} \\ &= \sigma_{\mathbf{A}}(\mathbf{G}) - \frac{1}{L} \text{tr} \left[(\mathbf{G} - \mathbf{A}) \cdot \frac{\mathbf{u}\mathbf{u}^\top}{\mathbf{u}^\top \mathbf{u}} \right]. \end{aligned} \quad (11)$$

Moreover, Eq. (11) trivially holds when $\mathbf{G}\mathbf{u} = \mathbf{A}\mathbf{u}$. Therefore, for a random direction \mathbf{u} , we only need $\mathbb{E}\mathbf{u}\mathbf{u}^\top/\mathbf{u}^\top \mathbf{u}$ to preserve some benign property, which leads to our assumption of the random update distribution.

$$\text{(Random Broyden)} \quad \mathbf{u} \sim \mathcal{D}, \text{ s.t. } \mathbb{E}_{\mathbf{u} \sim \mathcal{D}} \frac{\mathbf{u}\mathbf{u}^\top}{\mathbf{u}^\top \mathbf{u}} = \frac{1}{d} \mathbf{I}_d. \quad (12)$$

It is easy to verify that common distributions such as $\mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ and $\text{Unif}(\mathcal{S}^{d-1})$ satisfy our requirements. Based on Eq. (12) and update in Algorithm 1, we could show linear convergence of \mathbf{G}_k to \mathbf{A} under measure $\sigma_{\mathbf{A}}(\cdot)$. The proof of Theorem 6 is shown in Appendix B.1.

Theorem 6 *Under the update in Algorithm 1 with a randomly initialized \mathbf{G}_0 , such that $\mathbf{G}_0 \succeq \mathbf{A}$ always holds, we have that*

$$\forall k \geq 0, \mathbf{G}_k \succeq \mathbf{A} \text{ and } 0 \leq \mathbb{E}\sigma_k \leq \left(1 - \frac{1}{d\varkappa}\right)^k \mathbb{E}\sigma_0, \quad (13)$$

Therefore, $\mathbb{E}\sigma_{\mathbf{A}}(\mathbf{G}_k)$ converges to zero linearly.

Algorithm 2 Random quasi-Newton methods for quadratic minimization

- 1: Initialization: Choose $\mathbf{x}_0 \in \mathbb{R}^d$ and $\mathbf{G}_0 \succeq \mathbf{A}$.
 - 2: **for** $k \geq 0$ **do**
 - 3: Update $\mathbf{x}_{k+1} = \mathbf{x}_k - \mathbf{G}_k^{-1} \nabla f(\mathbf{x}_k)$.
 - 4: Choose $\tau_k \in [0, 1]$ and \mathbf{u}_k from distribution \mathcal{D} which satisfies Eq. (12).
 - 5: Compute $\mathbf{G}_{k+1} = \text{Broyd}_{\tau_k}(\mathbf{G}_k, \mathbf{A}, \mathbf{u}_k)$.
 - 6: **end for**
-

3.2 Unconstrained Quadratic Minimization

Based on the efficiency of random quasi-Newton updates in matrix approximation, we next turn to minimize the strongly convex quadratic function (with a fixed Hessian):

$$f(\mathbf{x}) = \frac{1}{2} \mathbf{x}^\top \mathbf{A} \mathbf{x} - \mathbf{b}^\top \mathbf{x}, \text{ where } \mu \mathbf{I}_d \preceq \mathbf{A} \preceq L \mathbf{I}_d \text{ with } L, \mu > 0. \quad (14)$$

The algorithm is shown in Algorithm 2. As classical quasi-Newton methods do, we need to use the quasi-Newton step for updating the parameters as well as approximating the true Hessian matrix \mathbf{A} . Moreover, Algorithm 2 is only for theoretical analysis, while we need to adopt the inverse update rules for \mathbf{G}_k^{-1} directly in practice.

We adopt $\{\lambda_k\}$ (defined in Eqs. (3) and (4)) to estimate the convergence rate of the objective in Eq. (14). Note that this measure of optimality is directly related to the functional residual. Indeed, note that $\mathbf{x}_* = \mathbf{A}^{-1} \mathbf{b}$ is the minimizer of Eq. (14). Then we obtain

$$f(\mathbf{x}) - f(\mathbf{x}_*) = \frac{1}{2} (\mathbf{x} - \mathbf{x}_*)^\top \mathbf{A} (\mathbf{x} - \mathbf{x}_*) = \frac{1}{2} (\mathbf{A} \mathbf{x} - \mathbf{b})^\top \mathbf{A}^{-1} (\mathbf{A} \mathbf{x} - \mathbf{b}) \stackrel{(3)}{=} \frac{1}{2} \lambda_f(\mathbf{x})^2.$$

The following lemma shows how $\lambda_f(\cdot)$ changes after one iteration of process in Algorithm 2.

Lemma 7 (Rodomanov and Nesterov, 2021b, Lemma 3.2) *Let $k \geq 0$, and $\eta_k \geq 1$ be such that $\mathbf{A} \preceq \mathbf{G}_k \preceq \eta_k \mathbf{A}$. Then we have $\lambda_{k+1} \leq \left(1 - \frac{1}{\eta_k}\right) \lambda_k \leq (\eta_k - 1) \lambda_k$.*

Thus, to estimate how fast $\{\lambda_k\}$ converges to zero, we need the upper bound η_k , which was already done in Theorem 6. Therefore, we can guarantee a superlinear convergence of $\{\lambda_k\}$ (under expectation) using the random quasi-Newton update. The proof of Theorem 8 can be found in Appendix C.1.

Theorem 8 *Under the update in Algorithm 2 with a randomly initialized \mathbf{G}_0 , such that $\mathbf{G}_0 \succeq \mathbf{A}$ always holds, we have that $\forall k \geq 0, \lambda_{k+1} \leq \rho_k \lambda_k$, where ρ_k is a certain nonnegative random variable such that*

$$\mathbb{E} \rho_k \leq \left(1 - \frac{1}{d\mathcal{X}}\right)^k \mathbb{E} \sigma_0, \quad \forall k \geq 0.$$

For better understanding the convergent behavior without expectation, we show the probabilistic version of Theorems 6 and 8 below, and leave the proof in Appendix C.2.

Corollary 9 *Under the same assumptions as Theorem 8, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over the random directions $\{\mathbf{u}_k\}$, we have for all $k \geq 0$,*

$$\sigma_k \leq \frac{2d^2 \mathcal{X}^2 \mathbb{E} \sigma_0}{\delta} \left(1 - \frac{1}{d\mathcal{X} + 1}\right)^k \quad \text{and} \quad \lambda_k \leq \left(\frac{2d^2 \mathcal{X}^2 \mathbb{E} \sigma_0}{\delta}\right)^k \left(1 - \frac{1}{d\mathcal{X} + 1}\right)^{k(k-1)/2} \lambda_0.$$

Algorithm 3 Random quasi-Newton methods for general strongly self-concordant objective

- 1: Initialization: Choose $\mathbf{x}_0 \in \mathbb{R}^d$ and $\mathbf{G}_0 \succeq \nabla^2 f(\mathbf{x}_0)$.
 - 2: **for** $k \geq 0$ **do**
 - 3: Update $\mathbf{x}_{k+1} = \mathbf{x}_k - \mathbf{G}_k^{-1} \nabla f(\mathbf{x}_k)$.
 - 4: Compute $r_k = \|\mathbf{x}_{k+1} - \mathbf{x}_k\|_{\mathbf{x}_k}$ and set $\tilde{\mathbf{G}}_k = (1 + Mr_k) \mathbf{G}_k$.
 - 5: Choose $\tau_k \in [0, 1]$ and \mathbf{u}_k from distribution \mathcal{D} which satisfies Eq. (12).
 - 6: Compute $\mathbf{G}_{k+1} = \text{Broyd}_{\tau_k}(\tilde{\mathbf{G}}_k, \nabla^2 f(\mathbf{x}_{k+1}), \mathbf{u}_k)$.
 - 7: **end for**
-

3.3 Minimization of General Functions

Next, we consider the optimization of a general machine learning objective: $\min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x})$, where $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is an M -strongly self-concordant, μ -strongly convex and L -smooth function with condition number $\varkappa = L/\mu$. Our goal is to extend the results in the previous sections, given that the methods can start from a sufficiently good initial point \mathbf{x}_0 . Unlike quadratic minimization, the true Hessian in each step varies. In order to ensure that $\mathbf{G}_{k+1} \succeq \nabla^2 f(\mathbf{x}_{k+1})$ holds for all $k \geq 0$, we adjust \mathbf{G}_k before doing quasi-Newton update. Instructed from the work of Rodomanov and Nesterov (2021b), we also use the *correction strategy*, which enlarges the approximation \mathbf{G}_k properly shown in Line 4 of Algorithm 3. Note that Algorithm 3 is only for theoretical analysis. We will use the inverse update rules for \mathbf{G}_k^{-1} and Hessian-vector products for $\text{Broyd}_{\tau_k}(\tilde{\mathbf{G}}_k, \nabla^2 f(\mathbf{x}_{k+1}), \mathbf{u}_k)$ in practice. For simplicity, we assume that the constants M and L are available, and $d \geq 2$. We first give convergent results in expectation in Lemma 10, and leave the proof in Appendix D.1.

Lemma 10 *Suppose in Algorithm 3, a random initialization \mathbf{G}_0 always satisfies $\nabla^2 f(\mathbf{x}_0) \preceq \mathbf{G}_0 \preceq \eta \nabla^2 f(\mathbf{x}_0)$ for some $\eta \geq 1$, and the initial point \mathbf{x}_0 is sufficiently close to the solution:*

$$M\lambda_0 \leq \frac{\ln 2}{4\eta(2d+1)}.$$

Then for all $k \geq 0$, we have $\nabla^2 f(\mathbf{x}_k) \preceq \mathbf{G}_k \preceq (1 + \delta_k) \nabla^2 f(\mathbf{x}_k)$, where δ_k is a certain nonnegative random variable such that

$$\mathbb{E}\delta_k \leq 2d\eta \left(1 - \frac{1}{d\varkappa}\right)^k,$$

and $\lambda_{k+1} \leq \rho_k \lambda_k$, where ρ_k is a certain nonnegative random variable such that

$$\mathbb{E}\rho_k \leq 2d\eta \left(1 - \frac{1}{d\varkappa}\right)^k.$$

We also show the probabilistic version of Lemma 10, which gives superlinear convergence of $\{\lambda_k\}$ and linear convergence of $\{\delta_k\}$ directly, and we leave the proof in Appendix D.2.

Theorem 11 *Under the same assumptions and notation as in Lemma 10, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over the random directions $\{\mathbf{u}_k\}$, we have for all $k \geq 0$,*

$$\delta_k \leq \frac{4d^3 \varkappa^2 \eta}{\delta} \left(1 - \frac{1}{d\varkappa + 1}\right)^k \quad \text{and} \quad \lambda_k \leq \left(\frac{4d^3 \varkappa^2 \eta}{\delta}\right)^k \left(1 - \frac{1}{d\varkappa + 1}\right)^{k(k-1)/2} \lambda_0.$$

Additionally, as mentioned by Rodomanov and Nesterov (2021b), if we adopt a weaker initialization of \mathbf{x}_0 , then the superlinear rate is valid only after certain iterations, i.e., the total iteration count $k' \geq k_0$ for some k_0 , while only linear convergence is guaranteed for $k' < k_0$. We combine both phases into the following corollary, and leave the proof in Appendix D.3.

Corollary 12 *Suppose in Algorithm 3, $\mathbf{G}_0 = L\mathbf{I}_d$ and \mathbf{x}_0 satisfies $M\lambda_0 \leq \frac{\ln \frac{3}{2}}{4\kappa}$, that is, the initial condition on λ_0 here is weaker than that in Lemma 10 with $\eta = \kappa$. Then we could obtain with probability at least $1 - \delta$ over the random directions $\{\mathbf{u}_k\}$,*

$$\lambda_{k_0+k} \leq \left(1 - \frac{1}{d\kappa + 1}\right)^{k(k-1)/2} \cdot \left(\frac{1}{2}\right)^k \cdot \left(1 - \frac{1}{2\kappa}\right)^{k_0} \cdot \lambda_0, \forall k \geq 0,$$

where $k_0 = O(d\kappa \ln(d\kappa/\delta))$.

Hence, we could see random quasi-Newton methods still have explicit superlinear convergence rates. The rate in Corollary 12 is slightly worse than the bounds in greedy methods (see Table 1) due to the probabilistic version, but is comparable overall.

4. Faster Rates for the BFGS and SR1 Methods

From Lemma 5, if $\mathbf{A} \preceq \mathbf{G} \preceq \eta\mathbf{A}$ for some $\eta \geq 1$, it follows that

$$\mathbf{A} \preceq \text{SR1}(\mathbf{G}, \mathbf{A}, \mathbf{u}) \preceq \text{BFGS}(\mathbf{G}, \mathbf{A}, \mathbf{u}) \preceq \text{DFP}(\mathbf{G}, \mathbf{A}, \mathbf{u}) \preceq \eta\mathbf{A}.$$

Intuitively, the approximation produced by SR1 is better than that produced by BFGS. And both of them are better than that produced by DFP. However, Rodomanov and Nesterov (2021b) reduced the analysis by casting all updates described by Broyden family ($\tau \in [0, 1]$) into the slowest DFP update ($\tau = 1$). Moreover, SR1 and BFGS methods also have faster numerical performance in practice. Therefore, Rodomanov and Nesterov (2021b) conjectured that SR1 and BFGS methods might have faster superlinear convergence rates. In this section, we will provide an affirmative answer to this conjecture.

4.1 Superlinear Convergence for SR1 Update

We first describe the SR1 update for approximating a fixed positive definite matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$. Let us now justify the efficiency of update Eq. (6) in ensuring convergence \mathbf{G} to \mathbf{A} . We adopt another measure $\tau_{\mathbf{A}}(\cdot)$ instead of $\sigma_{\mathbf{A}}(\cdot)$. According to $\tau_{\mathbf{A}}(\mathbf{G})$, one iteration update leads to

$$\tau_{\mathbf{A}}(\mathbf{G}_+) \stackrel{(2)(6)}{=} \tau_{\mathbf{A}}(\mathbf{G}) - \frac{\mathbf{u}^\top (\mathbf{G} - \mathbf{A})^2 \mathbf{u}}{\mathbf{u}^\top (\mathbf{G} - \mathbf{A}) \mathbf{u}}, \quad \mathbf{G}_+ = \text{SR1}(\mathbf{G}, \mathbf{A}, \mathbf{u}), \text{ if } \mathbf{G}\mathbf{u} \neq \mathbf{A}\mathbf{u}. \quad (15)$$

Now we revise greedy and random methods based on the progress of measure $\tau_{\mathbf{A}}(\cdot)$.

First, we introduce greedy method proposed in the work of Rodomanov and Nesterov (2021b), that greedily selects \mathbf{u} from the basis vectors to obtain the largest decrease of $\tau_{\mathbf{A}}(\mathbf{G}_+) - \tau_{\mathbf{A}}(\mathbf{G})$:

$$\bar{\mathbf{u}}_{\mathbf{A}}^{\text{raw}}(\mathbf{G}) := \arg \max_{\mathbf{u} \in \{\mathbf{e}_1, \dots, \mathbf{e}_d\}} \frac{\mathbf{u}^\top (\mathbf{G} - \mathbf{A})^2 \mathbf{u}}{\mathbf{u}^\top (\mathbf{G} - \mathbf{A}) \mathbf{u}}.$$

| Algorithm 4 Greedy/Random SR1 update | Algorithm 5 Greedy/Random BFGS update |
|---|---|
| 1: Initialization: Choose $\mathbf{G}_0 \succeq \mathbf{A}$. 2: for $k = 0, \dots, d - 1$ do 3: Choose \mathbf{u}_k from 1) <i>greedy method</i> : $\mathbf{u}_k = \bar{\mathbf{u}}_{\mathbf{A}}(\mathbf{G}_k)$, or 2) <i>random method</i> : $\mathbf{u}_k \sim \text{Unif}(\mathcal{S}^{d-1})$. 4: Compute $\mathbf{G}_{k+1} = \text{SR1}(\mathbf{G}_k, \mathbf{A}, \mathbf{u}_k)$. 5: end for | 1: Initialization: Set $\mathbf{G}_0 \succeq \mathbf{A}$, $\mathbf{L}_0^\top \mathbf{L}_0 = \mathbf{G}_0^{-1}$. 2: for $k \geq 0$ do 3: Compute $\mathbf{u}_k = \mathbf{L}_k^\top \tilde{\mathbf{u}}_k$ with $\tilde{\mathbf{u}}_k$ from 1) <i>greedy method</i> : $\tilde{\mathbf{u}}_k = \tilde{\mathbf{u}}_{\mathbf{A}}(\mathbf{L}_k)$, or 2) <i>random method</i> : $\tilde{\mathbf{u}}_k \sim \text{Unif}(\mathcal{S}^{d-1})$. 4: Compute $\mathbf{G}_{k+1} = \text{BFGS}(\mathbf{G}_k, \mathbf{A}, \mathbf{u}_k)$. 5: Compute \mathbf{L}_{k+1} based on Eq. (25). 6: end for |

However, we may encounter numerical overflow due to division by zero if $\mathbf{u}^\top (\mathbf{G} - \mathbf{A}) \mathbf{u}$ is nearly 0. Noting that $\mathbf{G} \succeq \mathbf{A}$, then from the Cauchy–Schwarz inequality, we have

$$\frac{\mathbf{u}^\top (\mathbf{G} - \mathbf{A})^2 \mathbf{u}}{\mathbf{u}^\top (\mathbf{G} - \mathbf{A}) \mathbf{u}} \geq \frac{\mathbf{u}^\top (\mathbf{G} - \mathbf{A}) \mathbf{u}}{\mathbf{u}^\top \mathbf{u}}. \quad (16)$$

Thus we employ a safer adjustment below:

$$(\text{Greedy SR1}) \quad \bar{\mathbf{u}}_{\mathbf{A}}(\mathbf{G}) := \arg \max_{\mathbf{u} \in \{\mathbf{e}_1, \dots, \mathbf{e}_d\}} \frac{\mathbf{u}^\top (\mathbf{G} - \mathbf{A}) \mathbf{u}}{\mathbf{u}^\top \mathbf{u}} = \arg \max_{\mathbf{u} \in \{\mathbf{e}_1, \dots, \mathbf{e}_d\}} \mathbf{u}^\top (\mathbf{G} - \mathbf{A}) \mathbf{u}. \quad (17)$$

Moreover, we only need to obtain the diagonal elements of \mathbf{A} (the current Hessian in practice), thus generally the total complexity is $O(d^2)$ in each iteration², which is acceptable and the same as the classical quasi-Newton methods.

Second, from the proof of the greedy method, we find that the random method by choosing \mathbf{u} from a spherically symmetric distribution, e.g.,

$$(\text{Random SR1}) \quad \mathbf{u} \sim \mathcal{N}(0, \mathbf{I}_d) \quad \text{or} \quad \mathbf{u} \sim \text{Unif}(\mathcal{S}^{d-1}), \quad (18)$$

also has similar performance and the same running complexity $O(d^2)$ in each iteration.

Next, we will show the convergence result below by estimating the decrease in the measure $\tau_{\mathbf{A}}(\cdot)$. In the following, the expectation considers all the randomness of the directions $\{\mathbf{u}_k\}$ during iterations, and when applied to the greedy method, we can view it with no randomness for the same notation. We leave the proof of Theorem 13 in Appendix B.2.

Theorem 13 *Suppose in Algorithm 4, a random initialization \mathbf{G}_0 always satisfies $\mathbf{G}_0 \succeq \mathbf{A}$. Then we obtain that for the greedy method defined in Eq. (17) or the random method defined in Eq. (18),*

$$\forall k \geq 0, \mathbf{G}_k \succeq \mathbf{A} \quad \text{and} \quad 0 \leq \mathbb{E} \tau_k \leq \left(1 - \frac{k}{d}\right)_+ \mathbb{E} \tau_0, \quad (19)$$

where $(x)_+ = \max\{0, x\}$. Hence, $\mathbb{E} \tau_{\mathbf{A}}(\mathbf{G}_k)$ converges to zero superlinearly. Particularly, $\forall k \geq d$, $\mathbf{G}_k = \mathbf{A}$ for greedy SR1 update, and $\mathbf{G}_k = \mathbf{A}$ almost surely for random SR1 update.

Previous work (Rodomanov and Nesterov, 2021b) adopted measure $\sigma_{\mathbf{A}}(\cdot)$, which only gives the same rate as general updates of the Broyden family (e.g., Theorem 6, and Rodomanov and Nesterov, 2021b, Theorem 2.5). Thus we employ a more precise measure $\tau_{\mathbf{A}}(\cdot)$.

2. Note that we can use the Hessian-vector product to obtain $\mathbf{A} \mathbf{u}$ (or $\nabla^2 f(\mathbf{x}) \cdot \mathbf{u}$) in practice. For most specific optimization problems, e.g., two problems in our experiments, one operation of the exact Hessian-vector product is tractable with $O(d)$ complexity.

4.2 Linear Convergence for BFGS Update

We now consider the classical BFGS update in the same scheme. Reusing the measure $\sigma_{\mathbf{A}}(\cdot)$, we obtain that

$$\sigma_{\mathbf{A}}(\mathbf{G}_+) \stackrel{(1)(8)}{=} \sigma_{\mathbf{A}}(\mathbf{G}) - \frac{\mathbf{u}^\top \mathbf{G} \mathbf{A}^{-1} \mathbf{G} \mathbf{u}}{\mathbf{u}^\top \mathbf{G} \mathbf{u}} + 1, \quad \mathbf{G}_+ = \text{BFGS}(\mathbf{G}, \mathbf{A}, \mathbf{u}), \text{ if } \mathbf{u} \neq \mathbf{0}. \quad (20)$$

If we directly apply the greedy or random method from the previous content, we could only obtain the same linear convergence rate as Rodomanov and Nesterov (2021b, Theorem 2.5). However, if we take advantage of the current \mathbf{G} , and choose a scaled direction such that $\mathbf{u} = \mathbf{L}^\top \tilde{\mathbf{u}}$ where \mathbf{L} is a square matrix satisfying $\mathbf{L}^\top \mathbf{L} = \mathbf{G}^{-1}$, then we could simplify the formulation and obtain a faster condition-number-free linear convergence rate. Specifically, after replacing \mathbf{u} with $\mathbf{L}^\top \tilde{\mathbf{u}}$ and \mathbf{G} with $\mathbf{L}^{-1} \mathbf{L}^{-\top}$, we get

$$\sigma_{\mathbf{A}}(\mathbf{G}_+) \stackrel{(20)}{=} \sigma_{\mathbf{A}}(\mathbf{G}) - \frac{\tilde{\mathbf{u}}^\top \mathbf{L}^{-\top} \mathbf{A}^{-1} \mathbf{L}^{-1} \tilde{\mathbf{u}}}{\tilde{\mathbf{u}}^\top \tilde{\mathbf{u}}} + 1, \quad \mathbf{G}_+ = \text{BFGS}(\mathbf{G}, \mathbf{A}, \mathbf{u}), \text{ if } \mathbf{u} \neq \mathbf{0}. \quad (21)$$

Thus our modified greedy BFGS update is as follows:

$$(\text{Greedy BFGS}) \quad \tilde{\mathbf{u}}_{\mathbf{A}}(\mathbf{L}) = \arg \max_{\tilde{\mathbf{u}} \in \{\mathbf{e}_1, \dots, \mathbf{e}_d\}} \tilde{\mathbf{u}}^\top \mathbf{L}^{-\top} \mathbf{A}^{-1} \mathbf{L}^{-1} \tilde{\mathbf{u}}. \quad (22)$$

Similar arguments apply to the random method used in Eq. (12):

$$(\text{Random BFGS}) \quad \tilde{\mathbf{u}} \sim \mathcal{D}, \text{ s.t. } \mathbb{E}_{\tilde{\mathbf{u}} \sim \mathcal{D}} \frac{\tilde{\mathbf{u}} \tilde{\mathbf{u}}^\top}{\tilde{\mathbf{u}}^\top \tilde{\mathbf{u}}} = \frac{1}{d} \mathbf{I}_d. \quad (23)$$

Now we give the linear convergence rate of the BFGS update under our modified method. We leave the proof of Theorem 14 in Appendix B.3.

Theorem 14 *Suppose in Algorithm 5, a random initialization \mathbf{G}_0 always satisfies $\mathbf{G}_0 \succeq \mathbf{A}$. Then we obtain that for the greedy method defined in Eq. (22) or the random method defined in Eq. (23),*

$$\forall k \geq 0, \mathbf{G}_k \succeq \mathbf{A} \text{ and } 0 \leq \mathbb{E} \sigma_k \leq \left(1 - \frac{1}{d}\right)^k \mathbb{E} \sigma_0. \quad (24)$$

Therefore, $\mathbb{E} \sigma_{\mathbf{A}}(\mathbf{G}_k)$ converges to zero linearly.

Remark 15 *Note that the complexity in Eq. (22) is $O(d^3)$ because we have multiplication-addition operations with (unknown) \mathbf{A}^{-1} . Hence we **do not** apply this greedy strategy in practice, but view it as a theoretical result similar to the random strategy. Moreover, the random method is still practical, and we will show the efficiency of our scaled direction compared to the original direction in our numerical experiments.*

Finally, we can employ an efficient way (with complexity $O(d^2)$) for updating \mathbf{L}_k at each step $k \geq 0$, and we leave the proof in Appendix E.

Proposition 16 *Suppose we already have $\mathbf{L}_k^\top \mathbf{L}_k = \mathbf{G}_k^{-1} \succ \mathbf{0}$, where \mathbf{L}_k is a square matrix, and $\mathbf{u}_k = \mathbf{L}_k^\top \tilde{\mathbf{u}}_k$. Then we can construct the square matrix \mathbf{L}_{k+1} which satisfies $\mathbf{L}_{k+1}^\top \mathbf{L}_{k+1} = \mathbf{G}_{k+1}^{-1} := [\text{BFGS}(\mathbf{G}_k, \mathbf{A}, \mathbf{u}_k)]^{-1}$ as below:*

$$\mathbf{L}_{k+1} = \mathbf{L}_k - \frac{[\mathbf{L}_k(\mathbf{A} \mathbf{u}_k) - \mathbf{v}_k] \mathbf{u}_k^\top}{\mathbf{u}_k^\top (\mathbf{A} \mathbf{u}_k)} \text{ with } \mathbf{v}_k = \sqrt{\mathbf{u}_k^\top \cdot (\mathbf{A} \mathbf{u}_k)} \cdot \frac{\tilde{\mathbf{u}}_k}{\|\tilde{\mathbf{u}}_k\|}. \quad (25)$$

Algorithm 6 Greedy/Random SR1/BFGS methods for quadratic minimization

- 1: Initialization: Choose $\mathbf{x}_0 \in \mathbb{R}^d$ and $\mathbf{G}_0 \succeq \mathbf{A}$, $\mathbf{L}_0^\top \mathbf{L}_0 = \mathbf{G}_0^{-1}$.
 - 2: **for** $k \geq 0$ **do**
 - 3: Update $\mathbf{x}_{k+1} = \mathbf{x}_k - \mathbf{G}_k^{-1} \nabla f(\mathbf{x}_k)$. Choose one of the following update rules:
 - 4: (i) SR1: Choose \mathbf{u}_k following Algorithm 4. Compute $\mathbf{G}_{k+1} = \text{SR1}(\mathbf{G}_k, \mathbf{A}, \mathbf{u}_k)$.
 - 5: (ii) BFGS: Choose \mathbf{u}_k following Algorithm 5. Compute $\mathbf{G}_{k+1} = \text{BFGS}(\mathbf{G}_k, \mathbf{A}, \mathbf{u}_k)$.
 - 6: **end for**
-

4.3 Unconstrained Quadratic Minimization

Based on the efficiency of the greedy/random SR1 and BFGS updates in matrix approximation, we next turn to minimize the strongly convex quadratic function in Eq. (14). We show the detail in Algorithm 6, which is only for theoretical analysis. In practice, we use the inverse update rules (Nocedal and Wright, 2006, Eqs. (6.17) and (6.25)) to update \mathbf{G}_k^{-1} :

$$\mathbf{G}_+^{-1} = \mathbf{G}^{-1} + \frac{(\mathbf{I}_d - \mathbf{G}^{-1}\mathbf{A})\mathbf{u}\mathbf{u}^\top(\mathbf{I}_d - \mathbf{A}\mathbf{G}^{-1})}{\mathbf{u}^\top(\mathbf{A} - \mathbf{A}\mathbf{G}^{-1}\mathbf{A})\mathbf{u}}, \quad \mathbf{G}_+ = \text{SR1}(\mathbf{G}, \mathbf{A}, \mathbf{u}); \quad (26)$$

$$\mathbf{G}_+^{-1} = \left(\mathbf{I}_d - \frac{\mathbf{u}\mathbf{u}^\top\mathbf{A}}{\mathbf{u}^\top\mathbf{A}\mathbf{u}}\right)\mathbf{G}^{-1}\left(\mathbf{I}_d - \frac{\mathbf{A}\mathbf{u}\mathbf{u}^\top}{\mathbf{u}^\top\mathbf{A}\mathbf{u}}\right) + \frac{\mathbf{u}\mathbf{u}^\top}{\mathbf{u}^\top\mathbf{A}\mathbf{u}}, \quad \mathbf{G}_+ = \text{BFGS}(\mathbf{G}, \mathbf{A}, \mathbf{u}). \quad (27)$$

Based on Lemma 7, we can guarantee a faster superlinear convergence of $\{\lambda_k\}$ (defined in Eqs. (3) and (4)) using the greedy/random SR1 or BFGS update. The proof of Theorem 17 can be found in Appendix C.1.

Theorem 17 *For Algorithm 6 with a randomly initialized \mathbf{G}_0 , such that $\mathbf{G}_0 \succeq \mathbf{A}$ always holds, we have that $\forall k \geq 0, \lambda_{k+1} \leq \rho_k \lambda_k$, where ρ_k is a certain nonnegative random variable such that for SR1 update,*

$$\mathbb{E}\rho_k \leq \left(1 - \frac{k}{d}\right)_+ \frac{\mathbb{E}\tau_0}{\mu}, \quad \forall k \geq 0,$$

and for BFGS update,

$$\mathbb{E}\rho_k \leq \left(1 - \frac{1}{d}\right)^k \mathbb{E}\sigma_0, \quad \forall k \geq 0.$$

We can also use a similar technique in Corollary 9 to give the probabilistic version of Theorem 17, but the differences from greedy/random quasi-Newton methods are clear. In particular, for the SR1 update, our bound recovers the classical result of Nocedal and Wright (2006, Theorem 6.1), showing that the update stops after finite steps because $\mathbf{G}_d = \mathbf{A}$ and $\lambda_{d+1} = 0$ almost surely. Moreover, we give an explicit rate during the entire optimization process. And the main decreasing term $\left(1 - \frac{k}{d}\right)_+$ for the SR1 update as well as $\left(1 - \frac{1}{d}\right)^k$ for the BFGS update in the k -th iteration are independent of the condition number \varkappa of \mathbf{A} , which improves the bound $\left(1 - \frac{1}{d\varkappa}\right)^k$ by Rodomanov and Nesterov (2021b, Theorem 3.4).

Algorithm 7 Greedy/Random SR1/BFGS methods for strongly self-concordant objective

- 1: Initialization: Choose $\mathbf{x}_0 \in \mathbb{R}^d$ and $\mathbf{G}_0 \succeq \nabla^2 f(\mathbf{x}_0)$, $\mathbf{L}_0^\top \mathbf{L}_0 = \mathbf{G}_0^{-1}$.
 - 2: **for** $k \geq 0$ **do**
 - 3: Update $\mathbf{x}_{k+1} = \mathbf{x}_k - \mathbf{G}_k^{-1} \nabla f(\mathbf{x}_k)$.
 - 4: Compute $r_k = \|\mathbf{x}_{k+1} - \mathbf{x}_k\|_{\mathbf{x}_k}$, $\tilde{\mathbf{G}}_k = (1 + Mr_k) \mathbf{G}_k$, $\tilde{\mathbf{L}}_k = \mathbf{L}_k / \sqrt{1 + Mr_k}$.
 - 5: (i) Greedy/Random SR1: Choose $\mathbf{u}_k = \tilde{\mathbf{u}}_{\nabla^2 f(\mathbf{x}_{k+1})}(\tilde{\mathbf{G}}_k)$, or $\mathbf{u}_k \sim \text{Unif}(\mathcal{S}^{d-1})$.
 Compute $\mathbf{G}_{k+1} = \text{SR1}(\tilde{\mathbf{G}}_k, \nabla^2 f(\mathbf{x}_{k+1}), \mathbf{u}_k)$.
 - 6: (ii) Greedy/Random BFGS: Choose $\mathbf{u}_k = \tilde{\mathbf{L}}_k^\top \tilde{\mathbf{u}}_k$ with $\tilde{\mathbf{u}}_k = \tilde{\mathbf{u}}_{\nabla^2 f(\mathbf{x}_{k+1})}(\tilde{\mathbf{L}}_k)$, or $\tilde{\mathbf{u}}_k \sim \text{Unif}(\mathcal{S}^{d-1})$.
 Compute $\mathbf{G}_{k+1} = \text{BFGS}(\tilde{\mathbf{G}}_k, \nabla^2 f(\mathbf{x}_{k+1}), \mathbf{u}_k)$, and \mathbf{L}_{k+1} based on Eq. (25).
 - 7: **end for**
-

4.4 Minimization of General Functions

Finally, we consider the optimization of an M -strongly self-concordant, μ -strongly convex and L -smooth objective as Subsection 3.3 does. We show the entire iteration coupled with our modified update rules in Algorithm 7. We underline that Algorithm 7 is only for theoretical analysis, and we will use the inverse update rules (Eqs. (26) and (27)) and Hessian-vector products in practice. Additionally, we assume that $d \geq 2$, and the constants M and L are available for simplicity. Using the same proof technique, we could obtain faster convergence rates of $\{\lambda_k\}$ (defined in Eqs. (3) and (4)) for greedy/random SR1 or BFGS method. The proof of Lemma 18 can be found in Appendix D.1.

Lemma 18 *Suppose in Algorithm 7, a randomly initialized \mathbf{G}_0 always satisfies $\nabla^2 f(\mathbf{x}_0) \preceq \mathbf{G}_0 \preceq \eta \nabla^2 f(\mathbf{x}_0)$ for some $\eta \geq 1$, and the initial point \mathbf{x}_0 is sufficiently close to the solution:*

$$M\lambda_0 \leq \frac{\ln 2}{4\eta(2cd + 1)},$$

where $c = 1$ for BFGS update and $c = \varkappa$ for SR1 update. Then for all $k \geq 0$, we have $\nabla^2 f(\mathbf{x}_k) \preceq \mathbf{G}_k \preceq (1 + \delta_k) \nabla^2 f(\mathbf{x}_k)$, where δ_k is a certain nonnegative random variable such that

$$\mathbb{E}\delta_k \leq 2cd\eta \left(1 - \frac{1}{d}\right)^k,$$

and $\lambda_{k+1} \leq \rho_k \lambda_k$, where ρ_k is a certain nonnegative random variable such that

$$\mathbb{E}\rho_k \leq 2cd\eta \left(1 - \frac{1}{d}\right)^k.$$

Similarly, we can give deterministic results of greedy methods and probabilistic results of randomized methods below. We leave the proof of Theorem 19 in Appendix D.2.

Theorem 19 *Under the same assumptions and notation as in Lemma 18, we have the explicit rates of $\{\lambda_k\}$ and $\{\delta_k\}$ shown in below:*

- for greedy BFGS/SR1 method, we have

$$\delta_k \leq 2cd\eta \left(1 - \frac{1}{d}\right)^k \quad \text{and} \quad \lambda_k \leq (2cd\eta)^k \left(1 - \frac{1}{d}\right)^{k(k-1)/2} \lambda_0, \forall k \geq 0;$$

- for random BFGS/SR1 method, with probability at least $1 - \delta$ over the random directions $\{\mathbf{u}_k\}$, we could obtain

$$\delta_k \leq \frac{4cd^3\eta}{\delta} \left(1 - \frac{1}{d+1}\right)^k \quad \text{and} \quad \lambda_k \leq \left(\frac{4cd^3\eta}{\delta}\right)^k \left(1 - \frac{1}{d+1}\right)^{k(k-1)/2} \lambda_0, \forall k \geq 0.$$

Finally, we combine with the linear convergence shown in Theorem 4.7 of Rodomanov and Nesterov (2021b) to give fair comparison of our superlinear convergence rates. Under the SR1 update, unlike the measure $\sigma_{\mathbf{A}}(\cdot)$ used by Rodomanov and Nesterov (2021b), we employ a different measure $\tau_{\mathbf{A}}(\cdot)$, requiring a stronger initial point condition to derive the convergence of $\{\lambda_k\}$ and $\{\delta_k\}$. Fortunately, we could obtain the same convergence bound with a slightly worse k_0 below. The proof of Corollary 20 is given in Appendix D.3.

Corollary 20 *Suppose in Algorithm 7, $\mathbf{G}_0 = L\mathbf{I}_d$ and \mathbf{x}_0 satisfies $M\lambda_0 \leq \frac{\ln \frac{3}{4\kappa}}{4\kappa}$, that is, the initial condition here is weaker than that in Lemma 18 when $\eta = \kappa$. Then we could obtain:*

- 1) for the greedy BFGS/SR1 method,

$$\lambda_{k_0+k} \leq \left(1 - \frac{1}{d}\right)^{k(k-1)/2} \cdot \left(\frac{1}{2}\right)^k \cdot \left(1 - \frac{1}{2\kappa}\right)^{k_0} \cdot \lambda_0, \text{ for all } k \geq 0,$$

where $k_0 = O((d + \kappa) \ln(d\kappa))$; 2) for the random BFGS/SR1 method, with probability at least $1 - \delta$ over the random directions $\{\mathbf{u}_k\}$, we have

$$\lambda_{k_0+k} \leq \left(1 - \frac{1}{d+1}\right)^{k(k-1)/2} \cdot \left(\frac{1}{2}\right)^k \cdot \left(1 - \frac{1}{2\kappa}\right)^{k_0} \cdot \lambda_0, \text{ for all } k \geq 0,$$

where $k_0 = O((d + \kappa) \ln(d\kappa/\delta))$.

Therefore, both the greedy and random methods have nonasymptotic superlinear convergence rates. Additionally, our superlinear rates are condition-number-free compared to the rates in Corollary 12 and the work of Rodomanov and Nesterov (2021b).

5. Discussion and Comparison

For better understanding the difference from the greedy quasi-Newton methods obtained in Rodomanov and Nesterov (2021b), we give detailed comparison from the scope of the local convergence region and superlinear rates with $\mathbf{G}_0 = L\mathbf{I}_d$, i.e., $\eta = \kappa$ in our results.

Local Convergence Region. Because we follow the proof of Rodomanov and Nesterov (2021b)'work, our linear convergence region is the same as theirs, i.e., $M\lambda_0 = O(\frac{1}{\kappa})$. Our superlinear convergence region of greedy/random BFGS (Lemma 18) and random Broyden (Lemma 10) is the same as the one obtained in Rodomanov and Nesterov (2021b, Theorem

| Quasi-Newton Methods with $\mathbf{G}_0 = L\mathbf{I}_d$ | Local Region ($M\lambda_0$) | Warm-up (K_1) | Starting Moment (K_2) |
|---|-------------------------------------|--------------------------|---|
| Greedy/Random Broyden (Rodomanov and Nesterov, 2021b) (Lemma 10 and Theorem 11) | $O\left(\frac{1}{d\kappa}\right)$ | $O(\kappa \ln d)$ | $O(d\kappa \ln(d\kappa))$ $O(d\kappa \ln(d\kappa/\delta))$ |
| Greedy/Random BFGS (Lemma 18 and Theorem 19) | $O\left(\frac{1}{d\kappa}\right)$ | $O(\kappa \ln d)$ | $O(d \ln(d\kappa))$ $O(d \ln(d\kappa/\delta))$ |
| Greedy/Random SR1 (Lemma 18 and Theorem 19) | $O\left(\frac{1}{d\kappa^2}\right)$ | $O(\kappa \ln(d\kappa))$ | $O(d \ln(d\kappa))$ $O(d \ln(d\kappa/\delta))$ |

Table 2: Comparison of 1) the local superlinear convergence region, 2) the warm-up iterations from linear rate region to superlinear rate region, and 3) the starting moment of superlinear rates at the local superlinear convergence region. We all adopt $\mathbf{G}_0 = L\mathbf{I}_d$ for brevity. For the randomized methods, the presented rates hold with probability at least $1 - \delta$.

4.9) for greedy Broyden method: $M\lambda_0 = O\left(\frac{1}{d\kappa}\right)$. While our greedy/random SR1 (Lemma 18) needs a slight worse local region $M\lambda_0 = O\left(\frac{1}{d\kappa^2}\right)$, because we use a different measure.

Different local regions show different warm-up iterations from linear rate region to superlinear rate region. Recall that the linear rates of these methods are the same as below:

$$\lambda_k \stackrel{(36)}{\leq} \left(1 - \frac{1}{2\kappa}\right)^k \lambda_0 \leq \exp\left\{-\frac{k}{2\kappa}\right\} \lambda_0, \forall k \geq 0.$$

Thus, with beginning from $M\lambda_0 = O\left(\frac{1}{\kappa}\right)$, the linear rate lasts for $K_1 = O(\kappa \ln d)$ iterations for greedy/random Broyden and BFGS methods to make $M\lambda_{K_1} = O\left(\frac{1}{d\kappa}\right)$, but a slight worse $K_1 = O(\kappa \ln(d\kappa))$ iterations for greedy/random SR1 methods to make $M\lambda_{K_1} = O\left(\frac{1}{d\kappa^2}\right)$.

Superlinear Rates. First, it is obvious that our greedy/random BFGS and SR1 methods have a faster rates than greedy/random Broyden methods because we improve the superlinear term from $(1 - \frac{1}{d\kappa})^{k(k-1)/2}$ to $(1 - \frac{1}{d})^{k(k-1)/2}$.

Second, let us consider the starting moment of superlinear convergence. For random Broyden methods, from Theorem 11, we have the superlinear convergence is valid after

$$K_2^{\text{R-Broyden}} := 2(d\kappa + 1) \ln \frac{4d^3 \kappa^3}{\delta} + 1 \quad (28)$$

iterations. Indeed, from Theorem 11, for all $k \geq K_2^{\text{R-Broyden}}$,

$$\lambda_k \leq \left(\frac{4d^3 \kappa^3}{\delta}\right)^k \left(1 - \frac{1}{d\kappa + 1}\right)^{k(k-1)/2} \lambda_0 \leq \left[\frac{4d^3 \kappa^3}{\delta} \exp\left\{-\frac{k-1}{2(d\kappa + 1)}\right\}\right]^k \lambda_0 \stackrel{(28)}{\leq} \lambda_0.$$

Similarly, from Theorem 19, we could obtain that the superlinear rates of our random BFGS and SR1 methods are valid after

$$K_2^{\text{R-BFGS}} := 2(d+1) \ln \frac{4d^3 \kappa}{\delta} + 1 \text{ and } K_2^{\text{R-SR1}} := 2(d+1) \ln \frac{4d^3 \kappa^2}{\delta} + 1$$

iterations, and the superlinear rates of our greedy BFGS and SR1 methods are valid after

$$K_2^{\text{G-BFGS}} := 2d \ln(2d\kappa) + 1 \text{ and } K_2^{\text{G-SR1}} := 2d \ln(2d\kappa^2) + 1$$

iterations. Moreover, based on Rodomanov and Nesterov (2021b, Theorem 4.9), we get

$$K_2^{\text{G-Broyden}} := 2d\kappa \ln(2d\kappa) + 1.$$

Thus, our proposed greedy/random BFGS and SR1 methods improve the factor $O(d\kappa \ln(d\kappa))$ and $O(d\kappa \ln(d\kappa/\delta))$ of greedy/random Broyden methods to $O(d \ln(d\kappa))$ and $O(d \ln(d\kappa/\delta))$.

Third, we note that the local convergence regions of these methods are different from the discussion. Thus, we consider the whole convergent phase when $M\lambda_0 = O(1/\kappa)$. Based on Corollary 12, Corollary 20 and Rodomanov and Nesterov (2021b, Theorem 4.9), the starting moment of superlinear rates of our proposed greedy/random BFGS and SR1 methods at this time need $O((d + \kappa) \ln(d\kappa))$ (or $O((d + \kappa) \ln(d\kappa/\delta))$), which improves $O(d\kappa \ln(d\kappa))$ (or $O(d\kappa \ln(d\kappa/\delta))$) of greedy/random Broyden methods. For brevity, we summarize the comparison discussed above to Tables 1 and 2.

6. Numerical Experiments

In this section, we verify our theorems through numerical results for quasi-Newton methods. Rodomanov and Nesterov (2021b, Section 5) have already compared their proposed greedy quasi-Newton methods with the classical quasi-Newton methods. They showed that GrDFP, GrBFGS, GrSR1 (greedy DFP, BFGS, SR1 methods) with directions based on $\hat{\mathbf{u}}_{\mathbf{A}}(\mathbf{G})$ (defined in Eq. (9)), have quite competitive convergence with the standard versions. They also presented the results for the randomized versions RaDFP, RaBFGS, RaSR1, which directly choose directions uniformly from the standard Euclidean sphere. They found that the randomized methods are slightly slower than the greedy versions. However, the difference is not really significant.

The difference between our algorithms and their methods mainly comes from the greedy strategy for SR1 and the random strategy for BFGS³. Hence, we mainly focus on exhibiting our validity in these schemes. We refer to GrSR1v2 as our revised method and GrSR1v1 as the previous method (by adopting $\hat{\mathbf{u}}_{\mathbf{A}}(\mathbf{G})$). Similarly, we denote RaBFGSv2 that uses scaled directions ($\mathbf{L}_k^\top \tilde{\mathbf{u}}$) and RaBFGSv1 that directly uses random directions $\tilde{\mathbf{u}}$ correspondingly. We choose the random directions from $\text{Unif}(\mathcal{S}^{d-1})$ in all randomized methods for brevity.

Matrix approximation. When using Algorithms 4 and 5 for approximating a matrix $\mathbf{A} \succ 0$, we show the measure as proved by Theorems 13 and 14 in Figures 1a, 1c, 1d and 1e. As Figure 1a depicts, our greedy and random SR1 updates (GrSR1v2 and RaSR1) share superlinear convergence rates under measure $\tau_{\mathbf{A}}(\cdot)$, while our theoretical bound matches them well. Moreover, Figures 1c, 1d and 1e describe the behavior of the random BFGS update under different condition numbers. Our theory matches the linear convergence of measure $\sigma_{\mathbf{A}}(\cdot)$ in our modified random BFGS update (RaBFGSv2) across different κ s. While directly choosing a direction without scaling (RaBFGSv1) fails to give such bounds. Particularly, a

3. There is no difference in the random SR1 method compared to Rodomanov and Nesterov (2021b), which directly selects random directions. And our greedy BFGS method is not efficient ($O(d^3)$ in each iteration) as we mentioned in Remark 15. Thus we leave it out.

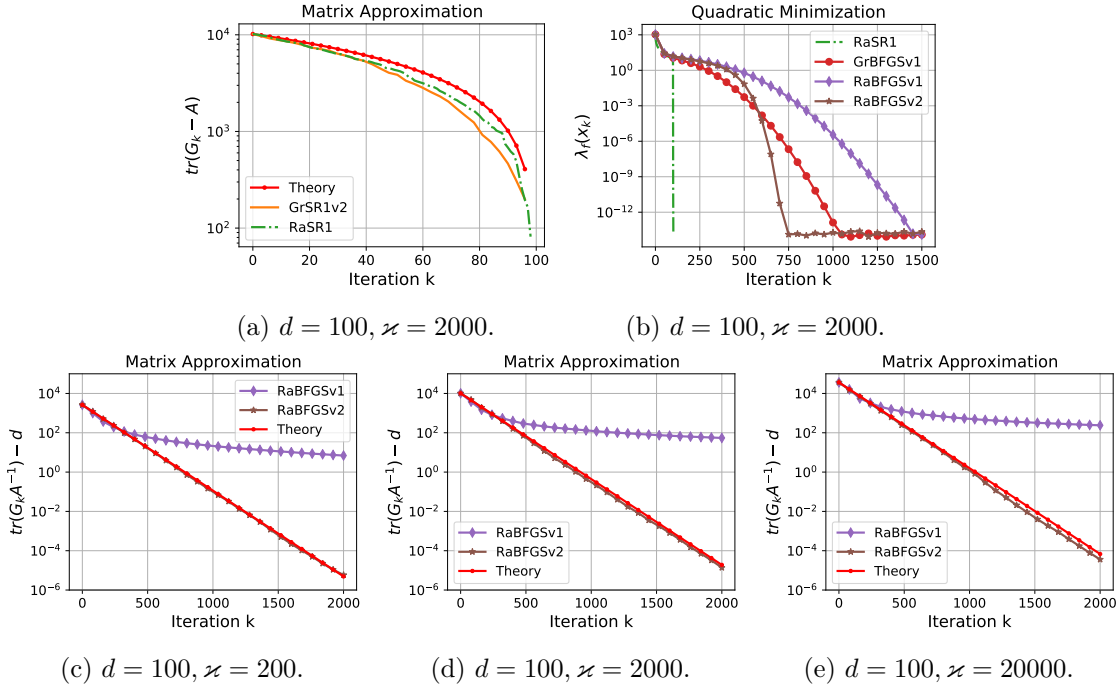


Figure 1: (a, c, d, e) Comparison of different direction choosing methods under the SR1 or BFGS update for approximating a matrix \mathbf{A} that $\mu\mathbf{I}_d \preceq \mathbf{A} \preceq L\mathbf{I}_d$ from $\mathbf{G}_0 = L\mathbf{I}_d$. (a) The variation of $\tau_{\mathbf{A}}(\mathbf{G}_k)$ during Random SR1 (RaSR1) and our Greedy SR1 (GrSR1v2) update with nearly matched upper bound. (c,d,e) The variation of $\sigma_{\mathbf{A}}(\mathbf{G}_k)$ during our Random BFGS (RaBFGSv2) update and the original random version (RaBFGSv1) under various condition numbers. (b) Comparison of SR1 and BFGS methods for quadratic objective. Here we only depict RaSR1 method, while the other SR1-type methods share similar behavior.

large condition number could cause slow convergence of RaBFGSv1. Hence, our methods provide effective ways of approaching a positive definite Hessian matrix.

Quadratic minimization. We also consider unconstrained quadratic minimization in Eq. (14) with the same positive definite matrix \mathbf{A} and a randomly selected vector $\mathbf{b} \in \mathbb{R}^d$. Running Algorithm 6 with SR1 and BFGS updates, we obtain the superlinear convergence of $\lambda_f(\cdot)$ shown in Figure 1b. Not surprisingly, our RaBFGSv2 runs faster than RaBFGSv1, while we also have the theoretical guarantee. At the same time, SR1-type methods converge to zero after $d + 1$ steps because of $\mathbf{G}_d = \mathbf{A}$ almost surely. Here, we only depict the RaSR1 update, while the other SR1-type methods share similar behavior. Although our theoretical bound can not directly match the experiments due to the related initial terms $\tau_{\mathbf{A}}(\mathbf{G}_0)$ and $\sigma_{\mathbf{A}}(\mathbf{G}_0)$, the decay terms: $(1 - k/d)$ vs. $(1 - 1/d)^k$ already show the superiority of the SR1 method over the BFGS method in the quadratic minimization problem.

Regularized Log-Sum-Exp. Following the work of Rodomanov and Nesterov (2021b), we present computational results for greedy and random quasi-Newton methods, applied to

the following test function with $\mathbf{C} = [\mathbf{c}_1, \dots, \mathbf{c}_m] \in \mathbb{R}^{d \times m}$, $b_1, \dots, b_m \in \mathbb{R}$, and $\gamma > 0$:

$$f(\mathbf{x}) := \ln \left(\sum_{j=1}^m e^{\mathbf{c}_j^\top \mathbf{x} - b_j} \right) + \frac{1}{2} \sum_{j=1}^m \left(\mathbf{c}_j^\top \mathbf{x} \right)^2 + \frac{\gamma}{2} \|\mathbf{x}\|^2, \mathbf{x} \in \mathbb{R}^d.$$

We need access to the gradient of function $f(\mathbf{x})$:

$$\nabla f(\mathbf{x}) = g(\mathbf{x}) + \sum_{j=1}^m \left(\mathbf{c}_j^\top \mathbf{x} \right) \mathbf{c}_j + \gamma \mathbf{x}, \text{ with } g(\mathbf{x}) := \sum_{j=1}^m \pi_j(\mathbf{x}) \mathbf{c}_j,$$

where

$$\pi_j(\mathbf{x}) := \frac{e^{\mathbf{c}_j^\top \mathbf{x} - b_j}}{\sum_{i=1}^m e^{\mathbf{c}_i^\top \mathbf{x} - b_i}} \in [0, 1], j = 1, \dots, m.$$

Moreover, given a point $\mathbf{x} \in \mathbb{R}^d$, we need to be able to perform the following two actions:

$$\mathbf{e}_i^\top [\nabla^2 f(\mathbf{x})] \mathbf{e}_i = \sum_{j=1}^m (\pi_j(\mathbf{x}) + 1) (\mathbf{c}_j^\top \mathbf{e}_i)^2 - (g(\mathbf{x})^\top \mathbf{e}_i)^2 + \gamma, \forall 1 \leq i \leq d,$$

and for a given direction $\mathbf{h} \in \mathbb{R}^d$,

$$\nabla^2 f(\mathbf{x}) \cdot \mathbf{h} = \sum_{j=1}^m (\pi_j(\mathbf{x}) + 1) \left(\mathbf{c}_j^\top \mathbf{h} \right) \mathbf{c}_j - \left(g(\mathbf{x})^\top \mathbf{h} \right) g(\mathbf{x}) + \gamma \mathbf{h}.$$

Thus both the above operations have a cost of $O(md)$. Thus, the cost of one iteration for all the methods is comparable. Furthermore, note that

$$\nabla^2 f(\mathbf{x}) = \sum_{j=1}^m (\pi_j(\mathbf{x}) + 1) \mathbf{c}_j \mathbf{c}_j^\top - g(\mathbf{x}) g(\mathbf{x})^\top + \gamma \mathbf{I}_d.$$

We get the Lipschitz constant of $\nabla f(\mathbf{x})$ can be taken as $L = 2\lambda_{\max}(\mathbf{C}\mathbf{C}^\top) + \gamma$, and $\varkappa = L/\gamma$. As mentioned in the work of Rodomanov and Nesterov (2021b, Section 5.1), the strong self-concordancy parameter is $M = 2$ with respect to the operator $\sum_{j=1}^m \mathbf{c}_j \mathbf{c}_j^\top$.

We also adopt the same synthetic data as used by Rodomanov and Nesterov (2021b, Section 5.1). First, we generate a collection of random vectors $\hat{\mathbf{c}}_1, \dots, \hat{\mathbf{c}}_m$ with entries, uniformly distributed in the interval $[-1, 1]$. Then we generate b_1, \dots, b_m from the same distribution. Using this data, we define

$$\forall 1 \leq j \leq m, \mathbf{c}_j := \hat{\mathbf{c}}_j - \nabla \hat{f}(\mathbf{0}), \text{ where } \hat{f}(\mathbf{x}) := \ln \left(\sum_{j=1}^m e^{\hat{\mathbf{c}}_j^\top \mathbf{x} - b_j} \right).$$

Note that by construction,

$$\nabla f(\mathbf{0}) = \frac{1}{\sum_{i=1}^m e^{-b_i}} \sum_{j=1}^m e^{-b_j} \left(\hat{\mathbf{c}}_j - \nabla \hat{f}(\mathbf{0}) \right) = \mathbf{0}.$$




Figure 2: Comparison of SR1 and BFGS updates for Regularized Log-Sum-Exp. The dimension d , the number m of linear functions, the regularization coefficient γ and condition number κ are displayed in the title of each graph. The lines of GrSR1v1 and GrSR1v2 are overlapped in each figure.

So the unique minimizer of our test function is $\mathbf{x}_* = \mathbf{0}$. The starting point \mathbf{x}_0 for all methods is the same and generated randomly from the uniform distribution on the standard Euclidean sphere of radius $1/d$ centered at the minimizer, i.e., $\mathbf{x}_0 \sim \text{Unif}(\frac{1}{d}\mathcal{S}^{d-1})$. We compare $\|\nabla f(\mathbf{x}_k)\|$ obtained by different methods.

As Figure 2 depicts, the BFGS-type methods are slower than the SR1-type methods, and the greedy algorithms converge more rapidly than the random algorithms. The only difference is that our RaBFGSv2 may have slower convergence behavior than RaBFGSv1 under a small κ in Figure 2a.

We consider our scaled direction is more suitable for a constant Hessian matrix as the quadratic objective has. Thus we still have a better convergence rate in the last few iterations when Hessians are nearly unchanged in Figure 2a. However, the Hessian varies drastically in the initial period. Thus there is less benefit under a more accurate Hessian approximation. When applied to the ill-conditioning setting with a large κ in Figures 2b and 2c, we find our RaBFGSv2 could be faster than GrBFGSv1 and RaBFGSv1. This implies that our proposed method has less dependence on the condition number κ .

Regularized Logistic Regression. Finally, we consider a common machine learning problem: ℓ_2 -regularized logistic regression, which has the objective as

$$f(\mathbf{w}) = \sum_{i=1}^n \ln(1 + e^{-y_i \mathbf{w}^\top \mathbf{x}_i}) + \frac{\gamma}{2} \|\mathbf{w}\|^2, \quad \mathbf{w} \in \mathbb{R}^d,$$

where $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n] \in \mathbb{R}^{d \times n}$ are training samples, the corresponding labels are $y_1, \dots, y_n \in \{+1, -1\}$, and $\gamma > 0$ is the regularization coefficient. The gradient of function $f(\mathbf{w})$ is

$$\nabla f(\mathbf{w}) = - \sum_{i=1}^n \frac{1}{1 + e^{y_i \mathbf{w}^\top \mathbf{x}_i}} \cdot y_i \mathbf{x}_i + \gamma \mathbf{w}, \quad \mathbf{w} \in \mathbb{R}^d.$$

Moreover, given a point $\mathbf{w} \in \mathbb{R}^d$, we need to be able to perform the following two actions:

$$\mathbf{e}_j^\top [\nabla^2 f(\mathbf{w})] \mathbf{e}_j = \sum_{i=1}^n \frac{e^{y_i \mathbf{w}^\top \mathbf{x}_i}}{(1 + e^{y_i \mathbf{w}^\top \mathbf{x}_i})^2} \cdot (\mathbf{x}_i^\top \mathbf{e}_j)^2 + \gamma, \quad \forall 1 \leq j \leq d,$$

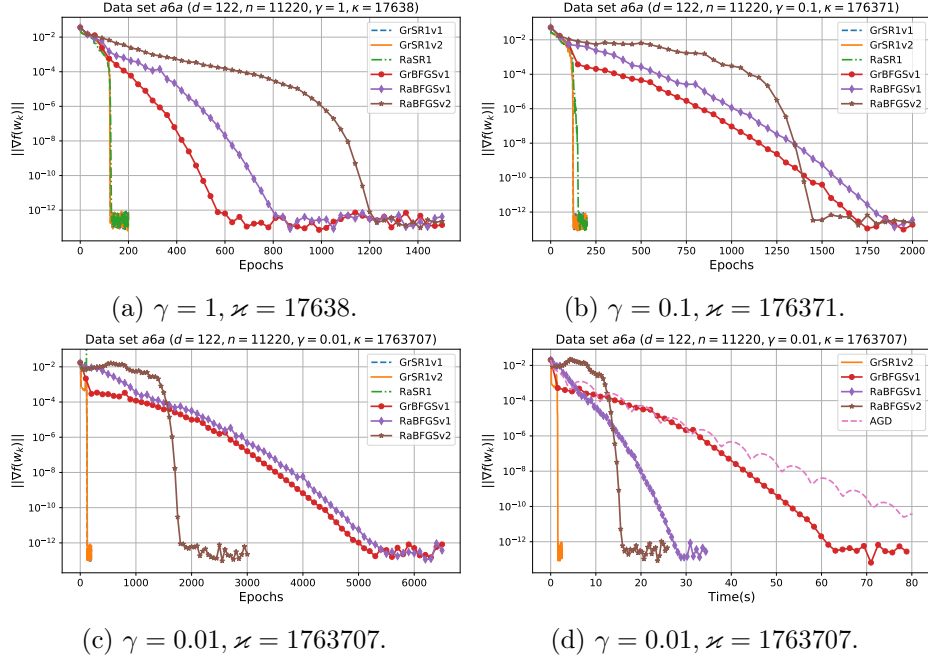


Figure 3: Comparison of SR1 and BFGS update for ℓ_2 -regularized logistic regression applied with ‘a6a’ data from the LIBSVM collection of real-world datasets. We list the name of dataset, the dimension d and the condition number κ under the corresponding γ in the title of each figure. (a,b,c) Comparison of convergence rates with various condition numbers. The lines of GrSR1v1 and GrSR1v2 are overlapped in some figures, and RaSR1 fails when $\gamma = 0.01$ due to the unsuitable initialization and search directions. (d) Comparison of running time with AGD when $\gamma = 0.01$.

and for a given $\mathbf{h} \in \mathbb{R}^d$,

$$\nabla^2 f(\mathbf{w}) \cdot \mathbf{h} = \sum_{i=1}^n \frac{e^{y_i \mathbf{w}^\top \mathbf{x}_i}}{(1 + e^{y_i \mathbf{w}^\top \mathbf{x}_i})^2} \cdot (\mathbf{x}_i^\top \mathbf{h}) \cdot \mathbf{x}_i + \gamma \mathbf{h}.$$

Thus, both the above operations have a cost of $O(nd)$. Furthermore, note that

$$\nabla^2 f(\mathbf{w}) = \sum_{i=1}^n \frac{e^{y_i \mathbf{w}^\top \mathbf{x}_i}}{(1 + e^{y_i \mathbf{w}^\top \mathbf{x}_i})^2} \cdot \mathbf{x}_i \mathbf{x}_i^\top + \gamma \mathbf{I}_d, \quad \mathbf{w} \in \mathbb{R}^d.$$

We obtain that the Lipschitz constant of $\nabla f(\mathbf{w})$ can be taken as $L = \frac{\lambda_{\max}(\mathbf{X}\mathbf{X}^\top)}{4} + \gamma$, and $\kappa = L/\gamma$. Additionally, we take data from the LIBSVM collection of real-world datasets for binary classification problems (Chang and Lin, 2011). And we *do not* apply the *correction strategy* (i.e., $\tilde{\mathbf{G}}_k = (1 + Mr_k)\mathbf{G}_k$) shown in Algorithm 7 recommended by Rodomanov and Nesterov (2021b, Section 5.2). In order to simulate the local convergence, we use the same initialization after running several standard Newton’s steps to make the measure $\|\nabla f(\mathbf{w}_0)\|$ small (around $10^{-2} \sim 10^0$). We compare $\|\nabla f(\mathbf{w}_k)\|$ obtained by different methods.

We show the results in Figure 3. As we can see, the general picture is the same as the Regularized Log-Sum-Exp. In particular, SR1-type methods are still faster than BFGS-type methods, and the greedy algorithms also converge more rapidly than the random algorithms. GrBFGSv1 and RaBFGSv1 are faster than RaBFGSv2 in a small condition number case in Figure 3a, but they become slower than RaBFGSv2 when the condition number becomes huge in Figures 3b and 3c. Therefore, we think our RaBFGSv2 which uses scaled direction indeed has less dependence on the condition number as our theory shows.

In addition, we also compare the running time of each method with a classical first-order method: accelerated gradient descent (AGD) following Nesterov (2003). We run the standard AGD algorithm for 25000 epochs with the same setting in Figure 3c. As Figure 3d shows, not surprisingly, we could discover the benefit of quasi-Newton methods in running time due to their superlinear convergence rates. Furthermore, we find that GrBFGSv1 takes more time compared to the other methods because of the greedy step for searching directions. Thus, the greedy method is time-consuming as Rodomanov and Nesterov (2021b) discussed. But the random method solves this problem through a random choice of directions. Moreover, we discover that random methods may fail for unsuitable initialization as the RaSR1 method in Figure 3c shows, since we may encounter bad random directions during iterations, and our theoretical guarantee is also a probabilistic description. Hence, we recommend a mixture of greedy and random strategies in practice to balance the convergence and running complexity.

Overall, our proposed methods do not lose the superlinear convergence rates particularly in the large condition number schemes, while we also present the theoretical guarantee for these algorithms.

7. Conclusion

In this work, we have addressed two open problems mentioned by Rodomanov and Nesterov (2021b). First, we have shown theoretical analysis of the random quasi-Newton methods, which also preserve a similar nonasymptotic superlinear convergence shown in the work of Rodomanov and Nesterov (2021b). Second, we have studied the behavior of two specific famous quasi-Newton methods: the SR1 and BFGS methods. We have presented different greedy methods in contrast to the work of Rodomanov and Nesterov (2021b), as well as the random version of these methods. In particular, we have provided the faster Hessian approximation behavior and the condition-number-free (local) superlinear convergence rates applied to quadratic or strongly self-concordant objectives. Moreover, the experiments match our analysis well. We hope that the theoretical analysis and the related work would be useful for understanding the explicit rates of quasi-Newton methods, and such convergence rates could benefit machine learning by developing new optimization methods.

Acknowledgments

We would like to thank the anonymous reviewers for their careful work and constructive comments that greatly help us improve the paper quality. We also thank an anonymous reviewer for pointing out efficient update in Proposition 16 and concise formulation in Lemmas 10 and 18, as well as providing a special lemma (Lemma 25) in Appendix.

Haishan Ye has been supported by the National Natural Science Foundation of China (No. 12101491).

Appendix A. Auxiliary Lemmas and Theorems

In the following, assume the objective $f(\mathbf{x})$ is an M -strongly self-concordant, μ -strongly convex and L -smooth function as Subsection 3.3 does.

Lemma 21 (Rodomanov and Nesterov, 2021b, Lemma 4.2) *Let $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, and $r := \|\mathbf{y} - \mathbf{x}\|_{\mathbf{x}}$. Then we have*

$$\frac{\nabla^2 f(\mathbf{x})}{1 + Mr} \preceq \nabla^2 f(\mathbf{y}) \preceq (1 + Mr) \nabla^2 f(\mathbf{x}). \quad (29)$$

Also, for $\mathbf{J} := \int_0^1 \nabla^2 f(\mathbf{x} + t(\mathbf{y} - \mathbf{x})) dt$ and any $\mathbf{v} \in \{\mathbf{x}, \mathbf{y}\}$, we have

$$\frac{\nabla^2 f(\mathbf{v})}{1 + \frac{Mr}{2}} \preceq \mathbf{J} \preceq \left(1 + \frac{Mr}{2}\right) \nabla^2 f(\mathbf{v}).$$

Lemma 22 (Rodomanov and Nesterov, 2021b, Lemmas 4.3 and 4.4) *Let $\mathbf{x} \in \mathbb{R}^d$, and a symmetric matrix \mathbf{G} , such that $\nabla^2 f(\mathbf{x}) \preceq \mathbf{G} \preceq \eta \nabla^2 f(\mathbf{x})$, for some $\eta \geq 1$. Let $\mathbf{x}_+ \in \mathbb{R}^d$, and $r = \|\mathbf{x}_+ - \mathbf{x}\|_{\mathbf{x}}$. Then we have*

$$\tilde{\mathbf{G}} := (1 + Mr) \mathbf{G} \succeq \nabla^2 f(\mathbf{x}_+), \quad (30)$$

and for all $\mathbf{u} \in \mathbb{R}^d$ and $\tau \in [0, 1]$, we have

$$\nabla^2 f(\mathbf{x}_+) \preceq \text{Broyd}_\tau \left(\tilde{\mathbf{G}}, \nabla^2 f(\mathbf{x}_+), \mathbf{u} \right) \preceq \left[(1 + Mr)^2 \eta \right] \nabla^2 f(\mathbf{x}_+). \quad (31)$$

More specifically, if $\mathbf{x}_+ = \mathbf{x} - \mathbf{G}^{-1} \nabla f(\mathbf{x})$, and letting $\lambda := \lambda_f(\mathbf{x})$ be such that $M\lambda \leq 2$, then,

$$r \leq \lambda, \text{ and } \lambda_f(\mathbf{x}_+) \leq \left(1 + \frac{M\lambda}{2}\right) \frac{\eta - 1 + \frac{M\lambda}{2}}{\eta} \cdot \lambda. \quad (32)$$

Theorem 23 (Extension of Rodomanov and Nesterov, 2021b, Theorem 4.7) *Suppose in Algorithm 8, a (random) initialization \mathbf{G}_0 satisfies*

$$\nabla^2 f(\mathbf{x}_0) \preceq \mathbf{G}_0 \preceq \eta \nabla^2 f(\mathbf{x}_0) \quad (33)$$

for some $\eta \geq 1$, and the (random) initial point \mathbf{x}_0 is sufficiently close to the solution:

$$M\lambda_0 \leq \frac{\ln \frac{3}{2}}{4\eta}. \quad (34)$$

Then, for all $k \geq 0$, we have

$$\nabla^2 f(\mathbf{x}_k) \preceq \mathbf{G}_k \preceq e^{2M \sum_{i=0}^{k-1} \lambda_i} \eta \nabla^2 f(\mathbf{x}_k) \preceq \frac{3\eta}{2} \nabla^2 f(\mathbf{x}_k), \quad (35)$$

and

$$\lambda_k \leq \left(1 - \frac{1}{2\eta}\right)^k \lambda_0, \quad (36)$$

where λ_k is defined in Eqs. (3) and (4).

Algorithm 8 Quasi-Newton Method (Rodomanov and Nesterov, 2021b, Scheme (4.17))

- 1: Initialization: Choose $\mathbf{G}_0 \succeq \nabla^2 f(\mathbf{x}_0)$ (such as $\mathbf{G}_0 = L\mathbf{I}_d$).
 - 2: **for** $k \geq 0$ **do**
 - 3: Update $\mathbf{x}_{k+1} = \mathbf{x}_k - \mathbf{G}_k^{-1} \nabla f(\mathbf{x}_k)$.
 - 4: Compute $r_k = \|\mathbf{x}_{k+1} - \mathbf{x}_k\|_{\mathbf{x}_k}$ and set $\tilde{\mathbf{G}}_k = (1 + Mr_k) \mathbf{G}_k$.
 - 5: Choose $\mathbf{u}_k \in \mathbb{R}^d$ and $\tau_k \in [0, 1]$.
 - 6: Compute $\mathbf{G}_{k+1} = \text{Broyd}_{\tau_k}(\tilde{\mathbf{G}}_k, \nabla^2 f(\mathbf{x}_{k+1}), \mathbf{u}_k)$.
 - 7: **end for**
-

Proof The proof is similar as Theorem 4.7 by Rodomanov and Nesterov (2021b). We give the detail for completeness.

From Eq. (33), both Eqs. (35) and (36) are satisfied for $k = 0$. Now let $k \geq 0$, and suppose Eqs. (35) and (36) have already been proved for all $0 \leq k' \leq k$. Denote

$$\eta_k := e^{2M \sum_{i=0}^{k-1} \lambda_i} \eta. \quad (37)$$

Then by inductive hypothesis, we have

$$\nabla^2 f(\mathbf{x}_k) \preceq \mathbf{G}_k \preceq \eta_k \nabla^2 f(\mathbf{x}_k). \quad (38)$$

Note that

$$M \sum_{i=0}^k \lambda_i \stackrel{(36)}{\leq} M \lambda_0 \sum_{i=0}^k \left(1 - \frac{1}{2\eta}\right)^i \leq 2\eta M \lambda_0 \stackrel{(34)}{\leq} \frac{\ln \frac{3}{2}}{2} < 1. \quad (39)$$

Hence, $M\lambda_k < 2$, and by Lemma 22, we have

$$r_k := \|\mathbf{x}_{k+1} - \mathbf{x}_k\|_{\mathbf{x}_k} \stackrel{(32)}{\leq} \lambda_k, \quad (40)$$

and

$$\lambda_{k+1} \stackrel{(32)(38)}{\leq} \left(1 + \frac{M\lambda_k}{2}\right) \frac{\eta_k - 1 + \frac{M\lambda_k}{2}}{\eta_k} \lambda_k \leq \left(1 + \frac{M\lambda_k}{2}\right) \left(1 - \frac{1 - \frac{M\lambda_k}{2}}{\eta_k}\right) \lambda_k. \quad (41)$$

Using the inequality $1 - t \geq e^{-2t}$ when $0 \leq t \leq 1/2$, we get

$$\frac{1 - \frac{M\lambda_k}{2}}{\eta_k} \stackrel{(39)}{\geq} e^{-M\lambda_k} \eta_k^{-1} \stackrel{(37)}{\geq} e^{-2M \sum_{i=0}^k \lambda_i} \eta^{-1} \stackrel{(39)}{\geq} \frac{2}{3\eta}.$$

Moreover, noting that $M\lambda_k \stackrel{(36)}{\leq} M\lambda_0 \stackrel{(34)}{\leq} \frac{1}{8\eta}$, we obtain

$$\left(1 + \frac{M\lambda_k}{2}\right) \left(1 - \frac{1 - \frac{M\lambda_k}{2}}{\eta_k}\right) \leq \left(1 + \frac{1}{16\eta}\right) \left(1 - \frac{2}{3\eta}\right) \leq 1 - \frac{2}{3\eta} + \frac{1}{16\eta} \leq 1 - \frac{1}{2\eta}.$$

Consequently,

$$\lambda_{k+1} \stackrel{(41)}{\leq} \left(1 - \frac{1}{2\eta}\right) \lambda_k \stackrel{(36)}{\leq} \left(1 - \frac{1}{2\eta}\right)^{k+1} \lambda_0.$$

Finally, from Eq. (31) in Lemma 22, it follows that

$$\begin{aligned}\nabla^2 f(\mathbf{x}_{k+1}) &\preceq \mathbf{G}_{k+1} \preceq (1 + Mr_k)^2 \eta_k \nabla^2 f(\mathbf{x}_{k+1}) \stackrel{(40)}{\preceq} (1 + M\lambda_k)^2 \eta_k \nabla^2 f(\mathbf{x}_{k+1}) \\ &\preceq e^{2M\lambda_k} \eta_k \nabla^2 f(\mathbf{x}_{k+1}) \stackrel{(37)}{=} e^{2M \sum_{i=0}^k \lambda_i} \eta \nabla^2 f(\mathbf{x}_{k+1}) \stackrel{(39)}{\preceq} \frac{3\eta}{2} \nabla^2 f(\mathbf{x}_{k+1}).\end{aligned}$$

Thus, Eqs. (35) and (36) are valid for index $k + 1$, and we can continue by induction. \blacksquare

Remark 24 Note that the choice of $\{\mathbf{u}_k\}$ and the update rule in Algorithm 8 are arbitrary, thus Algorithms 3 and 7 can be viewed as special cases of Algorithm 8. Therefore, Theorem 23 always holds as long as the initial point \mathbf{x}_0 is always sufficiently close to the solution based on the initial approximate matrix \mathbf{G}_0 , and Eq. (36) holds without any expectations.

Lemma 25 (Extension of Rodomanov and Nesterov, 2021b, Lemma 4.8)

Following the update in Algorithms 3 or 7, we can obtain

$$\forall k \geq 0, \nabla^2 f(\mathbf{x}_k) \preceq \mathbf{G}_k \preceq (1 + \delta_k) \nabla^2 f(\mathbf{x}_k), \quad (42)$$

where δ_k is a random sequence which satisfies the following recurrence:

$$\mathbb{E}_{\mathbf{u}_k} \delta_{k+1} \leq (1 - t^{-1}) (1 + Mr_k)^2 (\delta_k + 2cdMr_k) \quad (43)$$

for some constants $c, t \geq 1$. Particularly,

- 1) for the Broyden method in Algorithm 3, one has $\delta_k = \sigma_k, t = d\mathfrak{x}, c = 1$;
- 2) for the BFGS method in Algorithm 7, one has $\delta_k = \sigma_k, t = d, c = 1$;
- 3) for the SR1 method in Algorithm 7, one has $\delta_k = d\mathfrak{x}\tau_k / \text{tr}[\nabla^2 f(\mathbf{x}_k)], t = d, c = \mathfrak{x}$.

Here, $r_k := \|\mathbf{x}_{k+1} - \mathbf{x}_k\|_{\mathbf{x}_k}$, σ_k and τ_k are defined in Eq. (4).

Proof From the update rule in Algorithms 3 and 7, we apply Lemma 22 to obtain

$$\mathbf{G}_k \stackrel{(30)}{\succeq} \nabla^2 f(\mathbf{x}_k), \quad \forall k \geq 0. \quad (44)$$

Now for all $k \geq 0$, we define $\eta_k = \left\| [\nabla^2 f(\mathbf{x}_k)]^{-1/2} \mathbf{G}_k [\nabla^2 f(\mathbf{x}_k)]^{-1/2} \right\|$. Then $\eta_k \stackrel{(44)}{\geq} 1$ and $\mathbf{G}_k \preceq \eta_k \nabla^2 f(\mathbf{x}_k)$. Hence,

$$\begin{aligned}\eta_k - 1 &= \left\| [\nabla^2 f(\mathbf{x}_k)]^{-1/2} (\mathbf{G}_k - \nabla^2 f(\mathbf{x}_k)) [\nabla^2 f(\mathbf{x}_k)]^{-1/2} \right\| \\ &\leq \text{tr} \left([\nabla^2 f(\mathbf{x}_k)]^{-1/2} (\mathbf{G}_k - \nabla^2 f(\mathbf{x}_k)) [\nabla^2 f(\mathbf{x}_k)]^{-1/2} \right) \stackrel{(4)}{=} \sigma_k.\end{aligned}$$

Moreover, from $\mu \mathbf{I}_d \preceq \nabla^2 f(\mathbf{x}_k) \preceq L \mathbf{I}_d, \forall k \geq 0$, we also have

$$\begin{aligned}\eta_k - 1 &= \left\| [\nabla^2 f(\mathbf{x}_k)]^{-1/2} (\mathbf{G}_k - \nabla^2 f(\mathbf{x}_k)) [\nabla^2 f(\mathbf{x}_k)]^{-1/2} \right\| \leq \frac{\|\mathbf{G}_k - \nabla^2 f(\mathbf{x}_k)\|}{\mu} \\ &\leq \frac{\text{tr}[\mathbf{G}_k - \nabla^2 f(\mathbf{x}_k)]}{\mu} \stackrel{(4)}{=} \frac{\tau_k}{\mu} \leq \frac{dL}{\text{tr}[\nabla^2 f(\mathbf{x}_k)]} \cdot \frac{\tau_k}{\mu} = \frac{d\mathfrak{x}\tau_k}{\text{tr}[\nabla^2 f(\mathbf{x}_k)]}.\end{aligned}$$

Therefore, the choices of $\delta_k = \sigma_k$ and $\frac{d\mathcal{X}\tau_k}{\text{tr}[\nabla^2 f(\mathbf{x}_k)]}$ are valid to guarantee Eq. (42). Next, we deduce Eq. (43) based on the specific choice of δ_k .

1) For the Broyden method in Algorithm 3, by Theorem 6, one step update gives

$$\mathbb{E}_{\mathbf{u}_k} \sigma_{k+1} \stackrel{(4)}{=} \mathbb{E}_{\mathbf{u}_k} \sigma_{\nabla^2 f(\mathbf{x}_{k+1})} (\mathbf{G}_{k+1}) \stackrel{(13)}{\leq} \left(1 - \frac{1}{d\mathcal{X}}\right) \sigma_{\nabla^2 f(\mathbf{x}_{k+1})}(\tilde{\mathbf{G}}_k), \forall k \geq 0.$$

We deduce the result by bounding the last term as below:

$$\begin{aligned} \sigma_{\nabla^2 f(\mathbf{x}_{k+1})}(\tilde{\mathbf{G}}_k) &\stackrel{(1)}{=} \text{tr}([\nabla^2 f(\mathbf{x}_{k+1})]^{-1} \tilde{\mathbf{G}}_k) - d = (1 + Mr_k) \text{tr}([\nabla^2 f(\mathbf{x}_{k+1})]^{-1} \mathbf{G}_k) - d \\ &\stackrel{(29)}{\leq} (1 + Mr_k)^2 \text{tr}([\nabla^2 f(\mathbf{x}_k)]^{-1} \mathbf{G}_k) - d \stackrel{(4)}{=} (1 + Mr_k)^2 (\sigma_k + d) - d \\ &= (1 + Mr_k)^2 \sigma_k + d [2Mr_k + (Mr_k)^2] \leq (1 + Mr_k)^2 (\sigma_k + 2dMr_k). \end{aligned} \quad (45)$$

2) For the BFGS method in Algorithm 7, by Theorem 14, one step update gives

$$\mathbb{E}_{\mathbf{u}_k} \sigma_{k+1} \stackrel{(24)}{\leq} \left(1 - \frac{1}{d}\right) \sigma_{\nabla^2 f(\mathbf{x}_{k+1})}(\tilde{\mathbf{G}}_k), \forall k \geq 0.$$

The remaining proof is the same as Eq. (45).

3) For the SR1 method in Algorithm 7, from Theorem 13, one step update gives

$$\mathbb{E}_{\mathbf{u}_k} \tau_{k+1} \stackrel{(19)}{\leq} \left(1 - \frac{1}{d}\right) \tau_{\nabla^2 f(\mathbf{x}_{k+1})}(\tilde{\mathbf{G}}_k), \forall k \geq 0.$$

Now we can bound the last term as below:

$$\begin{aligned} \tau_{\nabla^2 f(\mathbf{x}_{k+1})}(\tilde{\mathbf{G}}_k) &\stackrel{(4)}{=} \text{tr}(\tilde{\mathbf{G}}_k - \nabla^2 f(\mathbf{x}_{k+1})) \stackrel{(29)}{\leq} \text{tr}\left((1 + Mr_k) \mathbf{G}_k - \frac{\nabla^2 f(\mathbf{x}_k)}{1 + Mr_k}\right) \\ &\stackrel{(4)}{=} (1 + Mr_k) \tau_k + \left(1 + Mr_k - \frac{1}{1 + Mr_k}\right) \text{tr}[\nabla^2 f(\mathbf{x}_k)] \\ &\stackrel{(i)}{\leq} (1 + Mr_k) \tau_k + 2Mr_k \text{tr}[\nabla^2 f(\mathbf{x}_k)] \\ &\leq (1 + Mr_k) \left(\frac{\tau_k}{\text{tr}[\nabla^2 f(\mathbf{x}_k)]} + 2Mr_k\right) \text{tr}[\nabla^2 f(\mathbf{x}_k)] \\ &\stackrel{(29)}{\leq} (1 + Mr_k)^2 \left(\frac{\tau_k}{\text{tr}[\nabla^2 f(\mathbf{x}_k)]} + 2Mr_k\right) \text{tr}[\nabla^2 f(\mathbf{x}_{k+1})], \end{aligned}$$

where (i) uses inequality $1 + a - \frac{1}{1+a} \leq 2a, a \geq 0$. Thus, by replacing $\delta_k = \frac{d\mathcal{X}\tau_k}{\text{tr}[\nabla^2 f(\mathbf{x}_k)]}$, we obtain

$$\mathbb{E}_{\mathbf{u}_k} \delta_{k+1} \leq \left(1 - \frac{1}{d}\right) (1 + Mr_k)^2 (\delta_k + 2\mathcal{X}dMr_k).$$

■

Lemma 26 *Suppose there exist some constants $a \geq 0, t > 1$, and a nonnegative random sequence $\{\mathbf{X}_k\}$ satisfies*

$$\mathbb{E}\mathbf{X}_k \leq a \left(1 - \frac{1}{t}\right)^k, \forall k \geq 0.$$

Then for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, we have

$$\forall k \geq 0, \mathbf{X}_k \leq \frac{at^2}{\delta} \left(1 - \frac{1}{1+t}\right)^k.$$

Proof If $a = 0$, then by $\mathbf{X}_k \geq 0, \forall k \geq 0$, we can see $\mathbf{X}_k = 0$, a.s. Then the results trivial hold. Now we consider $a > 0$. Noting that $\mathbf{X}_k \geq 0$ and using Markov's inequality, we have for any $\epsilon > 0$,

$$\mathbb{P}\left(\mathbf{X}_k > \frac{a}{\epsilon} \left(1 - \frac{1}{t}\right)^k\right) \leq \mathbb{P}\left(\mathbf{X}_k \geq \frac{a}{\epsilon} \left(1 - \frac{1}{t}\right)^k\right) \leq \frac{\mathbb{E}\mathbf{X}_k}{a \left(1 - \frac{1}{t}\right)^k} \cdot \epsilon \leq \epsilon. \quad (46)$$

Choosing $\epsilon_k = \delta(1 - q)q^k$ for some $q \in (0, 1)$ and applying the union bound, we obtain

$$\mathbb{P}\left(\exists k \geq 0, \mathbf{X}_k > \frac{a}{\epsilon_k} \left(1 - \frac{1}{t}\right)^k\right) \leq \sum_{k=0}^{+\infty} \mathbb{P}\left(\mathbf{X}_k > \frac{a}{\epsilon_k} \left(1 - \frac{1}{t}\right)^k\right) \stackrel{(46)}{\leq} \sum_{k=0}^{+\infty} \epsilon_k = \delta.$$

Therefore, with probability at least $1 - \delta$, we have

$$\forall k \geq 0, \mathbf{X}_k \leq \frac{a}{\delta(1 - q)q^k} \left(1 - \frac{1}{t}\right)^k.$$

If we let $q = 1 - \frac{1}{t^2}$, then we can simplify the above inequality into

$$\forall k \geq 0, \mathbf{X}_k \leq \frac{at^2}{\delta} \left(1 + \frac{1}{t}\right)^{-k} = \frac{at^2}{\delta} \left(1 - \frac{1}{t+1}\right)^k.$$

■

Appendix B. Missing Proofs of Matrix Approximation

B.1 Proof of Theorem 6

Proof From Lemma 5, we have $\mathbf{G}_k \succeq \mathbf{A}$ and $\sigma_k \stackrel{(4)}{=} \sigma_{\mathbf{A}}(\mathbf{G}_k) \geq 0, \forall k \geq 0$. Moreover, from Eqs. (11) and (12), we get

$$\begin{aligned} \mathbb{E}\mathbf{u}_k \sigma_{k+1} &\stackrel{(11)}{\leq} \sigma_k - \frac{1}{L} \text{tr} \left[(\mathbf{G}_k - \mathbf{A}) \cdot \mathbb{E}\mathbf{u}_k \frac{\mathbf{u}_k \mathbf{u}_k^\top}{\mathbf{u}_k^\top \mathbf{u}_k} \right] \stackrel{(12)}{=} \sigma_k - \frac{1}{dL} \text{tr}(\mathbf{G}_k - \mathbf{A}) \\ &\stackrel{(10)}{\leq} \sigma_k - \frac{\mu}{dL} \text{tr}[(\mathbf{G}_k - \mathbf{A})\mathbf{A}^{-1}] = \left(1 - \frac{1}{d\lambda}\right) \sigma_k. \end{aligned}$$

Finally, taking the expectation of all randomness, we get

$$\mathbb{E}\sigma_k \leq \left(1 - \frac{1}{d\alpha}\right) \mathbb{E}\sigma_{k-1} \leq \cdots \leq \left(1 - \frac{1}{d\alpha}\right)^k \mathbb{E}\sigma_0, \forall k \geq 1.$$

■

B.2 Proof of Theorem 13

Proof Denoting $\mathbf{G}_{k+1} := \text{SR1}(\mathbf{G}_k, \mathbf{A}, \mathbf{u}_k)$ and $\mathbf{R}_k := \mathbf{G}_k - \mathbf{A}, \forall k \geq 0$, we have the update rule:

$$\mathbf{R}_{k+1} \stackrel{(6)}{=} \begin{cases} \mathbf{R}_k - \frac{\mathbf{R}_k \mathbf{u}_k \mathbf{u}_k^\top \mathbf{R}_k}{\mathbf{u}_k^\top \mathbf{R}_k \mathbf{u}_k}, & \text{if } \mathbf{R}_k \mathbf{u}_k \neq \mathbf{0}; \\ \mathbf{R}_k, & \text{otherwise.} \end{cases} \quad (47)$$

We also have $\mathbf{R}_k \succeq \mathbf{0}$ by Lemma 5 since $\mathbf{R}_0 \succeq \mathbf{0}$. It is easily seen that 1) $\text{Ker}(\mathbf{R}_k) \subseteq \text{Ker}(\mathbf{R}_{k+1}) := \{\mathbf{v} : \mathbf{R}_{k+1} \mathbf{v} = \mathbf{0}\}$, and 2) $\mathbf{u}_k \in \text{Ker}(\mathbf{R}_{k+1})$ from Eq. (47).

For the greedy method, we denote $\bar{\mathbf{u}}_k = \bar{\mathbf{u}}_A(\mathbf{G}_k), \forall k \geq 0$. If for some $k' < d$, $\mathbf{R}_{k'} = \mathbf{0}$, then from Eq. (47), we have $\mathbf{R}_k = \mathbf{0}, \forall k \geq k'$. Thus, Eq. (19) trivially holds for $k \geq k'$.

Now we suppose $\forall 0 \leq k < d, \mathbf{R}_k \neq \mathbf{0}$. Then by $\mathbf{R}_k \succeq \mathbf{0}$ and $\mathbf{R}_k \neq \mathbf{0}$, we must have $\bar{\mathbf{u}}_k \notin \text{Ker}(\mathbf{R}_k)$ in view of Eq. (17) for all $0 \leq k < d$. Additionally, $\bar{\mathbf{u}}_i \in \text{Ker}(\mathbf{R}_k), \forall 0 \leq i < k$ by 1) and 2), so we can see $\bar{\mathbf{u}}_i \neq \bar{\mathbf{u}}_k, \forall i \neq k$. Thus, at least k of the diagonal elements of \mathbf{R}_k must be zero, leading to

$$\max_{\mathbf{u} \in \{\mathbf{e}_1, \dots, \mathbf{e}_d\}} \mathbf{u}^\top \mathbf{R}_k \mathbf{u} \geq \frac{1}{d-k} \cdot \text{tr}(\mathbf{R}_k) = \frac{1}{d-k} \cdot \tau_k. \quad (48)$$

Finally, for all $0 \leq k < d$, since $\bar{\mathbf{u}}_k \notin \text{Ker}(\mathbf{R}_k)$, then Eq. (15) is well-defined. Thus, we get

$$\tau_{k+1} \stackrel{(15)}{=} \tau_k - \frac{\bar{\mathbf{u}}_k^\top \mathbf{R}_k^2 \bar{\mathbf{u}}_k}{\bar{\mathbf{u}}_k^\top \mathbf{R}_k \bar{\mathbf{u}}_k} \stackrel{(16)}{\leq} \tau_k - \frac{\bar{\mathbf{u}}_k^\top \mathbf{R}_k \bar{\mathbf{u}}_k}{\bar{\mathbf{u}}_k^\top \bar{\mathbf{u}}_k} \stackrel{(17)}{=} \tau_k - \max_{\mathbf{u} \in \{\mathbf{e}_1, \dots, \mathbf{e}_d\}} \mathbf{u}^\top \mathbf{R}_k \mathbf{u} \stackrel{(48)}{\leq} \left(1 - \frac{1}{d-k}\right) \tau_k.$$

Consequently, we have for all $1 \leq k \leq d$,

$$\tau_k \leq \frac{d-k}{d-k+1} \tau_{k-1} \leq \cdots \leq \left[\prod_{j=1}^k \frac{d-j}{d-j+1} \right] \tau_0 = \left(1 - \frac{k}{d}\right) \tau_0.$$

Then $\tau_d = 0$, which leads to $\mathbf{G}_d = \mathbf{A}$. Further we obtain $\mathbf{G}_k = \mathbf{A}, \forall k \geq d$ following Eq. (47). Therefore, $\tau_k = 0, \forall k \geq d$. We conclude that for all $k \geq 0$, $\tau_k \leq \left(1 - \frac{k}{d}\right)_+ \tau_0$.

For the random method, $\forall k \geq 0, \mathbf{u}_k$ s are independently chosen from an identical spherically symmetric distribution, such as $\mathcal{N}(0, \mathbf{I}_d)$, $\text{Unif}(\mathcal{S}^{d-1})$.

We first consider $0 \leq k < d$. We have $\lambda_i := \lambda_i(\mathbf{R}_k) \geq 0$ since $\mathbf{R}_k \succeq \mathbf{0}$. Suppose $r_k := \text{rank}(\mathbf{R}_k) \geq 1$, i.e., $\mathbf{R}_k \neq \mathbf{0}$. We denote $\mathbf{R}_k = \mathbf{U}_k \mathbf{\Lambda}_k \mathbf{U}_k^\top$ as the spectral decomposition of \mathbf{R}_k with an orthogonal matrix \mathbf{U}_k and a diagonal matrix $\mathbf{\Lambda}_k = \text{diag}\{\lambda_1, \dots, \lambda_{r_k}, 0, \dots, 0\}$,

and $\mathbf{v}_k = (v_1, \dots, v_d)^\top := \mathbf{U}_k^\top \mathbf{u}_k$. Then we can derive that

$$\begin{aligned} & \mathbb{E}_{\mathbf{u}_k} \frac{\mathbf{u}_k^\top \mathbf{R}_k^2 \mathbf{u}_k}{\mathbf{u}_k^\top \mathbf{R}_k \mathbf{u}_k} \mathbb{1}_{\{\mathbf{R}_k \mathbf{u}_k \neq \mathbf{0}\}} \stackrel{(i)}{=} \mathbb{E}_{\mathbf{v}_k} \frac{\sum_{i=1}^{r_k} \lambda_i^2 v_i^2}{\sum_{i=1}^{r_k} \lambda_i v_i^2} \mathbb{1}_{\{\boldsymbol{\Lambda}_k \mathbf{v}_k \neq \mathbf{0}\}} \stackrel{(ii)}{\geq} \mathbb{E}_{\mathbf{v}_k} \frac{\sum_{i=1}^{r_k} \lambda_i v_i^2}{\sum_{i=1}^{r_k} v_i^2} \mathbb{1}_{\{\sum_{i=1}^{r_k} v_i^2 \neq 0\}} \\ & = \sum_{i=1}^{r_k} \lambda_i \mathbb{E}_{\mathbf{v}_k} \frac{v_i^2}{\sum_{j=1}^{r_k} v_j^2} \mathbb{1}_{\{\sum_{i=1}^{r_k} v_i^2 \neq 0\}} \stackrel{(iii)}{=} \frac{1}{r_k} \sum_{i=1}^{r_k} \lambda_i = \frac{\text{tr}(\mathbf{R}_k)}{r_k} \stackrel{(4)}{=} \frac{\tau_k}{r_k}, \end{aligned} \quad (49)$$

where (i) holds due to $\mathbf{u}_k^\top \mathbf{R}_k^i \mathbf{u}_k = \mathbf{v}_k^\top \boldsymbol{\Lambda}_k^i \mathbf{v}_k, \forall i \geq 1$ and $\mathbf{R}_k \mathbf{u}_k \neq \mathbf{0}$ is equivalent to $\boldsymbol{\Lambda}_k \mathbf{v}_k \neq \mathbf{0}$; (ii) holds due to the Cauchy-Schwarz inequality $(\sum_{i=1}^{r_k} \lambda_i^2 v_i^2) (\sum_{i=1}^{r_k} v_i^2) \geq (\sum_{i=1}^{r_k} \lambda_i v_i^2)^2$ with $\sum_{i=1}^{r_k} \lambda_i v_i^2 = \mathbf{u}_k^\top \mathbf{R}_k \mathbf{u}_k > 0$ since $\mathbf{R}_k \mathbf{u}_k \neq \mathbf{0}$ and $\mathbf{R}_k \succeq \mathbf{0}$, and the fact that

$$\boldsymbol{\Lambda}_k \mathbf{v}_k \neq \mathbf{0} \Leftrightarrow \exists 1 \leq i \leq r_k, \text{ s.t.}, \lambda_i v_i \neq 0 \Leftrightarrow \exists 1 \leq i \leq r_k, \text{ s.t.}, v_i \neq 0 \Leftrightarrow \sum_{j=1}^{r_k} v_j^2 \neq 0;$$

(iii) uses the fact that \mathbf{v}_k is still spherically symmetric, thus also permutation invariant:

$$\begin{aligned} & \mathbb{E}_{\mathbf{v}_k} \frac{v_1^2}{\sum_{j=1}^{r_k} v_j^2} \mathbb{1}_{\mathcal{V}_k} = \dots = \mathbb{E}_{\mathbf{v}_k} \frac{v_{r_k}^2}{\sum_{j=1}^{r_k} v_j^2} \mathbb{1}_{\mathcal{V}_k} = \frac{1}{r_k} \sum_{i=1}^{r_k} \mathbb{E}_{\mathbf{v}_k} \frac{v_i^2}{\sum_{j=1}^{r_k} v_j^2} \mathbb{1}_{\mathcal{V}_k} \\ & = \frac{1}{r_k} \mathbb{E}_{\mathbf{v}_k} \sum_{i=1}^{r_k} \frac{v_i^2}{\sum_{j=1}^{r_k} v_j^2} \mathbb{1}_{\mathcal{V}_k} = \frac{1}{r_k} \mathbb{E}_{\mathbf{v}_k} \mathbb{1}_{\mathcal{V}_k} \stackrel{(iv)}{=} \frac{1}{r_k}, \text{ with } \mathcal{V}_k := \{\mathbf{v}_k : \sum_{i=1}^{r_k} v_i^2 \neq 0\}, \end{aligned}$$

where (iv) holds because \mathcal{V}_k^c (the complementary event of \mathcal{V}_k) has zero Lebesgue measure. Therefore, the random choice of \mathbf{u}_k leads to

$$\begin{aligned} & \mathbb{E}_{\mathbf{u}_k} \tau_{k+1} = \mathbb{E}_{\mathbf{u}_k} \tau_{k+1} [\mathbb{1}_{\{\mathbf{R}_k \mathbf{u}_k \neq \mathbf{0}\}} + \mathbb{1}_{\{\mathbf{R}_k \mathbf{u}_k = \mathbf{0}\}}] \stackrel{(v)}{=} \mathbb{E}_{\mathbf{u}_k} \tau_{k+1} \mathbb{1}_{\{\mathbf{R}_k \mathbf{u}_k \neq \mathbf{0}\}} \\ & \stackrel{(15)}{=} \tau_k - \mathbb{E}_{\mathbf{u}_k} \frac{\mathbf{u}_k^\top \mathbf{R}_k^2 \mathbf{u}_k}{\mathbf{u}_k^\top \mathbf{R}_k \mathbf{u}_k} \mathbb{1}_{\{\mathbf{R}_k \mathbf{u}_k \neq \mathbf{0}\}} \stackrel{(49)}{\leq} \left(1 - \frac{1}{r_k}\right) \tau_k, \text{ if } \mathbf{R}_k \neq \mathbf{0}, \end{aligned}$$

where (v) uses the fact that $\mathbf{R}_k \mathbf{u}_k = \mathbf{0}$ is equal to $\sum_{i=1}^{r_k} v_i^2 = 0$, i.e., the event \mathcal{V}_k^c , which has zero measure.

Furthermore, if $\mathbf{u}_0, \dots, \mathbf{u}_{k-1}$ are linearly independent, then by 1) and 2), the dimension of $\text{Ker}(\mathbf{R}_i), i \leq k$ grows at least by one at every iteration, showing that $\text{Ker}(\mathbf{R}_k) \geq k$ and $r_k = \text{rank}(\mathbf{R}_k) = d - \text{Ker}(\mathbf{R}_k) \leq d - k$. Thus, we establish that

$$\mathbb{E}[\tau_{k+1} | \mathcal{M}_k, \mathbf{R}_k \neq \mathbf{0}] \leq \left(1 - \frac{1}{d-k}\right) \mathbb{E}[\tau_k | \mathcal{M}_k, \mathbf{R}_k \neq \mathbf{0}], \quad (50)$$

where $\mathcal{M}_k = \{\mathbf{u}_0, \dots, \mathbf{u}_{k-1} \text{ are linear independent}\}, k \geq 1$ and \mathcal{M}_0 is the full set. Besides, we note that $\mathbf{R}_k = \mathbf{0}$ gives $\mathbf{R}_{k+1} = \mathbf{0}$, then

$$\mathbb{E}[\tau_{k+1} | \mathcal{M}_k, \mathbf{R}_k = \mathbf{0}] = 0 = \left(1 - \frac{1}{d-k}\right) \mathbb{E}[\tau_k | \mathcal{M}_k, \mathbf{R}_k = \mathbf{0}]. \quad (51)$$

Using the law of total expectation conditioning on \mathcal{M}_k^4 , we obtain

$$\mathbb{E}[\tau_{k+1} | \mathcal{M}_k] \stackrel{(50)(51)}{\leq} \left(1 - \frac{1}{d-k}\right) \mathbb{E}[\tau_k | \mathcal{M}_k]. \quad (52)$$

4. $\mathbb{E}[\mathbf{X} | \mathbf{A}, \mathcal{M}] \cdot \mathbb{P}(\mathbf{A} | \mathcal{M}) + \mathbb{E}[\mathbf{X} | \mathbf{A}^c, \mathcal{M}] \cdot \mathbb{P}(\mathbf{A}^c | \mathcal{M}) = \mathbb{E}[\mathbf{X} | \mathcal{M}].$

Noting that $\mathbb{P}(\mathcal{M}_k^c) = \mathbb{P}(\exists 0 \leq t \leq k-1, \mathbf{u}_t \in \text{Span}\{\mathbf{u}_0, \dots, \mathbf{u}_{t-1}, \mathbf{u}_{t+1}, \dots, \mathbf{u}_{k-1}\}) = 0$ since the dimension of $\text{Span}\{\mathbf{u}_0, \dots, \mathbf{u}_{t-1}, \mathbf{u}_{t+1}, \dots, \mathbf{u}_{k-1}\}$ is at most $k-1 < d$, so by the law of total expectation again⁵, we conclude

$$\mathbb{E}\tau_{k+1} \stackrel{(52)}{\leq} \left(1 - \frac{1}{d-k}\right) \mathbb{E}\tau_k, \forall 0 \leq k < d. \quad (53)$$

Consequently, we have for all $1 \leq k \leq d$,

$$\mathbb{E}\tau_k \stackrel{(53)}{\leq} \frac{d-k}{d-k+1} \cdot \mathbb{E}\tau_{k-1} \leq \dots \leq \left[\prod_{j=1}^k \frac{d-j}{d-j+1} \right] \mathbb{E}\tau_0 = \left(1 - \frac{k}{d}\right) \mathbb{E}\tau_0.$$

That is, we obtain $\mathbb{E}\tau_d = 0$, showing that $\tau_d = 0$ a.s. and $\mathbf{G}_d - \mathbf{A} = \mathbf{0}$ a.s., since $\mathbf{G}_d - \mathbf{A} \succeq \mathbf{0}$. Furthermore, following Eq. (47), we obtain $\forall k \geq d, \mathbf{G}_k = \mathbf{A}$ a.s. Therefore, we derive that $\forall k \geq d, \tau_k = 0$ a.s. Then we conclude that $\mathbb{E}\tau_k \leq \left(1 - \frac{k}{d}\right)_+ \mathbb{E}\tau_0$. \blacksquare

B.3 Proof of Theorem 14

Proof For the greedy method, at step $k \geq 0$, since $\mathbf{G}_k^{-1} = \mathbf{L}_k^\top \mathbf{L}_k$, we obtain

$$\max_{\tilde{\mathbf{u}} \in \{e_1, \dots, e_d\}} \tilde{\mathbf{u}}^\top \mathbf{L}_k^{-\top} \mathbf{A}^{-1} \mathbf{L}_k^{-1} \tilde{\mathbf{u}} \geq \frac{1}{d} \text{tr} \left(\mathbf{L}_k^{-\top} \mathbf{A}^{-1} \mathbf{L}_k^{-1} \right) = \frac{1}{d} \text{tr} \left(\mathbf{L}_k^{-1} \mathbf{L}_k^{-\top} \mathbf{A}^{-1} \right) = \frac{1}{d} \text{tr}(\mathbf{G}_k \mathbf{A}^{-1}). \quad (54)$$

Therefore, the greedy choice of $\mathbf{u}_k = \mathbf{L}_k^\top \tilde{\mathbf{u}}_k$ with $\tilde{\mathbf{u}}_k$ following Eq. (22) leads to

$$\begin{aligned} \sigma_{k+1} &\stackrel{(21)}{=} \sigma_k - \frac{\tilde{\mathbf{u}}_k^\top \mathbf{L}_k^{-\top} \mathbf{A}^{-1} \mathbf{L}_k^{-1} \tilde{\mathbf{u}}_k}{\tilde{\mathbf{u}}_k^\top \tilde{\mathbf{u}}_k} + 1 \stackrel{(22)(54)}{\leq} \sigma_k - \frac{1}{d} \text{tr}(\mathbf{G}_k \mathbf{A}^{-1}) + 1 \\ &= \left(1 - \frac{1}{d}\right) \sigma_k \leq \dots \leq \left(1 - \frac{1}{d}\right)^{k+1} \sigma_0, \forall k \geq 0. \end{aligned}$$

For the random method, we have that $\mathbb{E}_{\tilde{\mathbf{u}}} \frac{\tilde{\mathbf{u}} \tilde{\mathbf{u}}^\top}{\tilde{\mathbf{u}}^\top \tilde{\mathbf{u}}} = \frac{1}{d} \mathbf{I}_d$. Hence, we obtain

$$\mathbb{E}_{\tilde{\mathbf{u}}} \frac{\tilde{\mathbf{u}}^\top \mathbf{L}_k^{-\top} \mathbf{A}^{-1} \mathbf{L}_k^{-1} \tilde{\mathbf{u}}}{\tilde{\mathbf{u}}^\top \tilde{\mathbf{u}}} = \text{tr} \left[\mathbf{L}_k^{-\top} \mathbf{A}^{-1} \mathbf{L}_k^{-1} \cdot \mathbb{E}_{\tilde{\mathbf{u}}} \frac{\tilde{\mathbf{u}} \tilde{\mathbf{u}}^\top}{\tilde{\mathbf{u}}^\top \tilde{\mathbf{u}}} \right] = \frac{1}{d} \text{tr}(\mathbf{L}_k^{-\top} \mathbf{A}^{-1} \mathbf{L}_k^{-1}) = \frac{1}{d} \text{tr}(\mathbf{G}_k \mathbf{A}^{-1}), \quad (55)$$

Therefore, the random choice of $\mathbf{u}_k = \mathbf{L}_k^\top \tilde{\mathbf{u}}_k$ leads to

$$\mathbb{E}_{\mathbf{u}_k} \sigma_{k+1} \stackrel{(21)}{=} \sigma_k - \mathbb{E}_{\tilde{\mathbf{u}}_k} \frac{\tilde{\mathbf{u}}_k^\top \mathbf{L}_k^{-\top} \mathbf{A}^{-1} \mathbf{L}_k^{-1} \tilde{\mathbf{u}}_k}{\tilde{\mathbf{u}}_k^\top \tilde{\mathbf{u}}_k} + 1 \stackrel{(55)}{=} \sigma_k - \frac{1}{d} \text{tr}(\mathbf{G}_k \mathbf{A}^{-1}) + 1 = \left(1 - \frac{1}{d}\right) \sigma_k,$$

Now taking expectation for all $\{\mathbf{u}_k\}$, we get

$$\mathbb{E}\sigma_{k+1} = \left(1 - \frac{1}{d}\right) \mathbb{E}\sigma_k = \dots = \left(1 - \frac{1}{d}\right)^{k+1} \mathbb{E}\sigma_0, \forall k \geq 0. \quad \blacksquare$$

5. $\mathbb{E}[\mathbf{X}|\mathcal{M}] \cdot \mathbb{P}(\mathcal{M}) + \mathbb{E}[\mathbf{X}|\mathcal{M}^c] \cdot \mathbb{P}(\mathcal{M}^c) = \mathbb{E}\mathbf{X}$.

Appendix C. Missing Proofs of Quadratic Objective

C.1 Proofs of Theorem 8 and Theorem 17

Proof From Lemma 5, we have $\forall k \geq 0, \mathbf{G}_k \succeq \mathbf{A}$. Now for all $k \geq 0$, we denote $\eta_k := \|\mathbf{A}^{-1/2} \mathbf{G}_k \mathbf{A}^{-1/2}\|$, then $\eta_k \geq 1$ and $\mathbf{G}_k \preceq \eta_k \mathbf{A}$. Hence,

$$\eta_k - 1 = \left\| \mathbf{A}^{-1/2} (\mathbf{G}_k - \mathbf{A}) \mathbf{A}^{-1/2} \right\| \leq \text{tr} \left(\mathbf{A}^{-1/2} (\mathbf{G}_k - \mathbf{A}) \mathbf{A}^{-1/2} \right) \stackrel{(1)}{=} \sigma_k \quad (56)$$

$$\stackrel{(14)}{\leq} \frac{\text{tr}(\mathbf{G}_k - \mathbf{A})}{\mu} \stackrel{(2)}{=} \frac{\tau_k}{\mu}. \quad (57)$$

By Lemma 7, we know that for all $k \geq 0$, $\lambda_{k+1} \leq (\eta_k - 1)\lambda_k$, then we can take $\rho_k = \eta_k - 1$.

1. For Broyden method, using Theorem 6, we get

$$\mathbb{E}\rho_k = \mathbb{E}(\eta_k - 1) \stackrel{(56)}{\leq} \mathbb{E}\sigma_k \stackrel{(13)}{\leq} \left(1 - \frac{1}{d\mathcal{X}}\right)^k \mathbb{E}\sigma_0, \forall k \geq 0.$$

2. For SR1 method, using Theorem 13, we obtain

$$\mathbb{E}\rho_k = \mathbb{E}(\eta_k - 1) \stackrel{(57)}{\leq} \frac{\mathbb{E}\tau_k}{\mu} \stackrel{(19)}{\leq} \left(1 - \frac{k}{d}\right)_+ \frac{\mathbb{E}\tau_0}{\mu}, \forall k \geq 0.$$

3. For BFGS method, using Theorem 14, we get

$$\mathbb{E}\rho_k = \mathbb{E}(\eta_k - 1) \stackrel{(56)}{\leq} \mathbb{E}\sigma_k \stackrel{(24)}{\leq} \left(1 - \frac{1}{d}\right)^k \mathbb{E}\sigma_0, \forall k \geq 0.$$

■

C.2 Proof of Corollary 9

Proof From Theorems 6 and 8, we can apply Lemma 26 with $\mathbf{X}_k = \sigma_k$ or $\rho_k, \forall k \geq 0$ and $a = \mathbb{E}\sigma_0, t = d\mathcal{X}$, i.e., with probability at least $1 - \delta/2$, we get

$$\sigma_k \leq \frac{2d^2 \mathcal{X}^2 \mathbb{E}\sigma_0}{\delta} \left(1 - \frac{1}{d\mathcal{X} + 1}\right)^k, \forall k \geq 0, \quad (58)$$

and with probability at least $1 - \delta/2$, we have

$$\rho_k \leq \frac{2d^2 \mathcal{X}^2 \mathbb{E}\sigma_0}{\delta} \left(1 - \frac{1}{d\mathcal{X} + 1}\right)^k, \forall k \geq 0. \quad (59)$$

Noting that $\lambda_{k+1} \leq \rho_k \lambda_k$ by definition, we further obtain with probability at least $1 - \delta/2$,

$$\lambda_{k+1} \leq \frac{2d^2 \mathcal{X}^2 \mathbb{E}\sigma_0}{\delta} \left(1 - \frac{1}{d\mathcal{X} + 1}\right)^k \lambda_k, \forall k \geq 0, \quad (60)$$

because the latter event (as a set) in Eq. (60) is contained in the former in Eq. (59).

Using the fact that for any sequences $\{\xi_k\}$ and $\{a_k\}$ of nonnegative random variables and nonnegative reals, respectively, it holds that

$$\mathbb{P}\left(\xi_{k+1} \leq \left(\prod_{i=0}^k a_i\right) \xi_0, \forall k \geq 0\right) \geq \mathbb{P}(\xi_{k+1} \leq a_k \xi_k, \forall k \geq 0), \quad (61)$$

because the latter event (as a set) is contained in the former. Thus we can telescope from $k-1$ to 0 in Eq. (60) for all $k \geq 1$. Then we get with probability at least $1 - \delta/2$,

$$\lambda_k \stackrel{(60)(61)}{\leq} \left(\frac{2d^2 \varkappa^2 \mathbb{E}\sigma_0}{\delta}\right)^k \left(1 - \frac{1}{d\varkappa + 1}\right)^{k(k-1)/2} \lambda_0, \forall k \geq 1, \quad (62)$$

and the above inequality trivially holds for $k = 0$.

Finally, applying the union bound again makes both Eqs. (58) and (62) hold with probability at least $1 - \delta$. \blacksquare

Appendix D. Missing Proofs of General Objective

D.1 Proofs of Lemma 10 and Lemma 18

Since the proofs of Lemma 10 and Lemma 18 have many overlapping parts, we recombine them into a lemma below.

Lemma 27 (Restating) *Suppose in Algorithm 3 or Algorithm 7, a random initialization \mathbf{G}_0 always satisfies $\nabla^2 f(\mathbf{x}_0) \preceq \mathbf{G}_0 \preceq \eta \nabla^2 f(\mathbf{x}_0)$ for some $\eta \geq 1$, and the (random) initial point \mathbf{x}_0 is sufficiently close to the solution:*

$$M\lambda_0 \leq \frac{\ln 2}{4\eta(2cd + 1)}. \quad (63)$$

Then for all $k \geq 0$, we have $\nabla^2 f(\mathbf{x}_k) \preceq \mathbf{G}_k \preceq (1 + \delta_k) \nabla^2 f(\mathbf{x}_k)$, where δ_k is a certain random variable such that

$$\mathbb{E}\delta_k \leq 2cd\eta \left(1 - \frac{1}{t}\right)^k, \quad (64)$$

and $\lambda_{k+1} \leq \rho_k \lambda_k$, where ρ_k is a certain random variable such that

$$\mathbb{E}\rho_k \leq 2cd\eta \left(1 - \frac{1}{t}\right)^k. \quad (65)$$

Here, the choices of δ_k, t, c are inherited from Lemma 25.

Proof The derivation is the same as Theorem 4.9 in the work of Rodomanov and Nesterov (2021b) by using Lemma 25. In order to providing better paper readability, we also show the detail below.

In view of Theorem 23, because the initial condition $\frac{\ln 2}{4\eta(2cd+1)} \leq \frac{\ln \frac{3}{2}}{4\eta}$, we get $\nabla^2 f(\mathbf{x}_k) \preceq \mathbf{G}_k, \forall k \geq 0$, and also

$$M \sum_{i=0}^k \lambda_i \stackrel{(36)}{\leq} M\lambda_0 \sum_{i=0}^k \left(1 - \frac{1}{2\eta}\right)^i \leq 2\eta M\lambda_0 \stackrel{(63)}{\leq} \frac{\ln 2}{2(2cd+1)}. \quad (66)$$

Moreover, we need to underline that Eq. (66) holds with no expectation, which is crucial in the following derivation. Next, let us show that $\forall k \geq 0$,

$$\mathbb{E}\theta_k \leq 2cd\eta \left(1 - \frac{1}{t}\right)^k, \text{ where } \theta_k := \delta_k + 2cdM\lambda_k. \quad (67)$$

Indeed, because $\nabla^2 f(\mathbf{x}_0) \preceq \mathbf{G}_0 \preceq \eta \nabla^2 f(\mathbf{x}_0)$, we have

$$\mathbb{E}\sigma_0 \stackrel{(4)}{=} \mathbb{E}\text{tr}(\nabla^2 f(\mathbf{x}_0)^{-1} \mathbf{G}_0) - d \leq d(\eta - 1),$$

and

$$\mathbb{E} \frac{d\boldsymbol{\varkappa}\tau_0}{\text{tr}[\nabla^2 f(\mathbf{x}_0)]} \stackrel{(4)}{=} d\boldsymbol{\varkappa} \cdot \frac{\mathbb{E}\text{tr}[\mathbf{G}_0 - \nabla^2 f(\mathbf{x}_0)]}{\text{tr}[\nabla^2 f(\mathbf{x}_0)]} \leq \boldsymbol{\varkappa}d(\eta - 1).$$

Hence, following the choice of δ_0 from Lemma 25, we derive that

$$\mathbb{E}\theta_0 = \mathbb{E}\delta_0 + 2cdM\lambda_0 \stackrel{(63)}{\leq} cd(\eta - 1) + \frac{2cd}{2cd+1} \cdot \frac{\ln 2}{4\eta} < cd(\eta - 1) + 1 \leq cd\eta. \quad (68)$$

Therefore, for $k = 0$, Eq. (67) is satisfied.

Now we consider the index $k + 1 \geq 1$. Because Eq. (66) shows that $M\lambda_k \leq 2$, we can employ Lemma 22, which leads to

$$r_k := \|\mathbf{x}_{k+1} - \mathbf{x}_k\|_{\mathbf{x}_k} \stackrel{(32)}{\leq} \lambda_k. \quad (69)$$

Then by Lemma 25, we have

$$\forall k \geq 0, \nabla^2 f(\mathbf{x}_k) \preceq \mathbf{G}_k \preceq (1 + \delta_k) \nabla^2 f(\mathbf{x}_k) \quad (70)$$

with

$$\mathbb{E}_{\mathbf{u}_k} \delta_{k+1} \stackrel{(43)}{\leq} \left(1 - \frac{1}{t}\right) (1 + Mr_k)^2 (\delta_k + 2cdMr_k) \stackrel{(69)}{\leq} \left(1 - \frac{1}{t}\right) e^{2M\lambda_k} \theta_k. \quad (71)$$

Moreover, using Lemma 22 again, we obtain

$$\lambda_{k+1} \stackrel{(32)(70)}{\leq} \left(1 + \frac{M\lambda_k}{2}\right) \frac{\delta_k + \frac{M\lambda_k}{2}}{1 + \delta_k} \lambda_k \leq \left(1 + \frac{M\lambda_k}{2}\right) \theta_k \lambda_k \leq e^{2M\lambda_k} \theta_k \lambda_k. \quad (72)$$

Note that $\frac{1}{2} \leq 1 - \frac{1}{t}$ because $t \geq d \geq 2$. Combining Eqs. (71) and (72), we obtain

$$\begin{aligned} \mathbb{E}_{\mathbf{u}_k} \theta_{k+1} &\leq \left(1 - \frac{1}{t}\right) e^{2M\lambda_k} \theta_k + 2cdM e^{2M\lambda_k} \theta_k \lambda_k \\ &\leq \left[\left(1 - \frac{1}{t}\right) + \left(1 - \frac{1}{t}\right) 4cdM\lambda_k \right] e^{2M\lambda_k} \theta_k = \left(1 - \frac{1}{t}\right) (1 + 4cdM\lambda_k) e^{2M\lambda_k} \theta_k \\ &\leq \left(1 - \frac{1}{t}\right) e^{2(2cd+1)M\lambda_k} \theta_k \stackrel{(36)}{\leq} \left(1 - \frac{1}{t}\right) e^{2(2cd+1)M\lambda_0 \left(1 - \frac{1}{2\eta}\right)^k} \theta_k. \end{aligned}$$

Therefore, by taking expectation for all randomness, we obtain

$$\begin{aligned}\mathbb{E}\theta_{k+1} &\leq \left(1 - \frac{1}{t}\right) e^{2(2cd+1)M\lambda_0\left(1-\frac{1}{2\eta}\right)^k} \mathbb{E}\theta_k \leq \dots \stackrel{(*)}{\leq} \left(1 - \frac{1}{t}\right)^{k+1} e^{2(2cd+1)M\lambda_0 \sum_{i=0}^k \left(1-\frac{1}{2\eta}\right)^i} \mathbb{E}\theta_0 \\ &\leq \left(1 - \frac{1}{t}\right)^{k+1} e^{4\eta(2cd+1)M\lambda_0} \mathbb{E}\theta_0 \stackrel{(66)(68)}{\leq} 2cd\eta \left(1 - \frac{1}{t}\right)^{k+1}.\end{aligned}\quad (73)$$

Thus, Eq. (67) is proved for the index $k + 1$. Therefore, Eq. (67) holds for all $k \geq 0$.

Now we prove the bound of δ_k and ρ_k based on Eq. (67). Since $\lambda_k \geq 0$, we have

$$\mathbb{E}\delta_k \leq \mathbb{E}\delta_k + 2cdM\lambda_k \stackrel{(67)}{=} \mathbb{E}\theta_k \stackrel{(67)}{\leq} 2cd\eta \left(1 - \frac{1}{t}\right)^k, \forall k \geq 0.$$

Finally, we adopt $\rho_k = e^{2M\lambda_k}\theta_k$ based on Eq. (72) for all $k \geq 0$. Then we get

$$\begin{aligned}\mathbb{E}\rho_k &= \mathbb{E}e^{2M\lambda_k}\theta_k \leq \mathbb{E}e^{2(2cd+1)M\lambda_k}\theta_k \stackrel{(36)}{\leq} e^{2(2cd+1)M\lambda_0\left(1-\frac{1}{2\eta}\right)^k} \mathbb{E}\theta_k \\ &\stackrel{(73)}{\leq} \left(1 - \frac{1}{t}\right)^k e^{2(2cd+1)M\lambda_0 \sum_{i=0}^k \left(1-\frac{1}{2\eta}\right)^i} \theta_0 \stackrel{(66)(68)}{\leq} 2cd\eta \left(1 - \frac{1}{t}\right)^k,\end{aligned}$$

where we use the inequality $(*)$ in Eq. (73) with subscript $k - 1$. Thus, Eqs. (64) and (65) are proved. \blacksquare

D.2 Proofs of Theorem 11 and Theorem 19

Proof The results of greedy methods directly follow Lemma 18, and the proof of randomized methods is the same as the proof of Corollary 9 by applying Lemma 26 with $\mathbf{X}_k = \delta_k$ or $\rho_k, \forall k \geq 0$, so we omit it. \blacksquare

D.3 Proofs of Corollary 12 and Corollary 20

Since the proofs of Corollary 12 and Corollary 20 also have many overlapping parts, we also recombine them into a corollary below.

Corollary 28 (Restating) *Suppose in Algorithm 3 or Algorithm 7, $\mathbf{G}_0 = L\mathbf{I}_d$ and \mathbf{x}_0 satisfies $M\lambda_0 \leq (\ln \frac{3}{2})/(4\mathcal{K})$. Then we could obtain for randomized methods, with probability at least $1 - \delta$ over the random directions $\{\mathbf{u}_k\}$,*

$$\lambda_{k_0+k} \leq \left(1 - \frac{1}{t+1}\right)^{k(k-1)/2} \cdot \left(\frac{1}{2}\right)^k \cdot \left(1 - \frac{1}{2\mathcal{K}}\right)^{k_0} \cdot \lambda_0, \forall k \geq 0,$$

where $k_0 = O((t + \mathcal{K}) \ln(cd\mathcal{K}t/\delta))$, and for greedy methods,

$$\lambda_{k_0+k} \leq \left(1 - \frac{1}{t}\right)^{k(k-1)/2} \cdot \left(\frac{1}{2}\right)^k \cdot \left(1 - \frac{1}{2\mathcal{K}}\right)^{k_0} \cdot \lambda_0, \text{ for all } k \geq 0,$$

where $k_0 = O((t + \mathcal{K}) \ln(cd\mathcal{K}))$. Here, the choices of t, c are inherited from Lemma 25.

Proof We note that $\mathbf{G}_0 = L\mathbf{I}_d$ gives $\nabla^2 f(\mathbf{x}_0) \preceq \mathbf{G}_0 \preceq \varkappa \nabla^2 f(\mathbf{x}_0)$. Since the initial point \mathbf{x}_0 is sufficiently close to the solution: $M\lambda_0 \leq \frac{\ln \frac{3}{2}}{4\varkappa}$, Theorem 23 holds for $\eta = \varkappa$.

Denote by $k_1 \geq 0$ the number of the first iteration, for which

$$\left(1 - \frac{1}{2\varkappa}\right)^{k_1} \leq \frac{2}{3} \cdot \frac{1}{2cd+1}.$$

Clearly, $k_1 \leq 2\varkappa \ln(3cd+2) + 1$. Then from Theorem 23, we obtain

$$\nabla^2 f(\mathbf{x}_{k_1}) \preceq \mathbf{G}_{k_1} \preceq \frac{3\varkappa}{2} \nabla^2 f(\mathbf{x}_{k_1}) \text{ and } M\lambda_{k_1} \stackrel{(36)}{\leq} M \left(1 - \frac{1}{2\varkappa}\right)^{k_1} \lambda_0 \leq \frac{\ln 2}{6\varkappa(2cd+1)},$$

which satisfies the initial condition in Lemma 27 with $\eta = \frac{3}{2}\varkappa$. That is, k_1 is the number of iterations for entering the region of superlinear convergence. Hence, from the (random) initialization \mathbf{G}_{k_1} and \mathbf{x}_{k_1} , by Lemma 27, we have $\forall k \geq k_1, \lambda_{k+1} \leq \rho_k \lambda_k$ with

$$\mathbb{E}\rho_k \leq 3cd\varkappa \left(1 - \frac{1}{t}\right)^{k-k_1}.$$

For randomized method, applying Lemma 26 with $\mathbf{X}_k = \rho_k, \forall k \geq 0$ and $a = 3cd\varkappa$, we can obtain with probability at least $1 - \delta$,

$$\rho_k \leq \frac{3cd\varkappa t^2}{\delta} \left(1 - \frac{1}{t+1}\right)^{k-k_1}, \forall k \geq k_1,$$

which leads to with probability at least $1 - \delta$,

$$\lambda_{k+1} \leq \rho_k \lambda_k \leq \frac{3cd\varkappa t^2}{\delta} \left(1 - \frac{1}{t+1}\right)^{k-k_1} \lambda_k, \forall k \geq k_1. \quad (74)$$

Denote by $k_2 \geq 0$ the number of the first iteration, for which

$$\frac{3cd\varkappa t^2}{\delta} \left(1 - \frac{1}{t+1}\right)^{k_2} \leq \frac{1}{2}.$$

Clearly, $k_2 \leq (t+1) \ln(6cd\varkappa t^2/\delta) + 1$, which is the number of iterations to make the superlinear rate ‘valid’. Applying Eq. (74) only to all $k \geq k_1 + k_2$ (which includes the event in Eq. (74)), we still get with probability at least $1 - \delta$,

$$\lambda_{k_1+k_2+k+1} \leq \frac{3cd\varkappa t^2}{\delta} \left(1 - \frac{1}{t+1}\right)^{k_2+k} \lambda_{k_1+k_2+k} \leq \left(\frac{1}{2}\right) \cdot \left(1 - \frac{1}{t+1}\right)^k \lambda_{k_1+k_2+k}, \forall k \geq 0.$$

Therefore, by the fact of Eq. (61), we also have with probability at least $1 - \delta$,

$$\lambda_{k_1+k_2+k} \leq \left(1 - \frac{1}{t+1}\right)^{k(k-1)/2} \left(\frac{1}{2}\right)^k \lambda_{k_1+k_2}, \forall k \geq 0. \quad (75)$$

Moreover, using Theorem 23 again, we have the deterministic result that

$$\lambda_{k_1+k_2} \stackrel{(36)}{\leq} \left(1 - \frac{1}{2\varkappa}\right)^{k_1+k_2} \lambda_0. \quad (76)$$

Noting that the event of Eq. (75) is contained in the following event based on Eq. (76), we finally obtain with probability at least $1 - \delta$,

$$\lambda_{k_0+k} \leq \left(1 - \frac{1}{t+1}\right)^{k(k-1)/2} \cdot \left(\frac{1}{2}\right)^k \cdot \left(1 - \frac{1}{2\mathcal{X}}\right)^{k_0} \lambda_0, \forall k \geq 0,$$

where $k_0 = k_1 + k_2 \leq 2\mathcal{X} \ln(3cd + 2) + (t+1) \ln(6cd\mathcal{X}t^2/\delta) + 2 = O((t + \mathcal{X}) \ln(cd\mathcal{X}t/\delta))$.

Similarly, we can get the results for greedy methods with $k_2 \leq t \ln(6cd\mathcal{X}) + 1$, leading to $k_0 = k_1 + k_2 \leq 2\mathcal{X} \ln(3cd + 2) + t \ln(6cd\mathcal{X}) + 2 = O((t + \mathcal{X}) \ln(cd\mathcal{X}))$. \blacksquare

Particularly, following the choices of t, c in Lemma 25, for the random Broyden method in Algorithm 3, one has $t = d\mathcal{X}, c = 1$, which leads to $k_0 = O(d\mathcal{X} \ln(d\mathcal{X}/\delta))$; for the greedy/random BFGS method in Algorithm 7, one has $t = d, c = 1$, which leads to $k_0 = O((d + \mathcal{X}) \ln(d\mathcal{X}))$ and $O((d + \mathcal{X}) \ln(d\mathcal{X}/\delta))$; for the greedy/random SR1 method in Algorithm 7, one has $t = d, c = \mathcal{X}$, which leads to $k_0 = O((d + \mathcal{X}) \ln(d\mathcal{X}))$ and $O((d + \mathcal{X}) \ln(d\mathcal{X}/\delta))$.

Appendix E. Proof of Proposition 16

Proof Denote $\mathbf{H}_k := \mathbf{G}_k^{-1}, \forall k \geq 0$. From the inverse update rule in Eq. (27), we obtain

$$\mathbf{H}_{k+1} = \mathbf{Q}_k^\top \mathbf{H}_k \mathbf{Q}_k = (\mathbf{L}_k \mathbf{Q}_k)^\top (\mathbf{L}_k \mathbf{Q}_k), \quad \mathbf{Q}_k := \mathbf{I}_d - \frac{\mathbf{A} \mathbf{u}_k \mathbf{u}_k^\top}{\mathbf{u}_k^\top \mathbf{A} \mathbf{u}_k} + \frac{\mathbf{G}_k \mathbf{u}_k \mathbf{u}_k^\top}{\sqrt{\mathbf{u}_k^\top \mathbf{A} \mathbf{u}_k \cdot \mathbf{u}_k^\top \mathbf{G}_k \mathbf{u}_k}}.$$

Indeed, we have

$$\begin{aligned} & \mathbf{Q}_k^\top \mathbf{H}_k \mathbf{Q}_k - \left(\mathbf{I}_d - \frac{\mathbf{u}_k \mathbf{u}_k^\top \mathbf{A}}{\mathbf{u}_k^\top \mathbf{A} \mathbf{u}_k} \right) \mathbf{H}_k \left(\mathbf{I}_d - \frac{\mathbf{A} \mathbf{u}_k \mathbf{u}_k^\top}{\mathbf{u}_k^\top \mathbf{A} \mathbf{u}_k} \right) \\ &= 2 \left(\mathbf{I}_d - \frac{\mathbf{u}_k \mathbf{u}_k^\top \mathbf{A}}{\mathbf{u}_k^\top \mathbf{A} \mathbf{u}_k} \right) \cdot \frac{\mathbf{H}_k \mathbf{G}_k \mathbf{u}_k \mathbf{u}_k^\top}{\sqrt{\mathbf{u}_k^\top \mathbf{A} \mathbf{u}_k \cdot \mathbf{u}_k^\top \mathbf{G}_k \mathbf{u}_k}} + \frac{\mathbf{u}_k \mathbf{u}_k^\top \mathbf{G}_k \mathbf{H}_k \mathbf{G}_k \mathbf{u}_k \mathbf{u}_k^\top}{\mathbf{u}_k^\top \mathbf{A} \mathbf{u}_k \cdot \mathbf{u}_k^\top \mathbf{G}_k \mathbf{u}_k} \\ &= 2 \left(\mathbf{I}_d - \frac{\mathbf{u}_k \mathbf{u}_k^\top \mathbf{A}}{\mathbf{u}_k^\top \mathbf{A} \mathbf{u}_k} \right) \mathbf{u}_k \cdot \frac{\mathbf{u}_k^\top}{\sqrt{\mathbf{u}_k^\top \mathbf{A} \mathbf{u}_k \cdot \mathbf{u}_k^\top \mathbf{G}_k \mathbf{u}_k}} + \frac{\mathbf{u}_k \mathbf{u}_k^\top}{\mathbf{u}_k^\top \mathbf{A} \mathbf{u}_k} = \frac{\mathbf{u}_k \mathbf{u}_k^\top}{\mathbf{u}_k^\top \mathbf{A} \mathbf{u}_k}, \end{aligned}$$

which is identical to Eq. (27).

Next, note that $\mathbf{H}_k \succ 0$. Thus the square matrix \mathbf{L}_k is also nonsingular, leading to $\mathbf{L}_k \mathbf{G}_k \mathbf{L}_k^\top = \mathbf{L}_k \mathbf{L}_k^{-1} \mathbf{L}_k^{-\top} \mathbf{L}_k^\top = \mathbf{I}_d$. Hence, from $\mathbf{u}_k = \mathbf{L}_k^\top \tilde{\mathbf{u}}_k$, we get $\mathbf{L}_k \mathbf{G}_k \mathbf{u}_k = \mathbf{L}_k \mathbf{G}_k \mathbf{L}_k^\top \tilde{\mathbf{u}}_k = \tilde{\mathbf{u}}_k$, and $\mathbf{u}_k^\top \mathbf{G}_k \mathbf{u}_k = \tilde{\mathbf{u}}_k^\top \mathbf{L}_k \mathbf{G}_k \mathbf{L}_k^\top \tilde{\mathbf{u}}_k = \|\tilde{\mathbf{u}}_k\|^2$. Therefore, by the expression of \mathbf{v}_k in Proposition 16, we obtain

$$\begin{aligned} \mathbf{L}_{k+1} &= \mathbf{L}_k \mathbf{Q}_k = \mathbf{L}_k - \frac{\mathbf{L}_k \mathbf{A} \mathbf{u}_k}{\mathbf{u}_k^\top \mathbf{A} \mathbf{u}_k} \cdot \mathbf{u}_k^\top + \frac{\mathbf{L}_k \mathbf{G}_k \mathbf{u}_k \mathbf{u}_k^\top}{\sqrt{\mathbf{u}_k^\top \mathbf{A} \mathbf{u}_k \cdot \mathbf{u}_k^\top \mathbf{G}_k \mathbf{u}_k}} \\ &= \mathbf{L}_k - \mathbf{L}_k \mathbf{A} \mathbf{u}_k \cdot \frac{\mathbf{u}_k^\top}{\mathbf{u}_k^\top \mathbf{A} \mathbf{u}_k} + \frac{\tilde{\mathbf{u}}_k}{\|\tilde{\mathbf{u}}_k\|} \cdot \sqrt{\mathbf{u}_k^\top \mathbf{A} \mathbf{u}_k} \cdot \frac{\mathbf{u}_k^\top}{\mathbf{u}_k^\top \mathbf{A} \mathbf{u}_k} = \mathbf{L}_k - \frac{(\mathbf{L}_k \mathbf{A} \mathbf{u}_k - \mathbf{v}_k) \mathbf{u}_k^\top}{\mathbf{u}_k^\top \mathbf{A} \mathbf{u}_k}. \end{aligned} \quad \blacksquare$$

References

- Antoine Bordes, Léon Bottou, and Patrick Gallinari. SGD-QN: Careful Quasi-Newton Stochastic Gradient Descent. *Journal of Machine Learning Research*, 10(59):1737–1754, 2009.
- Léon Bottou and Yann Le Cun. On-line learning for very large data sets. *Applied stochastic models in business and industry*, 21(2):137–151, 2005.
- Charles G Broyden. Quasi-Newton methods and their application to function minimisation. *Mathematics of Computation*, 21(99):368–381, 1967.
- Charles G Broyden. The convergence of a class of double-rank minimization algorithms: 2. the new algorithm. *IMA journal of applied mathematics*, 6(3):222–231, 1970a.
- Charles George Broyden. The convergence of a class of double-rank minimization algorithms 1. general considerations. *IMA Journal of Applied Mathematics*, 6(1):76–90, 1970b.
- Charles George Broyden, John E Dennis Jr, and Jorge J Moré. On the local and superlinear convergence of quasi-Newton methods. *IMA Journal of Applied Mathematics*, 12(3): 223–245, 1973.
- Richard H Byrd, Jorge Nocedal, and Ya-Xiang Yuan. Global convergence of a class of quasi-Newton methods on convex problems. *SIAM Journal on Numerical Analysis*, 24(5): 1171–1190, 1987.
- Chih-Chung Chang and Chih-Jen Lin. LIBSVM: a library for support vector machines. *ACM transactions on intelligent systems and technology (TIST)*, 2(3):1–27, 2011.
- William C Davidon. Variable metric method for minimization. *SIAM Journal on Optimization*, 1(1):1–17, 1991.
- John E Dennis and Jorge J Moré. A characterization of superlinear convergence and its application to quasi-Newton methods. *Mathematics of computation*, 28(126):549–560, 1974.
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7), 2011.
- Roger Fletcher. A new approach to variable metric algorithms. *The computer journal*, 13(3): 317–322, 1970.
- Roger Fletcher and Michael JD Powell. A rapidly convergent descent method for minimization. *The computer journal*, 6(2):163–168, 1963.
- Donald Goldfarb. A family of variable-metric methods derived by variational means. *Mathematics of computation*, 24(109):23–26, 1970.
- Andreas Griewank and Ph L Toint. Local convergence analysis for partitioned quasi-Newton updates. *Numerische Mathematik*, 39(3):429–448, 1982.

- Qiujiang Jin and Aryan Mokhtari. Non-asymptotic superlinear convergence of standard quasi-Newton methods. *arXiv preprint arXiv:2003.13607*, 2020.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR (Poster)*, 2015.
- Dmitry Kovalev, Robert M Gower, Peter Richtárik, and Alexander Rogozin. Fast linear convergence of randomized BFGS. *arXiv preprint arXiv:2002.11337*, 2020.
- Dong C Liu and Jorge Nocedal. On the limited memory BFGS method for large scale optimization. *Mathematical programming*, 45(1):503–528, 1989.
- Sifan Liu and Art B. Owen. Quasi-Monte Carlo Quasi-Newton in Variational Bayes. *Journal of Machine Learning Research*, 22(243):1–23, 2021.
- Aryan Mokhtari and Alejandro Ribeiro. A quasi-Newton method for large scale support vector machines. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8302–8306. IEEE, 2014.
- Aryan Mokhtari and Alejandro Ribeiro. Global convergence of online limited memory BFGS. *Journal of Machine Learning Research*, 16(1):3151–3181, 2015.
- Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2003.
- Jorge Nocedal and Stephen Wright. *Numerical optimization*. Springer Science & Business Media, 2006.
- MJD Powell. On the convergence of the variable metric algorithm. *IMA Journal of Applied Mathematics*, 7(1):21–36, 1971.
- Anton Rodomanov and Yurii Nesterov. New results on superlinear convergence of classical quasi-Newton methods. *Journal of Optimization Theory and Applications*, 188(3):744–769, Jan 2021a. ISSN 1573-2878. doi: 10.1007/s10957-020-01805-8. URL <http://dx.doi.org/10.1007/s10957-020-01805-8>.
- Anton Rodomanov and Yurii Nesterov. Greedy quasi-Newton methods with explicit superlinear convergence. *SIAM Journal on Optimization*, 31(1):785–811, 2021b.
- Anton Rodomanov and Yurii Nesterov. Rates of superlinear convergence for classical quasi-Newton methods. *Mathematical Programming*, pages 1–32, 2021c.
- Shai Shalev-Shwartz and Nathan Srebro. SVM optimization: inverse dependence on training set size. In *Proceedings of the 25th International Conference on Machine Learning*, pages 928–935, 2008.
- David F Shanno. Conditioning of quasi-Newton methods for function minimization. *Mathematics of computation*, 24(111):647–656, 1970.
- Andrzej Stachurski. Superlinear convergence of Broyden’s bounded θ -class of methods. *Mathematical Programming*, 20(1):196–212, 1981.

- Hiroshi Yabe and Naokazu Yamaki. Local and superlinear convergence of structured quasi-Newton methods for nonlinear optimization. *Journal of the Operations Research Society of Japan*, 39(4):541–557, 1996.
- Haishan Ye, Luo Luo, and Zhihua Zhang. Nesterov’s Acceleration for Approximate Newton. *Journal of Machine Learning Research*, 21(142):1–37, 2020. URL <http://jmlr.org/papers/v21/19-265.html>.
- Jin Yu, S.V.N. Vishwanathan, Simon Günter, and Nicol N. Schraudolph. A Quasi-Newton Approach to Nonsmooth Convex Optimization Problems in Machine Learning. *Journal of Machine Learning Research*, 11(39):1145–1200, 2010.
- Xiao-Tong Yuan and Ping Li. On Convergence of Distributed Approximate Newton Methods: Globalization, Sharper Bounds and Beyond. *Journal of Machine Learning Research*, 21: 206–1, 2020.