

IQN: AN INCREMENTAL QUASI-NEWTON METHOD WITH LOCAL SUPERLINEAR CONVERGENCE RATE*

ARYAN MOKHTARI[†], MARK EISEN[†], AND ALEJANDRO RIBEIRO[†]

Abstract. The problem of minimizing an objective that can be written as the sum of a set of n smooth and strongly convex functions is challenging because the cost of evaluating the function and its derivatives is proportional to the number of elements in the sum. The incremental quasi-Newton (IQN) method proposed here belongs to the family of stochastic and incremental methods that have a cost per iteration independent of n . IQN iterations are a stochastic version of BFGS iterations that use memory to reduce the variance of stochastic approximations. The method is shown to exhibit local superlinear convergence. The convergence properties of IQN bridge a gap between deterministic and stochastic quasi-Newton methods. Deterministic quasi-Newton methods exploit the possibility of approximating the Newton step using objective gradient differences. They are appealing because they have a smaller computational cost per iteration relative to Newton’s method and achieve a superlinear convergence rate under customary regularity assumptions. Stochastic quasi-Newton methods utilize stochastic gradient differences in lieu of actual gradient differences. This makes their computational cost per iteration independent of the number of objective functions n . However, existing stochastic quasi-Newton methods have sublinear or linear convergence at best. IQN is the first stochastic quasi-Newton method proven to converge superlinearly in a local neighborhood of the optimal solution. IQN differs from state-of-the-art incremental quasi-Newton methods in three aspects: (i) The use of aggregated information of variables, gradients, and quasi-Newton Hessian approximation matrices to reduce the noise of gradient and Hessian approximations. (ii) The approximation of each individual function by its Taylor’s expansion, in which the linear and quadratic terms are evaluated with respect to the same iterate. (iii) The use of a cyclic scheme to update the functions in lieu of a random selection routine. We use these fundamental properties of IQN to establish its local superlinear convergence rate. The presented numerical experiments match our theoretical results and justify the advantage of IQN relative to other incremental methods.

Key words. large-scale optimization, stochastic optimization, quasi-Newton methods, incremental methods, superlinear convergence

AMS subject classifications. 90C06, 90C25, 90C30, 90C52

DOI. 10.1137/17M1122943

1. Introduction. We study large-scale optimization problems with objective functions expressed as the sum of a set of components which arise often in application domains such as machine learning [3, 2, 34, 8], control [5, 7, 18], and wireless communications [31, 27, 28]. Formally, we consider a variable $\mathbf{x} \in \mathbb{R}^p$ and a function f which is defined as the average of n smooth and strongly convex functions labelled $f_i : \mathbb{R}^p \rightarrow \mathbb{R}$ for $i = 1, \dots, n$. We refer to individual functions f_i as sample functions and to the total number of functions n as the sample size. Our goal is to find the optimal argument \mathbf{x}^* that solves the strongly convex program

$$(1) \quad \mathbf{x}^* := \operatorname{argmin}_{\mathbf{x} \in \mathbb{R}^p} f(\mathbf{x}) := \operatorname{argmin}_{\mathbf{x} \in \mathbb{R}^p} \frac{1}{n} \sum_{i=1}^n f_i(\mathbf{x}).$$

We restrict attention to cases in which the component functions f_i are strongly convex and their gradients are Lipschitz continuous. We further focus on problems in

*Received by the editors March 28, 2017; accepted for publication (in revised form) January 31, 2018; published electronically May 22, 2018.

<http://www.siam.org/journals/siopt/28-2/M112294.html>

Funding: This work was supported by ONR N00014-12-1-0997.

[†]Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 (aryanm@seas.upenn.edu, maeisen@seas.upenn.edu, aribeiro@seas.upenn.edu).

which n is large enough so as to warrant application of stochastic or iterative methods. Our goal is to propose an iterative quasi-Newton method to solve (1) which is shown to exhibit a local superlinear convergence rate. This is achieved while performing local iterations with a cost of order $\mathcal{O}(p^2)$ independent of the number of samples n .

Temporarily setting aside the complications related to the number of component functions, the minimization of f in (1) can be carried out using iterative descent algorithms. A simple solution is to use gradient descent (GD) which iteratively descends along gradient directions $\nabla f(\mathbf{x}) = (1/n) \sum_{i=1}^n \nabla f_i(\mathbf{x})$. GD incurs a per iteration computational cost of order $\mathcal{O}(np)$ and is known to converge at a linear rate towards \mathbf{x}^* under the hypotheses we have placed on f . Whether the linear convergence rate of GD is acceptable depends on the desired accuracy and on the condition number of f , which, when large, can make the convergence constant close to one. As one or both of these properties often limit the applicability of GD, classical alternatives to improve convergence rates have been developed. Newton's method adapts to the curvature of the objective by computing Hessian inverses and converges at a quadratic rate in a local neighborhood of the optimal argument irrespective of the problem's condition number. To achieve this quadratic convergence rate, we must evaluate and invert Hessians resulting in a per iteration cost of order $\mathcal{O}(np^2 + p^3)$. Quasi-Newton methods build on the idea of approximating the Newton step using first-order information of the objective function and exhibit local superlinear convergence [4, 26, 13]. An important feature of quasi-Newton methods is that they have a per iteration cost of order $\mathcal{O}(np + p^2)$, where the term $\mathcal{O}(np)$ corresponds to the cost of gradient computation and the cost $\mathcal{O}(p^2)$ indicates the computational complexity of updating the approximate Hessian inverse matrix.

The combination of a local superlinear convergence rate and the smaller computational cost per iteration relative to Newton—a reduction by a factor of p operations per iteration—make quasi-Newton methods an appealing choice. In the context of optimization problems having the form in (1), quasi-Newton methods also have the advantage that curvature is estimated using gradient evaluations. To see why this is meaningful we must recall that the customary approach to avoiding the $\mathcal{O}(np)$ computational cost of GD iterations is to replace gradients $\nabla f(\mathbf{x})$ by their stochastic approximations $\nabla f_i(\mathbf{x})$, which can be evaluated with a cost of order $\mathcal{O}(p)$. One can then think of using stochastic versions of these gradients to develop stochastic quasi-Newton methods with per iteration cost of order $\mathcal{O}(p + p^2)$. This idea was demonstrated to be feasible in [33], which introduces a stochastic (online) version of the BFGS quasi-Newton method as well as a stochastic version of its limited memory variant. Although [33] provides numerical experiments illustrating significant improvements in convergence times relative to stochastic GD (SGD), theoretical guarantees are not established.

The issue of proving convergence of stochastic quasi-Newton methods is tackled in [22] and [23]. In [22] the authors show that stochastic BFGS may not be convergent because the Hessian approximation matrices can become close to singular. A regularized stochastic BFGS (RES) method is proposed by changing the proximity condition of BFGS to ensure that the eigenvalues of the Hessian inverse approximation are uniformly bounded. Enforcing this property yields a provably convergent algorithm. In [23] the authors show that the limited memory version of stochastic (online) BFGS proposed in [33] is almost surely convergent and has a sublinear convergence rate in expectation. This is achieved without using regularizations. An alternative provably convergent stochastic quasi-Newton method is proposed in [6]. This method differs

from those in [33, 22, 23] in that it collects (stochastic) second-order information to estimate the objective's curvature. This is in contrast to estimating curvature using the difference of two consecutive stochastic gradients. Similarly, stochastic quasi-Newton methods for non-convex problems are developed in [15, 35].

Although the methods in [33, 22, 23, 6] are successful in expanding the application of quasi-Newton methods to stochastic settings, their convergence rate is sublinear. This is not better than the convergence rate of SGD and, as is also the case in SGD, is a consequence of the stochastic approximation noise, which necessitates the use of diminishing stepsizes. The stochastic quasi-Newton methods in [19, 24] resolve this issue by using the variance reduction technique proposed in [14]. The fundamental idea of the work in [14] is to reduce the noise of the stochastic gradient approximation by computing the exact gradient in an outer loop to use it in an inner loop for gradient approximation. The methods in [19, 24], which incorporate the variance reduction scheme presented in [14] into the update of quasi-Newton methods, are successful in achieving a linear convergence rate.

At this point, we must remark on an interesting mismatch. The convergence rate of SGD is sublinear, and the convergence rate of deterministic GD is linear. The use of variance reduction techniques in SGD recovers the linear convergence rate of GD [14]. On the other hand, the convergence rate of stochastic quasi-Newton methods is sublinear, and the convergence rate of deterministic quasi-Newton methods is superlinear. The use of variance reduction in stochastic quasi-Newton methods achieves linear convergence but does not recover a superlinear rate. Hence, a fundamental question remains unanswered: Is it possible to design an incremental quasi-Newton (IQN) method that recovers the superlinear convergence rate of deterministic quasi-Newton algorithms? In this paper, we show that the answer to this open problem is positive by proposing an IQN method with a local superlinear convergence rate. This is the first quasi-Newton method to achieve superlinear convergence while having a per iteration cost independent of the number of functions n —the cost per iteration is of order $\mathcal{O}(p^2)$.

There are three major differences between the IQN method and state-of-the-art incremental (stochastic) quasi-Newton methods that lead to the former's superlinear convergence rate. First, the proposed IQN method uses the aggregated information of variables, gradients, and Hessian approximation matrices to reduce the noise of approximation for both gradients and Hessian approximation matrices. This is different to the variance-reduced stochastic quasi-Newton methods in [19, 24] that attempt to reduce only the noise of gradient approximations. Second, in IQN the index of the updated function is chosen in a cyclic fashion, rather than the random selection scheme used in the incremental methods in [33, 22, 23, 6]. The cyclic routine in IQN allows one to bound the error at each iteration as a function of the errors of the last n iterates, something that is not possible when using a random scheme. To explain the third and most important difference we point out that the form of quasi-Newton updates is the solution of a local second-order Taylor approximation of the objective. It is possible to understand stochastic quasi-Newton methods as an analogous approximation of individual sample functions. However, it turns out that the state-of-the-art stochastic quasi-methods evaluate the linear and quadratic terms of the Taylor's expansion at different points, yielding an inconsistent approximation (Remark 7). The IQN method utilizes a consistent Taylor series which yields a more involved update which we nonetheless show can be implemented with the same computational cost. These three properties together lead to an IQN method with a local superlinear convergence rate.

We start the paper by recapping the BFGS quasi-Newton method and the Dennis–Moré condition, which is sufficient and necessary to prove the superlinear convergence rate of the BFGS method (section 3). Then, we present the proposed IQN as an incremental aggregated version of the traditional BFGS method (section 4). We first explain the difference between the Taylor’s expansion used in IQN and state-of-the-art incremental (stochastic) quasi-Newton methods. Further, we explain the mechanism for aggregation of the functions’ information and the scheme for updating the stored information. Moreover, we present an efficient implementation of the proposed IQN method with a computational complexity of the order $\mathcal{O}(p^2)$ (section 4.1). The convergence analysis of the IQN method is then presented (section 5). We use the classic analysis of quasi-Newton methods to show that in a local neighborhood of the optimal solution the sequence of variables converges to the optimal argument \mathbf{x}^* linearly after each pass over the set of functions (Lemma 3). We use this result to show that for each component function f_i the Dennis–Moré condition holds (Proposition 4). However, this condition is not sufficient to prove superlinear convergence of the sequence of errors $\|\mathbf{x}^t - \mathbf{x}^*\|$, since it does not guarantee the Dennis–Moré condition for the global objective f . To overcome this issue we introduce a novel convergence analysis approach which exploits the local linear convergence of IQN to present a more general version of the Dennis–Moré condition for each component function f_i (Lemma 5). We exploit this result to establish superlinear convergence of the iterates generated by IQN (Theorem 6). In section 6, we present numerical simulation results, comparing the performance of IQN to that of first-order incremental and stochastic methods. We test the performance on a set of large-scale regression problems and observe strong numerical gain in total computation time relative to existing methods.

1.1. Notation. Vectors are written lowercase, $\mathbf{x} \in \mathbb{R}^p$, and matrices uppercase, $\mathbf{A} \in \mathbb{R}^{p \times p}$. We use $\|\mathbf{x}\|$ and $\|\mathbf{A}\|$ to denote the Euclidean norm of vector \mathbf{x} and matrix \mathbf{A} , respectively. Given a positive definite matrix \mathbf{M}_i , the weighted matrix norm $\|\mathbf{A}\|_{\mathbf{M}_i}$ is defined as $\|\mathbf{A}\|_{\mathbf{M}_i} := \|\mathbf{M}_i \mathbf{A} \mathbf{M}_i\|_{\mathbf{F}}$, where $\|\cdot\|_{\mathbf{F}}$ is the Frobenius norm. Given a function f , its gradient and Hessian at point \mathbf{x} are denoted by $\nabla f(\mathbf{x})$ and $\nabla^2 f(\mathbf{x})$, respectively.

2. Related works. Various methods have been studied in the literature to improve the performance of traditional full-batch optimization algorithms. The most famous method for reducing the computational complexity of gradient descent (GD) is stochastic gradient descent (SGD), which uses the gradient of a single randomly chosen function to approximate the full gradient [2]. Incremental gradient descent method (IGD) is similar to SGD except the function is chosen in a cyclic routine [1]. Both SGD and IGD suffer from slow sublinear convergence rate because of the noise of gradient approximation. The incremental aggregated methods, which use memory to aggregate the gradients of all n functions, are successful in reducing the noise of gradient approximation to achieve a linear convergence rate [30, 32, 9, 14]. The work in [30] suggests a random selection of functions, which leads to stochastic average gradient method (SAG), while the works in [1, 12, 21] use a cyclic scheme.

Moving beyond first-order information, there have been stochastic quasi-Newton methods to approximate Hessian information for convex [33, 22, 23, 6, 24, 11] and nonconvex problems [15, 35]. All of these stochastic quasi-Newton methods reduce the computational cost of quasi-Newton methods by updating only a randomly chosen single or small subset of gradients at each iteration. However, they are not able to recover the superlinear convergence rate of quasi-Newton methods [4, 26, 13]. The incremental Newton method (NIM) in [29] is the only incremental method shown

to have a superlinear convergence rate; however, the Hessian function is not always available or computationally feasible. Moreover, the implementation of NIM requires computation of the incremental aggregated Hessian inverse, which has computational complexity of the order $\mathcal{O}(p^3)$.

3. BFGS quasi-Newton method. Consider the problem in (1) for relatively large n . In a conventional optimization setting, this can be solved using a quasi-Newton method that iteratively updates a variable \mathbf{x}^t for $t = 0, 1, \dots$ based on the general recursive expression

$$(2) \quad \mathbf{x}^{t+1} = \mathbf{x}^t - \eta^t (\mathbf{B}^t)^{-1} \nabla f(\mathbf{x}^t),$$

where η^t is a scalar stepsize and \mathbf{B}^t is a positive definite matrix that approximates the exact Hessian of the objective function $\nabla^2 f(\mathbf{x}^t)$. The stepsize η^t is evaluated based on a line search routine for the global convergence of quasi-Newton methods. Our focus in this paper, however, is on the local convergence of quasi-Newton methods, which requires the unit stepsize $\eta^t = 1$. Therefore, throughout the paper we assume that the variable \mathbf{x}^t is close to the optimal solution \mathbf{x}^* —we will formalize the notion of being close to the optimal solution—and the stepsize is $\eta^t = 1$.

The goal of quasi-Newton methods is to compute the Hessian approximation matrix \mathbf{B}^t and its inverse $(\mathbf{B}^t)^{-1}$ by using only the first-order information, i.e., gradients, of the objective. Their use is widespread due to the many applications in which the Hessian information required in Newton's method is either unavailable or computationally intensive. There are various approaches to approximate the Hessian, but the common feature among quasi-Newton methods is that the Hessian approximation must satisfy the secant condition. To be more precise, consider \mathbf{s}^t and \mathbf{y}^t as the variable and gradient variations, explicitly defined as

$$(3) \quad \mathbf{s}^t := \mathbf{x}^{t+1} - \mathbf{x}^t, \quad \mathbf{y}^t := \nabla f(\mathbf{x}^{t+1}) - \nabla f(\mathbf{x}^t).$$

Then, given the variable variation \mathbf{s}^t and gradient variation \mathbf{y}^t , the Hessian approximation matrix in all quasi-Newton methods must satisfy the secant condition

$$(4) \quad \mathbf{B}^{t+1} \mathbf{s}^t = \mathbf{y}^t.$$

This condition is fundamental in quasi-Newton methods because the exact Hessian $\nabla^2 f(\mathbf{x}^t)$ satisfies this equality when the iterates \mathbf{x}^{t+1} and \mathbf{x}^t are close to each other. If we consider the matrix \mathbf{B}^{t+1} as the unknown matrix, the system of equations in (4) does not have a unique solution. Different quasi-Newton methods enforce different conditions on the matrix \mathbf{B}^{t+1} to come up with a unique update. This extra condition is typically a proximity condition that ensures that \mathbf{B}^{t+1} is close to the previous Hessian approximation matrix \mathbf{B}^t [4, 26, 13]. In particular, the Broyden–Fletcher–Goldfarb–Shanno (BFGS) method defines the update of Hessian approximation matrix as

$$(5) \quad \mathbf{B}^{t+1} = \mathbf{B}^t + \frac{\mathbf{y}^t \mathbf{y}^{tT}}{\mathbf{y}^{tT} \mathbf{s}^t} - \frac{\mathbf{B}^t \mathbf{s}^t \mathbf{s}^{tT} \mathbf{B}^t}{\mathbf{s}^{tT} \mathbf{B}^t \mathbf{s}^t}.$$

The BFGS method is popular not only for its strong numerical performance relative to the gradient descent method, but also because it is shown to exhibit a superlinear convergence rate [4], thereby providing a theoretical guarantee of superior performance. In fact, it can be shown that the BFGS update satisfies the condition

$$(6) \quad \lim_{t \rightarrow \infty} \frac{\|(\mathbf{B}^t - \nabla^2 f(\mathbf{x}^*)) \mathbf{s}^t\|}{\|\mathbf{s}^t\|} = 0,$$

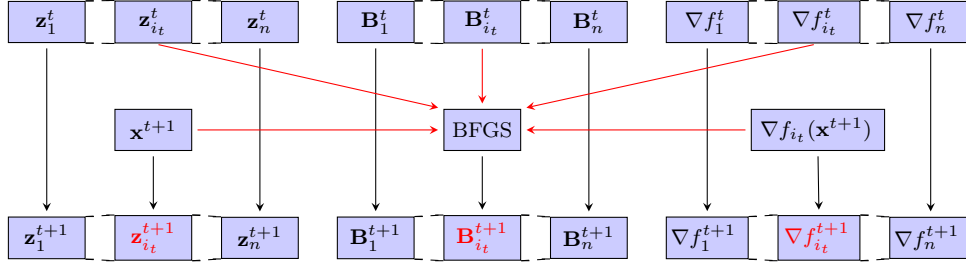


FIG. 1. The updating scheme for variables, gradients, and Hessian approximation matrices of function f_{i_t} at step t . The red arrows indicate the terms used in the update of $\mathbf{B}_{i_t}^{t+1}$ using the BFGS update in (15). The black arrows show the updates of all variables and gradients. The terms $\mathbf{z}_{i_t}^{t+1}$ and $\nabla f_{i_t}^{t+1}$ are updated as \mathbf{x}^{t+1} and $\nabla f_{i_t}(\mathbf{x}^{t+1})$, respectively. All other \mathbf{z}_j^{t+1} and ∇f_j^{t+1} are set as \mathbf{z}_j^t and ∇f_j^t , respectively.

known as the Dennis–Moré condition, which is both necessary and sufficient for superlinear convergence [13]. This result solidifies quasi-Newton methods as a strong alternative to first-order methods when exact second-order information is unavailable. However, implementation of the BFGS method is not feasible when the number of functions n is large, due to its high computational complexity of the order $\mathcal{O}(np + p^2)$. In the following section, we propose a novel incremental BFGS method that has the computational complexity of $\mathcal{O}(p^2)$ per iteration and converges at a superlinear rate.

4. IQN: Incremental aggregated BFGS. We propose an incremental aggregated BFGS algorithm, which we call the incremental quasi-Newton (IQN) method. The IQN method is incremental in that, at each iteration, only the information associated with a single function f_i is updated. The particular function is chosen by cyclicly iterating through the n functions. The IQN method is aggregated in that the aggregate of the most recently observed information of all functions f_1, \dots, f_n is used to compute the updated variable \mathbf{x}^{t+1} .

In the proposed method, we consider $\mathbf{z}_1^t, \dots, \mathbf{z}_n^t$ as the copies of the variable \mathbf{x} at time t associated with the functions f_1, \dots, f_n , respectively. Likewise, define $\nabla f_i(\mathbf{z}_i^t)$ as the gradient corresponding to the i th function. Further, consider \mathbf{B}_i^t as a positive definite matrix which approximates the i th component Hessian $\nabla^2 f_i(\mathbf{x}^t)$. We refer to $\mathbf{z}_i^t, \nabla f_i(\mathbf{z}_i^t)$, and \mathbf{B}_i^t as the information corresponding to the i th function f_i at step t . Note that the functions' information is stored in a shared memory, as shown in Figure 1 (see online version for color figures). To introduce the IQN method, we first explain the mechanism for computing the updated variable \mathbf{x}^{t+1} using the stored information $\{\mathbf{z}_i^t, \nabla f_i(\mathbf{z}_i^t), \mathbf{B}_i^t\}_{i=1}^n$. Then, we elaborate on the scheme for updating the information of the functions.

To derive the full variable update, consider the second-order approximation of the objective function $f_i(\mathbf{x})$ centered around its current iterate \mathbf{z}_i^t ,

$$(7) \quad f_i(\mathbf{x}) \approx f_i(\mathbf{z}_i^t) + \nabla f_i(\mathbf{z}_i^t)^T (\mathbf{x} - \mathbf{z}_i^t) + \frac{1}{2} (\mathbf{x} - \mathbf{z}_i^t)^T \nabla^2 f_i(\mathbf{z}_i^t) (\mathbf{x} - \mathbf{z}_i^t).$$

As in traditional quasi-Newton methods, we replace the i th Hessian $\nabla^2 f_i(\mathbf{z}_i^t)$ by \mathbf{B}_i^t . Using the approximation matrices in place of Hessians, the complete (aggregate) function $f(\mathbf{x})$ can be approximated with

$$(8) \quad f(\mathbf{x}) \approx \frac{1}{n} \sum_{i=1}^n \left[f_i(\mathbf{z}_i^t) + \nabla f_i(\mathbf{z}_i^t)^T (\mathbf{x} - \mathbf{z}_i^t) + \frac{1}{2} (\mathbf{x} - \mathbf{z}_i^t)^T \mathbf{B}_i^t (\mathbf{x} - \mathbf{z}_i^t) \right].$$

Note that the right-hand side of (8) is a quadratic approximation of the function f based on the available information at step t . Hence, the updated iterate \mathbf{x}^{t+1} can be defined as the minimizer of the quadratic program in (8), explicitly given by

$$(9) \quad \mathbf{x}^{t+1} = \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i^t \right)^{-1} \left[\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i^t \mathbf{z}_i^t - \frac{1}{n} \sum_{i=1}^n \nabla f_i(\mathbf{z}_i^t) \right].$$

First note that the update in (9) shows that the updated variable \mathbf{x}^{t+1} is a function of the stored information of all functions f_1, \dots, f_n . Furthermore, we use the aggregated information of variables, gradients, and the quasi-Newton Hessian approximations to evaluate the updated variable. This is done to vanish the noise in approximating both gradients and Hessians as the sequence approaches the optimal argument.

Remark 1. Given the BFGS Hessian approximation matrices $\{\mathbf{B}_i^t\}_{i=1}^n$ and gradients $\{\nabla f_i(\mathbf{z}_i^t)\}_{i=1}^n$, one may consider an update more akin to traditional descent-based methods, i.e.,

$$(10) \quad \mathbf{x}^{t+1} = \mathbf{x}^t - \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i^t \right)^{-1} \frac{1}{n} \sum_{i=1}^n \nabla f_i(\mathbf{z}_i^t).$$

To evaluate the advantage of the proposed update for IQN in (9) relative to the update in (10), we proceed to study the Taylor expansion that leads to the update in (10). It can be shown that the update in (10) is the outcome of the following approximation:

$$(11) \quad f(\mathbf{x}) \approx \frac{1}{n} \sum_{i=1}^n \left[f_i(\mathbf{z}_i^t) + \nabla f_i(\mathbf{z}_i^t)^T (\mathbf{x} - \mathbf{z}_i^t) + \frac{1}{2} (\mathbf{x} - \mathbf{x}^t)^T \mathbf{B}_i^t (\mathbf{x} - \mathbf{x}^t) \right].$$

Observe that the linear term in (11) is centered at \mathbf{z}_i^t , while the quadratic term is approximated near the iterate \mathbf{x}^t . Therefore, there is an inconsistency between the linear and quadratic terms in (11), whereas the linear and quadratic terms suggested in (8) are consistent and both centered at the same point \mathbf{z}_i^t .

Thus far we have discussed the procedure to compute the updated variable \mathbf{x}^{t+1} given the local iterates, gradients, and Hessian approximations at time t . Now, it remains to show how we update the local information of functions f_1, \dots, f_n using the variable \mathbf{x}^{t+1} . In each iteration of the IQN method, we update the local information of only a single function, chosen in a cyclic manner. Defining i_t to be the index of the function selected at time t , we update the local variables $\mathbf{z}_{i_t}^{t+1}$, $\nabla f_{i_t}(\mathbf{z}_{i_t}^{t+1})$, and $\mathbf{B}_{i_t}^{t+1}$ using the updated variable \mathbf{x}^{t+1} while all other local variables remain unchanged. In particular, the variables \mathbf{z}_i are updated as

$$(12) \quad \mathbf{z}_{i_t}^{t+1} = \mathbf{x}^{t+1}, \quad \mathbf{z}_i^{t+1} = \mathbf{z}_i^t \quad \text{for all } i \neq i_t.$$

Observe in the update in (12) that the variable associated with the function f_{i_t} is set to be the updated variable \mathbf{x}^{t+1} while the other iterates are simply kept as their previous value. Likewise, we update the table of gradients accordingly with the gradient of f_{i_t} evaluated at the new variable \mathbf{x}^{t+1} . The rest of the gradients stored in the memory will stay unchanged as in [9, 10, 20, 21], i.e.,

$$(13) \quad \nabla f_{i_t}(\mathbf{z}_{i_t}^{t+1}) = \nabla f_{i_t}(\mathbf{x}^{t+1}), \quad \nabla f_i(\mathbf{z}_i^{t+1}) = \nabla f_i(\mathbf{z}_i^t) \quad \text{for all } i \neq i_t.$$

To update the curvature information, it would be ideal to compute the Hessian $\nabla^2 f_{i_t}(\mathbf{x}^{t+1})$ and update the curvature information following the schemes for variables in (12) and gradients in (13). However, our focus is on applications for which the computation of the Hessian is either impossible or computationally expensive. Hence, to the update curvature approximation matrix $\mathbf{B}_{i_t}^t$ corresponding to the function f_{i_t} , we use the steps of BFGS in (5). To do so, we define variable and gradient variations associated with each individual function f_i as

$$(14) \quad \mathbf{s}_i^t := \mathbf{z}_i^{t+1} - \mathbf{z}_i^t, \quad \mathbf{y}_i^t := \nabla f_i(\mathbf{z}_i^{t+1}) - \nabla f_i(\mathbf{z}_i^t),$$

respectively. The Hessian approximation $\mathbf{B}_{i_t}^t$ corresponding to the function f_{i_t} can be computed using the update of BFGS as

$$(15) \quad \mathbf{B}_i^{t+1} = \mathbf{B}_i^t + \frac{\mathbf{y}_i^t \mathbf{y}_i^{tT}}{\mathbf{y}_i^{tT} \mathbf{s}_i^t} - \frac{\mathbf{B}_i^t \mathbf{s}_i^t \mathbf{s}_i^{tT} \mathbf{B}_i^t}{\mathbf{s}_i^{tT} \mathbf{B}_i^t \mathbf{s}_i^t} \quad \text{for } i = i_t.$$

Again, the Hessian approximation matrices for all other functions remain unchanged, i.e., $\mathbf{B}_i^{t+1} = \mathbf{B}_i^t$ for $i \neq i_t$. The system of updates in (12)–(15) explains the mechanism of updating the information of the function f_{i_t} at step t . Notice that to update the Hessian approximation matrix for the i_t th function there is no need to store the variations in (14), since the old variables \mathbf{z}_i^t and $\nabla f_i(\mathbf{z}_i^t)$ are available in memory and the updated versions $\mathbf{z}_i^{t+1} = \mathbf{x}^{t+1}$ and $\nabla f_i(\mathbf{z}_i^{t+1}) = \nabla f_i(\mathbf{x}^{t+1})$ are evaluated at step t ; see Figure 1 for more details.

Because of the cyclic update scheme, the set of iterates $\{\mathbf{z}_1^t, \mathbf{z}_2^t, \dots, \mathbf{z}_n^t\}$ is equal to the set $\{\mathbf{x}^t, \mathbf{x}^{t-1}, \dots, \mathbf{x}^{t-n+1}\}$, and, therefore, the set of variables used in the update of IQN is the set of the last n iterates. The update of IQN in (9) incorporates the information of all the functions f_1, \dots, f_n to compute the updated variable \mathbf{x}^{t+1} ; however, it uses delayed variables, gradients, and Hessian approximations rather than the the updated variable \mathbf{x}^{t+1} for all functions as in classic quasi-Newton methods. The use of delay allows IQN to update the information of a single function at each iteration, thus reducing the computational complexity relative to classic quasi-Newton methods.

Although the update in (9) is helpful in understanding the rationale behind the IQN method, it cannot be implemented at a low computational cost, since it requires computation of the sums $\sum_{i=1}^n \mathbf{B}_i^t$, $\sum_{i=1}^n \mathbf{B}_i^t \mathbf{z}_i^t$, and $\sum_{i=1}^n \nabla f_i(\mathbf{z}_i^t)$ as well as computing the inversion $(\sum_{i=1}^n \mathbf{B}_i^t)^{-1}$. In the following section, we introduce an efficient implementation of the IQN method that has the computational complexity of $\mathcal{O}(p^2)$.

4.1. Efficient implementation of IQN. To see that the updating scheme in (9) requires evaluation of only a single gradient and Hessian approximation matrix per iteration, consider writing the update as

$$(16) \quad \mathbf{x}^{t+1} = (\tilde{\mathbf{B}}^t)^{-1} (\mathbf{u}^t - \mathbf{g}^t),$$

where we define $\tilde{\mathbf{B}}^t := \sum_{i=1}^n \mathbf{B}_i^t$ as the aggregate Hessian approximation, $\mathbf{u}^t := \sum_{i=1}^n \mathbf{B}_i^t \mathbf{z}_i^t$ as the aggregate Hessian-variable product, and $\mathbf{g}^t := \sum_{i=1}^n \nabla f_i(\mathbf{z}_i^t)$ as the aggregate gradient. Then, given that at step t only a single index i_t is updated, we can evaluate these variables for step $t + 1$ as

$$(17) \quad \tilde{\mathbf{B}}^{t+1} = \tilde{\mathbf{B}}^t + (\mathbf{B}_{i_t}^{t+1} - \mathbf{B}_{i_t}^t),$$

$$(18) \quad \mathbf{u}^{t+1} = \mathbf{u}^t + (\mathbf{B}_{i_t}^{t+1} \mathbf{z}_{i_t}^{t+1} - \mathbf{B}_{i_t}^t \mathbf{z}_{i_t}^t),$$

$$(19) \quad \mathbf{g}^{t+1} = \mathbf{g}^t + (\nabla f_{i_t}(\mathbf{z}_{i_t}^{t+1}) - \nabla f_{i_t}(\mathbf{z}_{i_t}^t)).$$

Algorithm 1 Incremental quasi-Newton (IQN) method

Require: $\mathbf{x}^0, \{\nabla f_i(\mathbf{x}^0)\}_{i=1}^n, \{\mathbf{B}_i^0\}_{i=1}^n$
 1: Set $\mathbf{z}_1^0 = \dots = \mathbf{z}_n^0 = \mathbf{x}^0$
 2: Set $(\tilde{\mathbf{B}}^0)^{-1} = (\sum_{i=1}^n \mathbf{B}_i^0)^{-1}$, $\mathbf{u}^0 = \sum_{i=1}^n \mathbf{B}_i^0 \mathbf{x}^0$, $\mathbf{g}^0 = \sum_{i=1}^n \nabla f_i(\mathbf{x}^0)$
 3: **for** $t = 0, 1, 2, \dots$ **do**
 4: Set $i_t = (t \bmod n) + 1$
 5: Compute $\mathbf{x}^{t+1} = (\tilde{\mathbf{B}}^t)^{-1} (\mathbf{u}^t - \mathbf{g}^t)$ [cf. (16)]
 6: Compute $\mathbf{s}_{i_t}^{t+1}$, $\mathbf{y}_{i_t}^{t+1}$ [cf. (14)], and $\mathbf{B}_{i_t}^{t+1}$ [cf. (15)]
 7: Update \mathbf{u}^{t+1} [cf. (18)], \mathbf{g}^{t+1} [cf. (19)], and $(\tilde{\mathbf{B}}^{t+1})^{-1}$ [cf. (21), (22)]
 8: Update the functions' information tables as in (12), (13), and (15)
 9: **end for**

Thus, only $\mathbf{B}_{i_t}^{t+1}$ and $\nabla f_{i_t}(\mathbf{z}_{i_t}^{t+1})$ are required to be computed at step t .

Although the updates in (17)–(19) have low computational complexity, the update in (16) requires computing $(\tilde{\mathbf{B}}^t)^{-1}$, which has a computational complexity of $\mathcal{O}(p^3)$. This inversion can be avoided by simplifying the update in (17) as

$$(20) \quad \tilde{\mathbf{B}}^{t+1} = \tilde{\mathbf{B}}^t + \frac{\mathbf{y}_{i_t}^t \mathbf{y}_{i_t}^{tT}}{\mathbf{y}_{i_t}^{tT} \mathbf{s}_{i_t}^t} - \frac{\mathbf{B}_{i_t}^t \mathbf{s}_{i_t}^t \mathbf{s}_{i_t}^{tT} \mathbf{B}_{i_t}^t}{\mathbf{s}_{i_t}^{tT} \mathbf{B}_{i_t}^t \mathbf{s}_{i_t}^t}.$$

To derive the expression in (20) we have substituted the difference $\mathbf{B}_{i_t}^{t+1} - \mathbf{B}_{i_t}^t$ for its rank two expression in (15). Given the matrix $(\tilde{\mathbf{B}}^t)^{-1}$, by applying the Sherman–Morrison formula twice to the update in (20) we can compute $(\tilde{\mathbf{B}}^{t+1})^{-1}$ as

$$(21) \quad (\tilde{\mathbf{B}}^{t+1})^{-1} = \mathbf{U}^t + \frac{\mathbf{U}^t (\mathbf{B}_{i_t}^t \mathbf{s}_{i_t}^t) (\mathbf{B}_{i_t}^t \mathbf{s}_{i_t}^t)^T \mathbf{U}^t}{\mathbf{s}_{i_t}^{tT} \mathbf{B}_{i_t}^t \mathbf{s}_{i_t}^t - (\mathbf{B}_{i_t}^t \mathbf{s}_{i_t}^t)^T \mathbf{U}^t (\mathbf{B}_{i_t}^t \mathbf{s}_{i_t}^t)},$$

where the matrix \mathbf{U}^t is evaluated as

$$(22) \quad \mathbf{U}^t = (\tilde{\mathbf{B}}^t)^{-1} - \frac{(\tilde{\mathbf{B}}^t)^{-1} \mathbf{y}_{i_t}^t \mathbf{y}_{i_t}^{tT} (\tilde{\mathbf{B}}^t)^{-1}}{\mathbf{y}_{i_t}^{tT} \mathbf{s}_{i_t}^t + \mathbf{y}_{i_t}^{tT} (\tilde{\mathbf{B}}^t)^{-1} \mathbf{y}_{i_t}^t}.$$

The computational complexity of the updates in (21) and (22) is of the order $\mathcal{O}(p^2)$ rather than the $\mathcal{O}(p^3)$ cost of computing the inverse directly. Therefore, the overall cost of IQN is of the order $\mathcal{O}(p^2)$, which is substantially lower than the $\mathcal{O}(np^2)$ of deterministic quasi-Newton methods.

The complete IQN algorithm is outlined in Algorithm 1. Beginning with initial variable \mathbf{x}^0 and gradient and Hessian estimates $\nabla f_i(\mathbf{x}^0)$ and \mathbf{B}_i^0 for all i , each variable copy \mathbf{z}_i^0 is set to \mathbf{x}^0 in step 1 and initial values are set for \mathbf{u}^0 , \mathbf{g}^0 , and $(\tilde{\mathbf{B}}^0)^{-1}$ in step 2. For all t , in step 4 the index i_t of the next function to update is selected cyclically. The variable \mathbf{x}^{t+1} is computed according to the update in (16) in step 5. In step 6, the variable $\mathbf{s}_{i_t}^{t+1}$ and gradient $\mathbf{y}_{i_t}^{t+1}$ variations are evaluated as in (14) to compute the BFGS matrix $\mathbf{B}_{i_t}^{t+1}$ from the update in (15). This information, as well as the updated variable and its gradient, are used in step 7 to update \mathbf{u}^{t+1} and \mathbf{g}^{t+1} as in (18) and (19), respectively. The inverse matrix $(\tilde{\mathbf{B}}^{t+1})^{-1}$ is also computed by following the expressions in (21) and (22). Finally in step 8, we update the variable, gradient, and Hessian approximation tables based on the policies in (12), (13), and (15), respectively.

Remark 2. The proposed IQN method, similarly to other existing incremental aggregated methods, e.g., IAG [12], SAG [32], SAGA [9], Finito/MISO [10, 20], requires a memory of the order of the number of functions n , which in some large-scale

optimization problems might not be affordable. This issue may be resolved by grouping the functions into subsets and instead optimizing over the sum of the subsets. To be more specific, by combining every m functions and using their sum as the new function, the number of active functions reduces to n/m and consequently the required memory decreases to $\mathcal{O}(np/m)$. On the other hand, this process increases the computational complexity of each iteration from updating the information of a single function to that of m functions. Indeed, there is a trade-off between the memory and computational complexity per iteration which can be optimized based on the application of interest.

5. Convergence analysis. In this section, we study the convergence rate of the proposed IQN method. We first establish its local linear convergence rate, then demonstrate limit properties of the Hessian approximations, and finally show that in a region local to the optimal point the sequence of residuals converges at a superlinear rate. To prove these results we make two main assumptions, both of which are standard in the analysis of quasi-Newton methods.

Assumption 1. There exist positive constants $0 < \mu \leq L$ such that, for all i and $\mathbf{x}, \hat{\mathbf{x}} \in \mathbb{R}^p$, we can write

$$(23) \quad \mu \|\mathbf{x} - \hat{\mathbf{x}}\|^2 \leq (\nabla f_i(\mathbf{x}) - \nabla f_i(\hat{\mathbf{x}}))^T (\mathbf{x} - \hat{\mathbf{x}}) \leq L \|\mathbf{x} - \hat{\mathbf{x}}\|^2.$$

Assumption 2. There exists a positive constant $0 < \tilde{L}$ such that, for all i and $\mathbf{x}, \hat{\mathbf{x}} \in \mathbb{R}^p$, we can write

$$(24) \quad \|\nabla^2 f_i(\mathbf{x}) - \nabla^2 f_i(\hat{\mathbf{x}})\| \leq \tilde{L} \|\mathbf{x} - \hat{\mathbf{x}}\|.$$

The lower bound in (23) implies that the functions f_i are strongly convex with constant μ , and the upper bound shows that the gradients ∇f_i are Lipschitz continuous with parameter L . Note that Assumption 1 implies that the eigenvalues of the Hessians $\nabla^2 f_i$ are uniformly lower and upper bounded by μ and L , respectively.

The condition in Assumption 2 states that the Hessians $\nabla^2 f_i$ are Lipschitz continuous with constant \tilde{L} . This assumption is commonly made in the analyses of Newton's method [25] and quasi-Newton algorithms [4, 26, 13]. According to Lemma 3.1 in [4], Lipschitz continuity of the Hessians with constant \tilde{L} implies that for $i = 1, \dots, n$ and arbitrary vectors $\mathbf{x}, \tilde{\mathbf{x}}, \hat{\mathbf{x}} \in \mathbb{R}^p$ we can write

$$(25) \quad \|\nabla^2 f_i(\tilde{\mathbf{x}})(\mathbf{x} - \hat{\mathbf{x}}) - (\nabla f_i(\mathbf{x}) - \nabla f_i(\hat{\mathbf{x}}))\| \leq \tilde{L} \|\mathbf{x} - \hat{\mathbf{x}}\| \max \{\|\mathbf{x} - \tilde{\mathbf{x}}\|, \|\hat{\mathbf{x}} - \tilde{\mathbf{x}}\|\}.$$

Note that the inequality in (25) can also be achieved using the mean value theorem for gradients and the condition in (24). We use the inequality in (25) in the process of proving the convergence of IQN.

The goal of BFGS quasi-Newton methods is to approximate the objective function Hessian using the first-order information. Likewise, in the incremental BFGS method, we aim to show that the Hessian approximation matrices for all functions f_1, \dots, f_n are close to the exact Hessian. In the following lemma, we study the difference between the i th optimal Hessian $\nabla^2 f_i(\mathbf{x}^*)$ and its approximation \mathbf{B}_i^t over time.

LEMMA 1. Consider the proposed IQN method in (9). Further, let i be the index of the updated function at step t , i.e., $i = i_t$. Define the residual sequence for function f_i as $\sigma_i^t := \max\{\|\mathbf{z}_i^{t+1} - \mathbf{x}^*\|, \|\mathbf{z}_i^t - \mathbf{x}^*\|\}$ and set $\mathbf{M}_i = \nabla^2 f_i(\mathbf{x}^*)^{-1/2}$. If Assumptions 1 and 2 hold and the condition $\sigma_i^t < m/(3\tilde{L})$ is satisfied, then

$$(26) \quad \|\mathbf{B}_i^{t+1} - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} \leq \left[(1 - \alpha\theta_i^{t^2})^{1/2} + \alpha_3\sigma_i^t \right] \|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} + \alpha_4\sigma_i^t,$$

where α , α_3 , and α_4 are some positive bounded constants and

(27)

$$\theta_i^t = \frac{\|\mathbf{M}_i(\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*))\mathbf{s}_i^t\|}{\|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} \|\mathbf{M}_i^{-1} \mathbf{s}_i^t\|} \text{ for } \mathbf{B}_i^t \neq \nabla^2 f_i(\mathbf{x}^*), \quad \theta_i^t = 0 \text{ for } \mathbf{B}_i^t = \nabla^2 f_i(\mathbf{x}^*).$$

Proof. See Appendix A. □

The result in (26) establishes an upper bound for the weighted norm $\|\mathbf{B}_i^{t+1} - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i}$ with respect to its previous value $\|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i}$ and the sequence $\sigma_i^t := \max\{\|\mathbf{z}_i^{t+1} - \mathbf{x}^*\|, \|\mathbf{z}_i^t - \mathbf{x}^*\|\}$ when the variables are in a neighborhood of the optimal solution such that $\sigma_i^t < m/(3\tilde{L})$. Indeed, the result in (26) holds only for the index $i = i_t$ and for the rest of indices we have $\|\mathbf{B}_i^{t+1} - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} = \|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i}$, simply by definition of the cyclic update. Note that if the residual sequence σ_i^t associated with f_i approaches zero, we can simplify (26) as

$$(28) \quad \|\mathbf{B}_i^{t+1} - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} \lesssim (1 - \alpha\theta_i^{t^2})^{1/2} \|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i}.$$

By using the inequality $(1 - \alpha\theta_i^{t^2})^{1/2} \leq 1 - (\alpha/2)\theta_i^{t^2}$, regrouping terms, and summing both sides from $t = 0$ to $t = \infty$, we obtain that

(29)

$$\frac{\alpha}{2} \sum_{t=0}^{\infty} \theta_i^{t^2} \|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} \lesssim \|\mathbf{B}_i^0 - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} - \|\mathbf{B}_i^\infty - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} < \infty,$$

assuming that the sequence $\|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i}$ has a limit. Now if a subsequence of $\{\|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i}\}$ converges to zero, since the sequence has a limit, the whole sequence converges to zero, and, therefore, $\lim_{t \rightarrow \infty} \|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} = 0$. If not, then the sequence $\{\|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i}\}$ is bounded away from zero and θ_i^t converges to zero, which implies the Dennis–Moré condition from (6), i.e.,

$$(30) \quad \lim_{t \rightarrow \infty} \frac{\|(\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*))\mathbf{s}_i^t\|}{\|\mathbf{s}_i^t\|} = 0.$$

Therefore, under both conditions the result in (30) holds. This is true since the limit $\lim_{t \rightarrow \infty} \|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} = 0$ yields the result in (30).

Based on this intuition, we proceed to show that the sequence σ_i^t converges to zero for all $i = 1, \dots, n$. To do so, we show that the sequence $\|\mathbf{z}_i^t - \mathbf{x}^*\|$ is linearly convergent for all $i = 1, \dots, n$. To achieve this goal we first prove an upper bound for the error $\|\mathbf{x}^{t+1} - \mathbf{x}^*\|$ of IQN in the following lemma.

LEMMA 2. *Consider the proposed IQN method in (9). If the conditions in Assumptions 1 and 2 hold, then the sequence of iterates generated by IQN satisfies*

$$(31) \quad \|\mathbf{x}^{t+1} - \mathbf{x}^*\| \leq \frac{\tilde{L}\Gamma^t}{n} \sum_{i=1}^n \|\mathbf{z}_i^t - \mathbf{x}^*\|^2 + \frac{\Gamma^t}{n} \sum_{i=1}^n \|(\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)) (\mathbf{z}_i^t - \mathbf{x}^*)\|,$$

where $\Gamma^t := \|((1/n) \sum_{i=1}^n \mathbf{B}_i^t)^{-1}\|$ and \tilde{L} is the constant of Lipschitz continuity of the Hessians defined in (24).

Proof. See Appendix B. □

Lemma 2 shows that the residual $\|\mathbf{x}^{t+1} - \mathbf{x}^*\|$ is bounded above by a sum of quadratic and linear terms of the last n residuals. This can eventually lead to a superlinear convergence rate by showing the linear term converges to zero at a rate faster than linear, i.e., $\lim_{t \rightarrow \infty} \|(\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*))(\mathbf{z}_i^t - \mathbf{x}^*)\|/\|\mathbf{z}_i^t - \mathbf{x}^*\| = 0$, leaving us with an upper bound of quadratic terms only. First, however, we establish a local linear convergence rate in the following lemma to show that the sequence σ_i^t converges to zero.

LEMMA 3. *Consider the proposed IQN method in (9). If Assumptions 1 and 2 hold, then, for any $r \in (0, 1)$ there are positive constants $\epsilon(r)$ and $\delta(r)$ such that if $\|\mathbf{x}^0 - \mathbf{x}^*\| < \epsilon(r)$ and $\|\mathbf{B}_i^0 - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} < \delta(r)$ for $\mathbf{M}_i = \nabla^2 f_i(\mathbf{x}^*)^{-1/2}$ and $i = 1, 2, \dots, n$, the sequence of iterates generated by IQN satisfies*

$$(32) \quad \|\mathbf{x}^t - \mathbf{x}^*\| \leq r^{\lfloor \frac{t-1}{n} \rfloor + 1} \|\mathbf{x}^0 - \mathbf{x}^*\|.$$

Moreover, the sequences of norms $\{\|\mathbf{B}_i^t\|\}$ and $\{\|(\mathbf{B}_i^t)^{-1}\|\}$ are uniformly bounded.

Proof. See Appendix C. □

The result in Lemma 3 shows that the sequence of iterates generated by IQN has a local linear convergence rate after each pass over all functions. Consequently, we obtain that the i th residual sequence σ_i^t is linearly convergent for all i . Note that Lemma 3 can be considered as an extension of Theorem 3.2 of [4] for incremental settings. Note that, to achieve the results in Lemma 3, Assumptions 1 and 2 can be relaxed to only hold in a local neighborhood of the optimal solution. To be more precise, we only need these assumptions to hold in a neighborhood of the optimal solution defined in the statement of Lemma 3 in contrast to the original conditions in Assumptions 1 and 2, in which they hold for any point in \mathbb{R}^P .

Following the arguments in (28)–(30), we use the summability of the sequence $\sigma_i^t := \max\{\|\mathbf{z}_i^{t+1} - \mathbf{x}^*\|, \|\mathbf{z}_i^t - \mathbf{x}^*\|\}$ along with the result in Lemma 1 to prove Dennis–Moré condition for all functions f_i .

PROPOSITION 4. *Consider the proposed IQN method in (9). Assume that the hypotheses in Lemmata 1 and 3 are satisfied. If the optimum is not achieved at finitely many iterations, then, for all $i = 1, \dots, n$, it holds that*

$$(33) \quad \lim_{t \rightarrow \infty} \frac{\|(\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*))\mathbf{s}_i^t\|}{\|\mathbf{s}_i^t\|} = 0.$$

Proof. See Appendix D. □

The statement in Proposition 4 indicates that for each function f_i the Dennis–Moré condition holds. In the traditional quasi-Newton methods the Dennis–Moré condition is sufficient to show that the method is superlinearly convergent. However, the same argument does not hold for the proposed IQN method, since we can’t recover the Dennis–Moré condition for the global objective function f from the result in Proposition 4. In other words, the result in (33) does not imply the limit in (6) required in the superlinear convergence analysis of quasi-Newton methods. Therefore, here we pursue a different approach and seek to prove that the linear terms $(\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*))(\mathbf{z}_i^t - \mathbf{x}^*)$ in (31) converge to zero at a superlinear rate, i.e., for all i we can write $\lim_{t \rightarrow \infty} \|(\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*))(\mathbf{z}_i^t - \mathbf{x}^*)\|/\|\mathbf{z}_i^t - \mathbf{x}^*\| = 0$. If we establish this result, it follows from the result in Lemma 2 that the sequence of residuals $\|\mathbf{x}^t - \mathbf{x}^*\|$ converges to zero superlinearly.

We continue the analysis of the proposed IQN method by establishing a generalized limit property that follows from the Dennis–Moré criterion in (6). In the

following lemma, we leverage the local linear convergence of the iterates \mathbf{x}^t to show that the vector $\mathbf{z}_i^t - \mathbf{x}^*$ lies in the null space of $\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)$ as t approaches infinity.

LEMMA 5. *Consider the proposed IQN method in (9). Assume that the hypotheses in Lemmata 1 and 3 are satisfied. As t approaches infinity, for all i it holds that*

$$(34) \quad \lim_{t \rightarrow \infty} \frac{\|(\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*))(\mathbf{z}_i^t - \mathbf{x}^*)\|}{\|\mathbf{z}_i^t - \mathbf{x}^*\|} = 0.$$

Proof. See Appendix E. □

The result in Lemma 5 can thus be used in conjunction with Lemma 2 to show that the residual $\|\mathbf{x}^{t+1} - \mathbf{x}^*\|$ is bounded by a sum of quadratic terms of previous residuals and a term that converges to zero superlinearly. This result leads us to the local superlinear convergence of the sequence of residuals with respect to the average sequence, stated in the following theorem.

THEOREM 6. *Consider the proposed IQN method in (9). Suppose that the conditions in the hypotheses of Lemmata 1 and 3 are valid. If the optimum is not achieved at finitely many iterations, then the sequence of residuals $\|\mathbf{x}^t - \mathbf{x}^*\|$ satisfies*

$$(35) \quad \lim_{t \rightarrow \infty} \frac{\|\mathbf{x}^t - \mathbf{x}^*\|}{\frac{1}{n}(\|\mathbf{x}^{t-1} - \mathbf{x}^*\| + \dots + \|\mathbf{x}^{t-n} - \mathbf{x}^*\|)} = 0.$$

Proof. See Appendix F. □

The result in (35) shows a mean-superlinear convergence rate for the sequence of iterates generated by IQN. To be more precise, it shows that the ratio that captures the error at step t divided by the average of the last n errors converges to zero. This is not equivalent to the classic Q-superlinear convergence for full-batch quasi-Newton methods, i.e., $\lim_{t \rightarrow \infty} \|\mathbf{x}^{t+1} - \mathbf{x}^*\|/\|\mathbf{x}^t - \mathbf{x}^*\| = 0$. Although we cannot prove Q-superlinear convergence of the residuals $\|\mathbf{x}^t - \mathbf{x}^*\|$, we can show that there exists a subsequence of the sequence $\|\mathbf{x}^t - \mathbf{x}^*\|$ that converges to zero superlinearly. In addition, there exists a superlinearly convergent sequence that is an upper bound for the original sequence of errors $\|\mathbf{x}^t - \mathbf{x}^*\|$. We formalize these results in the following theorem.

THEOREM 7. *Consider the IQN method in (9). Suppose that the conditions in the hypotheses of Lemmata 1 and 3 hold. Define $\tilde{\mathbf{x}}^t = \operatorname{argmax}_{u \in \{tn, \dots, tn+n-1\}} \{\|\mathbf{x}^u - \mathbf{x}^*\|\}$ as the iterate that has the largest error among the iterates in the $t+1$ th pass. Then, the sequence of iterates $\{\tilde{\mathbf{x}}^t\}_{t=0}^\infty$, which is a subsequence of the sequence $\{\mathbf{x}^t\}_{t=0}^\infty$, converges to \mathbf{x}^* at a superlinear rate, i.e.,*

$$(36) \quad \lim_{t \rightarrow \infty} \frac{\|\tilde{\mathbf{x}}^{t+1} - \mathbf{x}^*\|}{\|\tilde{\mathbf{x}}^t - \mathbf{x}^*\|} = 0.$$

Moreover, there exists a sequence ζ^t such that $\|\mathbf{x}^t - \mathbf{x}^*\| \leq \zeta^t$ for all $t \geq 0$, and the sequence ζ^t converges to zero at a superlinear rate, i.e.,

$$(37) \quad \lim_{t \rightarrow \infty} \frac{\zeta^{t+1}}{\zeta^t} = 0.$$

Proof. See Appendix G. □

The first result in Theorem 7 states that although the whole sequence $\|\mathbf{x}^t - \mathbf{x}^*\|$ is not necessarily superlinearly convergent, there exists a subsequence of the sequence $\|\mathbf{x}^t - \mathbf{x}^*\|$ that converges at a superlinear rate. The second claim in Theorem 7 establishes the R-superlinear convergence rate of the whole sequence $\|\mathbf{x}^t - \mathbf{x}^*\|$. In other words, it guarantees that $\|\mathbf{x}^t - \mathbf{x}^*\|$ is upper bounded by a superlinearly convergent sequence.

6. Numerical results. We proceed by simulating the performance of IQN on a variety of machine learning problems on both artificial and real datasets. We compare the performance of IQN against a collection of well-known first-order stochastic and incremental algorithms—namely SAG [32], SAGA [9], and IAG [12]—and stochastic quasi-Newton methods—namely RES [22], SVRG-SQN [24], and SVRG-LBFGS [16]. To begin, we look at a simple quadratic program, also equivalent to the solution of a linear least squares estimation problem. Consider the objective function to be minimized,

$$(38) \quad \mathbf{x}^* = \operatorname{argmin}_{\mathbf{x} \in \mathbb{R}^p} f(\mathbf{x}) := \operatorname{argmin}_{\mathbf{x} \in \mathbb{R}^p} \frac{1}{n} \sum_{i=1}^n \frac{1}{2} \mathbf{x}^T \mathbf{A}_i \mathbf{x} + \mathbf{b}_i^T \mathbf{x}.$$

We generate $\mathbf{A}_i \in \mathbb{R}^{p \times p}$ as a random positive definite matrix and $\mathbf{b}_i \in \mathbb{R}^p$ as a random vector for all i . In particular we set the matrices $\mathbf{A}_i := \operatorname{diag}\{\mathbf{a}_i\}$ and generate random vectors \mathbf{a}_i with the first $p/2$ elements chosen from $[1, 10^\xi]$ and the last $p/2$ elements chosen from $[10^{-\xi}, 1]$. The parameter ξ is used to manually set the condition number for the quadratic program in (38), ranging from $\xi = 1$ (i.e., small condition number 10^2) and $\xi = 2$ (i.e., large condition number 10^4). The vectors \mathbf{b}_i are chosen uniformly and randomly from the box $[0, 10^3]^p$. The variable dimension is set to be $p = 10$ and the number of functions to be $n = 1000$. Given that we focus on local convergence, we use a constant step size of $\eta = 1$ for the proposed IQN method while choosing the largest stepsize allowable by the other methods to converge.

In Figure 2 (see online version for color figures) we present a simulation of the convergence path of the normalized error $\|\mathbf{x}^t - \mathbf{x}^*\| / \|\mathbf{x}^0 - \mathbf{x}^*\|$ for the quadratic program in comparison to first-order methods. In the the top images, we show a sample simulation path for all methods on the quadratic problem with a small condition number. Stepsizes of $\eta = 5 \times 10^{-5}$, $\eta = 10^{-4}$, and $\eta = 10^{-6}$ were used for SAG, SAGA, and IAG, respectively. These stepsizes are tuned to compare the best performance of these methods with IQN. The proposed method reaches a error of 10^{-10} after 10 passes through the data. Alternatively, SAG and SAGA achieve the same error after 47 and 40 passes, respectively, while IAG reaches an error of only 10^{-2} after 60 passes. However, because the computational complexity of IQN is higher, we show in the top right image the same convergence paths with respect to runtime. Here we see IQN achieve similar performance to that of SAG and SAGA for the well-conditioned problem.

In the bottom images of Figure 2 (see online version for color figures), we repeat the same simulation but with larger condition number. In this case, SAG uses stepsize $\eta = 2 \times 10^{-4}$ while others remain the same. Observe that while the performance of IQN does not degrade with larger condition number, the first-order methods all suffer large degradation. SAGA reaches an error of 10^{-2} after 60 passes, while the other two methods descend very little in 60 passes. This same difference holds true when comparing runtimes in the right-hand image as well. It can be seen that IQN significantly outperforms the first-order method for both condition number sizes, with




FIG. 2. Convergence results of the proposed IQN method in comparison to first-order SAG, SAGA, and IAG. In the top images, we present a sample convergence path of the normalized error on the quadratic program with a small condition number in terms of effective passes and time. In the bottom image, we show the convergence path for the quadratic program with a large condition number. For small condition number, IQN provides similar performance to SAGA, while providing significant improvement in convergence rates for larger condition number.

such improvement increasing for larger condition number. This is an expected result, as first-order methods often do not perform well for ill-conditioned problems.

We compare the performance of IQN against the other stochastic quasi-Newton methods (top) and full-batch BFGS (bottom) in Figure 3 on the ill-conditioned quadratic program. Observe that IQN outperforms all other simulated methods, with the strongest alternatives being SVRG-SQN, achieving an error of 10^{-8} after 60 passes, and BFGS, achieving an error of 10^{-10} after 25 passes. We also point out that, when comparing against other quasi-Newton methods, the relative benefits of IQN do not diminish when measured in terms of runtime because the computational complexity is similar across all stochastic quasi-Newton methods.

Remark 3. The proposed IQN method needs to store iterates, gradients, and approximations to the Hessian of n functions, which may not be affordable in some applications when n is extremely large. We highlight, however, that all stochastic and incremental methods—including first-order techniques—that recover convergence rates of their deterministic counterparts require a source of variance reduction that also comes at a cost. This cost is either incurred through the use of memory to keep track of the information of all n functions (as in SAG [32] and IQN), or through the need to compute the full gradient once in a while (as in SVRG [14] and SVRG-type quasi-Newton methods [24, 16]). The first approach suffers from high storage cost while the second requires computation of the full gradient, which may not be feasible in some applications. In our proposed IQN method, we have opted for the first approach, and the analysis that recovers superlinear convergence relies heavily on this storage. However, with further investigation it may be possible to design an




FIG. 3. Convergence results of proposed IQN method in comparison to second-order RES, SVRG-SQN, and SVRG-LBFGS (top) and full-batch BFGS (bottom). In the left-hand image, we present a sample convergence path of the normalized error on the quadratic program in terms of number of effective passes. In the right-hand image, we show the convergence path in terms of runtime. In both cases, IQN provides significant improvement over the other stochastic and full-batch quasi-Newton methods.

SVRG-type quasi-Newton method that recovers the same rate in a similar fashion.

Overall, we suggest that if one can afford a storage of order n , the superlinearly convergent IQN is the quasi-Newton method of choice; if memory of order n is not affordable but occasional computation of gradients is acceptable, the linearly-convergent SVRG-type quasi-Newton methods such as SVRG-oLBFGS [16] and SVRG-SQN [24] can be executed. If neither is feasible, then one may use low storage and computational cost stochastic quasi-Newton methods such as RES [22], SQN [6], and oLBFGS [33], which all converge at a sublinear rate of $\mathcal{O}(1/t)$.

6.1. Logistic regression. We proceed to numerically evaluate the performance of IQN relative to existing methods on the classification of handwritten digits in the MNIST database [17]. In particular, we solve the binary logistic regression problem. A logistic regression takes as inputs n training feature vectors $\mathbf{u}_i \in \mathbb{R}^p$ with associated labels $v_i \in \{-1, 1\}$ and outputs a linear classifier \mathbf{x} to predict the label of unknown feature vectors. For the digit classification problem, each feature vector \mathbf{u}_i represents a vectorized image and v_i represents its label as one of two digits. We evaluate for any training sample i the probability of a label $v_i = 1$ given image \mathbf{u}_i as $P(v = 1|\mathbf{u}) = 1/(1 + \exp(-\mathbf{u}^T \mathbf{x}))$. The classifier \mathbf{x} is chosen to be the vector which maximizes the log likelihood across all n samples. Given n images \mathbf{u}_i with associated labels v_i , the optimization problem for logistic regression is written as

$$(39) \quad \mathbf{x}^* = \underset{\mathbf{x} \in \mathbb{R}^p}{\operatorname{argmin}} f(\mathbf{x}) := \underset{\mathbf{x} \in \mathbb{R}^p}{\operatorname{argmin}} \frac{\lambda}{2} \|\mathbf{x}\|^2 + \frac{1}{n} \sum_{i=1}^n \log[1 + \exp(-v_i \mathbf{u}_i^T \mathbf{x})],$$

where the first term is a regularization term parametrized by $\lambda \geq 0$.




FIG. 4. *Convergence results for a sample convergence path for the logistic regression problem on classifying handwritten digits. IQN substantially outperforms the first-order methods.*

For our simulations we select from the MNIST dataset $n = 1000$ images with dimension $p = 784$ labelled as one of the digits “0” or “8” and fix the regularization parameter as $\lambda = 1/n$ and stepsize as $\eta = 0.01$ for all first-order methods. In Figure 4 we present the convergence path of IQN relative to existing methods in terms of the norm of the gradient. As in the case of the quadratic program, the IQN performs better than all the considered gradient-based methods. IQN reaches a gradient magnitude of 4.8×10^{-8} after 60 passes through the data while SAGA reaches only a magnitude of 7.4×10^{-5} (all other methods perform even worse). Further note that while the first-order methods begin to level out after 60 passes, the IQN method continues to descend. These results demonstrate the effectiveness of IQN on a practical machine learning problem with real-world data.

7. Conclusion. In this paper we proposed a novel incremental quasi-Newton (IQN) method for solving the average of a set of n smooth and strongly convex functions. The proposed IQN method recovers a superlinear convergence rate of deterministic quasi-Newton methods using aggregated information of variables, gradients, and quasi-Newton Hessian approximation matrices to reduce the noise of gradient and Hessian approximations, a consistent Taylor expansion of each individual function f_i , and a cyclic scheme to update the functions in lieu of a random selection routine. Numerical experiments showcase superior performance of IQN compared to existing stochastic quasi-Newton methods as well as variance reduced first-order methods.

As a future research direction, we aim to develop a limited memory version of the IQN method to reduce its required memory from $\mathcal{O}(np^2)$ to $\mathcal{O}(\tau np)$ and lower its computational complexity per iteration from $\mathcal{O}(p^2)$ to $\mathcal{O}(\tau p)$, where $\tau \ll p$ is the size of the memory. Another interesting future direction is investigating the possibility of proving a superlinear convergence rate for the stochastic variant of IQN, which updates functions in a stochastic manner in lieu of the cyclic update used in this paper. Indeed, the convergence guarantee of such a stochastic method in expectation should be better than the one for the (cyclic) IQN method, as has been observed for first-order cyclic and stochastic methods.

Appendix A. Proof of Lemma 1. To prove the claim in Lemma 1, we first prove the the following lemma, which is based on the result in [4, Lemma 5.2].

LEMMA 8. *Consider the proposed IQN method in (9). Let \mathbf{M}_i be a nonsingular symmetric matrix such that*

$$(40) \quad \|\mathbf{M}_i \mathbf{y}_i^t - \mathbf{M}_i^{-1} \mathbf{s}_i^t\| \leq \beta \|\mathbf{M}_i^{-1} \mathbf{s}_i^t\|$$

for some $\beta \in [0, 1/3]$ and vectors \mathbf{s}_i^t and \mathbf{y}_i^t in \mathbb{R}^p with $\mathbf{s}_i^t \neq \mathbf{0}$. Consider i as the index of the updated function at step t , i.e., $i = i_t$, and let \mathbf{B}_i^t be symmetric and computed according to the update in (15). Then, there exist positive constants α , α_1 , and α_2 such that, for any symmetric $\mathbf{A} \in \mathbb{R}^{p \times p}$, we have

$$(41) \quad \|\mathbf{B}_i^{t+n} - \mathbf{A}\|_{\mathbf{M}_i} \leq \left[(1 - \alpha\theta^2)^{\frac{1}{2}} + \alpha_1 \frac{\|\mathbf{M}_i \mathbf{y}_i^t - \mathbf{M}_i^{-1} \mathbf{s}_i^t\|}{\|\mathbf{M}_i^{-1} \mathbf{s}_i^t\|} \right] \|\mathbf{B}_i^t - \mathbf{A}\|_{\mathbf{M}_i} + \alpha_2 \frac{\|\mathbf{y}_i^t - \mathbf{A} \mathbf{s}_i^t\|}{\|\mathbf{M}_i^{-1} \mathbf{s}_i^t\|},$$

where $\alpha = (1 - 2\beta)/(1 - \beta^2) \in [3/8, 1]$, $\alpha_1 = 2.5(1 - \beta)^{-1}$, $\alpha_2 = 2(1 + 2\sqrt{\beta})\|\mathbf{M}_i\|_{\mathbf{F}}$, and

$$(42) \quad \theta = \frac{\|\mathbf{M}_i(\mathbf{B}_i^t - \mathbf{A})\mathbf{s}_i^t\|}{\|\mathbf{B}_i^t - \mathbf{A}\|_{\mathbf{M}_i} \|\mathbf{M}_i^{-1} \mathbf{s}_i^t\|} \quad \text{for } \mathbf{B}_i^t \neq \mathbf{A}, \quad \theta = 0 \quad \text{for } \mathbf{B}_i^t = \mathbf{A}.$$

Proof. First note that the Hessian approximation \mathbf{B}_i^{t+n} is equal to \mathbf{B}_i^{t+1} if the function f_i is updated at step t . Considering this observation and the result of Lemma 5.2 in [4], the claim in (41) follows. \square

The result in Lemma 8 provides an upper bound for the difference between the Hessian approximation matrix \mathbf{B}_i^{t+n} and any positive definite matrix \mathbf{A} with respect to the difference between the previous Hessian approximation \mathbf{B}_i^t and the matrix \mathbf{A} . The interesting choice for the arbitrary matrix \mathbf{A} is the Hessian of the i th function at the optimal argument, i.e., $\mathbf{A} = \nabla^2 f_i(\mathbf{x}^*)$, which allows us to capture the difference between the sequence of Hessian approximation matrices for function f_i and the Hessian $\nabla^2 f_i(\mathbf{x}^*)$ at the optimal argument. We proceed to use the result in Lemma 8 for $\mathbf{M}_i = \nabla^2 f_i(\mathbf{x}^*)^{-1/2}$ and $\mathbf{A} = \nabla^2 f_i(\mathbf{x}^*)$ to prove the claim in (26). To do so, we first need to show that the condition in (40) is satisfied. Note that according to the condition in Assumptions 1 and 2 we can write

$$(43) \quad \frac{\|\mathbf{y}_i^t - \nabla^2 f_i(\mathbf{x}^*) \mathbf{s}_i^t\|}{\|\nabla^2 f_i(\mathbf{x}^*)^{1/2} \mathbf{s}_i^t\|} \leq \frac{\tilde{L} \|\mathbf{s}_i^t\| \max\{\|\mathbf{z}_i^t - \mathbf{x}^*\|, \|\mathbf{z}_i^{t+1} - \mathbf{x}^*\|\}}{\sqrt{m} \|\mathbf{s}_i^t\|} = \frac{\tilde{L}}{\sqrt{m}} \sigma_i^t.$$

This observation implies that the left-hand side of the condition in (40) for $\mathbf{M}_i = \nabla^2 f_i(\mathbf{x}^*)^{-1/2}$ is bounded above by

$$(44) \quad \frac{\|\mathbf{M}_i \mathbf{y}_i^t - \mathbf{M}_i^{-1} \mathbf{s}_i^t\|}{\|\mathbf{M}_i^{-1} \mathbf{s}_i^t\|} \leq \frac{\|\nabla^2 f_i(\mathbf{x}^*)^{-1/2}\| \|\mathbf{y}_i^t - \nabla^2 f_i(\mathbf{x}^*) \mathbf{s}_i^t\|}{\|\nabla^2 f_i(\mathbf{x}^*)^{1/2} \mathbf{s}_i^t\|} \leq \frac{\tilde{L}}{m} \sigma_i^t.$$

Thus, the condition in (40) is satisfied since $\tilde{L} \sigma_i^t / m < 1/3$. Substituting the upper bounds in (43) and (44) into the expression in (41) implies the claim in (26) with

$$(45) \quad \beta = \frac{\tilde{L}}{m} \sigma_i^t, \quad \alpha = \frac{1 - 2\beta}{1 - \beta^2}, \quad \alpha_3 = \frac{5\tilde{L}}{2m(1 - \beta)}, \quad \alpha_4 = \frac{2(1 + 2\sqrt{\beta})\tilde{L}}{\sqrt{m}} \|\nabla^2 f_i(\mathbf{x}^*)^{-\frac{1}{2}}\|_{\mathbf{F}},$$

and the proof is complete.

Appendix B. Proof of Lemma 2. Start by subtracting \mathbf{x}^* from both sides of (9) to obtain

$$(46) \quad \mathbf{x}^{t+1} - \mathbf{x}^* = \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i^t \right)^{-1} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i^t \mathbf{z}_i^t - \frac{1}{n} \sum_{i=1}^n \nabla f_i(\mathbf{z}_i^t) - \frac{1}{n} \sum_{i=1}^n \mathbf{B}_i^t \mathbf{x}^* \right).$$

As the gradient of f at the optimal point is the vector zero, i.e., $(1/n) \sum_{i=1}^n \nabla f_i(\mathbf{x}^*) = \mathbf{0}$, we can subtract $(1/n) \sum_{i=1}^n \nabla f_i(\mathbf{x}^*)$ from the right-hand side of (46) and rearrange terms to obtain

$$(47) \quad \mathbf{x}^{t+1} - \mathbf{x}^* = \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i^t \right)^{-1} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i^t (\mathbf{z}_i^t - \mathbf{x}^*) - \frac{1}{n} \sum_{i=1}^n (\nabla f_i(\mathbf{z}_i^t) - \nabla f_i(\mathbf{x}^*)) \right).$$

The expression in (47) relates the residual at time $t + 1$ to the previous n residuals and the Hessian approximations \mathbf{B}_i^t . To analyze this further, we can replace the Hessian approximations \mathbf{B}_i^t with the actual Hessians $\nabla^2 f_i(\mathbf{x}^*)$ and the approximation difference $\nabla^2 f_i(\mathbf{x}^*) - \mathbf{B}_i^t$. To do so, we add and subtract $(1/n) \sum_{i=1}^n \nabla^2 f_i(\mathbf{x}^*) (\mathbf{z}_i^t - \mathbf{x}^*)$ to the right-hand side of (47) and rearrange terms to obtain

$$(48) \quad \begin{aligned} \mathbf{x}^{t+1} - \mathbf{x}^* &= \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i^t \right)^{-1} \left(\frac{1}{n} \sum_{i=1}^n [\nabla^2 f_i(\mathbf{x}^*) (\mathbf{z}_i^t - \mathbf{x}^*) - (\nabla f_i(\mathbf{z}_i^t) - \nabla f_i(\mathbf{x}^*))] \right) \\ &+ \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i^t \right)^{-1} \left(\frac{1}{n} \sum_{i=1}^n [\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)] (\mathbf{z}_i^t - \mathbf{x}^*) \right). \end{aligned}$$

We proceed to take the norms of both sides and use the triangle inequality to obtain an upper bound on the norm of the residual $\|\mathbf{x}^{t+1} - \mathbf{x}^*\|$,

$$(49) \quad \begin{aligned} \|\mathbf{x}^{t+1} - \mathbf{x}^*\| &\leq \left\| \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i^t \right)^{-1} \right\| \left\| \frac{1}{n} \sum_{i=1}^n \|\nabla^2 f_i(\mathbf{x}^*) (\mathbf{z}_i^t - \mathbf{x}^*) - (\nabla f_i(\mathbf{z}_i^t) - \nabla f_i(\mathbf{x}^*))\| \right\| \\ &+ \left\| \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i^t \right)^{-1} \right\| \left\| \frac{1}{n} \sum_{i=1}^n \|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\| (\mathbf{z}_i^t - \mathbf{x}^*) \right\|. \end{aligned}$$

To obtain the quadratic term in (31) from the first term in (49), we use the Lipschitz continuity of the Hessians $\nabla^2 f_i$, which leads to the inequality

$$(50) \quad \|\nabla^2 f_i(\mathbf{x}^*) (\mathbf{z}_i^t - \mathbf{x}^*) - (\nabla f_i(\mathbf{z}_i^t) - \nabla f_i(\mathbf{x}^*))\| \leq \tilde{L} \|\mathbf{z}_i^t - \mathbf{x}^*\|^2.$$

Replacing the expression $\|\nabla^2 f_i(\mathbf{x}^*) (\mathbf{z}_i^t - \mathbf{x}^*) - (\nabla f_i(\mathbf{z}_i^t) - \nabla f_i(\mathbf{x}^*))\|$ in (49) by the upper bound in (50), the claim in (31) follows.

Appendix C. Proof of Lemma 3. In this proof we use some steps from the proof of [4, Theorem 3.2]. To start we use the fact that in a finite-dimensional vector space there always exists a constant $\eta > 0$ such that $\|\mathbf{A}\| \leq \eta \|\mathbf{A}\|_{\mathbf{M}_i}$. Consider $\gamma = 1/m$ as an upper bound for the norm $\|\nabla^2 f(\mathbf{x}^*)^{-1}\|$. Assume that $\epsilon(r) = \epsilon$ and $\delta(r) = \delta$ are chosen such that

$$(51) \quad (2\alpha_3\delta + \alpha_4) \frac{\epsilon}{1-r} \leq \delta \quad \text{and} \quad \gamma(1+r)[\tilde{L}\epsilon + 2\eta\delta] \leq r.$$

Based on the assumption that $\|\mathbf{B}_i^0 - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} \leq \delta$ we can derive the upper bound $\|\mathbf{B}_i^0 - \nabla^2 f_i(\mathbf{x}^*)\| \leq \eta\delta$. This observation along with the inequality $\|\nabla^2 f_i(\mathbf{x}^*)\| \leq L$ implies that $\|\mathbf{B}_i^0\| \leq \eta\delta + L$. Therefore, we obtain $\|(1/n) \sum_{i=1}^n \mathbf{B}_i^0\| \leq \eta\delta + L$. The second inequality in (51) implies that $2\gamma(1+r)\eta\delta \leq r$. Based on this observation and the inequalities $\|\mathbf{B}_i^0 - \nabla^2 f_i(\mathbf{x}^*)\| \leq \eta\delta < 2\eta\delta$ and $\gamma \geq \|\nabla^2 f_i(\mathbf{x}^*)^{-1}\|$, we obtain from Banach’s lemma that $\|(\mathbf{B}_i^0)^{-1}\| \leq (1+r)\gamma$. Following the same argument for the matrix

$((1/n) \sum_{i=1}^n \mathbf{B}_i^0)^{-1}$ with the inequalities $\|(1/n) \sum_{i=1}^n \mathbf{B}_i^0 - (1/n) \sum_{i=1}^n \nabla^2 f_i(\mathbf{x}^*)\| \leq (1/n) \sum_{i=1}^n \|\mathbf{B}_i^0 - \nabla^2 f_i(\mathbf{x}^*)\| \leq \eta\delta$ and $\|\nabla^2 f(\mathbf{x}^*)^{-1}\| \leq \gamma$ we obtain that

$$(52) \quad \left\| \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i^0 \right)^{-1} \right\| \leq (1+r)\gamma.$$

This upper bound in conjunction with the result in (31) yields

$$(53) \quad \begin{aligned} \|\mathbf{x}^1 - \mathbf{x}^*\| &\leq (1+r)\gamma \left[\frac{\tilde{L}}{n} \sum_{i=1}^n \|\mathbf{z}_i^0 - \mathbf{x}^*\|^2 + \frac{1}{n} \sum_{i=1}^n \|\mathbf{B}_i^0 - \nabla^2 f_i(\mathbf{x}^*)\| \|\mathbf{z}_i^0 - \mathbf{x}^*\| \right] \\ &= (1+r)\gamma \left[\tilde{L} \|\mathbf{x}^0 - \mathbf{x}^*\|^2 + \frac{1}{n} \sum_{i=1}^n \|\mathbf{B}_i^0 - \nabla^2 f_i(\mathbf{x}^*)\| \|\mathbf{x}^0 - \mathbf{x}^*\| \right]. \end{aligned}$$

Considering the assumptions that $\|\mathbf{x}^0 - \mathbf{x}^*\| \leq \epsilon$ and $\|\mathbf{B}_i^0 - \nabla^2 f_i(\mathbf{x}^*)\| \leq \eta\delta < 2\eta\delta$ we can write

$$(54) \quad \begin{aligned} \|\mathbf{x}^1 - \mathbf{x}^*\| &\leq (1+r)\gamma[\tilde{L}\epsilon + 2\eta\delta]\|\mathbf{x}^0 - \mathbf{x}^*\| \\ &\leq r\|\mathbf{x}^0 - \mathbf{x}^*\|, \end{aligned}$$

where the second inequality follows from the second condition in (51). Without loss of generality, assume that $i_0 = 1$. Then, based on the result in (26) we obtain

$$(55) \quad \begin{aligned} \|\mathbf{B}_1^1 - \nabla^2 f_1(\mathbf{x}^*)\|_{\mathbf{M}_i} &\leq \left[(1 - \alpha\theta_1^{0^2})^{1/2} + \alpha_3\sigma_1^0 \right] \|\mathbf{B}_1^0 - \nabla^2 f_1(\mathbf{x}^*)\|_{\mathbf{M}_i} + \alpha_4\sigma_1^0 \\ &\leq (1 + \alpha_3\epsilon)\delta + \alpha_4\epsilon \\ &\leq \delta + 2\alpha_3\epsilon\delta + \alpha_4\epsilon \leq 2\delta. \end{aligned}$$

We proceed to the next iteration, which leads to the inequality

$$(56) \quad \begin{aligned} \|\mathbf{x}^2 - \mathbf{x}^*\| &\leq (1+r)\gamma \left[\frac{\tilde{L}}{n} \sum_{i=1}^n \|\mathbf{z}_i^1 - \mathbf{x}^*\|^2 + \frac{1}{n} \sum_{i=1}^n \|\mathbf{B}_i^1 - \nabla^2 f_i(\mathbf{x}^*)\| \|\mathbf{z}_i^1 - \mathbf{x}^*\| \right] \\ &\leq (1+r)\gamma \left[\tilde{L}\epsilon + 2\eta\delta \right] \left(\frac{n-1}{n} \|\mathbf{x}^0 - \mathbf{x}^*\| + \frac{1}{n} \|\mathbf{x}^1 - \mathbf{x}^*\| \right) \\ &\leq r \left(\frac{n-1}{n} \|\mathbf{x}^0 - \mathbf{x}^*\| + \frac{1}{n} \|\mathbf{x}^1 - \mathbf{x}^*\| \right) \\ &\leq r\|\mathbf{x}^0 - \mathbf{x}^*\|. \end{aligned}$$

And since the updated index is $i_1 = 2$ we obtain

$$(57) \quad \begin{aligned} \|\mathbf{B}_2^2 - \nabla^2 f_2(\mathbf{x}^*)\|_{\mathbf{M}_i} &\leq \left[(1 - \alpha\theta_2^{0^2})^{1/2} + \alpha_3\sigma_2^0 \right] \|\mathbf{B}_2^0 - \nabla^2 f_2(\mathbf{x}^*)\|_{\mathbf{M}_i} + \alpha_4\sigma_2^0 \\ &\leq (1 + \alpha_3\epsilon)\delta + \alpha_4\epsilon \\ &\leq \delta + 2\alpha_3\epsilon\delta + \alpha_4\epsilon \leq 2\delta. \end{aligned}$$

With the same argument we can show that $\|\mathbf{B}_t^t - \nabla^2 f_t(\mathbf{x}^*)\|_{\mathbf{M}_i} \leq 2\delta$ and $\|\mathbf{x}^t - \mathbf{x}^*\| \leq \epsilon$ for all iterates $t = 1, \dots, n$. Moreover, we have $\|\mathbf{x}^t - \mathbf{x}^*\| \leq r\|\mathbf{x}^0 - \mathbf{x}^*\|$ for $t = 1, \dots, n$.

Now we use the results for iterates $t = 1, \dots, n$ as the base of our induction argument. To be more precise, let's assume that for iterates $t = jn+1, jn+2, \dots, jn+$

n we know that the residuals are bounded above by $\|\mathbf{x}^t - \mathbf{x}^*\| \leq r^{j+1}\|\mathbf{x}^0 - \mathbf{x}^*\|$ and the Hessian approximation matrices \mathbf{B}_i^t satisfy the inequalities $\|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\| \leq 2\eta\delta$. Our goal is to show that for iterates $t = (j + 1)n + 1, (j + 1)n + 2, \dots, (j + 1)n + n$ the inequalities $\|\mathbf{x}^t - \mathbf{x}^*\| \leq r^{j+2}\|\mathbf{x}^0 - \mathbf{x}^*\|$ and $\|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\| \leq 2\eta\delta$ hold.

Based on the inequalities $\|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\| \leq 2\eta\delta$ and $\|\nabla^2 f_i(\mathbf{x}^*)^{-1}\| \leq \gamma$ we can show that for all $t = jn + 1, jn + 2, \dots, jn + n$ we have

$$(58) \quad \left\| \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i^t \right)^{-1} \right\| \leq (1 + r)\gamma.$$

Using (58) and the inequality in (26) for the iterate $t = (j + 1)n + 1$, we obtain

$$(59) \quad \begin{aligned} \|\mathbf{x}^{(j+1)n+1} - \mathbf{x}^*\| &\leq (1 + r)\gamma \frac{\tilde{L}}{n} \sum_{i=1}^n \left\| \mathbf{z}_i^{(j+1)n} - \mathbf{x}^* \right\|^2 \\ &+ (1 + r)\gamma \frac{1}{n} \sum_{i=1}^n \left\| \left[\mathbf{B}_i^{(j+1)n} - \nabla^2 f_i(\mathbf{x}^*) \right] \left(\mathbf{z}_i^{(j+1)n} - \mathbf{x}^* \right) \right\|. \end{aligned}$$

Since the variables are updated in a cyclic fashion, the set of variables $\{\mathbf{z}_i^{(j+1)n}\}_{i=1}^n$ is equal to the set $\{\mathbf{x}^{(j+1)n-i}\}_{i=0}^{i=n-1}$. By considering this relation and replacing the norms $\|\left[\mathbf{B}_i^{(j+1)n} - \nabla^2 f_i(\mathbf{x}^*) \right] \left(\mathbf{z}_i^{(j+1)n} - \mathbf{x}^* \right)\|$ by their upper bounds $2\eta\delta\|\mathbf{z}_i^{(j+1)n} - \mathbf{x}^*\|$ we can simplify the right-hand side of (59) as

$$(60) \quad \|\mathbf{x}^{(j+1)n+1} - \mathbf{x}^*\| \leq (1 + r)\gamma \left[\frac{\tilde{L}}{n} \sum_{i=1}^n \|\mathbf{x}^{jn+i} - \mathbf{x}^*\|^2 + \frac{2\eta\delta}{n} \sum_{i=1}^n \|\mathbf{x}^{jn+i} - \mathbf{x}^*\| \right].$$

Since $\|\mathbf{x}^{jn+i} - \mathbf{x}^*\| \leq \epsilon$ for all $j = 1, \dots, n$, we obtain

$$(61) \quad \|\mathbf{x}^{(j+1)n+1} - \mathbf{x}^*\| \leq (1 + r)\gamma \left[\tilde{L}\epsilon + 2\eta\delta \right] \left(\frac{1}{n} \sum_{i=1}^n \|\mathbf{x}^{jn+i} - \mathbf{x}^*\| \right).$$

According to the second inequality in (51) and the assumption that for iterates $t = jn + 1, jn + 2, \dots, jn + n$ we know that $\|\mathbf{x}^t - \mathbf{x}^*\| \leq r^{j+1}\|\mathbf{x}^0 - \mathbf{x}^*\|$, we can replace the right-hand side of (61) by the following upper bound:

$$(62) \quad \|\mathbf{x}^{(j+1)n+1} - \mathbf{x}^*\| \leq r^{j+2}\|\mathbf{x}^0 - \mathbf{x}^*\|.$$

Now we show that the updated Hessian approximation $\mathbf{B}_{i_t}^{(j+1)n+1}$ for $t = (j + 1)n + 1$ satisfies the inequality $\|\mathbf{B}_{i_t}^{(j+1)n+1} - \nabla^2 f_{i_t}(\mathbf{x}^*)\|_{\mathbf{M}_i} \leq 2\delta$. According to the result in (26), we can write

$$(63) \quad \begin{aligned} &\left\| \mathbf{B}_{i_t}^{(j+1)n+1} - \nabla^2 f_{i_t}(\mathbf{x}^*) \right\|_{\mathbf{M}_i} - \left\| \mathbf{B}_{i_t}^{jn+1} - \nabla^2 f_{i_t}(\mathbf{x}^*) \right\|_{\mathbf{M}_i} \\ &\leq \alpha_3 \sigma_{i_t}^{jn+1} \left\| \mathbf{B}_{i_t}^{jn+1} - \nabla^2 f_{i_t}(\mathbf{x}^*) \right\|_{\mathbf{M}_i} + \alpha_4 \sigma_{i_t}^{jn+1}. \end{aligned}$$

Now observe that $\sigma_{i_t}^{jn+1} = \max\{\|\mathbf{x}^{(j+1)n+1} - \mathbf{x}^*\|, \|\mathbf{x}^{jn+1} - \mathbf{x}^*\|\}$ is bounded above by $r^{j+1}\|\mathbf{x}^0 - \mathbf{x}^*\|$. Substituting this into (63) and considering the conditions

$$\|\mathbf{B}_{i_t}^{jn+1} - \nabla^2 f_{i_t}(\mathbf{x}^*)\|_{\mathbf{M}_i} \leq 2\delta \quad \text{and} \quad \|\mathbf{x}^0 - \mathbf{x}^*\| \leq \epsilon$$

leads to the inequality

$$(64) \quad \left\| \mathbf{B}_{i_t}^{(j+1)n+1} - \nabla^2 f_{i_t}(\mathbf{x}^*) \right\|_{\mathbf{M}_i} - \left\| \mathbf{B}_{i_t}^{jn+1} - \nabla^2 f_{i_t}(\mathbf{x}^*) \right\|_{\mathbf{M}_i} \leq r^{j+1} \epsilon (2\delta\alpha_3 + \alpha_4).$$

By writing the expression in (64) for previous iterations and using recursive logic we obtain that

$$(65) \quad \left\| \mathbf{B}_{i_t}^{(j+1)n+1} - \nabla^2 f_{i_t}(\mathbf{x}^*) \right\|_{\mathbf{M}_i} - \left\| \mathbf{B}_{i_t}^0 - \nabla^2 f_{i_t}(\mathbf{x}^*) \right\|_{\mathbf{M}_i} \leq \epsilon (2\delta\alpha_3 + \alpha_4) \frac{1}{1-r}.$$

Based on the first inequality in (51), the right-hand side of (65) is bounded above by δ . Moreover, the norm $\left\| \mathbf{B}_{i_t}^0 - \nabla^2 f_{i_t}(\mathbf{x}^*) \right\|_{\mathbf{M}_i}$ is also upper bounded by δ . These two bounds imply that

$$(66) \quad \left\| \mathbf{B}_{i_t}^{(j+1)n+1} - \nabla^2 f_{i_t}(\mathbf{x}^*) \right\|_{\mathbf{M}_i} \leq 2\delta,$$

and consequently $\left\| \mathbf{B}_{i_t}^{(j+1)n+1} - \nabla^2 f_{i_t}(\mathbf{x}^*) \right\| \leq 2\eta\delta$. By following the steps from (59) to (66), we can show for all iterates $t = (j+1)n+1, (j+1)n+2, \dots, (j+1)n+n$ that the inequalities $\left\| \mathbf{x}^t - \mathbf{x}^* \right\| \leq r^{j+2} \left\| \mathbf{x}^0 - \mathbf{x}^* \right\|$ and $\left\| \mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*) \right\| \leq 2\eta\delta$ hold. The induction proof is complete and (32) holds. Moreover, the inequality $\left\| \mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*) \right\| \leq 2\eta\delta$ holds for all i and steps t . Hence, the norms $\left\| \mathbf{B}_i^t \right\|$ and $\left\| (\mathbf{B}_i^t)^{-1} \right\|$, and consequently $\left\| (1/n) \sum_{i=1}^n \mathbf{B}_i^t \right\|$ and $\left\| ((1/n) \sum_{i=1}^n \mathbf{B}_i^t)^{-1} \right\|$, are uniformly bounded.

Appendix D. Proof of Proposition 4. According to the result in Lemma 3, we can show that the sequence of errors $\sigma_i^t = \max\{\left\| \mathbf{z}_i^{t+1} - \mathbf{x}^* \right\|, \left\| \mathbf{z}_i^t - \mathbf{x}^* \right\|\}$ is summable for all i . To do so, consider the sum of the sequence σ_i^t , which is upper bounded by

$$(67) \quad \sum_{t=0}^{\infty} \sigma_i^t = \sum_{t=0}^{\infty} \max\{\left\| \mathbf{z}_i^{t+1} - \mathbf{x}^* \right\|, \left\| \mathbf{z}_i^t - \mathbf{x}^* \right\|\} \leq \sum_{t=0}^{\infty} \left\| \mathbf{z}_i^{t+1} - \mathbf{x}^* \right\| + \sum_{t=0}^{\infty} \left\| \mathbf{z}_i^t - \mathbf{x}^* \right\|.$$

Note that the last time that the index i is chosen before time t should be in the set $\{t-1, \dots, t-n\}$. This observation in association with the result in (32) implies that

$$(68) \quad \sum_{t=0}^{\infty} \sigma_i^t \leq 2 \sum_{t=0}^{\infty} r^{\lceil \frac{t-n-1}{n} \rceil + 1} \left\| \mathbf{x}^0 - \mathbf{x}^* \right\| = 2 \sum_{t=0}^{\infty} r^{\lceil \frac{t-1}{n} \rceil} \left\| \mathbf{x}^0 - \mathbf{x}^* \right\|.$$

Simplifying the sum on the right-hand side of (68) yields

$$(69) \quad \sum_{t=0}^{\infty} \sigma_i^t \leq \frac{2 \left\| \mathbf{x}^0 - \mathbf{x}^* \right\|}{r} + 2n \left\| \mathbf{x}^0 - \mathbf{x}^* \right\| \sum_{t=0}^{\infty} r^t < \infty.$$

Thus, the sequence σ_i^t is summable for all $i = 1, \dots, n$. To complete the proof we use the following result from [13, Lemma 3.3].

LEMMA 9. *Let $\{\phi^t\}$ and $\{\delta^t\}$ be sequences of nonnegative numbers such that*

$$(70) \quad \phi^{t+1} \leq (1 + \delta^t)\phi^t + \delta^t \quad \text{and} \quad \sum_{t=1}^{\infty} \delta^t < \infty.$$

Then, the sequence $\{\phi^t\}$ converges.

According to the result in Lemma 1, we can write

$$(71) \quad \|\mathbf{B}_i^{t+1} - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} \leq \left[(1 - \alpha\theta_i^{t^2})^{1/2} + \alpha_3\sigma_i^t \right] \|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} + \alpha_4\sigma_i^t,$$

where α , α_3 , and α_4 are some positive bounded constants. Therefore, if we define $\alpha_{\max} = \max\{\alpha_3, \alpha_4\}$ it can be shown that

$$(72) \quad \|\mathbf{B}_i^{t+1} - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} \leq [1 + \alpha_{\max}\sigma_i^t] \|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} + \alpha_{\max}\sigma_i^t.$$

Now if we set $\phi^t = \|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i}$ and $\delta^t = \alpha_{\max}\sigma_i^t$, according to the inequality in (72) the first condition in (70) holds, i.e., $\phi^{t+1} \leq (1 + \delta^t)\phi^t + \delta^t$. Moreover, based on (69), we can write $\sum_{t=1}^{\infty} \delta^t = \sum_{t=1}^{\infty} \alpha_{\max}\sigma_i^t < \infty$, which implies that the second condition in (70) also holds. Hence, the conditions in Lemma 9 are satisfied and we obtain that the sequence $\{\phi^t\} = \{\|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i}\}$ for $\mathbf{M}_i := \nabla^2 f_i(\mathbf{x}^*)^{-1/2}$ is convergent and the following limit exists:

$$(73) \quad \lim_{t \rightarrow \infty} \|\nabla^2 f_i(\mathbf{x}^*)^{-1/2} \mathbf{B}_i^t \nabla^2 f_i(\mathbf{x}^*)^{-1/2} - \mathbf{I}\|_{\mathbf{F}} = l,$$

where l is a nonnegative constant. Moreover, following the proof of Theorem 3.4 in [13] we can show that

$$(74) \quad \begin{aligned} \alpha(\theta_i^t)^2 \|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} &\leq \|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} - \|\mathbf{B}_i^{t+1} - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} \\ &+ \sigma_i^t(\alpha_3 \|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} + \alpha_4), \end{aligned}$$

and, therefore, summing both sides implies

$$(75) \quad \sum_{t=0}^{\infty} (\theta_i^t)^2 \|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} < \infty.$$

Replacing θ_i^t in (75) by its definition in (27) results in

$$(76) \quad \sum_{t=0}^{\infty} \frac{\|\mathbf{M}_i(\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*))\mathbf{s}_i^t\|^2}{\|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i} \|\mathbf{M}_i^{-1}\mathbf{s}_i^t\|^2} < \infty.$$

Since the norm $\|\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)\|_{\mathbf{M}_i}$ is upper bounded and the eigenvalues of the matrix $\mathbf{M}_i = \nabla^2 f_i(\mathbf{x}^*)^{-1/2}$ are uniformly lower and upper bounded, we conclude from the result in (76) that

$$(77) \quad \lim_{t \rightarrow \infty} \frac{\|(\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*))\mathbf{s}_i^t\|^2}{\|\mathbf{s}_i^t\|^2} = 0,$$

which yields the claim in (33).

Appendix E. Proof of Lemma 5. Consider the sets of variable variations $\mathcal{S}_1 = \{\mathbf{s}_i^{t+n\tau}\}_{\tau=0}^T$ and $\mathcal{S}_2 = \{\mathbf{s}_i^{t+n\tau}\}_{\tau=0}^{\infty}$. It is trivial to show that $\mathbf{z}_i^t - \mathbf{x}^*$ is in the span of the set \mathcal{S}_2 , since the sequences of variables \mathbf{x}^t and \mathbf{z}_i^t converge to \mathbf{x}^* and we can write $\mathbf{x}^* - \mathbf{z}_i^t = \sum_{\tau=0}^{\infty} \mathbf{s}_i^{t+n\tau}$. We proceed to show that the vector $\mathbf{z}_i^t - \mathbf{x}^*$ is also in the span of the set \mathcal{S}_1 when T is sufficiently large. To do so, we use a contradiction argument. Let's assume that the vector $\mathbf{z}_i^t - \mathbf{x}^*$ does not lie in the span of the set \mathcal{S}_1 , and, therefore, it can be decomposed as the sum of two non-zero vectors given by

$$(78) \quad \mathbf{z}_i^t - \mathbf{x}^* = \mathbf{v}_{\parallel}^t + \mathbf{v}_{\perp}^t,$$

where \mathbf{v}_{\parallel}^t lies in the span of \mathcal{S}_1 and \mathbf{v}_{\perp}^t is orthogonal to the span of \mathcal{S}_1 . Since we assume that $\mathbf{z}_i^t - \mathbf{x}^*$ does not lie in the span of \mathcal{S}_1 , we obtain that $\mathbf{z}_i^{t+nT} - \mathbf{x}^*$ also does not lie in this span, since $\mathbf{z}_i^{t+nT} - \mathbf{x}^*$ can be written as the sum $\mathbf{z}_i^{t+nT} - \mathbf{x}^* = \mathbf{z}_i^t - \mathbf{x}^* + \sum_{\tau=0}^{T-1} \mathbf{s}_i^{t+n\tau}$. These observations imply that we can also decompose the vector $\mathbf{z}_i^{t+nT} - \mathbf{x}^*$ as

$$(79) \quad \mathbf{z}_i^{t+nT} - \mathbf{x}^* = \mathbf{v}_{\parallel}^{t+nT} + \mathbf{v}_{\perp}^{t+nT},$$

where $\mathbf{v}_{\parallel}^{t+nT}$ lies in the span of \mathcal{S}_1 and $\mathbf{v}_{\perp}^{t+nT}$ is orthogonal to the span of \mathcal{S}_1 . Moreover, we obtain that $\mathbf{v}_{\perp}^{t+nT}$ is equal to \mathbf{v}_{\perp}^t , i.e.,

$$(80) \quad \mathbf{v}_{\perp}^{t+nT} = \mathbf{v}_{\perp}^t.$$

This is true since $\mathbf{z}_i^{t+nT} - \mathbf{x}^*$ can be written as the sum of $\mathbf{z}_i^t - \mathbf{x}^*$ and a group of vectors that lie in the span of \mathcal{S}_1 . We assume that the norm $\|\mathbf{v}_{\perp}^{t+nT}\| = \|\mathbf{v}_{\perp}^t\| = \epsilon$, where $\epsilon > 0$ is a strictly positive constant. According to the linear convergence of the sequence $\|\mathbf{x}^t - \mathbf{x}^*\|$ in Lemma 3 we know that

$$(81) \quad \|\mathbf{z}_i^{t+nT} - \mathbf{x}^*\| \leq r^{\lceil \frac{t+nT-1}{n} \rceil + 1} \|\mathbf{x}^0 - \mathbf{x}^*\| \leq r^T \|\mathbf{x}^0 - \mathbf{x}^*\|.$$

If we pick large enough T such that $r^T \|\mathbf{x}^0 - \mathbf{x}^*\| < \epsilon$, then we obtain $\|\mathbf{z}_i^{t+nT} - \mathbf{x}^*\| < \epsilon$, which contradicts the assumption $\|\mathbf{v}_{\perp}^t\| = \epsilon$. Thus, we obtain that the vector $\mathbf{z}_i^t - \mathbf{x}^*$ is also in the span of set \mathcal{S}_1 .

Since the vector $\mathbf{z}_i^t - \mathbf{x}^*$ is in the span of \mathcal{S}_1 , we can write the normalized vector $(\mathbf{z}_i^t - \mathbf{x}^*) / \|\mathbf{z}_i^t - \mathbf{x}^*\|$ as a linear combination of the set of normalized vectors $\{\mathbf{s}_i^{t+n\tau} / \|\mathbf{s}_i^{t+n\tau}\|\}_{\tau=0}^{T-1}$. This property allows us to write

$$(82) \quad \begin{aligned} \lim_{t \rightarrow \infty} \frac{\|(\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*))(\mathbf{z}_i^t - \mathbf{x}^*)\|}{\|\mathbf{z}_i^t - \mathbf{x}^*\|} &= \lim_{t \rightarrow \infty} \left\| \left(\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*) \right) \frac{(\mathbf{z}_i^t - \mathbf{x}^*)}{\|\mathbf{z}_i^t - \mathbf{x}^*\|} \right\| \\ &= \lim_{t \rightarrow \infty} \left\| \left(\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*) \right) \sum_{\tau=0}^{T-1} a_{\tau} \frac{\mathbf{s}_i^{t+n\tau}}{\|\mathbf{s}_i^{t+n\tau}\|} \right\|, \end{aligned}$$

where a_{τ} is the coefficient of the vector $\mathbf{s}_i^{t+n\tau}$ when we write $(\mathbf{z}_i^t - \mathbf{x}^*) / \|\mathbf{z}_i^t - \mathbf{x}^*\|$ as the linear combination of the normalized vectors $\{\mathbf{s}_i^{t+n\tau} / \|\mathbf{s}_i^{t+n\tau}\|\}_{\tau=0}^{T-1}$. Now since the index of the difference $\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)$ does not match with the descent directions $\mathbf{s}_i^t + n\tau$, we add and subtract the term $\mathbf{B}_i^{t+n\tau}$ to the expression $\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)$ and use the triangle inequality to write

$$(83) \quad \begin{aligned} &\lim_{t \rightarrow \infty} \frac{\|(\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*))(\mathbf{z}_i^t - \mathbf{x}^*)\|}{\|\mathbf{z}_i^t - \mathbf{x}^*\|} \\ &\leq \lim_{t \rightarrow \infty} \left\| \sum_{\tau=0}^{T-1} a_{\tau} \frac{(\mathbf{B}_i^{t+n\tau} - \nabla^2 f_i(\mathbf{x}^*))\mathbf{s}_i^{t+n\tau}}{\|\mathbf{s}_i^{t+n\tau}\|} \right\| + \left\| \sum_{\tau=0}^{T-1} a_{\tau} \frac{(\mathbf{B}_i^t - \mathbf{B}_i^{t+n\tau})\mathbf{s}_i^{t+n\tau}}{\|\mathbf{s}_i^{t+n\tau}\|} \right\|. \end{aligned}$$

We first simplify the first limit in the right-hand side of (83). Using the Cauchy-Schwarz inequality and the result in Proposition 4 we can write

$$(84) \quad \begin{aligned} \lim_{t \rightarrow \infty} \left\| \sum_{\tau=0}^{T-1} a_{\tau} \frac{(\mathbf{B}_i^{t+n\tau} - \nabla^2 f_i(\mathbf{x}^*))\mathbf{s}_i^{t+n\tau}}{\|\mathbf{s}_i^{t+n\tau}\|} \right\| &\leq \lim_{t \rightarrow \infty} \sum_{\tau=0}^{T-1} a_{\tau} \left\| \frac{(\mathbf{B}_i^{t+n\tau} - \nabla^2 f_i(\mathbf{x}^*))\mathbf{s}_i^{t+n\tau}}{\|\mathbf{s}_i^{t+n\tau}\|} \right\| \\ &= \sum_{\tau=0}^{T-1} a_{\tau} \lim_{t \rightarrow \infty} \left\| \frac{(\mathbf{B}_i^{t+n\tau} - \nabla^2 f_i(\mathbf{x}^*))\mathbf{s}_i^{t+n\tau}}{\|\mathbf{s}_i^{t+n\tau}\|} \right\| = 0. \end{aligned}$$

Based on the results in (83) and (84), to prove the claim in (34) it remains to show that

$$(85) \quad \lim_{t \rightarrow \infty} \left\| \sum_{\tau=0}^T a_\tau \frac{(\mathbf{B}_i^t - \mathbf{B}_i^{t+n\tau}) \mathbf{s}_i^{t+n\tau}}{\|\mathbf{s}_i^{t+n\tau}\|} \right\| = 0.$$

To reach this goal, we first study the limit of the difference between two consecutive update Hessian approximation matrices $\lim_{t \rightarrow \infty} \|\mathbf{B}_i^t - \mathbf{B}_i^{t+n}\|$. Note that if we set $\mathbf{A} = \mathbf{B}_i^t$ in (41), we obtain that

$$(86) \quad \|\mathbf{B}_i^{t+n} - \mathbf{B}_i^t\|_{\mathbf{M}_i} \leq \alpha_2 \frac{\|\mathbf{y}_i^t - \mathbf{B}_i^t \mathbf{s}_i^t\|}{\|\mathbf{M}_i^{-1} \mathbf{s}_i^t\|},$$

where $\mathbf{M}_i = (\nabla^2 f_i(\mathbf{x}^*))^{-1/2}$. By adding and subtracting the term $\nabla^2 f_i(\mathbf{x}^*) \mathbf{s}_i^t$ and using the result in (33), we can show that the difference $\|\mathbf{B}_i^{t+n} - \mathbf{B}_i^t\|_{\mathbf{M}_i}$ approaches zero asymptotically. In particular,

$$(87) \quad \begin{aligned} \lim_{t \rightarrow \infty} \|\mathbf{B}_i^{t+n} - \mathbf{B}_i^t\|_{\mathbf{M}_i} &\leq \alpha_2 \lim_{t \rightarrow \infty} \frac{\|\mathbf{y}_i^t - \mathbf{B}_i^t \mathbf{s}_i^t\|}{\|\mathbf{M}_i^{-1} \mathbf{s}_i^t\|} \\ &\leq \alpha_2 \lim_{t \rightarrow \infty} \frac{\|\mathbf{y}_i^t - \nabla^2 f_i(\mathbf{x}^*) \mathbf{s}_i^t\|}{\|\mathbf{M}_i^{-1} \mathbf{s}_i^t\|} + \alpha_2 \lim_{t \rightarrow \infty} \frac{\|(\nabla^2 f_i(\mathbf{x}^*) - \mathbf{B}_i^t) \mathbf{s}_i^t\|}{\|\mathbf{M}_i^{-1} \mathbf{s}_i^t\|}. \end{aligned}$$

Since $\|\mathbf{y}_i^t - \nabla^2 f_i(\mathbf{x}^*) \mathbf{s}_i^t\|$ is bounded above by $\tilde{L} \|\mathbf{s}_i^t\| \max\{\|\mathbf{z}_i^t - \mathbf{x}^*\|, \|\mathbf{z}_i^{t+1} - \mathbf{x}^*\|\}$ and the eigenvalues of the matrix \mathbf{M}_i are uniformly bounded we obtain that the first limit on the right-hand side of (87) converges to zero. Further, the result in (33) shows that the second limit on the right-hand side of (87) also converges to zero. Therefore,

$$(88) \quad \lim_{t \rightarrow \infty} \|\mathbf{B}_i^{t+n} - \mathbf{B}_i^t\|_{\mathbf{M}_i} = 0.$$

Following the same argument we can show that for any two consecutive Hessian approximation matrices the difference approaches zero asymptotically. Thus, we obtain

$$(89) \quad \begin{aligned} \lim_{t \rightarrow \infty} \|\mathbf{B}_i^t - \mathbf{B}_i^{t+n\tau}\|_{\mathbf{M}_i} &\leq \lim_{t \rightarrow \infty} \left\| \sum_{u=0}^{\tau-1} (\mathbf{B}_i^{t+nu} - \mathbf{B}_i^{t+n(u+1)}) \right\|_{\mathbf{M}_i} \\ &\leq \sum_{u=0}^{\tau-1} \lim_{t \rightarrow \infty} \|\mathbf{B}_i^{t+nu} - \mathbf{B}_i^{t+n(u+1)}\|_{\mathbf{M}_i} = 0. \end{aligned}$$

Observing the result in (89) we can show that

$$(90) \quad \begin{aligned} \lim_{t \rightarrow \infty} \left\| \sum_{\tau=0}^T a_\tau \frac{(\mathbf{B}_i^t - \mathbf{B}_i^{t+n\tau}) \mathbf{s}_i^{t+n\tau}}{\|\mathbf{s}_i^{t+n\tau}\|} \right\| &\leq \sum_{\tau=0}^T a_\tau \lim_{t \rightarrow \infty} \left\| \frac{(\mathbf{B}_i^t - \mathbf{B}_i^{t+n\tau}) \mathbf{s}_i^{t+n\tau}}{\|\mathbf{s}_i^{t+n\tau}\|} \right\| \\ &\leq \sum_{\tau=0}^T a_\tau \lim_{t \rightarrow \infty} \|\mathbf{B}_i^t - \mathbf{B}_i^{t+n\tau}\| = 0. \end{aligned}$$

Therefore, the result in (85) holds. The claim in (34) follows by combining the results in (83), (84), and (85).

Appendix F. Proof of Theorem 6. The result in Lemma 2 implies

$$(91) \quad \|\mathbf{x}^{t+1} - \mathbf{x}^*\| \leq \frac{\tilde{L}\Gamma^t}{n} \sum_{i=1}^n \|\mathbf{z}_i^t - \mathbf{x}^*\|^2 + \frac{\Gamma^t}{n} \sum_{i=1}^n \|(\mathbf{B}_i^t - \nabla^2 f_i(\mathbf{x}^*)) (\mathbf{z}_i^t - \mathbf{x}^*)\|.$$

Divide both sides of (91) by $(1/n) \sum_{i=1}^n \|z_i^t - x^*\|$ to obtain

$$(92) \quad \frac{\|x^{t+1} - x^*\|}{\frac{1}{n} \sum_{i=1}^n \|z_i^t - x^*\|} \leq \tilde{L} \Gamma^t \sum_{i=1}^n \frac{\|z_i^t - x^*\|^2}{\sum_{i=1}^n \|z_i^t - x^*\|} + \Gamma^t \sum_{i=1}^n \frac{\|(\mathbf{B}_i^t - \nabla^2 f_i(x^*)) (z_i^t - x^*)\|}{\sum_{i=1}^n \|z_i^t - x^*\|}.$$

Since the error $\|z_i^t - x^*\|$ is a lower bound for the sum of errors $\sum_{i=1}^n \|z_i^t - x^*\|$, we can replace $\|z_i^t - x^*\|$ for $\sum_{i=1}^n \|z_i^t - x^*\|$ in (92), which implies

$$(93) \quad \begin{aligned} \frac{\|x^{t+1} - x^*\|}{\frac{1}{n} \sum_{i=1}^n \|z_i^t - x^*\|} &\leq \tilde{L} \Gamma^t \sum_{i=1}^n \frac{\|z_i^t - x^*\|^2}{\|z_i^t - x^*\|} + \Gamma^t \sum_{i=1}^n \frac{\|(\mathbf{B}_i^t - \nabla^2 f_i(x^*)) (z_i^t - x^*)\|}{\|z_i^t - x^*\|} \\ &= \tilde{L} \Gamma^t \sum_{i=1}^n \|z_i^t - x^*\| + \Gamma^t \sum_{i=1}^n \frac{\|(\mathbf{B}_i^t - \nabla^2 f_i(x^*)) (z_i^t - x^*)\|}{\|z_i^t - x^*\|}. \end{aligned}$$

By considering the fact that Γ^t is bounded above and using the result in Lemma 5, if we compute the limit of both sides of (93), the following expression holds:

$$(94) \quad \lim_{t \rightarrow \infty} \frac{\|x^{t+1} - x^*\|}{\frac{1}{n} \sum_{i=1}^n \|z_i^t - x^*\|} = 0.$$

The result in (94) in association with the simplification for the sum $\sum_{i=1}^n \|z_i^t - x^*\| = \sum_{i=0}^{n-1} \|x^{t-i} - x^*\|$ leads to the claim in (35).

Appendix G. Proof of Theorem 7. Recall the definition of the sequence $\tilde{x}^t = \operatorname{argmax}_{u \in \{tn, \dots, tn+n-1\}} \{\|x^u - x^*\|\}$, which is a subsequence of the sequence $\{x^t\}_{t=0}^\infty$. Our goal is to show that this subsequence converges superlinearly to x^* , i.e., $\lim_{t \rightarrow \infty} \|\tilde{x}^{t+1} - x^*\| / \|\tilde{x}^t - x^*\| = 0$. To do so, first note that the result in Theorem 6 implies that

$$(95) \quad \lim_{t \rightarrow \infty} \frac{\|x^t - x^*\|}{\max\{\|x^{t-1} - x^*\|, \dots, \|x^{t-n} - x^*\|\}} = 0,$$

which follows from the inequality

$$\max\{\|x^{t-1} - x^*\|, \dots, \|x^{t-n} - x^*\|\} \geq \frac{1}{n} (\|x^{t-1} - x^*\| + \dots + \|x^{t-n} - x^*\|).$$

Based on the limit in (95), there exists a large enough t_0 such that for all $t \geq t_0$ the following inequality holds:

$$(96) \quad \|x^t - x^*\| < \max\{\|x^{t-1} - x^*\|, \dots, \|x^{t-n} - x^*\|\}.$$

Combining the inequality in (96) with the inequalities

$$\|x^{t-i} - x^*\| \leq \max\{\|x^{t-1} - x^*\|, \dots, \|x^{t-n} - x^*\|\}$$

for $i = 1, \dots, n-1$ yields

$$(97) \quad \max\{\|x^t - x^*\|, \dots, \|x^{t-n+1} - x^*\|\} \leq \max\{\|x^{t-1} - x^*\|, \dots, \|x^{t-n} - x^*\|\},$$

and consequently we can generalize this result to obtain

$$(98) \quad \max\{\|x^t - x^*\|, \dots, \|x^{t-n+1} - x^*\|\} \leq \max\{\|x^{t-\tau} - x^*\|, \dots, \|x^{t-\tau-n+1} - x^*\|\}$$

for any positive integer τ such that $t - \tau \geq t_0$.

We use the result in (98) to build a superlinearly convergent subsequence of the residuals sequence $\|\mathbf{x}^t - \mathbf{x}^*\|$. If we define $\mathbf{x}^{tn+u_i^*}$ as the iterate that has the largest error among the iterates in the $t + 1$ th pass, i.e.,

$$(99) \quad \mathbf{x}^{tn+u_i^*} = \operatorname{argmax}_{u \in \{tn, \dots, tn+n-1\}} \{\|\mathbf{x}^u - \mathbf{x}^*\|\},$$

then it follows that $\tilde{\mathbf{x}}^t = \mathbf{x}^{tn+u_i^*}$, where $u_i^* \in \{0, 1, \dots, n-1\}$. Moreover, we obtain

$$(100) \quad \begin{aligned} \frac{\|\tilde{\mathbf{x}}^{t+1} - \mathbf{x}^*\|}{\|\tilde{\mathbf{x}}^t - \mathbf{x}^*\|} &= \frac{\|\mathbf{x}^{(t+1)n+u_{i+1}^*} - \mathbf{x}^*\|}{\max\{\|\mathbf{x}^{tn} - \mathbf{x}^*\|, \dots, \|\mathbf{x}^{tn+n-1} - \mathbf{x}^*\|\}} \\ &\leq \frac{\|\mathbf{x}^{tn+n+u_{i+1}^*} - \mathbf{x}^*\|}{\max\{\|\mathbf{x}^{tn+u_{i+1}^*-1} - \mathbf{x}^*\|, \dots, \|\mathbf{x}^{tn+n+u_{i+1}^*-1} - \mathbf{x}^*\|\}}. \end{aligned}$$

The equality follows from the definition of the iterate $\tilde{\mathbf{x}}^t$ and the definition in (99), and the inequality holds because of the result in (98). Considering the result in (95), computing the limit of both sides leads to the conclusion that the sequence $\|\tilde{\mathbf{x}}^t - \mathbf{x}^*\|$ is superlinearly convergent. In other words, we obtain that the subsequence $\{\|\mathbf{x}^{tn+u_i^*} - \mathbf{x}^*\|\}_{i=0}^{t=\infty}$ superlinearly converges to zero.

Let's define the sequence q^t such that $q^{kn} = \dots = q^{kn+n-1} = \|\tilde{\mathbf{x}}^k - \mathbf{x}^*\|$ for $k = 0, 1, 2, \dots$, which means that the value of the sequence q^t is fixed for each pass and is equal to the maximum error of the corresponding pass. Therefore, it is trivial to show that q^t is always larger than or equal to $\|\mathbf{x}^t - \mathbf{x}^*\|$, i.e., $\|\mathbf{x}^t - \mathbf{x}^*\| \leq q^t$ for all $t \geq 0$. Now define the sequence ζ^t such that $\zeta^t = q^t$ for $t = 0, \dots, n-1$, and, for $t \geq n$,

$$(101) \quad \zeta^{kn+i} = q^{kn-1} \left(\frac{q^{kn+n-1}}{q^{kn-1}} \right)^{\frac{i+1}{n}} \quad \text{for } i = 0, \dots, n-1, k \geq 1.$$

According to this definition we can verify that ζ^t is an upper bound for the sequence q^t and, consequently, an upper bound for the sequence of errors $\|\mathbf{x}^t - \mathbf{x}^*\|$. Based on the definition of the sequence ζ^t in (101), the ratio ζ^{t+1}/ζ^t is given by $(q^{\lfloor \frac{t+1}{n} \rfloor n+n-1} / q^{\lfloor \frac{t+1}{n} \rfloor n-1})^{1/n}$. This simplification in association with the definitions of the sequences $\|\mathbf{x}^t - \mathbf{x}^*\|$ and $\|\tilde{\mathbf{x}}^t - \mathbf{x}^*\|$ implies that

$$(102) \quad \lim_{t \rightarrow \infty} \frac{\zeta^{t+1}}{\zeta^t} = \lim_{t \rightarrow \infty} \left(\frac{q^{\lfloor \frac{t+1}{n} \rfloor n+n-1}}{q^{\lfloor \frac{t+1}{n} \rfloor n-1}} \right)^{\frac{1}{n}} = \lim_{t \rightarrow \infty} \left(\frac{\|\tilde{\mathbf{x}}^{\lfloor \frac{t+1}{n} \rfloor} - \mathbf{x}^*\|}{\|\tilde{\mathbf{x}}^{\lfloor \frac{t+1}{n} \rfloor-1} - \mathbf{x}^*\|} \right)^{\frac{1}{n}} = 0,$$

which leads to the claim in (37).

REFERENCES

[1] D. BLATT, A. O. HERO, AND H. GAUCHMAN, *A convergent incremental gradient method with a constant step size*, SIAM J. Optim., 18 (2007), pp. 29–51.
 [2] L. BOTTOU, *Large-scale machine learning with stochastic gradient descent*, in Proceedings of COMPSTAT'2010: The 19th International Conference on Computational Statistics, Y. Lechevallier and G. Saporta, eds., Springer, Berlin, 2010, pp. 177–186.
 [3] L. BOTTOU AND Y. LE CUN, *On-line learning for very large datasets*, Statist. Learn., 21 (2005), pp. 137–151.
 [4] C. G. BROYDEN, J. E. DENNIS, JR., AND J. J. MORÉ, *On the local and superlinear convergence of quasi-Newton methods*, IMA J. Appl. Math., 12 (1973), pp. 223–245.

- [5] F. BULLO, J. CORTÉS, AND S. MARTÍNEZ, *Distributed Control of Robotic Networks: A Mathematical Approach to Motion Coordination Algorithms*, Princeton Ser. Appl. Math., Princeton University Press, Princeton, NJ, 2009.
- [6] R. H. BYRD, S. L. HANSEN, J. NOCEDAL, AND Y. SINGER, *A stochastic quasi-Newton method for large-scale optimization*, SIAM J. Optim., 26 (2016), pp. 1008–1031.
- [7] Y. CAO, W. YU, W. REN, AND G. CHEN, *An overview of recent progress in the study of distributed multi-agent coordination*, IEEE Trans. Ind. Inform., 9 (2013), pp. 427–438.
- [8] V. CEVHER, S. BECKER, AND M. SCHMIDT, *Convex optimization for big data: Scalable, randomized, and parallel algorithms for big data analytics*, IEEE Signal Process. Mag., 31 (2014), pp. 32–43.
- [9] A. DEFAZIO, F. R. BACH, AND S. LACOSTE-JULIEN, *SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives*, in Adv. Neural Inf. Process. Syst. 27, MIT Press, Cambridge, MA, 2014, pp. 1646–1654.
- [10] A. DEFAZIO, J. DOMKE, AND T. S. CAETANO, *Finito: A faster, permutable incremental gradient method for big data problems*, Proceedings of the 31st International Conference on Machine Learning, Proc. Mach. Learn. Res. 32, 2014, pp. 1125–1133; available at <http://proceedings.mlr.press/v32/defazio14.html>.
- [11] R. M. GOWER, D. GOLDFARB, AND P. RICHTÁRIK, *Stochastic Block BFGS: Squeezing More Curvature Out of Data*, preprint, arXiv:1603.09649, 2016.
- [12] M. GÜRBÜZBALABAN, A. OZDAGLAR, AND P. PARRILO, *On the Convergence Rate of Incremental Aggregated Gradient Algorithms*, preprint, arXiv:1506.02081, 2015.
- [13] J. E. DENNIS, JR., AND J. J. MORE, *A characterization of superlinear convergence and its application to quasi-Newton methods*, Math. Comp., 28 (1974), pp. 549–560.
- [14] R. JOHNSON AND T. ZHANG, *Accelerating stochastic gradient descent using predictive variance reduction*, in Adv. Neural Inf. Process. Syst. 26, MIT Press, Cambridge, MA, 2013, pp. 315–323.
- [15] N. S. KESKAR AND A. S. BERAHAS, *adaQN: An adaptive quasi-Newton algorithm for training RNNs*, in Machine Learning and Knowledge Discovery in Databases, P. Frasconi, N. Landwehr, G. Manco, and J. Vreeken, eds., Springer, New York, 2016, pp. 1–16.
- [16] R. KOLTE, M. ERDOGDU, AND A. ÖZGÜR, *Accelerating SVRG via second-order information*, in NIPS Workshop on Optimization for Machine Learning, 2015, <https://erdogdu.github.io/papers/asvrg-opt.pdf>.
- [17] Y. LECUN, C. CORTES, AND C. J. C. BURGESS, *The MNIST Database of Handwritten Digits*, <http://yann.lecun.com/exdb/mnist>, 2010.
- [18] C. G. LOPES AND A. H. SAYED, *Diffusion least-mean squares over adaptive networks: Formulation and performance analysis*, IEEE Trans. Signal Process., 56 (2008), pp. 3122–3136.
- [19] A. LUCCHI, B. MCWILLIAMS, AND T. HOFMANN, *A variance reduced stochastic Newton method*, preprint, arXiv:1503.08316, 2015.
- [20] J. MAIRAL, *Incremental majorization-minimization optimization with application to large-scale machine learning*, SIAM J. Optim., 25 (2015), pp. 829–855.
- [21] A. MOKHTARI, M. GÜRBÜZBALABAN, AND A. RIBEIRO, *Surpassing Gradient Descent Provably: Acyclic Incremental Method with Linear Convergence Rate*, preprint, arXiv:1611.00347, 2016.
- [22] A. MOKHTARI AND A. RIBEIRO, *RES: Regularized stochastic BFGS algorithm*, IEEE Trans. Signal Process., 62 (2014), pp. 6089–6104.
- [23] A. MOKHTARI AND A. RIBEIRO, *Global convergence of online limited memory BFGS*, J. Mach. Learn. Res., 16 (2015), pp. 3151–3181.
- [24] P. MORITZ, R. NISHIHARA, AND M. I. JORDAN, *A linearly-convergent stochastic L-BFGS algorithm*, in Proceedings of the Nineteenth International Conference on Artificial Intelligence and Statistics, Proc. Mach. Learn. Res. 51, 2016, pp. 249–258; available at <http://proceedings.mlr.press/v51/moritz16.pdf>.
- [25] Y. NESTEROV, *Introductory Lectures on Convex Optimization: A Basic Course*, Springer Sci. Bus. Media 87, Springer, New York, 2013.
- [26] M. J. D. POWELL, *Some Global Convergence Properties of a Variable Metric Algorithm for Minimization Without Exact Line Search*, 2nd ed., Academic Press, London, 1971.
- [27] A. RIBEIRO, *Ergodic stochastic optimization algorithms for wireless communication and networking*, IEEE Trans. Signal Process., 58 (2010), pp. 6369–6386.
- [28] A. RIBEIRO, *Optimal resource allocation in wireless communication and networking*, EURASIP J. Wireless Commun. Netw., 2012 (2012), <https://doi.org/10.1186/1687-1499-2012-272>.

- [29] A. RODOMANOV AND D. KROPOTOV, *A superlinearly-convergent proximal Newton-type method for the optimization of finite sums*, in Proceedings of the International Conference on Machine Learning, Proc. Mach. Learn. Res. 48, 2016, pp. 2597–2605; available at <http://proceedings.mlr.press/v48/>.
- [30] N. LE ROUX, M. W. SCHMIDT, AND F. R. BACH, *A stochastic gradient method with an exponential convergence rate for finite training sets*, in Adv. Neural Inf. Process. Syst. 25, MIT Press, Cambridge, MA, 2012, pp. 2672–2680.
- [31] I. SCHIZAS, A. RIBEIRO, AND G. GIANNAKIS, *Consensus in ad hoc WSNs with noisy links—Part I: Distributed estimation of deterministic signals*, IEEE Trans. Signal Process., 56 (2008), pp. 350–364.
- [32] M. SCHMIDT, N. LE ROUX, AND F. BACH, *Minimizing finite sums with the stochastic average gradient*, Math. Program., 162 (2016), pp. 83–112, <http://dx.doi.org/10.1007/s10107-016-1030-6>.
- [33] N. N. SCHRAUDOLPH, J. YU, AND S. GÜNTER, *A stochastic quasi-Newton method for online convex optimization*, in Proceedings of the Eleventh International Conference on Artificial Intelligence and Statistics, Proc. Mach. Learn. Res. 2, 2007, pp. 436–443; available at <http://proceedings.mlr.press/v2>.
- [34] S. SHALEV-SHWARTZ AND N. SREBRO, *SVM optimization: Inverse dependence on training set size*, in Proceedings of the 25th Annual International Conference on Machine Learning (ICML 2008), ACM Publications, New York, 2008, pp. 928–935.
- [35] X. WANG, S. MA, D. GOLDFARB, AND W. LIU, *Stochastic quasi-Newton methods for nonconvex stochastic optimization*, SIAM J. Optim., 27 (2017), pp. 927–956.