Subsampled cubic regularization method with distinct sample sizes for function, gradient, and Hessian

Max L.N. Gonçalves[1]

[1]IME, Universidade Federal de Goiás, Rua Jacarandá, Goiânia, CEP 74001-970, GO, Brazil, Email: maxlng@ufg.br.

Abstract

We develop and study a subsampled cubic regularization method for finite-sum composite optimization problems, in which the function, gradient, and Hessian are estimated using possibly different sample sizes. By allowing each quantity to have its own sampling strategy, the proposed method offers greater flexibility to control the accuracy of the model components and to better balance computational effort and estimation quality. Such flexibility is particularly valuable in large-scale settings where the relative cost of evaluating these quantities can vary significantly. We establish iteration-complexity bounds for computing approximate first-order critical points and prove global convergence properties. In addition, we present numerical experiments that illustrate the practical performance of the proposed method.

Keywords: Cubic regularization method, Subsampled technique, Composite optimization problem, Iteration-complexity.

1 Introduction

In this paper, we focus on finite-sum composite optimization problems characterized by the formulation

$$\min_{x\in\mathbb{R}^n} F(x) := f(x) + h(x), \quad \text{where} \quad f(x) := \frac{1}{N}\sum_{i=1}^{N} f_i(x), \tag{1}$$

in which each function $f_i : \mathbb{R}^n \to \mathbb{R}$ is twice continuously differentiable, possibly nonconvex, and $h : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ is a proper, closed, and convex function that may be nondifferentiable. We assume that there exists a constant F_{low} such that $F(x) \geq F_{\text{low}}$ for all $x \in Q := \text{dom}(h)$. We are interested in the large-scale regime in which the number of component functions N is very large. A point $\bar{x} \in Q$ is said to be a first-order stationary (or critical) point of (1) if $0 = \nabla f(\bar{x}) + h'(\bar{x})$ for some $h'(\bar{x}) \in \partial h(\bar{x})$, where ∂h denotes the subdifferential of h. Problem (1) captures a wide range of modern optimization models, including empirical risk minimization, regularized regression, and constrained optimization. In these settings, $f_i(x)$ typically represents the loss associated with a single data sample, while $h(x)$ encodes structure-promoting regularization such as sparsity, smoothness, or feasibility constraints. This formulation has become central to large-scale learning and signal recovery applications, motivating the development of algorithms that can efficiently exploit the finite-sum structure.

A variety of first- and second-order methods have been proposed for solving problems of the form (1). Classical full-batch methods, while theoretically well understood, are often computationally prohibitive when N is large. To address this issue, subsampling techniques have emerged as an effective way to approximate function, gradient, or Hessian information using only a subset of the component functions at each iteration. This strategy significantly reduces the computational burden while retaining desirable convergence properties; see, for instance, [4–8, 10–12, 17, 19–21].

Among second-order approaches, Cubic Regularization Methods (CRMs) have received considerable attention since the seminal work of Nesterov and Polyak [18], who established their global convergence and optimal iteration-complexity properties for nonconvex smooth optimization. Subsequent research extended these methods to composite and stochastic settings, demonstrating their robustness and practical efficiency. In particular, subsampled or inexact variants of CRMs have proven well-suited to large-scale problems, where exact second-order information is unavailable or too costly to compute; see, for example, [4–6, 10, 11, 13, 15, 17, 19, 21].

Here, we develop and analyze a subsampled CRM specifically designed for problem (1). A distinctive feature of the proposed method is that it allows for different sample sizes to be employed in estimating the function value, the gradient, and the Hessian. In particular, this represents a significant extension of the work [10], where the function and gradient shared the same sample while the Hessian was estimated using an independent, possibly smaller one. By allowing each quantity to have its own sampling strategy, the present framework offers greater flexibility to control the accuracy of the model components and to better balance computational effort and estimation quality. Such flexibility is particularly valuable in large-scale settings where the relative cost of evaluating these quantities can vary significantly. For instance, in many applications, function evaluations can be inexpensive while gradient or Hessian computations are costly, whereas in others, the opposite may hold. By decoupling the sampling strategies, the proposed CRM can be effectively adapted to such heterogeneous computational environments. It is worth noting that, in contrast to work [10], which focused on the smooth case $h \equiv 0$, the proposed method is designed to handle the composite formulation (1), thereby covering both unconstrained problems and those involving simple convex constraints or additive regularization terms. In addition to introducing this new method, we establish iteration-complexity guarantees for computing approximate critical points and discuss global convergence properties. We also present numerical experiments showing that the proposed approach not only preserves the favorable theoretical properties of CRMs but also delivers significant computational savings in practice.

The rest of the paper is organized as follows. Section 2 describes the proposed subsampled CRM and presents its theoretical results, whose proofs are postponed to Subsection 2.1. Section 3 reports numerical results demonstrating the practical performance of the method. Concluding remarks are given in Section 4.

2 The method and its convergence analysis

In this section, we formally describe the proposed Subsampled Cubic Regularization Method (S-CRM) for computing approximate stationary points of problem (1), and present its iteration-complexity analysis and global convergence results. Detailed proofs are deferred to Subsection 2.1.

To reduce the overall computational cost associated with solving (1), we allow the function f, its gradient, and its Hessian to be approximated using distinct subsamples of the finite-sum components. Specifically, we define

$$f_{\mathcal{D}}(x) := \frac{1}{|\mathcal{D}|}\sum_{j\in\mathcal{D}} f_j(x),\ \nabla f_{\mathcal{G}}(x) := \frac{1}{|\mathcal{G}|}\sum_{j\in\mathcal{G}} \nabla f_j(x),\ \nabla^2 f_{\mathcal{H}}(x) := \frac{1}{|\mathcal{H}|}\sum_{j\in\mathcal{H}} \nabla^2 f_j(x), \tag{2}$$

where $\mathcal{D},\mathcal{G},\mathcal{H}\subset\{1,\ldots,N\}$ are the corresponding subsample index sets, and $|\mathcal{D}|$, $|\mathcal{G}|$, and $|\mathcal{H}|$ denote their cardinalities.

We next present the formal description of the proposed algorithm.

Algorithm 1. ***(S-CRM)***

S0. Choose an initial point $x_0\in Q$ and parameters $\theta\geq 0$, $\sigma_0>0$, $c>6$, and growth factors $\gamma_{\mathcal{D}^1},\gamma_{\mathcal{D}^2},\gamma_{\mathcal{G}},\gamma_{\mathcal{H}}\geq\alpha>1$. Initialize the subsamples $\tilde{\mathcal{D}}_0^1,\tilde{\mathcal{D}}_0^2,\tilde{\mathcal{G}}_0,\tilde{\mathcal{H}}_0\subset\mathcal{F}_0\subset\{1,\ldots,N\}$, and set $t:=0$.

S1. Find the smallest integer $i\geq 0$ such that $\alpha^{i-1}\sigma_t\geq\sigma_0$. Select subsamples $\mathcal{D}_{t,i}^1,\mathcal{D}_{t,i}^2,\mathcal{G}_{t,i},\mathcal{H}_{t,i}\subset\mathcal{F}_t$ satisfying $|\mathcal{D}_{t,i}^1|\geq|\tilde{\mathcal{D}}_0^1|$, $|\mathcal{D}_{t,i}^2|\geq|\tilde{\mathcal{D}}_0^2|$, $|\mathcal{G}_{t,i}|\geq|\tilde{\mathcal{G}}_0|$ and $|\mathcal{H}_{t,i}|\geq|\tilde{\mathcal{H}}_0|$.

S1.1. Compute $\nabla f_{\mathcal{G}_{t,i}}(x_t)$ and $\nabla^2 f_{\mathcal{H}_{t,i}}(x_t)$ as in (2), and consider the cubic model

$$M_{x_t,\alpha^i\sigma_t}^{\mathcal{G}_{t,i},\mathcal{H}_{t,i}}(y) := \langle\nabla f_{\mathcal{G}_{t,i}}(x_t), y-x_t\rangle + \frac{1}{2}\langle\nabla^2 f_{\mathcal{H}_{t,i}}(x_t)(y-x_t), y-x_t\rangle + \frac{\alpha^i\sigma_t}{6}\|y-x_t\|^3. \tag{3}$$

Compute an approximate solution $x_{t,i}^+$ of the subproblem

$$\min_{y\in\mathbb{R}^n} M_{x_t,\alpha^i\sigma_t}^{\mathcal{G}_{t,i},\mathcal{H}_{t,i}}(y) + h(y), \tag{4}$$

such that

$$M_{x_t,\alpha^i\sigma_t}^{\mathcal{G}_{t,i},\mathcal{H}_{t,i}}(x_{t,i}^+) + h(x_{t,i}^+) \le h(x_t), \quad \|\nabla M_{x_t,\alpha^i\sigma_t}^{\mathcal{G}_{t,i},\mathcal{H}_{t,i}}(x_{t,i}^+) + h'(x_{t,i}^+)\| \le \theta\|x_{t,i}^+ - x_t\|^2, \tag{5}$$

for some $h'(x_{t,i}^+) \in \partial h(x_{t,i}^+)$.

S1.2. Compute $f_{\mathcal{D}_{t,i}^1}(x_t)$, $f_{\mathcal{D}_{t,i}^2}(x_{t,i}^+)$ and $\nabla f_{\mathcal{G}_{t,i}}(x_{t,i}^+)$ as in (2). If

$$(f_{\mathcal{D}_{t,i}^1} + h)(x_t) - (f_{\mathcal{D}_{t,i}^2} + h)(x_{t,i}^+) \ge \frac{\alpha^i\sigma_t}{c}\|x_{t,i}^+ - x_t\|^3 \tag{6}$$

and

$$\|\nabla f_{\mathcal{G}_{t,i}}(x_{t,i}^+) + h'(x_{t,i}^+)\| \le \left(\frac{(c-3)}{c}\alpha^i\sigma_t + \sigma_0 + \theta\right)\|x_{t,i}^+ - x_t\|^2 \tag{7}$$

hold, set $i_t = i$, $\mathcal{D}_t^1 = \mathcal{D}_{t,i_t}^1$, $\mathcal{D}_t^2 = \mathcal{D}_{t,i_t}^2$, $\mathcal{G}_t = \mathcal{G}_{t,i_t}$, $\mathcal{H}_t = \mathcal{H}_{t,i_t}$, and prooced to **S2**. Otherwise, enlarge the subsamples by computing: $\mathcal{D}_{t,i+1}^1, \mathcal{D}_{t,i+1}^2, \mathcal{G}_{t,i+1}, \mathcal{H}_{t,i+1} \subset \mathcal{F}_t$ such that

$$|\mathcal{D}_{t,i+1}^1| = \min\{\lceil\gamma_{\mathcal{D}^1}^i\sigma_t\rceil|\mathcal{D}_{t,i}^1|, |\mathcal{F}_t|\}, \quad |\mathcal{D}_{t,i+1}^2| = \min\{\lceil\gamma_{\mathcal{D}^2}^i\sigma_t\rceil|\mathcal{D}_{t,i}^2|, |\mathcal{F}_t|\},$$

$$|\mathcal{G}_{t,i+1}| = \min\{\lceil\gamma_{\mathcal{G}}^i\sigma_t\rceil|\mathcal{G}_{t,i}|, |\mathcal{F}_t|\}, \quad |\mathcal{H}_{t,i+1}| = \min\{\lceil\gamma_{\mathcal{H}}^i\sigma_t\rceil|\mathcal{H}_{t,i}|, |\mathcal{F}_t|\},$$

set $i := i + 1$ and return to **S1.1**.

S2. Choose a new subsample $\mathcal{F}_t^+ \subset \{1, \ldots, N\}$ such that $|\mathcal{F}_t^+| \ge |\mathcal{F}_0|$. Update

$$x_{t+1} := x_{t,i_t}^+, \quad \mathcal{F}_{t+1} := \mathcal{F}_t^+, \quad \sigma_{t+1} := \alpha^{i_t-1}\sigma_t,$$

set $t := t + 1$, and return to **S1**.

Remark 1 (i) In Algorithm 1, the ordering among the subsamples $\mathcal{D}_{t,i+1}^1, \mathcal{D}_{t,i+1}^2, \mathcal{G}_{t,i+1}, \mathcal{H}_{t,i+1}$ is not fixed, which provides significant flexibility. For instance, one may consider $\mathcal{H}_{t,i+1} \subset \mathcal{G}_{t,i+1} \subset \mathcal{D}_{t,i+1}^1 \subset \mathcal{D}_{t,i+1}^2$ when gradient and Hessian evaluations are more expensive than function evaluations, or, conversely, $\mathcal{D}_{t,i+1}^1 \subset \mathcal{D}_{t,i+1}^2 \subset \mathcal{H}_{t,i+1} \subset \mathcal{G}_{t,i+1}$ when function evaluations are the dominant cost. This flexibility allows the proposed scheme to adapt to a wide range of computational environments with different relative costs for evaluating f_i, ∇f_i, and $\nabla^2 f_i$. In practice, one may fix the outer subsample $\mathcal{F}_t$ and progressively enlarge the others until all reach comparable cardinalities. Moreover, no specific update rule is imposed for the sequence $\{\mathcal{F}_t\}$ in Step S2. In particular, the subsample size $|\mathcal{F}_t|$ does not necessarily increase at every outer iteration, although it is expected that the full sample will be reached after finitely many iterations. For a detailed discussion of update strategies for subsample sequences, we refer the reader to [7, 8, 20]. (ii) Conditions in (5) imply that $x_{t,i}^+$ yields a decrease in the cubic model plus the function h, and that it is an approximate first-order stationary point of (4). These requirements are in line with those adopted in related cubic regularization frameworks (see, e.g., [9, 13]). The parameter $\theta \in (0, 1)$ controls the level of inexactness allowed in solving the cubic subproblem (3). Larger values of θ relax the condition in (5), reducing the computational effort per iteration, whereas smaller values enforce higher accuracy and may yield faster convergence. (iii) Note that the sequence of parameters $\{\sigma_t\}$ can be nonmonotone. Indeed, if $i_t = 0$, then $\sigma_{t+1} = 2^{i_t-1}\sigma_t = \sigma_t/2 \leq \sigma_t$.

We now discuss the iteration-complexity bounds and global convergence properties of Algorithm 1, with detailed proofs deferred to the next section. Throughout the analysis, we impose the following standard assumption:

(A1) The Hessian $\nabla^2 f_{\mathcal{H}}$ is L-Lipschitz continuous for every $\mathcal{H} \subset \{1,\dots,N\}$, that is,

$$\|\nabla^2 f_{\mathcal{H}}(y) - \nabla^2 f_{\mathcal{H}}(x)\| \leq L\|y-x\|, \quad \forall x, y \in \mathbb{R}^n.$$

We begin by presenting a bound in terms of outer iterations.

Theorem 1 *Let $\{x_t\}_{t=0}^T$ be generated by Algorithm 1 and define*

$$\nu_t := \max\{|f_{\mathcal{D}_t^2}(x_{t+1}) - f(x_{t+1})|, |f_{\mathcal{D}_t^1}(x_t) - f(x_t)|\}, \quad \forall t \geq 0, \tag{8}$$

and

$$\sigma_{\max} := \max\left\{\frac{c(L+6\sigma_0)}{c-6}, \frac{N}{\min\{|\tilde{\mathcal{D}}_0^1|, |\tilde{\mathcal{D}}_0^2|, |\tilde{\mathcal{G}}_0|, |\tilde{\mathcal{H}}_0|\}}\right\}. \tag{9}$$

Assume that

$$\lim_{T\to\infty} \frac{1}{T}\sum_{t=0}^{T-1} \nu_t = 0. \tag{10}$$

Given $\varepsilon \in (0,1)$, let $T_0(\varepsilon)$ be the smallest non-negative integer such that:

$$T \geq T_0(\varepsilon) \Longrightarrow \frac{1}{T}\sum_{t=0}^{T-1} \nu_t \leq \frac{\alpha\sigma_0 c^{1/2}\varepsilon^{\frac{3}{2}}}{4\left(\alpha(c-3)\sigma_{max} + c(\sigma_0+\theta)\right)^{\frac{3}{2}}}. \tag{11}$$

If

$$T \geq \max\left\{\frac{2(F(x_0) - F_{\text{low}})\left(\alpha(c-3)\sigma_{max} + c(\sigma_0+\theta)\right)^{\frac{3}{2}}}{\alpha\sigma_0 c^{1/2}}\varepsilon^{-\frac{3}{2}}, T_0(\varepsilon)\right\}, \tag{12}$$

then

$$\min_{t=0,\dots,T-1} \|\nabla f_{\mathcal{G}_t}(x_{t+1}) + h'(x_{t+1})\| \leq \varepsilon. \tag{13}$$

Remark 2 The assumption in (10) is weaker than requiring the sequence $\{\nu_t\}$ to be summable, a condition typically imposed in the iteration-complexity analyses of deterministic subsampled methods such as the subsampled inexact Newton and subsampled spectral gradient methods (see [7, 20]). Indeed, if $\{\nu_t\}$ is summable, then

$$0 \leq \frac{1}{T}\sum_{t=0}^{T-1} \nu_t \leq \frac{1}{T}\sum_{t=0}^{\infty} \nu_t,$$

and, by letting $T \to \infty$, condition (10) follows immediately. Conversely, as observed in [14], there exist sequences $\{\nu_t\}$ that satisfy (10) but are not summable. A relevant class of such examples consists of sequences for which $\nu_t \to 0$ as $t \to \infty$ (see [14, Corollary 2]).

If the sequence $\{\nu_t\}$ is assumed to be summable, then Theorem 1 implies an iteration-complexity bound of $\mathcal{O}\left(\left(N\varepsilon^{-1}/\min\{|\tilde{\mathcal{D}}_0^1|, |\tilde{\mathcal{D}}_0^2|, |\tilde{\mathcal{G}}_0|, |\tilde{\mathcal{H}}_0|\}\right)^{\frac{3}{2}}\right)$, in terms of outer iterations (or equivalently, calls of a certain oracle), for Algorithm 1 to produce an ε-approximate stationary point of (1).

Corollary 2 Let $\{x_t\}_{t=0}^T$ be generated by Algorithm 1. Assume that $\sum_{t=0}^{\infty}\nu_t < +\infty$, where ν_t is as in (8). Given $\varepsilon \in (0,1)$, if

$$
\begin{aligned}
T &\geq \frac{2\left(\alpha(c-3)\sigma_{max} + c(\sigma_0+\theta)\right)^{\frac{3}{2}}}{\alpha\sigma_0 c^{1/2}} \max\left\{F(x_0) - F_{\text{low}}, 2\sum_{t=0}^{\infty}\nu_t\right\}\varepsilon^{-\frac{3}{2}} \qquad (14)\\
&= \mathcal{O}\left(\left(\frac{N}{\min\{|\tilde{\mathcal{D}}_0^1|, |\tilde{\mathcal{D}}_0^2|, |\tilde{\mathcal{G}}_0|, |\tilde{\mathcal{H}}_0|\}}\right)^{\frac{3}{2}}\varepsilon^{-\frac{3}{2}}\right),
\end{aligned}
$$

then

$$
\min_{t=0,\ldots,T-1}\|\nabla f_{\mathcal{G}_t}(x_{t+1}) + h'(x_{t+1})\| \leq \varepsilon. \qquad (15)
$$

As a consequence, Algorithm 1 needs at most $\mathcal{O}\left(\left(N/\min\{|\tilde{\mathcal{D}}_0^1|, |\tilde{\mathcal{D}}_0^2|, |\tilde{\mathcal{G}}_0|, |\tilde{\mathcal{H}}_0|\}\right)^{\frac{3}{2}}\varepsilon^{-\frac{3}{2}}\right)$ calls of the oracle[1] to generate an iterate x_t such that $\|\nabla f_{\mathcal{G}_t}(x_t) + h'(x_{t+1})\| \leq \varepsilon$.

In the particular case where exact function, gradient, and Hessian evaluations are employed, that is, when $\mathcal{D}_t^1 = \mathcal{D}_t^2 = \mathcal{G}_t = \mathcal{H}_t = \{1,\ldots,N\}$ for all $t \geq 0$, the bound of $\mathcal{O}(\varepsilon^{-\frac{3}{2}})$ for the standard cubic regularization method follows directly from Corollary 2.

We conclude this section with a result showing that, under suitable conditions, all limit points of the sequence generated by Algorithm 1, if any, are stationary points of problem (1).

Corollary 3 Let $\{x_t\}_{t=0}^T$ be generated by Algorithm 1. Assume that $\sum_{t=0}^{\infty}\nu_t < +\infty$, where ν_t is as in (8). Then,

$$
\lim_{t\to+\infty}\left\|\nabla f_{\mathcal{G}_t}(x_t) + h'(x_t)\right\| = 0. \qquad (16)
$$

If, additionally, $\lim_{t\to\infty}\|\nabla f_{\mathcal{G}_t}(x_t) - \nabla f(x_t)\| = 0$, we have

$$
\lim_{t\to+\infty}\left\|\nabla f(x_t) + h'(x_t)\right\| = 0.
$$

Under the assumptions of Corollary 3, and noting that $\{\nabla f(x_t) + h'(x_t)\} \to 0$, together with

$$
\nabla f(x_t) + h'(x_t) \in \nabla f(x_t) + \partial h(x_t) \subset \partial F(x_t),
$$

we conclude that, if the sequence $\{x_t\}$ admits an accumulation point $\bar{x}$, then, by the closedness of the graph of ∂F, we have $0 \in \partial F(\bar{x})$, that is, $\bar{x}$ is a stationary point of problem (1).

[1] Throughout this work, *a call of the oracle* means the partial or total evaluation of one of following terms $f(\cdot)$, $\nabla f(\cdot)$ and $\nabla^2 f(\cdot)$.

2.1 Proof of Theorem 1 and Corollaries 2 and 3

We now proceed with the proofs of Theorem 1 and Corollaries 2 and 3. To this end, we first recall that, as a direct consequence of assumption **(A1)**, for any $\mathcal{H} \subset \{1, \ldots, N\}$ and all $x, y \in \mathbb{R}^n$, the following inequalities hold:

$$f_{\mathcal{H}}(y) \leq f_{\mathcal{H}}(x) + \langle \nabla f_{\mathcal{H}}(x), y - x \rangle + \frac{1}{2} \langle \nabla^2 f_{\mathcal{H}}(x)(y - x), y - x \rangle + \frac{L}{6} \|y - x\|^3 \quad (17)$$

and

$$\left\| \nabla f_{\mathcal{H}}(y) - \nabla f_{\mathcal{H}}(x) - \nabla^2 f_{\mathcal{H}}(x)(y - x) \right\| \leq \frac{L}{2} \|y - x\|^2. \quad (18)$$

Let $\phi \in \{f, \nabla f, \nabla^2 f\}$ and let $\mathcal{A}, \mathcal{B} \subseteq \{1, \ldots, N\}$ be index sets. We denote the sample-based difference approximation of ϕ by

$$\phi_{\mathcal{A}/\mathcal{B}}(x) := \frac{1}{|\mathcal{A}|} \sum_{i \in \mathcal{A}} \phi_i(x) - \frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \phi_i(x).$$

We next present an auxiliary result that will play a central role in the forthcoming analysis. For clarity, we omit the indices (t, i) from the iterates generated by Algorithm 1.

Lemma 4 For given $x \in Q$, $\theta \geq 0$, $\sigma_0 > 0$, $c > 6$ and $\mathcal{D}^1, \mathcal{D}^2, \mathcal{G}, \mathcal{H} \subset \{1, \ldots, N\}$. Assume that $x^+ \in Q$ satisfies

$$M_{x,\sigma}^{\mathcal{G},\mathcal{H}}(x^+) + h(x^+) \leq h(x) \quad \text{and} \quad \|\nabla M_{x,\sigma}^{\mathcal{G},\mathcal{H}}(x^+) + h'(x^+)\| \leq \theta \|x^+ - x\|^2, \quad (19)$$

for a certain $h'(x^+) \in \partial h(x^+)$. If $\sigma \geq c(L + 6\sigma_0)/(c - 6)$,

$$f_{\mathcal{D}^2/\mathcal{D}^1}(x) + \langle \nabla f_{\mathcal{D}^2/\mathcal{G}}(x), x^+ - x \rangle + \frac{1}{2} \langle \nabla^2 f_{\mathcal{D}^2/\mathcal{H}}(x)(x^+ - x), x^+ - x \rangle \leq \sigma_0 \|x^+ - x\|^3, \quad (20)$$

and

$$\|\nabla^2 f_{\mathcal{G}/\mathcal{H}}(x)\| \leq \sigma_0 \|x^+ - x\|, \quad (21)$$

then

$$(f_{\mathcal{D}^1} + h)(x) - (f_{\mathcal{D}^2} + h)(x^+) \geq \frac{\sigma}{c} \|x^+ - x\|^3, \quad (22)$$

and

$$\|\nabla f_{\mathcal{G}}(x^+) + h'(x^+)\| \leq \left(\frac{(c - 3)}{c} \sigma + \sigma_0 + \theta \right) \|x^+ - x\|^2. \quad (23)$$

Proof Using (17) with $y := x^+$ and $\mathcal{H} := \mathcal{D}^2$, the definition of $M_{x,\sigma}^{\mathcal{G},\mathcal{H}}$ in (3), and the first inequality in (19), (20) and (21), we have

$$\begin{aligned} f_{\mathcal{D}^2}(x^+) + h(x^+) &\leq f_{\mathcal{D}^2}(x) + \langle \nabla f_{\mathcal{D}^2}(x), x^+ - x \rangle + \frac{1}{2} \langle \nabla^2 f_{\mathcal{D}^2}(x)(x^+ - x), x^+ - x \rangle \\ &\quad + \frac{L}{6} \|x^+ - x\|^3 + h(x^+) \\ &\leq f_{\mathcal{D}^2}(x) + M_{x,\sigma}^{\mathcal{G},\mathcal{H}}(x^+) + h(x^+) + \langle \nabla f_{\mathcal{D}^2/\mathcal{G}}(x), x^+ - x \rangle \end{aligned}$$

$$
\begin{aligned}
&+\ \frac{1}{2}\langle \nabla^2 f_{\mathcal{D}^2/\mathcal{H}}(x)(x^+ - x), x^+ - x\rangle + \frac{L-\sigma}{6}\|x^+ - x\|^3 \\
&\leq\ f_{\mathcal{D}^1}(x) + h(x) + f_{\mathcal{D}^2/\mathcal{D}^1}(x) + \langle \nabla f_{\mathcal{D}^2/\mathcal{G}}(x), x^+ - x\rangle \\
&+\ \frac{1}{2}\langle \nabla^2 f_{\mathcal{D}^2/\mathcal{H}}(x)(x^+ - x), x^+ - x\rangle + \frac{L-\sigma}{6}\|x^+ - x\|^3 \\
&\leq\ f_{\mathcal{D}^1}(x) + h(x) + \left(\sigma_0 + \frac{L-\sigma}{6}\right)\|x^+ - x\|^3,
\end{aligned}
$$

which, combined with the fact that $\sigma \geq c(L+6\sigma_0)/(c-6)$, proves the first inequality in (22).

Now, using the definition of $M_{x,\sigma}^{\mathcal{G},\mathcal{H}}$ in (3), the second inequalities in (19), (21) and (18) with $y := x^+$ and $\mathcal{H} := \mathcal{G}$, we get

$$
\begin{aligned}
\|\nabla f_{\mathcal{G}}(x^+) + h'(x^+)\| &\leq \|\nabla f_{\mathcal{G}}(x^+) - \nabla M_{x,\sigma}^{\mathcal{G},\mathcal{H}}(x^+)\| + \|\nabla M_{x,\sigma}^{\mathcal{G},\mathcal{H}}(x^+) + h'(x^+)\| \\
&= \left\|\nabla f_{\mathcal{G}}(x^+) - \nabla f_{\mathcal{G}}(x) - \nabla^2 f_{\mathcal{H}}(x)(x^+ - x) - \frac{\sigma}{2}\|x^+ - x\|(x^+ - x)\right\| \\
&+\ \theta\|x^+ - x\|^2 \\
&\leq \left\|\nabla f_{\mathcal{G}}(x^+) - \nabla f_{\mathcal{G}}(x) - \nabla^2 f_{\mathcal{H}}(x)(x^+ - x)\right\| + \left(\frac{\sigma}{2} + \theta\right)\|x^+ - x\|^2 \\
&\leq \left\|\nabla f_{\mathcal{G}}(x^+) - \nabla f_{\mathcal{G}}(x) - \nabla^2 f_{\mathcal{G}}(x)(x^+ - x)\right\| + \|\nabla^2 f_{\mathcal{G}/\mathcal{H}}(x)\|\|x^+ - x\|| \\
&+\ \left(\frac{\sigma}{2} + \theta\right)\|x^+ - x\|^2 \\
&\leq \left(\frac{L+\sigma}{2} + \sigma_0 + \theta\right)\|x^+ - x\|^2,
\end{aligned}
$$

which, combined with the fact that $\sigma \geq c(L+6\sigma_0)/(c-6)$, proves the second inequality in (23). □

The following lemma ensures that the inner loop in Step S1 terminates after finitely many trials and establishes bounds on the number of (partial or full) function, gradient, and Hessian evaluations up to a given iteration.

Lemma 5 The sequence of regularization parameters $\{\sigma_t\}$ generated by Algorithm 1 satisfies

$$\sigma_0 \leq \sigma_t \leq \sigma_{\max}, \quad \forall\, t \geq 0, \tag{24}$$

where $\sigma_{\max}$ is as in (9). Moreover, the total number O_T of oracle calls after T iterations satisfies

$$O_T \leq 5\big[\,2T + \log_\alpha(\sigma_{\max}) - \log_\alpha(\sigma_0)\,\big]. \tag{25}$$

Proof Since $c > 6$, it follows trivially from (9) that (24) is true for $t = 0$. Suppose that (24) is true for some $t \geq 0$. If $i_t = 0$, it follows from S2 and the induction assumption that

$$\sigma_0 \leq \sigma_{t+1} = \frac{1}{\alpha}\sigma_t \leq \sigma_t \leq \sigma_{max},$$

which proves that (24) holds for $t+1$. Now, if $i_t \geq 1$, then we must have

$$\alpha^{i_t-1}\sigma_t \leq \sigma_{max} = \max\left\{c(L+6\sigma_0)/(c-6), N/\min\{|\tilde{\mathcal{D}}_0^1|, |\tilde{\mathcal{D}}_0^2|, |\tilde{\mathcal{G}}_0|, |\tilde{\mathcal{H}}_0|\}\right\}. \tag{26}$$

Indeed, assuming by contradiction that (26) is not true, that is

$$\alpha^{i_t-1}\sigma_t > c(L+6\sigma_0)/(c-6), \quad \alpha^{i_t-1}\sigma_t > N/\min\{|\tilde{\mathcal{D}}_0^1|, |\tilde{\mathcal{D}}_0^2|, |\tilde{\mathcal{G}}_0|, |\tilde{\mathcal{H}}_0|\}. \tag{27}$$

Hence, as $|\mathcal{G}_{t,i_t-2}| \geq |\tilde{\mathcal{G}}_0| \geq \min\{|\tilde{\mathcal{D}}_0^1|, |\tilde{\mathcal{D}}_0^2|, |\tilde{\mathcal{G}}_0|, |\tilde{\mathcal{H}}_0|\}$ and $\gamma_{\mathcal{G}} \geq \alpha > 1$, the second inequality in (27) yields that

$$\gamma_{\mathcal{G}}^{i_t-1}\sigma_t|\mathcal{G}_{t,i_t-2}| \geq \alpha^{i_t-1}\sigma_t|\tilde{\mathcal{G}}_0| > N \geq |\mathcal{F}_t|,$$

which in turn implies that $|\mathcal{G}_{t,i_t-1}| = |\mathcal{F}_t|$. With similar arguments, we also have $|\mathcal{D}^1_{t,i_t-1}| = |\mathcal{D}^2_{t,i_t-1}| = |\mathcal{H}_{t,i_t-1}| = |\mathcal{F}_t|$. So, as $\mathcal{D}^1_{t,i+1}, \mathcal{D}^2_{t,i+1}, \mathcal{H}_{t,i+1}, \mathcal{G}_{t,i+1} \subset \mathcal{F}_t$, it follows that $f_{\mathcal{D}^2_{t,i_t-1}/\mathcal{D}^1_{t,i_t-1}}(x_t) = 0$, $\nabla f_{\mathcal{D}^2_{t,i_t-1}/\mathcal{G}_{t,i_t-1}}(x_t) = 0$ and $\nabla^2 f_{\mathcal{D}^2_{t,i_t-1}/\mathcal{H}_{t,i_t-1}}(x_t) = 0$ and $\|\nabla^2 f_{\mathcal{G}_{t,i_t-1}/\mathcal{H}_{t,i_t-1}}(x^+_{t,i})\| = 0$. Hence, by combining the last equalities, the first equality in (27) and Lemma 4 with $x := x_t$, $x^+ := x^+_{t,i}$, $\mathcal{D}^1 := \mathcal{D}^1_{t,i_t-1}$, $\mathcal{D}^2 := \mathcal{D}^2_{t,i_t-1}$, $\mathcal{G} := \mathcal{G}_{t,i_t-1}$, $\mathcal{H} := \mathcal{H}_{t,i_t-1}$ and $\sigma := \alpha^{i_t-1}\sigma_t$, we obtain that the inequalities in (6)-(7) would have been satisfied for $i = i_t - 1$, contradicting the minimality of i_t. Therefore, (26) is true.

Finally, note that at the t-th iteration of Algorithm 1 the number of calls of the oracle is bounded by $5(i_t + 1)$ times. On the other hand,

$$\sigma_{t+1} = \alpha^{i_t-1}\sigma_t \Longrightarrow 5(i_t+1) = 5\left[2 + \log_\alpha(\sigma_{t+1}) - \log_\alpha(\sigma_t)\right].$$

Thus,

$$\begin{aligned} O_T &\leq \sum_{t=0}^{T}[5(i_t+1)] \leq \sum_{t=0}^{T} 10 + 5\log_\alpha(\sigma_{T+1}) - 5\log_\alpha(\sigma_0) \\ &\leq 10T + 5[\log_\alpha(\sigma_{max}) - \log_\alpha(\sigma_0)], \end{aligned}$$

where the last inequality is due to (24). □

It follows from (25) that

$$\frac{O_T}{T} \leq 10 + \frac{5}{T}[\log_\alpha(\sigma_{max}) - \log_\alpha(\sigma_0)],$$

which implies that the average number of oracle calls per inner iteration, up to the T-th outer iteration, is asymptotically bounded by 10.

We next establish a key result concerning the summability of the sequence $\{\|\nabla f_{\mathcal{G}_t}(x_{t+1}) + h'(x_{t+1})\|\}$.

Lemma 6 Let $\{x_t\}_{t=0}^T$ be generated by Algorithm 1. Then,

$$F(x_t) - F(x_{t+1}) \geq \frac{\alpha\sigma_{t+1}}{c}\|x_{t+1} - x_t\|^3 - 2\nu_t, \quad t = 0, \dots, T-1, \tag{28}$$

and

$$\sum_{t=0}^{T-1} \|\nabla f_{\mathcal{G}_t}(x_{t+1}) + h'(x_{t+1})\|^{\frac{3}{2}} \leq \frac{\left(F(x_0) - F_{\text{low}} + 2\sum_{t=0}^{T-1}\nu_t\right)\left(\alpha(c-3)\sigma_{max} + c(\sigma_0+\theta)\right)^{\frac{3}{2}}}{\alpha\sigma_0 c^{1/2}}, \tag{29}$$

where ν_t and $\sigma_{\max}$ are as in (8) and (9), respectively.

Proof From (6) and the relation $\sigma_{t+1} = \alpha^{i_t-1}\sigma_t$, we have

$$(f_{\mathcal{D}^1_t} + h)(x_t) - (f_{\mathcal{D}^2_t} + h)(x_{t+1}) \geq \frac{\alpha\sigma_{t+1}}{c}\|x_{t+1} - x_t\|^3, \quad t = 0, \dots, T-1.$$

On the other hand, it follows from (8) that

$$\begin{aligned} f_{\mathcal{D}_t^1}(x_t) - f_{\mathcal{D}_t^2}(x_{t+1}) &= f(x_t) - f(x_{t+1}) + [f_{\mathcal{D}_t^1}(x_t) - f(x_t)] - [f_{\mathcal{D}_t^2}(x_{t+1}) - f(x_{t+1})] \\ &\leq f(x_t) - f(x_{t+1}) + |f_{\mathcal{D}_t^1}(x_t) - f(x_t)| + |f_{\mathcal{D}_t^2}(x_{t+1}) - f(x_{t+1})| \\ &\leq f(x_t) - f(x_{t+1}) + 2\nu_t. \end{aligned}$$

Combining the above inequalities and using $F = f + h$, we obtain

$$F(x_t) - F(x_{t+1}) + 2\nu_t \geq \frac{\alpha\sigma_{t+1}}{c}\|x_{t+1} - x_t\|^3, \quad t = 0, \ldots, T-1,$$

which corresponds to inequality (28).

Summing the inequalities in (28) over $t = 0, \ldots, T-1$ and using the definition of F_{low} together with (24) , we obtain

$$\begin{aligned} F(x_0) - F_{\text{low}} + 2\sum_{t=0}^{T-1}\nu_t \geq F(x_0) - F(x_T) + 2\sum_{t=0}^{T-1}\nu_t &\geq \sum_{t=0}^{T-1}\frac{\alpha\sigma_{t+1}}{c}\|x_{t+1} - x_t\|^3 \\ &\geq \frac{\alpha\sigma_0}{c}\sum_{t=0}^{T-1}\|x_{t+1} - x_t\|^3. \end{aligned}$$

Hence,

$$\sum_{t=0}^{T-1}\|x_{t+1} - x_t\|^3 \leq \frac{c\big(F(x_0) - F_{\text{low}} + 2\sum_{t=0}^{T-1}\nu_t\big)}{\alpha\sigma_0}. \tag{30}$$

Moreover, from (7), the relation $\sigma_{t+1} = \alpha^{i_t - 1}\sigma_t$, and (24), it follows that

$$\begin{aligned} \|\nabla f_{\mathcal{G}_t}(x_{t+1}) + h'(x_{t+1})\| &\leq \left(\frac{\alpha(c-3)}{c}\sigma_{t+1} + \sigma_0 + \theta\right)\|x_{t+1} - x_t\|^2 \\ &\leq \left(\frac{\alpha(c-3)}{c}\sigma_{\max} + \sigma_0 + \theta\right)\|x_{t+1} - x_t\|^2. \end{aligned}$$

Combining this inequality with (30) yields the desired result (29). □

We are now ready to prove Theorem 1.

Proof of Theorem 1 First, note that the assumption in (10) ensures that $T_0(\varepsilon)$ is well-defined for any given $\varepsilon > 0$. Let

$$t_* := \arg\min_{j\in\{0,\ldots,T-1\}}\|\nabla f_{\mathcal{G}_j}(x_{j+1}) + h'(x_{j+1})\|.$$

Then, from (29), we obtain

$$\begin{aligned} \|\nabla f_{\mathcal{G}_{t_*}}(x_{t_*+1}) + h'(x_{t_*+1})\|^{3/2} &\leq \frac{\big(F(x_0) - F_{\text{low}} + 2\sum_{t=0}^{T-1}\nu_t\big)\,[\alpha(c-3)\sigma_{\max} + c(\sigma_0+\theta)]^{3/2}}{T\alpha\sigma_0 c^{1/2}} \\ &= \frac{(F(x_0) - F_{\text{low}})\,[\alpha(c-3)\sigma_{\max} + c(\sigma_0+\theta)]^{3/2}}{T\alpha\sigma_0 c^{1/2}} \\ &\quad + \frac{2[\alpha(c-3)\sigma_{\max} + c(\sigma_0+\theta)]^{3/2}}{T\alpha\sigma_0 c^{1/2}}\sum_{t=0}^{T-1}\nu_t. \end{aligned}$$

Since (12) holds, we have $T \geq T_0(\varepsilon)$. Therefore, using (11), it follows that

$$\frac{2[\alpha(c-3)\sigma_{\max} + c(\sigma_0+\theta)]^{3/2}}{T\alpha\sigma_0 c^{1/2}}\sum_{t=0}^{T-1}\nu_t \leq \frac{\varepsilon^{3/2}}{2}.$$

In addition, by (12) again,

$$\frac{(F(x_0) - F_{\text{low}})\,[\alpha(c-3)\sigma_{\max} + c(\sigma_0+\theta)]^{3/2}}{T\alpha\sigma_0 c^{1/2}} \leq \frac{\varepsilon^{3/2}}{2}.$$

Combining the above inequalities yields

$$\|\nabla f_{\mathcal{G}_{t_*}}(x_{t_*+1}) + h'(x_{t_*+1})\|^{3/2} \leq \varepsilon^{3/2},$$

which proves (13). □

We next prove Corollary 2.

Proof of Corollary 2 As discussed in Remark 2, if the sequence $\{\nu_t\}$ is summable, then condition (10) automatically holds. Moreover, by defining

$$T_0(\varepsilon) := \frac{4[\alpha(c-3)\sigma_{\max} + c(\sigma_0+\theta)]^{3/2}}{\alpha\sigma_0 c^{1/2}\varepsilon^{3/2}} \sum_{t=0}^{\infty} \nu_t < \infty,$$

we obtain

$$T \geq T_0(\varepsilon) \implies \frac{1}{T}\sum_{t=0}^{T-1} \nu_t \leq \frac{\alpha\sigma_0 c^{1/2}\varepsilon^{3/2}}{4[\alpha(c-3)\sigma_{\max} + c(\sigma_0+\theta)]^{3/2}},$$

which ensures that condition (11) is satisfied. Consequently, from (14) we have

$$\begin{aligned} T &\geq \frac{2[\alpha(c-3)\sigma_{\max} + c(\sigma_0+\theta)]^{3/2}}{\alpha\sigma_0 c^{1/2}} \max\left\{F(x_0) - F_{\text{low}},\ 2\sum_{t=0}^{\infty}\nu_t\right\}\varepsilon^{-3/2} \\ &\geq \max\left\{\frac{2(F(x_0) - F_{low})[\alpha(c-3)\sigma_{\max} + c(\sigma_0+\theta)]^{3/2}}{\alpha\sigma_0 c^{1/2}}\varepsilon^{-3/2},\ T_0(\varepsilon)\right\}. \end{aligned}$$

Thus, inequality (15) follows directly from Theorem 1. The second part of the corollary then follows immediately from the first one, inequality in (25) and the definition of $\sigma_{\max}$ in (9). □

We next prove Corollary 3.

Proof of Corollary 3 By taking the limit in (29) as T goes to infinity, we obtain

$$\sum_{t=0}^{\infty} \|\nabla f_{\mathcal{G}_t}(x_{t+1}) + h(x_{t+1})\|^{\frac{3}{2}} < \infty$$

which in turn implies (16). Now, note that

$$0 \leq \|\nabla f(x_t) + h(x_t)\| \leq \|\nabla f(x_t) - \nabla f_{\mathcal{G}_t}(x_t)\| + \|\nabla f_{\mathcal{G}_t}(x_t) + h(x_t)\|.$$

Hence, the desired result follows trivially from the assumption that $\lim_{t\to+\infty}\|\nabla f_{\mathcal{G}_t}(x_t) - \nabla f(x_t)\| = 0$ and (16). □

3 Numerical Experiments

In this section, we report numerical results of Algorithm 1 (referred to as complete-S-CRM) for solving the problem

$$\min_{x\in\mathbb{R}^N} f(x) \equiv \frac{1}{N}\sum_{i=1}^{N} f_i(x), \tag{31}$$

where

$$f_i(x) := \left(N - \sum_{j=1}^{N} \cos(x_j) + i\big(1 - \cos(x_i)\big) - \sin(x_i) \right)^2 .$$

Problem (31) is an instance of (1) with $n = N$ and $h \equiv 0$. The proposed algorithm was compared with two other schemes: (i) the subsampled CRM of [10] (referred to as partial-S-CRM), which is essentially equivalent to Algorithm 1 with $\mathcal{D}_t^1 = \mathcal{D}_t^2 = \mathcal{G}_t$ for all $t \geq 0$; and (ii) a standard CRM variant (referred to as standard-CRM), which corresponds to Algorithm 1 with $\mathcal{D}_t^1 = \mathcal{D}_t^2 = \mathcal{G}_t = \mathcal{H}_t = \{1, 2, \ldots, N\}$ for all $t \geq 0$. The algorithms were initialized with the parameters

$$x_0 = (1, \ldots, 1), \quad \theta = 5, \quad \sigma_0 = 1, \quad c = 12, \quad \gamma_{\mathcal{D}^1} = \gamma_{\mathcal{D}^2} = \gamma_{\mathcal{G}} = \gamma_{\mathcal{H}} = \alpha = 2,$$

and

$$\tilde{\mathcal{D}}_0^1 = \tilde{\mathcal{D}}_0^2 = \tilde{\mathcal{G}}_0 = \tilde{\mathcal{H}}_0 = \mathcal{F}_0 = 0.05N,$$

with $\mathcal{F}_0$ chosen randomly. In step (S.1) of Algorithm 1, the subsamples in complete-S-CRM were selected according to

$$\mathcal{D}_{t,i}^1 = \mathcal{F}_t, \quad \mathcal{D}_{t,i}^2 = 0.9\,|\mathcal{D}_{t,i}^1|, \quad |\mathcal{G}_{t,i}| = 0.95\,|\mathcal{D}_{t,i}^1|, \quad |\mathcal{H}_{t,i}| = 0.1\,|\mathcal{G}_{t,i}|,$$

whereas in the partial-S-CRM the choices were $\mathcal{D}_{t,i}^1 = \mathcal{D}_{t,i}^2 = \mathcal{G}_{t,i} = \mathcal{F}_t$ and $|\mathcal{H}_{t,i}| = 0.1\,|\mathcal{F}_t|$. For both methods, the set $\mathcal{F}_t$ was updated as

$$|\mathcal{F}_t| = \min\big\{N, \lceil 1.25^t |\mathcal{F}_0| \rceil\big\}, \qquad t \geq 0.$$

For all methods, each cubic subproblem was approximately solved using the Barzilai–Borwein gradient (BBG) method [3], combined with the nonmonotone line search of [16]. In order to have a fair stopping criterion, we terminate the algorithms when

$$\|\nabla f(x_k)\| / \|\nabla f(x_0)\| \leq 10^{-6}. \tag{32}$$

It is worth noting that the evaluation of the full gradient $\nabla f(x_k)$ in the above stopping criterion were not included in the performance measures reported below. The experiments were conducted using the Python programming language, which was installed on a machine equipped with a 3.5 GHz Dual-Core Intel Core i5 processor and 16 GB of 2400 MHz DDR4 memory.

We adopted the total computational cost as the performance measure. The total cost of an algorithm up to iteration T is defined as

$$\text{Cost}_T(\cdot) := \frac{1}{N} \sum_{j=1}^{T} \left(f_i e(j) + 3 \times g_i e(j) + p g_i v(j) + p H_i v(j) \right),$$

where, for each iteration $t \geq 0$: (a) $f_i e(t)$ denotes the total number of evaluations of the component functions $f_i(\cdot)$; (b) $g_i e(t)$ represents the total number of evaluations of the component gradients $\nabla f_i(\cdot)$; (c) $p g_i v(t)$ and $p H_i v(t)$ denote, respectively, the total number of component gradient–vector and component Hessian–vector products computed inside the BBG method. This cost function reflects the relative computational burden of gradient evaluations compared to function evaluations by assigning a weight of 3 to each gradient evaluation. This choice is motivated by the fact that, when using reverse-mode automatic differentiation, computing $\nabla f_i(\cdot)$ typically requires about three times the effort of a single evaluation of $f_i(\cdot)$; see, e.g., [1, 2]. In addition, the definition incorporates the cost of the gradient-vector and Hessian-vector products arising in the BBG method.

Table 1 reports the total cost incurred by complete-S-CRM, partial-S-CRM, and standard-CRM to compute an iterate x_t satisfying the stopping criterion (32), for the function $f(\cdot)$ defined in (31) with different values of N. The table also shows the reductions (Reduc), in terms of $Cost_T$, obtained by complete-S-CRM relative to the other methods. As observed, complete-S-CRM was the most efficient method, achieving the lowest cost in all instances. The reductions achieved by complete-S-CRM compared to standard-CRM were 66.05%, 45.50%, 44.48%, and 42.42% for $N = 100$, 500, 1000, and 3000, respectively, clearly demonstrating the benefit of using subsampling techniques. Moreover, the proposed algorithm also showed an impressive reduction compared to partial-S-CRM, mainly because, in order to accept a new step (see Eq. (6) and (7)), the sample sizes involved are smaller than those required by the latter method. Figure 1 illustrates the behavior of the sample sizes $|\mathcal{D}_t^2|$, $|\mathcal{G}_t|$, and $|\mathcal{H}_t|$ for the subsampled methods during the iterations for solving the problem (31) with $N = 1000$ and $N = 3000$. The sample size $|\mathcal{D}_t^1|$ is omitted because it is identical for both methods. For the subsamples compared, complete-S-CRM employs smaller sets in almost every iteration, leading to substantial cost savings.

N	$Cost_T$(Complete-S-CRM)	$Cost_T$(partial-S-CRM)/Reduc	$Cost_T$(standard-CRM)/Reduc
100	382.61	425.71/10.12%	1127.0/66.05%
500	629.43	745.94/15.62 %	1155.0/45.50%
1,000	718.47	1011.87/29.00%	1294.0/44.48%
3,000	814.22	1183.56/31.21%	1414.0/42.42%

Table 1 Comparison of $Cost_T$ among complete-S-CRM, partial-S-CRM and standard-CRM across different problem sizes.

4 Conclusion

We proposed a subsampled cubic regularization method for solving finite-sum composite optimization problems, in which the function, gradient, and Hessian are estimated using possibly different sample sizes. This flexibility enables the algorithm to independently control the accuracy of each model component, achieving a more effective balance between computational cost and estimation quality. We established iteration-complexity guarantees for computing approximate first-order critical points and proved global convergence of the generated sequence. Numerical experiments further confirmed the efficiency and practical advantages of the proposed approach.

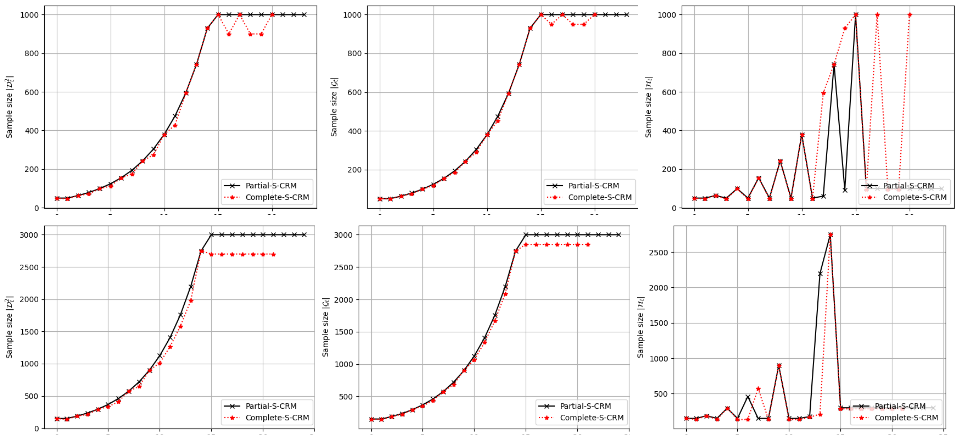

Fig. 1 Evolution of the samples sizes $|\mathcal{D}_t^2|$, $|\mathcal{G}_t|$ and $|\mathcal{H}_t|$ during the iterations for problem (31). Top row: $N = 1000$; bottow row line: $N = 3000$.

Acknowledgments

The author is grateful to Philippe Toint (naXys, University of Namur, Belgium) for valuable suggestions and discussions regarding this work. This work was funded by FAPEG and FAPESC [grant number 2024TR002238]

References

[1] A. Griewank. Evaluationg derivatives: Principles and techniques of algorithmic differentiation. SIAM, Philadelphia, 2008.

[2] Michael Bartholomew-Biggs, Steven Brown, Bruce Christianson, and Laurence Dixon. Automatic differentiation of algorithms. *Journal of Computational and Applied Mathematics*, 124(1):171–190, 2000. Numerical Analysis 2000. Vol. IV: Optimization and Nonlinear Equations.

[3] J. Barzilai and J. M. Borwein. Two-Point Step Size Gradient Methods. *IMA J. Numer. Anal.*, 8(1):141–148, 1988.

[4] S. Bellavia and Gurioli. G. Stochastic analysis of an adaptive cubic regularization method under inexact gradient evaluations and dynamic hessian accuracy. *Optimization*, 71(1):227–261, 2022.

[5] S. Bellavia, G. Gurioli, and B. Morini. Adaptive cubic regularization methods with dynamic inexact Hessian information and applications to finite-sum minimization. *IMA J. Numer. Anal.*, 41(1):764–799, 04 2020.

[6] S. Bellavia, G. Gurioli, B. Morini, and P. L. Toint. Quadratic and cubic regularisation methods with inexact function and random derivatives for finite-sum minimisation. In *2021 21st International Conference on Computational Science*

and Its Applications (ICCSA), pages 258–267, 2021.

[7] S. Bellavia, N. Krejić, and N. K. Jerinkić. Subsampled inexact Newton methods for minimizing large sums of convex functions. *IMA J. Numer. Anal.*, 40(4):2309–2341, 07 2019.

[8] S. Bellavia, N. Krejić, and B. Morini. Inexact restoration with subsampled trust-region methods for finite-sum minimizations. *Comput. Optim. Appl.*, 76(2):701–736, 2020.

[9] E. G. Birgin, J. L. Gardenghi, J. M. Martínez, S. A. Santos, and Ph. L. Toint. Worst-case evaluation complexity for unconstrained nonlinear optimization using high-order regularized models. *Math. Program.*, 163(1):359–368, 2017.

[10] M. L. N. Gonçalves. Subsampled cubic regularization method for finite-sum minimization. *Optimization*, 74(7):1591–1614, 2025.

[11] X. Chen, B. Jiang, T. Lin, and S. Zhang. Accelerating adaptive cubic regularization of Newton's method via random sampling. *J. Mach. Learn. Res.*, 23(90):1–38, 2022.

[12] D. di Serafino, N. Krejić, N. K. Jerinkić, and M. Viola. Lsos: Line-search second-order stochastic optimization methods for nonconvex finite sums. *Math. Comput.*, 92(341):1273–1299, 2023.

[13] N. Doikov and G. N. Grapiglia. First and zeroth-order implementations of the regularized Newton method with lazy approximated hessians. *J Sci Comput*, 103(32):1–36, 2025.

[14] G. N. Grapiglia and E.W. Sachs. A generalized worst-case complexity analysis for non-monotone line searches. *Numer. Algorithms*, 87(2):779–796, 2021.

[15] G.N. Grapiglia, M.L.N. Gonçalves, and G.N. Silva. A cubic regularization of Newton's method with finite difference hessian approximations. *Numer. Algorithms*, 90:607–630, 2022.

[16] L. Grippo, F. Lampariello, and S. Lucidi. A nonmonotone line search technique for Newton's method. *SIAM J. Numer. Anal.*, 23(4):707–716, 1986.

[17] J. Kohler and A. Lucchi. Sub-sampled cubic regularization for non-convex optimization. In *In Proceedings of the 34th International Conference on Machine Learning (ICML 2017)*, volume 70, pages 1895 – 1904, 2017.

[18] Y. Nesterov and B. T. Polyak. Cubic regularization of Newton method and its global performance. *Math. Program.*, 184(1):177–205, 2006.

[19] S. Park, S. H. Jung, and P. M. Pardalos. Combining stochastic adaptive cubic regularization with negative curvature for nonconvex optimization. *J. Optim.*

Theory Appl., 184(3):953–971, 2020.

[20] B. Stefania, N. K. Jerinkić, and G. Malaspina. Subsampled nonmonotone spectral gradient methods. *Commun. Appl. Ind. Math.*, 11(1):19–34, 2020.

[21] Z. Wang, Y. Zhou, Y. Liang, and G. Lan. Cubic regularization with momentum for nonconvex optimization. In *Uncertainty in Artificial Intelligence*, pages 313–322. PMLR, 2020.