On Structured Filtering-Clustering: Global Error Bound and Optimal First-Order Algorithms

Nhat Ho[⋆,⋄] Tianyi Lin[⋆,‡] Michael I. Jordan[⋄,†]

Department of Electrical Engineering and Computer Sciences[⋄]
Department of Industrial Engineering and Operations Research[‡]
Department of Statistics[†]
University of California, Berkeley

August 4, 2019

Abstract

In recent years, the filtering-clustering problems have been a central topic in statistics and machine learning, especially the ℓ_1-trend filtering and ℓ_2-convex clustering problems. In practice, such structured problems are typically solved by first-order algorithms despite the extremely ill-conditioned structures of difference operator matrices. Inspired by the desire to analyze the convergence rates of these algorithms, we show that for a large class of filtering-clustering problems, a *global error bound* condition is satisfied for the dual filtering-clustering problems when a certain regularization is chosen. Based on this result, we show that many first-order algorithms attain the *optimal rate of convergence* in different settings. In particular, we establish a generalized dual gradient ascent (GDGA) algorithmic framework with several subroutines. In deterministic setting when the subroutine is accelerated gradient descent (AGD), the resulting algorithm attains the linear convergence. This linear convergence also holds for the finite-sum setting in which the subroutine is the Katyusha algorithm. We also demonstrate that the GDGA with stochastic gradient descent (SGD) subroutine attains the optimal rate of convergence up to the logarithmic factor, shedding the light to the possibility of solving the filtering-clustering problems efficiently in online setting. Experiments conducted on ℓ_1-trend filtering problems illustrate the favorable performance of our algorithms over other competing algorithms.

1 Introduction

Trend filtering and convex clustering are instances of the general *filtering-clustering problem*, a class of problems that has been widely studied in machine learning and statistics. Examples of trend-filtering applications include nonparametric regression [14, 33, 17, 21, 11], adaptive estimators in graphs [39, 22], and time series analysis [16]. Convex clustering has been proposed as an alternative to traditional clustering methods such as K-means clustering and hierarchical clustering that has appealing robustness and stability properties [12, 44, 31, 40, 26].

There has been much recent work on theoretical and algorithmic aspects of trend filtering and convex clustering. Statistically, solutions to these problem formulations have been shown to possess desirable optimality properties [33, 44, 31, 39, 40, 26, 22, 11]. As for the algorithmic problem of finding optimal solutions, a variety of algorithms have been investigated—for trend filtering these include primal-dual interior-point algorithms (PDIP) [14], the alternating direction method of multipliers (ADMM) [28] and Newton's algorithm [39] and for convex clustering they include ADMM, an alternating minimization algorithm (AMA) [6], projected

⋆ Nhat Ho and Tianyi Lin contributed equally to this work.

dual gradient ascent [38], and semismooth Newton's algorithms [41]. There remains a gap, however, in the theoretical understanding of these algorithms as applied to trend filtering and convex clustering. Indeed, while these algorithms can be successful in practice, there is currently a paucity of theoretical complexity analysis to help explain and guide that progress [6, 28, 38].

The general filtering-clustering problem can be formulated abstractly as follows:

$$\min_{\beta \in \mathbb{R}^d} \ \Phi(\beta) \ := \ f(\beta) + \lambda \sum_{j=1}^{n} \|D_j \beta\|_p \,, \tag{1}$$

where $f : \mathbb{R}^d \to \mathbb{R}$ is a strongly convex loss function with Lipschitz-continuous gradient, $D_j \in \mathbb{R}^{m \times d}$ are *discrete difference operator* matrices for $1 \leq j \leq n$, $\lambda > 0$ is a *regularization parameter* and p is a *regularization index*. In specific applications of the filtering-clustering problem in (1), including ℓ_1-trend filtering and ℓ_2-convex clustering, the loss function f is chosen to be $\|\cdot\|^2$ while the matrices D_j are often extremely poorly conditioned. The ill-conditioning poses several challenges for developing stable and efficient first-order optimization algorithms to find optimal solutions of filtering-clustering problems.

Linear convergence of first-order optimization algorithms is commonly established under some additional assumptions on problem structure; e.g., strong convexity [20]. In the optimization literature, the *local error bound condition* is well known as a relaxation of the strong convexity assumption, providing a guarantee for the *asymptotic* linear convergence of feasible descent algorithms [18, 19, 35] and the conditional gradient method [3]. Notably, Zhou et al. [43] developed a broadly useful approach to the local error bound condition using the *upper Lipschitz continuity* of the underlying set-valued mappings. On the other hand, results also exist for *global error bounds*. For example, Theorem 3.1 in Pang [24] provided a detailed analysis but one that still requires the strong convexity of the objective. Wang and Lin [37] partially relaxed the strong convexity and derived a clean form of a global error bound for a class of structured non-strongly convex problems. Recently, Drusvyatskiy and Lewis [7] have presented a systematic study of the error bound condition and its relationship with quadratic growth and application to the convergence analysis of proximal gradient methods. To the best of our knowledge, no global error bound analysis has been obtained for the filtering-clustering formulation in (1).

Another line of related work focuses on first-order primal-dual optimization methods for convex-concave saddle-point problems (see, for instance, [4, 23, 36, 42, 8] and the references therein). Working in a continuous-time setting, Cherukuri et al. [5] obtained a convergence result under mild conditions but without any results for the rate. For the discrete-time dynamics, some recent work has assumed either a strongly convex-concave structure [4, 23] or full column rank of the coupling matrix [8], together with efficient proximal mappings for non-smooth terms. Unfortunately, however, these assumptions are not satisfied by general filtering-clustering problems. Another interesting approach has constructed a potential function which decreases at a linear rate [36, 42]; however, this function relies heavily on the proximal mapping and can not be used for analyzing filtering-clustering problems.

Our contributions. The contributions of the paper are three-fold.

1. We analyze the structure of general filtering-clustering problems and prove that a *global* error bound condition is satisfied for its dual formulation when $p = 1$ or $p \in [2, +\infty]$. It is worth noting that the result is nontrivial; in particular, it is not amenable to the standard techniques developed in [37] which require the nonsmooth term in the dual objective to have a polyhedral epigraph, which corresponds to $p = 1$ and $p = +\infty$. Additionally,

since the filtering-clustering problem in (1) can not be formulated as an $\ell_{1,p}$-regularized problem for some $p \geq 1$, the proof techniques in [43] is not directly applicable.

2. We propose a class of deterministic first-order algorithms for solving filtering-clustering problems with a linear rate of convergence, which is known to be optimal in terms of ε for the deterministic settings [20]. There are two fundamental reasons for the non-triviality of the result: (i) The dual objective function of filtering-clustering problems is not strongly convex since matrices $D_j^\top$ are not full column rank; (ii) the gradient of the dual objective function is not accessible in general, so vanilla projected gradient descent is not applicable. Facing these challenges, we propose a class of efficient first-order deterministic algorithms for filtering-clustering problems with provably optimal linear convergence.

3. In addition to deterministic first-order algorithms, we also propose and analyze a class of stochastic first-order gradient-type optimization algorithms for filtering-clustering problems. For the finite-sum versions of these problems, stochastic first-order algorithms based on variance reduction attain the optimal linear convergence rate [1]. Moving to the online setting, a similar analysis is applied to show that our stochastic first-order algorithms achieve the optimal rate up to a logarithmic factor [27].

Paper Organization. The remainder of the paper is organized as follows. In Section 2, we present definitions and the main assumptions made throughout this paper. Besides that, specific examples of filtering-clustering problems and several different forms of problem (1) are presented to provide insight into the scope of the problem. In Section 3, we derive a global error bound for the dual form when $p = 1$ or $p \in [2, +\infty]$. In Section 4, we develop a unified algorithmic framework of generalized dual gradient descent for solving problem (1) with a rigorous theoretical guarantee. Some specializations of the general algorithmic framework to different settings are also analyzed where we provide complexity bounds for these problems. We present some numerical results on the ℓ_1-trend filtering problem in Section 6. A few detailed proofs are presented in Section 7. We conclude in Section 8.

Notation. Throughout the paper, we let $\sigma_{\max}(A)$ denote the largest eigenvalue of matrix $A \in \mathbb{R}^{m\times m}$. Additionally, $\|\mathbf{x}\|_p$, where $p \in [1, +\infty]$, denotes ℓ_p-norm of $\mathbf{x}$ and $\|\mathbf{x}\|$ denotes the standard Euclidean norm. For all $q \geq 1$, $\mathbb{B}_q = \{\alpha \in \mathbb{R}^m \mid \|\alpha\|_q \leq 1\}$ refers to a ℓ_q-norm unit ball in $\mathbb{R}^m$ and $D_q = \max_{\mathbf{x},\mathbf{y}\in\mathbb{B}_q} \|\mathbf{x} - \mathbf{y}\|$ refers to a diameter of ℓ_q-norm unit ball in ℓ_2-norm. We also denote $\mathbb{B}_q^n$ as the product of n unit balls in ℓ_q-norm. For a convex function f, ∂f refers to the subdifferential of f. If f is differentiable, $\partial f = \{\nabla f\}$ where ∇f is the gradient vector of f. For any closed set $\mathcal{S}$, we let $d(\mathbf{x}, \mathcal{S})$ denote the distance between $\mathbf{x}$ and $\mathcal{S}$. If $\mathcal{S}$ is convex, we let $\mathcal{N}_{\mathcal{S}}(\mathbf{x})$ denote the normal cone to $\mathcal{S}$ at $\mathbf{x}$. Given a scalar tolerance $\varepsilon \in (0, 1)$, the notation $n = \mathcal{O}(m(\varepsilon))$ stands for the upper bound $n \leq Cm(\varepsilon)$ in which $C > 0$ is independent of ε.

2 Background

In this section, we first flesh out the basic filtering-clustering problem in (1). Then, specific examples of filtering-clustering problems are presented in Section 2.2. Finally, we proceed to discuss various forms of the general filtering-clustering problem in Section 2.3 and discuss how first-order optimization methods can be applied to their solution.

2.1 Filtering-clustering problems

Our goal is to find an optimal solution to problem (1):

Definition 1. *β^* is an optimal solution to problem (1) if $\forall \beta \in \mathbb{R}^d$, $\Phi(\beta^*) \leq \Phi(\beta)$.*

Since convergence of algorithms to an optimal solution will depend on the gradient in a neighborhood of a global optimal solution, it is necessary to impose smoothness conditions on the gradient. Furthermore, since f refers to a loss function for the filtering-clustering problem, it is reasonable to impose the strong convexity on f. For stochastic first-order algorithms, we impose unbiased and bounded variance conditions on the stochastic gradient oracle.

Definition 2. *f is ℓ-gradient Lipschitz if $\forall \beta, \beta' \in \mathbb{R}^d$, $\|\nabla f(\beta) - \nabla f(\beta')\| \leq \ell \|\beta - \beta'\|$.*

Definition 3. *f is μ-strongly convex if $\forall \beta, \beta' \in \mathbb{R}^d$, $\|\nabla f(\beta) - \nabla f(\beta')\| \geq \mu \|\beta - \beta'\|$.*

Definition 4. *$G(\cdot, \xi)$ is unbiased if $\forall \beta \in \mathbb{R}^d$, $\mathbb{E}[G(\beta, \xi)] = \nabla f(\beta)$.*

Definition 5. *$G(\cdot, \xi)$ is bounded if $\forall \beta \in \mathbb{R}^d$, $\mathbb{E}[\|G(\beta, \xi)\|^2] \leq C^2$ for a universal $C > 0$.*

Throughout this paper, we make the following main assumption.

Assumption 2.1. *$f : \mathbb{R}^d \to \mathbb{R}$ is ℓ-gradient Lipschitz and μ-strongly convex. The stochastic gradient oracle $G(\cdot, \xi)$ is unbiased and bounded if available. The optimal set is nonempty.*

Since the objective is strongly convex, the filtering-clustering problem in (1) has a unique optimal solution β^*. For a finite-time algorithm, we cannot expect to find an exact optimal solution in general and we therefore aim for an *ε-optimal solution*.

Definition 6. *$\beta \in \mathbb{R}^d$ is a ε-optimal solution to problem (1) if $\|\beta - \beta^*\|^2 \leq \varepsilon$.*

Given these definitions, our goal in the paper is to develop efficient first-order optimization algorithms that find a ε-optimal solution to problem (1) under the Lipschitz assumptions of f.

2.2 Specific instances of filtering-clustering problems

In this section, we provide some examples of filtering-clustering problems in real applications and comment on the existing algorithms developed for solving them.

2.2.1 Univariate ℓ_1-trend filtering

Trend filtering [14, 33] has been proposed as a new approach to nonparametric regression. In particular, for $1 \leq i \leq \bar{n}$, $y_i = f_0(\mathbf{x}_i) + w_i$, where $(\mathbf{x}_i, y_i)$ are an input/response pair and the random variables $w_1, \ldots, w_{\bar{n}}$ are independent and identically distributed.

Given an integer $k \geq 0$, the k-th order ℓ_1-trend filtering is implemented by solving the following ℓ_1-regularized least-squares problem:

$$\min_{\beta \in \mathbb{R}^{\bar{n}}} \left\{ \frac{1}{2} \|y - \beta\|^2 + \lambda \|D^{(k+1)} \beta\|_1 \right\}, \tag{2}$$

where $\lambda \geq 0$ is a regularization parameter and $D^{(k+1)} \in \mathbb{R}^{(\bar{n}-k-1) \times \bar{n}}$ is a discrete difference operator of order $k+1$. Some typical examples are presented as follows,

$$D^{(1)} = \begin{bmatrix} -1 & 1 & 0 & \ldots & 0 & 0 \\ 0 & -1 & 1 & \ldots & 0 & 0 \\ \vdots & \vdots & \vdots & \ldots & \vdots & \vdots \\ 0 & 0 & 0 & \ldots & -1 & 1 \end{bmatrix}, \qquad D^{(2)} = \begin{bmatrix} 1 & -2 & 1 & 0 & \ldots & 0 & 0 \\ 0 & -1 & 2 & 1 & \ldots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ldots & \vdots & \\ 0 & 0 & 0 & 0 & \ldots & -2 & 1 \end{bmatrix}.$$

In general, the nonzero elements in each row of the matrix $D^{(k+1)}$ are the $(k+1)$-th row of Pascal's triangle with alternating signs. When $k=0$, the ℓ_1-trend filtering problem (2) can be interpreted as a special instance of one-dimensional total variation denoising [30] or fused lasso problem [32]. As for the algorithmic problem of finding optimal solutions, two representative algorithms for ℓ_1-trend filtering problems (2) are specialized PDIP [14] and specialized ADMM [28]. Despite their superior practical performance, these specialized algorithms lack theoretical guarantees. To the best of our knowledge, no linearly convergent first-order algorithm has been proposed for ℓ_1-trend filtering problems in (2).

2.2.2 Graph ℓ_1-trend filtering

Graph trend filtering [39] is an interesting extension of the trend filtering problem to graphs. Let $G=(V,E)$ denote a graph consisting of a set of nodes $V=\{1,2,\ldots,\bar{n}\}$ and undirected edges $E=(e_1,\ldots,e_{\bar{m}})$. Given an integer $k\geq 0$ and observations associated with the nodes, $y=(y_1,\ldots,y_{\bar{n}})\in\mathbb{R}^{\bar{n}}$, the k-th order graph ℓ_1-trend filtering is implemented by solving the following ℓ_1-regularized least squares problem:

$$\min_{\beta\in\mathbb{R}^{\bar{n}}}\left\{\frac{1}{2}\|y-\beta\|^2+\lambda\|\Delta^{(k+1)}\beta\|_1\right\}, \tag{3}$$

where $\lambda\geq 0$ is a regularization parameter and $\Delta^{(k+1)}$ is graph-difference operator of order $k+1$. The explicit form of $\Delta^{(1)}\in\mathbb{R}^{\bar{m}\times\bar{n}}$ is given by:

$$\Delta_l^{(1)} \;=\; \left(0,\ldots,\underbrace{-1}_{i},\ldots,\underbrace{1}_{j},\ldots,0\right), \quad \text{if } e_l=(i,j),\ 1\leq l\leq m.$$

In other words, $\Delta_l^{(1)}$ denotes the l-th row of matrix $\Delta^{(1)}$. Based on this, the structure of $\Delta^{(k+1)}$ can be represented as

$$\Delta^{(k+1)} \;=\; \begin{cases}(\Delta^{(1)})^\top\Delta^{(k)} & \text{if } k \text{ is odd},\\ \Delta^{(1)}\Delta^{(k)} & \text{if } k \text{ is even}.\end{cases}$$

When $k=0$, the graph ℓ_1-trend filtering problem is a special fused-lasso problem over the graph [34]. State-of-the-art algorithms for graph ℓ_1-trend filtering problem include ADMM and a Newton algorithm [39]. However, the theory for these algorithms is limited, and in particular there is no analysis of a first-order algorithm that establishes linear convergence.

2.2.3 ℓ_2-Convex clustering

Convex clustering has been proposed as an alternative to traditional clustering models and leads to a convex optimization problem [12]. In particular, given a number of observations $\{x_i\}_{i=1}^{\bar{n}}\subseteq\mathbb{R}^{\bar{n}}$, the ℓ_2-convex clustering is achieved by solving the following ℓ_2-regularized least square problem:

$$\min_{\{\beta_i\}_{i=1}^{\bar{n}}\subseteq\mathbb{R}^{\bar{n}}}\left\{\frac{1}{2}\sum_{i=1}^{\bar{n}}\|x_i-\beta_i\|^2+\lambda\sum_{1\leq i<j\leq\bar{n}}w_{ij}\|\beta_i-\beta_j\|_2\right\}, \tag{4}$$

where $\lambda\geq 0$ is a regularization parameter and $w_{ij}\geq 0$ are weight parameters. Standard optimization algorithms that have been used heuristically for this problem include ADMM and

alternating minimization algorithm (AMA) [6]. but neither ADMM nor AMA is rigorously justified in this setting. Recently Yuan et al. [41] have proposed a semismooth Newton algorithm for solving convex clustering problem (4). This algorithm enjoys a solid theoretical guarantee and shows favorable performance on several real datasets. However, the semismooth Newton algorithm is quite complex to implement, requiring the tuning of several hyperparameters that have a strong effect on performance.

2.3 Filtering-clustering problems in different forms

In this section, we discuss alternative forms of the filtering-clustering problem, beginning with a convex-concave saddle-point formulation. Letting $q = \frac{p}{p-1}$, problem (1) can be rewritten as

$$\min_{\beta\in\mathbb{R}^d} \max_{\alpha\in\mathbb{B}_q^n} \left\{ f(\beta) - \lambda\alpha^\top D\beta \right\}, \qquad \alpha := \begin{bmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{bmatrix}, \; D := \begin{bmatrix} D_1 \\ \vdots \\ D_n \end{bmatrix}, \tag{5}$$

where $\mathbb{B}_q^n$ denotes the product of n unit balls in ℓ_q-norm in $\mathbb{R}^m$. Note that (5) is different from the formulations in [15, 42]. Indeed, these formulations are built on the conjugate of f [29] and given by $\min_{\beta\in\mathbb{R}^d} \max_{\alpha\in\mathbb{R}^d} \left\{ \alpha^\top\beta - f^*(\alpha) + \lambda \sum_{j=1}^n \|D_j\beta\|_p \right\}$. While the algorithms [15, 42] appear suitable for this form of structured loss function f, in fact they are not applicable in general because the proximal mapping of $\|D_j\beta\|_p$ is difficult to compute [25]. Furthermore, (5) also differs from the formulation in [28], given by

$$\min_{\beta\in\mathbb{R}^d, \alpha_j\in\mathbb{R}^m} \left\{ f(\beta) + \lambda \sum_{j=1}^n \|\alpha_j\|_p \right\}, \quad \text{s.t. } \alpha_j = D_j\beta, \; \forall j.$$

Based on the convex-concave saddle point formulation of filtering-clustering problems in (5), the dual form can be obtained as follows:

$$\min_{\alpha\in\mathbb{B}_q^n} \; \bar{f}(\alpha) := f^*\left(\lambda D^\top \alpha\right), \tag{6}$$

where $f^* : \mathbb{R}^{mn} \to \mathbb{R}$ is the convex conjugate of f. More precisely, problem (6) is derived by,

$$\begin{aligned} \min_{\beta\in\mathbb{R}^d} \max_{\alpha\in\mathbb{B}_q^n} \left\{ f(\beta) - \lambda\alpha^\top D\beta \right\} &\iff \max_{\alpha\in\mathbb{B}_q^n} \min_{\beta\in\mathbb{R}^d} \left\{ f(\beta) - \lambda\alpha^\top D\beta \right\} \iff \max_{\alpha\in\mathbb{B}_q^n} \; -f^*(\lambda\alpha^\top D) \\ &\iff \min_{\alpha\in\mathbb{B}_q^n} \; f^*(\lambda\alpha^\top D). \end{aligned}$$

Compared to the filtering-clustering problem in the primal form (1), the filtering-clustering problem in the dual form (6) admits a very special structure. Indeed, f^* is a smooth and strongly convex function and $\mathbb{B}_q$ is a simple and bounded convex set with efficient projection for $q = 1, 2, +\infty$. In the sequel, we demonstrate that a global error bound condition is satisfied for problem (6), allowing for the development of a class of optimal first-order gradient-type optimization algorithms in general settings.

3 Global Error Bound Condition

In this section, we define and prove *a global error bound* (GEB) condition for the filtering-clustering problem in dual form (6). In particular, we first analyze the structure of problem (6),

including the objective function and the optimal set. Then, we introduce a *upper Lipschitz continuity* (ULC) property of a set-valued mapping and borrow some techniques from [43] to provide a sufficient condition for the GEB condition. Finally, we prove that the GEB condition is satisfied by dual problem (6) with $q \in [1,2] \cup \{+\infty\}$, which corresponds to the filtering-clustering problem (1) with $p = \{1\} \cup [2, +\infty]$.

3.1 Problem structure

In this section we develop some structural properties of the objective function in problem (6). We begin by recalling some basic facts on conjugate functions f^*, defined as the minimizer $\beta^*(\alpha) = \operatorname{argmin}_{\beta \in \mathbb{R}^d} \left\{ f(\beta) - \lambda \alpha^\top D\beta \right\}$ and the dual objective function $\bar{f}$.

Lemma 3.1. *f^* is $\frac{1}{\mu}$-gradient Lipschitz and $\frac{1}{\ell}$-strongly convex.*

Proof. We first show that $\alpha \in \partial f(\beta) \Leftrightarrow \beta \in \partial f^*(\alpha)$. Indeed, if $\alpha \in \partial f(\beta)$, then $f^*(\alpha) = \alpha^\top \beta - f(\beta)$. This implies that for all $\alpha' \in \mathbb{R}^m$, we have

$$f^*(\alpha') - f^*(\alpha) \;\geq\; (\alpha')^\top \beta - f(\beta) - \left(\alpha^\top \beta - f(\beta)\right) \;=\; \left(\alpha' - \alpha\right)^\top \beta.$$

Hence, we achieve that $\beta \in \partial f^*(\alpha)$. This also implies that $\beta \in \partial f^*(\alpha) \Rightarrow \alpha \in \partial f^{**}(\beta)$. Since f is proper and convex, $f = f^{**}$ [29, Theorem 12.2] and $\beta \in \partial f^*(\alpha) \Rightarrow \alpha \in \partial f(\beta)$. Based on this result, we prove that f^* is $1/\mu$-gradient Lipschitz. Indeed, since f is μ-strongly convex and differentiable, then $\partial f(\beta) = \{\nabla f(\beta)\}$ and ∇f is one-to-one. This implies that $\partial f^*(\alpha) = \{\nabla f^*(\alpha)\}$ and

$$\|\nabla f^*(\alpha_1) - \nabla f^*(\alpha_2)\| \;=\; \|\beta_1 - \beta_2\| \;\leq\; \frac{\|\nabla f(\beta_1) - \nabla f(\beta_2)\|}{\mu} \;=\; \frac{\|\alpha_1 - \alpha_2\|}{\mu}.$$

Similar arguments yield that f^* is $1/\ell$-strongly convex. As a consequence, we reach the conclusion of the lemma. □

The second lemma presents some properties of $\beta^*(\alpha) = \operatorname{argmin}_{\beta \in \mathbb{R}^d} \{f(\beta) - \lambda \alpha^\top D\beta\}$.

Lemma 3.2. *$\beta^*(\alpha)$ is well-defined and the following statement holds true,*

$$\|\beta^*(\alpha_1) - \beta^*(\alpha_2)\| \;\leq\; \frac{\lambda \|D^\top \alpha_1 - D^\top \alpha_2\|}{\mu}, \quad \forall\, \alpha_1, \alpha_2 \in \mathbb{B}_q^n. \tag{7}$$

Furthermore, $\beta^*(\alpha)$ is $\frac{\lambda\sqrt{\sigma}}{\mu}$-Lipschitz over $\mathbb{B}_q^n$.

Proof. Since f is μ-strongly convex, $\beta^*(\alpha)$ is uniquely determined given $\alpha \in \mathbb{B}_q^n$ and hence well-defined. By the optimality condition, $\nabla f(\beta^*(\alpha)) = \lambda D^\top \alpha$. Putting these pieces together yields that, for all $\alpha_1, \alpha_2 \in \mathbb{B}_q^n$,

$$\|\beta^*(\alpha_1) - \beta^*(\alpha_2)\| \;\leq\; \frac{\|\nabla f(\beta^*(\alpha_1)) - \nabla f(\beta^*(\alpha_2))\|}{\mu} \;\leq\; \frac{\lambda \left\|D^\top \alpha_1 - D^\top \alpha_2\right\|}{\mu}.$$

Furthermore, since $\|D^\top \alpha_1 - D^\top \alpha_2\| \leq \sqrt{\sigma}\|\alpha_1 - \alpha_2\|$, we have

$$\|\beta^*(\alpha_1) - \beta^*(\alpha_2)\| \;\leq\; \frac{\lambda\sqrt{\sigma}\|\alpha_1 - \alpha_2\|}{\mu}, \quad \forall\, \alpha_1, \alpha_2 \in \mathbb{B}_q^n.$$

This completes the proof of the lemma. □

The third lemma shows that $\bar{f}$ is differentiable and gradient Lipschitz.

Lemma 3.3. *$\bar{f}$ is differentiable with $\nabla \bar{f}(\alpha) = \lambda D \beta^*(\alpha)$ and $\frac{\sigma\lambda^2}{\mu}$-gradient Lipschitz.*

Proof. By the definition, $\bar{f}(\alpha) = f^*(\lambda D^\top \alpha)$. This implies that $\bar{f}$ is differentiable and $\nabla \bar{f}(\alpha) = \lambda D \nabla f^*(\lambda D^\top \alpha)$. Since $\nabla f(\beta^*(\alpha)) = \lambda D^\top \alpha$, then $\alpha \in \partial f(\beta) \Leftrightarrow \beta \in \partial f^*(\alpha)$ implies that $\beta^*(\alpha) = \nabla f^*(\lambda D^\top \alpha)$. Putting these pieces together yields that $\nabla \bar{f}(\alpha) = \lambda D \beta^*(\alpha)$. For all $\alpha_1, \alpha_2 \in \mathbb{B}_q$, it holds true that

$$\left\|\nabla \bar{f}(\alpha_1) - \nabla \bar{f}(\alpha_2)\right\| \;\leq\; \left\|\lambda D \left(\nabla f^*(\lambda D^\top \alpha_1) - \nabla f^*(\lambda D^\top \alpha_2)\right)\right\| \;\leq\; \frac{\sigma\lambda^2 \left\|\alpha_1 - \alpha_2\right\|}{\mu}.$$

Therefore, we achieve the conclusion of the lemma. □

Lemmas 3.1, 3.2 and 3.3 shed light on the special structures of f^* and $\bar{f}$. This inspires us to ask if the optimal set of problem (6), denoted as Ω^*, has some special structure. The following proposition gives an affirmative answer by showing that Ω^* is convex and compact.

Proposition 3.4. *There exists a pair of vector $(\mathbf{s}^*, \mathbf{g}^*) \in \mathbb{R}^d \times \mathbb{R}^{mn}$ with $\mathbf{g}^* = \lambda D \nabla f^*(\lambda \mathbf{s}^*)$ such that $\alpha^* \in \Omega^*$ if and only if*

$$D^\top \alpha^* = \mathbf{s}^*, \quad \nabla \bar{f}(\alpha^*) = \mathbf{g}^*, \quad \alpha^* \in \mathbb{B}_q^n.$$

Proof. Firstly, we show that, if $\alpha^* \in \Omega^*$, then $D^\top \alpha^* = \mathbf{s}^*$ and $\alpha^* \in \mathbb{B}_q^n$. Indeed, since $\Omega^* \subseteq \mathbb{B}_q^n$, it is obvious that $\alpha^* \in \Omega^*$ implies that $\alpha^* \in \mathbb{B}_q^n$. Furthermore, $\bar{f}(\alpha) = f^*(\lambda D^\top \alpha)$ and f^* is $1/\ell$-strongly convex (cf. Lemma 3.3). This implies that $\lambda D^\top \alpha^*$ remains the same for all $\alpha^* \in \Omega^*$. Putting these pieces together yields that there is a pair of vector $(\mathbf{s}^*, \mathbf{g}^*) \in \mathbb{R}^d \times \mathbb{R}^{mn}$ with $\mathbf{g}^* = \lambda D \nabla f^*(\lambda \mathbf{s}^*)$ such that $\alpha^* \in \Omega^*$ implies that $D^\top \alpha^* = \mathbf{s}^*$, $\nabla \bar{f}(\alpha^*) = \mathbf{g}^*$ and $\alpha^* \in \mathbb{B}_q^n$.

On the other hand, if $D^\top \alpha^* = \mathbf{s}^*$, $\nabla \bar{f}(\alpha^*) = \mathbf{g}^*$ and $\alpha^* \in \mathbb{B}_q^n$, then $\lambda D^\top \alpha^*$ remains the same for all $\alpha^* \in \Omega^*$ and $\bar{f}(\alpha^*) = f^*(\lambda \mathbf{s}^*)$ achieves the optimal objective value of problem (6). Putting these pieces together yields that there is a pair of vector $(\mathbf{s}^*, \mathbf{g}^*) \in \mathbb{R}^d \times \mathbb{R}^{mn}$ with $\mathbf{g}^* = \lambda D \nabla f^*(\lambda \mathbf{s}^*)$ such that $D^\top \alpha^* = \mathbf{s}^*$, $\nabla \bar{f}(\alpha^*) = \mathbf{g}^*$ and $\alpha^* \in \mathbb{B}_q^n$ implies that $\alpha^* \in \Omega^*$. As a consequence, the proof of the proposition is achieved. □

3.2 GEB condition and ULC property

In this section, we introduce the ULC property of a set-valued mapping and provide a sufficient condition for the GEB condition, which forms the basis for our subsequent analysis.

In the convergence analysis of optimization algorithms for solving problem (6), it is essential to measure the distance between any iterate α_t and the optimal set Ω^*, i.e., $d(\alpha_t, \Omega^*)$. However, such a quantity is not easily accessible since Ω^* is unknown. As an alternative, we define a function $R : \mathbb{R}^{mn} \to \mathbb{R}^{mn}$, which is called *residual function*, given by

$$R(\alpha) \;:=\; \mathrm{proj}_{\mathbb{B}_q^n}\left(\alpha - \nabla \bar{f}(\alpha)\right) - \alpha. \tag{8}$$

We can verify that $R(\alpha) = 0$ if and only if $\alpha \in \Omega^*$. Moreover, given any $\alpha \in \mathbb{R}^{mn}$, the function $R(\alpha)$ is much easier to be evaluated than $d(\alpha, \Omega^*)$. This suggests that $\|R(\alpha)\|$ can serve as a surrogate measure for the optimality of α. In this case, the remaining thing is to establish a relationship between $\|R(\alpha)\|$ and $d(\alpha, \Omega^*)$.

Definition 7 (GEB Condition)**.** *Problem (6) satisfies a* global error bound *(GEB) condition if there exists a constant $\tau > 0$ such that $d(\alpha, \Omega^*) \leq \tau \|R(\alpha)\|$ for all $\alpha \in \mathbb{B}_q^n$.*

The GEB condition can be interpreted as a relaxed notion of global strong convexity [24]. Indeed, after removing the constraint set $\mathbb{B}_q^n$, we see that $R(\alpha) = -\nabla \bar{f}(\alpha)$ and the GEB condition is satisfied when $\bar{f}$ is strongly convex. However, the residual function R can be very difficult to analyze for problem (6), inspiring us to pursue some other approaches. An useful alternative approach is based on the notion of ULC property of set-valued mappings, which has been used in [43] to analyze the local error bound condition for $\ell_{1,p}$-norm regularized problems.

We begin with some definitions. Let $\mathcal{Y}$ and $\mathcal{Z}$ be two Euclidean spaces. A mapping $\Gamma : \mathcal{Y} \to \mathcal{Z}$ is said to be a *set-valued mapping*, or equivalently, a *multifunction* if for any $y \in \mathcal{Y}$, then $\Gamma(y)$ is a subset of $\mathcal{Z}$. The *graph* of Γ is a subset defined by $\{(y, z) \in \mathcal{Y} \times \mathcal{Z} \mid z \in \Gamma(y)\}$. In what follows, we can define a notion of continuity as follows:

Definition 8 (ULC Property)**.** *A set-valued mapping $\Gamma : \mathcal{Y} \to \mathcal{Z}$ has the* upper Lipschitz continuity *(ULC) property at $y \in \mathcal{Y}$ if $\Gamma(y)$ is nonempty and closed, and there exist constants $\kappa > 0$ and $\delta > 0$ such that for any $y \in \mathcal{Y}$ with $\|y' - y\| \leq \delta$, $\Gamma(y') \subseteq \Gamma(y) + \kappa\|y' - y\|\mathcal{B}$ where $\mathcal{B}$ is the unit ℓ_2-norm ball of $\mathcal{Z}$ and "+" is the Minkowski sum of two sets.*

To proceed, we prove a sufficient condition for the GEB condition to hold. In particular, let Σ be the set-valued mapping defined by

$$\Sigma(\mathbf{s}, \mathbf{g}) \;:=\; \left\{\alpha \in \mathbb{B}_q^n \mid D^\top \alpha = \mathbf{s},\; -\mathbf{g} \in \mathcal{N}_{\mathbb{B}_q^n}(\alpha)\right\} \tag{9}$$

for any $(\mathbf{s}, \mathbf{g}) \in \mathbb{R}^d \times \mathbb{R}^{mn}$. The following proposition characterizes the relationship between the set-valued mapping Σ and the optimal set Ω^*.

Proposition 3.5. *Let $(\mathbf{s}^*, \mathbf{g}^*)$ be given in Proposition 3.4, then $\Omega^* = \Sigma(\mathbf{s}^*, \mathbf{g}^*)$.*

Proof. Since problem (6) is convex, the first-order optimality condition is both necessary and sufficient. Hence, we have

$$\Omega^* \;=\; \left\{\alpha^* \in \mathbb{R}^{mn} \mid \mathbf{0} \in \nabla \bar{f}(\alpha^*) + \mathcal{N}_{\mathbb{B}_q^n}(\alpha^*)\right\}. \tag{10}$$

In what follows, we show that $\alpha^* \in \Omega^* \Rightarrow \alpha^* \in \Sigma(\mathbf{s}^*, \mathbf{g}^*)$. By Proposition 3.4, we have $D^\top \alpha^* = \mathbf{s}^*$, $\nabla \bar{f}(\alpha^*) = \mathbf{g}^*$ and $\alpha^* \in \mathbb{B}_q^n$. Combining this with (10) yields that $\alpha \in \Sigma(\mathbf{s}^*, \mathbf{g}^*)$. Conversely, since $\alpha^* \in \Sigma(\mathbf{s}^*, \mathbf{g}^*)$, then $\mathbf{g}^* = \lambda D \nabla f^*(\lambda \mathbf{s}^*) = \lambda D \nabla f^*(\lambda D^\top \alpha^*) = \nabla \bar{f}(\alpha^*)$. Therefore, we conclude from $-\mathbf{g}^* \in \mathcal{N}_{\mathbb{B}_q^n}(\alpha^*)$ that that $\mathbf{0} \in \nabla \bar{f}(\alpha^*) + \mathcal{N}_{\mathbb{B}_q^n}(\alpha^*)$ and $\alpha^* \in \Omega^*$. This completes the proof of the proposition. □

Given the result of Proposition 3.5, we present the main result showing that the ULC property of Σ implies the GEB condition for problem (6) in the following theorem.

Theorem 3.6. *Let $(\mathbf{s}^*, \mathbf{g}^*) \in \mathbb{R}^d \times \mathbb{R}^{mn}$ be given in Proposition 3.4. Then, the GEB condition is satisfied by problem (6) if the set-valued mapping Σ has the ULC property at $(\mathbf{s}^*, \mathbf{g}^*)$.*

Proof. The proof is based on the techniques in [43] and the property of $\mathbb{B}_q^n$. For the sake of completeness, we provide the details of the proof. In particular, since Σ has the ULC property at $(\mathbf{s}^*, \mathbf{g}^*)$, there exist constants $\kappa > 0$ and $\delta > 0$ such that for all $(\mathbf{s}, \mathbf{g}) \in \mathbb{R}^d \times \mathbb{R}^{mn}$ with $\|(\mathbf{s}, \mathbf{g}) - (\mathbf{s}^*, \mathbf{g}^*)\| \leq \delta$, we have

$$\Sigma(\mathbf{s}, \mathbf{g}) \;\subseteq\; \Sigma(\mathbf{s}^*, \mathbf{g}^*) + \kappa\|(\mathbf{s}, \mathbf{g}) - (\mathbf{s}^*, \mathbf{g}^*)\|\mathbb{B}_2. \tag{11}$$

Recalling the residual function in (8) that $R(\alpha) = \mathrm{proj}_{\mathbb{B}_q^n}\left(\alpha - \nabla \bar{f}(\alpha)\right) - \alpha$ and defining two functions $\mathbf{s}^+ : \mathbb{B}_q^n \to \mathbb{R}^d$ and $\mathbf{g}^+ : \mathbb{B}_q^n \to \mathbb{R}^{mn}$ given by

$$\mathbf{s}^+(\alpha) \;:=\; D^\top\left(\alpha + R(\alpha)\right), \qquad \mathbf{g}^+(\alpha) \;:=\; \nabla \bar{f}(\alpha) + R(\alpha). \tag{12}$$

Since $\mathbb{B}_q^n$ is convex and compact, R is Lipschitz continuous [29]. Additionally, $\nabla \bar{f}$ is Lipschitz continuous. Thus, we conclude that $\mathbf{s}^+$ and $\mathbf{g}^+$ are both Lipschitz continuous. This together with Proposition 3.5 implies that there exists a constant $\rho > 0$ such that,

$$\left\|\left(\mathbf{s}^+(\alpha), \mathbf{g}^+(\alpha)\right) - (\mathbf{s}^*, \mathbf{g}^*)\right\| \;\le\; \delta, \quad \forall \alpha \in \mathbb{B}_q^n \cap \{d(\alpha, \Omega^*) \le \rho\}. \tag{13}$$

By the definition of the residual function R, we have

$$\alpha + R(\alpha) \;=\; \underset{z \in \mathbb{B}_q^n}{\mathrm{argmin}} \left\{ \left\langle \nabla \bar{f}(\alpha), z \right\rangle + \frac{1}{2}\left\|z - \alpha\right\|^2 \right\}.$$

By the optimality condition, we have

$$-\nabla \bar{f}(\alpha) - R(\alpha) \;\in\; \mathcal{N}_{\mathbb{B}_q^n}\left(\alpha + R(\alpha)\right). \tag{14}$$

Combining (12) and (14) yields that $\alpha + R(\alpha) \in \Sigma(\mathbf{s}^+(\alpha), \mathbf{g}^+(\alpha))$ for all $\alpha \in \mathbb{B}_q^n$. This together with (11) and (13) yields that

$$d\left(\alpha + R(\alpha), \Sigma(\mathbf{s}^*, \mathbf{g}^*)\right) \;\le\; \kappa \|(\mathbf{s}^+(\alpha), \mathbf{g}^+(\alpha)) - (\mathbf{s}^*, \mathbf{g}^*)\|, \quad \forall \alpha \in \mathbb{B}_q^n \cap \{d(\alpha, \Omega^*) \le \rho\}.$$

Recalling the fact that $\nabla \bar{f}(\alpha) = \lambda D \nabla f^*(\lambda D^\top \alpha)$ and $\mathbf{g}^* = \lambda D \nabla f^*(\lambda \mathbf{s}^*)$, we have

$$\begin{aligned} \|\mathbf{s}^+(\alpha) - \mathbf{s}^*\| &\le \|D^\top \alpha - \mathbf{s}^*\| + \sqrt{\sigma}\|R(\alpha)\|, \\ \|\mathbf{g}^+(\alpha) - \mathbf{g}^*\| &\le \frac{\sqrt{\sigma}\lambda^2}{\mu}\|D^\top \alpha - \mathbf{s}^*\| + \|R(\alpha)\|. \end{aligned}$$

Furthermore, in view of Proposition 3.5, $d(\alpha, \Omega^*) \le d(\alpha + R(\alpha), \Sigma(\mathbf{s}^*, \mathbf{g}^*)) + \|R(\alpha)\|$. Putting these pieces together yields that, for all $\alpha \in \mathbb{B}_q^n \cap \{d(\alpha, \Omega^*) \le \rho\}$,

$$d(\alpha, \Omega^*) \;\le\; \left(\kappa + \frac{\sqrt{\sigma}\lambda^2\kappa}{\mu}\right)\|D^\top \alpha - \mathbf{s}^*\| + \left(\sqrt{\sigma}\kappa + \kappa + 1\right)\|R(\alpha)\|.$$

Letting $\kappa_0 = \max\{\kappa + \frac{\sqrt{\sigma}\lambda^2\kappa}{\mu}, \sqrt{\sigma}\kappa + \kappa + 1\}$ and using the inequality $(a+b)^2 \le 2(a^2 + b^2)$ yields that, for all $\alpha \in \mathbb{B}_q^n \cap \{d(\alpha, \Omega^*) \le \rho\}$,

$$d^2(\alpha, \Omega^*) \;\le\; 2\kappa_0^2\left(\|D^\top \alpha - \mathbf{s}^*\|^2 + \|R(\alpha)\|^2\right). \tag{15}$$

Since f^* is $\frac{1}{\ell}$-strongly convex, we have

$$\begin{aligned} \|D^\top \alpha - \mathbf{s}^*\|^2 \;=\; \frac{\|\lambda D^\top \alpha - \lambda \mathbf{s}^*\|^2}{\lambda^2} &\le \frac{\ell}{\lambda^2}\left\langle \nabla f^*(\lambda D^\top \alpha) - \nabla f^*(\lambda \mathbf{s}^*), \lambda D^\top \alpha - \lambda \mathbf{s}^* \right\rangle \\ &= \frac{\ell}{\lambda^2}\left\langle \nabla \bar{f}(\alpha) - \mathbf{g}^*, \alpha - \alpha^* \right\rangle, \end{aligned} \tag{16}$$

where α^* is the projection of α onto Ω^*. Furthermore, by the definition of the normal cone, for all $u \in \mathcal{N}_{\mathbb{B}_q^n}(\alpha + R(\alpha))$ and for all $v \in \mathcal{N}_{\mathbb{B}_q^n}(\alpha^*)$, we have

$$\langle u - v, \alpha + R(\alpha) - \alpha^* \rangle \;\ge\; 0.$$

Taking $u = -\nabla \bar{f}(\alpha) - R(\alpha)$ and $v = -\mathbf{g}^*$ together with the optimality of α^* and (14) yields that

$$\left\langle \nabla \bar{f}(\alpha) - \mathbf{g}^*, \alpha - \alpha^* \right\rangle + \|R(\alpha)\|^2 \;\leq\; \left\langle \mathbf{g}^* - \nabla \bar{f}(\alpha) + \alpha^* - \alpha, R(\alpha) \right\rangle.$$

Since $\|R(\alpha)\|^2 \geq 0$ and $\nabla \bar{f}$ is Lipschitz continuous, there exists a constant $\kappa_1 > 0$ such that $\left\langle \nabla \bar{f}(\alpha) - \mathbf{g}^*, \alpha - \alpha^* \right\rangle \leq \kappa_1 \|\alpha - \alpha^*\| \|R(\alpha)\|$. Combining this with (15) and (16) yields that there exists a constant $\kappa_2 > 0$ such that

$$d^2(\alpha, \Omega^*) \;\leq\; \kappa_2 \left(\|\alpha - \alpha^*\| \|R(\alpha)\| + \|R(\alpha)\|^2 \right), \quad \forall \alpha \in \mathbb{B}_q^n \cap \{d(\alpha, \Omega^*) \leq \rho\}.$$

Upon solving this quadratic inequality yields that there exists a constant $\kappa_3 > 0$ such that

$$d(\alpha, \Omega^*) \;\leq\; \kappa_3 \|R(\alpha)\|, \quad \forall \alpha \in \mathbb{B}_q^n \cap \{d(\alpha, \Omega^*) \leq \rho\}.$$

Furthermore, since $\mathbb{B}_q^n$ is convex and compact and the function $h(\alpha) = \frac{d(\alpha,\Omega^*)}{\|R(\alpha)\|}$ is finite and continuous over $\alpha \in \mathbb{B}_q^n \cap \{d(\alpha, \Omega^*) > \rho\}$, there exists a constant $\kappa_4 > 0$ such that

$$d(\alpha, \Omega^*) \;\leq\; \kappa_4 \|R(\alpha)\|, \quad \forall \alpha \in \mathbb{B}_q^n \cap \{d(\alpha, \Omega^*) > \rho\}.$$

Letting $\tau = \max\{\kappa_3, \kappa_4\}$, we conclude that $d(\alpha, \Omega^*) \leq \tau \|R(\alpha)\|$ for all $\alpha \in \mathbb{B}_q^n$ and hence the GEB condition is satisfied by problem (6). As a consequence, we achieve the conclusion of the theorem. □

Equipped with the result of Theorem 3.6, in Section 3.3, we will establish the GEB condition for the dual filtering-clustering problem (6) when $q \in [1, 2] \cup \{+\infty\}$.

3.3 GEB condition holds when $q = 1$ or $q = +\infty$

In this section, we show that Σ has the ULC property when $q = 1$ or $q = +\infty$.

Lemma 3.7. *Assume that $q \in \{1, +\infty\}$. Then, the set-valued mapping Σ is a polyhedral multifunction.*

Proof. Since $q \in \{1, +\infty\}$, $\mathbb{B}_q^n$ is a polyhedron. Therefore, the indicator function for $\mathbb{B}_q^n$ has a polyhedral epigraph. In addition, by the definition, the normal cone $\mathcal{N}_{\mathbb{B}_q^n}$ is the subdifferential of the indicator function for $\mathbb{B}_q^n$. Putting these pieces together with [43, Lemma 2] yields that the set-valued mapping Σ is a polyhedral multifunction. □

Based on the ULC property of Σ, we are ready to prove the existence of global error bound for problem (6).

Theorem 3.8. *The GEB condition is satisfied by problem (6) when $q = 1$ or $q = +\infty$.*

Proof. By the definition, both $\mathbb{B}_1$ and $\mathbb{B}_\infty$ are polyhedra. Thus, it follows from Lemma 3.7 that the set-valued mapping Σ is a polyhedral multifunction when $q = 1$ and $q = +\infty$. Hence, by [43, Lemma 1], Σ has the ULC property at $(\mathbf{s}^*, \mathbf{g}^*) \in \mathbb{R}^d \times \mathbb{R}^{mn}$ if $\Sigma(\mathbf{s}^*, \mathbf{g}^*)$ is nonempty. Putting these pieces together with Theorem 3.6 yields the desired result. □

3.4 GEB condition holds when $q \in (1, 2]$

Now, we show that the GEB condition holds for problem (6) when $q \in (1, 2]$. In particular, it suffices to establish the ULC property of the set-valued mapping Σ. Before that, we state several technical results that will be used for proving the main theorem. The first lemma is concerning the linear regularity of a collection of polyhedral sets; see the detailed proof in [2, Corollary 5.26].

Lemma 3.9. *Let $\mathcal{S}_1, \ldots, \mathcal{S}_M$ be a collection of polyhedra in $\mathbb{R}^{mn}$. Then, there exists a constant $\bar{\kappa} > 0$ such that $d\left(\alpha, \bigcap_{i=1}^M \mathcal{S}_i\right) \leq \bar{\kappa} \sum_{i=1}^M d(\alpha, \mathcal{S}_i)$ for any $\alpha \in \mathbb{R}^{mn}$.*

The next proposition provides a detailed representation of $\Omega^* = \Sigma(\mathbf{s}^*, \mathbf{g}^*)$. In particular, we consider two cases: $\mathbf{g}_j^* = \mathbf{0}$ or $\mathbf{g}_j^* \neq \mathbf{0}$ and let $\mathcal{J} = \{j \in \{1, \ldots, n\} \mid \mathbf{g}_j^* = \mathbf{0}\}$.

Proposition 3.10. *Suppose that the set-valued mapping Σ is defined in (9) and $(\mathbf{s}^*, \mathbf{g}^*) \in \mathbb{R}^d \times \mathbb{R}^{mn}$ be given in Proposition 3.4. If $j \in \mathcal{J}$, we have $\alpha_j^* = -v(\mathbf{g}_j^*)/\|v(\mathbf{g}_j^*)\|_q$ where the function $v : \mathbb{R}^m \to \mathbb{R}^m$ is defined by $v(\mathbf{g}) := \left(\mathrm{sign}(g_1)|g_1|^{\frac{p}{q}}, \ldots, \mathrm{sign}(g_m)|g_m|^{\frac{p}{q}}\right)$. Otherwise, we have $\{\alpha_j^* \in \mathbb{R}^m \mid D^\top \alpha^* = \mathbf{s}^*\} \subseteq \mathbb{B}_q$. That is to say,*

$$\Sigma(\mathbf{s}^*, \mathbf{g}^*) \;=\; \left\{ \alpha^* = \begin{bmatrix} \alpha_1^* \\ \vdots \\ \alpha_n^* \end{bmatrix} \in \mathbb{R}^{mn} \;\middle|\; \begin{array}{l} D^\top \alpha^* = \mathbf{s}^*, \\ \alpha_j^* = -\frac{v(\mathbf{g}_j^*)}{\|v(\mathbf{g}_j^*)\|_q}, \;\; \forall j \in \mathcal{J} \end{array} \right\},$$

Finally, for all $q \in (1, 2]$, there exist constants $\delta > 0$ and $\nu > 0$ such that for all $\mathbf{g} \in \mathbb{R}^m$ satisfying $\|\mathbf{g} - \mathbf{g}^*\| \leq \delta$, we have

$$\|v(\mathbf{g}) - v(\mathbf{g}^*)\| \;\leq\; \nu \|\mathbf{g} - \mathbf{g}^*\|.$$

Proof. By the definition of $(\mathbf{s}^*, \mathbf{g}^*) \in \mathbb{R}^d \times \mathbb{R}^{mn}$ in Proposition 3.4, $\mathbf{g}_j^* = \mathbf{0}$ refers to the case that the optimal set of problem (6) without the constraint over the j-th block is still contained in $\mathbb{B}_q^n$. Putting these pieces together with Proposition 3.5 yields that

$$\{\alpha_j^* \in \mathbb{R}^m \mid D^\top \alpha^* = \mathbf{s}^*\} \;\subseteq\; \mathbb{B}_q.$$

If $\mathbf{g}_j^* \neq \mathbf{0}$, then $\alpha^* \in \Omega^*$ if and only if there exists $\mu \geq 0$ such that

$$\begin{aligned} 1 - \|\alpha_j^*\|_q &\geq 0, \\ \mathbf{s}^* - D^\top \alpha^* &= \mathbf{0}, \\ \mathbf{g}_j^* + \mu \cdot \frac{\left(\mathrm{sign}((\alpha_j^*)_1)|(\alpha_j^*)_1|^{q-1}, \ldots, \mathrm{sign}((\alpha_j^*)_m)|(\alpha_j^*)_m|^{q-1}\right)}{\|\alpha_j^*\|_q^{q/p}} &= \mathbf{0}, \\ \mu \left(\|\alpha_j^*\|_q - 1\right) &= 0. \end{aligned}$$

Since $\mathbf{g}_j^* \neq \mathbf{0}$, we have $\mu > 0$ and $\|\alpha_j^*\|_q = 1$. Then, we can solve α_j^* in terms of $\mathbf{g}_j^*$ by using the above relationship and obtain that

$$\mathbf{s}^* - D^\top \alpha^* \;=\; \mathbf{0}, \qquad \alpha_j^* + \frac{v(\mathbf{g}_j^*)}{\|v(\mathbf{g}_j^*)\|_q} \;=\; \mathbf{0}.$$

Putting these pieces together yields that

$$\Sigma(\mathbf{s}^*, \mathbf{g}^*) \;=\; \left\{ \alpha^* = \begin{bmatrix} \alpha_1^* \\ \vdots \\ \alpha_n^* \end{bmatrix} \in \mathbb{R}^{mn} \;\middle|\; \begin{array}{l} D^\top \alpha^* = \mathbf{s}^* \\ \alpha_j^* = -\frac{v(\mathbf{g}_j^*)}{\|v(\mathbf{g}_j^*)\|_q}, \; \forall j \in \mathcal{J} \end{array} \right\}.$$

Finally, if $q \in (1, 2]$, then $p/q \geq 1$. In this case, the function $s \mapsto \text{sign}(s)|s|^{\frac{p}{q}}$ is continuously differentiable and hence locally Lipschitz. Thus, there exist constants $\nu > 0$ and $\delta > 0$ such that for all $s_1, s_2 \in \mathbb{R}$ satisfying $|s_1 - s_2| \leq \delta$, we achieve that

$$\left| \text{sign}(s_1)|s_1|^{\frac{p}{q}} - \text{sign}(s_2)|s_2|^{\frac{p}{q}} \right| \;\leq\; \nu \left| s_1 - s_2 \right|.$$

Therefore, we conclude that $\|v(\mathbf{g}) - v(\mathbf{g}^*)\| \;\leq\; \nu \|\mathbf{g} - \mathbf{g}^*\|$ for all $g \in \mathbb{R}^m$. As a consequence, we reach the conclusion of the proposition. □

The above proposition shows that $\Sigma(\mathbf{s}^*, \mathbf{g}^*)$ is closed. Since $\Sigma(\mathbf{s}^*, \mathbf{g}^*) \subseteq \mathbb{B}_q^n$, then $\Sigma(\mathbf{s}^*, \mathbf{g}^*)$ is bounded where $(\mathbf{s}^*, \mathbf{g}^*) \in \mathbb{R}^d \times \mathbb{R}^{mn}$ be given in Proposition 3.4. Given the above results, we are ready to study the ULC property of the set-valued mapping Σ.

Theorem 3.11. *The GEB condition is satisfied by problem (6) when $q \in (1, 2]$.*

Proof. Define the sets $\mathcal{S}_1$ and $\mathcal{S}_2$ as

$$\mathcal{S} \;:=\; \left\{ \alpha^* \in \mathbb{R}^{mn} \mid D^\top \alpha^* = \mathbf{s}^* \right\}, \quad \mathcal{S}_j \;:=\; \left\{ \alpha^* = \begin{bmatrix} \alpha_1^* \\ \vdots \\ \alpha_n^* \end{bmatrix} \in \mathbb{R}^{mn} \;\middle|\; \alpha_j^* = -\frac{v(\mathbf{g}_j^*)}{\|v(\mathbf{g}_j^*)\|_q} \right\}.$$

Then, by Proposition 3.10, we have

$$\Sigma(\mathbf{s}^*, \mathbf{g}^*) \;=\; \mathcal{S} \cap (\cap_{j \in \mathcal{J}} \mathcal{S}_j)$$

Moreover, $\mathcal{S}$ and $\{\mathcal{S}_j, \; \forall j \in \mathcal{J}\}$ are all polyhedral subsets of $\mathbb{R}^{mn}$. Thus, by Lemma 3.9, there exists a constant $\bar{\kappa} > 0$ such that for any $\alpha \in \mathbb{R}^{mn}$,

$$d\left(\alpha, \Sigma(\mathbf{s}^*, \mathbf{g}^*)\right) \;\leq\; \bar{\kappa} \left(d(\alpha, \mathcal{S}) + \sum_{j \in \mathcal{J}} d(\alpha, \mathcal{S}_j) \right). \tag{17}$$

Thus, it suffices to bound the right-hand side of (17) for all $\alpha \in \Sigma(\mathbf{s}, \mathbf{g})$ satisfying that $(\mathbf{s}, \mathbf{g})$ lies in the neighborhood of $(\mathbf{s}^*, \mathbf{g}^*) \in \mathbb{R}^d \times \mathbb{R}^{mn}$ and $\Sigma(\mathbf{s}, \mathbf{g})$ is nonempty. Towards that end, we discuss the bound on $d(\alpha, \mathcal{S})$ and $d(\alpha, \mathcal{S}_j)$ separately.

The bound on $d(\alpha, \mathcal{S}_1)$ follows from the well-known Hoffman bound [13]. In particular, there exists a constant $\theta > 0$ such that $d(\alpha, \mathcal{S}_1) \leq \theta \|D^\top \alpha - \mathbf{s}^*\|$ for any $\alpha \in \mathbb{R}^{mn}$. In addition, for all $\alpha \in \Sigma(\mathbf{s}, \mathbf{g})$ with $\mathbf{g} \neq \mathbf{0}$, we have $D^\top \alpha = \mathbf{s}$. Putting these pieces together yields that

$$d(\alpha, \mathcal{S}) \;\leq\; \theta \|\mathbf{s} - \mathbf{s}^*\|, \quad \forall \alpha \in \Sigma(\mathbf{s}, \mathbf{g}).$$

Since $\mathbf{g}_j^* \neq \mathbf{0}$ for $\forall j \in \mathcal{J}$, there exists a constant $\delta_j > 0$ such that $\|(\mathbf{s}, \mathbf{g}) - (\mathbf{s}^*, \mathbf{g}^*)\| \leq \delta_j$ implies $\mathbf{g}_j \neq \mathbf{0}$. Therefore, for any $\alpha \in \Sigma(\mathbf{s}, \mathbf{g})$ that satisfies $\|(\mathbf{s}, \mathbf{g}) - (\mathbf{s}^*, \mathbf{g}^*)\| \leq \delta_j$ must satisfy the following conditions:

$$\mathbf{s} - D^\top \alpha \;=\; \mathbf{0}, \qquad \alpha_j + \frac{v(\mathbf{g}_j)}{\|v(\mathbf{g}_j)\|_q} \;=\; \mathbf{0}.$$

Since $\mathbf{g}_j, \mathbf{g}_j^* \neq \mathbf{0}$, we have $\|v(\mathbf{g}_j)\|_q, \|v(\mathbf{g}_j^*)\|_q > 0$ and

$$d(\alpha, \mathcal{S}_j) \;\leq\; \left\| \frac{v(\mathbf{g}_j)}{\|v(\mathbf{g}_j)\|_q} - \frac{v(\mathbf{g}_j^*)}{\|v(\mathbf{g}_j^*)\|_q} \right\|.$$

This implies that

$$\left\| \frac{v(\mathbf{g}_j)}{\|v(\mathbf{g}_j)\|_q} - \frac{v(\mathbf{g}_j^*)}{\|v(\mathbf{g}_j^*)\|_q} \right\| \;=\; \frac{\|v(\mathbf{g}_j) - v(\mathbf{g}_j^*)\|_q \|v(\mathbf{g}_j)\| + \|v(\mathbf{g}_j) - v(\mathbf{g}_j^*)\| \|v(\mathbf{g}_j)\|_q}{\|v(\mathbf{g}_j)\|_q \|v(\mathbf{g}_j^*)\|_q}.$$

Furthermore, since $\|(\mathbf{s}, \mathbf{g}) - (\mathbf{s}^*, \mathbf{g}^*)\| \leq \delta_j$, then $\|v(\mathbf{g}_j)\|$ and $\|\|v(\mathbf{g}_j)\|_q$ are both bounded. Since $q \in (1, 2]$, then the Hólder inequality implies that $\|v(\mathbf{g}_j) - v(\mathbf{g}_j^*)\|_q \leq m^{\frac{2-q}{2q}} \|v(\mathbf{g}_j) - v(\mathbf{g}_j^*)\|$. Putting these pieces together yields that

$$d(\alpha, \mathcal{S}_j) \;\leq\; C\left(m^{\frac{2-q}{2q}} + 1\right) \|v(\mathbf{g}_j) - v(\mathbf{g}_j^*)\| \;\overset{\text{Proposition 3.10}}{\leq}\; C\nu_j \left(m^{\frac{2-q}{2q}} + 1\right) \|\mathbf{g}_j - \mathbf{g}_j^*\|.$$

Therefore, we obtain that

$$d\left(\alpha, \Sigma(\mathbf{s}^*, \mathbf{g}^*)\right) \;\leq\; \left(\bar{\kappa} \max\left\{\theta, C\left(m^{\frac{2-q}{2q}} + 1\right)\left(\sum_{j \in \mathcal{J}} \nu_j\right)\right\}\right) \|\,(\mathbf{s}, \mathbf{g}) - (\mathbf{s}^*, \mathbf{g}^*)\,\|$$

for any $\alpha \in \Sigma(\mathbf{s}, \mathbf{g})$ with $\|(\mathbf{s}, \mathbf{g}) - (\mathbf{s}^*, \mathbf{g}^*)\| \leq \min_{j \in \mathcal{J}} \delta_j$. This implies that Σ has the ULC property at $(\mathbf{s}^*, \mathbf{g}^*) \in \mathbb{R}^d \times \mathbb{R}^{mn}$, which completes the proof the theorem. □

4 Algorithmic Framework

In this section, we analyze a *generalized dual gradient ascent* (GDGA) algorithm for solving the filtering-clustering problems (1). In particular, we prove that the proposed approach is linear convergent without considering the number of gradient or stochastic gradient oracles used in the subroutine, which will be carefully analyzed for different scenarios in Section 5.

4.1 Generalized dual gradient ascent

The GDGA algorithmic framework is summarized in Algorithm 1. It is worthy noting that the information we can only access is the gradient of f, the matrix D and the parameter λ. Roughly speaking, this framework can be simply seen as an inexact gradient ascent for solving the dual problem (6). Indeed, since f^* is inaccessible for general f, we need to design a subroutine and approximately solve $f(\beta) - \lambda\alpha_t^\top D\beta$ to get an inexact gradient of f^*. In what follows, we provide some comments on the GDGA algorithmic framework.

Firstly, this framework has a solid theoretical guarantee. In particular, the algorithm has linear convergence without considering the number of gradient or stochastic gradient oracles used in the subroutine (see Section 4.2). This complexity analysis is based on the fact that problem (6) has a global error bound (see Section 3).

Secondly, the subroutine can be constructed based on the different efficient algorithms as mentioned before. Indeed, we show that the total complexity of the GDGA algorithm with the subroutines is near-optimal in deterministic, finite-sum, and online settings. Furthermore, if f is in some special form, such as least squared loss, this subroutine can even be removed since the exact minimizer of $f(\beta) - \lambda\alpha_t^\top D\beta$ is available.

Algorithm 1: Generalized Dual Gradient Ascent (GDGA)

Input: (β_0, α_0), learning rates $\eta > 0$.
for $t = 0, 1, \ldots, T$ **do**
  Find $\beta_{t+1} \in \mathbb{R}^d$ such that β_{t+1} is an $\hat{\varepsilon}$-minimizer of $f(\beta) - \lambda \alpha_t^\top D\beta$ using $(\nabla f, D, \beta_t)$.
  $\alpha_{t+1} \leftarrow \mathrm{proj}_{\mathbb{B}_q^n}(\alpha_t - \eta\lambda D\beta_{t+1})$.
end for
Return: β_{T+1}.

Finally, the proposed algorithm is simple and *matrix-free*. Thus, it can be easily implemented in distributed setting and amenable to large-scale filtering-clustering problems. This makes the proposed algorithm intrinsically different from PDIP [14], the projected Newton algorithm [39], ADMM [6, 28] and the semismooth Newton algorithm [41].

4.2 Complexity of GDGA algorithmic framework

In this subsection, we establish the main result on the linear convergence of the GDGA algorithmic framework without considering the number of gradient or stochastic gradient oracles used in the subroutine. The proof is based on the global error bound of problem (6) (cf. Theorem 3.8 and Theorem 3.11). Note that, this result is surprising for the filtering-clustering problem due to the following reasons:

1. Despite the regularity of f (cf. Assumption 2.1), problem (1) is nonsmooth without a computationally friendly nonsmooth terms. More specifically, the proximal mapping of $\|D_j\beta\|_p$ can not be easily computed in general and hence the proximal algorithms [25] are not applicable in practice.

2. Despite the regularity of f^* (cf. Lemma 3.3), problem (6) is not strongly convex. Furthermore, ∇f^* is not available. Thus, it is not obvious if the dual gradient ascent can be applicable with linear convergence.

3. Despite the objective in problem (5) is convex-concave, it is not strongly convex-concave yet. Thus, the gradient descent ascent (GDA) can not be proven linearly convergent using the existing theory [23, 36, 8].

We denote $\bar{\alpha}_t$ as the projection of α_t onto the optimal set of problem (6), i.e., Ω^*, and trace the objective gap between $\bar{f}(\alpha_t)$ and $\bar{f}(\bar{\alpha}_t)$ by defining

$$\Delta_t := \bar{f}(\alpha_t) - \bar{f}(\bar{\alpha}_t).$$

Theorem 4.1 (Complexity of GDGA algorithm)**.** *Let $\eta \in \left(0, \frac{\mu}{4\sigma \max\{1,\lambda^2\}}\right)$ in Algorithm 1 and τ be defined in Theorem 3.8 and Theorem 3.11. Given any tolerance $\varepsilon > 0$, let $\hat{\varepsilon}$ satisfy*

$$\hat{\varepsilon} \;\leq\; \min\left\{\frac{\sqrt{\varepsilon}}{2},\; \frac{\sqrt{\varepsilon}\lambda\mu}{4\sqrt{\ell}}\sqrt{\frac{\tau}{C(17\tau^2 + (14+\eta\lambda^2)\tau + 1)}}\right\} \tag{18}$$

where $C = \frac{2\tau^2\lambda^4\sigma}{17\tau^2+14\tau+1} + \frac{\sigma}{2}$. Then the number of iterations required by the GDGA algorithm to find an ε-optimal solution is bounded by

$$N \;\leq\; \left(\frac{17\tau^2 + 14\tau + 1}{\tau\lambda^2\eta}\right)\log\left(\frac{16\ell\Delta_0}{\mu^2\varepsilon}\right),$$

Proof. Denote $\delta_t = \|\beta_t - \beta^*(\alpha_t)\|$ and β^* as an optimal solution to filtering-clustering problem (1). Then we have

$$\|\beta_t - \beta^*\|^2 \;\le\; 2\left(\|\beta_t - \beta^*(\alpha_t)\|^2 + \|\beta^*(\alpha_t) - \beta^*\|^2\right) \;=\; 2\left(\delta_t^2 + \|\beta^*(\alpha_t) - \beta^*\|^2\right). \tag{19}$$

Letting $\bar{\alpha}_t$ be the projection of α_t onto the optimal set. By the uniqueness of the optimal solution to problem (1), we have $\beta^* = \beta^*(\bar{\alpha}_t)$. Then we have

$$\begin{aligned}
\|\beta^*(\alpha_t) - \beta^*\|^2 &\overset{\text{Lemma 3.2}}{\le} \frac{\left\|\lambda D^\top \alpha_t - \lambda D^\top \bar{\alpha}_t\right\|^2}{\mu^2} \\
&\overset{\text{Lemma 3.1}}{\le} \frac{2\ell\left(f^*(\lambda D^\top \alpha_t) - f^*(\lambda D^\top \bar{\alpha}_t)\right)}{\mu^2} \\
&= \frac{2\ell\Delta_t}{\mu^2}.
\end{aligned}$$

Since $\Delta_t \le \frac{\mu^2\varepsilon}{8\ell}$ if $t \ge \left(\frac{17\tau^2+14\tau+1}{\tau\lambda^2\eta}\right)\log\left(\frac{16\ell\Delta_0}{\mu^2\varepsilon}\right)$ (cf. Lemma 7.3), we have

$$\|\beta^*(\alpha_t) - \beta^*\|^2 \;\le\; \frac{\varepsilon}{4}.$$

By the definition of $\hat{\varepsilon}$, we have $\delta_t^2 \le \frac{\varepsilon}{4}$. Putting these pieces together with (19) implies that $\|\beta_t - \beta^*\|^2 \le \varepsilon$. We achieve the conclusion of the theorem. □

Finally, we proceed to a corollary concerning about the stochastic setting. In particular, the GDGA algorithm is intrinsically stochastic if the subroutines are based on the stochastic gradient-type algorithms, e.g., Katyusha and SGD. However, it does not affect the complexity bound of the iteration numbers without considering the number of stochastic gradient oracles used in the subroutine.

Corollary 4.2. *Under the same setting as Theorem 4.1 but the subroutine is based on stochastic algorithm, the number of iterations required by the GDGA algorithm to find a ε-optimal solution is bounded by*

$$N \;\le\; \left(\frac{17\tau^2 + 14\tau + 1}{\tau\lambda^2\eta}\right)\log\left(\frac{16\ell\Delta_0}{\mu^2\varepsilon}\right),$$

Proof. The proof is nearly the same as that in Theorem 4.1. In particular, we take expectation of both sides of (19) and use Lemma 7.6 instead of Lemma 7.3. □

5 Different Variants of GDGA Algorithm

In this section, we consider the subroutines in the GDGA algorithmic framework in different scenarios. In particular, we remark that the subroutines are unnecessary if f is in the special form. For example, the minimizer of $f(\beta) - \alpha_t^\top D\beta$ is available if f is $\|\cdot\|^2$, which is commonly used in real applications of filtering-clustering problems, such as trend filtering and ℓ_1 convex clustering. This leads to the simplified GDGA algorithm (Algorithm 2) with the direct complexity analysis; see Section 5.1.

For the general case of f, we can design the subroutines by applying the appropriate gradient-type optimization methods in different scenarios, e.g., the accelerated gradient descent (AGD) [20] for the deterministic setting, the accelerated stochastic variance reduced gradient (Katyusha) [1] for the finite-sum setting and the stochastic gradient descent (SGD) [27] for the online setting. The complexity analyses based on the regularity of f are presented in Sections 5.2, 5.3 and 5.4.

Algorithm 2: Simplified GDGA algorithm

Input: (β_{-1}, α_0), learning rates $\eta > 0$.
for $t = 0, 1, \ldots$ **do**
  $\beta_t = 0.5\lambda D^\top \alpha_t$.
  $\alpha_{t+1} \leftarrow \mathrm{proj}_{\mathbb{B}_q^n}(\alpha_t - \eta\lambda D\beta_t)$.
end for
Return: β_T.

5.1 Simplified GDGA algorithm

In this subsection, we focus on the case that f is the squared ℓ_2 norm, i.e., $f(\beta) = (1/2)\|\beta\|^2$. This particular case is crucial since the squared loss function arises from the specific setting of trend filtering and convex clustering. By exploiting this special structure of f, we arrive at the simplified GDGA algorithm by removing the subroutine in the GDGA algorithm but using the exact minimizer instead. The resulting algorithm for convex clustering with ℓ_2-regularization recovers the algorithm in [38]. We present the *simplified GDGA* algorithm in Algorithm 2.

The complexity bound of the simplified GDGA algorithm can be directly obtained by Theorem 4.1. In particular, the subroutine is removed since we do not need to approximately solve $f(\beta) - \alpha_t^\top \lambda D\beta$. This implies that the per-iteration computational cost will not depend on $\hat{\varepsilon}$. We summarize the result in the following theorem.

Theorem 5.1 (Complexity bound of simplified GDGA algorithm)**.** *Let the step size $\eta > 0$ in Algorithm 2 satisfy $\eta \in \left(0, \frac{\mu}{4\sigma \max\{1,\lambda^2\}}\right)$. Then, for any $\varepsilon > 0$, the number of iterations $\widetilde{N}_{total}$ for Algorithm 2 to to find an ε-optimal solution is bounded by*

$$\widetilde{N}_{total} \;\leq\; \left(\frac{17\tau^2 + 14\tau + 1}{\tau\lambda^2\eta}\right) \log\left(\frac{16\ell\Delta_0}{\mu^2\varepsilon}\right).$$

The proof of Theorem 5.1 is similar to that used for Theorem 4.1; therefore, it is omitted. The result of Theorem 5.1 implies that the simplified GDGA algorithm (Algorithm 2) has linear convergence. Later, in experiment section (Section 6), we provide careful simulation studies with simplified GDGA algorithm on various real datasets with applications to trend filtering and convex clustering and compare its performance with state-of-the-art baseline optimization methods in these problems.

5.2 Deterministic GDGA algorithm

In this subsection, we focus on the deterministic setting in which the gradient oracle ∇f is used in each iteration of the subroutine in Algorithm 1. Different from the simplified case discussed in Section 5.1, the minimizer of $f(\beta) - \alpha_t^\top D\beta$ is not available. Instead, we obtain an $\hat{\varepsilon}$-minimizer by applying the subroutine based on the accelerated gradient descent (AGD) algorithm [20]. The resulting algorithm based on that subroutine is termed as *deterministic GDGA* algorithm. We provide the pseudocode of that algorithm in Algorithm 3.

The complexity bound of the deterministic GDGA algorithm is obtained by combining Theorem 4.1 and the complexity bound of the AGD algorithm in terms of gradient oracles. We summarize the complexity bound of the subroutine based on the AGD algorithm in the following lemma.

Algorithm 3: Deterministic GDGA algorithm

Input: (β_{-1}, α_0), learning rates $\eta > 0$.
for $t = 0, 1, \ldots$ **do**
  Find $\beta_t \in \mathbb{R}^d$ such that β_t is an $\hat{\varepsilon}$-minimizer of $f(\beta) - \lambda \alpha_t^\top D\beta$ using AGD algorithm.
  $\alpha_{t+1} \leftarrow \mathrm{proj}_{\mathbb{B}_q^n}(\alpha_t - \eta\lambda D\beta_t)$.
end for
Return: β_T.

Lemma 5.2. *Let $\hat{\varepsilon} \in (0,1)$ be given in deterministic GDGA algorithm. Then, the number of gradient oracles to reach $\|\beta_t - \beta^*(\alpha_t)\| \le \hat{\varepsilon}$ is bounded by*

$$N_t \le \begin{cases} \sqrt{\kappa}\log\left(\frac{\|\beta_{-1}-\beta^*(\alpha_0)\|}{\hat{\varepsilon}}\right) & t = 0, \\ \sqrt{\kappa}\log\left(\frac{1+\lambda\sqrt{\sigma}D_q/\mu}{\hat{\varepsilon}}\right) & t \ge 1. \end{cases} \tag{20}$$

The proof of Lemma 5.2 is deferred to Section 7.2. Equipped with the result of that lemma, we are ready to present the main result on the complexity bound of the deterministic GDGA algorithm in terms of the number of gradient oracles.

Theorem 5.3 (Complexity of deterministic GDGA algorithm)**.** *Let the step size $\eta > 0$ satisfy $\eta \in (0, \mu/4\sigma)$ in the deterministic GDGA algorithm. Then, for any $\varepsilon > 0$, the number of gradient oracles $\widetilde{N}_{grad}$ for the deterministic GDGA algorithm to to find an ε-optimal solution is bounded by*

$$\widetilde{N}_{grad} \le \left(\frac{\sqrt{\kappa}\left(17\tau^2 + 14\tau + 1\right)}{\tau\lambda^2\eta}\right)\log\left(\frac{16\ell\Delta_0}{\mu^2\varepsilon}\right)\log\left(\frac{1+\lambda\sqrt{\sigma}D_q/\mu}{\hat{\varepsilon}}\right) + \sqrt{\kappa}\log\left(\frac{\|\beta_{-1}-\beta^*(\alpha_0)\|}{\hat{\varepsilon}}\right). \tag{21}$$

where $\hat{\varepsilon}$ satisfies condition (18).

The result of Theorem 5.3 guarantees the linear convergence of the deterministic GDGA algorithm for solving filtering-clustering problems. We now proceed to the proof of that theorem.

Proof. By the definition of $\widetilde{N}_{\mathrm{grad}}$, we get $\widetilde{N}_{\mathrm{grad}} = N_0 + \sum_{t=1}^{N} N_t$. Therefore, we conclude that

$$\begin{aligned} \widetilde{N}_{\mathrm{grad}} &\overset{\text{Lemma 5.2}}{\le} N\sqrt{\kappa}\log\left(\frac{1+\lambda\sqrt{\sigma}D_q/\mu}{\hat{\varepsilon}}\right) + \sqrt{\kappa}\log\left(\frac{\|\beta_{-1}-\beta^*(\alpha_0)\|}{\hat{\varepsilon}}\right) \\ &\overset{\text{Theorem 4.1}}{\le} \left(\frac{\sqrt{\kappa}\left(17\tau^2+14\tau+1\right)}{\tau\lambda^2\eta}\right)\log\left(\frac{16\ell\Delta_0}{\mu^2\varepsilon}\right)\log\left(\frac{1+\lambda\sqrt{\sigma}D_q/\mu}{\hat{\varepsilon}}\right) \\ &\qquad + \sqrt{\kappa}\log\left(\frac{\|\beta_{-1}-\beta^*(\alpha_0)\|}{\hat{\varepsilon}}\right), \end{aligned}$$

where $\hat{\varepsilon}$ is defined in (18). This completes the proof. □

5.3 Stochastic variance reduced GDGA algorithm

In this subsection, we concentrate on the finite-sum setting of filtering-clustering problems in which the loss function f is of the form $\frac{1}{n_{\mathsf{sam}}}\sum_{i=1}^{n_{\mathsf{sam}}} f_i$ and the component gradient oracle ∇f_i is used in each iteration of the subroutine in Algorithm 1. To ease the ensuing presentation,

Algorithm 4: Stochastic Variance Reduced GDGA algorithm

Input: (β_{-1}, α_0), learning rates $\eta > 0$.
for $t = 0, 1, \ldots$ **do**
  Find $\beta_t \in \mathbb{R}^d$ such that β_t is an $\hat{\varepsilon}$-minimizer of $f(\beta) - \lambda \alpha_t^\top D\beta$ using Katyusha algorithm [1] with $f = (\sum_{i=1}^{n_{\mathsf{sam}}} f_i)/n_{\mathsf{sam}}$.
  $\alpha_{t+1} \leftarrow \mathrm{proj}_{\mathbb{B}_q^n}(\alpha_t - \eta\lambda D\beta_t)$.
end for
Return: β_T.

we denote n_{sam} the total number of samples. To this end, we obtain an $\hat{\varepsilon}$-minimizer of $f(\beta) - \alpha_t^\top \lambda D\beta$ by applying the subroutine based on the Katyusha algorithm [1]. This procedure results in *stochastic variance reduced GDGA* algorithm where its pseudocode is summarized in Algorithm 4.

The complexity bound of the stochastic variance reduced GDGA algorithm is obtained by combining the complexity of GDGA algorithm from Theorem 4.1 and the complexity bound of the Katyusha algorithm in terms of component gradient oracles. We summarize the complexity bound of the subroutine based on the Katyusha algorithm in the following lemma.

Lemma 5.4. *Let $\hat{\varepsilon} \in (0,1)$ be given in stochastic variance reduced GDGA algorithm (Algorithm 4). Then the number of component gradient oracles to reach $\mathbb{E}\left[\|\beta_t - \beta^*(\alpha_t)\|^2\right] \le \hat{\varepsilon}^2$ is bounded by*

$$N_t \;\le\; C_{\mathsf{Kat}} \cdot \begin{cases} n_{\mathsf{sam}} + \sqrt{\kappa n_{\mathsf{sam}}} \log\left(\frac{\kappa\|\beta_{-1} - \beta^*(\alpha_0)\|}{\hat{\varepsilon}}\right) & t = 0, \\ n_{\mathsf{sam}} + \sqrt{\kappa n_{\mathsf{sam}}} \log\left(\frac{\kappa + \lambda\kappa\sqrt{\sigma} D_q/\mu}{\hat{\varepsilon}}\right) & t \ge 1. \end{cases} \tag{22}$$

where C_{Kat} is a constant defined in [1, Theorem 2.1] and independent of ℓ, μ, σ and ε.

The proof of Lemma 5.4 is provided in Section 7.3. Drawing on the result of Lemma 5.4, we are ready to present the main result on the complexity bound of the stochastic variance reduced GDGA algorithm in terms of the number of component gradient oracles.

Theorem 5.5 (Complexity of stochastic variance reduced GDGA algorithm)**.** *Let the step size $\eta > 0$ satisfy $\eta \in \left(0, \frac{\mu}{4\sigma \max\{1, \lambda^2\}}\right)$ in the stochastic variance GDGA algorithm. Then, for any $\varepsilon > 0$, the number of component gradient oracles $\widetilde{N}_{cgrad}$ for the stochastic variance reduced GDGA algorithm to find an ε-optimal solution is bounded by*

$$\begin{aligned} \widetilde{N}_{cgrad} \;\le\; & C_{\mathsf{Kat}} n_{\mathsf{sam}} \left(\left(\frac{17\tau^2 + 14\tau + 1}{\tau\lambda^2\eta} \right) \log\left(\frac{16\ell\Delta_0}{\mu^2\varepsilon} \right) + 1 \right) \\ & + C_{\mathsf{Kat}} \sqrt{\kappa n_{\mathsf{sam}}} \left(\frac{17\tau^2 + 14\tau + 1}{\tau\lambda^2\eta} \right) \log\left(\frac{16\ell\Delta_0}{\mu^2\varepsilon} \right) \log\left(\frac{\kappa + \lambda\kappa\sqrt{\sigma} D_q/\mu}{\hat{\varepsilon}} \right) \\ & + C_{\mathsf{Kat}} \sqrt{\kappa n_{\mathsf{sam}}} \log\left(\frac{\kappa \|\beta_{-1} - \beta^*(\alpha_0)\|}{\hat{\varepsilon}} \right) \end{aligned}$$

where $\hat{\varepsilon}$ satisfies condition (18) and C_{Kat} is a constant defined in [1, Theorem 2.1].

Proof. By the definition of $\widetilde{N}_{\mathrm{cgrad}}$, we get $\widetilde{N}_{\mathrm{cgrad}} = N_0 + \sum_{t=1}^{N} N_t$. Therefore, we conclude

Algorithm 5: Stochastic GDGA algorithm

Input: (β_{-1}, α_0), learning rates $\eta > 0$.
for $t = 0, 1, \ldots$ **do**
  Find $\beta_t \in \mathbb{R}^d$ such that β_t is an $\hat{\varepsilon}$-minimizer of $f(\beta) - \lambda \alpha_t^\top D\beta$ based on SGD algorithm [27] with $f(\beta) = \mathbb{E}_\xi [F(\beta, \xi)]$.
  $\alpha_{t+1} \leftarrow \mathrm{proj}_{\mathbb{B}_q^n} (\alpha_t - \eta\lambda D\beta_t)$.
end for
Return: β_T.

that

$$
\begin{aligned}
\widetilde{N}_{\mathrm{cgrad}} &\overset{\text{Lemma 5.4}}{\leq} C_{\mathsf{Kat}} n_{\mathsf{sam}}(N+1) + C_{\mathsf{Kat}} N \sqrt{\kappa n_{\mathsf{sam}}} \log\left(\frac{\kappa + \lambda\kappa\sqrt{\sigma} D_q/\mu}{\hat{\varepsilon}}\right) \\
&\quad + C_{\mathsf{Kat}} \sqrt{\kappa n_{\mathsf{sam}}} \log\left(\frac{\kappa \left\|\beta_{-1} - \beta^*(\alpha_0)\right\|}{\hat{\varepsilon}}\right) \\
&\overset{\text{Corollary 4.2}}{\leq} C_{\mathsf{Kat}} n_{\mathsf{sam}} \left(\left(\frac{17\tau^2 + 14\tau + 1}{\tau\lambda^2\eta}\right) \log\left(\frac{16\ell\Delta_0}{\mu^2\varepsilon}\right) + 1\right) \\
&\quad + C_{\mathsf{Kat}} \sqrt{\kappa n_{\mathsf{sam}}} \left(\frac{17\tau^2 + 14\tau + 1}{\tau\lambda^2\eta}\right) \log\left(\frac{16\ell\Delta_0}{\mu^2\varepsilon}\right) \log\left(\frac{\kappa + \lambda\kappa\sqrt{\sigma} D_q/\mu}{\hat{\varepsilon}}\right), \\
&\quad + C_{\mathsf{Kat}} \sqrt{\kappa n_{\mathsf{sam}}} \log\left(\frac{\kappa \left\|\beta_{-1} - \beta^*(\alpha_0)\right\|}{\hat{\varepsilon}}\right)
\end{aligned}
$$

where $\hat{\varepsilon}$ is defined in (18) and C_{Kat} is a constant defined in [1]. This completes the proof. □

The result of Theorem 5.5 guarantees the linear convergence of the stochastic variance reduced GDGA algorithm for solving filtering-clustering problems. Additionally, the complexity bound of stochastic variance GDGA algorithm outperforms that of the deterministic GDGA algorithm in terms of the number of component gradient oracles. In particular, we can also apply the deterministic GDGA algorithm in the finite-sum setting. By Theorem 5.3, $\sqrt{\kappa} n_{\mathsf{sam}}$ appears in the complexity bound based on the number of the number of component gradient oracles. In contrast, since C_{Kat} does not depend on κ, only $\sqrt{\kappa n_{\mathsf{sam}}}$ appears in the complexity bound of stochastic variance reduced GDGA algorithm (cf. Theorem 5.5). This also matches the recognized superiority of the Katyusha algorithm over the AGD algorithm [1].

5.4 Stochastic GDGA algorithm

In this subsection, we focus on the online setting in which f is in the form of $\mathbb{E}_\xi [F(\cdot, \xi)]$ and the stochastic gradient oracle $G(\cdot, \xi)$ is used in each iteration of the subroutine in Algorithm 1. To this end, we obtain an $\hat{\varepsilon}$-minimizer of $f(\beta) - \lambda\alpha_t^\top D\beta$ by applying the subroutine based on the stochastic gradient descent (SGD) algorithm [27]. The resulting algorithm based on that subroutine is called *stochastic GDGA* algorithm. The pseudocode of that algorithm is presented in Algorithm 5.

The complexity bound of the resulting stochastic GDGA algorithm is obtained by combining the complexity bound of GDGA algorithm in Theorem 4.1 and the complexity bound of the SGD algorithm in terms of stochastic gradient oracles. We summarize the complexity bound of the subroutine based on the SGD algorithm in the following lemma.

Lemma 5.6. *Let $\hat{\varepsilon} \in (0,1)$ be given in stochastic GDGA algorithm (Algorithm 5). Then the number of stochastic gradient oracles to reach $\mathbb{E}\left[\|\beta_t - \beta^*(\alpha_t)\|^2\right] \leq \hat{\varepsilon}^2$ is bounded by*

$$N_t \;\leq\; \frac{4C^2}{\mu^2\hat{\varepsilon}^2}, \quad \forall t \geq 0. \tag{23}$$

The proof of Lemma 5.6 is provided in Section 7.4. Based on the result of Lemma 5.6, we are ready to present the main result on the complexity bound of the stochastic GDGA algorithm in terms of the number of stochastic gradient oracles.

Theorem 5.7 (Complexity of Stochastic GDGA algorithm)**.** *Let the step size $\eta > 0$ satisfy $\eta \in (0, \mu/4\sigma)$ in the stochastic GDGA algorithm. Then, for any $\varepsilon > 0$, the number of stochastic gradient oracles $\widetilde{N}_{sgrad}$ for the stochastic GDGA algorithm to find an ε-optimal solution is bounded by*

$$\widetilde{N}_{sgrad} \;\leq\; \frac{4C^2}{\mu^2\hat{\varepsilon}^2}\left(\left(\frac{17\tau^2+14\tau+1}{\tau\lambda^2\eta}\right)\log\left(\frac{16\ell\Delta_0}{\mu^2\varepsilon}\right)+1\right).$$

where $\hat{\varepsilon}$ is defined in (18).

Proof. By the definition of $\widetilde{N}_{\text{sgrad}}$, we get $\widetilde{N}_{\text{sgrad}} = \sum_{t=0}^{N} N_t$. Therefore, we conclude that

$$\begin{aligned}\widetilde{N}_{\text{sgrad}} &\overset{\text{Lemma 5.6}}{\leq} \frac{4C^2(N+1)}{\mu^2\hat{\varepsilon}^2} \\ &\overset{\text{Corollary 4.2}}{\leq} \frac{4C^2}{\mu^2\hat{\varepsilon}^2}\left(\left(\frac{17\tau^2+14\tau+1}{\tau\lambda^2\eta}\right)\log\left(\frac{16\ell\Delta_0}{\mu^2\varepsilon}\right)+1\right),\end{aligned}$$

where $\hat{\varepsilon}$ is defined in (18). This completes the proof. □

The result of Theorem 5.7 guarantees the sublinear convergence of the stochastic GDGA algorithm for solving filtering-clustering problems. Furthermore, by the definition of $\hat{\varepsilon}$ in (18), we obtain that $\hat{\varepsilon} = \Omega(\sqrt{\varepsilon})$. Therefore, the complexity bound of the stochastic GDGA algorithm in terms of the number of stochastic gradient oracles is $O\left(\log(1/\varepsilon)/\varepsilon\right)$ (cf. Theorem 5.7). This complexity bound is slightly worse than the optimal complexity bound of $O\left(1/\varepsilon\right)$. It is unclear if the improvement of the complexity bound of stochastic GDGA algorithm to $O\left(1/\varepsilon\right)$ is possible by further exploring the filtering-clustering problems structure. We leave this direction for the future work.

6 Experiments

In this section, we conduct extensive simulation studies of the GDGA algorithm for ℓ_1-trend filtering problem. The simplified GDGA algorithm (Algorithm 2) with Barzilai-Borwein step size [9] are applied since the loss functions are both squared ℓ_2-norm. The baseline algorithms include standard ADMM algorithm, specialized ADMM algorithm [39, 28], and projected Newton algorithm [39].

Datasets: We consider three real images with various sizes: 128 by 128 pixels (small image), 256 by 256 pixels (medium image), and 512 by 512 pixels (large image)[1].

Experimental settings: We present comparative experiments between the simplified GDGA, ADMM, specialized ADMM, and projected Newton algorithms as the order k varies

[1]These images can be found at: http://sipi.usc.edu/database/database.php?volume=misc

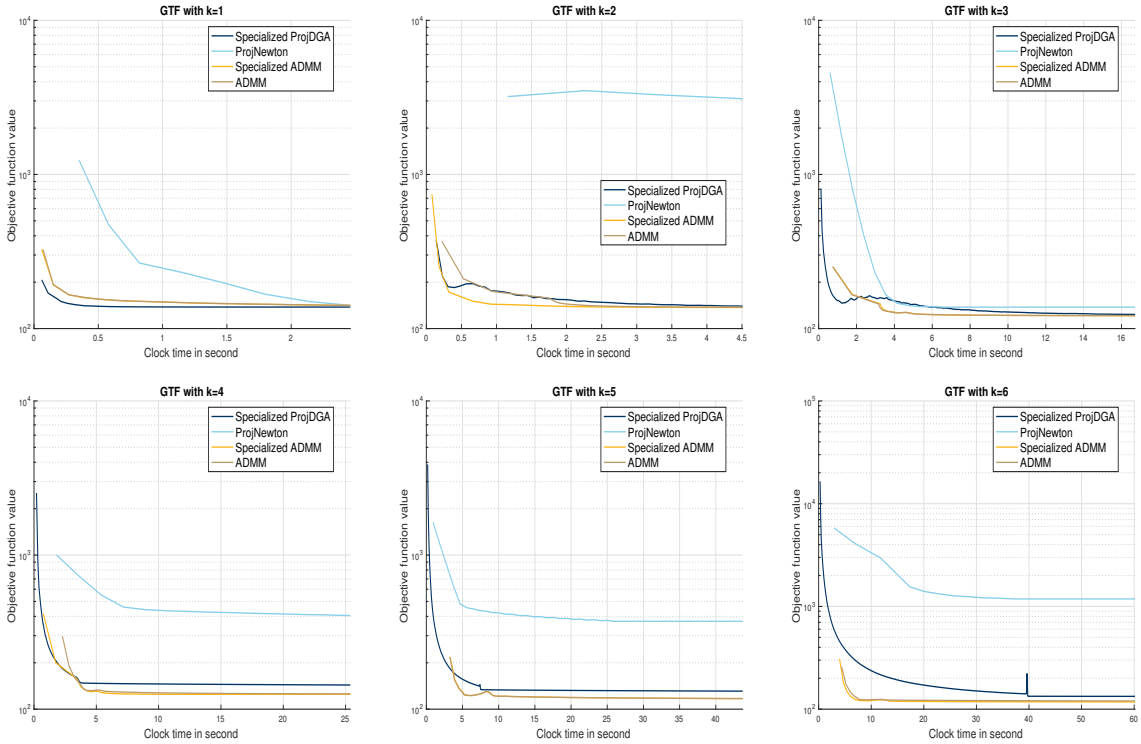

Figure 1: Performance of GDGA, ADMM, specialized ADMM, and projected Newton algorithms for small image.

in the discrete difference operator $D^{(k+1)}$ (see Section 2.2.1 for the details). The evaluation metric for our comparison is the objective function value of the ℓ_1-trend-filtering problem.

Experimental results: Figure 1-3 present the experimental results for different settings of images. The simplified GDGA algorithm and the two ADMM algorithms are consistently comparable, all of which significantly outperform the projected Newton algorithm. In particular, the performance of simplified GDGA algorithm is the best consistently among all the algorithms when $k = 1$. As k increases, the performance of the projected Newton algorithm deteriorates quickly. This makes sense since the subroutine based on the conjugate gradient algorithm is known to be inefficient if the conditioning of the Hessian is bad. In contrast, the simplified GDGA algorithm and the two ADMM algorithms remain effective while the ADMM algorithms are slightly better in general. The good performance of these two ADMM algorithms arises from the use of a Cholesky factorization, which alleviates the ill-conditioning of the matrix $D^{(k+1)}$. On massive-scale problems, however, this advantage becomes a liability since the computational and memory requirements for ADMM become severe. By way of contrast, our GDGA algorithm with Barzilai-Borwein step size is purely matrix-free, in which no matrix factorization is required. Moreover, the usage of Barzilai-Borwein step size accelerates the algorithm by exploring the curvature information and alleviates the ill-conditioning.

7 Proofs

In this section, we first present the proof for the complexity bounds of the GDGA algorithmic framework. Then we combine it with the complexity bounds for the best-known algorithms [20,

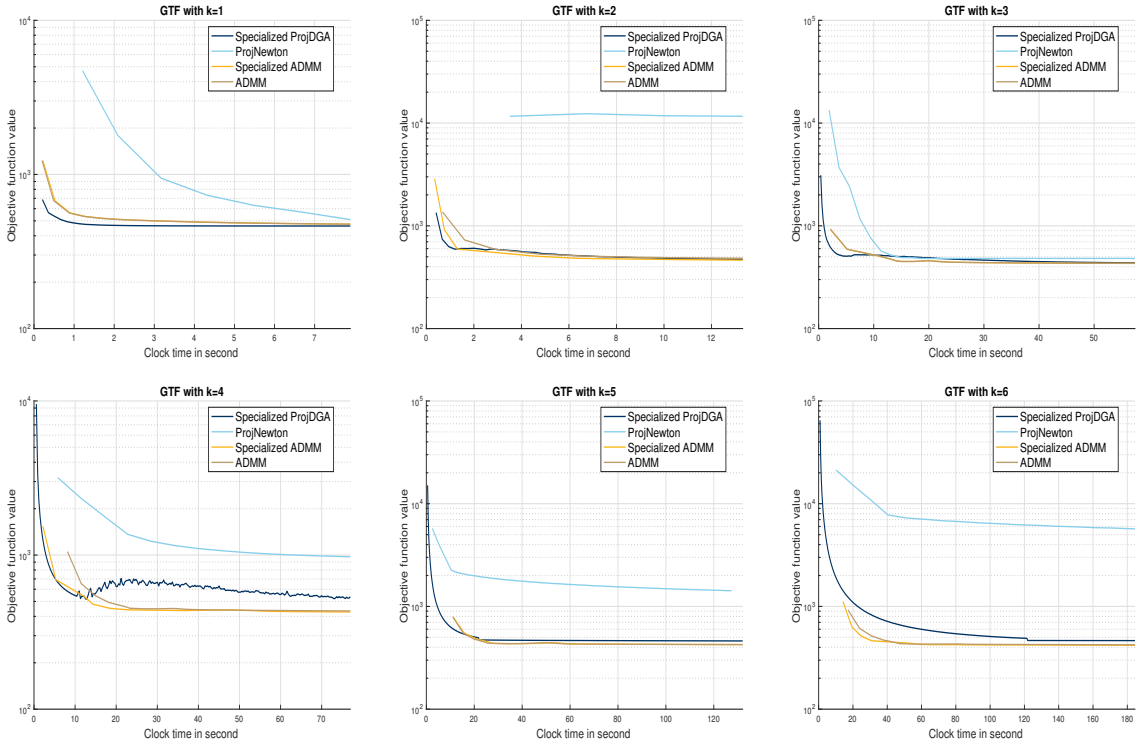

Figure 2: Performance of GDGA, ADMM, specialized ADMM, and projected Newton algorithms for medium image.

1, 27] and obtain the complexity bounds of the specifications of the GDGA algorithmic framework to different settings of filtering-clustering problems.

7.1 Technical lemmas for complexity bounds

In this subsection, we prove several technical lemmas used for establishing complexity bounds in the paper. To ease the ensuing proof argument, we trace the distance between β_t and $\beta^*(\alpha_t)$ by defining

$$\delta_t \;:=\; \|\beta_t - \beta^*(\alpha_t)\|\,.$$

Additionally, we denote $\bar{\alpha}_t$ as the projection of α_t onto the optimal set of problem (6), i.e., Ω^*, and trace the objective gap between $\bar{f}(\alpha_t)$ and $\bar{f}(\bar{\alpha}_t)$ by defining

$$\Delta_t := \bar{f}(\alpha_t) - \bar{f}(\bar{\alpha}_t).$$

The first lemma provides a key lower bound for the iterative objective gap, i.e., $\Delta_t - \Delta_{t+1}$.

Lemma 7.1. *Let $(\alpha_t, \beta_t)_{t\geq 0}$ be the iterates generated by Algorithm 1 with a stepsize $\eta \in \left(0,\ \frac{\mu}{4\sigma\max\{1,\lambda^2\}}\right)$. Then, for any $t \geq 0$, the following holds*

$$\Delta_t - \Delta_{t+1} \;\geq\; \frac{\lambda^2\,\|\alpha_t - \alpha_{t+1}\|^2}{4\eta} - \frac{\eta\sigma\hat{\varepsilon}^2}{2}. \tag{24}$$

Proof. Since $\alpha_{t+1} = \mathrm{proj}_{\mathbb{B}_q^n}(\alpha_t - \eta\lambda D\beta_t)$, we obtain from the definition of the projection operator that

$$(\alpha - \alpha_{t+1})^\top(\alpha_{t+1} - \alpha_t + \eta\lambda D\beta_t) \;\geq\; 0, \quad \forall \alpha \in \mathbb{B}_q^n.$$

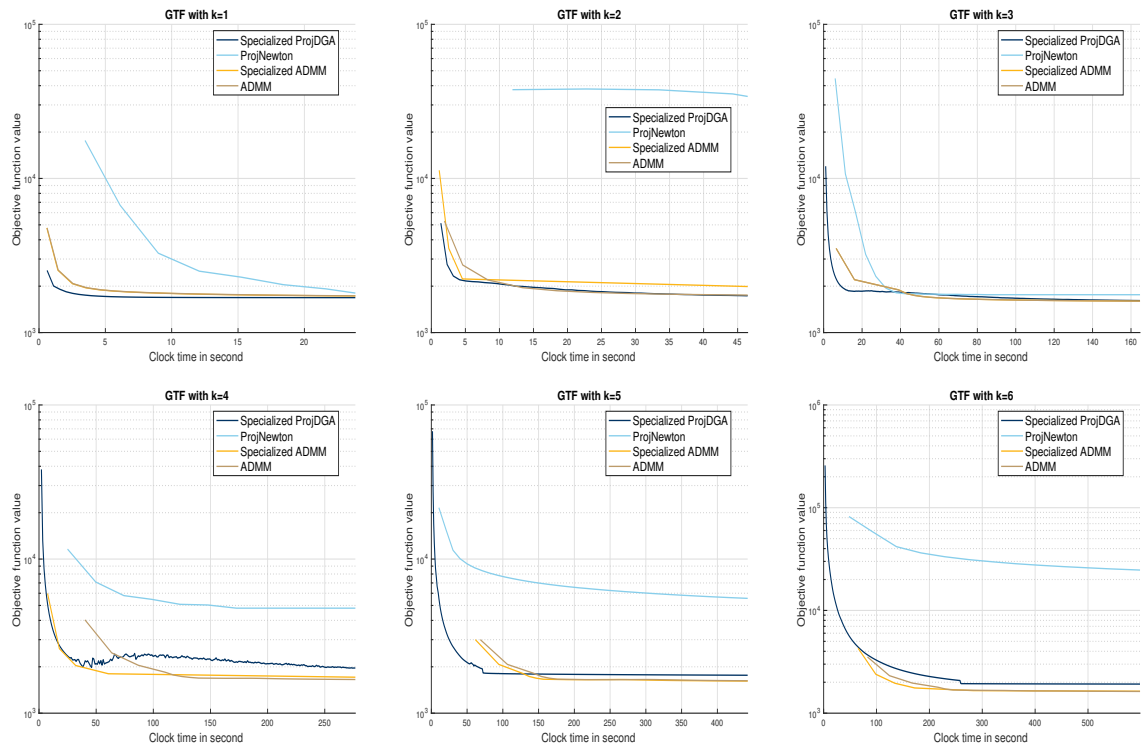

Figure 3: Performance of GDGA, ADMM, specialized ADMM, and projected Newton algorithms for large image.

Let $\alpha = \alpha_t$, then

$$(\alpha_t - \alpha_{t+1})^\top \lambda D\beta_t \;\geq\; \frac{\|\alpha_t - \alpha_{t+1}\|^2}{\eta}. \tag{25}$$

Furthermore, we have the following inequalities

$$\begin{aligned}(\alpha_t - \alpha_{t+1})^\top \lambda D\beta_t &= (\alpha_t - \alpha_{t+1})^\top \nabla \bar{f}(\alpha_t) + (\alpha_t - \alpha_{t+1})^\top \left(\lambda D\beta_t - \nabla \bar{f}(\alpha_t)\right) \\ &\leq \bar{f}(\alpha_t) - \bar{f}(\alpha_{t+1}) + \frac{\sigma\lambda^2 \|\alpha_t - \alpha_{t+1}\|^2}{\mu} + (\alpha_t - \alpha_{t+1})^\top \left(\lambda D\beta_t - \nabla \bar{f}(\alpha_t)\right)\end{aligned} \tag{26}$$

where the last inequality holds since $\bar{f}$ is $\frac{\sigma\lambda^2}{\mu}$-gradient Lipschiz (cf. Lemma 3.3). Furthermore, since $\nabla \bar{f}(\alpha_t) = \lambda D\beta^*(\alpha_t)$ (cf. Lemma 3.2), then an application of Cauchy-Schwarz's inequality yields that

$$\begin{aligned}(\alpha_t - \alpha_{t+1})^\top \left(\lambda D\beta_t - \nabla \bar{f}(\alpha_t)\right) &= (\alpha_t - \alpha_{t+1})^\top \left(\lambda D\beta_t - \lambda D\beta^*(\alpha_t)\right) \\ &\leq \frac{\lambda^2}{2\eta} \|\alpha_t - \alpha_{t+1}\|^2 + \frac{\eta}{2} \|D\beta_t - D\beta^*(\alpha_t)\|^2 \\ &\leq \frac{\lambda^2}{2\eta} \|\alpha_t - \alpha_{t+1}\|^2 + \frac{\eta\sigma\hat{\varepsilon}^2}{2}.\end{aligned} \tag{27}$$

Plugging (26) and (27) into (25) yields that

$$\begin{aligned}\bar{f}(\alpha_t) - \bar{f}(\alpha_{t+1}) &\geq \left(\frac{\lambda^2}{2\eta} - \frac{\sigma\lambda^2}{\mu}\right)\|\alpha_t - \alpha_{t+1}\|^2 - \frac{\eta\sigma\hat{\varepsilon}^2}{2} \\ &\geq \frac{\lambda^2\|\alpha_t - \alpha_{t+1}\|^2}{4\eta} - \frac{\eta\sigma\hat{\varepsilon}^2}{2}.\end{aligned}$$

To this end, we complete the proof by concluding (24). □

The second lemma presents an upper bound for Δ_{t+1} based on $\|\alpha_t - \alpha_{t+1}\|^2$ using the global error bound (cf. Theorem 3.8 and Theorem 3.11).

Lemma 7.2. *Let $(\alpha_t, \beta_t)_{t\geq 0}$ be the iterates generated by Algorithm 1 with a stepsize $\eta \in \left(0, \frac{\mu}{4\sigma\max\{1,\lambda^2\}}\right)$. Then, for any $t \geq 0$, the following holds*

$$\|\alpha_{t+1} - \alpha_t\|^2 \geq \frac{4\tau\eta^2\Delta_{t+1}}{17\tau^2 + 14\tau + 1} - \frac{8\tau^2\lambda^2\sigma\eta^2\hat{\varepsilon}^2}{17\tau^2 + 14\tau + 1}. \tag{28}$$

Proof. We observe that

$$\begin{aligned}\left\|\alpha_t - \mathrm{proj}_{\mathbb{B}_q^n}\left(\alpha_t - \eta\nabla\bar{f}(\alpha_t)\right)\right\| &= \left\|\alpha_t - \mathrm{proj}_{\mathbb{B}_q^n}\left(\alpha_t - \eta\lambda D\beta^*(\alpha_t)\right)\right\| \\ &\leq \|\alpha_t - \alpha_{t+1}\| + \left\|\alpha_{t+1} - \mathrm{proj}_{\mathbb{B}_q^n}\left(\alpha_t - \eta\lambda D\beta^*(\alpha_t)\right)\right\|\end{aligned}$$

where the second inequality is due to triangle inequality. Since $\alpha_{t+1} = \mathrm{proj}_{\mathbb{B}_q^n}(\alpha_t - \eta\lambda D\beta_t)$ and the projection operator is nonexpansive, we achieve that

$$\left\|\alpha_{t+1} - \mathrm{proj}_{\mathbb{B}_q^n}\left(\alpha_t - \eta\lambda D\beta^*(\alpha_t)\right)\right\| \leq \eta\lambda\sqrt{\sigma}\delta_t. \tag{29}$$

From [10, Lemma 1], the term $\left\|\alpha_t - \mathrm{proj}_{\mathbb{B}_q^n}\left(\alpha_t - \eta\nabla\bar{f}(\alpha_t)\right)\right\|/\eta$ is monotonically decreasing for $\forall\eta > 0$. Since $\eta \in (0,1)$, we find that

$$\eta\left\|\alpha_t - \mathrm{proj}_{\mathbb{B}_q^n}\left(\alpha_t - \nabla\bar{f}(\alpha_t)\right)\right\| \leq \left\|\alpha_t - \mathrm{proj}_{\mathbb{B}_q^n}\left(\alpha_t - \eta\nabla\bar{f}(\alpha_t)\right)\right\|. \tag{30}$$

Recall that, $\bar{\alpha}_t$ is the projection of α_t onto Ω^*, the following inequalities hold

$$\begin{aligned}\|\alpha_t - \bar{\alpha}_t\| &\overset{\text{Theorems 3.8, 3.11}}{\leq} \tau\left\|\alpha - \mathrm{proj}_{\mathbb{B}_q^n}\left(\alpha_t - \nabla\bar{f}(\alpha_t)\right)\right\| \\ &\overset{(30)}{\leq} \frac{\tau\left\|\alpha_t - \mathrm{proj}_{\mathbb{B}_q^n}\left(\alpha_t - \eta\nabla\bar{f}(\alpha_t)\right)\right\|}{\eta} \\ &\overset{(29)}{\leq} \frac{\tau\|\alpha_t - \alpha_{t+1}\|}{\eta} + \tau\lambda\sqrt{\sigma}\hat{\varepsilon}.\end{aligned} \tag{31}$$

Finally, we bound the term Δ_{t+1}. More specifically, since $\alpha_{t+1} = \mathrm{proj}_{\mathbb{B}_q^n}(\alpha_t - \eta\lambda D\beta_t)$, we obtain from the definition of the projection operator that

$$(\alpha - \alpha_{t+1})^\top(\alpha_{t+1} - \alpha_t + \eta\lambda D\beta_t) \geq 0, \quad \forall\alpha \in \mathbb{B}_q^n.$$

Let $\alpha = \bar{\alpha}_t$, then

$$(\bar{\alpha}_t - \alpha_{t+1})^\top\lambda D\beta_t \geq \frac{(\bar{\alpha}_t - \alpha_{t+1})^\top(\alpha_t - \alpha_{t+1})}{\eta}. \tag{32}$$

Furthermore, we obtain from the convexity of $\bar{f}$ that

$$\Delta_{t+1} \;=\; \bar{f}(\alpha_{t+1}) - \bar{f}(\bar{\alpha}_{t+1}) \;=\; \bar{f}(\alpha_{t+1}) - \bar{f}(\bar{\alpha}_t) \;\leq\; (\alpha_{t+1} - \bar{\alpha}_t)^\top \nabla \bar{f}(\alpha_{t+1}). \tag{33}$$

Combining (31), (32), (33) and $\nabla \bar{f}(\alpha_t) = \lambda D \beta^*(\alpha_t)$ yields that

$$\begin{aligned}
& (\alpha_{t+1} - \bar{\alpha}_t)^\top \nabla \bar{f}(\alpha_{t+1}) \\
= \;& (\alpha_{t+1} - \bar{\alpha}_t)^\top \left(\nabla \bar{f}(\alpha_{t+1}) - \nabla \bar{f}(\alpha_t)\right) + (\alpha_{t+1} - \bar{\alpha}_t)^\top (\lambda D \beta^*(\alpha_t) - \lambda D \beta_t) + (\alpha_{t+1} - \bar{\alpha}_t)^\top \lambda D \beta_t \\
\leq \;& \|\alpha_{t+1} - \bar{\alpha}_t\| \left(\frac{\sigma \lambda^2 \|\alpha_{t+1} - \alpha_t\|}{\mu} + \lambda \sqrt{\sigma} \delta_t + \frac{\|\alpha_{t+1} - \alpha_t\|}{\eta} \right) \\
\leq \;& \left(\frac{(\tau + \eta) \|\alpha_{t+1} - \alpha_t\|}{\eta} + \tau \lambda \sqrt{\sigma} \hat{\varepsilon} \right) \left(\lambda \sqrt{\sigma} \hat{\varepsilon} + \left(\frac{\sigma \lambda^2}{\mu} + \frac{1}{\eta} \right) \|\alpha_{t+1} - \alpha_t\| \right).
\end{aligned}$$

Since $\eta \leq 1$ and $\sigma \lambda^2 / \mu < 1/4\eta$, we have

$$(\alpha_{t+1} - \bar{\alpha}_t)^\top \nabla \bar{f}(\alpha_{t+1}) \;\leq\; \frac{2(\tau + 1) \|\alpha_{t+1} - \alpha_t\|^2}{\eta^2} + \frac{(3\tau + 1) \lambda \sqrt{\sigma} \hat{\varepsilon} \|\alpha_{t+1} - \alpha_t\|}{\eta} + \tau \lambda^2 \sigma \hat{\varepsilon}^2. \tag{34}$$

Applying the Young's inequality to the term $\hat{\varepsilon} \|\alpha_{t+1} - \alpha_t\|$ yields that

$$\hat{\varepsilon} \|\alpha_{t+1} - \alpha_t\| \;\leq\; \frac{\tau \eta \lambda \sqrt{\sigma} \hat{\varepsilon}^2}{(3\tau + 1)} + \frac{(3\tau + 1) \|\alpha_{t+1} - \alpha_t\|^2}{4 \tau \eta \lambda \sqrt{\sigma}}. \tag{35}$$

Plugging (35) into (34) yields that

$$(\alpha_{t+1} - \bar{\alpha}_t)^\top \nabla \bar{f}(\alpha_{t+1}) \;\leq\; \frac{(17\tau^2 + 14\tau + 1) \|\alpha_{t+1} - \alpha_t\|^2}{4 \tau \eta^2} + 2 \tau \lambda^2 \sigma \hat{\varepsilon}^2.$$

Combining the above bound with (33) yields that

$$\|\alpha_{t+1} - \alpha_t\|^2 \;\geq\; \frac{4 \tau \eta^2 \Delta_{t+1}}{17\tau^2 + 14\tau + 1} - \frac{8 \tau^2 \lambda^2 \sigma \eta^2 \hat{\varepsilon}^2}{17\tau^2 + 14\tau + 1}.$$

As a consequence, we reach the conclusion of the lemma. □

Equipped with the bounds of iterative objective gap $\Delta_t - \Delta_{t+1}$ and objective gap Δ_{t+1} in Lemma 7.1 and 7.2, we are ready to prove the main lemma for the number of iterations of GDGA algorithm to reach a certain threshold with objective gap Δ_t. Before stating that result, we assume the following key technical assumption with approximation error $\hat{\varepsilon}$:

$$\hat{\varepsilon} \;\leq\; \min \left\{ \frac{\sqrt{\varepsilon}}{2}, \; \frac{\sqrt{\varepsilon} \lambda \mu}{4 \sqrt{\ell}} \sqrt{\frac{\tau}{C(17\tau^2 + (14 + \eta \lambda^2)\tau + 1)}} \right\}, \tag{36}$$

where $C > 0$ is defined as

$$C \;:=\; \frac{2 \tau^2 \lambda^4 \sigma}{17\tau^2 + 14\tau + 1} + \frac{\sigma}{2}.$$

Lemma 7.3. *Let $(\alpha_t, \beta_t)_{t \geq 0}$ be the iterates generated by Algorithm 1 with $\eta \in \left(0, \; \frac{\mu}{4\sigma \max\{1, \lambda^2\}}\right)$. Given the bound (36) with $\hat{\varepsilon}$, the number of iterations to reach $\Delta_t \leq \frac{\mu^2 \varepsilon}{8\ell}$ is*

$$N \;\leq\; \left(\frac{17\tau^2 + 14\tau + 1}{\tau \lambda^2 \eta} \right) \log \left(\frac{16 \ell \Delta_0}{\mu^2 \varepsilon} \right).$$

where $\Delta_0 \geq 0$ is the distance between α_0 and the optimal solution set of problem (6).

Proof. Invoking the results from (24) (cf. Lemma 7.1) and (28) (cf. Lemma 7.2) yields that

$$\begin{aligned}\Delta_t - \Delta_{t+1} &\geq \frac{\lambda^2}{4\eta}\left[\frac{4\tau\eta^2\Delta_{t+1}}{17\tau^2+14\tau+1} - \frac{8\tau^2\lambda^2\sigma\eta^2\hat{\varepsilon}^2}{17\tau^2+14\tau+1}\right] - \frac{\eta\sigma\hat{\varepsilon}^2}{2} \\ &= \frac{\tau\lambda^2\eta\Delta_{t+1}}{17\tau^2+14\tau+1} - \left(\frac{2\tau^2\lambda^4\sigma}{17\tau^2+14\tau+1} + \frac{\sigma}{2}\right)\eta\hat{\varepsilon}^2.\end{aligned}$$

Let $\rho > 0$ be defined as

$$\rho = \left(1 + \frac{\tau\lambda^2\eta}{17\tau^2+14\tau+1}\right)^{-1}.$$

Then, for any $t \geq 0$, we find that

$$\Delta_{t+1} \leq \rho\left(\Delta_t + C\eta\hat{\varepsilon}^2\right).$$

Recursively performing the above inequality yields that

$$\Delta_t \leq \rho^t\Delta_0 + \left(\sum_{j=0}^{t-1}\rho^{t-1-j}\right)\cdot C\eta\hat{\varepsilon}^2 \leq \rho^t\Delta_0 + \frac{C\eta\hat{\varepsilon}^2}{1-\rho}.$$

By using the definition of ρ, we have

$$\frac{C\eta}{1-\rho} = C\eta + \frac{C\left(17\tau^2+14\tau+1\right)}{\tau\lambda^2}.$$

By the definition of $\hat{\varepsilon}$ in (36), we have $\Delta_t \leq \rho^t\Delta_0 + \frac{\mu^2\varepsilon}{16\ell}$. Therefore, the number of iterations to reach $\Delta_t \leq \mu^2\varepsilon/8\ell$ is

$$N \leq \left(\frac{17\tau^2+14\tau+1}{\tau\lambda^2\eta}\right)\log\left(\frac{16\ell\Delta_0}{\mu^2\varepsilon}\right).$$

As a consequence, we achieve the conclusion of the lemma. □

Finally, we consider the lemmas in the stochastic setting. More specifically, the GDGA algorithm is intrinsically stochastic if the subroutines are based on the stochastic gradient-type algorithms, e.g., Katyusha and SGD algorithms. Since the proofs of lemmas with the stochastic setting are nearly the same as those from deterministic setting, we present these lemmas but omit their proofs.

Lemma 7.4. *Let $(\alpha_t, \beta_t)_{t\geq 0}$ be the iterates generated by Algorithm 1 with stochastic subroutine, then*

$$\mathbb{E}\left[\Delta_t\right] - \mathbb{E}\left[\Delta_{t+1}\right] \geq \frac{\lambda^2}{4\eta}\mathbb{E}\left[\|\alpha_t - \alpha_{t+1}\|^2\right] - \frac{\eta\sigma\hat{\varepsilon}^2}{2}. \tag{37}$$

Lemma 7.5. *Let $(\alpha_t, \beta_t)_{t\geq 0}$ be the iterates generated by Algorithm 1 with stochastic subroutine and $\eta \in (0, \mu/4\sigma)$, then*

$$\mathbb{E}\left[\|\alpha_{t+1} - \alpha_t\|^2\right] \geq \frac{4\tau\eta^2\mathbb{E}\left[\Delta_{t+1}\right]}{17\tau^2+14\tau+1} - \frac{8\tau^2\lambda^2\sigma\eta^2\hat{\varepsilon}^2}{17\tau^2+14\tau+1}. \tag{38}$$

Lemma 7.6. *Let $(\alpha_t, \beta_t)_{t\geq 0}$ be the iterates generated by Algorithm 1 with $\eta \in \left(0, \frac{\mu}{4\sigma\max\{1,\lambda^2\}}\right)$. Given the bound (36) with $\hat{\varepsilon}$, the number of iterations to reach $\mathbb{E}[\Delta_t] \leq \frac{\mu^2\varepsilon}{8\ell}$ is*

$$N \leq \left(\frac{17\tau^2+14\tau+1}{\tau\lambda^2\eta}\right)\log\left(\frac{16\ell\Delta_0}{\mu^2\varepsilon}\right).$$

where $\Delta_0 \geq 0$ is the distance between α_0 and the optimal solution set of problem (6).

7.2 Proof of Lemma 5.2

We establish our result by using the existing complexity bound of the AGD algorithm with the step size $1/\ell$ [20, Theorem 2.2.2]. Since f is μ-strongly convex and ℓ-gradient Lipschitz, it holds true that $f(\beta) - \alpha_t^\top \lambda D\beta$ is μ-strongly convex and ℓ-gradient Lipschiz with the condition number $\frac{\ell}{\mu}$. For $t = 0$, the initial distance is $\|\beta_{-1} - \beta^*(\alpha_0)\|$ so N_0 is bounded by

$$N_0 \;\leq\; \sqrt{\kappa} \log\left(\frac{\|\beta_{-1} - \beta^*(\alpha_0)\|}{\hat{\varepsilon}}\right).$$

For $t \geq 1$, the initial distance is $\|\beta_{t-1} - \beta^*(\alpha_t)\|$ so N_t is bounded by

$$N_t \;\leq\; \sqrt{\kappa} \log\left(\frac{\|\beta_{t-1} - \beta^*(\alpha_t)\|}{\hat{\varepsilon}}\right).$$

Furthermore, by using the triangle inequality, we have

$$\begin{aligned} \|\beta_{t-1} - \beta^*(\alpha_t)\| &\leq \|\beta_{t-1} - \beta^*(\alpha_{t-1})\| + \|\beta^*(\alpha_{t-1}) - \beta^*(\alpha_t)\| \\ &\leq 1 + \|\beta^*(\alpha_{t-1}) - \beta^*(\alpha_t)\|. \end{aligned}$$

Since $\beta^*(\alpha)$ is $\frac{\lambda\sqrt{\sigma}}{\mu}$-Lipschitz over $\mathbb{B}_q^n$ (cf. Lemma 3.2), we have

$$\|\beta^*(\alpha_{t-1}) - \beta^*(\alpha_t)\| \;\leq\; \frac{\lambda\sqrt{\sigma}\,\|\alpha_t - \alpha_{t-1}\|}{\mu} \;\leq\; \frac{\lambda\sqrt{\sigma} D_q}{\mu}$$

Therefore, we have

$$N_t \;\leq\; \sqrt{\kappa} \log\left(\frac{1 + \lambda\sqrt{\sigma} D_q/\mu}{\hat{\varepsilon}}\right).$$

This completes the proof of the lemma.

7.3 Proofs of Lemma 5.4

We establish our result by using the existing complexity bound of the Katyusha algorithm with the step size $\max\left\{2/3\ell, 1/\sqrt{3n\mu\ell}\right\}$ [1, Theorem 2.1]. Since f is μ-strongly convex and ℓ-gradient Lipschitz, it holds true that $f(\beta) - \alpha_t^\top \lambda D\beta$ is μ-strongly convex and ℓ-gradient Lipschiz with the condition number $\kappa := \ell/\mu$. For $t = 0$, the initial distance is $\|\beta_{-1} - \beta^*(\alpha_0)\|$ so N_0 is bounded by

$$N_0 \;\leq\; C_{\mathsf{Kat}}\left(n_{\mathsf{sam}} + \sqrt{\kappa n_{\mathsf{sam}}} \log\left(\frac{\kappa\,\|\beta_{-1} - \beta^*(\alpha_0)\|}{\hat{\varepsilon}}\right)\right).$$

For $t \geq 1$, the initial distance is $\|\beta_{t-1} - \beta^*(\alpha_t)\|$ so N_t is bounded by

$$N_t \;\leq\; C_{\mathsf{Kat}}\left(n_{\mathsf{sam}} + \sqrt{\kappa n_{\mathsf{sam}}} \log\left(\frac{\kappa\,\|\beta_{t-1} - \beta^*(\alpha_t)\|}{\hat{\varepsilon}}\right)\right).$$

By applying the similar argument as that in the proof of Lemma 5.2 in Section 7.2, we find that

$$N_t \;\leq\; C_{\mathsf{Kat}}\left(n_{\mathsf{sam}} + \sqrt{\kappa n_{\mathsf{sam}}} \log\left(\frac{\kappa + \lambda\kappa\sqrt{\sigma} D_q/\mu}{\hat{\varepsilon}}\right)\right).$$

As a consequence, we achieve the conclusion of the lemma.

7.4 Proofs of Lemma 5.6

We establish our result by using the existing complexity bound of the SGD algorithm with the diminishing step size $1/\mu k$ [27, Lemma 1]. Since f is μ-strongly convex and ℓ-gradient Lipschitz, it holds true that $f(\beta) - \alpha_t^\top \lambda D\beta$ is μ-strongly convex and ℓ-gradient Lipschiz. Also, the stochastic gradient oracle is unbiased and bounded by a constant $C > 0$ (cf. Assumption 2.1). Therefore, we conclude that

$$N_t \;\leq\; \frac{4C^2}{\mu^2 \hat{\varepsilon}^2}$$

for all $t \geq 0$. This completes the proof of the lemma.

8 Discussion

In the paper, we have proposed and analyzed a class of first-order gradient-type optimization algorithms to solve the filtering-clustering problems (1). In particular, deterministic generalized dual gradient ascent (GDGA) algorithms are shown to have optimal linear convergence rates for finding a global optimal solution of the filtering-clustering problems. The favorable convergence of GDGA is based on a crucial global error bound of the dual form of these problems. Furthermore, stochastic versions of GDGA algorithm, including stochastic GDGA algorithm and accelerated stochastic variance reduced GDGA algorithm, have been proposed to deal with the finite sum setting or online setting of filtering clustering problems. These algorithms are demonstrated to have the optimal convergence rates in their respective settings. Finally, careful experiments with ℓ_1-trend filtering show that our GDGA algorithms have competitive performance with several state-of-the-art algorithms for these problems.

Acknowledgements

This work was supported in part by the Mathematical Data Science program of the Office of Naval Research under grant number N00014-18-1-2764.

References

[1] Z. Allen-Zhu. Katyusha: The first direct acceleration of stochastic gradient methods. In *STOC*, pages 1200–1205. ACM, 2017. (Cited on pages 3, 16, 19, 20, 23, and 28.)

[2] H. H. Bauschke and J. M. Borwein. On projection algorithms for solving convex feasibility problems. *SIAM Review*, 38(3):367–426, 1996. (Cited on page 12.)

[3] A. Beck and S. Shtern. Linearly convergent away-step conditional gradient for non-strongly convex functions. *Mathematical Programming*, 164(1-2):1–27, 2017. (Cited on page 2.)

[4] G. H. G. Chen and R. T. Rockafellar. Convergence rates in forward–backward splitting. *SIAM Journal on Optimization*, 7(2):421–444, 1997. (Cited on page 2.)

[5] A. Cherukuri, B. Gharesifard, and J. Cortes. Saddle-point dynamics: conditions for asymptotic stability of saddle points. *SIAM Journal on Control and Optimization*, 55(1):486–511, 2017. (Cited on page 2.)

[6] E. C. Chi and K. Lange. Splitting methods for convex clustering. *Journal of Computational and Graphical Statistics*, 24(4):994–1013, 2015. (Cited on pages 1, 2, 6, and 15.)

[7] D. Drusvyatskiy and A. S. Lewis. Error bounds, quadratic growth, and linear convergence of proximal methods. *Mathematics of Operations Research*, 2018. (Cited on page 2.)

[8] S. S. Du and W. Hu. Linear convergence of the primal-dual gradient method for convex-concave saddle point problems without strong convexity. *ArXiv Preprint: 1802.01504*, 2018. (Cited on pages 2 and 15.)

[9] R. Fletcher. On the Barzilai-Borwein method. In *Optimization and control with applications*, pages 235–256. Springer, 2005. (Cited on page 21.)

[10] E. M. Gafni and D. P. Bertsekas. Two-metric projection methods for constrained optimization. *SIAM Journal on Control and Optimization*, 22(6):936–964, 1984. (Cited on page 25.)

[11] A. Guntuboyina, D. Lieu, S. Chatterjee, and B. Sen. Adaptive risk bounds in univariate total variation denoising and trend filtering. *Annals of Statistics*, To appear. (Cited on page 1.)

[12] T. Hocking, J.-P. Vert, F. Bach, and A. Joulin. Clusterpath: an algorithm for clustering using convex fusion penalties. In *ICML*, 2011. (Cited on pages 1 and 5.)

[13] A. Jourani. Hoffman's error bound, local controllability, and sensitivity analysis. *SIAM Journal on Control and Optimization*, 38(3):947–970, 2000. (Cited on page 13.)

[14] S.-J. Kim, K. Koh, S. Boyd, and D. Gorinevsky. ℓ_1 trend filtering. *SIAM Review*, 51(2):339–360, 2009. (Cited on pages 1, 4, 5, and 15.)

[15] G. Lan and Y. Zhou. An optimal randomized incremental gradient method. *Mathematical Programming*, pages 1–49, 2017. (Cited on page 6.)

[16] C. Leser. A simple method of trend construction. *Journal of the Royal Statistical Society: Series B*, 23:91–107, 1961. (Cited on page 1.)

[17] K. Lin, J. L. Sharpnack, A. Rinaldo, and R. J. Tibshirani. A sharp error analysis for the fused lasso, with application to approximate changepoint screening. In *NIPS*, 2017. (Cited on page 1.)

[18] Z.-Q. Luo and P. Tseng. On the linear convergence of descent methods for convex essentially smooth minimization. *SIAM Journal on Control and Optimization*, 30(2):408–425, 1992. (Cited on page 2.)

[19] Z.-Q. Luo and P. Tseng. Error bounds and convergence analysis of feasible descent methods: a general approach. *Annals of Operations Research*, 46(1):157–178, 1993. (Cited on page 2.)

[20] Y. Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2013. (Cited on pages 2, 3, 16, 17, 23, and 28.)

[21] O. H. M. Padilla, J. Sharpnack, Y. Chen, and D. M. Witten. Adaptive non-parametric regression with the K-NN fused lasso. *Arxiv Preprint: 1807.11641*, 2018. (Cited on page 1.)

[22] O. H. M. Padilla, J. Sharpnack, J. G. Scott, and R. J. Tibshirani. The DFS fused lasso: Linear-time denoising over general graphs. *Journal of Machine Learning Research*, 18(1):136, 2018. (Cited on page 1.)

[23] B. Palaniappan and F. Bach. Stochastic variance reduction methods for saddle-point problems. In *NeurIPS*, pages 1416–1424, 2016. (Cited on pages 2 and 15.)

[24] J.-S. Pang. A posteriori error bounds for the linearly-constrained variational inequality problem. *Mathematics of Operations Research*, 12(3):474–484, 1987. (Cited on pages 2 and 9.)

[25] N. Parikh, S. Boyd, et al. Proximal algorithms. *Foundations and Trends® in Optimization*, 1(3):127–239, 2014. (Cited on pages 6 and 15.)

[26] P. Radchenko and G. Mukherjee. Convex clustering via ℓ_1 fusion penalization. *Journal of the Royal Statistical Society: Series B*, 79(5):15271546, 2017. (Cited on page 1.)

[27] A. Rakhlin, O. Shamir, and K. Sridharan. Making gradient descent optimal for strongly convex stochastic optimization. In *ICML*, pages 1571–1578. Omnipress, 2012. (Cited on pages 3, 16, 20, 23, and 29.)

[28] A. Ramdas and R. J. Tibshirani. Fast and flexible admm algorithms for trend filtering. *Journal of Computational and Graphical Statistics*, 25(3):839–858, 2016. (Cited on pages 1, 2, 5, 6, 15, and 21.)

[29] R. T. Rockafellar. *Convex analysis*. Princeton University Press, 2015. (Cited on pages 6, 7, and 10.)

[30] L. I. Rudin, S. Osher, and E. Faterni. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60:259–268, 1992. (Cited on page 5.)

[31] K. M. Tan and D. Witten. Statistical properties of convex clustering. *Electronic Journal of Statistics*, 9(2):2324–2347, 2015. (Cited on page 1.)

[32] R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B*, 67(1):91–108, 2005. (Cited on page 5.)

[33] R. J. Tibshirani. Adaptive piecewise polynomial estimation via trend filtering. *Annals of Statistics*, 42(1):285–323, 2014. (Cited on pages 1 and 4.)

[34] R. J. Tibshirani and J. Taylor. The solution path of the generalized lasso. *Annals of Statistics*, 39(3):13351371, 2011. (Cited on page 5.)

[35] P. Tseng. Approximation accuracy, gradient methods, and error bound for structured convex optimization. *Mathematical Programming*, 125(2):263–295, 2010. (Cited on page 2.)

[36] J. Wang and L. Xiao. Exploiting strong convexity from data with primal-dual first-order algorithms. In *ICML*, pages 3694–3702, 2017. (Cited on pages 2 and 15.)

[37] P.-W. Wang and C.-J. Lin. Iteration complexity of feasible descent methods for convex optimization. *The Journal of Machine Learning Research*, 15(1):1523–1548, 2014. (Cited on page 2.)

[38] Q. Wang, P. Gong, S. Chang, T. S. Huang, and J. Zhou. Robust convex clustering analysis. In *ICDM*, pages 1263–1268. IEEE, 2016. (Cited on pages 2 and 17.)

[39] Y. Wang, J. Sharpnack, A. J. Smola, and R. J. Tibshirani. Trend filtering on graphs. *The Journal of Machine Learning Research*, 17(1):3651–3691, 2016. (Cited on pages 1, 5, 15, and 21.)

[40] C. Wu, S. Kwon, X. Shen, and W. Pan. A new algorithm and theory for penalized regression-based clustering. *Journal of Machine Learning Research*, 17:125, 2016. (Cited on page 1.)

[41] Y. Yuan, D. Sun, and K. Toh. An efficient semismooth Newton based algorithm for convex clustering. In *ICML*, 2018. (Cited on pages 2, 6, and 15.)

[42] Y. Zhang and L. Xiao. Stochastic primal-dual coordinate method for regularized empirical risk minimization. *The Journal of Machine Learning Research*, 18(1):2939–2980, 2017. (Cited on pages 2 and 6.)

[43] Z. Zhou, Q. Zhang, and A. M.-C. So. $\ell_{1,p}$-norm regularization: error bounds and convergence rate analysis of first-order methods. In *ICML*, pages 1501–1510. JMLR. org, 2015. (Cited on pages 2, 3, 7, 9, and 11.)

[44] C. Zhu, H. Xu, C. Leng, and S. Yan. Convex optimization procedure for clustering: Theoretical revisit. In *NIPS*, 2014. (Cited on page 1.)