

Composite convex minimization involving self-concordant-like cost functions

Quoc Tran-Dinh, Yen-Huan Li *and* Volkan Cevher
Laboratory for Information and Inference Systems (LIONS)
EPFL, Lausanne, Switzerland

Abstract. The self-concordant-like property of a smooth convex function is a new analytical structure that generalizes the self-concordant notion. While a wide variety of important applications feature the self-concordant-like property, this concept has heretofore remained unexploited in convex optimization. To this end, we develop a variable metric framework of minimizing the sum of a “simple” convex function and a self-concordant-like function. We introduce a new analytic step-size selection procedure and prove that the basic gradient algorithm has improved convergence guarantees as compared to “fast” algorithms that rely on the Lipschitz gradient property. Our numerical tests with real-data sets show that the practice indeed follows the theory.

1 Introduction

In this paper, we consider the following composite convex minimization problem:

$$F^* := \min_{\mathbf{x} \in \mathbb{R}^n} \{F(\mathbf{x}) := f(\mathbf{x}) + g(\mathbf{x})\}, \quad (1)$$

where f is a nonlinear smooth convex function, while g is a “simple” possibly nonsmooth convex function. Such composite convex problems naturally arise in many applications of machine learning, data sciences, and imaging science. Very often, f measures a data fidelity or a loss function, and g encodes a form of low-dimensionality, such as sparsity or low-rankness.

To trade-off accuracy and computation optimally in large-scale instances of (1), existing optimization methods invariably invoke the additional assumption that the smooth function f also has an L -Lipschitz continuous gradient (cf., [11] for the definition). A highlight is the recent developments on proximal gradient methods, which feature (nearly) dimension-independent, global sublinear convergence rates [3,9,11]. When the smooth f in (1) also has strong regularity [15], the problem (1) is also within the theoretical and practical grasp of proximal-(quasi) Newton algorithms with linear, superlinear, and quadratic convergence rates [5,8,17]. These algorithms specifically exploit second order information or its principled approximations (e.g., via BFGS or L-BFGS updates [13]).

In this paper, we do away with the Lipschitz gradient assumption and instead focus on another structural assumption on f in developing an algorithmic framework for (1), which is defined below.

Definition 1. A convex function $f \in C^3(\mathbb{R}^n)$ is called a self-concordant-like function $f \in \mathcal{F}_{\text{scl}}$, if:

$$|\varphi'''(t)| \leq M_f \varphi''(t) \|\mathbf{u}\|_2, \quad (2)$$

for $t \in \mathbb{R}$ and $M_f > 0$, where $\varphi(t) := f(\mathbf{x} + t\mathbf{u})$ for any $\mathbf{x} \in \text{dom}(f)$ and $\mathbf{u} \in \mathbb{R}^n$.

Definition 1 mimics the standard self-concordance concept ([10, Definition 4.1.1]) and was first discussed in [1] for model consistency in logistic regression. For composite convex minimization, self-concordant-like functions abound in machine learning, including but not limited to logistic regression, multinomial logistic regression, conditional random fields, and robust regression (cf., the references in [2]). In addition, special instances of geometric programming [6] can also be recast as (1) where $f \in \mathcal{F}_{\text{scl}}$.

The importance of the assumption $f \in \mathcal{F}_{\text{scl}}$ in (1) is twofold. First, it enables us to derive an explicit step-size selection strategy for proximal variable metric methods, enhancing backtracking-line search operations with improved theoretical convergence guarantees. For instance, we can prove that our proximal gradient method can automatically adapt to the local strong convexity of f near the optimal solution to feature linear convergence under mild conditions. This theoretical result is backed up by great empirical performance on real-life problems where the fast Lipschitz-based methods actually exhibit sublinear convergence (cf. Section 4). Second, the self-concordant-like assumption on f also helps us provide scalable numerical solutions of (1) for specific problems where f does not have Lipschitz continuous gradient, such as special forms of geometric programming problems.

Contributions. Our specific contributions can be summarized as follows:

1. We propose a new *variable metric* framework for minimizing the sum $f + g$ of a self-concordant-like function f and a convex, possibly nonsmooth function g . Our approach relies on the solution of a convex subproblem obtained by linearizing and regularizing the first term f , and uses an *analytical* step-size to achieve descent in three classes of algorithms: first order methods, second order methods, and quasi-Newton methods.
2. We establish both the global and the local convergence of different variable metric strategies. We pay particular attention to diagonal variable metrics since in this case many of the proximal subproblems can be solved exactly. We derive conditions on when and where these variants achieve locally linear convergence. When the variable metric is the Hessian of f at each iteration, we show that the resulting algorithm locally exhibits quadratic convergence without requiring any globalization strategy such as a backtracking line-search.
3. We apply our algorithms to large-scale real-world and synthetic problems to highlight the strengths and the weaknesses of our variable-metric scheme.

Relation to prior work. Many of the composite problems with self-concordant-like f , such as regularized logistics and multinomial logistics, also have Lipschitz continuous gradient. In those specific instances, many theoretically efficient algorithms are applicable [3,5,8,9,11,17]. Compared to these works, our framework has theoretically stronger local convergence guarantees thanks to the specific step-size strategy matched with $f \in \mathcal{F}_{\text{scl}}$. The authors of [18] consider composite problems where f is standard self-concordant and proposes a proximal Newton algorithm optimally exploiting this structure. Our structural assumptions and algorithmic emphasis here are different.

Paper organization. We first introduce the basic definitions and optimality conditions before deriving the variable metric strategy in Section 2. Section 3 proposes our new variable metric framework, describes its step-size selection

procedure, and establishes the convergence theory of its variants. Section 4 illustrates our framework in real and synthetic data.

2 Preliminaries

We adopt the notion of self-concordant functions in [10,12] to a different smooth function class. Then we present the optimality condition of problem (1).

2.1 Basic definitions

Let $g : \mathbb{R}^n \rightarrow \mathbb{R}$ be a proper, lower semicontinuous convex function [16] and $\text{dom}(g)$ denote the domain of g . We use $\partial g(\mathbf{x})$ to denote the subdifferential of g at $\mathbf{x} \in \text{dom}(g)$ if g is nondifferentiable at \mathbf{x} and $\nabla g(\mathbf{x})$ to denote its gradient, otherwise. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a $\mathcal{C}^3(\text{dom}(f))$ function (i.e., f is three times continuously differentiable). We denote by $\nabla f(\mathbf{x})$ and $\nabla^2 f(\mathbf{x})$ the gradient and the Hessian of f at \mathbf{x} , respectively. Suppose that, for a given $\mathbf{x} \in \text{dom}(f)$, $\nabla^2 f(\mathbf{x})$ is positive definite (i.e., $\nabla^2 f(\mathbf{x}) \in \mathcal{S}_{++}^n$), we define the local norm of a given vector $\mathbf{u} \in \mathbb{R}^n$ as $\|\mathbf{u}\|_{\mathbf{x}} := [\mathbf{u}^T \nabla^2 f(\mathbf{x}) \mathbf{u}]^{1/2}$. The corresponding dual norm of \mathbf{u} , $\|\mathbf{u}\|_{\mathbf{x}}^*$ is defined as $\|\mathbf{u}\|_{\mathbf{x}}^* := \max \{ \mathbf{u}^T \mathbf{v} \mid \|\mathbf{v}\|_{\mathbf{x}} \leq 1 \} = [\mathbf{u}^T \nabla^2 f(\mathbf{x})^{-1} \mathbf{u}]^{1/2}$.

2.2 Composite self-concordant-like minimization

Let $f \in \mathcal{F}_{\text{scl}}(\mathbb{R}^n)$ and g be proper, closed and convex. The optimality condition for (1) can be concisely written as follows:

$$0 \in \nabla f(\mathbf{x}^*) + \partial g(\mathbf{x}^*). \quad (3)$$

Let us denote by \mathbf{x}^* as an optimal solution of (1). Then, the condition (3) is necessary and sufficient. We also say that \mathbf{x}^* is *nonsingular* if $\nabla^2 f(\mathbf{x}^*)$ is positive definite. We now establish the existence and uniqueness of the solution \mathbf{x}^* of (1), whose proof can be found in [19].

Lemma 1. *Suppose that $f \in \mathcal{F}_{\text{scl}}(\mathbb{R}^n)$ satisfies Definition 1 for some $M_f > 0$. Suppose further that $\nabla^2 f(\mathbf{x}) \succ 0$ for some $\mathbf{x} \in \text{dom}(f)$. Then the solution \mathbf{x}^* of (1) exists and is unique.*

For a given symmetric positive definite matrix \mathbf{H} , we define a generalized proximal operator $\text{prox}_{\mathbf{H}^{-1}g}$ as:

$$\text{prox}_{\mathbf{H}^{-1}g}(\mathbf{x}) := \arg \min_{\mathbf{z}} \{ g(\mathbf{z}) + (1/2) \|\mathbf{z} - \mathbf{x}\|_{\mathbf{H}^{-1}}^2 \}. \quad (4)$$

Due to the convexity of g , this operator is well-defined and single-valued. If we can compute $\text{prox}_{\mathbf{H}^{-1}g}$ efficiently (e.g., by a closed form or by polynomial time algorithms), then we say that g is *proximally tractable*. Examples of proximal tractability convex functions can be found, e.g., in [14]. Using $\text{prox}_{\mathbf{H}^{-1}g}$, we can write condition (1) as:

$$\mathbf{x}^* - \mathbf{H}^{-1} \nabla f(\mathbf{x}^*) \in (\mathbb{I} + \mathbf{H}^{-1} \partial g)(\mathbf{x}^*) \iff \mathbf{x}^* = \text{prox}_{\mathbf{H}^{-1}g}(\mathbf{x}^* - \mathbf{H}^{-1} \nabla f(\mathbf{x}^*)).$$

This expression shows that \mathbf{x}^* is a fixed point of $\mathcal{R}_{\mathbf{H}}(\cdot) := \text{prox}_{\mathbf{H}^{-1}g}(\cdot) - \mathbf{H}^{-1} \nabla f(\cdot)$. Based on the fixed point principle, one can expect that the iterative sequence $\{\mathbf{x}^k\}_{k \geq 0}$ generated by $\mathbf{x}^{k+1} := \mathcal{R}_{\mathbf{H}}(\mathbf{x}^k)$ converges to \mathbf{x}^* . This observation is made rigorous below.

3 Our variable metric framework

We first present a generic variable metric proximal framework for solving (1). Then, we specify this framework to obtain three variants: proximal gradient, proximal Newton and proximal quasi-Newton algorithms.

3.1 Generic variable metric proximal algorithmic framework

Given $\mathbf{x}^k \in \text{dom}(F)$ and an appropriate choice $\mathbf{H}_k \in \mathcal{S}_{++}^n$, since $f \in \mathcal{F}_{\text{scl}}$, one can approximate f at \mathbf{x}^k by the following quadratic model:

$$Q_{\mathbf{H}_k}(\mathbf{x}, \mathbf{x}^k) := f(\mathbf{x}^k) + \langle \nabla f(\mathbf{x}^k), \mathbf{x} - \mathbf{x}^k \rangle + \frac{1}{2} \langle \mathbf{H}_k(\mathbf{x} - \mathbf{x}^k), \mathbf{x} - \mathbf{x}^k \rangle. \quad (5)$$

Our algorithmic approach uses the variable metric forward-backward framework to generate a sequence $\{\mathbf{x}^k\}_{k \geq 0}$ starting from $\mathbf{x}^0 \in \text{dom}(F)$ and update:

$$\mathbf{x}^{k+1} := \mathbf{x}^k + \alpha_k \mathbf{d}^k \quad (6)$$

where $\alpha_k \in (0, 1]$ is a given step-size and \mathbf{d}^k is a search direction defined by:

$$\mathbf{d}^k := \mathbf{s}^k - \mathbf{x}^k, \quad \text{with } \mathbf{s}^k := \underset{\mathbf{x}}{\text{argmin}} \{Q_{\mathbf{H}_k}(\mathbf{x}, \mathbf{x}^k) + g(\mathbf{x})\}. \quad (7)$$

In the rest of this section, we explain how to determine the step size α_k in the iterative scheme (6) optimally for special cases of \mathbf{H}_k . For this, we need the following definitions:

$$\lambda_k := \|\mathbf{d}^k\|_{\mathbf{x}^k}, \quad r_k := M_f \|\mathbf{d}^k\|_2, \quad \text{and } \beta_k := \|\mathbf{d}^k\|_{\mathbf{H}_k} = \langle \mathbf{H}_k \mathbf{d}^k, \mathbf{d}^k \rangle^{1/2}. \quad (8)$$

3.2 Proximal-gradient algorithm

When the variable matrix \mathbf{H}_k is *diagonal* and g is proximally tractable, we can efficiently obtain the solution of the subproblem (7) in a distributed fashion or even in a closed form. Hence, we consider $\mathbf{H}_k = \mathbf{D}_k := \text{diag}(\mathbf{D}_{k,1}, \dots, \mathbf{D}_{k,n})$ with $\mathbf{D}_{k,i} > 0$, for $i = 1, \dots, n$. Lemma 2, whose proof is in [19], provides a step-size selection procedure and proves the global convergence of this proximal-gradient algorithm.

Lemma 2. *Let $\{\mathbf{x}^k\}_{k \geq 0}$ be a sequence generated by (6) and (7) starting from $\mathbf{x}^0 \in \text{dom}(F)$. For λ_k , r_k and β_k defined by (8), we consider the step-size α_k as:*

$$\alpha_k := \frac{1}{r_k} \ln \left(1 + \frac{\beta_k^2 r_k}{\lambda_k^2} \right), \quad (9)$$

If $\beta_k^2 r_k \leq (e^{r_k} - 1) \lambda_k^2$, then $\alpha_k \in (0, 1]$ and:

$$F(\mathbf{x}^{k+1}) \leq F(\mathbf{x}^k) - \frac{\beta_k^2}{r_k} \left[\left(1 + \frac{\lambda_k^2}{r_k \beta_k^2} \right) \ln \left(1 + \frac{\beta_k^2 r_k}{\lambda_k^2} \right) - 1 \right]. \quad (10)$$

Moreover, this step-size α_k is optimal (w.r.t. the worst-case performance).

Algorithm 1 (Proximal-gradient algorithm with a *diagonal variable metric*)

Initialization: Given $\mathbf{x}^0 \in \text{dom}(F)$, and a tolerance $\varepsilon > 0$.

for $k = 0$ **to** k_{\max} **do**

1. Choose $\mathbf{D}_k \in \mathcal{S}_{++}^n$ (e.g., using $\mathbf{D}_k := L_k \mathbb{I}$, where L_k is given by (11)).
2. Compute the proximal-gradient search direction \mathbf{d}^k as (7).
3. Compute $\beta_k := \|\mathbf{d}^k\|_{\mathbf{D}_k}$, $r_k := M_f \|\mathbf{d}^k\|_2$ and $\lambda_k := \|\mathbf{d}^k\|_{\mathbf{x}^k}$.
4. If $\beta_k \leq \varepsilon$ then terminate.
5. If $\beta_k^2 r_k \leq (e^{r_k} - 1) \lambda_k^2$, then compute $\alpha_k := \frac{1}{r_k} \ln \left(1 + \frac{\beta_k^2 r_k}{\lambda_k^2} \right)$ and update $\mathbf{x}^{k+1} := \mathbf{x}^k + \alpha_k \mathbf{d}^k$. Otherwise, set $\mathbf{x}^{k+1} := \mathbf{x}^k$ and update \mathbf{D}_{k+1} from \mathbf{D}_k .

end for

By our condition, the second term on the right-hand side of (10) is always positive, establishing that the sequence $\{F(\mathbf{x}^k)\}$ is decreasing. Moreover, as $e^{r_k} - 1 \geq r_k$, the condition $\beta_k^2 r_k \leq (e^{r_k} - 1) \lambda_k^2$ can be simplified to $\beta_k \leq \lambda_k$. It is easy to verify that this is satisfied whenever $\mathbf{D}_k \preceq \nabla^2 f(\mathbf{x}^k)$. In such cases, our step-size selection ensures the best decrease of the objective value regarding the self-concordant-like structure of f (and not the actual objective instance). When $\beta_k > \lambda_k$, we scale down \mathbf{D}_k until $\beta_k \leq \lambda_k$. It is easy to prove that the number of backtracking steps to find $\mathbf{D}_{k,i}$ is time constant.

Now, by using our step-size (9), we can describe the proximal-gradient algorithm as in Algorithm 1.

We combine the above analysis to obtain the following proximal gradient algorithm for solving (1). The main step in Algorithm 1 is to compute the search direction \mathbf{d}^k at Step 2, which is equivalent to the solution of the convex sub-problem (7). The second main step is to compute $\lambda_k = \langle \nabla^2 f(\mathbf{x}^k) \mathbf{d}^k, \mathbf{d}^k \rangle^{1/2}$. This quantity requires the product of Hessian $\nabla^2 f(\mathbf{x}^k)$ of f and \mathbf{d}^k , but not the full-Hessian. It is clear that if $\beta_k = 0$ then $\mathbf{d}^k = 0$ and $\mathbf{x}^{k+1} \equiv \mathbf{x}^k$ and we obtain the solution of (1), i.e., $\mathbf{x}^k \equiv \mathbf{x}^*$. The diagonal matrix \mathbf{D}_k can be updated as $\mathbf{D}_{k+1} := c \mathbf{D}_k$ for a given factor $c > 1$.

We now explain how the new theory enhances the standard backtracking linesearch approaches. For simplicity, let us assume $\mathbf{D}_k := L_k \mathbb{I}$, where \mathbb{I} is the identity matrix. By a careful inspection of (10), we see that $L_k = \sigma_{\max}(\nabla^2 f(\mathbf{x}^k))$ achieves the maximum guaranteed decrease (in the worst case sense) in the objective. There are many principled ways of approximating this constant based on the secant equation underlying the quasi-Newton methods. In Section 4, we use Barzilai-BenTal's rule:

$$L_k := \frac{\|\mathbf{y}^k\|_2^2}{\langle \mathbf{y}^k, \mathbf{s}^k \rangle}, \text{ where } \mathbf{s}^k := \mathbf{x}^k - \mathbf{x}^{k-1} \text{ and } \mathbf{y}^k := \nabla f(\mathbf{x}^k) - \nabla f(\mathbf{x}^{k-1}). \quad (11)$$

We then deviate from the standard backtracking approaches. As opposed to, for instance, checking the Armijo-Goldstein condition, we use a *new analytic condition* (i.e., Step 5 of Algorithm 1), which is computationally cheaper in many cases. Our analytic step-size then further refines the solution based on the worst-case problem structure, even if the backtracking update satisfies the Armijo-Goldstein condition.

Surprisingly, our analysis also enables us to also establish local linear convergence as described in Theorem 1 under mild assumptions. The proof can be found in [19].

Theorem 1. Let $\{\mathbf{x}^k\}_{k \geq 0}$ be a sequence generated by Algorithm 1. Suppose that the sub-level set $\mathcal{L}_F(F(\mathbf{x}^0)) := \{\mathbf{x} \in \text{dom}(F) : F(\mathbf{x}) \leq F(\mathbf{x}^0)\}$ is bounded and $\nabla^2 f$ is nonsingular at some $\mathbf{x} \in \text{dom}(f)$. Suppose further that $\mathbf{D}_k := L_k \mathbb{I} \succeq \tau \mathbb{I}_n$ for given $\tau > 0$. Then, $\{\mathbf{x}^k\}$ converges to \mathbf{x}^* the solution of (1). Moreover, if $\rho_* := \max\{L_k/\sigma_{\min}^* - 1, 1 - L_k/\sigma_{\max}^*\} < \frac{1}{2}$ for k sufficiently large then the sequence $\{\mathbf{x}^k\}$ locally converges to \mathbf{x}^* at a linear rate, where σ_{\min}^* and σ_{\max}^* are the smallest and the largest eigenvalues of $\nabla^2 f(\mathbf{x}^*)$, respectively.

Linear convergence: According to Theorem 1, linear convergence is only possible when the condition number κ of the Hessian at the true solution satisfies $\kappa = \sigma_{\max}^*/\sigma_{\min}^* < 3$. While this seems too imposing, we claim that, for most f and g , this requirement is not too difficult to satisfy (see also the empirical evidence in Section 4). This is because the proof of Theorem 1 only needs the smallest and the largest eigenvalues of $\nabla^2 f(\mathbf{x}^*)$, restricted to the subspaces of the union of $\mathbf{x}^* - \mathbf{x}^k$ for k sufficiently large, to satisfy the conditions imposed by ρ_* . For instance, when g is based on the ℓ_1 -norm/the nuclear norm, the differences $\mathbf{x}^* - \mathbf{x}^k$ have at most twice the sparsity/rank of \mathbf{x}^* near convergence. Given such subspace restrictions, one can prove, via probabilistic assumptions on f (cf., [1]), that the restricted condition number is not only dramatically smaller than the full condition number κ of the Hessian $\nabla^2 f(\mathbf{x}^*)$, but also it can even be dimension independent with high probability.

3.3 Proximal-Newton algorithm

The case $\mathbf{H}_k \equiv \nabla^2 f(\mathbf{x}^k)$ deserves a special attention as the step-size selection rule becomes explicit and backtracking-free. The resulting method is a *proximal-Newton* method and can be computationally attractive in certain big data problems due to its low iteration count.

The main step of the proximal-Newton algorithm is to compute the proximal-Newton search direction \mathbf{d}^k as:

$$\mathbf{d}^k := \mathbf{s}^k - \mathbf{x}^k, \text{ where } \mathbf{s}^k := \underset{\mathbf{x}}{\text{argmin}} \{Q_{\nabla^2 f(\mathbf{x}^k)}(\mathbf{x}, \mathbf{x}^k) + g(\mathbf{x})\}. \quad (12)$$

Then, it updates the sequence $\{\mathbf{x}^k\}$ by:

$$\mathbf{x}^{k+1} := \mathbf{x}^k + \alpha_k \mathbf{d}^k = (1 - \alpha_k) \mathbf{x}^k + \alpha_k \mathbf{s}^k, \quad (13)$$

where $\alpha_k \in (0, 1]$ is the step size. If we set $\alpha_k = 1$ for all $k \geq 0$, then (13) is called the full-step proximal-Newton method. Otherwise, it is a damped-step proximal-Newton method.

First, we show how to compute the step size α_k in the following lemma, which is a direct consequence of Lemma 2 by taking $\mathbf{H}_k \equiv \nabla^2 f(\mathbf{x}^k)$.

Lemma 3. Let $\{\mathbf{x}^k\}_{k \geq 0}$ be a sequence generated by the proximal-Newton scheme (13) starting from $\mathbf{x}^0 \in \text{dom}(F)$. Let λ_k and r_k be as defined by (8). If we choose the step-size $\alpha_k = r_k^{-1} \ln(1 + r_k)$ then:

$$F(\mathbf{x}^{k+1}) \leq F(\mathbf{x}^k) - r_k^{-1} \lambda_k^2 [(1 + r_k^{-1}) \ln(1 + r_k) - 1]. \quad (14)$$

Moreover, this step-size α_k is optimal (w.r.t. the worst-case performance).

Next, Theorem 2 proves the local quadratic convergence of the full-step proximal-Newton method, whose proof can be found in [19].

Theorem 2. *Suppose that the sequence $\{\mathbf{x}^k\}_{k \geq 0}$ is generated by (13) with full-step, i.e., $\alpha_k = 1$ for $k \geq 0$. If $r_k \leq \ln(4/3) \approx 0.28768207$ then it holds that:*

$$\left(\lambda_{k+1} / \sqrt{\sigma_{\min}^{k+1}} \right) \leq 2M_f \left(\lambda_k / \sqrt{\sigma_{\min}^k} \right)^2, \quad (15)$$

where σ_{\min}^k is the smallest eigenvalue of $\nabla^2 f(\mathbf{x}^k)$. Consequently, if we choose \mathbf{x}^0 such that $\lambda_0 \leq \sigma_{\min}(\nabla^2 f(\mathbf{x}^0)) \ln(4/3)$, then the sequence $\left\{ \lambda_k / \sqrt{\sigma_{\min}^k} \right\}$ converges to zero at a quadratic rate.

Theorem 2 rigorously establishes where we can take full steps and still have quadratic convergence. Based on this information, we propose the proximal-Newton algorithm as in Algorithm 2.

Algorithm 2 (Prototype proximal-Newton algorithm)

Initialization: Given $\mathbf{x}^0 \in \text{dom}(F)$ and $\sigma \in (0, \sigma_{\min}(\nabla^2 f(\mathbf{x}^0)) \ln(4/3)]$.

for $k = 0$ **to** k_{\max} **do**

1. Compute \mathbf{s}^k by (12). Then, define $\mathbf{d}^k := \mathbf{s}^k - \mathbf{x}^k$ and $\lambda_k := \|\mathbf{d}^k\|_{\mathbf{x}^k}$.
2. If $\lambda_k \leq \varepsilon$, then terminate.
3. If $\lambda_k > \sigma$, then compute $r_k := M_f \|\mathbf{d}^k\|_2$ and $\alpha_k := \frac{1}{r_k} \ln(1 + r_k)$; else $\alpha_k := 1$.
4. Update $\mathbf{x}^{k+1} := \mathbf{x}^k + \alpha_k \mathbf{d}^k$.

end for

The most remarkable feature of Algorithm 2 is that it does not require any globalization strategy such as backtracking line search for global convergence.

Complexity analysis. First, we estimate the number of iterations needed when $\lambda_k \leq \sigma$ to reach the solution \mathbf{x}^k such that $\frac{\lambda_k}{\sqrt{\sigma_k}} \leq \varepsilon$ for a given tolerance $\varepsilon > 0$. Based on the conclusion of Theorem 2, we can show that the number of iterations of Algorithm 2 when $\lambda_k > \sigma$ does not exceed $k_{\max} := \left\lceil \log_2 \left(\frac{\ln(2M_f \varepsilon)}{\ln(2\sigma)} \right) \right\rceil$. Finally, we estimate the number of iterations needed when $\lambda_k > \sigma$. From Lemma 3, we see that for all $k \geq 0$ we have $\lambda_k \geq \sigma$ and $r_k \geq \sigma$. Therefore, the number of iterations is $\left\lceil \frac{F(\mathbf{x}^0) - F(\mathbf{x}^*)}{\psi(\sigma)} \right\rceil$, where $\psi(\tau) := \tau((1 + \tau^{-1}) \ln(1 + \tau) - 1) > 0$.

3.4 Proximal quasi-Newton algorithm

In many applications, estimating the Hessian $\nabla^2 f(\mathbf{x}^k)$ can be costly even though the Hessian is given in a closed form (cf., Section 4). In such cases, variable metric strategies employing approximate Hessian can provide computation-accuracy tradeoffs. Among these approximations, applying quasi-Newton methods with BFGS updates for \mathbf{H}_k would ensure its positive definiteness. Our analytic step-size procedures with backtracking automatically applies to the BFGS proximal-quasi Newton method, whose algorithm details and convergence analysis are omitted here.

4 Numerical experiments

We use a variety of different real-data problems to illustrate the performance of our variable metric framework using a MATLAB implementation. We pick two advanced solvers for comparison: TFOCS [4] and PNOPT [8]. TFOCS hosts accelerated first order methods. PNOPT provides a several proximal-(quasi)

Newton implementations, which has been shown to be quite successful in logistic regression problems [8]. Both use sophisticated backtracking linesearch enhancements. We benchmark all algorithms with performance profiles [7].

A performance profile is built based on a set \mathcal{S} of n_s algorithms (solvers) and a collection \mathcal{P} of n_p problems. We first build a profile based on computational time. We denote by $T_{p,s} :=$ *computational time required to solve problem p by solver s* . We compare the performance of algorithm s on problem p with the best performance of any algorithm on this problem; that is we compute the performance ratio $r_{p,s} := \frac{T_{p,s}}{\min\{T_{p,\tilde{s}} : \tilde{s} \in \mathcal{S}\}}$. Now, let $\tilde{\rho}_s(\tilde{\tau}) := \frac{1}{n_p} \text{size} \{p \in \mathcal{P} : r_{p,s} \leq \tilde{\tau}\}$ for $\tilde{\tau} \in \mathbb{R}_+$. The function $\tilde{\rho}_s : \mathbb{R} \rightarrow [0, 1]$ is the probability for solver s that a performance ratio is within a factor $\tilde{\tau}$ of the best possible ratio. We use the term ‘‘performance profile’’ for the distribution function $\tilde{\rho}_s$ of a performance metric. In the following numerical examples, we plotted the performance profiles in \log_2 -scale, i.e. $\rho_s(\tau) := \frac{1}{n_p} \text{size} \{p \in \mathcal{P} : \log_2(r_{p,s}) \leq \tau := \log_2 \tilde{\tau}\}$.

4.1 Sparse logistic regression

We consider the classical logistic regression problem of the form [20]:

$$\min_{\mathbf{x}, \mu} \left\{ N^{-1} \sum_{j=1}^N \log \left(1 + e^{-y_j (\langle \mathbf{w}^{(j)}, \mathbf{x} \rangle + \mu)} \right) + \rho N^{-1/2} \|\mathbf{x}\|_1 \right\}, \quad (16)$$

where $\mathbf{x} \in \mathbb{R}^p$ is an unknown vector, μ is an unknown bias, and $y^{(j)}$ and \mathbf{w}^j are observations where $j = 1, \dots, N$. The logistic term in (16) is self-concordant-like with $M_f := \max \|\mathbf{w}^{(j)}\|_2$ [1]. In this case, the smooth term in (16) has Lipschitz gradient, hence several fast algorithms are applicable.

Figure 1 illustrates the performance profiles for computational time (left) and the number of prox-operations (right) using the 36 medium size problems¹. For comparison, we use TFOCS-N07, which is Nesterov’s 2007 two prox-method; and TFOCS-AT, which is Auslender and Teboulle’s accelerated method, PNOPT with L-BFGS updates, and our algorithms: proximal gradient and proximal-Newton. From these performance profiles, we can observe that our proximal gradient is the best one in terms of computational time and the number of prox-operations. In terms of time, proximal-gradient solves upto 83.3% of problems with the best performance, while these numbers in TFOCS-N07 and PNOPT-LBFGS are 2.7%. Proximal Newton algorithm solves 11.1% problems with the best performance. In prox-operations, proximal-gradient is also the best one in 75% of problems.

We now show an example convergence behavior of our proximal-gradient algorithm via two large-scale problems with $\rho = 0.1$. The first problem is `rcv1_train.binary` with the size $p = 20242$ and $N = 47236$ and the second one is `real-sim` with the size $p = 72309$ and $N = 20958$. For comparison, we use TFOCS-N07 and TFOCS-AT. For this example, PNOPT (with Newton, BFGS, and L-BFGS options) and our proximal-Newton do not scale and are omitted.

Figure 2 shows that our simple gradient algorithm locally exhibits linear convergence whereas the fast method TFOCS-AT shows a sublinear convergence rate. The variant TFOCS-N07 is the Nesterov’s dual proximal algorithm, which exhibits oscillations but performs comparable to our proximal gradient method

¹ Available at <http://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>.

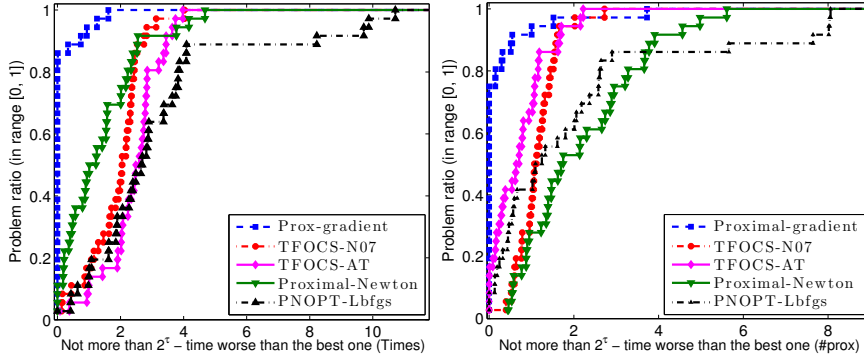


Fig. 1. Computational time (*left*) and number of prox-operations (*right*)




Fig. 2. *Left*: rcv1_train.binary, and *Right*: real-sim.

in terms of accuracy, time, and the total number of prox operations. The computational time and the number of prox-operations in these both problems are given as follows: Proximal-gradient: (15.67s, 698), (13.71s, 152); TFOCS-AT: (20.57s, 678), (33.82s, 466); TFOCS-N07: (17.09s, 1049), (22.08s, 568), respectively. For these data sets, the relative performance of the algorithms is surprisingly consistent across various regularization parameters.

4.2 Restricted condition number in practice

The convergence plots in Figure 2 indicate that the linear convergence condition in Theorem 1 may be satisfied. *In fact, in all of our tests, the proximal gradient algorithm exhibits locally linear convergence.* Hence, to see if Remark 1 is grounded in practice, we perform the following test on the `a#a` dataset¹, consisting of small to medium problems. We first solve each problem with the proximal-Newton method up to 16 digits of accuracy to obtain \mathbf{x}^* , and we calculate $\nabla^2 f(\mathbf{x}^*)$. We then run our proximal gradient algorithm until convergence, and during its linear convergence, we record $\|\nabla^2 f(\mathbf{x}^*)(\mathbf{x}^* - \mathbf{x}^k)\|^2 / \|\mathbf{x}^* - \mathbf{x}^k\|_2^2$, and take the ratios of the maximum and the minimum to estimate the restricted condition number for each problem.

Figure 3 illustrates that while the condition number of the Hessian $\nabla^2 f(\mathbf{x}^*)$ can be extremely large, as the algorithm navigates to the optimal solution \mathbf{x}^* through sparse subspaces, the restricted condition number estimates are in fact very close to 3. Given that algorithm still exhibit linear convergence for the cases $\# = 2, 3, 4, 5, 6, 8$ (where our condition cannot be met), we believe that




Fig. 3. Restricted condition number (*left*), and condition number (*right*) estimates

the tightness of our convergence condition is an artifact of our proof and may be improved.

4.3 Sparse multinomial logistic regression

For sparse multimonomial logistic regression, the underlying problem is formulated in the form of (1), which the objective function f is given as:

$$f(\mathbf{X}) := N^{-1} \sum_{j=1}^N \left[\log \left(1 + \sum_{i=1}^m e^{\langle \mathbf{w}^{(j)}, \mathbf{X}^{(i)} \rangle} \right) - \sum_{i=1}^m \mathbf{y}_i^{(j)} \langle \mathbf{w}^{(j)}, \mathbf{X}^{(i)} \rangle \right]. \quad (17)$$

where \mathbf{X} can be considered as a matrix variable of size $m \times p$ formed from $\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(m)}$. Other vectors, $\mathbf{y}^{(j)}$ and $\mathbf{w}^{(j)}$ are given as input data for $j = 1, \dots, N$. The function f has closed form gradient as well as Hessian. However, forming a full hessian matrix $\nabla^2 f(\mathbf{x})$ is especially costly in large scale problems when $N \gg 1$. In this case, proximal-quasi-Newton methods are more suitable. First, we show in Lemma 4 that f satisfies Definition 1, whose proof is in [19].

Lemma 4. *The function f defined by (17) is convex and self-concordant-like in the sense of Definition 1 with the parameter $M_f := \sqrt{6}N^{-1} \max_{j=1, \dots, N} \|\mathbf{w}^{(j)}\|_2$.*

The performance profiles of 20 small-to-medium size problems¹ are shown in Figure 4 in terms of computational time (left) as well as number of prox-operations (right), respectively. Both proximal-gradient method and proximal-Newton method with BFGS have good performance. They can solve unto 55% and 45% problems with the best time performance, respectively. These methods are also the best in terms of prox-operations (70% and 30%).

4.4 A sytlized example of a non-Lipschitz gradient function for (1)

We consider the following convex composite minimization problem by modifying one of the canonical examples of geometric programming [6]:

$$\min_{\mathbf{x} \in \Omega} \left\{ f(\mathbf{x}) := \sum_{i=1}^m e^{\mathbf{a}_i^T \mathbf{x} + b_i} + \mathbf{c}^T \mathbf{x} \right\} + g(\mathbf{x}), \quad (18)$$

where Ω is a simple convex set, $\mathbf{a}_i, \mathbf{c} \in \mathbb{R}^n$ and $b_i \in \mathbb{R}$ are random, and g is the ℓ_1 -norm. After some algebra, we can show that f satisfies Definition 1 with $M_f := \max \{\|\mathbf{a}_i\|_2 : 1 \leq i \leq m\}$. Unfortunately, f does not have Lipschitz continuous gradient in \mathbb{R}^n .




Fig. 4. Computational time (*left*), and number of prox-operations (*right*)

We implement our proximal-gradient algorithm and compare it with TFOCS and PNOPT-LBFGS. However, TFOCS breaks down in running this example due to the estimation of Lipschitz constant, while PNOPT is rather slow. Several tests on synthetic data show that our algorithm outperforms PNOPT-LBFGS. As an example, we show the convergence behavior of both these methods in Figure 5 where we plot the accuracy of the objective values w.r.t. the number of prox-operators for two cases of $\varepsilon = 10^{-6}$ and $\varepsilon = 10^{-12}$, respectively. As we can




Fig. 5. Relative objective values w.r.t. #prox: *left*: $\varepsilon = 10^{-6}$, and *right*: $\varepsilon = 10^{-12}$.

see from this figure that our prox-gradient method requires many fewer prox-operations to achieve a very high accuracy compared to PNOPT. Moreover, our method is also 20 to 40 times faster than PNOPT in this numerical test.

5 Conclusions

Convex optimization efficiency relies significantly on the structure of the objective functions. In this paper, we propose a variable metric method for minimizing the sum of a self-concordant-like convex function and a proximally tractable convex function. Our framework is applicable in several interesting machine learning problems and do not rely on the usual Lipschitz gradient assumption on the smooth part for its convergence theory. A highlight of this work is the new analytic step-size selection procedure that enhances backtracking procedures. Thanks to this new approach, we can prove that the basic gradient variant of our framework has improved local convergence guarantees under certain conditions while the tuning-free proximal Newton method has locally quadratic convergence. While our assumption on the restricted condition number in Theorem

1 is not deterministically verifiable *a priori*, we provide empirical evidence that it can hold in many practical problems. Numerical experiments on different applications that have both self-concordant-like and Lipschitz gradient properties demonstrate that the gradient algorithm based on the former assumption can be more efficient than the fast algorithms based on the latter assumption. As a result, we plan to look into fast versions of our gradient scheme as future work.

Acknowledgments: This work is supported in part by the European Commission under the grants MIRG-268398 and ERC Future Proof, and by the Swiss Science Foundation under the grants SNF 200021-132548, SNF 200021-146750 and SNF CRSII-147633.

References

1. F. Bach. Self-concordant analysis for logistic regression. *Electron. J. Statist.*, 4:384–414, 2010.
2. F. Bach. Adaptivity of averaged stochastic gradient descent to local strong convexity for logistic regression. 2013.
3. A. Beck and M. Teboulle. A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems. *SIAM J. Imaging Sciences*, 2(1):183–202, 2009.
4. S. Becker, E. J. Candès, and M. Grant. Templates for convex cone problems with applications to sparse signal recovery. *Mathematical Programming Computation*, 3(3):165–218, 2011.
5. S. Becker and M.J. Fadili. A quasi-Newton proximal splitting method. In *Adv. Neural Information Processing Systems*, 2012.
6. S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
7. E.D. Dolan and J.J. Moré. Benchmarking optimization software with performance profiles. *Math. Program.*, 91:201–213, 2002.
8. J.D. Lee, Y. Sun, and M.A. Saunders. Proximal Newton-type methods for convex optimization. *SIAM J. Optim.*, 24(3):1420–1443, 2014.
9. H. Mine and M. Fukushima. A minimization method for the sum of a convex function and a continuously differentiable function. *J. Optim. Theory Appl.*, 33:9–23, 1981.
10. Y. Nesterov. *Introductory lectures on convex optimization: a basic course*, volume 87 of *Applied Optimization*. Kluwer Academic Publishers, 2004.
11. Y. Nesterov. Gradient methods for minimizing composite objective function. *Math. Program.*, 140(1):125–161, 2013.
12. Y. Nesterov and A. Nemirovski. *Interior-point Polynomial Algorithms in Convex Programming*. Society for Industrial Mathematics, 1994.
13. J. Nocedal and S.J. Wright. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer, 2 edition, 2006.
14. N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):123–231, 2013.
15. S. M. Robinson. Strongly Regular Generalized Equations. *Mathematics of Operations Research*, Vol. 5, No. 1 (Feb., 1980), pp. 43–62, 5:43–62, 1980.
16. R. T. Rockafellar. *Convex Analysis*, volume 28 of *Princeton Mathematics Series*. Princeton University Press, 1970.
17. M. Schmidt, N.L. Roux, and F. Bach. Convergence rates of inexact proximal-gradient methods for convex optimization. *NIPS, Granada, Spain*, 2011.
18. Q. Tran-Dinh, A. Kyrillidis, and V. Cevher. Composite self-concordant minimization. *J. Mach. Learn. Res. (accepted)*, 15:1–54, 2014.

19. Q. Tran-Dinh, Y.-H. Li, and V. Cevher. Composite convex minimization involving self-concordant-like cost functions. *LIONS-Tech. Report.*, 1-19, 2015. Online at: <http://arxiv.org/abs/1502.01068>.
20. Y. X. Yuan. Recent advances in numerical methods for nonlinear equations and nonlinear least squares. *Numerical Algebra, Control and Optimization*, 1(1):15–34, 2011.

A Appendix: Composite convex minimization involving self-concordant-like cost functions

We derive some fundamental properties of self-concordant-like functions, introduce the notion of scaled proximal operators, and provide the full-proofs of the technical results in the main text.

A.1 Properties of self-concordant-like functions

We define $D^2 f(\mathbf{x})[\mathbf{u}, \mathbf{u}] := \|\mathbf{u}\|_{\mathbf{x}}^2$ and $D^3 f(\mathbf{x})[\mathbf{u}, \mathbf{u}, \mathbf{u}] := \langle D^3 f(\mathbf{x})[\mathbf{u}]\mathbf{u}, \mathbf{u} \rangle$, following the notations in [10,12]. An equivalent definition of self-concordant-like functions is provided by the following theorem [1].

Theorem 3. *A convex function $f \in \mathcal{C}^3 : \mathbb{R}^n \rightarrow \mathbb{R}$ is M_f -self-concordant-like if and only if for any $\mathbf{x}, \mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3 \in \mathbb{R}^n$, we have:*

$$|D^3 f(\mathbf{x})[\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3]| \leq M_f \|\mathbf{u}_1\|_2 \|\mathbf{u}_2\|_{\mathbf{x}} \|\mathbf{u}_3\|_{\mathbf{x}}.$$

Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be an M_f -self-concordant-like function. Then:

- The function $f_q(\mathbf{x}) := \alpha + \langle \mathbf{a}, \mathbf{x} \rangle + (1/2)\langle \mathbf{A}\mathbf{x}, \mathbf{x} \rangle + f(\mathbf{x})$ is M_f -self-concordant-like, for any $\alpha \in \mathbb{R}$, $\mathbf{a} \in \mathbb{R}^n$, and positive symmetric matrix $\mathbf{A} : \mathbb{R}^n \rightarrow \mathbb{R}^n$.
- The function αf is M_f -self-concordant-like for any $\alpha \geq 0$.
- Let g be an M_g -self-concordant-like function. Then the function $(f + g)$ is M -self-concordant-like, where $M := \max\{M_f, M_g\}$.
- The function $f(\mathbf{A}\mathbf{x} + \mathbf{b})$ is $(M_f \|\mathbf{A}\|_2)$ -self-concordant-like, for any matrix $\mathbf{A} : \mathbb{R}^m \rightarrow \mathbb{R}^n$, $\mathbf{x} \in \mathbb{R}^m$, and $\mathbf{b} \in \mathbb{R}^n$.

We will repeatedly use the following inequalities in the rest of this appendix.

Theorem 4. *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be an M_f -self-concordant-like function, and define $\lambda_{\mathbf{x}}(\mathbf{y}) := \|\mathbf{y} - \mathbf{x}\|_{\mathbf{x}}$ and $r_{\mathbf{x}}(\mathbf{y}) := M_f \|\mathbf{y} - \mathbf{x}\|_2$ for all $\mathbf{x}, \mathbf{y} \in \text{dom}(f)$. We have the following inequalities for any $\mathbf{x}, \mathbf{y} \in \text{dom}(f)$:*

- Bounds on the local norm:

$$e^{-\frac{r_{\mathbf{x}}(\mathbf{y})}{2}} \lambda_{\mathbf{x}}(\mathbf{y}) \leq \lambda_{\mathbf{y}}(\mathbf{x}) \leq e^{\frac{r_{\mathbf{x}}(\mathbf{y})}{2}} \lambda_{\mathbf{x}}(\mathbf{y}). \quad (19)$$

- Bounds on the Hessian matrix:

$$e^{-r_{\mathbf{x}}(\mathbf{y})} \nabla^2 f(\mathbf{x}) \preceq \nabla^2 f(\mathbf{y}) \preceq e^{r_{\mathbf{x}}(\mathbf{y})} \nabla^2 f(\mathbf{x}). \quad (20)$$

- Bounds on the gradient vector:

$$\gamma_*(r_{\mathbf{x}}(\mathbf{y})) \lambda_{\mathbf{x}}(\mathbf{y})^2 \leq \langle \nabla f(\mathbf{y}) - \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle \leq \gamma(r_{\mathbf{x}}(\mathbf{y})) \lambda_{\mathbf{x}}(\mathbf{y})^2, \quad (21)$$

where $\gamma_*(\tau) := -\left(\frac{e^{-\tau}-1}{\tau}\right)$, and $\gamma(\tau) := \frac{e^{\tau}-1}{\tau}$.

- Bounds on the function value:

$$\omega_*(r_{\mathbf{x}}(\mathbf{y})) \lambda_{\mathbf{x}}(\mathbf{y})^2 \leq f(\mathbf{y}) - f(\mathbf{x}) - \nabla f(\mathbf{x})^T(\mathbf{y} - \mathbf{x}) \leq \omega(r_{\mathbf{x}}(\mathbf{y})) \lambda_{\mathbf{x}}(\mathbf{y})^2, \quad (22)$$

where $\omega_*(\tau) := \frac{e^{-\tau} + \tau - 1}{\tau^2}$ and $\omega(\tau) := \frac{e^{\tau} - \tau - 1}{\tau^2}$ are both strictly convex and increasing.

Proof. We denote $\mathbf{y}_t := \mathbf{x} + t(\mathbf{y} - \mathbf{x})$ for notational convenience, where $t \in [0, 1]$. First we prove (19). Consider the function $\phi(t) := \log(\langle \nabla^2 f(\mathbf{y}_t) \mathbf{u}, \mathbf{u} \rangle)$ for some $\mathbf{u} \in \mathbb{R}^n$. We have, by Definition 1, that:

$$|\phi'(t)| = \frac{|D^3 f(\mathbf{y}_t)[\mathbf{u}, \mathbf{u}, \mathbf{u}]|}{D^2 f(\mathbf{y}_t)[\mathbf{u}, \mathbf{u}]} \leq M_f \|\mathbf{u}\|_2.$$

Set $\mathbf{u} = \mathbf{y} - \mathbf{x}$, and then we have $\phi(0) = \log(\|\mathbf{y} - \mathbf{x}\|_{\mathbf{x}}^2)$ and $\phi(1) = \log(\|\mathbf{y} - \mathbf{x}\|_{\mathbf{y}}^2)$. Integrating $\phi'(\cdot)$ over the interval $[0, 1]$, we obtain:

$$\left| \log(\|\mathbf{y} - \mathbf{x}\|_{\mathbf{y}}^2) - \log(\|\mathbf{y} - \mathbf{x}\|_{\mathbf{x}}^2) \right| \leq M_f \|\mathbf{y} - \mathbf{x}\|_2,$$

which leads to (19).

Next, we prove (20). Consider the function $\psi(t) := \langle \nabla^2 f(\mathbf{y}_t) \mathbf{u}, \mathbf{u} \rangle = \|\mathbf{u}\|_{\mathbf{y}_t}^2$ for some $\mathbf{u} \in \mathbb{R}^n$. We have $\psi(0) = \|\mathbf{u}\|_{\mathbf{x}}^2$ and $\psi(1) = \|\mathbf{u}\|_{\mathbf{y}}^2$. By Theorem 3, we obtain:

$$|\psi'(t)| = |D^3 f(\mathbf{y}_t)[\mathbf{y} - \mathbf{x}, \mathbf{u}, \mathbf{u}]| \leq M_f \|\mathbf{y} - \mathbf{x}\|_2 \psi(t),$$

or, equivalently,

$$\left| \frac{d \ln \psi(t)}{dt} \right| \leq M_f \|\mathbf{y} - \mathbf{x}\|_2.$$

We get (20) by integrating both sides over $[0, 1]$.

Now, we prove (21). By the mean-value theorem, we have:

$$\langle f(\mathbf{y}) - \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle = \int_0^1 \langle \nabla^2 f(\mathbf{y}_t) \mathbf{x}, \mathbf{x} \rangle dt.$$

Applying the right-hand side of (20), we obtain:

$$\int_0^1 \langle \nabla^2 f(\mathbf{y}_t) \mathbf{x}, \mathbf{x} \rangle dt \leq \int_0^1 \exp(M_f \|\mathbf{y}_t - \mathbf{x}\|_2) \langle \nabla^2 f(\mathbf{x}) \mathbf{x}, \mathbf{x} \rangle dt,$$

which leads to the right-hand side of (21). Similarly, we can prove the left-hand side of (21).

Finally, (22) is a direct consequence of (21), since:

$$f(\mathbf{y}) - f(\mathbf{x}) - \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle = \int_0^1 \frac{1}{t} \langle \nabla f(\mathbf{y}_t) - \nabla f(\mathbf{x}), \mathbf{y}_t - \mathbf{x} \rangle dt.$$

Hence, all the statements of Theorem 4 are proved. \square

A.2 Proof of Lemma 1: The existence and uniqueness of \mathbf{x}^*

Consider the level set $\mathcal{L}_F(\mathbf{x}) := \{\mathbf{y} \in \text{dom}(F) : F(\mathbf{y}) \leq F(\mathbf{x})\}$. By (22) and the convexity of g , for any $\mathbf{y} \in \mathcal{L}_F(\mathbf{x})$ and $\mathbf{v} \in \partial g(\mathbf{x})$, we have:

$$F(\mathbf{x}) \geq F(\mathbf{y}) \geq F(\mathbf{x}) + \langle \nabla f(\mathbf{x}) + \mathbf{v}, \mathbf{y} - \mathbf{x} \rangle + \omega_*(r_{\mathbf{x}}(\mathbf{y})) \lambda_{\mathbf{x}}(\mathbf{y})^2,$$

where we use the notations $r_{\mathbf{x}}(\mathbf{y})$ and $\lambda_{\mathbf{x}}(\mathbf{y})$ in Theorem 4. Applying the Cauchy-Schwarz inequality and we obtain:

$$\omega_*(r_{\mathbf{x}}(\mathbf{y})) \lambda_{\mathbf{x}}(\mathbf{y}) \leq \|\nabla f(\mathbf{x}) + \mathbf{v}\|_{\mathbf{x}}^*.$$

Note that $\lambda_{\mathbf{x}}(\mathbf{y}) \geq \sigma_{\min} \|\mathbf{x} - \mathbf{y}\|_2$, where σ_{\min} is the smallest eigenvalue of $\nabla^2 f(\mathbf{x})$. This inequality implies that:

$$\omega_*(r_{\mathbf{x}}(\mathbf{y}))r_{\mathbf{x}}(\mathbf{y}) \leq (M_f/\sigma_{\min}) \|\nabla f(\mathbf{x}) + \mathbf{v}\|_{\mathbf{x}}^*.$$

The function $\psi(t) := t\omega_*(t)$ is increasing in $[0, +\infty)$ and its values are also in $[0, +\infty)$, so its inverse ψ^{-1} is also increasing in $[0, +\infty)$. This implies that the level set $\mathcal{L}_f(\mathbf{x})$ is bounded, and thus problem (1) has a solution \mathbf{x}^* .

Let $\mathbf{y} \neq \mathbf{x}^*$ be a point in $\text{dom}(f)$. By (22) and the optimality condition (3), for any $\mathbf{v}^* \in \partial g(\mathbf{x}^*)$, we have:

$$\begin{aligned} f(\mathbf{y}) &\geq f(\mathbf{x}^*) + \langle \nabla f(\mathbf{x}^*), \mathbf{y} - \mathbf{x}^* \rangle + \omega_*(r_{\mathbf{x}^*}(\mathbf{y}))\lambda_{\mathbf{x}^*}(\mathbf{y})^2, \\ g(\mathbf{y}) &\geq g(\mathbf{x}^*) + \langle \mathbf{v}^*, \mathbf{y} - \mathbf{x}^* \rangle \stackrel{(3)}{=} g(\mathbf{x}^*) - \langle \nabla f(\mathbf{x}^*), \mathbf{y} - \mathbf{x}^* \rangle. \end{aligned} \quad (23)$$

Summing up the two inequalities, we obtain

$$F(\mathbf{y}) \geq F(\mathbf{x}^*) + \omega_*(r_{\mathbf{x}^*}(\mathbf{y}))\lambda_{\mathbf{x}^*}(\mathbf{y})^2. \quad (24)$$

By (20) and the non-singularity of $\nabla^2 f(\mathbf{x})$ for some $\mathbf{x} \in \text{dom}(f)$, the function f is strictly convex, and the uniqueness of \mathbf{x}^* follows. \square

A.3 Proof of Lemma 2: Step-size selection strategy

Since \mathbf{s}^k is the solution of the convex subproblem (7), we have $0 \in \nabla f(\mathbf{x}^k) + \mathbf{D}_k \mathbf{d}^k + \partial g(\mathbf{s}^k)$, or, equivalently:

$$- (\nabla f(\mathbf{x}^k) + \mathbf{D}_k \mathbf{d}^k) \in \partial g(\mathbf{s}^k). \quad (25)$$

Using (25) and (22), we can derive:

$$f(\mathbf{x}^{k+1}) \leq f(\mathbf{x}^k) + \alpha_k \langle \nabla f(\mathbf{x}^k), \mathbf{d}^k \rangle + \omega(\alpha_k r_k) \alpha_k^2 \lambda_k^2. \quad (26)$$

Since $\mathbf{x}^{k+1} = (1 - \alpha_k)\mathbf{x}^k + \alpha_k \mathbf{s}^k$, it follows by the convexity of g and (25) that:

$$g(\mathbf{x}^{k+1}) \leq g(\mathbf{x}^k) - \alpha_k \langle \nabla f(\mathbf{x}^k) + \mathbf{D}_k \mathbf{d}^k, \mathbf{d}^k \rangle. \quad (27)$$

Summing up (26) and (27), we obtain the following estimate:

$$F(\mathbf{x}^{k+1}) \leq F(\mathbf{x}^k) - \psi_k(\alpha_k),$$

where $\psi_k(\tau) := \beta_k^2 \tau - \lambda_k^2 \omega(r_k \tau) \tau^2$. It is easy to check that the function ψ_k is concave, and attains the maximum at $\tau_k^* = \frac{1}{r_k} \ln \left(1 + \frac{r_k \beta_k^2}{\lambda_k^2} \right)$ with the maximum value:

$$\psi_k(\tau_k^*) = \frac{\beta_k^2}{r_k} \left[\left(1 + \frac{\lambda_k^2}{r_k \beta_k^2} \right) \ln \left(1 + \frac{\beta_k^2 r_k}{\lambda_k^2} \right) - 1 \right].$$

Moreover, $\tau_k^* \leq 1$ due to the condition $\beta_k^2 r_k \leq (e^{r_k} - 1) \lambda_k^2$. By choosing $\alpha_k = \tau_k^*$, we obtain (10). Since α_k maximizes ψ_k , it is optimal in the sense of the worst-case performance. \square

A.4 Proof of Theorem 1: Local linear convergence

Scaled proximal operators. In order to prove Theorem 1 and Theorem 2, we introduce the notion of scaled proximal operators here.

Definition 2. Let $\mathbf{H} \in \mathcal{S}_{++}^n$ be a positive definite matrix, and g be a proper, lower semi-continuous convex function. We define the operator $\mathcal{P}_{\mathbf{H}}^g(\mathbf{u})$ as:

$$\mathcal{P}_{\mathbf{H}}^g(\mathbf{u}) = (\mathbf{H} + \partial g)^{-1} := \underset{\mathbf{x}}{\operatorname{argmin}} \{g(\mathbf{x}) + (1/2)\langle \mathbf{H}\mathbf{x}, \mathbf{x} \rangle - \langle \mathbf{u}, \mathbf{x} \rangle\}.$$

We refer to $\mathcal{P}_{\mathbf{H}}^g$ as a scaled proximity operator. Note that if \mathbf{H} is the identity matrix, $\mathcal{P}_{\mathbf{H}}^g$ collapses to the standard proximal operator [16]. For $\mathbf{H} \in \mathcal{S}_{++}^n$, we define the weighted norm of \mathbf{x} as $\|\mathbf{x}\|_{\mathbf{H}} := \langle \mathbf{H}\mathbf{x}, \mathbf{x} \rangle^{1/2}$ and its dual norm of \mathbf{y} as $\|\mathbf{y}\|_{\mathbf{H}}^* := \langle \mathbf{H}^{-1}\mathbf{y}, \mathbf{y} \rangle^{1/2}$.

Lemma 5. Let g be a proper, lower semi-continuous convex function, and let \mathbf{H} be a positive definite matrix. The mapping $\mathcal{P}_{\mathbf{H}}^g$ is non-expansive in terms of the norm defined by \mathbf{H} , i.e.:

$$\|\mathcal{P}_{\mathbf{H}}^g(\mathbf{u}) - \mathcal{P}_{\mathbf{H}}^g(\mathbf{v})\|_{\mathbf{H}} \leq \|\mathbf{u} - \mathbf{v}\|_{\mathbf{H}}^*, \quad \forall \mathbf{u}, \mathbf{v}. \quad (28)$$

Proof. Let $p := \mathcal{P}_{\mathbf{H}}^g(\mathbf{u})$ and $q := \mathcal{P}_{\mathbf{H}}^g(\mathbf{v})$. We have $\mathbf{u} - \mathbf{H}p \in \partial g(p)$ and $\mathbf{v} - \mathbf{H}q \in \partial g(q)$. By the convexity of g , we have $\langle \mathbf{u} - \mathbf{v} - \mathbf{H}p + \mathbf{H}q, p - q \rangle \geq 0$. This implies that $\langle p - q, \mathbf{u} - \mathbf{v} \rangle \geq \|p - q\|_{\mathbf{H}}^2$. By the Cauchy-Schwarz inequality, we obtain $\|\mathbf{u} - \mathbf{v}\|_{\mathbf{H}}^* \geq \|p - q\|_{\mathbf{H}}$, which proves the theorem. \square

Let us consider the distance between \mathbf{x}^{k+1} and \mathbf{x}^* measured by $\|\mathbf{x}^{k+1} - \mathbf{x}^*\|_{\mathbf{x}^*}$. By the definition of \mathbf{x}^{k+1} , we have:

$$\|\mathbf{x}^{k+1} - \mathbf{x}^*\|_{\mathbf{x}^*} \leq (1 - \alpha_k) \|\mathbf{x}^k - \mathbf{x}^*\|_{\mathbf{x}^*} + \alpha_k \|\mathbf{s}^k - \mathbf{x}^*\|_{\mathbf{x}^*}. \quad (29)$$

We then derive an upper bound of $\|\mathbf{s}^k - \mathbf{x}^*\|_{\mathbf{x}^*}$ in terms of $\|\mathbf{x}^k - \mathbf{x}^*\|_{\mathbf{x}^*}$. We define $\mathcal{P}_{\mathbf{x}^*}(\mathbf{u}) := \mathcal{P}_{\mathbf{H}_*}^g(\mathbf{u})$ with $\mathbf{H}_* := \nabla^2 f(\mathbf{x}^*)$, $S_{\mathbf{x}^*}(\mathbf{u}) := \nabla^2 f(\mathbf{x}^*)\mathbf{u} - \nabla f(\mathbf{u})$, and $e_{\mathbf{x}^*}(\mathbf{u}, \mathbf{v}) := [\nabla^2 f(\mathbf{x}^*) - \mathbf{D}_k](\mathbf{v} - \mathbf{u})$. It follows from the optimality conditions (3) and (25) that $\mathbf{s}^k = \mathcal{P}_{\mathbf{x}^*}(S_{\mathbf{x}^*}(\mathbf{x}^k) + e_{\mathbf{x}^*}(\mathbf{s}^k, \mathbf{x}^k))$ and $\mathbf{x}^* = \mathcal{P}_{\mathbf{x}^*}(S_{\mathbf{x}^*}(\mathbf{x}^*))$. By Lemma 5 and the triangle inequality, we obtain:

$$\|\mathbf{s}^k - \mathbf{x}^*\|_{\mathbf{x}^*} \leq \|S_{\mathbf{x}^*}(\mathbf{x}^k) - S_{\mathbf{x}^*}(\mathbf{x}^*)\|_{\mathbf{x}^*}^* + \|e_{\mathbf{x}^*}(\mathbf{x}^k, \mathbf{s}^k)\|_{\mathbf{x}^*}^*. \quad (30)$$

Let $\tilde{r}_k := M_f \|\mathbf{x}^k - \mathbf{x}^*\|_2$, we first bound the term $\|S_{\mathbf{x}^*}(\mathbf{x}^k) - S_{\mathbf{x}^*}(\mathbf{x}^*)\|_{\mathbf{x}^*}^*$ as follows:

$$\|S_{\mathbf{x}^*}(\mathbf{x}^k) - S_{\mathbf{x}^*}(\mathbf{x}^*)\|_{\mathbf{x}^*}^* \leq \frac{e^{\tilde{r}_k} - \tilde{r}_k - 1}{\tilde{r}_k} \|\mathbf{x}^k - \mathbf{x}^*\|_{\mathbf{x}^*}. \quad (31)$$

Indeed, let us write, for notational convenience, $\mathbf{G}_k := \nabla^2 f(\mathbf{x}^* + t(\mathbf{x}^k - \mathbf{x}^*)) - \nabla^2 f(\mathbf{x}^*)$ and $\mathbf{H}_k := \nabla^2 f(\mathbf{x}^*)^{-1/2} \mathbf{G}_k \nabla^2 f(\mathbf{x}^*)^{1/2}$. By the definition of $S_{\mathbf{x}^*}$, we have:

$$S_{\mathbf{x}^*}(\mathbf{x}^k) - S_{\mathbf{x}^*}(\mathbf{x}^*) = \int_0^1 \mathbf{G}_k(\mathbf{x}^k - \mathbf{x}^*) dt. \quad (32)$$

By applying the bound (20), we get:

$$\left(\frac{1 - e^{-\tilde{r}_k}}{\tilde{r}_k} - 1 \right) \nabla^2 f(\mathbf{x}^*) \preceq G_k \preceq \left(\frac{e^{\tilde{r}_k} - 1}{\tilde{r}_k} - 1 \right) \nabla^2 f(\mathbf{x}^*),$$

which implies:

$$\|\mathbf{H}_k\| \leq \max \left\{ \frac{1 - e^{-\tilde{r}_k}}{\tilde{r}_k} - 1, \frac{e^{\tilde{r}_k} - 1}{\tilde{r}_k} - 1 \right\} = \frac{e^{\tilde{r}_k} - \tilde{r}_k - 1}{\tilde{r}_k}. \quad (33)$$

From (32), we can easily show that:

$$\|S_{\mathbf{x}^*}(\mathbf{x}^k) - S_{\mathbf{x}^*}(\mathbf{x}^*)\|_{\mathbf{x}^*}^* \leq \|\mathbf{H}_k\| \|\mathbf{x}^k - \mathbf{x}^*\|_{\mathbf{x}^*}, \quad (34)$$

which is combined with (33) to obtain (31).

Next, we bound the second term $\|e_{\mathbf{x}^*}(\mathbf{x}^k, \mathbf{s}^k)\|_{\mathbf{x}^*}^*$ of (30) as follows:

$$\|e_{\mathbf{x}^*}(\mathbf{x}^k, \mathbf{s}^k)\|_{\mathbf{x}^*}^* \leq \rho_* (\|\mathbf{s}^k - \mathbf{x}^*\|_{\mathbf{x}^*} + \|\mathbf{x}^k - \mathbf{x}^*\|_{\mathbf{x}^*}), \quad (35)$$

where $\rho_* := \max \{L_k/\sigma_{\min}^* - 1, 1 - L_k/\sigma_{\max}^*\}$. Indeed, let us define the matrix:

$$\begin{aligned} \tilde{\mathbf{H}}_* &:= \nabla^2 f(\mathbf{x}^*)^{-1/2} [\nabla^2 f(\mathbf{x}^*) - \mathbf{D}_k] \nabla^2 f(\mathbf{x}^*)^{-1/2} \\ &= \mathbb{I} - \nabla^2 f(\mathbf{x}^*)^{-1/2} \mathbf{D}_k \nabla^2 f(\mathbf{x}^*)^{-1/2}. \end{aligned}$$

Then ρ_* is the largest singular value of $\tilde{\mathbf{H}}_*$, and:

$$\|e_{\mathbf{x}^*}(\mathbf{x}^k, \mathbf{s}^k)\|_{\mathbf{x}^*}^* \leq \rho_* \|\mathbf{s}^k - \mathbf{x}^k\|_{\mathbf{x}^*} \leq \rho_* (\|\mathbf{s}^k - \mathbf{x}^*\|_{\mathbf{x}^*} + \|\mathbf{x}^k - \mathbf{x}^*\|_{\mathbf{x}^*}),$$

which proves (35).

Finally, suppose that $\rho_* \in (0, 1)$, by (29), (30), (31), and (35), we have the upper bound $\|\mathbf{x}^{k+1} - \mathbf{x}^*\|_{\mathbf{x}^*} \leq \gamma_k \|\mathbf{x}^k - \mathbf{x}^*\|_{\mathbf{x}^*}$, where:

$$\gamma_k := \left\{ 1 - \alpha_k + \alpha_k \left[\frac{e^{\tilde{r}_k} - \tilde{r}_k - 1}{(1 - \rho_*)\tilde{r}_k} + \frac{\rho_*}{1 - \rho_*} \right] \right\}.$$

Therefore, with a proper choice of L_k such that $\rho_* \in [0, 1/2)$, and for \tilde{r}_k sufficiently small such that $\gamma_k < 1$, the sequence $\{\mathbf{x}^k\}_{k \geq 0}$ generated by Algorithm 1 converges linearly to \mathbf{x}^* . \square

A.5 Proof of Theorem 2: Local quadratic convergence

Let us define $\mathcal{P}_{\mathbf{x}^k}(\mathbf{u}) := P_{\nabla^2 f(\mathbf{x}^k)}^g(\mathbf{u})$ with $\mathbf{H}_k := \nabla^2 f(\mathbf{x}^k)$, $S_{\mathbf{x}^k}(\mathbf{u}) := \nabla^2 f(\mathbf{x}^k)\mathbf{u} - \nabla f(\mathbf{u})$, and $e_{\mathbf{x}^k}(\mathbf{u}, \mathbf{v}) := [\nabla^2 f(\mathbf{x}^k) - \nabla^2 f(\mathbf{u})](\mathbf{v} - \mathbf{u})$. Since \mathbf{s}^k is the minimizer of (12), we have $0 \in \nabla f(\mathbf{x}^k) + \nabla^2 f(\mathbf{x}^k)(\mathbf{s}^k - \mathbf{x}^k) + \partial g(\mathbf{s}^k)$. Using this optimality condition and the condition $\alpha_k = 1$, we can write:

$$\begin{aligned} \mathbf{x}^{k+1} &= \mathcal{P}_{\mathbf{x}^k}(S_{\mathbf{x}^k}(\mathbf{x}^k) + e_{\mathbf{x}^k}(\mathbf{x}^k, \mathbf{s}^k)) \equiv \mathbf{s}^k \\ \mathbf{s}^{k+1} &= \mathcal{P}_{\mathbf{x}^k}(S_{\mathbf{x}^k}(\mathbf{x}^{k+1}) + e_{\mathbf{x}^k}(\mathbf{x}^{k+1}, \mathbf{s}^{k+1})). \end{aligned} \quad (36)$$

We define $\tilde{\lambda}_{k+1} := \|\mathbf{s}^{k+1} - \mathbf{x}^{k+1}\|_{\mathbf{x}^k}$. It follows by Lemma 5 and the triangle inequality that:

$$\tilde{\lambda}_{k+1} \leq \|S_{\mathbf{x}^k}(\mathbf{x}^{k+1}) - S_{\mathbf{x}^k}(\mathbf{x}^k)\|_{\mathbf{x}^k}^* + \|e_{\mathbf{x}^k}(\mathbf{s}^{k+1}, \mathbf{s}^{k+1}) - e_{\mathbf{x}^k}(\mathbf{x}^k, \mathbf{s}^k)\|_{\mathbf{x}^k}^* \quad (37)$$

We can prove, in a similar way as in the proof of (31), that the first term of (37) can be upper-bounded as:

$$\|S_{\mathbf{x}^k}(\mathbf{x}^{k+1}) - S_{\mathbf{x}^k}(\mathbf{x}^k)\|_{\mathbf{x}^k}^* \leq \frac{e^{r_k} - r_k - 1}{r_k} \|\mathbf{x}^{k+1} - \mathbf{x}^k\|_{\mathbf{x}^k}. \quad (38)$$

By using (20) and the fact that $e_{\mathbf{x}_k}(\mathbf{x}^k, \mathbf{s}^k) = 0$, the second term of (37) can be upper-bounded as:

$$\begin{aligned} \|e_{\mathbf{x}^k}(\mathbf{x}^{k+1}, \mathbf{s}^{k+1})\|_{\mathbf{x}^k}^* &\leq \max\{e^{-r_k} - 1, e^{r_k} - 1\} \|\mathbf{s}^{k+1} - \mathbf{x}^{k+1}\|_{\mathbf{x}^k} \\ &= (e^{r_k} - 1) \tilde{\lambda}_{k+1}. \end{aligned} \quad (39)$$

Combining (37), (38), and (39) and assuming that $e^{r_k} < 2$, we have:

$$\tilde{\lambda}_{k+1} \leq \frac{e^{r_k} - r_k - 1}{(2 - e^{r_k}) r_k} \lambda_k.$$

By using (20), we estimate λ_{k+1} as:

$$\lambda_{k+1}^2 := \|\mathbf{s}^{k+1} - \mathbf{x}^{k+1}\|_{\mathbf{x}^{k+1}} \leq e^{r_k} \tilde{\lambda}_{k+1}^2,$$

and thus, provided that $r_k \leq \ln(2)$, we obtain an upper bound for λ_{k+1} as:

$$\lambda_{k+1} \leq \frac{e^{r_k/2} (e^{r_k} - r_k - 1)}{r_k (2 - e^{r_k})} \lambda_k.$$

Denote by σ_{\min}^k the smallest eigenvalue of $\nabla^2 f(\mathbf{x}^k)$. We have $(\sigma_{\min}^{k+1})^{-1} \leq e^{r_k} (\sigma_{\min}^k)^{-1}$ by (20), which, combining with (40), gives us:

$$\frac{\lambda_{k+1}}{\sqrt{\sigma_{\min}^{k+1}}} \leq \frac{e^{r_k} (e^{r_k} - r_k - 1)}{r_k \sqrt{\sigma_{\min}^k} (2 - e^{r_k})} \lambda_k, \quad (40)$$

provided that $r_k \leq \ln(2)$. Finally, it is easy to check that if $r_k \leq \ln(4/3) \approx 0.28768207$, then:

$$\frac{e^{r_k} (e^{r_k} - r_k - 1)}{r_k (2 - e^{r_k})} \leq 2r_k.$$

Furthermore, we note that $\lambda_k \geq \sqrt{\sigma_{\min}^k} (r_k/M_f)$. Substituting these estimates into (40) we obtain the conclusions of Theorem 2. \square

A.6 Proof of Lemma 4: Self-concordant-like property

The concavity of f is straightforward. We now prove that f is self-concordant-like. Consider the function $\psi(t) := \log(\sum_{i=1}^m e^{\mathbf{a}_i t + \mu_i})$, where $\mathbf{a} := (\mathbf{a}_1, \dots, \mathbf{a}_m)$ is a fixed m -dimensional real vector. We define the polynomial $P(t; \mathbf{a}^k) := \mathbf{a}_1^k e^{\mathbf{a}_1 t + \mu_1} + \dots + \mathbf{a}_n^k e^{\mathbf{a}_n t + \mu_n}$. Then we have $\psi(t) = \log P(t, \mathbf{a}^0)$, and for any $k \geq 0$, $P(t; \mathbf{a}^k)'_t := \frac{dP(t; \mathbf{a}^k)}{dt} = P(t, \mathbf{a}^{k+1})$.

Applying the definition of $P(t, \mathbf{a}^k)$, it is straightforward to obtain the following expressions.

$$\psi'(t) = \frac{P(t; \mathbf{a}^0)'_t}{P(t; \mathbf{a}^0)} = \frac{P(t; \mathbf{a}^1)}{P(t; \mathbf{a}^0)}, \quad \psi''(t) = \frac{P(t; \mathbf{a}^2)P(t; \mathbf{a}^0) - P(t; \mathbf{a}^1)^2}{P(t; \mathbf{a}^0)^2},$$

and

$$\psi'''(t) = \frac{P(t; \mathbf{a}^3)P(t; \mathbf{a}^0)^2 - 3P(t; \mathbf{a}^2)P(t; \mathbf{a}^1)P(t; \mathbf{a}^0) + 2P(t; \mathbf{a}^1)^3}{P(t; \mathbf{a}^0)^3}. \quad (41)$$

Let us denote $b_i := e^{\mathbf{a}_i t + \mu_i}$. Then we can write $\psi''(t)$ as:

$$\psi''(t) := \frac{\sum_{i < j} (\mathbf{a}_i - \mathbf{a}_j)^2 b_i b_j}{(\sum_{i=1}^n b_i)^2} \geq 0,$$

and $\psi'''(t)$ as:

$$\psi'''(t) = \frac{\sum_{i < j} (\mathbf{a}_i - \mathbf{a}_j)^2 b_i b_j [\sum_{k=1}^n (\mathbf{a}_i + \mathbf{a}_j - 2\mathbf{a}_k) b_k]}{(\sum_{i=1}^n b_i)^3}. \quad (42)$$

We note that $|\mathbf{a}_i + \mathbf{a}_j - 2\mathbf{a}_k| \leq \sqrt{6} \sqrt{\mathbf{a}_i^2 + \mathbf{a}_j^2 + \mathbf{a}_k^2} \leq \sqrt{6} \|\mathbf{a}\|_2$ for $i, j, k = 1, \dots, m$, and $b_k \geq 0$ for all $k = 1, \dots, m$. Thus we have:

$$\left| \sum_{i=1}^m (\mathbf{a}_i + \mathbf{a}_j - 2\mathbf{a}_k) b_k \right| \leq \sqrt{6} \|\mathbf{a}\|_2 \sum_{i=1}^m b_i.$$

Substituting this inequality into (42) we obtain:

$$|\psi'''(t)| \leq \sqrt{6} \|\mathbf{a}\|_2 \frac{\sum_{i < j} (\mathbf{a}_i - \mathbf{a}_j)^2 b_i b_j}{(\sum_{i=1}^n b_i)^2} = \sqrt{6} \|\mathbf{a}\|_2 \psi''(t). \quad (43)$$

Let $\tilde{\mathbf{X}}$ be a matrix of size $(m+1) \times p$, formed from $\tilde{\mathbf{X}}^{(1)}, \dots, \tilde{\mathbf{X}}^{(m+1)}$. Now, we define the function $f_j(\tilde{\mathbf{X}}) := \log \left(\sum_{i=1}^{m+1} e^{\langle \mathbf{w}^{(j)}, \mathbf{X}^{(i)} \rangle} \right)$ for $j = 1, \dots, N$, and consider $\psi_j(t) := f_j(\mathbf{X} + t\mathbf{d}) = \log \left(\sum_{i=1}^{m+1} e^{\mathbf{a}_i t + \mu_i} \right)$ for given vectors \mathbf{X} and \mathbf{d} , where $\mathbf{a}_i := \langle \mathbf{w}^{(j)}, \mathbf{d}^{(i)} \rangle$ and $\mu_i := \langle \mathbf{w}^{(j)}, \mathbf{X}^{(i)} \rangle$. We note that:

$$\mathbf{a}_2 := \left(\sum_{i=1}^{m+1} \mathbf{a}_i^2 \right)^{1/2} \leq \|\mathbf{w}^{(j)}\|_2 \|\mathbf{d}\|_2. \quad (44)$$

By (43) and (44), we obtain:

$$|\psi_j'''(t)| \leq \sqrt{6} \|\mathbf{a}\|_2 \psi_j''(t) = \sqrt{6} \|\mathbf{w}^{(j)}\|_2 \|\mathbf{d}\|_2 \psi_j''(t).$$

This inequality shows that f_j is M_j -self-concordant-like, where $M_j = \sqrt{6} \|\mathbf{w}^{(j)}\|_2$. By the definition of f we have $f(\mathbf{X}) = N^{-1} \left[\mathcal{A}(\tilde{\mathbf{X}}) - \sum_{j=1}^N f_j(\tilde{\mathbf{X}}) \right]$ by setting $\tilde{\mathbf{X}}^{(i)} = \mathbf{X}^{(i)}$ for $i = 1, \dots, m$, and restricting $\tilde{\mathbf{X}}^{m+1} \equiv 0$, where \mathcal{A} is an affine operator. This implies that f is M_f -self-concordant-like with the constant $M_f := \sqrt{6} N^{-1} \max \{ \|\mathbf{w}^{(j)}\|_2 : j = 1, \dots, N \}$. \square