

ACCELERATED BREGMAN DIVERGENCE OPTIMIZATION WITH SMART: AN INFORMATION GEOMETRIC POINT OF VIEW

MAREN RAUS¹, YARA ELSHIATY¹, STEFANIA PETRA^{2,*}

¹*Institute for Mathematics, Heidelberg University, Im Neuenheimer Feld 205, 69120 Heidelberg, Germany*

²*Institute for Mathematics, Augsburg University, Universitätsstraße 14, 86159 Augsburg, Germany*

Abstract. We investigate the problem of minimizing Kullback-Leibler divergence between a linear model Ax and a positive vector b in different convex domains (positive orthant, n -dimensional box, probability simplex). Our focus is on the SMART method that employs efficient multiplicative updates. We explore the exponentiated gradient method, which can be viewed as a Bregman proximal gradient method and as a Riemannian gradient descent on the parameter manifold of a corresponding distribution of the exponential family. This dual interpretation enables us to establish connections and achieve accelerated SMART iterates while smoothly incorporating constraints. The performance of the proposed acceleration schemes is demonstrated by large-scale numerical examples.

Keywords. Bregman proximal gradient; Multiplicative algebraic reconstruction; Mirror descent; Information geometry; Optimization on manifolds; non-Lipschitz objectives.

1. INTRODUCTION

1.1. Overview and Motivation. We consider basic *SMART* (*Simultaneous Multiplicative Algebraic Reconstruction Technique*) iterations [1] for solving

$$\min_{x \in \mathbb{R}_+^n \cap \text{dom } \phi} f(x), \quad f(x) = \text{KL}(Ax, b), \quad A \in \mathbb{R}_+^{m \times n}, \quad b \in \mathbb{R}_{++}^m, \quad (1.1)$$

where the generalized *Kullback-Leibler* (*KL*) divergence (aka relative entropy and I-divergence [2]) is the Bregman divergence D_{ϕ_f} induced by the Bregman kernel ϕ_f ,

$$\text{KL}(y, y') := D_{\phi_f}(y, y') = \langle y, \log y - \log y' \rangle - \langle \mathbb{1}, y - y' \rangle, \quad y \in \mathbb{R}_+^m, \quad y' \in \mathbb{R}_{++}^m \quad (1.2a)$$

$$\phi_f(y) = \langle y, \log y \rangle - \langle \mathbb{1}, y \rangle. \quad (1.2b)$$

The KL divergence plays a distinguished role among all divergence functions [3, Section 3.4]. The SMART iteration with parameter $\tau_k > 0$ reads

$$(x_{k+1})_j = (x_k)_j \prod_{i \in [m]} \left(\frac{b_i}{(Ax_k)_i} \right)^{\tau_k A_{ij}}, \quad j \in [n], \quad x_0 \in \mathbb{R}_+^n, \quad k = 0, 1, \dots \quad (1.3)$$

*Corresponding author.

E-mail address: raus@stud.uni-heidelberg.de (M. Raus), elshiaty@math.uni-heidelberg.de (Y. Elshiaty), stefania.petra@uni-a.de (S. Petra).

Received 29 December 2023; Accepted 2 February 2024; Published online 5 March 2024.

The objective function $f(x) = \text{KL}(Ax, b)$ is motivated by inverse problems with inconsistent linear systems, where no $x \in \mathbb{R}_+^n$ satisfies $Ax = b$. Taking nonnegativity into account, the KL divergence can be employed instead of the usual least-squares approach with the squared Euclidean norm. Since the KL divergence lacks symmetry, the order of arguments matters. The reverse order $\text{KL}(b, Ax)$ is applied, for example, when measurements arise from counting discrete events (e.g., photons, electrons) subject to noise described by a Poisson process [4, 5].

Besides nonnegativity $x \in \mathbb{R}_+^n$, problem (1.1) takes also into account convex constraints $x \in \text{dom } \phi$ in a generic way, in terms of another Bregman kernel ϕ used for proximal regularization. In this paper, the domain $\text{dom } \phi$ of feasible solutions is turned into a Riemannian manifold using basic information geometry. This enables us to relate the basic multiplicative updates of the SMART iteration to a Riemannian gradient flow and also paves the way for accelerated optimization, as briefly reviewed below.

1.2. Related Work. Regarding the SMART iteration (1.6) introduced in [6], significant contributions have been made by Censor in [7, 8] and by Byrne in [1, 9]. The authors of [10] considered SMART as a *mirror descent* (MD) method and derived a convergence rate without the common L -smoothness assumption which does not hold for (1.1).

MD has a rich history in optimization and has been extensively studied. Initially proposed by [11], its basic convergence property was studied in [12] from the viewpoint of the Bregman Proximal Gradient (BPG) method. Recently, the connection to information geometry was explored in [13, 14]. The paper [13] does not cover the methods employed in the present paper, however, including the systematic use of e-geodesics. Moreover, advancements like Accelerated Bregman Proximal Gradient (ABPG) methods [15, 16] and unified algorithms for both MD and dual averaging (DA) [17] have shown superior performance in specific scenarios. Other innovations, such as new regularization methods unifying additive and multiplicative updates [18] and efforts to substantiate discrete schemes by continuous-time ODEs and numerical integration [19], have expanded the understanding of MD methods, even though some of them become computationally more expensive when the objective lacks L -smoothness, as in (1.1).

Below, MD is shown to be closely related to Riemannian gradient descent involving a particular retraction. Riemannian optimization includes foundational work by [20] which emphasizes the use of retractions and vector transport for efficient optimization. Recent advancements along this line include conjugate gradient methods in the Riemannian setting [21, 22] and approaches for optimization on manifolds with lower-bounded curvature [23]. Acceleration within the Riemannian setting is considered in [24]. Additionally, first-order methods for geodesically convex optimization [25] contribute to the broader landscape of Riemannian optimization techniques, but do not align precisely with the scenarios studied in this paper.

1.3. Contribution and Organisation. Our main contribution can be summarized as follows.

- Adopting the mirror descent point of view, we characterize SMART as Riemannian gradient descent with the exponential map induced by the e-connection of information geometry [26] as retraction.
- Our generic approach is exemplified using three different domains

$$\text{dom } \phi = \mathbb{R}_+^n, \quad \text{dom } \phi = [0, 1]^n, \quad \text{dom } \phi = \Delta_n, \quad (1.4)$$

where Δ_n denotes the probability simplex (2.2). The (relative) interior of these domains are turned into Riemannian manifolds (\mathcal{M}, g) , endowed with the Fisher-Rao geometry

g , and specifically denoted by

$$(\mathcal{P}_n, g), \quad (\mathcal{B}_n, g), \quad (\mathcal{S}_n, g). \quad (1.5)$$

The basic SMART iteration (1.3) valid for the nonnegative orthant then takes the general form

$$x_{k+1} = \frac{1}{Z(x_k)} x_k \cdot e^{-\tau_k \partial f(x_k)}, \quad x_0 \in \mathcal{M} \in \{\mathcal{P}_n, \mathcal{B}_n, \mathcal{S}_n\} \quad (\text{SMART}) \quad (1.6a)$$

where \cdot denotes componentwise vector multiplication, ∂f denotes the Euclidean gradient of the objective function f , and

$$Z(x_k) = \begin{cases} 1, & \text{if } \mathcal{M} = \mathcal{P}_n, \\ \mathbb{1} - x_k + x_k \cdot e^{-\tau_k \partial f(x_k)}, & \text{if } \mathcal{M} = \mathcal{B}_n, \\ \langle x_k, e^{-\tau_k \partial f(x_k)} \rangle, & \text{if } \mathcal{M} = \mathcal{S}_n. \end{cases} \quad (1.6b)$$

In addition, the basic ingredients for first- and second-order geometric optimization, e-geodesics and vector transport, are computed in each case which extends the conference paper [14].

- We also apply and empirically study the acceleration methods introduced by [15].
- The underlying geometry is exploited to enhance numerical techniques like (non-monotone) Riemannian line search strategies and Riemannian Conjugate Gradient (CG).

Section 2 introduces basic notation and the information geometry of statistical manifolds in order to specify the Riemannian geometry of the specific manifolds (1.5). Section 3 discusses SMART's acceleration from the viewpoint of Bregman Proximal Gradient, whereas Section 4 exploits SMART's Riemannian geometry from Section 2. This enables to design advanced Riemannian updates and to explore various line search strategies as well as Riemannian conjugate gradient. In Section 5, the results of large-scale experiments such as tomographic reconstruction, image deblurring and sparse signal recovery are reported to validate SMART's competitiveness with state-of-the-art Bregman first-order methods. On the other hand, the accelerated $\mathcal{O}(1/k^2)$ rate of the accelerated version *could not* be certified numerically, whereas the Riemannian variants showcase accelerated convergence, but at a notably higher cost per iterate along e-geodesics. Corresponding open problems for future research on optimization problems that lack L -smoothness, are outlined in Section 6.

2. PRELIMINARIES

2.1. Basic Notation. We set $[n] = \{1, 2, \dots, n\}$ for $n \in \mathbb{N}$ and write $\mathbb{1} = (1, 1, \dots, 1)^\top$ for the constant one-vector whose dimension depends on the context. $\langle \cdot, \cdot \rangle$ denotes the Euclidean inner product with induced norm $\|\cdot\| = \sqrt{\langle \cdot, \cdot \rangle}$. We also use the ℓ^1 and ℓ^∞ norm indicated by subscripts $\|\cdot\|_1$ and $\|\cdot\|_\infty$. For a vector x , the diagonal *matrix* with the components of x as entries is denoted by $\text{Diag}(x)$. Matrix norms $\|A\|, \|A\|_1, \|A\|_\infty$ always are the norms induced by the respective vector norms. The nonnegative orthant in \mathbb{R}^n is denoted by \mathbb{R}_+^n and its interior by $\mathbb{R}_{++}^n = \text{int}(\mathbb{R}_+^n)$. The relative interior of a set $S \subset \mathbb{R}^n$ contained in an affine subspace of \mathbb{R}^n is denoted by $\text{rint}(S)$.

Unary functions like the natural logarithm \log and the exponential function $e^{(\cdot)}$ apply *componentwise* to vectors, e.g.

$$\log x = (\log x_1, \dots, \log x_n), \quad e^x = (e^{x_1}, \dots, e^{x_n}). \quad (2.1a)$$

Likewise, we write *componentwise multiplication and subdivision* of two vectors as

$$z \cdot x = (z_1 x_1, \dots, z_n x_n) \quad \text{and} \quad \frac{z}{x} = \left(\frac{z_1}{x_1}, \dots, \frac{z_n}{x_n} \right) \quad \text{when} \quad x_i > 0, \forall i \in [n]. \quad (2.1b)$$

The *probability simplex* is denoted by

$$\Delta_n = \{x \in \mathbb{R}^n : x \geq 0, \langle \mathbb{1}, x \rangle = 1\} \quad (2.2)$$

and the *unit simplex* by

$$\Delta_{n-1}^0 = \{x \in \mathbb{R}^{n-1} : x \geq 0, \mathbb{1}^\top x \leq 1\}. \quad (2.3)$$

We denote the conjugate function of a function $h : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ with

$$h^*(p) = \sup_{x \in \mathbb{R}^n} \{\langle p, x \rangle - h(x)\} \quad (2.4)$$

and the gradient of a continuously differentiable function h is denoted by

$$\partial h(x) = \left(\frac{\partial}{\partial x_1} h(x), \frac{\partial}{\partial x_2} h(x), \dots \right)^\top. \quad (2.5)$$

The Hessian of the smooth function h is denoted by

$$\partial^2 h(x) \in \mathbb{R}^{n \times n}, \quad (\partial^2 h(x))_{ij} = \frac{\partial^2}{\partial x_i \partial x_j} h(x), \quad i, j \in [n]. \quad (2.6)$$

For a smooth Riemannian manifold (\mathcal{M}, g) with metric g , $T_x \mathcal{M}$ denotes the tangent space at $x \in \mathcal{M}$ and $d_x h : T_x \mathcal{M} \rightarrow \mathbb{R}$ denotes the differential of a smooth function $h : \mathcal{M} \rightarrow \mathbb{R}$. For the Riemannian metric g , we use the notation

$$g_x : T_x \mathcal{M} \times T_x \mathcal{M} \rightarrow \mathbb{R}, \quad \langle u, v \rangle_x := g_x(u, v), \quad x \in \mathcal{M} \quad (2.7a)$$

with the induced Riemannian norm

$$\|v\|_x = \sqrt{g_x(v, v)}. \quad (2.7b)$$

The Riemannian gradient $\text{grad} h(x) \in T_x \mathcal{M}$ of h is uniquely defined by

$$d_x h(v) = g_x(\text{grad} h(x), v), \quad \forall v \in T_x \mathcal{M}. \quad (2.8)$$

2.2. Bregman Divergences and Related Definitions.

Definition 2.1. (Legendre functions [27, Chap. 26]) A lower semicontinuous proper convex function $\varphi : \mathbb{R}^n \rightarrow (-\infty, \infty]$ is called

- (i) *essentially smooth* if φ is differentiable on $\text{int}(\text{dom } \varphi) \neq \emptyset$ and $\|\varphi(x_k)\| \rightarrow \infty$ for every sequence $(x_k)_{k \in \mathbb{N}} \subseteq \text{int}(\text{dom } \varphi)$ converging to a boundary point of $\text{dom } \varphi$ for $k \rightarrow \infty$;
- (ii) of *Legendre type* if φ is essentially smooth and strictly convex on $\text{int}(\text{dom } \varphi)$.

Proposition 2.1 (Legendre functions on affine subspaces [28, Prop. 5.3]). *Let*

$$\mathcal{A} = \{x \in \mathbb{R}^n : Ax = b\}, \quad A \in \mathbb{R}^{m \times n}, \quad b \in \text{Im}(A), \quad (2.9)$$

and suppose the convex function $\varphi : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ is of Legendre type. If $\text{int}(\text{dom } \varphi) \cap \mathcal{A} \neq \emptyset$, then the restriction $\varphi|_{\mathcal{A}}$ of φ to \mathcal{A} is of Legendre type and $\text{int}_{\mathcal{A}} \text{dom } \varphi|_{\mathcal{A}} = \text{int}(\text{dom } \varphi) \cap \mathcal{A}$ (where $\text{int}_{\mathcal{A}} \text{dom } \varphi|_{\mathcal{A}}$ is the interior of $\text{dom } \varphi|_{\mathcal{A}}$ in \mathcal{A} as a topological subspace of \mathbb{R}^n).

Theorem 2.1. [27, Thm 26.5] *A convex function φ is of Legendre type if and only if its conjugate φ^* is. In this case, the gradient mapping $\partial\varphi : \text{int}(\text{dom } \varphi) \mapsto \text{int}(\text{dom } \varphi^*)$ is one-to-one and*

$$(\partial\varphi)^{-1} = \partial\varphi^*, \quad (2.10a)$$

$$\varphi^*(\partial\varphi(x)) = \langle x, \partial\varphi(x) \rangle - \varphi(x), \quad \forall x \in \text{int}(\text{dom } \varphi). \quad (2.10b)$$

A given Legendre function φ induces a corresponding *Bregman divergence*,

$$D_{\varphi}(x, y) = \varphi(x) - \varphi(y) - \langle \partial\varphi(y), x - y \rangle, \quad (2.11)$$

as studied, e.g., by [29, 30].

Lemma 2.1 (Bregman projection [31, Thm. 3.12]). *Suppose φ is closed proper convex and differentiable on $\text{int}(\text{dom } \varphi)$, C is closed convex with $C \cap \text{int}(\text{dom } \varphi) \neq \emptyset$, and $y_0 \in \text{int}(\text{dom } \varphi)$. If φ is Legendre, then the Bregman projection \tilde{y} of y_0 defined by*

$$\underset{y \in C \cap \text{dom } \varphi}{\text{argmin}} D_{\varphi}(y, y_0) = \{\tilde{y}\}, \quad \tilde{y} \in \text{int}(\text{dom } \varphi). \quad (2.12)$$

is unique and contained in $\text{int}(\text{dom } \varphi)$.

Lemma 2.2 (Three-points identity [32, Lem. 3.1]). *Let φ closed proper convex and differentiable on $\text{int}(\text{dom } \varphi)$. For any $r \in \text{dom } \varphi$, and $p, q \in \text{int}(\text{dom } \varphi)$ the following identity holds*

$$D_{\varphi}(r, q) + D_{\varphi}(q, p) - D_{\varphi}(r, p) = \langle r - q, \partial\varphi(p) - \partial\varphi(q) \rangle. \quad (2.13)$$

2.3. Fisher-Rao Geometry. We summarize basic concepts of the information geometry of regular probability distributions of the exponential family, given by densities of the form

$$p(y; \theta) = h(y) e^{\langle \theta, t(y) \rangle - \psi(\theta)}, \quad \theta \in \Theta, \quad (2.14)$$

where h is the base measure, t is the sufficient statistic, and ψ is the *log-partition function*

$$\psi(\theta) = \log Z(\theta) \quad (2.15a)$$

with the normalizing constant

$$Z(\theta) = \int h(y) e^{\langle \theta, t(y) \rangle} dy, \quad (2.15b)$$

and Θ parameter space. Basic references include [26, 33].

The density p is *regular* if there are no dependencies among the functions forming the vector $y \mapsto (1, t_1(y), t_2(y), \dots)$. As a consequence, the domain Θ of the *exponential* (a.k.a. *natural, canonical*) *parameters* θ is a open convex set with nonempty interior $\overset{\circ}{\Theta} \neq \emptyset$. Moreover, the log-partition function $\psi(\theta)$ is lower-semicontinuous, essentially smooth and convex such that the classical Legendre transformation applies (cf. [27, Section 26]).

An alternative parametrization of any regular density p of the exponential family is provided by the so-called *mean-parametrization*

$$\eta = \eta(\theta) = \partial\psi(\theta), \quad (2.16a)$$

$$\mathcal{M} := \partial\psi(\Theta), \quad (2.16b)$$

where ‘mean’ refers to the expected value of the sufficient statistics, which defines η by

$$\eta \stackrel{(2.15)}{=} \frac{1}{Z(\theta)} \int t(y)p(y; \theta)dy = \mathbb{E}_p[t(Y)]. \quad (2.17)$$

Denoting the conjugate function of ψ by

$$\phi(\eta) = \psi^*(\eta) \quad (2.18)$$

and assuming η and θ are related by (2.16a), one has by the classical Legendre transform

$$\theta = \theta(\eta) = \partial\phi(\eta), \quad (2.19a)$$

$$(\partial^2\psi(\theta))^{-1} = \partial^2\phi(\eta). \quad (2.19b)$$

The Fisher-Rao metric is given by

$$H(\theta) \in \mathbb{R}^{n \times n}, \quad H_{ij}(\theta) = \mathbb{E}_p \left[\left(\frac{\partial}{\partial \theta_i} \log p(Y; \theta) \right) \left(\frac{\partial}{\partial \theta_j} \log p(Y; \theta) \right) \right] = (\partial^2\psi(\theta))_{ij}. \quad (2.20)$$

Under the coordinate change (2.16a), it covariantly transforms to (cf. [34, Section 1.4])

$$H_{ij}(\theta) = \sum_{k,l} G_{kl}(\eta) \left(\frac{\partial}{\partial \theta_i} (\partial\psi(\theta))_k \right) \left(\frac{\partial}{\partial \theta_j} (\partial\psi(\theta))_l \right) = \left((\partial^2\psi(\theta)) G(\eta) (\partial^2\psi(\theta)) \right)_{ij}. \quad (2.21)$$

Hence, taking into account (2.19b) and (2.20), one has

$$G(\eta) = \partial^2\phi(\eta). \quad (2.22)$$

As a result, an exponential family of densities p of the form (2.14) may be identified with the Riemannian manifolds

$$(\Theta, h) \quad \text{and} \quad (\mathcal{M}, g), \quad (2.23)$$

with metrics h and g given by the metric tensors $H(\theta)$ and $G(\eta)$ as defined by (2.20) and (2.22), respectively.

2.4. Specific Manifolds. We will consider the three sets

$$\mathcal{P}_n = \mathbb{R}_{++}^n, \quad \mathcal{B}_n = (0, 1)^n, \quad \mathcal{S}_n = \text{rint}(\Delta_n) = \Delta_n \cap \mathbb{R}_{++}^n \quad (2.24)$$

and regard them as instances of a Riemannian manifold \mathcal{M} , respectively, as specified in the subsequent sections.

2.4.1. *Positive Orthant.* The Poisson distribution of a discrete random variable $Y \in \mathbb{N}_0$ is given by the probability mass function

$$p(y; \eta) = \frac{\eta^y e^{-\eta}}{y!}, \quad \eta \in (0, \infty). \quad (2.25)$$

Rewriting this as density of the exponential family gives

$$p(y; \theta) = \frac{1}{y!} e^{\theta y - \psi(\theta)}, \quad \theta = \theta(\eta) = \log \eta, \quad \psi(\theta) = e^\theta. \quad (2.26)$$

We note that denoting the parameter of (2.25) with the symbol η used for the mean parameter is consistent since $\mathbb{E}_p[Y] = \psi'(\theta) = e^\theta = \eta$. Computing the conjugate function (2.18) yields

$$\phi(\eta) = \eta \log \eta - \eta \quad (2.27)$$

and $\phi''(\eta) = \frac{1}{\eta}$.

In the n -dimensional case, one has the density in product form

$$p(y; \eta) = \prod_{i \in [n]} \frac{\eta_i^{y_i} e^{-\eta_i}}{y_i!}, \quad \eta \in \mathcal{P}_n = \mathbb{R}_{++}^n, \quad (2.28)$$

and it is straightforward to check that one obtains the Riemannian manifold

$$(\mathcal{P}_n, g), \quad G(\eta) = \text{Diag} \left(\frac{\mathbb{1}_n}{\eta} \right) \quad (2.29)$$

with metric g and corresponding metric tensor $G(\eta)$ given by (2.22), yielding the diagonal matrix (2.29).

2.4.2. *Unit Box.* The Bernoulli distribution of random variable $Y \in \{0, 1\}$ reads

$$p(y; \eta) = \eta^y (1 - \eta)^{1-y}, \quad \eta \in (0, 1). \quad (2.30)$$

Rewriting this in exponential form gives

$$p(y; \theta) = e^{\theta y - \psi(\theta)}, \quad \theta = \log \frac{\eta}{1 - \eta}, \quad \psi(\theta) = \log(1 + e^\theta) \quad (2.31)$$

with conjugation function

$$\phi(\eta) = \psi^*(\eta) = \eta \log \eta + (1 - \eta) \log(1 - \eta) \quad (2.32)$$

and $\phi''(\eta) = \frac{1}{\eta(1-\eta)}$. Thus, in the n -dimensional case $\mathcal{B}_n = (0, 1)^n$, we obtain by (2.22) the Riemannian manifold

$$(\mathcal{B}_n, g), \quad G(\eta) = \text{Diag} \left(\frac{\mathbb{1}_n}{\eta \cdot (\mathbb{1}_n - \eta)} \right), \quad (2.33)$$

with a metric g and the corresponding metric tensor $G(\eta)$ given by (2.33), which is a diagonal matrix.

2.4.3. *Probability Simplex.* Every point $p \in \mathcal{S}_n$ is a discrete (a.k.a. categorical) distribution governing a random variable $Y \in [n]$ by $p_i = \Pr(Y = i)$, $i \in [n]$. Introducing local coordinates

$$\vartheta: \mathcal{S}_n \rightarrow \mathbb{R}_{++}^{n-1}, \quad \eta := \vartheta(p) = (p_1, \dots, p_{n-1})^\top, \quad \eta_i = \mathbb{E}_p[Y = i], \quad i \in [n-1] \quad (2.34)$$

and

$$\vartheta(\mathcal{S}_n) = \{\eta \in \mathbb{R}_{++}^{n-1} : \langle \mathbb{1}_{n-1}, \eta \rangle < 1\} \quad (2.35)$$

one has

$$p = \vartheta^{-1}(\eta) = \left(\begin{array}{c} \eta \\ 1 - \langle \mathbb{1}_{n-1}, \eta \rangle \end{array} \right). \quad (2.36)$$

Note that

$$\vartheta(\mathcal{S}_n) = \text{int}(\Delta_{n-1}^0). \quad (2.37)$$

Defining the mapping

$$\delta: [n] \rightarrow \{0, 1\}^{n-1}, \quad \delta_i(y) := \begin{cases} 1, & \text{if } y = i, \\ 0, & \text{otherwise,} \end{cases} \quad i \in [n-1], \quad (2.38)$$

a minimal parameterization of the categorical distribution of a random variable $Y \in [n]$ reads

$$p(y; \eta) = \prod_{i=1}^{n-1} p_i^{\delta_i(y)} p_n^{1 - \langle \mathbb{1}_{n-1}, \delta(y) \rangle}. \quad (2.39)$$

Rewriting this as distribution of the exponential family [33] gives

$$p(y; \theta) = e^{\langle \theta, \delta(y) \rangle - \psi(\theta)}, \quad \psi(\theta) = \log(1 + \langle \mathbb{1}_{n-1}, e^\theta \rangle) \quad (2.40)$$

with the natural parameters

$$\theta_i := \log \frac{\eta_i}{1 - \langle \mathbb{1}_{n-1}, \eta \rangle}, \quad i \in [n-1]. \quad (2.41)$$

Computing conjugate function (2.18), we obtain

$$\phi(\eta) = \psi^*(\eta) = \langle \eta, \log \eta \rangle + (1 - \langle \mathbb{1}_{n-1}, \eta \rangle) \log(1 - \langle \mathbb{1}_{n-1}, \eta \rangle). \quad (2.42)$$

Using (2.22) yields the metric tensor

$$G(\eta) = \text{Diag} \left(\frac{\mathbb{1}_{n-1}}{\eta} \right) + \frac{1}{1 - \langle \mathbb{1}_{n-1}, \eta \rangle} \mathbb{1}_{n-1} \mathbb{1}_{n-1}^\top \quad (2.43)$$

and in turn the Riemannian manifold (cf. (2.37))

$$(\vartheta(\mathcal{S}_n), g) \quad (2.44)$$

with the Fisher-Rao metric g given in coordinates by (2.43).

It will be convenient to express (2.44) in terms of the ambient coordinates $p = \vartheta^{-1}(\eta)$, see (2.36).

Lemma 2.3 (probability simplex, Fisher-Rao metric). *The manifold (2.44) is isomorphic to*

$$(\mathcal{S}_n, g) \quad (2.45)$$

where the Fisher-Rao metric is given in ambient coordinates by

$$G(p) = \text{Diag} \left(\frac{\mathbb{1}_n}{p} \right), \quad p \in \mathcal{S}_n. \quad (2.46)$$

Proof. The constraint $\langle \mathbb{1}_n, p \rangle = 1$ yields for any smooth curve $t \mapsto p(t) \in \mathcal{S}_n$ the condition $\langle \mathbb{1}_n, \dot{p} \rangle = 0$ for tangent vectors, that is the tangent space

$$T_0 := \{v \in \mathbb{R}^n : \langle \mathbb{1}_n, v \rangle = 0\} \quad (2.47)$$

and the trivial tangent bundle $T\mathcal{S}_n = \mathcal{S}_n \times T_0$.

A basis of T_0 is given by $d\vartheta^{-1}(e_i) = e_i - e_n$, $i \in [n-1]$, where e_1, \dots, e_n denote the canonical basis vectors of the ambient space \mathbb{R}^n . Collecting this basis of T_0 as column vectors of a matrix yields the coordinate representation of tangent vectors

$$v = B_0 v_\eta, \quad B_0 := (e_1 - e_n, \dots, e_{n-1} - e_n) = \begin{pmatrix} I_{n-1} \\ -\mathbb{1}_{n-1}^\top \end{pmatrix} \in \mathbb{R}^{n \times (n-1)} \quad (2.48)$$

and the equation

$$g_\eta(u_\eta, v_\eta) = \langle u_\eta, G(\eta)v_\eta \rangle = \langle u_\eta, B_0^\top G(p)B_0 v_\eta \rangle = \langle u, G(p)v \rangle = g_p(u, v), \quad u, v \in T_0 \quad (2.49a)$$

with

$$G(p) = \text{Diag} \left(\frac{\mathbb{1}_n}{p} \right), \quad p \in \mathcal{S}_n. \quad (2.49b)$$

□

We conclude this section by deriving the expression for Riemannian gradients [34, p. 89] of functions specified in *ambient* coordinates.

Lemma 2.4 (Riemannian gradient in ambient coordinates). *Let $f: \mathcal{S}_n \rightarrow \mathbb{R}$ be a smooth function specified in ambient coordinates $f(p)$ with gradient $\partial f(p) = (\partial_{p_1} f(p), \dots, \partial_{p_n} f(p))$. Then the Riemannian gradient of f with respect to the Fisher-Rao metric reads*

$$\text{grad} f(p) = \Pi_p \partial f(p), \quad p \in \mathcal{S}_n \quad (2.50a)$$

with

$$\Pi_p := \text{Diag}(p) - pp^\top. \quad (2.50b)$$

Proof. Writing $f(p) = f(\vartheta^{-1}(\eta))$ by using (2.36), the chain rule gives

$$\partial_\eta f(p) = B_0^\top \partial f(p), \quad p = \vartheta^{-1}(\eta). \quad (2.51)$$

A straightforward computation yields the inverse of the metric tensor (2.43)

$$G(\eta)^{-1} = \text{Diag}(\eta) - \eta\eta^\top := \Pi_\eta, \quad (2.52)$$

that we denote by Π_η and, in turn, the Riemannian gradient in *local* coordinates

$$\text{grad}_\eta f(p) = \Pi_\eta B_0^\top \partial f(p), \quad p = \vartheta^{-1}(\eta), \quad (2.53)$$

which by definition is a tangent vector in local coordinates. Adopting the coordinate representation (2.48) finally yields

$$\text{grad} f(p) = B_0 \Pi_\eta B_0^\top \partial f(p) = \Pi_p \partial f(p) \quad (2.54)$$

with Π_p given by (2.50b). □

2.4.4. *Summary.* The following tables summarize the basic expressions derived above. The relation of the last two rows of Table 2 will be examined in Remark 2.1.

TABLE 1. Elementary Riemannian manifolds \mathcal{M} , the corresponding distributions $p(y)$, metric tensors G and inverse metric tensors G^{-1} . For the probability simplex \mathcal{S}_n , working with the ambient coordinates is a convenient alternative (last row).

\mathcal{M}	$p(y)$	G	G^{-1}
\mathcal{P}_n	$\prod_{i \in [n]} \frac{\eta_i^{y_i} e^{-\eta_i}}{y_i!}$	(2.29)	$\text{Diag}(\eta)$
\mathcal{B}_n	$\prod_{i \in [n]} \eta_i^{y_i} (1 - \eta_i)^{1 - y_i}$	(2.33)	$\text{Diag}(\eta \cdot (\mathbb{1} - \eta))$
$\mathfrak{D}(\mathcal{S}_n)$	$\prod_{i=1}^{n-1} \eta_i^{\delta_i(y)} (1 - \langle \mathbb{1}_{n-1}, \eta \rangle)^{1 - \langle \mathbb{1}_{n-1}, \delta(y) \rangle}$	(2.43)	(2.52)
\mathcal{S}_n	$\prod_{i \in [n-1]} p_i^{\delta_i(y)} p_n^{1 - \langle \mathbb{1}_{n-1}, \delta(y) \rangle}$	(2.49b)	(2.50b)

TABLE 2. The conjugate potentials ψ, ϕ corresponding to the elementary Riemannian manifolds \mathcal{M} of Table 1. Using the convention $0 \log 0 = 0$, the domain of each ϕ is extended to the closure $\overline{\text{dom } \phi}$.

\mathcal{M}	ψ	$\text{dom } \psi$	$\phi = \psi^*$	$\overline{\text{dom } \phi}$
\mathcal{P}_n	$\langle \mathbb{1}_n, e^\theta \rangle$	\mathbb{R}^n	$\langle \eta, \log \eta \rangle - \langle \mathbb{1}_n, \eta \rangle$	\mathbb{R}_+^n
\mathcal{B}_n	$\sum_{i \in [n]} \log(1 + e^{\theta_i})$	\mathbb{R}^n	$\langle \eta, \log \eta \rangle + \langle \mathbb{1}_{n-1} - \eta, \log(\mathbb{1}_{n-1} - \eta) \rangle$	$[0, 1]^n$
$\mathfrak{D}(\mathcal{S}_n)$	$\log(1 + \langle \mathbb{1}_{n-1}, e^\theta \rangle)$	\mathbb{R}^{n-1}	$\langle \eta, \log \eta \rangle + (1 - \langle \mathbb{1}_{n-1}, \eta \rangle) \log(1 - \langle \mathbb{1}_{n-1}, \eta \rangle)$	Δ_{n-1}^0
\mathcal{S}_n	$\frac{e^q}{\langle \mathbb{1}_n, e^q \rangle}$	T_0 (2.47)	$\langle p, \log p \rangle$	Δ_n

Remark 2.1 (Bregman kernels). Each function ϕ listed in Table 2 is Legendre. This includes in particular the correspondence $\phi = \psi^*$ in the last row due to [35, Example 11.12], with the lack of *strict* convexity of the log-exponential function $\psi = \frac{e^q}{\langle \mathbb{1}_n, e^q \rangle}$ removed by restricting ψ to the tangent space T_0 . Proposition 2.1 with the affine subspace $\{x \in \mathbb{R}^n : \langle \mathbb{1}_n, x \rangle = 1\}$ and $\phi = \phi$, where $\text{dom } \phi = \mathbb{R}_+^n$, yields the assertion.

Accordingly, we denote the functions listed in Table 2 by

$$\phi = \begin{cases} \phi_+ & \text{if } \mathcal{M} = \mathcal{P}_n, \\ \phi_\square & \text{if } \mathcal{M} = \mathcal{B}_n, \\ \phi_\Delta^0 & \text{if } \mathcal{M} = \mathfrak{D}(\mathcal{S}_n), \\ \phi_\Delta & \text{if } \mathcal{M} = \mathcal{S}_n. \end{cases} \quad (2.55)$$

We compute directly using (2.11)

$$D_{\phi_+}(x, y) = \langle x, \log x - \log y \rangle - \langle \mathbb{1}, x - y \rangle =: \text{KL}(x, y), \quad (2.56a)$$

and

$$D_{\phi_{\square}}(x, y) = D_{\phi_{+}}(x, y) + D_{\phi_{+}}(\mathbb{1}_n - x, \mathbb{1}_n - y), \quad (2.56b)$$

$$D_{\phi_{\Delta}^0}(\eta, \eta') = D_{\phi_{\Delta}}(p, p'), \quad (\text{Lemma 2.5}) \quad (2.56c)$$

$$D_{\phi_{\Delta}}(p, p') = \left\langle p, \log \frac{p}{p'} \right\rangle. \quad (2.56d)$$

It remains to establish Equation (2.56c).

Lemma 2.5 (Bregman divergence using ambient coordinates). *One has*

$$D_{\phi_{\Delta}^0}(\eta, \eta') = D_{\phi_{\Delta}}(p, p') = \text{KL}(p, p') = \left\langle p, \log \frac{p}{p'} \right\rangle, \quad \eta = \vartheta(p), \eta' = \vartheta(p'), \quad \forall p, p' \in \mathcal{S}_n. \quad (2.57)$$

Proof. Setting $\phi = \phi_{\Delta}^0$, we directly compute using (2.36)

$$\phi(\eta) = \langle p, \log p \rangle|_{p=\vartheta^{-1}(\eta)} \quad (2.58a)$$

$$= \langle \eta, \log \eta \rangle + (1 - \langle \mathbb{1}_{n-1}, \eta \rangle) \log(1 - \langle \mathbb{1}_{n-1}, \eta \rangle), \quad (2.58b)$$

$$\partial \phi(\eta) = B^{\top} \log p|_{p=\vartheta^{-1}(\eta)} = \log \eta - \log(1 - \langle \mathbb{1}_{n-1}, \eta \rangle) \mathbb{1}_{n-1} \quad (2.58c)$$

Hence

$$D_{\phi}(\eta, \eta') = \phi(\eta) - \phi(\eta') - \langle \partial \phi(\eta'), \eta - \eta' \rangle \quad (2.58d)$$

$$= \langle \eta, \log \eta \rangle + p_n \log p_n - (\langle \eta', \log \eta' \rangle + p'_n \log p'_n) \quad (2.58e)$$

$$- \langle \log \eta' - \log p'_n \mathbb{1}_{n-1}, \eta - \eta' \rangle \quad (2.58f)$$

$$= \langle \eta, \log \eta \rangle + p_n \log p_n - p'_n \log p'_n - \langle \eta, \log \eta' \rangle + \log p'_n \langle \mathbb{1}_{n-1}, \eta - \eta' \rangle \quad (2.58g)$$

$$= \langle p, \log p \rangle - p'_n \log p'_n - \langle \eta, \log \eta' \rangle + \log p'_n \langle \mathbb{1}_{n-1}, \eta - \eta' \rangle + \underbrace{\log p'_n - \log p'_n}_{=0} \quad (2.58h)$$

$$= \langle p, \log p \rangle - \langle \eta, \log \eta' \rangle - \underbrace{(p'_n + \langle \mathbb{1}_{n-1}, \eta' \rangle - 1)}_{=0} \log p'_n - \underbrace{(1 - \langle \mathbb{1}_{n-1}, \eta \rangle)}_{=p_n} \log p'_n \quad (2.58i)$$

$$= \langle p, \log p - \log p' \rangle = \text{KL}(p, p'). \quad (2.58j)$$

□

3. SMART: CONVERGENCE AND ACCELERATION

In this section, we adopt the viewpoint of the Bregman Proximal Gradient (BPG) (Section 3.1) to understand the convergence rate of SMART (Section 3.2). Furthermore, we detail various accelerated variants of the BPG iteration (Section 3.3) which are supposed to accelerate the BPG iteration at minimal cost.

3.1. Bregman Proximal Gradient (BPG). The classical *mirror descent (MD)* method [11] for solving the convex optimization problem

$$\min_{x \in C_F} F(x), \quad \text{dom } F =: C_F \subset \mathbb{R}^n, \quad (3.1)$$

with F continuously differentiable convex and C_F nonempty closed convex, was investigated by Beck and Teboulle [12] in more general form after replacing the squared Euclidean norm by general distance-like functions as proximal term. Specifically, using a Bregman divergence (2.11) as distance function, the update scheme with stepsize $\tau_k > 0$ reads

$$x_{k+1} = \underset{x \in C_\varphi}{\operatorname{argmin}} \langle \partial F(x_k), x \rangle + \frac{1}{\tau_k} D_\varphi(x, x_k), \quad k = 0, 1, \dots, \quad (3.2a)$$

$$=: \operatorname{Mirr}_\varphi(\tau_k \partial F(x_k)), \quad x_0 \in C_\varphi := C_F \cap \operatorname{int}(\operatorname{dom} \varphi). \quad (3.2b)$$

This may also be seen as a linearization of the classical proximal point iteration [36] and using a Bregman divergence D_φ as proximity measure instead of the squared Euclidean distance and, in this sense, as a linearization of the approach of Censor and Zenios [37]. Rewriting (3.2) in the form (cf., e.g., [12])

$$\operatorname{Mirr}_\varphi(\tau_k \partial F(x_k)) = \arg \min_{x \in C_\varphi} D_\varphi(x, x_{k+1}) \quad (3.3a)$$

$$= x_{k+1} = \partial \varphi^*(\partial \varphi(x_k) - \tau_k \partial F(x_k)) \in C_\varphi \quad (3.3b)$$

reveals that the *mirror descent map* $\operatorname{Mirr}_\varphi$ equals the *interior Bregman gradient step* studied, e.g., in [38]. Evaluating (3.3) yields the specific update expressions (1.6).

3.2. Convergence Analysis. In this section, we establish the convergence rate $\mathcal{O}(\frac{1}{k})$ of the SMART iteration for solving problem (1.1), by recognizing it as special case of the BPG scheme (3.2). Since the objective function has a gradient which is *not* L -Lipschitz, the following weaker notion is instrumental.

Definition 3.1 (relative L -smoothness). [39, Prop. 1] A convex function F is called *L -smooth relative* to the Legendre function φ if $\operatorname{dom} \varphi \subset \operatorname{dom} F$ and there exists a constant $L > 0$ such that

$$D_F(x, y) \leq L D_\varphi(x, y), \quad \forall (x, y) \in \operatorname{dom} \varphi \times \operatorname{int}(\operatorname{dom} \varphi). \quad (3.4)$$

We adopt the following

Assumptions A:

- (i) the Bregman kernel φ is of Legendre type;
- (ii) the objective function $F : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ is a convex function and continuously differentiable on $\operatorname{int} \operatorname{dom} \varphi$, with $\operatorname{dom} \varphi \subset \operatorname{dom} F$;
- (iii) the objective is bounded from below, i.e. $-\infty < \inf_{x \in \operatorname{dom} \varphi} F(x) =: F_*$;
- (iv) the objective F is L -smooth relative to φ .

Remark 3.1 (well posedness). Assumption A imply that the mirror descent map (3.2) is a well-defined, single-valued map $\operatorname{Mirr}_\varphi : \operatorname{int}(\operatorname{dom} \varphi) \rightarrow \operatorname{int}(\operatorname{dom} \varphi)$ (Lemma 2.1).

Now we consider the problem (1.1),

$$f(x) = \operatorname{KL}(Ax, b) =: D_{\phi_f}(Ax, b) \quad (3.5a)$$

with regularizing proximal term

$$D_\phi(x, y) \tag{3.5b}$$

given by (2.56) and ϕ given by any function listed in (2.55). In each case, (f, ϕ) plays the role of (F, φ) in the Assumptions A.

Proposition 3.1 (Existence). *Consider ϕ_f in (1.2b), f defined by (3.5a), with A and b as in (1.1) and ϕ given by any function listed in (2.55). Set $K = \{Ax: x \in \text{dom } \phi \subset \mathbb{R}_+^n\} \subset \mathbb{R}_+^m$ and assume $b \in \text{int}(K)$ and $\text{dom } \phi_f \cap K \neq \emptyset$. Then*

$$\underset{\mathbb{R}_+^n \cap \text{dom } \phi}{\text{argmin}} f(x) \neq \emptyset. \tag{3.6}$$

Proof. Problem (3.6) satisfies the assumptions of Lemma 2.1 with the specified K and $\varphi = \phi_f$. \square

We continue with Lemmata preparing the main result, Theorem 3.1.

Lemma 3.1. [14, Lemma 1] *Let f be defined by (3.5a) and consider ϕ_f in (1.2b). Then*

$$D_f(x, y) = D_{\phi_f}(Ax, Ay) = \text{KL}(Ax, Ay). \tag{3.7}$$

Lemma 3.2. *Suppose $A \in \mathbb{R}_+^{m \times n}$ and consider ϕ_+ from (2.55). Then*

$$D_{\phi_f}(Ax, Ay) \leq \|A\|_1 D_{\phi_+}(x, y), \quad \forall x, y \in \mathbb{R}_+^n. \tag{3.8}$$

Proof. We apply the log-sum inequality [40, Thm. 2.7.1]

$$\forall a, b \in \mathbb{R}_+^n, \quad \left\langle a, \log \frac{a}{b} \right\rangle \geq \langle \mathbb{1}, a \rangle \log \frac{\langle \mathbb{1}, a \rangle}{\langle \mathbb{1}, b \rangle}, \tag{3.9}$$

which naturally extends by continuity to the boundary of \mathbb{R}_+^n . By assumption, $\|A\|_1 \geq \sum_{i \in [m]} A_{ij}$, $j \in [n]$. We compute

$$\|A\|_1 D_{\phi_+}(x, y) \geq \sum_{j \in [n]} \left(\sum_{i \in [m]} A_{ij} \right) \left(x_j \log \frac{x_j}{y_j} - x_j + y_j \right) \tag{3.10a}$$

$$= \sum_{i \in [m]} \left(\sum_{j \in [n]} A_{ij} x_j \log \frac{A_{ij} x_j}{A_{ij} y_j} - \sum_{j \in [n]} A_{ij} (x_j - y_j) \right) \tag{3.10b}$$

$$\stackrel{(3.9)}{\geq} \sum_{i \in [m]} \left[\left(\sum_{j \in [n]} A_{ij} x_j \right) \log \frac{\sum_{j \in [n]} A_{ij} x_j}{\sum_{j \in [n]} A_{ij} y_j} - \sum_{j \in [n]} A_{ij} (x_j - y_j) \right] \tag{3.10c}$$

$$= \left\langle (Ax) \log \frac{Ax}{Ay} \right\rangle - \langle \mathbb{1}_m, Ax - Ay \rangle \tag{3.10d}$$

$$= D_{\phi_f}(Ax, Ay). \tag{3.10e}$$

\square

Next, we check the relative L -smoothness (Definition 3.1) of the SMART objective function for the three relevant scenarios.

Proposition 3.2. *Let $f(x) = D_{\phi_f}(Ax, b)$. Then f is $\|A\|_1$ -smooth relative to any $\phi \in \{\phi_+, \phi_\square, \phi_\Delta\}$ specified by (2.55) and Table 2.*

Proof. We check each case.

- The case ϕ_+ results from Lemma 3.1 and Lemma 3.2.
- The case ϕ_\square follows from the explicit expression of D_{ϕ_\square} given by (2.56b), from the case ϕ_+ just established, and from the nonnegativity $D_{\phi_+}(\mathbb{1}_n - x, \mathbb{1}_n - y) \geq 0$ for any $x, y \in \mathcal{B}_n$.
- The case ϕ_Δ follows from $\Delta_n \subset \mathbb{R}_+^n$ and the case ϕ_+ .

□

The following main result in this section draws upon [1, Thm. 2], [41, Thm. 3.1], and [10, Thm. 2]. These references utilize properties of Legendre functions, specifically Prop. 3.2 and Lem. 2.2, alongside the optimality criteria of (3.3) and the convexity of f . Compare also [Thm.1, Thm.2 (i)][39].

Theorem 3.1 (convergence rate). *Let X_* be the solution set of (1.1) with f and $\phi \in \{\phi_+, \phi_\square, \phi_\Delta\}$ as in Proposition 3.2 and $L = \|A\|_1$. For $(x_k)_{k \in \mathbb{N}}$ generated by (1.6), with starting point $x_0 \in \text{rint}(\text{dom } \phi)$ and $\tau_k = \tau \leq 1/L$, one has the following.*

- (i) *The sequence $(x_k)_{k \in \mathbb{N}}$ converges to a unique point in $x_* \in X_*$,*

$$x_* := \arg \min_{x \in X_*} D_\phi(x, x_0). \quad (3.11)$$

- (ii) *For every k ,*

$$f(x_k) - f(x_*) \leq \frac{LD_\phi(x_*, x_0)}{k}. \quad (3.12)$$

3.3. Accelerated Bregman Proximal Gradient (ABPG). This section briefly reviews recent related work (Section 3.3.1) and corresponding algorithms (Section 3.3.2), which served as baselines for the experimental evaluation reported in Section 5.

3.3.1. Acceleration and Related Work. In [42], the optimality of the $\mathcal{O}(1/k)$ rate for a wide range of Bregman proximal gradient (BPG) algorithms was established, under *general* assumptions regarding the objective function F and the Bregman function ϕ . In particular, L -smoothness of the objective function F relative to the Bregman function ϕ is merely required, rather than L -Lipschitz continuity of the gradient ∂F .

The derivation of *accelerated* Bregman proximal gradient (ABPG) algorithms, on the other hand, require additional assumptions. The authors of [15] consider assumptions which implies a $\mathcal{O}(1/k^\gamma)$ rate with $\gamma \in [1, 2]$:

- The *triangle-scaling property* with *uniform triangle-scaling exponent (TSE)* γ ,

$$D_\phi((1-\theta)x + \theta z, (1-\theta)x + \theta \tilde{z}) \leq \theta^\gamma D_\phi(z, \tilde{z}), \quad \forall \theta \in [0, 1], \quad (\text{uniform TSE}) \quad (3.13a)$$

$$\forall x, z, \tilde{z} \in \text{int dom } \phi. \quad (3.13b)$$

Major examples are provided by *jointly* convex Bregman divergences D_ϕ which satisfy the inequality with TSE $\gamma = 1$. In particular, the KL-divergence is jointly convex and each of the related Bregman divergences (2.56).

- The *intrinsic* TSE defined as

$$\gamma_{\text{in}} := \limsup_{\theta \searrow 0} \frac{D_\phi((1-\theta)x + \theta z, (1-\theta)x + \theta \tilde{z})}{\theta^\gamma} < \infty, \quad (\text{intrinsic TSE}) \quad (3.14a)$$

$$\forall x, z, \tilde{z} \in \text{int dom } \phi. \quad (3.14b)$$

A broad class of Bregman divergences has $\gamma_{\text{in}} = 2$ which constitutes a tight upper bound for any uniform TSE.

The analysis in [15] rests upon the *triangle-scaling gain* $G(x, z, \tilde{z})$ defined by the relaxed triangle-scaling inequality

$$D_\phi((1 - \theta)x + \theta z, (1 - \theta)x + \theta \tilde{z}) \leq G(x, z, \tilde{z}) \theta^\gamma D_\phi(z, \tilde{z}), \quad \forall \theta \in [0, 1]. \quad (3.15)$$

$G(x, z, \tilde{z})$ is bounded based on the relative scaling of the Hessian of ϕ at different points. In particular, using (3.15), *adaptive* ABPG algorithms are proposed of the form

$$\min_{x \in C} F(x) + \Psi(x), \quad (3.16)$$

with F being L -smooth relative to ϕ , C closed and C, Ψ convex and simple, in the sense that the key step of the ABPG method

$$z_{k+1} = \operatorname{argmin}_{x \in C} \left\{ F(y_k) + \langle \partial F(y_k), x - y_k \rangle + \theta_k^{\gamma-1} L D_\phi(x, z_k) + \Psi(x) \right\}, \quad (3.17)$$

can be solved efficiently. The convergence analysis of ABPG uses basic relations derived by [32] and [43] in order to relate two subsequent updates.

3.3.2. Algorithms. We specify three accelerated ABPG algorithms and briefly characterize their properties. Here, the *local* triangle-scaling property [15] is relevant which in view of Lemma 3.1 takes the form

$$D_{\phi_f}(Ax_{k+1}, Ay_{k+1}) < \theta_k^{\gamma_k} L D_\phi(z_{k+1}, z_k), \quad k \in \mathbb{N} \quad (3.18)$$

where the definitions of $x_{k+1}, y_{k+1}, z_{k+1}$ follow below in (3.19).

- *ABPG-e algorithm* [15]. Employing *exponent adaption* with an initial value $\gamma_0 > 2$ in (3.15), γ_k is reduced by a fixed factor $\delta < 1$ in subsequent iterations until the local triangle-scaling property (3.18) is met. The resulting value γ_k serves as a *posterior certificate* of accelerated convergence which cannot be guaranteed beforehand, however.
- *ABPG-g algorithm* [15]. Here the local triangle-scaling property is gradually achieved by *gain adaption*, that is by adjusting the gain $G_k = G(x_k, z_k, \tilde{z}_k)$ within an inner iterative loop.
- *FSMART algorithm* [10]. According to Proposition 3.2, the objective function f in (1.1) is L -smooth relative to ϕ in (2.55) with $L = \|A\|_1$. The ABPG iteration (3.17) leads with $\gamma = 1$ and $\Psi \equiv 0$ to the F(ast)-SMART iteration [10]

$$y_k = (1 - \theta_k)x_k + \theta_k z_k \quad (3.19a)$$

$$z_{k+1} = z_k \exp \left(-A^\top \log \left(\frac{Ay_k}{b} \right) / L \right) \quad (3.19b)$$

$$x_{k+1} = (1 - \theta_k)x_k + \theta_k z_{k+1}, \quad (3.19c)$$

where $x_0 = z_0 \in \operatorname{int}(\operatorname{dom} \phi)$ and $\theta_k \in (0, 1]$ satisfies $\frac{1 - \theta_{k+1}}{\theta_{k+1}^2} \leq \frac{1}{\theta_k^2}$. As the uniform TSE γ equals 1 for our choices of ϕ in (2.55), a $\mathcal{O}(1/k)$ rate can only be guaranteed theoretically according to [15, Thm. 1]. Convergence of the sequence $(x_k)_{k \in \mathbb{N}}$ generated by FSMART has remained an open issue. Empirically, however, FSMART effectively accelerates the SMART iteration. In addition, no additional cost per iteration are encountered, as opposed to the inner loop for gain adaption of ABPG-g.

In [15, Section 4.1], the authors address challenges associated with establishing a $\mathcal{O}(k^{-2})$ rate. This requires bounding the geometric mean $\bar{G} := (G_0^\gamma G_1 \cdots G_k)^{\frac{1}{k+\gamma}}$ of gains at each step without additional assumptions. The authors conjecture that in practical scenarios a specific reference function ϕ can be utilized, which may possess inherent properties contributing to rapid convergence. Exploiting such properties for the considered ϕ remains a subject for further research.

4. SMART: A GEOMETRIC PERSPECTIVE

In this section, we adopt the viewpoint of information geometry. First, we identify suitable retractions on each of our smooth manifolds (Section 4.1) which enables to describe SMART as Riemannian gradient flow with a fixed stepsize. Stepsize selection adapted by line search is studied in Section 4.2. Finally, we detail various geometric variants of the conjugate gradient iteration (Section 4.3) based on vector transport, which are supposed to accelerate Riemannian first-order gradient iterations.

4.1. Information Geometry. In this section, we further explore the geometry of the specific manifolds introduced in Section 2.4.

4.1.1. E-Geodesics, Exponential Maps. A basic concept of Riemannian geometry [34, Def. 1.4.3] is the exponential map $\exp: T\mathcal{M} \rightarrow \mathcal{M}$ defined by $\exp_p(t) = \gamma_v(t)$, where $\gamma_v(t)$ is the geodesic curve through $\gamma_v(0) = p \in \mathcal{M}$ with $v \in T_p\mathcal{M}$ with respect to the Riemannian connection. Since we work with the e-connection, we compute below the analogous *auto-parallel curves* called *e-geodesics*, and define the corresponding exponential maps.

A basic result of information geometry [26, Section 2.3] is that any statistical manifold \mathcal{M} corresponding to a distribution of the exponential family (2.14) is *flat* with respect to the e-connection and that the natural parameters θ constitute a corresponding affine coordinate system, which means that the e-geodesics have the simple form

$$\theta_u(t) = \theta + tu. \quad (4.1)$$

In particular, the natural parameters are defined everywhere, as $\text{dom } \psi$ in Table 2 displays. Both facts are convenient for numerical computation.

Since we model optimization problems using the η -coordinates, however, we transform the e-geodesics accordingly and still call them e-geodesics. For each specific manifold \mathcal{M} , this transformation is determined via the mappings $\eta = \eta(\theta)$ and its inverse $\theta = \theta(\eta)$ given by (2.16a) and (2.19a), respectively, while taking into account that \mathcal{M} is also flat with respect to the m -connection, with corresponding affine coordinates η and auto-parallel curves of the form

$$\eta_v(t) = \eta + tv. \quad (4.2)$$

The transformation reads

$$\gamma_v(t) := \eta(\theta_u(t)) \Big|_{\theta=\theta(\eta), u=u(\eta,v)} \quad (4.3a)$$

$$= \eta(\theta(\eta) + tu(\eta, v)) \quad (4.3b)$$




FIGURE 1. E-geodesics (4.5) emanating from three points η (black dots) in all directions $v = (\cos \omega, \sin \omega)$, $\omega \in [0, 2\pi)$ with $t \in \{0.02, 0.05, 0.1, 0.15, 0.2\}$.

with

$$u(\eta, v) = d\theta(\eta)v = \frac{d}{dt}\theta(\eta_v(t))\Big|_{t=0} \stackrel{(4.2)}{=} \frac{d}{dt}\theta(\eta + tv)\Big|_{t=0} \stackrel{(2.22)}{=} G(\eta)v. \quad (4.3c)$$

$$(4.3d)$$

Defining

$$\text{Exp}_\eta(tv) := \gamma_v(t), \quad (4.4)$$

which by construction is well-defined for any $t \in \mathbb{R}$ and tangent vector v (or v_η), we obtain for each of the specific manifolds from Section 2.4 the formula for the corresponding **e-geodesic** (see Figure 1 for an illustration)

$$\mathcal{P}_n: \quad \text{Exp}_\eta(tv) = \eta \cdot e^{\frac{tv}{\eta}}, \quad (4.5a)$$

$$\mathcal{B}_n: \quad \text{Exp}_\eta(tv) = \frac{\eta \cdot e^{\frac{tv}{\mathbb{1}-\eta}}}{\mathbb{1} - \eta + \eta \cdot e^{\frac{tv}{\mathbb{1}-\eta}}}, \quad (4.5b)$$

$$\mathcal{S}_n: \quad \text{Exp}_\eta(tv_\eta) = \frac{\eta \cdot e^{tG(\eta)v_\eta}}{1 - \langle \mathbb{1}_{n-1}, \eta \rangle + \langle \eta, e^{tG(\eta)v_\eta} \rangle}, \quad G(\eta) \text{ from (2.43)} \quad (4.5c)$$

$$\mathcal{S}_n: \quad \text{Exp}_p(tv) = \frac{p \cdot e^{\frac{tv}{p}}}{\langle p, e^{\frac{tv}{p}} \rangle}, \quad p = \vartheta^{-1}(\eta), \quad v = Bv_\eta \in T_0, \quad (4.5d)$$

where (4.5d) results from (4.5c) by applying ϑ^{-1} and the relations $G(\eta) = B^\top G(p)B$, $v = Bv_\eta$ and $B^\top z = \bar{z} - z_n \mathbb{1}_{n-1}$, $\forall z \in \mathbb{R}^n$ with B given by (2.48).

Lemma 4.1. *For the e-geodesics defined by (4.4), we have*

$$\text{Exp}_\eta(tv) = \partial\phi^*(\partial\phi(\eta) + tG(\eta)v). \quad (4.6)$$

Proof. The definition (4.4) yields with (4.3b)

$$\text{Exp}_\eta(v) = \eta(\theta(\eta) + u(\eta, v)) \stackrel{(4.3c)}{=} \eta(\theta(\eta) + G(\eta)v), \quad (4.7)$$

and substituting η, θ as given by (2.16a), (2.19a) together with (2.18) shows (4.6). \square

E-geodesics and, in particular, the mappings (4.5), provide *retractions* in the sense of [20, Definition 4.1.1], that are surrogates for the exponential map with respect to the Riemannian (Levi Civita) connection. This fact follows generally from [20, Prop. 5.4.1], but we include a short proof below.

Lemma 4.2 (e-geodesics provide retractions). *Let $\mathcal{M} \in \{\mathcal{P}_n, \mathcal{B}_n, \mathcal{S}_n\}$ denote either manifold and denote by*

$$x = \begin{cases} \eta \in \mathcal{M} & \text{if } \mathcal{M} = \mathcal{P}_n \text{ or } \mathcal{M} = \mathcal{B}_n, \\ p \in \mathcal{M} & \text{if } \mathcal{M} = \mathcal{S}_n, \end{cases} \quad (4.8)$$

and by $v \in T_x \mathcal{M}$ a corresponding tangent vector. Each of the mappings (4.5) provides a retraction

$$R_x: T_x \mathcal{M} \rightarrow \mathcal{M}, \quad R_x := \text{Exp}_x, \quad (4.9a)$$

i.e. it satisfies the defining conditions

$$R_x(0) = x, \quad 0 \in T_x \mathcal{M}, \quad (4.9b)$$

$$dR_x(0) = \text{id}_{T_x \mathcal{M}}, \quad \forall x \in \mathcal{M}. \quad (4.9c)$$

Proof. Condition (4.9b) is immediate from (4.6),

$$\text{Exp}_x(0) = \partial \phi^*(\partial \phi(x) + 0) = x. \quad (4.10)$$

Furthermore, we have

$$d\text{Exp}_x(u)v = \frac{d}{dt} \text{Exp}_x(u+tv) \Big|_{t=0} \stackrel{(4.6)}{=} \frac{d}{dt} \partial \phi^*(\partial \phi(x) + G(x)(u+tv)) \Big|_{t=0} \quad (4.11a)$$

$$= \partial^2 \phi^*(\partial \phi(x) + G(x)u)G(x)v. \quad (4.11b)$$

Thus

$$d\text{Exp}_x(0)v = \partial^2 \phi^*(\partial \phi(x))G(x)v \stackrel{(2.18)}{=} \stackrel{(2.19a)}{=} \partial^2 \psi(\theta)G(x)v \stackrel{(2.19b)}{=} v, \quad (4.12)$$

which establishes condition (4.9b). \square

Remark 4.1 (notation). In the remainder of this paper, we adopt the notation (4.8).

In connection with first-order optimization, the tangent vector v in (4.5) and in (4.6), respectively, will be a Riemannian gradient. Adopting the notation (4.8) and introducing the following shorthands for the corresponding e-geodesics will be convenient.

$$\mathcal{P}_n: \quad \exp_x(\partial f) := \text{Exp}_x(G(x)^{-1}\partial f) = x \cdot e^{\partial f}, \quad \partial f = \partial f(x), \quad (4.13a)$$

$$\mathcal{B}_n: \quad \exp_x(\partial f) := \text{Exp}_x(G(x)^{-1}\partial f) = \frac{x \cdot e^{\partial f}}{\mathbb{1} - x + x \cdot e^{\partial f}}, \quad \partial f = \partial f(x), \quad (4.13b)$$

$$\mathcal{S}_n: \quad \exp_x(\partial f) := \text{Exp}_x(\Pi_x \partial f) = \frac{x \cdot e^{\partial f}}{\langle x, e^{\partial f} \rangle}, \quad \partial f = \partial f(x). \quad (4.13c)$$

4.1.2. *Vector Transport.* In order to exploit *second-order* information on a manifold for geometric optimization, a *vector transport* is required [20, Def. 8.1.1]. Such a mapping provides a surrogate for *parallel transport* [34, Def. 3.1.2], analogously to replacing the exponential map with respect to the Riemannian connection by a retraction.

Given a retraction, a convenient way for obtaining a vector transport \mathcal{T} is to take the differential of the retraction [20, Section 8.1.2],

$$\mathcal{T} : T\mathcal{M} \rightarrow T\mathcal{M}, \quad \mathcal{T}_{(x,u)}(v) = d\text{Exp}_x(u)v \in T_{x'}\mathcal{M}, \quad u, v \in T_x\mathcal{M}, \quad x' = \text{Exp}_x(u). \quad (4.14)$$

Based on the characterization of e-geodesics given in Lemma 4.1, we can express the induced vector transports in closed form.

Lemma 4.3. *For the e-geodesics defined by (4.4), we have*

$$d\text{Exp}_x(u)v = G(x')^{-1}G(x)v, \quad u, v \in T_x\mathcal{M}, \quad x' = \text{Exp}_x(u). \quad (4.15)$$

Proof. With $x' = \text{Exp}_x(u)$ and $\theta' := \partial\phi(x')$, we obtain

$$\theta' = \partial\phi(x) + G(x)u, \quad (4.16)$$

due to (4.6) and $\partial\phi$ and $\partial\phi^*$ being one-to-one. Further, we have

$$d\text{Exp}_x(u)v \stackrel{(4.11)}{=} \partial^2\phi^*(\partial\phi(x) + G(x)u)G(x)v \stackrel{(4.16)}{\stackrel{(2.18)}}{=} \partial^2\psi(\theta')G(x)v \stackrel{(2.19a)}{\stackrel{(2.19b)}}{=} G(x')^{-1}G(x)v. \quad (4.17)$$

□

Applying, the previous result to the concrete manifolds from Section 2.3, we obtain the vector transports corresponding to the mappings (4.5),

$$\mathcal{P}_n: \quad \mathcal{T}_{(x,u)}(v) = \frac{x'}{x} \cdot v, \quad x' = \text{Exp}_x(u), \quad (4.18a)$$

$$\mathcal{B}_n: \quad \mathcal{T}_{(x,u)}(v) = \frac{x' \cdot (\mathbb{1} - x')}{x \cdot (\mathbb{1} - x)} \cdot v, \quad x' = \text{Exp}_x(u), \quad (4.18b)$$

$$\mathfrak{V}(\mathcal{S}_n): \quad \mathcal{T}_{(\eta, u_\eta)}(v_\eta) = \Pi_{\eta'}G(\eta)v_\eta, \quad \eta' = \text{Exp}_\eta(u_\eta), \quad (4.18c)$$

$$\mathcal{S}_n: \quad \mathcal{T}_{(x,u)}(v) = \Pi_{x'}\left(\frac{v}{x}\right), \quad v = Bv_\eta, \quad x' = \text{Exp}_x(u), \quad (4.18d)$$

where the mapping $\Pi_{\eta'}$ and $\Pi_{x'}$ are defined by (2.52) and (2.50b), respectively, and (4.18d) is equivalent to (4.18c) due to $B(\Pi_{\eta'}G(\eta)v_\eta) = B\Pi_{\eta'}B^\top G(x)Bv_\eta = \Pi_{x'}G(x)v$.

Below, vector transports \mathcal{T} will be applied to Riemannian gradients $v = \text{grad} f(x)$. Defining corresponding shorthands, analogous to the expressions (4.13) for e-geodesics emanating in the direction of Riemannian gradients, will be convenient.

Lemma 4.4 (vector transport of Riemannian gradients). *Let*

$$\mathcal{T}_{(x,u)}^g(f) := \mathcal{T}_{(x,u)}(\text{grad} f) \quad (4.19)$$

denote the vector transport of Riemannian gradients. Then

$$\mathcal{T}_{(x,u)}^g(f) = d\text{exp}_x(u)\partial f, \quad x' = \text{exp}_x(u), \quad \partial f = \partial f(x), \quad (4.20)$$

with exp_x defined by (4.13).

Proof. Evaluating the right-hand side of (4.20) for the expressions (4.13) is a routine calculation (cf. (4.21) below). Substituting the argument $\text{grad} f(x)$ into the corresponding expressions (4.18) shows equality. \square

Evaluating (4.19) or (4.20) yields the following closed-form expressions.

$$\mathcal{P}_n: \quad \mathcal{T}_{(x,u)}^g(f) = x' \cdot \partial f(x), \quad x' = \exp_x(u), \quad (4.21a)$$

$$\mathcal{B}_n: \quad \mathcal{T}_{(x,u)}^g(f) = x' \cdot (\mathbb{1} - x') \cdot \partial f(x), \quad x' = \exp_x(u), \quad (4.21b)$$

$$\mathcal{S}_n: \quad \mathcal{T}_{(x,u)}^g(f) = \Pi_{x'} \partial f(x), \quad x' = \exp_x(u). \quad (4.21c)$$

4.2. SMART As Riemannian Gradient Descent. In the context of minimizing an objective function f defined on a Riemannian manifold \mathcal{M} , any numerical first-order update must effectively utilize the Riemannian gradient $\text{grad} f(x) \in T_x \mathcal{M}$ in order to update $x \in \mathcal{M}$ (recall the notation: Remark 4.1). This typically involves the exponential map with respect to the Levi-Civita connection. However, often a retraction is used, because this is computationally cheaper than evaluating the exponential map or if the latter is not globally defined, as is the case for the simplex manifold, \mathcal{S}_n .

Therefore, in this paper, we employ throughout the exponential maps (4.5) corresponding to the e-geodesics on each of the three manifolds $\mathcal{P}_n, \mathcal{B}_n$ and \mathcal{S}_n . These retractions (Lemma 4.2), when applied to the Riemannian gradient, yield the expressions (4.13).

Proposition 4.1 (equivalence of e-geodesic and mirror descent map). *Application of the mappings Exp_η defined by (4.5), to the Riemannian gradient of an objective function f , yields the corresponding mirror descent mappings (3.3) and (1.6).*

Proof. The assertion follows directly from the expressions (4.13). \square

Corollary 4.1. *Let (\mathcal{M}, g) be endowed with the Riemannian metrics introduced in Section 2.3. Then the SMART iteration (1.6) equals*

$$x_{k+1} = \text{Exp}_{x_k}(-\tau_k \text{grad} f(x_k)) = \exp_{x_k}(-\tau_k \partial f(x_k)). \quad (4.22)$$

The step size τ_k can now be adapted to the current iterate by line search and we still obtain a global convergence result.

4.2.1. Armijo Line Search. Much like the Euclidean space's line search method, we can determine the step size τ_k through a curvilinear search on the manifold.

Theorem 4.1. [20, Thm. 4.3.1] *Let $(x_k)_{k \in \mathbb{N}}$ be a sequence generated by the iteration (4.22) with step-size $\tau_k = \beta^m \tau$ and scalars $\tau > 0, \beta, \sigma \in (0, 1)$, where m is the smallest nonnegative integer such that*

$$f(x_k) - f(x_{k+1}) \geq \sigma \tau_k \|\text{grad} f(x_k)\|_{x_k}^2. \quad (4.23)$$

Then every cluster point \hat{x} is a critical point of f , i.e. $\text{grad} f(\hat{x}) = 0$.

The Riemannian gradient method with the monotone geometric Armijo line search strategy (4.23), specialized to our specific setup, is stated in Algorithm 1.

Algorithm 1: SMART with geometric monotone Armijo line search [20]

initialization: Set $k = 0$, pick an initial point $x_0 \in \mathcal{M}$ and choose initial step size τ , parameters $\sigma \in (0, 1)$, $\beta \in (0, 1)$ and $\varepsilon > 0$. Set $f_0 = f(x_0)$, $v_0 = -\partial f(x_0)$ and $M_0 = G(x_0)^{-1}$ (cf. Table 1).

while $\|Mv_k\| > \varepsilon$ **do**

$$x_{k+1} = \exp_{x_k}(\tau v_k)$$

$$f_{k+1} = f(x_{k+1})$$

while $f_{k+1} - f_k > \sigma \tau v_k^\top M_k v_k$ **do**

$$\tau = \beta \tau$$

$$x_{k+1} = \exp_{x_k}(\tau v_k) \text{ by (4.13)}$$

$$f_{k+1} = f(x_{k+1}) \text{ with } f(x_{k+1}) = \text{KL}(Ax_{k+1}, b)$$

$$v_{k+1} = -\partial f(x_{k+1}) \text{ with } \partial f(x_{k+1}) = A^\top \log \frac{Ax_{k+1}}{b}$$

$$M_{k+1} = G(x_{k+1})^{-1} \text{ with inverse metric tensor as listed in Table 1}$$

$$x_{k+1} = \exp_{x_k}(\tau v_k)$$

Increment $k \leftarrow k + 1$.

4.2.2. *Hager-Zhang-type Line Search.* In [44], a novel approach to non-monotone line search was introduced, building on [45] where non-monotone strategies were first introduced for Newton's method. For a descent direction v_k , a step size τ_k is determined by checking either the non-monotone Wolfe conditions

$$f(x_k + \tau_k v_k) \leq C_k + \rho_1 \tau_k \partial f(x_k)^\top v_k, \quad (4.24)$$

$$\partial f(x_k + \tau_k)^\top v_k \geq \rho_2 \partial f(x_k)^\top v_k, \quad (4.25)$$

with scalars $\rho_1, \rho_2 > 0$, or, alternatively, the non-monotone Armijo condition

$$f(x_k + \tau_k v_k) \leq C_k + \sigma \tau_k \text{grad } f(x_k)^\top v_k, \quad (4.26)$$

where $\tau_k = \beta^m \tau$ and scalars $\tau > 0$, $\beta, \sigma \in (0, 1)$, with m the smallest nonnegative integer such that (4.26) holds.

Unlike monotone strategies that strictly ensure a decrease in the sequence of function values $(f(x_k))_{k \in \mathbb{N}}$ with each iteration, this approach does not require $f(x_{k+1}) < f(x_k)$ at every step. Sacrificing monotonicity, as seen in accelerated first-order convex optimization, is a deliberate choice made to achieve faster convergence and allows for more aggressive steps. For example, in [45], a maximum of recent function values must decrease, while in [44], the average of successive function values decreases.

Geometric Zhang-Hager-type line search has been explored in several works, including [21, 22, 46, 47]. In this context, we focus on the approach presented in [21], which builds upon the ideas of the Zhang-Hager technique [44].

Specifically, this approach involves determining the step size $\tau_k > 0$ such that it satisfies

$$f(R_{x_k}(\tau v_k)) \leq C_k + \rho_1 \tau \langle \text{grad } f(x_k), v_k \rangle_{x_k} - \rho_2 \tau^2 \|v_k\|_{x_k}^2, \quad (4.27)$$

where $\rho_1, \rho_2 > 0$. Additionally, each reference value C_k is calculated as a convex combination of C_{k-1} and $f(x_k)$. This allows a comparison of the current function value with a weighted average of previous values and enables an average decrease in the function value.

We now specialize [21, Alg. 1] to the negative Riemannian gradient and our particular manifolds. This customized version is presented as Algorithm 2.

Algorithm 2: SMART with geometric Hager-Zhang-type line search [21, Alg. 1]

initialization: Set $k = 0$, pick an initial point $x_0 \in \mathcal{M}$ and choose initial step size τ , parameters $0 < \rho_1 < \rho_2 < 1, \beta \in (0, 1), \rho \in [0, 1)$. Set $C_0 = f(x_0), Q_0 = 1, v_0 = -\partial f(x_0)$ and $M_0 = G(x_0)^{-1}$ (cf. Table 1).

while $\|Mv_k\| > \varepsilon$ **do**

$$x_{k+1} = \exp_{x_k}(\tau v_k)$$

$$f_{k+1} = f(x_{k+1})$$

while $f_{k+1} - C_k > \tau(\rho_1 - \tau\rho_2)v_k^\top M_k v_k$ **do**

$$\tau = \beta\tau$$

$$x_{k+1} = \exp_{x_k}(\tau v_k) \text{ by (4.13)}$$

$$f_{k+1} = f(x_{k+1}) \text{ with } f(x_{k+1}) = \text{KL}(Ax_{k+1}, b)$$

$$v_{k+1} = -\partial f(x_{k+1}) \text{ with } \partial f(x_{k+1}) = A^\top \log \frac{Ax_{k+1}}{b}$$

$$M_{k+1} = G(x_{k+1})^{-1} \text{ with inverse metric tensor as listed in Table 1}$$

$$x_{k+1} = \exp_{x_k}(\tau v_k)$$

$$Q_{k+1} = \rho Q_k + 1$$

$$C_{k+1} = \frac{\rho Q_k C_k + f(x_{k+1})}{Q_{k+1}}$$

Increment $k \leftarrow k + 1$.

The convergence of Algorithm 2 is analyzed in [21]. As $v_k = -\tau_k \text{grad} f(x_k)$ satisfies the conditions

$$\langle \text{grad} f(x_k), v_k \rangle_{x_k} \leq -c_1 \|\text{grad} f(x_k)\|_{x_k}^2, \quad (4.28)$$

$$\|v_k\|_{x_k} \leq c_2 \|\text{grad} f(x_k)\|_{x_k}, \quad (4.29)$$

where $c_1, c_2 > 0$, convergence of the sequence $(x_k)_{k \in \mathbb{N}}$ generated by Algorithm 2 to a stationary point follows under the following conditions ([21, Assumption 1]):

- (1) $f: \mathcal{M} \rightarrow \mathbb{R}$ is continuously differentiable on \mathcal{M} and is bounded below in the level set $\mathcal{L}_{x_0} = \{x \in \mathcal{M} : f(x) \leq f(x_0)\}$;
- (2) the differential $d\hat{f}_x$ of the function

$$\hat{f}_x: T_x \mathcal{M} \rightarrow \mathbb{R}, \quad \hat{f}_x(v) := f(R_x(v)), \quad v \in T_x \mathcal{M}, \quad (4.30)$$

represented by the tangent vector $u_{\hat{f}(v)}$ and the identification $T_v(T_x \mathcal{M}) \simeq T_x \mathcal{M}$ such that

$$d\hat{f}_x(v)w = \langle u_{\hat{f}(v)}, w \rangle_x, \quad \forall v, w \in T_x \mathcal{M}, \quad (4.31)$$

is required to be uniformly Lipschitz continuous at $0 \in T_x \mathcal{M}$, i.e. there exist $\kappa, L > 0$ such that

$$\|d\hat{f}_x(v) - d\hat{f}_x(0)\|_x = \|u_{\hat{f}(v)} - u_{\hat{f}(0)}\|_x \leq L\|v\|_x, \quad x \in \mathcal{L}_{x_0}, \quad v \in T_x \mathcal{M} \quad (4.32)$$

holds, where $\|v\|_x \leq \kappa$.

We elaborate condition (2). By the chain rule and the general relation

$$df(x)v = \langle \text{grad } f(x), v \rangle_x, \quad x \in \mathcal{M}, \quad v \in T_x \mathcal{M}, \quad (4.33)$$

one has

$$d\widehat{f}_x(v)w = df(R_x(v))dR_x(v)w = \langle \text{grad } f(R_x(v)), dR_x(v)w \rangle_x. \quad (4.34)$$

Thus, by (4.31),

$$u_{\widehat{f}(v)} = dR_x(v)^\top \text{grad } f(R_x(v)) \quad (4.35)$$

and in particular, due to (4.9b), (4.9c),

$$u_{\widehat{f}(0)} = \text{grad } f(x). \quad (4.36)$$

Taking now into account (4.9a) and (4.3), condition (4.32) becomes

$$\|d\widehat{f}_x(v) - d\widehat{f}_x(0)\|_x = \|dR_x(v)^\top \text{grad } f(R_x(v)) - \text{grad } f(x)\|_x \quad (4.37a)$$

$$= \|G(x)G(x')^{-2}\partial f(x') - G(x)^{-1}\partial f(x)\|_x, \quad x' = R_x(v) \quad (4.37b)$$

$$\leq L\|v\|_x = L\|R_x^{-1}(x')\|_x. \quad (4.37c)$$

As a simple concrete scenario, consider the orthant $\mathcal{P}^2 \subset \mathbb{R}^2$ with objective function $f(x) = D_\phi(x, \mathbb{1})$ given by (2.56a). Fixing x and subtracting the left-hand side from the right-hand side of (4.37), defines the function

$$h_x(x') := L\|R_x^{-1}(x')\|_x - \|G(x)G(x')^{-2}\partial f(x') - G(x)^{-1}\partial f(x)\|_x \quad (4.38a)$$

where

$$R_x^{-1}(x') = x \cdot \log \left(\frac{x'}{x} \right) \quad (4.38b)$$

by (4.5a) and (4.9a). Using $G(x) = \text{Diag}(x)^{-1}$ and $\partial f(x) = \log x$, we thus have

$$h_x(x') = L \left\| x \cdot \log \frac{x'}{x} \right\|_x - \left\| \frac{x' \cdot x'}{x} \cdot \log x' - x \cdot \log x \right\|_x \quad (4.39)$$

and specifically for $x = \mathbb{1}$

$$h_{\mathbb{1}}(x') = L\|\log x'\| - \|x' \cdot x' \cdot \log x'\|. \quad (4.40)$$

Since $((x'_i)^2 \log x'_i - L \log x'_i) \rightarrow \infty$ for $x'_i \rightarrow \infty$ and fixed $L > 0$, we conclude that $h_x(x') \geq 0$ and thus also condition (2) above *might not* hold in general.

Remark 4.2 (conditions for convergence of Algorithm 2). The above reasoning indicates that specifying *general* conditions for the convergence of Algorithm 2 for *each* manifold (2.24) considered in this paper is a subtle issue. Although we indicated how condition (2) above might be violated, we did not take into account that v is bounded, $\|v\|_x \leq \kappa$, which in turn bounds x' by $x' = R_x(v)$. Thus, we regard this issue as open and leave it for future work, motivated also by the fact that Algorithm 2 converged in all our numerical experiments.

4.2.3. *Barzilai-Borwein Method.* An alternative non-monotone line search approach expected to perform well is the Barzilai-Borwein (BB) method, introduced in [48] for the Euclidean setting.

Canonical line search methods in Euclidean optimization aim to approximate the exact line search problem given by

$$\tau_k = \arg \min_{\tau} f(x_k + \tau v_k) \quad (4.41)$$

where v_k is a descent direction.

The BB method approaches this problem differently. Let $g_k = \partial f(x_k)$ denote the Euclidean gradient of f at x_k , $s_k = x_{k+1} - x_k$ and $y_k = g_{k+1} - g_k$. The approach is motivated by a quasi-Newton iteration, where the update rule is given by $x_{k+1} = x_k - B_k^{-1} g_k$. Here, B_k is an approximation of the Hessian of the objective function, which satisfies the secant equation

$$B_{k+1} s_k = y_k. \quad (4.42)$$

Barzilai and Borwein simplify B_{k+1} to a scalar $\frac{1}{\tau_k}$, which typically cannot precisely fulfill the secant equation (4.42). Instead, if $\langle s_k, y_k \rangle > 0$, they approximate it as

$$\tau_{k+1} = \arg \min_{\tau} \left\| \frac{1}{\tau} s_k - y_k \right\|^2, \quad (4.43)$$

resulting in the so-called *long BB step length*, characterized by the optimality conditions of (4.43) that give

$$\tau_{k+1} = \frac{\langle s_k, s_k \rangle}{\langle s_k, y_k \rangle}. \quad (4.44)$$

We can now translate this condition into the geometric setting by using the Riemannian metric, Riemannian gradient and the retraction corresponding to the constraint manifold. The resulting Riemannian gradient method with geometric Barzilai-Borwein line step is introduced in [49]. Instead of the difference $x_{k+1} - x_k$ they consider the vector

$$v_k = -\tau_k \operatorname{grad} f(x_k) \in T_{x_k} \mathcal{M}$$

and transport it to $T_{x_{k+1}} \mathcal{M}$. This yields

$$s_k := \mathcal{T}_{(x_k, v_k)}(v_k) = \mathcal{T}_{(x_k, v_k)}(-\tau_k \operatorname{grad} f(x_k)) = -\tau_k \mathcal{T}_{(x_k, v_k)}^g(f). \quad (4.45)$$

The difference of the Riemannian gradients is defined as in [49]

$$y_k := \operatorname{grad} f(x_{k+1}) - \mathcal{T}_{(x_k, v_k)}^g(f) \stackrel{(4.45)}{=} \operatorname{grad} f(x_{k+1}) + \frac{1}{\tau_k} \mathcal{T}_{(x_k, v_k)}(v_k). \quad (4.46)$$

Analogously to the Euclidean setting the Riemannian long BB step length writes

$$\tau_{k+1} = \frac{\langle s_k, s_k \rangle_{x_{k+1}}}{\langle s_k, y_k \rangle_{x_{k+1}}}. \quad (4.47)$$

To guarantee convergence, the step size additionally needs to fulfill the non-monotone Armijo condition. Similar to [44] the update of the cost C_k is replaced by a combination of the last $m_k \leq 10$ function values

$$C_k = \max_{0 \leq j \leq m_k} f(x_{k-j}),$$

where $m_0 = 0$ and $m_k = \min\{m_{k-1} + 1, 10\}$ for $k > 0$. The algorithm is summarized in Algorithm 3.

Algorithm 3: Riemannian Gradient Descent with Barzilai-Borwein Line Search [49]

Data: initial point $x_0 \in \mathcal{M}$, function f , initial step size τ_0 , $\gamma_{\min} \in [0, 1]$, $\gamma_{\max} \geq 1$, $\rho, \beta \in (0, 1)$.

Result: sequence of (x_k) towards $\arg \min f$.

$C_0 = f(x_0)$

for $k = 0, 1, 2, \dots$ **do**

$v_k = -\text{grad } f(x_k)$

$x_{k+1} = R_x(\tau_k v_k)$

calculate $s_k = -\tau_k \mathcal{T}_{(x_k, v_k)}(\text{grad } f(x_k))$ and $y_k = \text{grad } f(x_{k+1}) + \frac{1}{\tau_k} s_k$

calculate $\gamma_k = \frac{\langle s_k, s_k \rangle_{x_{k+1}}}{|\langle s_k, y_k \rangle_{x_{k+1}}|}$

set $\gamma_k = \max(\gamma_{\min}, \min(\gamma_k, \gamma_{\max}))$

find step size $\tau_k = \gamma_k \beta^m$ where m is the smallest integer such that

$$f(R_{x_k}(\tau_k v_k)) \leq C_k + \rho \tau_k \langle \text{grad } f(x_k), v_k \rangle_{x_k}$$

update the iterate $x_{k+1} = R_{x_k}(\tau_k v_k)$

update $C_{k+1} = \max_{0 \leq j \leq m_k} f(x_{k+1-j})$

end

Theorem 4.2. *Let $(x_k)_{k \in \mathbb{N}}$ be a sequence generated by Algorithm 3. Then every cluster point \hat{x} is a critical point of f , i.e. $\text{grad } f(\hat{x}) = 0$.*

Proof. Since the retractions considered in (4.5) are globally defined, the result can be proved analogously to [50, Thm. 3]. For details see [51, Thm. 4.13], where the more general update of C_k (as used in Algorithm 3) is considered. \square

4.3. Accelerating SMART: Riemannian Conjugate Gradient. The algorithms introduced above only exploit the negative Riemannian gradient $v_k = -\text{grad } f(x_k)$ as descent direction at every iteration. Alternative choices of descent directions which proved to be efficient in Euclidean scenarios, may be generalized to geometric scenarios as well. [21] extends the conjugate gradient iteration to a Riemannian method by updating the search direction at every iteration to

$$v_{k+1} = -\text{grad } f(x_{k+1}) + \beta_k \mathcal{T}_{(x_k, \alpha_k v_k)}(v_k),$$

with the vector transport \mathcal{T} defined by (4.14) and parameter β_k yet to be chosen.

Some of the common choices for the parameter β_k for existing Riemannian conjugate gradient methods include (see [22])

$$\beta_{k+1}^{FR} = \frac{\|\text{grad } f(x_{k+1})\|_{x_{k+1}}^2}{\|\text{grad } f(x_k)\|_{x_k}^2}, \quad (\text{Fletcher-Reeves}) \tag{4.48a}$$

$$\beta_{k+1}^{PR} = \frac{\langle \text{grad } f(x_{k+1}), y_k \rangle_{x_{k+1}}}{\|\text{grad } f(x_k)\|_{x_k}^2}, \quad (\text{Polak-Ribière}) \quad (4.48b)$$

$$y_k = \text{grad } f(x_{k+1}) - \mathcal{T}_{(x_k, \alpha_k v_k)}^g(f), \quad (\mathcal{T}^g \text{ by (4.21)}) \quad (4.48c)$$

$$\beta_{k+1}^{DY} = \frac{\|\text{grad } f(x_{k+1})\|_{x_{k+1}}^2}{\langle \text{grad } f(x_{k+1}), \mathcal{T}_{(x_k, \alpha_k v_k)}(v_k) \rangle_{x_{k+1}} - \langle \text{grad } f(x_k), v_k \rangle_{x_k}}, \quad (\text{Dai-Yuan}) \quad (4.48d)$$

$$\beta_{k+1}^{HS} = \frac{\langle \text{grad } f(x_{k+1}), y_k \rangle_{x_{k+1}}}{\langle \text{grad } f(x_{k+1}), \mathcal{T}_{(x_k, \alpha_k v_k)}(v_k) \rangle_{x_{k+1}} - \langle \text{grad } f(x_k), v_k \rangle_{x_k}}, \quad (\text{Hestenes-Stiefel}) \quad (4.48e)$$

$$\beta_{k+1}^{HZ} = \beta_{k+1}^{HS} - \mu \frac{\|y_k\|_{x_{k+1}}^2 \langle \text{grad } f(x_{k+1}), \mathcal{T}_{(x_k, \alpha_k v_k)}(v_k) \rangle_{x_{k+1}}}{(\langle \text{grad } f(x_{k+1}), \mathcal{T}_{(x_k, \alpha_k v_k)}(v_k) \rangle_{x_{k+1}} - \langle \text{grad } f(x_k), v_k \rangle_{x_k})^2}, \quad (\text{Hager-Zhang}) \quad (4.48f)$$

$$\beta_k^{OV} = \mu_k \frac{\langle \text{grad } f(x_{k+1}), \mathcal{T}_{(x_k, \alpha_k v_k)}(v_k) \rangle_{x_{k+1}}}{-\|v_k\|_{x_k}^2}, \quad [21] \quad (4.48g)$$

where $\mu > 0$ in (4.48f) and $\{\mu_k\}$ in (4.48g) is a bounded positive sequence such that $0 < \mu_k < \infty, \forall k$.

Algorithm 4 summarizes the Riemannian CG iteration.

Algorithm 4: Riemannian Conjugate Gradient [22]

Data: initial point x_0 , objective function f , initial step size α_0 , $\sigma \in (0, 1)$, $\beta \in (0, 1)$.

Result: sequence of $(x_k)_{k \in \mathbb{N}}$ converging to a minimum of f .

Set the initial search direction to $v_0 = -\text{grad } f(x_0)$.

for $k = 0, 1, 2, \dots$ **do**

$\bar{\alpha} = \alpha_k$

Find a step size $\alpha_k = \bar{\alpha} \beta^m$, where m is the smallest integer such that (with R_{x_k} by (4.9a))

$$f(R_{x_k}(\alpha_k v_k)) - f(x_k) \leq \sigma \alpha_k \langle \text{grad } f(x_k), v_k \rangle_{x_k}.$$

Update the iterate to $x_{k+1} = R_{x_k}(\alpha_k v_k)$.

Determine β_{k+1} by evaluating one of the expressions (4.48).

Update the search direction to

$$v_{k+1} = -\text{grad } f(x_{k+1}) + \beta_k \mathcal{T}_{(x_k, \alpha_k v_k)}(v_k).$$

$k = k + 1$

end

Remark 4.3 (convergence conditions for Algorithm 4). Conditions [22, Ass. 3.1, 3.2] that imply convergence can be extended to a retraction. For instance, condition [22, Ass. 3.2] leads to (4.38), which unfortunately does not hold in general for our manifolds. While we consistently

observe convergence in all our numerical experiments, we defer the detailed analysis for future investigation.

Table 3 summarizes the algorithms explored in this paper.

Algorithm	Convergence rate	Convergence of iterates	Boundary
SMART	$O(1/k)$ (Thm. 3.1 (ii))	towards solution (Thm. 3.1 (i))	✓
FSMART	emp. $O(1/k^\gamma)$ [15, Thm. 1]	open	✓
FSMART-e	emp. $O(1/k^\gamma)$ [15, Thm. 2]	open	✓
FSMART-g	emp. $O(1/k^\gamma)$ [15, Thm. 3]	open	✓
RG-Armijo	open	conv. subsequence (Thm. 4.1)	✗
RG-HZ	open	open	✗
RG-BB	open	conv. subsequence (Thm. 4.2)	✗
CG	open	open	✗

TABLE 3. **Overview. Algorithms and their properties:** convergence rate, convergence of iterates (vs. function values), boundary behavior. The first four algorithms are well-defined on the boundary of the feasible sets whereas the latter four are not. The FSMART algorithms exhibit empirically acceleration with a convergence rate $\mathcal{O}(1/k^\gamma)$, $\gamma \in (1, 2)$. For FSMART-e (ABPG-e [15, Alg. 2] adapted to our scenario) and FSMART-g (ABPG-g [15, Alg. 3] adapted to our scenario) only this can be checked a posteriori. In summary, SMART is a well understood method that achieves optimal convergence rate in terms of function values and converges to the minimum Bregman distance solution from any starting point. The table lists a range of algorithms which emerged from SMART. They show empirically a faster convergence rate and stimulate research on various open points relevant to non-Lipschitz scenarios of optimization.

5. EXPERIMENTS

Due to the variety of algorithms and multiple manifolds, we opt to focus solely on the Bernoulli manifold, detailed in Section 2.4.2. Our study includes an extensive comparison between SMART and the recent accelerated Bregman proximal gradient methods discussed in [15]. Our adaptation of these methods, as explained in Section 3, result in several distinct variations of SMART, namely FSMART, FSMART-e, and FSMART-g. Additionally, we characterize SMART as a Riemannian gradient descent scheme on the parameter manifold induced by the Fisher-Rao geometry. This characterization allows us to incorporate a retraction, as specified in equation (4.5b), into a geometric line search strategy. We refer to this variant of SMART as Riemannian gradient (RG). We explore three well-known line search methods within our geometric setup: Armijo, Hager-Zhang, and Barzilai-Borwein. Furthermore, we include a Riemannian conjugate gradient from [21] in our assessment. We consider relative-entropy regression over the box in three diverse applications: image deblurring, discrete tomography, and sparse recovery with expanders. For each application, we compare the algorithms examined in this paper and demonstrate significant performance improvements over the basic version of SMART. For

detailed results and discussions, please refer to Figures 2 – 12. Across all three applications we generate problem data uniformly. We start with either a binary signal $\hat{x} \in \{0, 1\}^n$ or a signal $\hat{x} \in [0, 1]^n$ that undergoes *undersampling* by $A \in \mathbb{R}_+^{m \times n}$, where $m \leq n$. This process results in $b = A\hat{x} + e$ with $e \in \mathbb{R}^m$ representing added noise. Thus, problem (1.1) is applicable to both the inconsistent (or underdetermined) system $Ax \approx b$. Furthermore, we are interested in analyzing the algorithms' behavior when the solution resides at the boundary. Therefore, observing both scenarios is crucial for our analysis.

5.1. Implementation Details. We implemented the algorithms listed in Table 3 for solving (1.1) numerically. The Bregman kernel is ϕ_{\square} from (2.55). We chose the uninformative barycenter $\frac{1}{2}$ as the natural initialization for all algorithms. For each algorithm the same set of parameters was used across all experiment problem instances. The algorithms and corresponding parameter choices are listed below:

SMART: solely performs the multiplicative update specified in (1.6) with its stepsize fixed to $\tau_k = \frac{1}{L}$, where again f is L-smooth relative to ϕ_{\square} .

FSMART: based on the iteration suggested in [10], where initially $\theta_0 = 1$ is chosen, which is then subsequently updated via $\theta_{k+1} = \frac{\sqrt{\theta_k^4 + 4\theta_k^2} - \theta_k^2}{2}$, as suggested in [43].

FSMART-e: as described in [15, Algorithm 2] was applied to (1.1) with parameters $\gamma_{\min} = 1$, $\gamma_0 = 5$ and $\delta = 0.05$. The choices for γ_0 and δ deviate slightly from the recommendation in [15], but were chosen to facilitate fastest possible convergence on the selected problem instances. To ensure comparability restarting mechanisms and stopping criteria based on the divergence of iterates were foregone. Updates for θ were conducted via Newton's method.

FSMART-g: specified in [15, Algorithm 3] is used with parameters: $\rho = 1.2$, $\gamma = 2$ and $G_{\min} = 10^{-3}$. Restarting, stopping criteria and updating θ was handled analogously to ABPG-e.

RG-Armijo: a SMART iteration with Armijo line search for choosing the step size τ_k via the retraction in (4.5b) to iterate according to (4.22), see Algorithm 1. The line search parameters are $\sigma = 10^{-3}$, $\beta = 0.8$ and initial stepsize $\tau = 0.2$.

RG-HZ: a SMART iteration with Hager-Zhang line search which includes the Armijo line search condition implemented with the same parameters as stated in **RG-Armijo** for comparison. Additional parameters are $\sigma_2 = 10^{-3}$ and $\rho = 0.5$.

RG-BB: is a SMART iteration with Barzilai-Borwein line search which includes the Armijo line search condition implemented with the same parameters as stated in **RG-Armijo** for comparison. Additional parameters are $\gamma_{\min} = 10^{-7}$ and $\gamma_{\max} = 1.0$. For details see Algorithm 3.

CG: is the Riemannian conjugate gradient, Algorithm 4. It includes the Armijo line search condition implemented with the same parameters as stated in **RG-Armijo** for comparison.

PG: is the projected gradient algorithm with Armijo linesearch.

5.2. Toy Example. We first conduct a toy experiment with

$$\hat{x} = (1 \ 1)^{\top}, \quad A = (0.25 \ 0.75), \quad b = 1. \quad (5.1)$$

In this specific underdetermined scenario, signal \hat{x} , can be uniquely reconstructed from a single measurement b when confined to the box. We demonstrate the smooth trajectory of iterates in contrast to the behavior observed in projected gradient descent in Figure 2 (left).

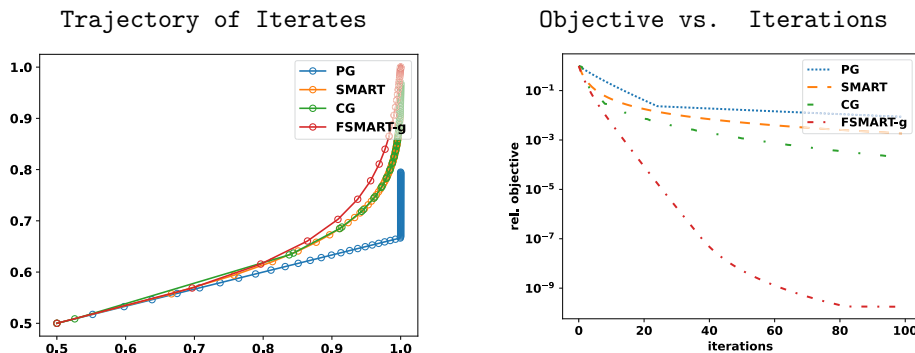


FIGURE 2. **Toy example:** Trajectories of iterates, starting at $x_0 = (0.5, 0.5)^\top$ and converging to the unique solution $\hat{x} = (1, 1)^\top$, are displayed in the **left** panel. The projected gradient method (PG) shows a kink due to its nonsmooth projection step. Methods aware of the constraints geometry depict smoother, more optimal iterate trajectories by exploiting smooth Riemannian geometry to accommodate box constraints. In the **right** panel, FSMART-g stands out with larger initial steps and a rapid decrease in the relative objective value.

5.3. Expander Graphs. We consider binary measurement matrices $A \in \{0, 1\}^{m \times n}$ with varying row numbers, $m \in \{40, 70, 100\}$, and a fixed number of columns $n = 200$ with $\|A\|_1 = 12$. These matrices correspond to adjacency matrices of expander graphs as detailed in [52]. The sparsity of $\hat{x} \in \{0, 1\}^{200}$ is set to 20. With high probability, all algorithms should converge to the same solution, given the chosen sparsity and structure of A , ensuring likely uniqueness.

First, we compare different CG variants that differ in the β_k parameters choice in (4.48). Results are depicted in Figure 3. Additionally, Figure 4 illustrates a comparison of various methods described in Section 5.1, where the DY mode (4.48d) is employed for the conjugate gradient (CG) approach. The original signal \hat{x} and the resulting reconstructions by different algorithms are presented in Figure 5. Finally, Figure 6 displays a-posteriori certificates obtained from FSMART-e, illustrating γ_k observations across iterations in the left panel, while visualizing average matrix vector operations on the right panel.

5.4. Tomographic Reconstruction. We consider large scale tomographic reconstruction as a problem class, where we reconstruct the three phantoms shown in Fig. 7. Tomographic projection matrices A were generated using the ASTRA-toolbox, with parallel beam geometry and equidistant angles in the range $[0, \pi]$. Each entry in A is nonnegative as it corresponds to the length of the intersection of a ray with a pixel. The undersampling rate was chosen to be 2% which corresponds to 20 projection angles. None of the images in Fig. 7 is expected to be the unique solution to $Ax = b$ within the $[0, 1]^n$ box due to the large undersampling ratio [53]. Hence, different algorithms might converge to different solutions within $[0, 1]^n$. Figure 9 shows

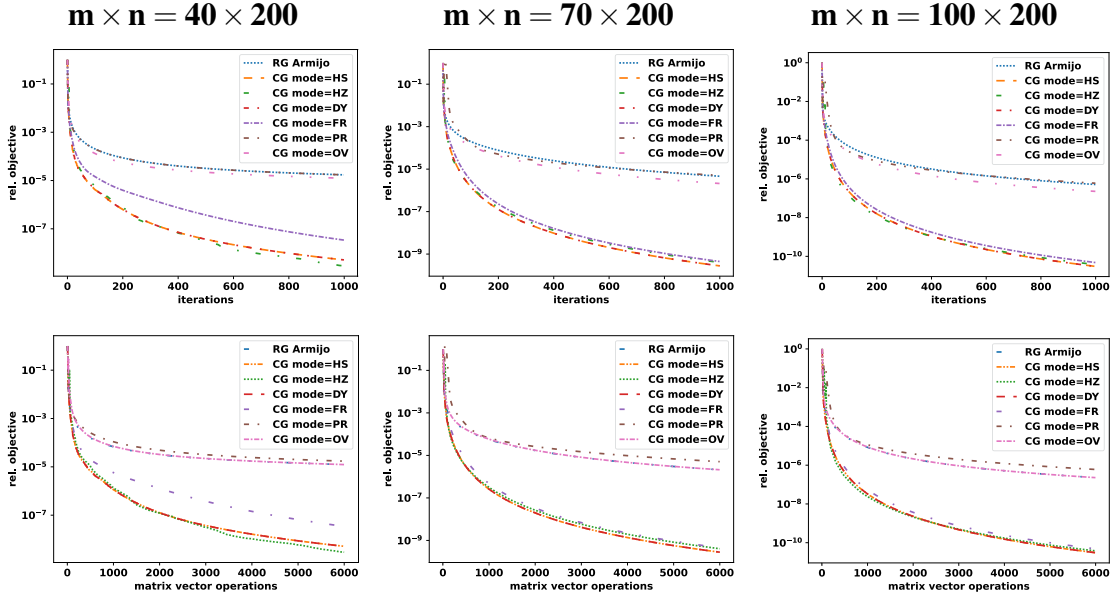


FIGURE 3. **Comparison of CG variants on expander graphs** across three instances with varying levels of ill-posedness ($m \in \{40, 70, 100\}$ from left to right). The plots illustrate the relative decrease in objective values over iterations (**first row**) and costly operations (**second row**) for different approaches detailed in (4.48) used to select β_k for updating the search direction in the conjugate gradient approach. Additionally, Riemannian gradient descent with Armijo line search is considered as baseline. We observe a rough clustering of the variants into two groups: $\{\text{OV}, \text{PR}\}$ perform similarly to the Armijo baseline, while $\{\text{DY}, \text{FR}, \text{HS}, \text{HZ}\}$ exhibit notably faster convergence.

the a-posteriori certificates obtained from FSMART-e by observing γ_k over the iterations in the left panel, while the right panel visualizes the average matrix vector operations.

5.5. Deblurring. We evaluated the algorithms listed in Table 3 for the deblurring task using the images displayed in the first row of Figure 10. The first image, 'Kitten,' has pixel values within the range $[0, 1]$, while 'Tiger' and 'QR-Code' admit binary pixel values. We degraded these images by severe Gaussian blur with a 33×33 mask and $\sigma = 10$, as illustrated in the second row of Figure 10. Subsequent rows display the resulting reconstructions after 1000 iterations using the SMART, FSMART, and CG (with DY mode) algorithms. We refer to Figure 11 for the plot illustrating the relative decrease in the objective. Finally, Figure 12 shows a-posteriori certificates obtained from FSMART-e, depicting observations of γ_k over the iterations in the left panel, and average matrix vector operations in the right panel.

5.6. Discussion. We discuss our observations according to the following aspects:

- (1) Observations about acceleration in the Bregman framework;
- (2) Observations about line search strategies;
- (3) Bregman vs. Riemannian framework: empirical findings;

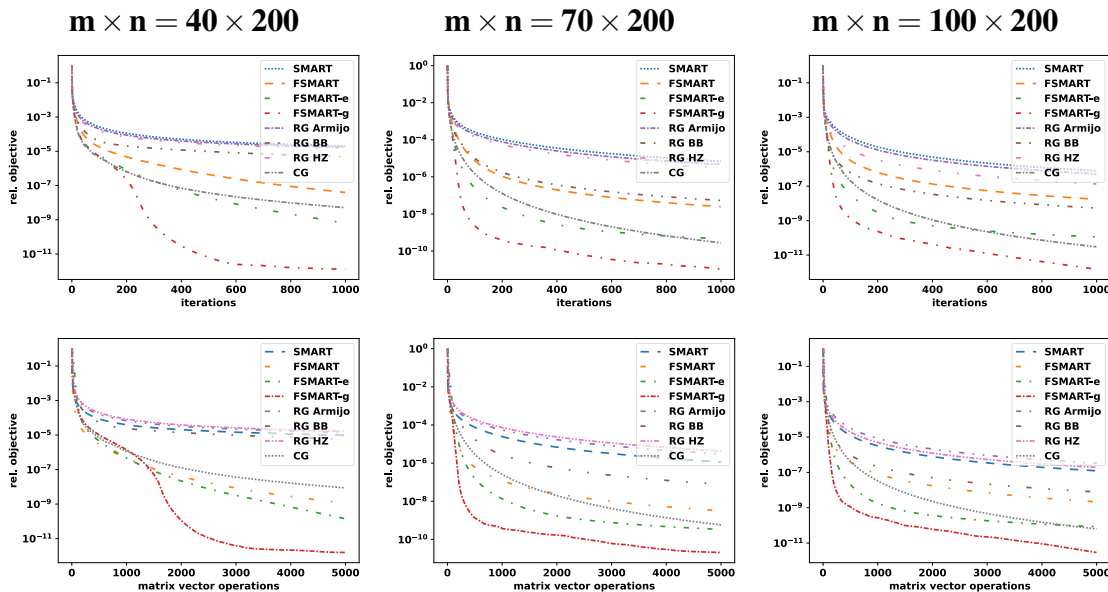


FIGURE 4. **Comparison of different algorithms on expander graphs** across instances with varying levels of ill-posedness ($m \in \{40, 70, 100\}$ from left to right). The plots illustrate the relative decrease in objective values over iterations (**first row**) and in terms of costly operations (**second row**) for the different approaches detailed in Section 5.1. Among the accelerated algorithms within the Bregman regime, FSMART-e and particularly FSMART-g demonstrate superior performance. Within the Riemannian optimization regime, the conjugate gradient (CG) implemented with the DY mode for β_k -update rule (4.48d) achieves results comparable to FSMART-e.

- (1) (a) *Accelerated Bregman proximal gradient: convergence rate.* In all experiments, the ABPG methods (FSMART, FSMART-e, FSMART-g), especially the adaptive variants (FSMART-e, FSMART-g), demonstrate superior performance compared with the SMART method. Moreover, we obtain numerical certificates for the empirical $O(1/k^\gamma)$ rate, where $\gamma \in (1, 2)$ in all our experiments. Specifically, the adaptive variants can automatically search for the largest possible γ_k for which the convergence rate holds for finite k , even though $\gamma = 1$ in the TSE (3.13) for the KL based divergences considered in this paper. Notably, $\gamma_k \geq 2$ when far from the solution and consistently $\gamma_k = 1$ as we approach the solution. Consequently, this causes the convergence rate to slow down as we near the solution.
- (b) *Accelerated Bregman proximal gradient: performance.* The comparison involves three algorithms across three problem instances: In the case of expander graphs, FSMART-g is superior to FSMART and FSMART-e. However, for tomography, FSMART and FSMART-g perform comparably and outperform FSMART-e. A similar pattern emerges in image deblurring, except for the QR-code instance, where FSMART-e notably leads over FSMART-g and FSMART.

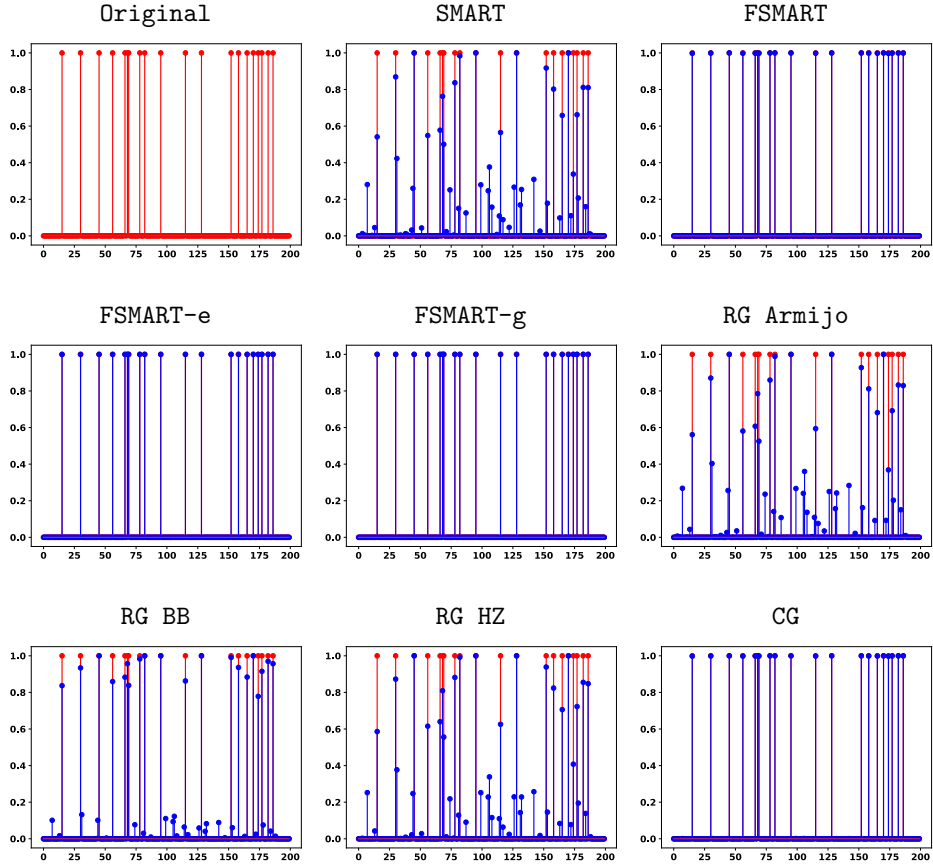


FIGURE 5. **Sparse spike reconstructions** from adjacency matrices corresponding to expander graphs, generated by different algorithms, are displayed for $m = 40$ after 1000 iterations. The original binary signal \hat{x} with sparsity 20 is presented in the top left. Faster algorithms (see Figure 4) demonstrate perfect signal reconstruction. However, it is worth noting that reconstructions from slower algorithms can be thresholded entrywise at 0.5, transforming them into perfect reconstructions.

- (2) (a) *Line search strategies: overall performance.* As the SMART iteration represents a Riemannian gradient (RG) step with a fixed steplength, it serves as our baseline for comparison against RG with Armijo, Hager-Zhang (HZ), or Barzilai-Borwein (BB) steplength selection. When not considering computational cost, BB line search outperforms HZ, with HZ slightly surpassing Armijo. Meanwhile, Armijo itself surpasses SMART (without line search) on expander graphs. A similar trend is observed in the tomography examples, with HZ displaying better performance than Armijo on the non-binary image. However, in deblurring, Armijo and HZ perform comparably but are notably surpassed by BB. Interestingly, SMART outperforms the line search methods in this scenario. It is worth noting that certain sets of line search parameters could reverse this scenario. Nevertheless, our aim was to avoid tuning line search parameters individually for each problem instance.




FIGURE 6. **A-posteriori certificates for spike recovery** are obtained from FSMART-e by observing γ_k -values over iterations shown in the (left) panel for all problem instances. We observe that γ_k drops to 1 in all instances. We explored the drop-down-point for each instance and observed that it occurs when FSMART-e approaches the solution. **Average of matrix vector operations** are shown in the (right) panel for each algorithm over iterates and instances. By definition SMART, FSMART constantly employ two matrix vector operations per iteration. The other algorithms require potentially more operations as they employ line search strategies.

- (b) *Line search strategies: computational costs.* Due to the specific line search parameters chosen, the average number of matrix-vector operations per iteration ranges between 4 and 7 for all RG methods using Armijo, HZ, or BB. Among these, Armijo consistently demonstrates the highest computational efficiency and robustness across all problem instances. This contradicts the empirical findings in [46], where HZ was performing the best, albeit on a different type of manifold.
- (3) (a) *Bregman vs. Riemannian framework: convergence.* Only SMART and Riemannian Gradient with Armijo line search guarantee a monotonically decreasing objective function. A non-monotonic decrease is observed for RG with HZ and BB line searches, particularly noticeable in the tomographic instances. There's no significant difference in convergence rates when the solution lies at the boundary. These trends are consistent across algorithms from both the Bregman and Riemannian gradient frameworks.
- (b) *Bregman vs. Riemannian framework: best performance is achieved by accelerated BPG and by CG.* In tomographic recovery, the conjugate gradient method with DY mode (4.48d) from the Riemannian optimization framework demonstrates superior performance, followed by FSMART-g from the Bregman framework. Moreover, in deblurring and spike recovery tasks, the accelerated Bregman methods alongside the Riemannian CG method notably outperform the Riemannian gradient methods.




FIGURE 7. **Phantoms and reconstructions:** The first row shows large-scale phantoms of size 1024×1024 used for numerical evaluation. Subsequent rows illustrate the resulting reconstructions from 2% undersampling from the algorithms SMART, FSMART, and CG (with DY mode), after 400 iterations. We refer to Figure 8 for the plot depicting the relative decrease of the objective and remark that reconstructions from faster algorithms (last 2 rows) resemble the original more closely.

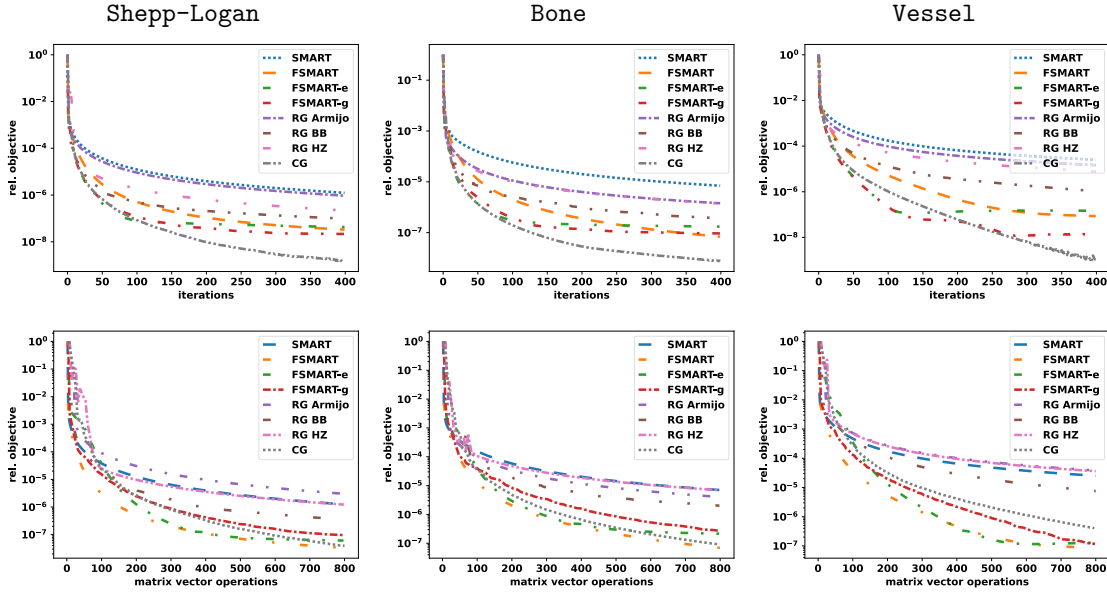



FIGURE 8. **Comparison of different algorithms on tomographic reconstruction** of phantoms shown in Figure 7. The plots show the relative decrease in objective values over iterations (**first row**) and in terms of costly operations (**second row**) for the different approaches detailed in Section 5.1. In these instances, the conjugate gradient method with DY mode (4.48d) from the Riemannian optimization framework exhibits the best performance, followed by FSMART-g from the Bregman framework.

A-posteriori certificates



Average Matrix Vector Operations




FIGURE 9. **A-posteriori certificates for tomographic reconstruction** are obtained from FSMART-e by tracking γ_k values across iterations, depicted in the (**left**) panel for all problem instances. It is notable that γ_k consistently drops to 1 in all instances. We investigated this drop-point for each instance and found it aligns with FSMART-e approaching the solution. The (**right**) panel illustrates the **average number of matrix-vector operations** for each algorithm across iterations and instances. SMART and FSMART constantly employ two matrix-vector operations per iteration by definition. However, other algorithms require more operations due to their utilization of line search or triangle gain adaption strategies.




FIGURE 10. **Original images and deblurring:** These plots display the images used for numerical evaluation in the first row and their degraded versions due to severe Gaussian blur with a mask of size 33×33 and $\sigma = 10$ in the second row. The rows below, show the resulting reconstructions after 1000 iterations from the algorithms SMART, FSMART, and CG (with DY mode).

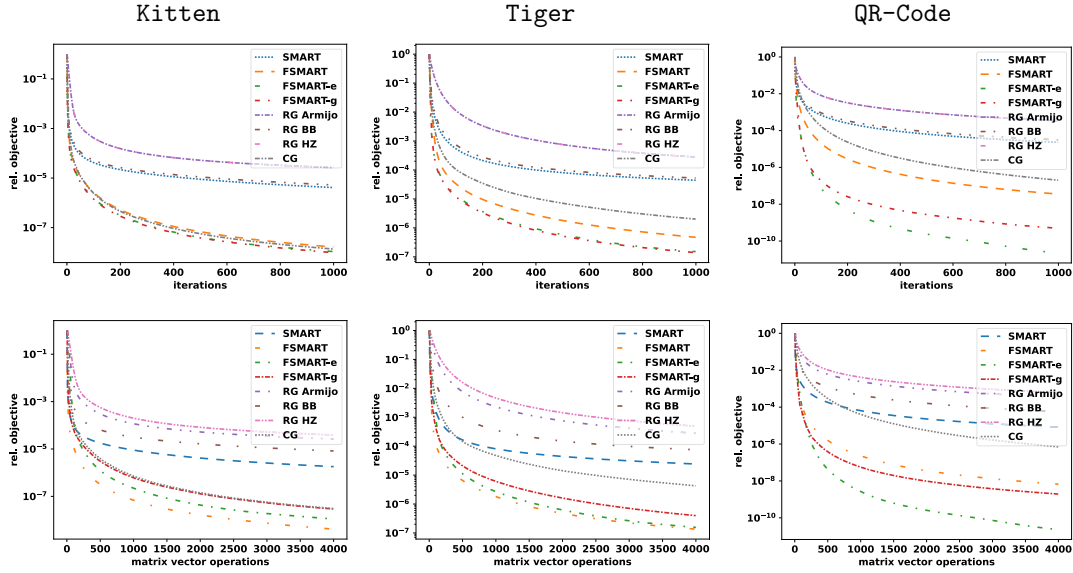


FIGURE 11. **Comparison of different algorithms on deblurring** of the instances shown in Figure 10. The plots illustrate the relative decrease in objective values over iterations (**first row**) and as a function of costly operations (**second row**) for the various approaches detailed in Section 5.1. Among the accelerated algorithms within the Bregman framework, FSMART, FSMART-e, and FSMART-g perform the best. Within the Riemannian optimization group, the conjugate gradient (CG) implemented with DY mode as the β_k -update rule (4.48d) the best results.

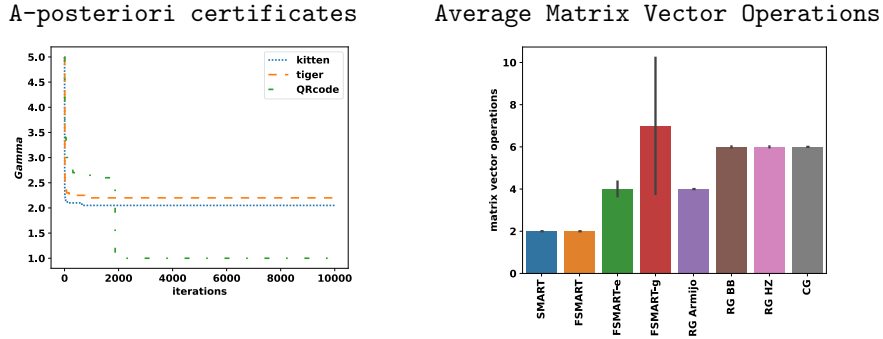


FIGURE 12. **A-posteriori certificates for deblurring** are obtained from FSMART-e by observing γ_k values over iterations, shown in the (**left**) panel for all problem instances. It is observed that γ_k does not always drop to 1 in all instances, especially in the two more challenging cases (kitten, tiger). In these instances, the algorithms are still far from the sought solutions even after 10000 iterations. **Average of matrix vector operations** are shown in the (**right**) panel for each algorithm over iterates and instances. By definition SMART, FSMART constantly employ two matrix vector operations per iteration. The other algorithms require more operations as they employ line search or triangle gain adaption strategies.

6. CONCLUSION

We studied recent acceleration techniques derived from Bregman proximal gradient methods for SMART and conducted numerical a-posteriori certification for several large-scale problems. Although we could not certify the $\mathcal{O}(1/k^2)$ rate, the accelerated variants of SMART proved to be remarkably efficient.

Moreover, we characterized SMART as a Riemannian gradient descent scheme on parameter manifolds induced by the Fisher-Rao geometry. This approach utilizes a retraction based on geodesics. We computed closed-form expressions for retractions and the corresponding vector transports for three manifolds as case studies. Additionally, we explored various Riemannian line search strategies such as geometric Hager-Zhang-type or Barzilai-Borwein line search, and investigated acceleration using Riemannian conjugate gradient.

Conditions ensuring convergence in geometric settings relate to Lipschitz continuity, extending the well-known gradient Lipschitz condition from Euclidean settings. This geometric L -smoothness condition does not precisely align with the scenarios studied in our paper, however. On the other hand, in Section 3.1, we showed how convergence of SMART, considered as Riemannian gradient descent iteration with a fixed step size, can be proven under relative L -smoothness. This result might pave the way for proving convergence in the geometric context even when geometric L -smoothness conditions are not met.

Our further work will explore possibilities for connecting the local triangle scaling property used by accelerated Bregman proximal gradient, to certify convergence rates with line search methods based on suitable retractions.

Acknowledgments

This work was funded by Deutsche Forschungsgemeinschaft (DFG) under Germany's Excellence Strategy EXC-2181/1 - 390900948 (the Heidelberg STRUCTURES Excellence Cluster).

REFERENCES

- [1] C. Byrne, Iterative image reconstruction algorithms based on cross-entropy minimization, *IEEE Trans. Image Process.* 2 (1993), 96-103.
- [2] I. Csiszar, I-divergence geometry of probability distributions and minimization problems, *Ann. Probab.* 3 (1976), 146-158.
- [3] S. Amari, A. Cichocki, Information geometry of divergence functions, *Bull. Polish Acad. Sci.* 58 (2010), 183-195.
- [4] G. Herman, Y. Censor, D. Gordon, R. Lewitt, Comment on a statistical model for positron emission tomography, *J. Amer. Statist. Assoc.* 80 (1985), 22-25.
- [5] M. Bertero, P. Boccacci, G. Desiderà, G. Vicidomini, Image deblurring with Poisson data: from cells to galaxies, *Inverse Probl.* 25 (2009), 123006.
- [6] J. Darroch, D. Ratcliff, Generalized iterative scaling for Log-linear models, *Ann. Math. Statist.* 43 (1972), 1470-1480.
- [7] A. Lent, Y. Censor, The primal-dual algorithm as a constraint-set-manipulation device, *Math. Program.* 50 (1991), 343-357.
- [8] Y. Censor, J. Segman, On block-iterative entropy maximization, *J. Inform. Optim. Sci.* 8 (1987), 275-291.
- [9] C. Byrne, *Iterative Optimization in Inverse Problems*, CRC Press, 2014.
- [10] S/ Petra, C. Schnörr, F. Becker, F. Lenzen, B-SMART: Bregman-based first-order algorithms for non-negative compressed sensing problems, *Proc. SSVM, LNCS.* 7893 (2013), 110-124.
- [11] A. Nemirovski, D. Yudin, *Problem complexity and method efficiency in optimization*, Wiley, 1983.

- [12] A. Beck, M. Teboulle, Mirror descent and nonlinear projected subgradient methods for convex optimization, *Oper. Res. Lett.* 31 (2003), 167-175.
- [13] G. Raskutti, S. Mukherjee, The information geometry of mirror descent, *IEEE Trans. Info. Theory* 61 (2015), 1451-1457.
- [14] M. Kahl, S. Petra, C. Schnörr, G. Steidl, M. Zisler, On the remarkable efficiency of the SMART iteration, *Proc. SSVN, LNCS 14009* (2023), 418-430.
- [15] F. Hanzely, P. Richtárik, L. Xiao, Accelerated Bregman proximal gradient methods for relatively smooth convex optimization, *Comput. Optim. Appl.* 79 (2021), 405-440.
- [16] D. Gutman, J. Peña, Perturbed fenchel duality and first-order methods, *Math. Program.* 198 (2023), 443-469.
- [17] A. Juditsky, J. Kwon, É. Moulines, Unifying mirror descent and dual averaging, *Math. Programm.* 199 (2022), 793-830.
- [18] U. Ghai, E. Hazan, Y. Singer, Exponentiated gradient vs. meets gradient descent, *Proc. Mach. Learning Res.* 117 (2020), 1-23.
- [19] W. Krichene, A. Bayen, P. Bartlett, Accelerated mirror descent in continuous and discrete time, *Advances in Neural Information Processing Systems 28 (NIPS 2015)*.
- [20] P. Absil, R. Mahony, R. Sepulchre, *Optimization Algorithms on Matrix Manifolds*, Princeton University Press, Princeton, 2008.
- [21] H. Oviedo, Global convergence of Riemannian line search methods with a Zhang-Hager-type condition, *Numer. Algor.* 91 (2022), 1183-1203.
- [22] H. Sakai, H. Sato, H. Iiduka, Global convergence of Hager-Zhang type Riemannian conjugate gradient method, *Appl. Math. Comput.* 441 (2023), 127685.
- [23] O. Ferreira, M. Louzeiro, L. Prudente, Gradient method for optimization on Riemannian manifolds with lower bounded curvature, *SIAM J. Optim.* 29 (2019), 2517-2541.
- [24] J. Jin, S. Sra, Understanding Riemannian acceleration via a proximal extragradient framework, *Proc. Mach. Learning Res.* 178 (2022), 1-39.
- [25] H. Zhang, S. Sra, First-order methods for geodesically convex optimization, arXiv:1602.06053, 2016.
- [26] S. Amari, H. Nagaoka, *Methods of Information Geometry*, Amer. Math. Soc. 2000.
- [27] R. Rockafellar, *Convex Analysis*, Princeton University Press, Princeton, 1997.
- [28] F. Alvarez, J. Bolte, O. Brahic, Hessian Riemannian gradient flows in convex programming, *SIAM J. Control Optim.* 43 (2004), 477-501.
- [29] Y. Censor, Row-action methods for huge and sparse systems and their applications, *SIAM Rev.* 23 (1981), 444-466.
- [30] Y. Censor, S.A. Zenios, *Parallel Optimization: Theory, Algorithms, and Applications*, Oxford University Press, New York/Oxford, 1997.
- [31] H. Bauschke, J. Borwein, Legendre functions and the method of random Bregman projections, *J. Convex Anal.* 4 (1997), 27-67.
- [32] G. Chen, M. Teboulle, Convergence analysis of a proximal-like minimization algorithm using Bregman functions, *SIAM J. Optim.* 3 (1993), 538-543.
- [33] L. Brown, *Fundamentals of Statistical Exponential Families with Applications in Statistical Decision Theory*, IMS Lecture Notes Monogr. Ser. 9, 1986. DOI: 10.1214/lnms/1215466757.
- [34] J. Jost, *Riemannian Geometry and Geometric Analysis*, Springer, 2017.
- [35] R. Rockafellar, R. Wets, *Variational Analysis* 3rd ed., Springer, 2010.
- [36] R. Rockafellar, Monotone operators and the proximal point algorithm, *SIAM J. Control Optim.* 14 (1976), 877-898.
- [37] Y. Censor, S. Zenios, Proximal minimization algorithm with D-functions, *J. Optim. Theory Appl.* 73 (1992), 451-464.
- [38] A. Auslender, M. Teboulle, Interior gradient and proximal methods for convex and conic optimization, *SIAM J. Optim.* 16 (2006), 697-725.
- [39] H. Bauschke, J. Bolte, M. Teboulle, A descent lemma beyond Lipschitz gradient continuity: First-order methods revisited and applications, *Math. Oper. Res.* 42 (2017), 330-348.
- [40] T. Cover, J. Thomas, *Elements of Information Theory* (2. ed.). Wiley, 2006.

- [41] C. Byrne, Iterative algorithms for deblurring and deconvolution with constraints, *Inverse Probl.* 14 (1998), 1455.
- [42] R. Dragomir, A. Taylor, A. D'Aspremont, J. Bolte, Optimal complexity and certification of Bregman first-order methods, *Math. Program.* 194 (2022), 41-83.
- [43] P. Tseng, On accelerated proximal gradient methods for convex-concave optimization, 2008. unpublished manuscript
- [44] H. Zhang, W. Hager, A nonmonotone line search technique and its application to unconstrained optimization, *SIAM J. Optim.* 14 (2004), 1043-1056.
- [45] L. Grippo, F. Lampariello, S. Lucidi, A nonmonotone line search technique for Newton's method, *SIAM J. Numer. Anal.* 23 (1986), 707-716.
- [46] M. Sutti, B. Vandereycken, Riemannian multigrid line search for low-rank problems, *SIAM J. Sci. Comput.* 43 (2021), A1803-A1831.
- [47] S. Müller, S. Petra, M. Zisler, Multi-level geometric optimization for regularised constrained linear inverse problems, *Pure Appl. Funct. Anal.* 8 (2022), 855–880.
- [48] J. Barzilai, J. Borwein, Two-point step size gradient methods, *IMA J. Numer. Anal.* 8 (1988), 141-148.
- [49] B. Iannazzo, M. Porcelli, The Riemannian Barzilai–Borwein Method with nonmonotone line search and the matrix geometric mean computation, *IMA J. Numer. Anal.* 38 (2017), 495-517.
- [50] J. Hu, X. Liu, Z. Wen, Y. Yuan, A brief introduction to manifold optimization, arXiv:1906.05450, 2019.
- [51] M. Raus, Master's thesis, Heidelberg University, 2023.
- [52] R. Berinde, P. Indyk, Sparse Recovery Using Sparse Random Matrices, Technical Report, MIT, 2008.
- [53] S. Petra, C. Schnörr, Average case recovery analysis of tomographic compressive sensing, *Linear Algebra Appl.* 441 (2014), 168-198.