

MULTILEVEL GEOMETRIC OPTIMIZATION FOR REGULARISED CONSTRAINED LINEAR INVERSE PROBLEMS

SEBASTIAN MÜLLER, STEFANIA PETRA, AND MATTHIAS ZISLER

Dedicated to Professor Lev Bregman on occasion of his 80th birthday

ABSTRACT. We present a geometric multilevel optimization approach that smoothly incorporates box constraints. Given a box constrained optimization problem, we consider a hierarchy of models with varying discretization levels. Finer models are accurate but expensive to compute, while coarser models are less accurate but cheaper to compute. When working at the fine level, multi-level optimisation computes the search direction based on a coarser model which speeds up updates at the fine level. Moreover, exploiting geometry induced by the hierarchy the feasibility of the updates is preserved. In particular, our approach extends classical components of multigrid methods like restriction and prolongation to the Riemannian structure of our constraints.

1. INTRODUCTION

1.1. Overview, Motivation. Multiscale representations have a long tradition in numerical analysis [16, 34] and in signal and image processing [13, 24], but much less so in numerical optimization if we ignore unconstrained convex quadratic optimization problems that boil down to solving a linear symmetric positive-definite system. In this paper, we consider constrained convex optimization problems of the form

$$(1.1) \quad \min_{y \in [0,1]^n} f(y)$$

with smooth objective functions $f: \mathbb{R}^n \rightarrow \mathbb{R}$. The main goal is to accelerate a first-order iterative algorithm for solving (1.1) by adopting a multilevel problem representation analogous to numerical multigrid, in order to compute efficiently descent directions for f at the original fine level using a much smaller number set of variables at coarser levels.

A key ingredient of the approach is to replace the feasible set $[0, 1]^n$ by the interior $(0, 1)^n$ that is turned into a Riemannian manifold (\mathcal{B}^n, g) using information geometry [3]. This geometry differs from the more common geometry underlying interior point methods [28] (see Section 1.2). Transfer of variables from fine to coarse levels and vice versa are performed by restriction and prolongation mappings, that are adapted to the geometry and whose differentials are used for pushing tangent

2020 *Mathematics Subject Classification.* 65K10, 49J40, 49M37, 68U10, 74P20, 90C06.

Key words and phrases. Multilevel geometric optimization, nonlinear prolongation and restriction, information geometry, Riemannian optimization.




FIGURE 1. **Left.** Incidence geometry of projection rays and cells indexed by I_n , that cover a subset in \mathbb{R}^d (here: $d = 2$ for illustration). The corresponding line integrals define the matrix A that linearly maps cell values y_i , $i \in I_n$ to projections b_j , $j \in \{1, \dots, p\}$. This gives rise to an underdetermined linear system $Ay = b$. The task is to recover y from b by solving (1.2) using the constraints and the regularizer to obtain a well-posed problem. **Right.** Reducing the spatial resolution yields a coarse level representation of the projection matrix, that enables to evaluate a surrogate of the objective function on the coarse level which does not involve the data vector b recorded on the fine level. This coarse level representation allows to compute descent directions on the fine level efficiently using a smaller subset of variables at the coarse level.

vectors between levels. Criteria for invoking coarse level computation and certifying resulting search directions are provided. Numerical results illustrate the effectiveness of our approach for the specific instance of (1.1)

$$(1.2) \quad \min_{y \in [0,1]^n} f(y), \quad f(y) := D_\varphi(Ay, b) + J(y),$$

where $b \in \mathbb{R}_{++}^p$ is a positive measurement vector, $A \in \mathbb{R}^{p \times n}$ is a nonnegativity-preserving linear mapping, J is a smoothed total variation regularizer [11, eq. (2.5)] and D_φ is the Bregman divergence with respect to the convex function

$$(1.3) \quad \varphi: \mathbb{R}^n \rightarrow \mathbb{R}, \quad y \mapsto \varphi(y) = \langle y, \log y \rangle - \langle \mathbf{1}, y \rangle.$$

See Figure 1 and the caption for further explanation.

Although we consider the specific problem (1.1) in this paper, we believe that our approach generalizes to other constrained convex programs, analogous to the way how open convex parameter sets of probability distributions are turned into statistical manifolds [3, 23].

1.2. Related Work. *Multilevel optimization* has been introduced by Nash [25]. Multilevel smooth unconstrained optimization is the setting of [15, 25, 37], that includes trust-region and line search optimization. The first multilevel optimization method that covers the nonsmooth convex composite case was introduced in [29]. It is a multilevel version of the well-known proximal gradient method. An accelerated variant was proposed in [18]. Both approaches in [18, 29] can only handle smoothable functions in the sense of [6]. A multilevel optimization approach for box-constraints

(non-smoothable functions) was proposed in [21]. The design of the coarse model in [21] was inspired by the trust-region method in [15] and the work reported in [20], and was adopted also in [30].

Multilevel approaches for general *non-convex* optimization problems, as seen in [37], have not been extended to the non-Euclidean context, so there are currently no theoretical results available. Recently, a geometric multilevel optimization approach has been proposed by [33] for the specific case of low-rank matrix manifolds. Our approach differs in that we focus on information geometry for constructing multilevel problem representations, which should be applicable to various constrained programs beyond the case study (1.2) considered in this paper.

Since Nesterov's work [26], *accelerated* first-order convex optimization has become an active research field [14, 31, 35, 38]. Due to the performance of Nesterov's discrete scheme and the work [32], a major line of research concerns the understanding and generation of accelerated discrete scheme using continuous-time ODEs, where an adaptive 'friction term' causes acceleration. In our approach, acceleration is solely achieved by employing a multilevel problem representation. Acceleration of the first-order Riemannian dynamics is an interesting problem of future research.

The use of *geometry* for convex programming has a long history [1, 5, 28]. The commonly adopted interior point geometry is based on Hessian metrics generated by self-concordant barrier functions, due to the provable optimality in connection with Newton-like second-order optimization [27], see e.g. [36] for recent related work. Closer to our work is the comprehensive paper by [2] where Hessian metrics generated by convex functions of Legendre type are studied for a class of convex programs that include affine subspace constraints.

In the present paper, we also consider such a Riemannian metric on a simpler structured bounded open convex feasible set, in order to focus on multilevel representation and accelerated first-order optimization that copes with large problem sizes. This necessitates, in particular, to devise restriction and prolongation operators not only for points on the manifold but also for tangent vectors. To this end, we employ information geometry in order to design smooth nonlinear mappings based on geometric averaging that can be efficiently computed in closed form. Numerical experiments demonstrate that our method outperforms the recent state-of-the-art method [17].

1.3. Contribution, Organization. We introduce basic notation and briefly recall required properties of Bregman divergences in Section 2. Section 3 summarizes the basic scheme of two-level optimization in Euclidean spaces. The core part of this paper, Section 4, generalizes this scheme to a Riemannian setting. In particular, information geometry is employed in a way that indicates how other convex programs could be handled in the same way, in principle. While we indicate multilevel extensions in Sections 3.2 and 4.7, respectively, we mainly focus on the core problem in this paper, that is two-level geometric optimization. Numerical experiments are reported in Section 5 for a range of problem instances and compared to a recent state-of-the-art method [17]. We conclude in Section 6.

2. PRELIMINARIES

2.1. **Notation.** We set $[n] = \{1, 2, \dots, n\}$ for $n \in \mathbb{N}$ and

$$(2.1a) \quad \mathcal{B} = (0, 1),$$

$$(2.1b) \quad \mathcal{B}_n = \mathcal{B} \times \dots \times \mathcal{B} \subset \mathbb{R}_{++}^n. \quad (\text{open } n\text{-box})$$

$\mathbb{1} = (1, 1, \dots, 1)^\top$ denotes the one-vector with dimension depending on the context. Sometimes we indicate the dimension by a subscript, e.g. $\mathbb{1}_n \in \mathbb{R}^n$. The componentwise application of functions like \log is simply denoted as

$$(2.2) \quad \log y = (\log y_1, \dots, \log y_n)^\top, \quad y \in \mathbb{R}_{++}^n.$$

Likewise, the componentwise multiplication (Hadamard product) and division of vectors y, y' is written as

$$(2.3) \quad y \cdot y' = (y_1 y'_1, \dots, y_n y'_n)^\top, \quad \frac{y}{y'} = \left(\frac{y_1}{y'_1}, \dots, \frac{y_n}{y'_n} \right)^\top.$$

The Euclidean vector inner product is denoted by $\langle y, y' \rangle = y^\top y'$. ∂f denotes the gradient of a differentiable function $f: \mathbb{R}^n \rightarrow \mathbb{R}$, whereas $\text{grad } f$ denotes the Riemannian gradient of a function $f: \mathcal{M} \rightarrow \mathbb{R}$ on a manifold (\mathcal{M}, g) with metric g .

Components y_j of vectors $y \in \mathcal{B}_n$ are indexed by $[n]$ that we specifically denote in this context by

$$(2.4) \quad I_n = [n]. \quad (n\text{-grid})$$

The notion “ n -grid” reminds of the fact that indices j typically refer to locations in an underlying physical domain (cf. Section 1). In this sense, we will synonymously refer to j as ‘point’ and ‘index’. A disjoint partition

$$(2.5) \quad I_n = I_m \dot{\cup} I_m^c, \quad m < n$$

defines a coarse grid I_m indexing coarse grid vectors $x \in \mathcal{B}_m$ and fine grid points indexed by the complement I_m^c . I_n is called fine grid and vectors $y \in \mathcal{B}_n$ are called fine grid vectors.

Regarding problem (1.2), we summarize our assumptions.

$$(A) \quad \begin{aligned} & b_i > 0, \quad i = 1, \dots, p & (1) \\ & A_{ij} \geq 0, \quad i = 1, \dots, p, j = 1, \dots, n & (2) \\ & \forall y > 0: \quad \text{supp}(Ay) = p, & (3) \\ & \text{rank } A = \min\{p, n\}, & (4) \end{aligned}$$

where $\text{supp}(y)$ denotes the number of nonzero components of a vector y .

2.2. **Bregman Divergences.**

Definition 2.1 (Bregman divergence [8], [10, Section 2], [7, Def. 7.6.1]). Let E be a Euclidean space and $\phi: E \rightarrow (-\infty, +\infty]$ be a convex function of Legendre type with non-empty domain $\text{dom } \phi \subseteq E$ [7, Def. 7.1.1]. The Bregman divergence corresponding to ϕ is defined by

$$(2.6) \quad D_\phi: E \times \text{int dom } \phi \rightarrow [0, \infty]: \quad (x, y) \mapsto \phi(x) - \phi(y) - \langle \partial\phi(y), x - y \rangle.$$

Bregman divergences have the properties $D_\phi(x, y) \geq 0$ and $D_\phi(x, y) = 0$ if and only if $x = y$. A basic example is the Kullback-Leibler (KL) divergence

$$(2.7a) \quad D_\varphi(y, y') = \left\langle \mathbb{1}, y \cdot \log \frac{y}{y'} + y' - y \right\rangle$$

corresponding to the function

$$(2.7b) \quad \varphi(y) = \langle y, \log y \rangle - \langle \mathbb{1}, y \rangle, \quad y \in \mathbb{R}_+^n.$$

Lemma 2.2 ([12, Lem 3.1]). *Let $S \subset \mathbb{R}^n$ be an open set with closure \bar{S} , and let $\phi: \bar{S} \rightarrow \mathbb{R}$ be a convex function as in Definition 2.1. Then, for any three points $a, b \in S$ and $c \in \bar{S}$, the identity*

$$(2.8) \quad D_\phi(c, a) + D_\phi(a, b) - D_\phi(c, b) = \langle \partial\phi(b) - \partial\phi(a), c - a \rangle$$

holds.

Now we consider the data term of (1.2)

$$(2.9) \quad h(y) := D_\varphi(Ay, b)$$

with the KL-divergence D_φ given by (2.7) and gradient

$$(2.10) \quad \partial h(y) = A^\top \log \frac{Ay}{b}.$$

Lemma 2.3. *Assume $b \in \mathbb{R}_{++}^p$, $y \in \mathbb{R}_+^n$, $y' \in \mathbb{R}_{++}^n$ and $Ay' \in \mathbb{R}_{++}^p$. Then the Bregman divergence corresponding to the function h given by (2.9) reads*

$$(2.11) \quad D_h(y, y') = D_\varphi(Ay, Ay').$$

Proof. By definition (2.6), we have

$$(2.12) \quad \begin{aligned} D_h(y, y') &= h(y) - h(y') - \langle \partial h(y'), y - y' \rangle \\ &\stackrel{(2.9)}{=} D_\varphi(Ay, b) - D_\varphi(Ay', b) - \langle \partial\varphi(Ay') - \partial\varphi(b), Ay - Ay' \rangle. \end{aligned}$$

Substituting $\phi \leftarrow \varphi$, $b \leftarrow b$, $a \leftarrow Ay'$ and $c \leftarrow Ay$ in Lemma 2.2 yields

$$D_\varphi(Ay, Ay') + D_\varphi(Ay', b) - D_\varphi(Ay, b) = \langle \partial\varphi(b) - \partial\varphi(Ay'), Ay - Ay' \rangle.$$

Comparing this equation and (2.12) is equivalent to (2.11). □

3. TWO-GRID EUCLIDEAN OPTIMIZATION

In this section, we summarize prior work on two-level optimization in Euclidean spaces [25, 37]. Our main contribution in this paper, the generalization to a Riemannian setting, is described in Section 4.

3.1. Unconstrained Two-Grid Euclidean Optimization. As motivated in Section 1, we assume a generic objective function f to be given that can be evaluated at different discretization levels. In this section, we consider two discretization levels called fine grid and coarse grid, respectively, and use the following notation.

$$(3.1a) \quad \text{fine grid variable: } y \in \mathbb{R}^n, \quad \text{current iterate: } y_0 \in \mathbb{R}^n,$$

$$(3.1b) \quad \text{coarse grid variable: } x \in \mathbb{R}^m, \quad \text{restriction of } y_0: x_0 = Ry_0 \in \mathbb{R}^m.$$

Here, the restriction map R and its counterpart, the prolongation map P ,

$$(3.2a) \quad R: \mathbb{R}^m \rightarrow \mathbb{R}^n,$$

$$(3.2b) \quad P: \mathbb{R}^n \rightarrow \mathbb{R}^m,$$

transfer fine and coarse grid variables to the coarse and fine grid, respectively, typically via linear interpolation or simple injection as in classical multigrid methods [34].

We assume that the *argument determines* if the objective function f is evaluated on the fine grid or the coarse grid, respectively. Thus, in view of (3.1), $f(y)$ implies $f: \mathbb{R}^n \rightarrow \mathbb{R}$ whereas $f(x)$ implies $f: \mathbb{R}^m \rightarrow \mathbb{R}$, with $m < n$.

We next describe a two grid cycle that computes an update y_1 of y_0 at a fine grid. This is done either along a search direction efficiently obtained from a surrogate objective function defined on the coarse grid using a much smaller number of variables: *coarse correction*. Alternatively, whenever a coarse correction is not effective, the update is determined using a standard local objective function approximation defined on the fine grid: *fine correction*.

The general two grid approach is summarized as Algorithm 1. It uses a representation ψ of the objective function, specified in Section 3.1.1, that is *not* equal to the evaluation of the objective $f: \mathbb{R}^m \rightarrow \mathbb{R}$ on the coarse grid.

Algorithm 1: Two Level Euclidean Optimization

input : initial point y_0 , mappings R, P (cf. (3.2)), coarse grid model ψ (cf. (3.3)).

output: approximate global minimizer y^* .

1 Initialization: $k = 0$

2 **repeat**

3 **if** *condition to invoke the coarse model is satisfied at* y_0 **then**

4 set $x_0 = Ry_0$

5 determine a coarse grid iterate x such that $\psi(x) < f(x_0)$

6 set $d = P(x - x_0)$

7 find $\alpha > 0$ such that $f(y_0 + \alpha d) < f(y_0)$ /* line search */

8 update $y_0 \leftarrow y_0 + \alpha d$

9 **else**

10 apply one iteration of a monotone fine level algorithm to find y_1 with $f(y_1) < f(y_0)$ and update $y_0 \leftarrow y_1$

11 increment $k \leftarrow k + 1$

12 **until** *a stopping rule is met*.

The condition in line 3 is specified in Section 3.1.3. The role of the coarse grid model ψ is further discussed in the following section.

3.1.1. *Coarse Grid Model.* The core ingredient of two level optimization is a *coarse grid model*, that is a coarse grid representation of the fine grid problem in terms of the objective function f evaluated at the coarse grid, the current iterate y_0 and its restriction $x_0 = Ry_0$ to the coarse grid. It reads

$$(3.3a) \quad \psi(x) = f(x) - \langle x - x_0, \kappa(x_0, y_0) \rangle,$$

with the linear correction term

$$(3.3b) \quad \kappa(x_0, y_0) = \partial f(x_0) - R\partial f(y_0).$$

The rationale behind the coarse grid model is to determine efficiently a gradient-like descent direction on the *fine* grid using a much smaller number of variables on the coarse grid, as first proposed in [25]. The linear correction term (3.3a) guarantees that a fine grid stationary point is also a coarse grid stationary point. More precisely, the gradient of ψ equals

$$(3.4a) \quad \partial\psi(x) = \partial f(x) - \kappa(x_0, y_0)$$

$$(3.4b) \quad = \partial f(x) - \partial f(x_0) + R\partial f(y_0)$$

$$(3.4c) \quad = \partial (f(x) - f(x_0)) + R\partial f(y_0),$$

and hence the *first-order coherence condition* [25]

$$(3.5) \quad \partial\psi(x_0) = R\partial f(y_0),$$

holds at $x = x_0$.

3.1.2. *Descent Directions Based on the Coarse Grid Model.* We next show that a sufficient decrease of ψ given by (3.3) by updating the smaller number of coarse grid variables x determines a descent direction for the fine grid objective. To this end, we rewrite ψ in terms of the Bregman divergence corresponding to the coarse grid objective f ,

$$(3.6a) \quad \psi(x) = f(x) - \langle x - x_0, \kappa(x_0, y_0) \rangle$$

$$(3.6b) \quad = f(x) - \langle x - x_0, \partial f(x_0) - R\partial f(y_0) \rangle$$

$$(3.6c) \quad = f(x) - f(x_0) - \langle x - x_0, \partial f(x_0) - R\partial f(y_0) \rangle + f(x_0)$$

$$(3.6d) \quad = D_f(x, x_0) + \langle x - x_0, R\partial f(y_0) \rangle + f(x_0).$$

Now, assuming R and P satisfy the Galerkin condition

$$(3.7) \quad R = P^\top$$

that is a standard condition in multigrid literature [34], we get

$$(3.8) \quad \psi(x) = D_f(x, x_0) + \langle P(x - x_0), \partial f(y_0) \rangle + f(x_0).$$

Hence, whenever

$$(3.9) \quad \psi(x) < f(x_0) = \psi(x_0)$$

at a coarse iterate x and the coarse objective f is convex, then the vector

$$(3.10) \quad d = P(x - x_0)$$

is a descent direction of the fine grid objective f , since (3.9) and $D_f \geq 0$ implies

$$(3.11) \quad \langle P(x - x_0), \partial f(y_0) \rangle < 0$$

by (3.8). To find such a coarse grid iterate x in practice, one typically employs few iterations of a monotone algorithm for minimizing ψ at the coarse level.

Remark 3.1. The coarse grid model rewritten in the form (3.8) reveals that it incorporates both second-order information of the coarse grid objective (first term) and first-order information of the fine grid objective (second term). Indeed, ignoring the constant $f(x_0)$ and assuming f is twice differentiable, the first term can be equivalently rewritten as

$$(3.12) \quad D_f(x, x_0) = \frac{1}{2} \langle x - x_0, \partial^2 f(\tilde{x})(x - x_0) \rangle$$

for some $\tilde{x} \in \{(1 - t)x + tx_0 : t \in [0, 1]\}$, where $\partial^2 f(\tilde{x})$ is the Hessian of f at \tilde{x} .

3.1.3. Coarse Grid Correction Criteria. We employ the coarse grid model whenever the criteria

$$(3.13) \quad \|R\partial f(y_0)\| \geq \eta \|\partial f(y_0)\| \quad \text{and} \quad \|y_0 - y_c\| > \varepsilon$$

are satisfied, where $\eta \in (0, \min(1, \|R\|))$, $\varepsilon \in (0, 1)$ and y_c is the last point that initiated a coarse grid correction.

- (1) The first condition in (3.13) is adopted from [37] and is widely used in the multilevel optimization literature. We discuss this first condition in view of the first-order coherence condition (3.5).

If the left-hand side $\|R\partial f(y_0)\|$ is too small at the current iterate y_0 , then there is nothing to optimize on the coarse grid – as $\|\partial\psi(x_0)\|$ is small too. The entire first condition says that y_0 is non-optimal in a twofold sense: $\|\partial f(y_0)\| > 0$ (otherwise both sides vanish) and on a coarser scale f is sufficiently non-flat as well such that invoking the coarse grid model pays off.

- (2) Violation of the second condition says that the current iterate y_0 is too close to a point that already initiated a coarse grid correction.

3.1.4. Application to a Regularized Inverse Problem. In order to apply two-level optimization to the regularized inverse problem (1.2), the coarse grid model function ψ given by (3.8) has to be computed. Similar to the evaluation of the objective function at both levels, we assume that the operator A can be directly evaluated at both levels. For example, in discrete tomography, the projection matrix A represents the incidence relation of projection rays and cells centered at the grid points, and can be evaluated on every grid.

It suffices to consider the Bregman divergence since the remaining term only involves fine grid variables (we ignore the constant $f(x_0)$).

Lemma 3.2. *The Bregman divergence corresponding to the objective function (1.2), evaluated at the coarse level, reads*

$$(3.14) \quad D_f(x, x_0) = D_\varphi(Ax, Ax_0) + \lambda D_J(x, x_0).$$

Proof. Lemma 2.3 shows $D_h(x, x_0) = D_\varphi(Ax, Ax_0)$ with h given by (2.9). The result follows from linearity of the Bregman divergence with respect to the inducing function, i.e. $D_{h+\lambda J}(x, x_0) = D_h(x, x_0) + \lambda D_J(x, x_0)$. \square

Remark 3.3. Lemma 3.2 shows, in particular, that the specific ray geometry (depicted in Fig. 1) at the coarse level can be selected independently of the ray geometry at the fine level, since the projection data b drop out and hence have not to be transferred between levels.

3.2. Multilevel Extension. We introduce the notation

$$(3.15a) \quad x^{(\ell)} \in \mathbb{R}^{n_\ell}$$

for the variables at levels $\ell \in \{0, 1, 2, \dots\}$, where

$$(3.15b) \quad x^{(0)} := x \in \mathbb{R}^n =: \mathbb{R}^{n_0}$$

are the original fine grid variables. The components $x_i^{(\ell)}$ of these vectors are indexed by the nested index sets

$$(3.15c) \quad i \in I_{n_\ell} \subset I_{n_{\ell-1}} \subset \dots \subset I_n = I_{n_0}.$$

Accordingly, the coarse grid model (3.3) applied at level ℓ reads

$$(3.16a) \quad \psi(x^{(\ell)}) = f(x^{(\ell)}) - \langle x^{(\ell)} - x_0^{(\ell)}, \kappa(x_0^{(\ell)}, x_0^{(0)}) \rangle,$$

$$(3.16b) \quad \kappa(x_0^{(\ell)}, x_0^{(0)}) = \partial f(x_0^{(\ell)}) - R_{(\ell)}^{(0)} \partial f(x_0^{(0)}),$$

with $x_0^{(\ell)} = R_{(\ell)}^{(0)} x_0^{(0)}$. Substituting $\kappa(x_0^{(\ell)}, x_0^{(0)})$ into $\psi(x^{(\ell)})$, we have

$$(3.17) \quad \begin{aligned} \psi(x^{(\ell)}) - f(x_0^{(\ell)}) &= f(x^{(\ell)}) - f(x_0^{(\ell)}) - \langle x^{(\ell)} - x_0^{(\ell)}, \partial f(x_0^{(\ell)}) - R_{(\ell)}^{(0)} \partial f(x_0^{(0)}) \rangle \\ &= D_f(x^{(\ell)}, x_0^{(\ell)}) + \langle P_{(0)}^{(\ell)}(x^{(\ell)} - x_0^{(\ell)}), \partial f(x_0^{(0)}) \rangle \\ &\stackrel{\text{Lem. 2.3}}{=} D_\varphi(A^{(\ell)}x^{(\ell)}, A^{(\ell)}x_0^{(\ell)}) + \lambda D_J(x^{(\ell)}, x_0^{(\ell)}) + \langle P_{(0)}^{(\ell)}(x^{(\ell)} - x_0^{(\ell)}), \partial f(x_0^{(0)}) \rangle, \end{aligned}$$

where $A^{(\ell)} \in \mathbb{R}^{p_\ell \times n_\ell}$ is the discretization of the linear operator at level ℓ , $R_{(\ell)}^{(0)} = (P_{(0)}^{(\ell)})^\top$ and $P_{(0)}^{(\ell)}$ prolongates vectors at level ℓ to level 0 by interpolation, e.g. as through the composition $P_{(0)}^{(\ell)} = P_{(\ell-1)}^{(\ell)} \circ \dots \circ P_{(0)}^{(1)}$ of interpolation operators between adjacent levels.

We note that the data vector $b \in \mathbb{R}^{p_0} := \mathbb{R}^p$ in (1.2) is only required on the finest (original) level. In addition, comparing (3.8) and (3.17) makes evident that the discussion of (3.8) applies also to (3.17): an iterate $x^{(\ell)}$ satisfying $\psi(x^{(\ell)}) - f(x_0^{(\ell)}) < 0$ yields a descent direction $P_{(0)}^{(\ell)}(x^{(\ell)} - x_0^{(\ell)})$ at level 0. Clearly, the criterion analogous to (3.13) (with $R_{(\ell)}^{(0)}$ in place of R) for employing a coarse correction will be less likely satisfied at level ℓ than a finer level $\ell' < \ell$. On the other hand, if it is satisfied, then determining a descent direction is computationally (very) cheap.

Optimizing the costs for invoking coarse grid models vs. saving computations through evaluating these models with much smaller numbers of variables, is left for future work.

4. MULTILEVEL GEOMETRIC OPTIMIZATION

4.1. Riemannian Geometry of the Box. In this section, we represent the open n -box \mathcal{B}_n (2.1b) as a Riemannian manifold. To this end, we turn the open interval (2.1a) into a manifold (\mathcal{B}, g) with metric g and define (\mathcal{B}_n, g) as the corresponding product manifold.

In order to specify (\mathcal{B}, g) , we apply basic information geometry [3]. Let points $\eta \in \mathcal{B}$ parametrize the Bernoulli distribution

$$(4.1) \quad p(z; \eta) = \eta^z (1 - \eta)^{1-z}$$

of a binary random variable $Z \in \{0, 1\}$. Then the metric tensor of the Fisher-Rao metric is simply a scalar function of η given by

$$(4.2) \quad G(\eta) = 4 \sum_{z \in \{0,1\}} \left(\frac{d}{d\eta} \sqrt{p(z; \eta)} \right)^2 = \frac{1}{1 - \eta} + \frac{1}{\eta} = \frac{1}{\eta(1 - \eta)}.$$

In view of (2.1b), this naturally extends to the metric g on $T\mathcal{B}_n$ in terms of the diagonal matrix

$$(4.3) \quad G_n(y) = \text{Diag} \left(\frac{1}{y_1(1 - y_1)}, \dots, \frac{1}{y_n(1 - y_n)} \right)$$

as metric tensor, using again the symbol ‘ G ’ for simplicity. We naturally identify

$$(4.4) \quad T_y \mathcal{B}_n \cong \mathbb{R}^n, \quad \forall y \in \mathcal{B}_n$$

and denote this metric interchangeably by

$$(4.5) \quad g_y(v, v') = \langle v, v' \rangle_y = \langle v, G_n(y)v' \rangle, \quad \forall v, v' \in T_y \mathcal{B}_n.$$

4.2. Retraction and its Inverse. Retractions [1, Def. 4.1.1] and their inverses are basic ingredients of first-order optimization algorithms on Riemannian manifolds. The main motivation is to replace the exponential map with respect to the metric (Levi Civita) connection by an approximation that can be efficiently evaluated or even in closed form. Below, we compute the exponential map with respect to the e -connection of information geometry [3] and show subsequently that it is a retraction.

Proposition 4.1. *The exponential maps on \mathcal{B} resp. \mathcal{B}_n with respect to the e -connection are given by*

$$(4.6a) \quad \exp: \mathcal{B} \times T\mathcal{B} \rightarrow \mathcal{B}, \quad \exp_\eta(tv) = \frac{\eta e^{t \frac{v}{\eta(1-\eta)}}}{1 - \eta + \eta e^{t \frac{v}{\eta(1-\eta)}}, \quad t > 0$$

$$(4.6b) \quad \exp: \mathcal{B}_n \times T\mathcal{B}_n \rightarrow \mathcal{B}_n, \quad \exp_y(tv) = \left(\exp_{y_j}(tv_j) \right)_{j \in I_n}$$

with inverses

$$(4.7a) \quad \exp^{-1}: \mathcal{B} \times \mathcal{B} \rightarrow T\mathcal{B}, \quad \exp_\eta^{-1}(\eta') = \eta(1 - \eta) \log \frac{(1 - \eta)\eta'}{\eta(1 - \eta')}$$

$$(4.7b) \quad \exp^{-1}: \mathcal{B}_n \times \mathcal{B}_n \rightarrow T\mathcal{B}_n, \quad \exp_y^{-1}(y') = \left(\exp_{y_j}^{-1}(y'_j) \right)_{j \in I_n}.$$

Proof. It suffices to show (4.6a) from which (4.7a) follows from an elementary calculation. (4.6b) and (4.7b) result from the componentwise application of these maps.

A key concept of information geometry is to replace the metric connection by a pair of connections that are dual to each other with respect to the Riemannian metric g [3, Section 3.1]. In particular, under suitable assumptions, the parameter space of a probability distribution becomes a Riemannian manifold that is dually flat, i.e. two distinguished coordinate systems (the so-called m- and e-coordinates) exists with affine geodesics. We consider the simple case (\mathcal{B}, g) .

First, we rewrite the distribution (4.1) as distribution of the exponential family [9]

$$(4.8a) \quad p(z; \theta) = \exp(z\theta - \psi(\theta))$$

$$(4.8b) \quad \theta = \theta(\eta) = \log \frac{\eta}{1-\eta}, \quad \eta = \eta(\theta) = \frac{e^\theta}{1+e^\theta}$$

with exponential parameter θ and log-partition function

$$(4.9) \quad \psi(\theta) = \log Z(\theta), \quad Z(\theta) = 1 + e^\theta$$

that is convex and of Legendre type [7, Def. 7.1.1]. The aforementioned distinguished two coordinates are η and θ with affine geodesics

$$(4.10a) \quad t \mapsto \eta_v(t) = \eta + tv \in \mathcal{B},$$

$$(4.10b) \quad t \mapsto \theta_u(t) = \theta + tu \in \mathbb{R}.$$

Note that unlike η, v , the coordinate θ and the tangent u are unconstrained. Using (4.8b), the e-geodesic reads

$$(4.11) \quad \eta(\theta_u(t)) = \frac{e^{\theta_u(t)}}{1 + e^{\theta_u(t)}} = \frac{e^\theta e^{tu}}{1 + e^\theta e^{tu}} \in \mathcal{B}.$$

It remains to replace θ by $\eta = \eta(\theta)$ due to (4.8b) and u by η and v , respectively. The latter results from substituting the affine geodesic (4.10a) into $\theta = \theta(\eta)$ and to express u by

$$(4.12a) \quad u = u(\eta, v) = \left. \frac{d}{dt} \theta(\eta_v(t)) \right|_{t=0} = \left. \frac{d}{dt} \theta(\eta + tv) \right|_{t=0}$$

$$(4.12b) \quad = \left. \frac{d}{dt} \log \frac{\eta + tv}{1 - \eta - tv} \right|_{t=0} = \frac{v}{\eta(1 - \eta)}.$$

Substituting into (4.11) yields

$$(4.13) \quad \exp_\eta(tv) := \frac{e^{\theta(\eta)} e^{tu(\eta,v)}}{1 + e^{\theta(\eta)} e^{tu(\eta,v)}} = \frac{\eta e^{\frac{t v}{\eta(1-\eta)}}}{1 - \eta + \eta e^{\frac{t v}{\eta(1-\eta)}}}$$

which is (4.6a). □

Remark 4.2 (g is a Hessian metric). The dual nature of the exponential parametrization (4.8) is also highlighted by recovering the metric tensor (4.2) as Hessian metric from the potential $\phi(\eta)$ that is conjugate to the log-partition function (4.9),

$$(4.14) \quad \phi(\eta) = \psi^*(\eta) = \eta \log \eta + (1 - \eta) \log(1 - \eta),$$

that is

$$(4.15) \quad \phi''(\eta) = G(\eta) = \frac{1}{\eta(1 - \eta)}.$$

The dual coordinate and potential yield the inverse metric tensor

$$(4.16) \quad \psi''(\theta) = \frac{e^\theta}{(1 + e^\theta)^2} = \frac{1}{G(\eta)} \Big|_{\eta=\eta(\theta)}.$$

The following formulas will be used later on.

Lemma 4.3 (Differential of \exp and \exp^{-1}). *The differentials of the mappings \exp_η and \exp_η^{-1} at u and η' , respectively, are given by*

$$(4.17a) \quad d \exp_\eta(u): T_\eta \mathcal{B} \rightarrow T_{\eta'} \mathcal{B},$$

$$(4.17b) \quad v \mapsto d \exp_\eta(u)v = \frac{e^{\frac{u}{\eta(1-\eta)}}}{(1 - \eta + \eta e^{\frac{u}{\eta(1-\eta)}})^2} v = \frac{\eta'(1 - \eta')}{\eta(1 - \eta)} v, \quad \eta' = \exp_\eta(u)$$

$$(4.17c) \quad d \exp_\eta^{-1}(\eta'): T_{\eta'} \mathcal{B} \rightarrow T_\eta \mathcal{B},$$

$$(4.17d) \quad v' \mapsto d \exp_\eta^{-1}(\eta')v' = \frac{\eta(1 - \eta)}{\eta'(1 - \eta')} v'.$$

Proof. By direction computation,

$$(4.18) \quad d \exp_\eta(u)v = \frac{d}{dt} \exp_\eta(u + tv) \Big|_{t=0}$$

using (4.6a) and similarly for (4.17d) using (4.7a). Clearly, $d(\exp_\eta^{-1}(\eta')) = (d \exp_\eta(u))^{-1}$. □

Retractions provide a proper class of surrogate mappings for replacing the canonical exponential map corresponding to the metric connection.

Proposition 4.4 (\exp is a retraction). *The mapping $\exp: T\mathcal{B} \rightarrow \mathcal{B}$ is a retraction in the sense of [1, Def. 4.1.1.].*

Proof. We check the two criteria that characterize retractions. First,

$$(4.19) \quad \exp_\eta(0) \stackrel{(4.6a)}{=} \eta, \quad \forall \eta \in \mathcal{B}.$$

Second, the so-called local rigidity condition

$$(4.20) \quad d \exp_\eta(0) \stackrel{(4.17b)}{=} 1 = \text{id}_{T_\eta \mathcal{B}}, \quad \forall \eta \in \mathcal{B}$$

holds as well. □

Remark 4.5 (Vector transport). In the same way as retractions may properly replace the canonical exponential map, the class of vector transport mappings

$$(4.21) \quad \Pi: T\mathcal{B} \times T\mathcal{B} \rightarrow T\mathcal{B}, \quad (u_\eta, v_\eta) \mapsto \Pi_{u_\eta}(v_\eta)$$

may properly replace the canonical parallel transport of tangent vectors [1, Section 8.1]. Given a retraction, like in Proposition 4.4, a vector transport map can be conveniently obtained by differentiation [1, Section 8.1.2],

$$(4.22) \quad \Pi_u(v) = d \exp_\eta(u)v \stackrel{(4.17b)}{=} \frac{\eta'(1 - \eta')}{\eta(1 - \eta)} v, \quad \eta' = \exp_\eta(u).$$

The above proof of Proposition 4.1 revealed the representation of tangents v and u of (4.10) in the form

$$(4.23a) \quad T_\eta^{(m)} = \{v \in \mathbb{R} : \eta + v \in \mathcal{B}\}, \quad \eta \in \mathcal{B},$$

$$(4.23b) \quad T_\eta^{(e)} = \left\{u = \frac{v_u}{\eta(1-\eta)} : v_u \in \mathbb{R}\right\}, \quad \eta \in \mathcal{B}.$$

Thus, in view of (4.23b), the vector transport (4.22) may be interpreted as sending $u = \frac{v}{\eta(1-\eta)} \in T_\eta^{(e)}$ to $u' = \frac{v'}{\eta'(1-\eta')} \in T_{\eta'}^{(e)}$ with $v' = \Pi_u(v)$.

4.3. Line Search. We now adopt the simplest approach to optimizing a differentiable function $f: \mathcal{B}_n \rightarrow \mathbb{R}$ on the manifold (\mathcal{B}_n, g) , by continuously translating a current iterate y in the direction of steepest descent, $-\text{grad } f(y)$. With the retraction calculated in Proposition 4.1, we will now use Armijo line search [1, Def. 4.2.2] to move within our Riemannian manifold. For minimizing coarse level approximations, see Section 4.6, we will employ the same techniques.

In particular, given f , a point $y \in \mathcal{B}_n$ and $\eta := -\text{grad } f(y) \in T\mathcal{B}_n$, the goal is to determine an update of y along the steepest descent direction η . To this end, we set

$$(4.24) \quad y_+ = f(\exp_y(\alpha\eta)), \quad \alpha \in \mathbb{R}_+,$$

with $\exp_y(\cdot)$ given by (4.6a). The appropriate value for the step size parameter $\alpha > 0$ determines the update $\exp_y(\alpha\eta)$ of y . An exact line search, which would find $\alpha = \text{argmin}_{t>0} f(\exp_y(t\eta))$, is often not feasible. A commonly used alternative is the Armijo line search, which starts with a large estimate for α and gradually reduces it until a value is found that results in a sufficient decrease of the objective function f . Algorithm 2 details the the Armijo line search strategy, and Theorem 4.6 guarantees convergence of the overall procedure.

Theorem 4.6. [1, Thm. 4.3.1] *Any cluster point of the sequence generated by the Algorithm (2) is a critical point of f .*

Algorithm 2: Riemannian Gradient with Armijo Line Search [1]

initialization: initial point y , initial step size $\alpha_0 > 0$, $\sigma, \beta \in (0, 1)$.

set search direction to $\eta = -\text{grad } f(y)$;

repeat

set $\alpha = \alpha_0$;

if $f(\exp_y(\alpha\eta)) - f(y) > \sigma\alpha\langle \text{grad } f(y), \eta \rangle_y$ **then**

└ set $\alpha = \beta\alpha$;

update the iterate $y = \exp_y(\alpha\eta)$;

update the search direction to $\eta = -\text{grad } f(y)$;

until *termination criterion holds.*

4.4. Geometric Means. Intergrid transfer operators, used to transfer information between grids, are usually calculated using linear interpolation or injection. As interpolation requires computing averages, we need to generalize such operations to our manifolds in order to devise e.g. prolongation operators. Proposition 4.1 enables to average points on \mathcal{B} in a computationally convenient way.

Proposition 4.7. *Let $(\eta_i)_{i \in I} \subset \mathcal{B}$ be arbitrary numbers and let*

$$(4.25) \quad \Omega = \{\omega_i > 0 : i \in I\}, \quad \sum_{i \in I} \omega_i = 1$$

be given corresponding weights. Then the weighted geometric mean of these numbers is

$$(4.26) \quad \bar{\eta}_\Omega := \frac{\prod_{i \in I} \left(\frac{\eta_i}{1-\eta_i}\right)^{\omega_i}}{1 + \prod_{i \in I} \left(\frac{\eta_i}{1-\eta_i}\right)^{\omega_i}} \in \mathcal{B}.$$

Likewise, the weighted geometric mean of vectors $\{y^i : i \in I\} \subset \mathcal{B}_n$ is given by

$$(4.27) \quad \bar{y}_\Omega = (\bar{y}_{j;\Omega})_{j \in I_n},$$

that is, each component $(\bar{y}_\Omega)_j$ is the weighted geometric average of the corresponding components $\{y_j^i : i \in I\}$ of the given vectors.

Proof. We define $\bar{\eta}_\Omega$ by the condition that characterizes the center of mass on a Riemannian manifold [19, Lemma 6.9.4],

$$(4.28) \quad 0 = \sum_{i \in I} \omega_i \exp_{\bar{\eta}_\Omega}^{-1}(\eta_i),$$

except for using the exponential map (4.6) instead of the exponential map corresponding to the metric connection. Using (4.7) we get

$$(4.29a) \quad 0 = \sum_{i \in I} \omega_i \bar{\eta}_\Omega (1 - \bar{\eta}_\Omega) \log \frac{(1 - \bar{\eta}_\Omega) \eta_i}{\bar{\eta}_\Omega (1 - \eta_i)}.$$

Subdividing by $\bar{\eta}_\Omega (1 - \bar{\eta}_\Omega) \in \mathcal{B}$ and taking into account (4.25) gives (4.26),

$$(4.29b) \quad \log \frac{1 - \bar{\eta}_\Omega}{\bar{\eta}_\Omega} = - \sum_{i \in I} \omega_i \log \frac{\eta_i}{1 - \eta_i}$$

$$(4.29c) \quad \bar{\eta}_\Omega = \frac{e^{\sum_{i \in I} \omega_i \log \frac{\eta_i}{1 - \eta_i}}}{1 + e^{\sum_{i \in I} \omega_i \log \frac{\eta_i}{1 - \eta_i}}} = \frac{\prod_{i \in I} \left(\frac{\eta_i}{1 - \eta_i}\right)^{\omega_i}}{1 + \prod_{i \in I} \left(\frac{\eta_i}{1 - \eta_i}\right)^{\omega_i}}.$$

□

Remark 4.8 (Geometric averaging in closed form). We point out that the computation of Riemannian means typically requires to solve an optimality condition of the form (4.28) by a fixed-point iteration. In the present case, the chosen geometry yields the corresponding geometric mean (4.26) in closed form.

4.5. Grid Transfer Operators.

4.5.1. *Prolongation.* We associate with every fine grid point $j \in I_m^c$ – recall the notation and nomenclature in Section 2.1 – a neighborhood N_j and corresponding weights Ω_j ,

$$(4.30a) \quad N_j \subset I_m, \quad j \in I_m^c$$

$$(4.30b) \quad \Omega_j = \{\omega_i : i \in N_j\}.$$

The prolongation map is defined by keeping coarse grid components and by assigning geometric averages of coarse grid components to fine grid components, i.e.

$$(4.31) \quad P: \mathcal{B}_m \rightarrow \mathcal{B}_n, \quad P(x)_j = \begin{cases} x_j & \text{if } j \in I_m, \\ \bar{x}_{\Omega_j}, & \text{if } j \in I_m^c. \end{cases}$$

Tangent vectors are transferred from coarse to fine by the differential of the prolongation map.

Lemma 4.9. *The differential of the prolongation map (4.31) is given by*

$$(4.32a) \quad dP_x: T_x \mathcal{B}_m \rightarrow T_{P(x)} \mathcal{B}_n,$$

$$(4.32b) \quad (dP_x u)_j = \begin{cases} u_j, & \text{if } j \in I_m, \\ \bar{x}_{\Omega_j} (1 - \bar{x}_{\Omega_j}) \sum_{i \in N_j} \frac{\omega_i}{x_i(1-x_i)} u_i, & \text{if } j \in I_m^c, \\ 0, & \text{otherwise.} \end{cases}$$

Proof. Put

$$(4.33a) \quad \kappa_j(x; \Omega) = \sum_{i \in N_j} \omega_i \log \frac{x_i}{1-x_i},$$

$$(4.33b) \quad \varphi_j(x; \Omega) = \bar{x}_{\Omega_j} \stackrel{(4.26)}{=} \frac{e^{\kappa_j(x; \Omega)}}{1 + \kappa_j(x; \Omega)}.$$

Then

$$(4.34a) \quad d\kappa_j(x; \Omega)[u] = \frac{d}{dt} \kappa_j(x + tu; \Omega)|_{t=0} = \sum_{i \in N_j} \frac{\omega_i}{x_i(1-x_i)} u_i,$$

$$(4.34b) \quad d\varphi_j(x; \Omega)[u] = \frac{e^{\kappa_j(x; \Omega)}}{(1 + \kappa_j(x; \Omega))^2} d\kappa_j(x; \Omega)[u] \stackrel{(4.33b)}{=} \bar{x}_{\Omega_j} (1 - \bar{x}_{\Omega_j}) d\kappa_j(x; \Omega)[u]$$

and hence from (4.31)

$$(4.35) \quad (dP_x u)_j = (dP(x)[u])_j = \begin{cases} u_j, & \text{if } j \in I_m, \\ d\varphi_j(x; \Omega)[u], & \text{if } j \in I_m^c, \\ 0, & \text{otherwise,} \end{cases}$$

which is (4.32). □

4.5.2. Restriction. In view of Definition 4.31, it is obvious that the only coarse grid vector x that conforms to the fine grid vector $y = P(x)$ is the trivial restriction

$$(4.36) \quad R: \mathcal{B}_n \rightarrow \mathcal{B}_m, \quad x = Ry = [y]_{I_m} = (y_j)_{j \in I_m}.$$

Unlike the prolongation map P (4.31), this restriction map R is linear. Concerning the restriction of tangent vectors, using the differential $dR = R$ analogous to (4.32) would be a less sensible choice, however. Rather, we define for this purpose an individual operator as follows.

Let $x = Ry$, $v \in T_y\mathcal{B}_n$ and $u \in T_x\mathcal{B}_m$. We define

$$(4.37a) \quad TR_y: T_y\mathcal{B}_n \rightarrow T_x\mathcal{B}_m,$$

$$(4.37b) \quad TR_y = G_m(x)^{-1}dP_x^\top G_n(y), \quad x \in Ry$$

$$(4.37c) \quad (TR_y)_{ij} = \begin{cases} (dP_x^\top)_{ij}, & \text{if } j = i \in I_m \\ \frac{x_i(1-x_i)}{y_j(1-y_j)}(dP_x^\top)_{ij}, & \text{if } j \in I_m^c \text{ and } i \in N_j, \\ 0, & \text{otherwise.} \end{cases}$$

$$(4.37d) \quad \text{for } j \in I_n = I_m \dot{\cup} I_m^c, \quad i \in I_m.$$

The rationale behind this definition is the equation

$$(4.38) \quad \langle u, TR_y v \rangle_x = \langle dP_x u, v \rangle_y$$

that mimics in terms of TR and the differential dP the common Galerkin condition $R = P^\top$ in the case of Euclidean scenarios and linear mappings R, P , while taking into account the coarse and fine-grid metrics. Expanding this equation using (4.5) gives

$$(4.39) \quad \langle u, G_m(x)TR_y v \rangle = \langle dP_x u, G_n(y)v \rangle$$

which implies (4.37b). Inspecting the componentwise specification (4.37c) of TR_y , we observe the aforementioned relation $TR_y = dP_x^\top$ in the case of coinciding coarse and fine grid points indexed by i and j , whereas the second line of (4.37c) shows that vector transport from the fine grid to the coarse grid is involved – cf. Remark 4.5 and (4.22) – if $j \notin I_m \subset I_n$ differs from $i \in I_m$.

As a result, pushing tangent vectors to the coarse grid using the nonlinear mapping TR better conforms to the nonlinear prolongation map (4.31), than merely using $dR = R$ based on the restriction map (4.36), which is a reasonable choice for transferring vectors of \mathcal{B}_n (rather than tangent vectors of $T\mathcal{B}_n$).

4.6. Two-Level Geometric Optimization. We generalize the two-level Euclidean optimization approach of Section 3 to our Riemannian setting.

4.6.1. Coarse Objective Function. We adapt the coarse grid model (3.3) of the objective function at $x_0 = Ry_0$ to the present geometric setting,

$$(4.40a) \quad \psi(x) = f(x) - \langle \exp_{x_0}^{-1}(x), \kappa(x_0, y_0) \rangle_{x_0},$$

$$(4.40b) \quad \kappa(x_0, y_0) = \text{grad } f(x_0) - TR_{y_0} \text{grad } f(y_0),$$

with the Riemannian gradients $\text{grad } f(x_0) = G_m(x_0)^{-1}\partial f(x_0)$, $\text{grad } f(y_0) = G_n(y_0)^{-1}\partial f(y_0)$. We check the first-order coherence condition corresponding to (3.5).

Lemma 4.10 (First-order coherence condition). *The coarse grid model (4.40) satisfies*

$$(4.41) \quad \text{grad } \psi(x_0) = TR_{y_0} \text{grad } f(y_0).$$

Proof. By Lemma 4.3, we have

$$(4.42) \quad \frac{\partial}{\partial x_i} \langle \exp_{x_0}^{-1}(x), \kappa(x_0, y_0) \rangle = \frac{x_{0i}(1-x_{0i})}{x_i(1-x_i)} \kappa_i(x_0, y_0)$$

and hence

$$(4.43a) \quad \partial \langle \exp_{x_0}^{-1}(x), \kappa(x_0, y_0) \rangle_{x_0} = G_m(x) G_m(x_0)^{-1} G_m(x_0) \kappa(x_0, y_0)$$

$$(4.43b) \quad = G_m(x) \kappa(x_0, y_0).$$

Substituting into (4.40) gives

$$(4.44a) \quad \text{grad } \psi(x) = \text{grad } f(x) - \kappa(x_0, y_0)$$

$$(4.44b) \quad \stackrel{(4.40b)}{=} \text{grad } f(x) - \text{grad } f(x_0) + TR_{y_0} \text{grad } f(y_0).$$

Setting $x = x_0$ yields (4.41). □

4.6.2. *Coarse Grid Descent Step, Algorithm.* We examine the effect of minimizing the coarse grid function (4.40), analogous to Section 3.1.2. Rearranging terms, we have

$$(4.45a) \quad \psi(x) - f(x_0) = f(x) - f(x_0) - \langle \exp_{x_0}^{-1}(x), \text{grad } f(x_0) \rangle_{x_0}$$

$$(4.45b) \quad + \langle \exp_{x_0}^{-1}(x), TR_{y_0} \text{grad } f(y_0) \rangle_{x_0}$$

$$(4.45c) \quad \stackrel{(4.38)}{=} f(x) - f(x_0) - \langle \exp_{x_0}^{-1}(x), \partial f(x_0) \rangle$$

$$(4.45d) \quad + \langle dP_{x_0} \exp_{x_0}^{-1}(x), \text{grad } f(y_0) \rangle_{y_0}.$$

Following the reasoning of Section 3.1.2 shows that, if x can be determined such that $\psi(x) < f(x_0)$, then $dP_{x_0} \exp_{x_0}^{-1}(x)$ is a descent direction with respect to f on the *fine* grid if the first three terms on the right-hand side, i.e. $f(x) - f(x_0) - \langle \exp_{x_0}^{-1}(x), \partial f(x_0) \rangle$, are nonnegative.

In the Euclidean case, with $x - x_0$ in place of $\exp_{x_0}^{-1}(x)$, the latter automatically holds since $f(x) - f(x_0) - \langle \exp_{x_0}^{-1}(x), \partial f(x_0) \rangle = D_f(x, x_0) \geq 0$ (cf. Remark 3.1). In the present geometric setting, however, this reasoning can only be guaranteed if $\|x - x_0\|$ is sufficiently small, as stated in the following Lemma.

Lemma 4.11. *For the mapping defined by (4.7), one has*

$$(4.46) \quad \exp_{x_0}^{-1}(x) = x - x_0 + \mathcal{R}(x - x_0) \quad \text{with} \quad \lim_{x \rightarrow x_0} \frac{\mathcal{R}(x - x_0)}{\|x - x_0\|} = 0.$$

Furthermore,

$$(4.47) \quad (x - x_0)_i (\exp_{x_0}^{-1}(x))_i \geq 0, \quad \forall i \in I_m.$$

Proof. The expansion (4.46) is easy to compute. As for (4.47), we have by (4.17d) that, for each component $i \in I_m$, the derivative of the function $x_i \mapsto \exp_{x_{0i}}(x_i)$ is nonnegative. Consequently, this function is monotonously increasing as is the function $x_i \mapsto x_i - x_{0i}$, and both functions only attain the value 0 at x_{0i} . □

Expansion (4.46) shows that for x sufficiently close to x_0 , $dP_{x_0} \exp_{x_0}^{-1}(x)$ is a descent direction on the fine grid whenever $\psi(x) < f(x_0)$. Relation (4.47), on the other hand, shows that the vector $\exp_{x_0}^{-1}(x)$ points into the positive half space defined by $x - x_0$ and, in this sense, is aligned with $x - x_0$. Yet, for increasing $\|x - x_0\|$, the nonlinearity of this term increases as well, such that the inner product $\langle \exp_{x_0}^{-1}(x), \partial f(x_0) \rangle$ may not allow to certify a descent direction by merely checking

the sign of $\psi(x) - f(x_0)$. Therefore, besides the geometric version of condition (3.13),

$$(4.48) \quad \|TR_{y_0} \text{grad } f(y_0)\|_{Ry_0} \geq \eta \|\text{grad } f(y_0)\|_{Ry_0} \quad \wedge \quad \|y - y_c\| \geq \varepsilon,$$

the following algorithm involves the additional test

$$(4.49) \quad \tilde{D}_f(x, x_0) := f(x) - f(x_0) - \langle \exp_{x_0}^{-1}(x), \partial f(x_0) \rangle \geq 0,$$

that still can be cheaply done at the coarse level, rather than checking directly and more costly the vector $dP_{x_0} \exp_{x_0}^{-1}(x)$ on the fine level.

Algorithm 3: Two-Level Geometric Optimization

initialization: initial fine grid point y , corresponding coarse grid point $x = Ry$.

repeat

- set $y_0 = y, x_0 = Ry$;
- if** condition (4.48) is satisfied at x_0 **then**
 - find x such that $\psi(x) < f(x_0)$; /* coarse grid update */
 - if** condition (4.49) holds **then**
 - set $h = dP_{x_0} \exp_{x_0}^{-1}(x)$;
 - find $\alpha > 0$: $f(\exp_{y_0}(\alpha h)) < f(y_0)$; /* line search */
 - $y \leftarrow \exp_{y_0}(\alpha h)$ /* update */
- else**
 - apply one fine level iteration to compute y such that $f(y) < f(y_0)$.
 - /* update */

until termination criterion holds.

4.7. MultiLevel Extension. Analogous to Section 3.2, we introduce the notation

$$(4.50a) \quad x^{(\ell)} \in \mathcal{B}_{n_\ell}^{(\ell)}$$

for the variables at levels $\ell \in \{0, 1, 2, \dots\}$, where

$$(4.50b) \quad x^{(0)} = x \in \mathcal{B}_n = \mathcal{B}_{n_0}^{(0)}$$

are the original fine grid variables. The components $x_i^{(\ell)}$ of these vectors are indexed by the nested index sets

$$(4.50c) \quad i \in I_{n_\ell} \subset I_{n_{\ell-1}} \subset \dots \subset I_n = I_{n_0}.$$

Accordingly, the coarse grid model (4.40) applied at level ℓ reads

$$(4.51a) \quad \psi(x^{(\ell)}) = f(x^{(\ell)}) - \langle \exp_{x_0^{(\ell)}}^{-1}(x^{(\ell)}), \kappa(x_0^{(\ell)}, x_0^{(0)}) \rangle_{x_0^{(\ell)}},$$

$$(4.51b) \quad \kappa(x_0^{(\ell)}, x_0^{(0)}) = \text{grad } f(x_0^{(\ell)}) - TR_{(\ell)}^{(0)}(x_0^{(0)}) \text{grad } f(x_0^{(0)}),$$

with $x_0^{(\ell)} = R_{(\ell)}^{(0)} x_0^{(0)}$. Prolongation operators $P_{(0)}^{(\ell)}$, $\ell > 1$ can be defined in the same way as $P_0^{(1)} = P$ is defined by (4.31), with increasingly larger neighborhoods (4.30).

By this, the mappings $TR_{(\ell)}^{(0)}$ for pushing tangent vectors from level 0 to ℓ follow analogous to (4.38).

This formulation adapts to the geometric setting the Euclidean multilevel approach outlined in Section 3.2, and similar comments concerning the trade-off of invoking coarse grid models at various levels apply. We leave a detailed study of this multilevel extension for future work.

5. EXPERIMENTS

In this section we numerically evaluate the proposed approach. In a first experiment we compare the Riemannian Gradient (RG) descent, see Algorithm 2, with a state-of-the-art Accelerated Bregmann Proximal Gradient (ABPG) method [17]. In a second experiment we evaluate one-level vs. two-level (2L) schemes and compare our approach using Riemannian Gradient descent on both fine and coarse grids with a state-of-the-art first-order multilevel approach [21] based on Projected Gradients (PG).

5.1. Data Setup. For our numerical evaluation we considered the phantom (1024×1024 pixels) shown by Figure 2. The tomographic projection matrices were computed by means of the ASTRA-toolbox¹, using parallel beam geometry with equidistant angles in the range $[0, \pi]$. The undersampling rate at the fine grid was chosen to be 2% and 20%, respectively, corresponding to 20 and 200 projection angles. Each entry a_{ij} of the projection matrix A corresponds to the length of the line segment of the i -th projection ray passing through the j -th pixel in the image domain (Figure 1). At every level the width of the detector-array was set to the grid size, so that at each scale every pixel intersects with at least a single projection ray.




FIGURE 2. The phantoms (1024×1024) used for the numerical evaluation exhibit both fine and large scale structures and various shapes.

5.2. Implementation Details. The coarse grid update in Algorithm 3 was performed by Armijo line search with a modification, such that if the calculated feasible stepsize was numerically close to 0, then the coarse step was discarded and a normal step was performed instead.

For an efficient implementation of the geometric grid transfer operators we used vectorized versions. For the experiments in this work we use uniform weights ω_i ,

¹<https://www.astra-toolbox.com>

therefore the geometric mean, defined in Eq. (4.26), can be reformulated as

$$(5.1) \quad \bar{\eta}_\Omega = \frac{\prod_{i \in I} \left(\frac{\eta_i}{1-\eta_i}\right)^{\omega_i}}{1 + \prod_{i \in I} \left(\frac{\eta_i}{1-\eta_i}\right)^{\omega_i}} = \frac{e^{\sum_{i \in I} \omega_i \log \frac{\eta_i}{1-\eta_i}}}{1 + e^{\sum_{i \in I} \omega_i \log \frac{\eta_i}{1-\eta_i}}} = \frac{e^{BI\left(\log \frac{\eta}{1-\eta}\right)}}{1 + e^{BI\left(\log \frac{\eta}{1-\eta}\right)}},$$

where $BI(v)$ is the standard bilinear interpolation operator. We used PyLops², the linear operator library for Python, for performing bilinear interpolation and its transpose operation. Consequently, the grid transfer operators (4.31), (4.32) and (4.37) can be vectorized by using bilinear interpolation, resulting in

$$(5.2a) \quad P(x) = \bar{x}_\Omega$$

$$(5.2b) \quad dP_x u = G_n(P(x))^{-1} BI(G_m(x)u)$$

$$(5.2c) \quad TR_y v = BI^\top(G_n(P(R(y))))^{-1} G_n(y)v,$$

where $BI^\top(v)$ denotes the transposed bilinear interpolation operator. Note that in case of non-uniform weights, the operation $BI(v)$ can be still represented by a corresponding linear operator.

In order to avoid numerical issues very close to the boundary of the box manifold, we clip each component y_i to $[\varepsilon, 1-\varepsilon]$ with $\varepsilon = 10^{-10}$. This clipping step is performed directly after evaluating the exponential map (4.6).

For all algorithms in our numerical evaluation we used the same set of parameters. The maximum number of iterations was set to 50, which also serves as termination criterion. We used the natural initialization at $y^0 = 0.5 \mathbf{1}$ (uninformative choice). The parameters for the Armijo line search [1, Def. 4.2.2] were set to $\sigma = 10^{-4}$, $\beta = 0.6$, $\alpha_0 = 1/\beta$. For condition (3.13) and its geometric version (4.48) for optimization on the coarse model, we selected $\eta = 0.49$ and $\varepsilon = 10^{-3}$. Regarding the objective function f defined in Eq. (1.2), we set $\lambda = 0.5$ for weighting the smoothed total variation with smoothing parameter $\rho = 0.5$.

5.3. Experiments and Discussion. In a first experiment, we studied the impact of the Riemannian geometry on first-order optimization. We compared Riemannian gradient descent with a state-of-the-art accelerated proximal gradient method, the ABPG method [17, Alg. 4], that is suited to minimize objectives that are not Lipschitz-smooth, like the KL-divergence based objective function (1.2). Results are shown as Figure 3. By choosing the two quite different under-sampling scenarios, we created one problem instance with 200 angles (20% undersampling) that is likely to have a unique optimum, see [22], and a second problem with multiple solutions due to the high undersampling ratio of 2% with 20 angles. The proposed method performed on par or better than ABPG and was significantly superior for that latter three phantoms shown in Figure 2 that have more homogeneous shapes.

In a second experiment we studied the influence of using a second coarse level for computing efficiently descent directions on a 512×512 grid, as summarized by Algorithm 3. Figures 4 and 5 show the results for the Riemannian gradient descent: using one level versus two levels compared with state of the art ABPG. We observed the following. Firstly, the impact of the two-level acceleration is much stronger in the highly underdetermined case. Secondly, we concede that although

²<https://github.com/PyLops/pylops>

ABPG performed worse in both scenarios, the cost of the iterations on the fine level are more expensive when using our method due to line search. On the coarser level however, the higher cost of the line search is compensated by using a much smaller number of variables for computing search directions.

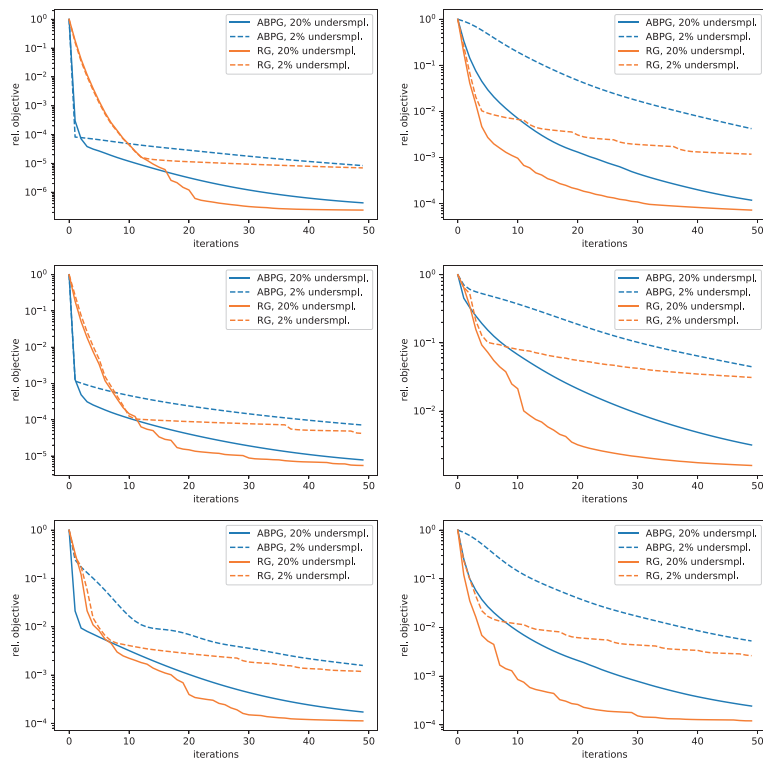


FIGURE 3. **ABPG vs. Riemannian gradient (RG) descent** in terms of relative objective values is compared on 2% and 20% undersampled projection data. The left column corresponds to the first three phantoms: **gear**, **bone**, **vessel** in Figure 2, the right columns to the last three: **batenburg**, **roux**, **skulls**. The Riemannian gradient descent achieves lower objective values due to proper step-size selection (line search). For the more homogeneous phantoms corresponding to the right column, the proposed method significantly outperforms ABPG [17].

Finally, we compared the proposed method to the Euclidean state-of-the-art multilevel method in [21], that is capable on handling box constraints. The results are shown as Figure 6. We call the latter method two-level projected gradient (2L PG) as it uses the projected gradient method on both levels for optimizing the fine grid objective as well as the coarse grid model. Two-level projected gradient uses a scaled-down box that is adapted to the current iterate in such a way that feasibility is achieved after prolongation of the coarse level correction. Such an adaption of the constraints is not needed in our proposed approach and therefore makes our approach more flexible. We notice however that the impact of the second level in

the case of the projected gradient is stronger than in the geometric case. We believe that this is due to the fact that our coarse model loses its convexity when the iterate is close to the boundary of the feasible set.

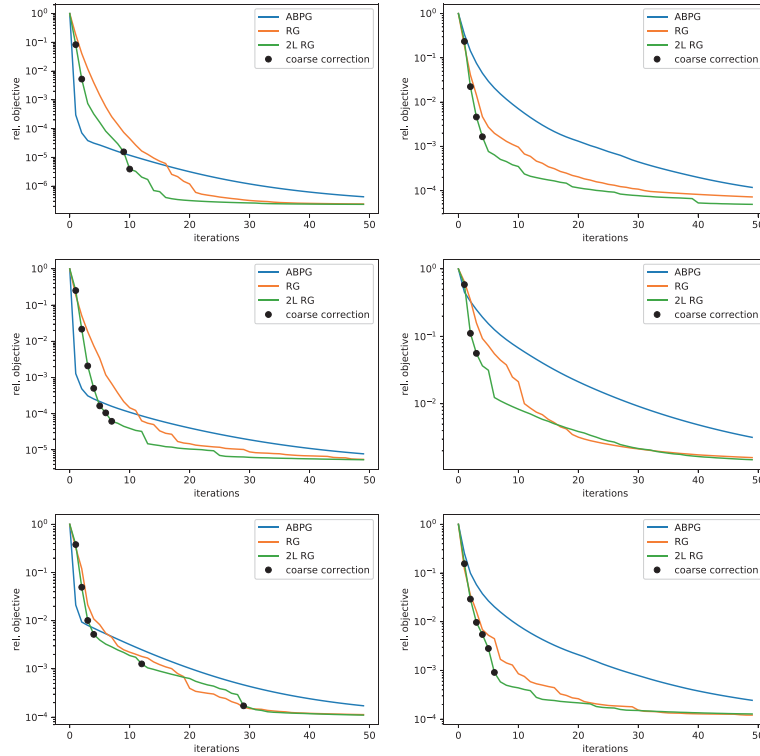


FIGURE 4. **20% undersampling: comparison of relative objective function values** for single-level resp. two-level (2L) Riemannian gradient descent (RG, 2L RG) and ABPG [17]. The left column corresponds to the first three phantoms: **gear**, **bone**, **vessel** in Figure 2, the right columns to the last three: **batenburg**, **roux**, **skulls**. Black dots indicate when descent directions were computed on coarser grids. On the coarse grid we have 40% undersampling and the coarse problem has a unique solution [22]. As a consequence we quickly approach the boundary of the box. This results in an inefficient update using the Armijo line search and more similar curves for single-level resp. two-level (2L) Riemannian gradient descent.

6. CONCLUSIONS

We introduced a novel approach to geometric multilevel optimization. The approach employs information geometry in order to devise all ingredients of the iterative multilevel scheme. Invoking coarse level representations for computing descent directions effectively accelerates convergence. Experiments conducted for a range of

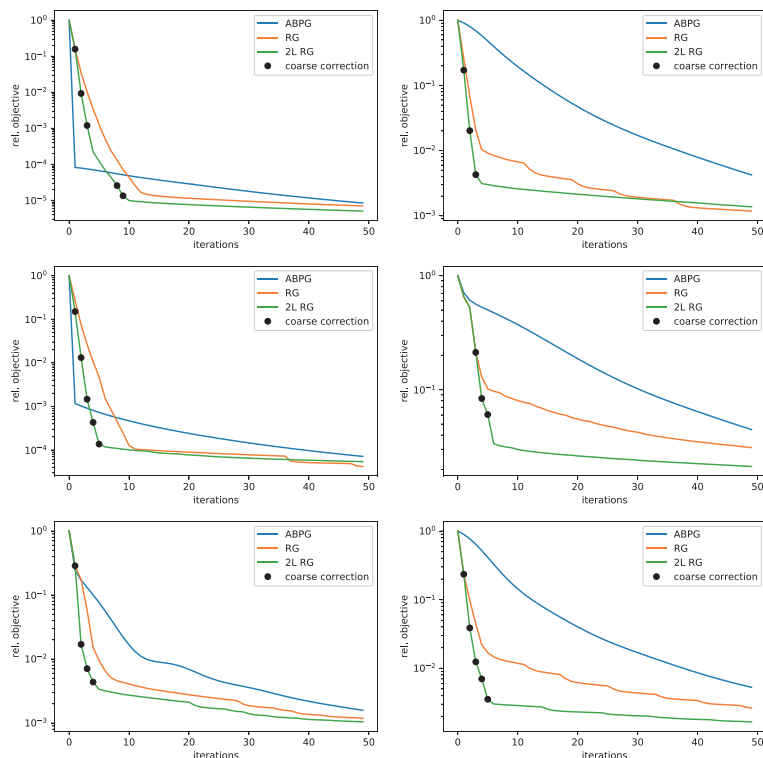


FIGURE 5. **2% undersampling: comparison of relative objective function values** for single-level resp. two-level (2L) Riemannian gradient descent (RG, 2L RG) and ABPG [17]. The left column corresponds to the first three phantoms: **gear**, **bone**, **vessel** in Figure 2, the right columns to the last three: **batenburg**, **roux**, **skulls**. Black dots indicate when descent directions were computed on coarser grids. The two-level schemes - where we now have 4% undersampling - aggressively decreases the objective, in particular for more homogeneous phantoms in the right column.

instances of an ill-posed linear inverse problem with a non-quadratic convex regularizer and box constraints demonstrate promising performance of the method relative to the state-of-the-art.

The derivation of the approach for boxed-constrained convex programs can be transferred to other convex programs with simply structured feasible sets, analogous to turning parameter spaces of probability distributions into Riemannian manifolds in information geometry. Simplices instead of boxes as feasible sets provide an example [4].

This paper mainly focuses on the case of two-grid geometric optimization, which constitutes the core problem of multilevel optimization. Our future work will examine the multilevel case in detail, the selection of an appropriate number of resolution levels and related problems, and possible refinements. The latter includes machine

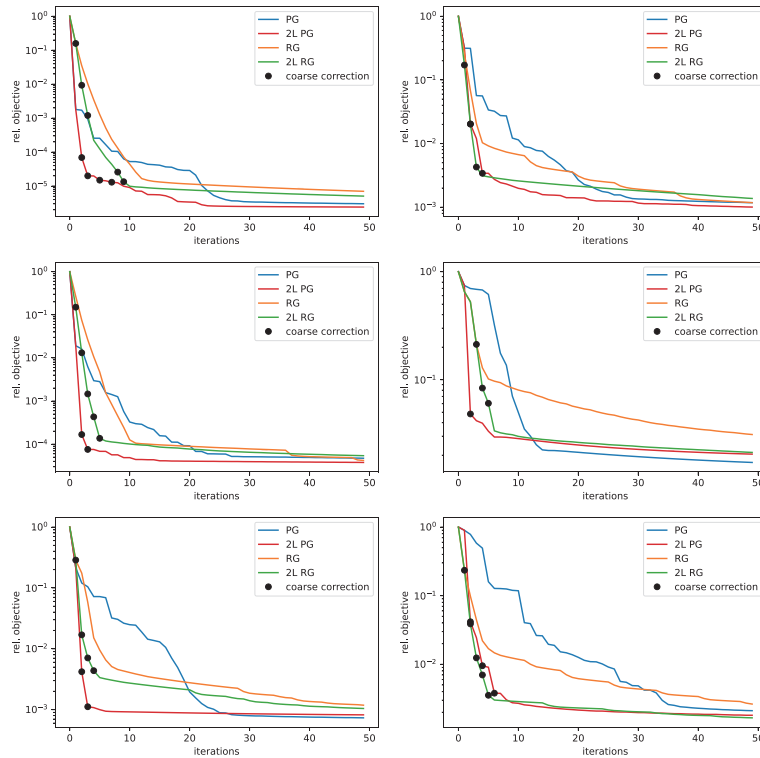


FIGURE 6. **Projected gradient (PG) vs. Riemannian gradient (RG) descent** in terms of relative objective values is compared for the single-level and two-level (2L) scenario with 2% undersampled projection data. The left column corresponds to the first three phantoms: `gear`, `bone`, `vessel` in Figure 2, the right columns to the last three: `batenburg`, `roux`, `skulls`. Black dots indicate when descent directions were computed on coarser grids. Overall, the Riemannian version of the gradient descent decreases faster the objective than projected gradient, since the descent direction already respects the constraints. Both two-level approaches show in addition more rapid reduction of the objective function.

learning components that can be controlled precisely, in order to optimize the transfer between levels, to adapt to each concrete problem instance and to given input data, for further accelerating the overall optimization process.

ACKNOWLEDGEMENT

We thank Jan Plier (Heidelberg University) for simulation code that efficiently evaluates our objective. We greatly benefited from fruitful discussions with Christoph Schnörr (Heidelberg University). To Oana Curtef (University of Würzburg) we are indebted for her observation concerning line search. SP and MZ gratefully acknowledge the generous and invaluable support of the Klaus Tschira Foundation.

REFERENCES

- [1] P. A. Absil, R. Mahony, and R. Sepulchre, *Optimization Algorithms on Matrix Manifolds*, Princeton University Press, 2008.
- [2] F. Alvarez, J. Bolte, and O. Brahic, *Hessian Riemannian Gradient Flows in Convex Programming*, SIAM Journal on Control and Optimization **43** (2004), no. 2, 477–501.
- [3] S.-I. Amari and H. Nagaoka, *Methods of Information Geometry*, Amer. Math. Soc. and Oxford Univ. Press, 2000.
- [4] F. Åström, S. Petra, B. Schmitzer, and C. Schnörr, *Image Labeling by Assignment*, Journal of Mathematical Imaging and Vision **58** (2017), no. 2, 211–238.
- [5] D. A. Bayer and J. C. Lagarias, *The Nonlinear Geometry of Linear Programming. I. Affine and Projective Scaling Trajectories*, Trans. Amer. Math. Soc. **314** (1989), no. 2, 499–526.
- [6] A. Beck and M. Teboulle, *Smoothing and First Order Methods: A Unified Framework*, SIAM J. Optim. **22** (2012), 557–580.
- [7] J. M. Borwein and J. D. Vanderwerff, *Convex Functions: Constructions, Characterizations and Counterexamples*, Cambridge Univ. Press, 2010.
- [8] L. M. Bregman, *A Relaxation Method of Finding a Common Point of Convex Sets and its Application to the Solution of Problems in Convex Programming*, U.S.S.R. Comput. Math. Phys. **7** (1967), 200–217.
- [9] L. D. Brown, *Fundamentals of Statistical Exponential Families*, Institute of Mathematical Statistics, Hayward, CA, 1986.
- [10] Y. Censor and A. Lent, *An Iterative Row-Action Method for Interval Convex Programming*, J. Optimiz. Theory Appl. **34** (1981), no. 3, 321–353.
- [11] Y. Censor, S. Petra, and C. Schnörr, *Superiorization vs. Accelerated Convex Optimization: The Superiorized/Regularized Least Squares Case*, J. Appl. Numer. Optimization **2** (2020), no. 1, 15–62.
- [12] G. Chen and M. Teboulle, *Convergence Analysis of a Proximal-Like Minimization Algorithm Using Bregman Functions*, SIAM Journal on Optimization **3** (1993), no. 3, 538–543.
- [13] O. Christensen, H. G. Feichtinger, and S. Paukner, *Gabor Analysis for Imaging*, Handbook of Mathematical Methods in Imaging, Scherzer, O., 2015, pp. 1717–1757.
- [14] M. Fazlylab, A. Ribeiro, M. Morari, and V. M. Preciado, *Analysis of Optimization Algorithms via Integral Quadratic Constraints: Nonstrongly Convex Problems*, SIAM J. Optim. **28** (2018), no. 3, 2654–2689.
- [15] S. Gratton, A. Sartenaer, and P. L. Toint, *Recursive Trust-Region Methods for Multiscale Nonlinear Optimization*, SIAM J. Optim. **19** (2008), 414–444.
- [16] W. Hackbusch, *Multi-Grid Methods and Applications*, Springer, 1985.
- [17] F. Hanzely, P. Richtárik, and L. Xiao, *Accelerated Bregman Proximal Gradient Methods for Relatively Smooth Convex Optimization*, Computational Optimization and Applications **79** (2021), 405–440.
- [18] V. Hovhannisyán, P. Parpas, and S. Zafeiriou, *MAGMA: Multi-level accelerated gradient mirror descent algorithm for large-scale convex composite minimization*, SIAM J. Imaging Sci. **9** (2016), no. 4, 1829–1857.
- [19] J. Jost, *Riemannian Geometry and Geometric Analysis*, 7th ed., Springer-Verlag Berlin Heidelberg, 2017.
- [20] R. Kornhuber, *Monotone Multigrid Methods for Elliptic Variational Inequalities II*, I. Numer. Math **69** (1993), 18–4.
- [21] M. Kočvara and S. Mohammed, *A First-Order Multigrid Method for Bound-Constrained Convex Optimization*, Optim. Method. Softw. **31** (2016), no. 3, 622–644.
- [22] J. Kuske and S. Petra, *Performance Bounds For Co-/Sparse Box Constrained Signal Recovery*, An. St. Univ. Ovidius Constanta **1** (2019), 79–109.
- [23] S. L. Lauritzen, *Chapter 4: Statistical Manifolds*, Differential Geometry in Statistical Inference (Shanti S. Gupta, S. I. Amari, O. E. Barndorff-Nielsen, R. E. Kass, S. L. Lauritzen, and C. R. Rao, eds.), Institute of Mathematical Statistics, Hayward, CA, 1987, pp. 163–216.
- [24] S. Mallat, *A Wavelet Tour of Signal Processing: The Sparse Way*, Academic Press, 2009.

- [25] S.G. Nash, *A Multigrid Approach to Discretized Optimization Problems*, Optim. Method. Softw. **14** (2000), 99–119.
- [26] Y. Nesterov, *Smooth Minimization of Non-Smooth Functions*, Math. Programming, Ser. A **103** (2005), 127–152.
- [27] Y. Nesterov and A. Nemirovskii, *Interior Point Polynomial Algorithms in Convex Programming*, SIAM, 1994.
- [28] Y. E. Nesterov and M. J. Todd, *On the Riemannian Geometry Defined by Self-Concordant Barriers and Interior-Point Methods*, Found. Comp. Math. **2** (2002), 333–361.
- [29] P. Parpas, *A Multilevel Proximal Gradient Algorithm for a Class of Composite Optimization Problems*, SIAM J. Sci. Comput. **39** (2017), no. 5, 681–701.
- [30] J. Plier, F. Savarino, M. Kočvara, and S. Petra, *First-Order Geometric Multilevel Optimization for Discrete Tomography*, Scale Space and Variational Methods in Computer Vision: 8th International Conference, SSVM 2021, Virtual Event, May 16–20, 2021, Proceedings, Springer-Verlag, 2021, pp. 191–203.
- [31] S. Salzo and S. Villa, *Inexact and Accelerated Proximal Point Algorithms*, J. Convex Analysis **19** (2012), no. 4, 1167–1192.
- [32] W. Su, S. Boyd, and E.J. Candes, *A Differential Equation for Modeling Nesterov’s Accelerated Gradient Method: Theory and Insights*, J. Mach. Learning Res. **17** (2016), 1–43.
- [33] M. Sutti and B. Vandereycken, *Riemannian Multigrid Line Search for Low-Rank Problems*, SIAM Journal on Scientific Computing **43** (2021), no. 3, A1803–A1831.
- [34] U. Trottenberg, C. Oosterlee, and A. Schüller, *Multigrid*, Academic Press, 2001.
- [35] P. Tseng, *On Accelerated Proximal Gradient Methods for Convex-Concave Optimization*, unpublished manuscript, available at <http://www.mit.edu/~dimitrib/PTseng/papers/apgm.pdf>, 2008.
- [36] N. Tsipinakis and P. Parpas, *A Multilevel Method for Self-Concordant Minimization*, 2021, <https://arxiv.org/abs/2106.13690>.
- [37] Z. Wen and D. Goldfarb, *A Line Search Multigrid Method for Large-Scale Nonlinear Optimization*, SIAM J. Optim. **20** (2009), no. 3, 1478–1503.
- [38] A. Wibisono, A. C. Wilson, and M. I. Jordan, *A Variational Perspective on Accelerated Methods in Optimization*, Proceedings of the National Academy of Sciences **113** (2016), no. 47, E7351–E7358.

Manuscript received July 1 2022

revised February 23 2021

S. MÜLLER

Dept. Mathematics and Comp. Science, Heidelberg University, Germany

E-mail address: `sebastian_mueller@stud.uni-heidelberg.de`

S. PETRA

Dept. Mathematics and Comp. Science, Heidelberg University, Germany

E-mail address: `petra@math.uni-heidelberg.de`

M. ZISLER

Dept. Mathematics and Comp. Science, Heidelberg University, Germany

E-mail address: `zisler@math.uni-heidelberg.de`