

Riemannian optimization with finite-difference gradient approximations

Timothé Taminiau^{*1}, Estelle Massart^{†1}, and Geovani Nunes Grapiglia^{‡1}

¹Université Catholique de Louvain, ICTEAM/INMA, Avenue Georges Lemaître 4, B-1348, Louvain-la-Neuve, Belgium

January 27, 2026

Abstract

Derivative-free Riemannian optimization (DFRO) aims to minimize an objective function using only function evaluations, under the constraint that the decision variables lie on a Riemannian manifold. The rapid increase in problem dimensions over the years calls for computationally cheap DFRO algorithms, that is, algorithms requiring as few function evaluations and retractions as possible. We propose a novel DFRO method based on finite-difference gradient approximations that relies on an adaptive selection of the finite-difference accuracy and stepsize that is novel even in the Euclidean setting. When endowed with an intrinsic finite-difference scheme, that measures variations of the objective in tangent directions using retractions, our proposed method requires $O(d\epsilon^{-2})$ function evaluations and retractions to find an ϵ -critical point, where d is the manifold dimension. We then propose a variant of our method when the search space is a Riemannian submanifold of an n -dimensional Euclidean space. This variant relies on an extrinsic finite-difference scheme, approximating the Riemannian gradient directly in the embedding space, assuming that the objective function can be evaluated outside of the manifold. This approach leads to worst-case complexity bounds of $O(d\epsilon^{-2})$ function evaluations and $O(\epsilon^{-2})$ retractions. We also present numerical results showing that the proposed methods achieve superior performance over existing derivative-free methods on various problems in both Euclidean and Riemannian settings.

1 Introduction

In this work, we consider the problem of minimizing a continuously differentiable function over a d -dimensional Riemannian manifold \mathcal{M} . Riemannian optimization, which has expanded considerably in the last two decades, is the standard approach for addressing this problem [3, 8, 39]. In this work, we assume that the derivatives of the objective function are not available, preventing the use of gradient-based Riemannian optimization methods. The lack of derivative information typically arises when the objective involves a black-box component, as in the design of adversarial attacks for neural network classifiers. These attacks seek an input perturbation that leads to a failure of

^{*}Email: timothe.taminiau@uclouvain.be. Supported by “Fonds spéciaux de Recherche”, UCLouvain.

[†]Email: estelle.massart@uclouvain.be.

[‡]Email: geovani.grapiglia@uclouvain.be. Partially supported by FRS-FNRS, Belgium (Grant CDR J.0081.23).

the classifier, without having access to its architecture, and are naturally formulated as a derivative-free *Euclidean* optimization problems [32, 42]. The advent of geometric deep learning and the design of neural networks with manifold-valued inputs [12, 27] calls for derivative-free *Riemannian* optimization (DFRO) algorithms for adversarial attacks design and robustness quantification. Other applications of DFRO arise, e.g., in the control of robotic systems [28], and in blind source separation [41].

While in the Euclidean setting, derivative-free optimization is a rather mature topic ([6, 13, 31]), considerably fewer works address DFRO. Several works extended direct search methods to Riemannian manifolds [11, 17, 18, 19, 30, 41]. These methods search for a new candidate by evaluating the objective on a mesh of trial points around the current iterate. When the objective function is sufficiently smooth, an alternative strategy extends classical derivative-based optimization algorithms by replacing derivatives by their finite-difference approximations. The local linear structure of the manifold is then used to define local sets of coordinates, allowing for the generalization of finite-differences schemes to Riemannian manifolds. Since this approximation is based on local variations along curves on the manifold, we refer to it as an *intrinsic finite-difference gradient*. It has been used in several DFRO algorithms [26, 34, 35] and generalized to simplex gradients, in which the reference directions form a spanning set of the tangent space at the reference point [37]. Other works addressed Gaussian smoothing, which estimates the gradient from a two-point function evaluation scheme [32], and finite-difference Hessian approximations [7].

As the problem dimension increases, ensuring the computational efficiency of DFRO algorithms becomes a major challenge. Their computational cost is governed by their total number of function evaluations and retractions. Computationally expensive function evaluations are common in derivative-free optimization; they may for example require solving a set of partial differential equations [29]. As for retractions, note that, while the design of numerically cheap retractions has been the focus of dedicated research [4, 5, 38, 40], their computational cost remains typically high for several manifolds of interest, which motivated the design of retraction-free Riemannian optimization methods, see, e.g., [1, 2, 20, 22, 33, 43]. As a whole, it is therefore crucial to derive DFRO algorithms with a worst-case complexity in terms of function evaluations and retractions that is as low as possible.

Unfortunately, the above-cited finite-difference DFRO methods are either not equipped with complexity results, or are equipped with results derived under stronger problem assumptions, such as the knowledge of some smoothness constants of the objective [26, 32], an assumption that is too strong for many problems of interest. In derivative-free *Euclidean* optimization, a strategy to avoid the use of smoothness constants is to rely on adaptive finite-difference accuracy and stepsize parameters [23, 24]. Under the assumption that $f(\cdot)$ is lower bounded and has Lipschitz continuous gradient, these methods need at most $O(d\epsilon^{-2})$ function evaluations, with d the dimension of the search space, to find an ϵ -critical point, i.e. a point \bar{x} such that $\|\nabla f(\bar{x})\| \leq \epsilon$, with no information on the smoothness constants of the objective. Note that, for direct-search methods, a quadratic complexity in problem dimension was shown to be unavoidable [15]. To our knowledge, the design of DFRO algorithms with adaptive parameter selection (i.e., that do not relying on the knowledge of parameter constants), with associated complexity analysis, is still an open question.

The contributions of this paper are twofold. First, assuming that we have no access to the smoothness constants of the objective, we propose a finite-difference DFRO algorithm that relies on an adaptive selection scheme for the stepsize and finite-difference accuracy, which is novel even in the Euclidean setting. This scheme is motivated by the observation that the method proposed

in [24] (in the Euclidean setting) tunes the stepsize and finite-difference accuracy based on one single parameter that approximates the Lipschitz constant of the gradient of the objective. However, the level of accuracy required in this estimation differs depending on the goal considered: while more conservative estimates should be used in the finite-difference computations, the use of sharper estimates in the stepsize selection may result in larger steps and speed up convergence. Inspired by [14], we propose a novel DFRO algorithm, based on intrinsic finite-difference gradients, that relies on two adaptive estimates of the Lipschitz constant of the gradient of the objective, used respectively in the finite-difference accuracy and stepsize computation. We prove that our algorithm finds an ϵ -critical point (i.e., a point at which the norm of the Riemannian gradient is no greater than ϵ) after at most $O(d\epsilon^{-2})$ function evaluations and $O(d\epsilon^{-2})$ retractions, and show numerically an improvement over the method proposed in [24] in the Euclidean setting.

As a second contribution, we decrease the retraction complexity of our DFRO algorithm by proposing a variant that relies on an alternative *extrinsic finite-difference scheme*, for optimization over Riemannian submanifolds of a Euclidean space. This scheme approximates the gradient from local variations along tangent directions *in the embedding space*, allowing the saving of d retractions per finite-difference gradient estimation. We prove that the worst-case complexities of this variant of our DFRO algorithm are $\mathcal{O}(d\epsilon^{-2})$ functions evaluations and $\mathcal{O}(\epsilon^{-2})$ retractions, with the latter being independent of the problem dimension. Numerical results are provided that show a substantial computational cost reduction compared to the intrinsic scheme.

The structure of the paper is as follows. In Section 2, we provide a short reminder of key concepts in Riemannian optimization, specify the problem assumptions and give preliminary lemmas that we will need in the convergence analysis of our methods. We present in Section 3 our DFRO method based on the intrinsic finite-difference gradients, with associated worst-case complexity bounds, and in Section 4 our variant that relies on an extrinsic finite-difference scheme for optimization on Riemannian submanifolds of a Euclidean space. Finally, Section 5 contains numerical experiments that compare our proposed methods with state-of-the-art methods regarding both function evaluations and running time on benchmark problems.

Notation: We write $T_x\mathcal{M}$ the tangent space to the manifold \mathcal{M} at $x \in \mathcal{M}$, $\langle \cdot, \cdot \rangle_x$ the Riemannian metric at x , and $\|\cdot\|_x$ the associated Riemannian norm. The tangent bundle of \mathcal{M} is written $T\mathcal{M}$. Let $R : T\mathcal{M} \rightarrow \mathcal{M}$ be a retraction on \mathcal{M} . For each $x \in \mathcal{M}$, the mapping $R_x : T_x\mathcal{M} \rightarrow \mathcal{M}$ is defined by $R_x(\eta) := R(x, \eta)$. The Riemannian gradient of $f : \mathcal{M} \rightarrow \mathbb{R}$ is written $\text{grad } f(x) \in T_x\mathcal{M}$. When the manifold is a Riemannian submanifold of a Euclidean space \mathcal{E} with respect to the canonical metric $\langle \cdot, \cdot \rangle$, we denote $\text{Proj}_{T_x\mathcal{M}} : \mathcal{E} \rightarrow T_x\mathcal{M}$ the orthogonal projection of any vector $\eta \in \mathcal{E}$ onto $T_x\mathcal{M}$.

2 Preliminary material

We start this section with a short reminder of the key concepts from Riemannian optimization used in this work, and describe next our problem assumptions and a couple of lemmas required for the convergence analyses.

2.1 Basic tools and definitions in Riemannian Optimization

A Riemannian manifold is a (smooth) manifold whose tangent space (i.e., local linear approximation) at any $x \in \mathcal{M}$ (written $T_x \mathcal{M}$) is endowed with an inner product $\langle \cdot, \cdot \rangle_x$ that varies smoothly with x , see, e.g., [3, 8]. This Riemannian metric allows to define the Riemannian gradient of f at x , as the unique vector $\text{grad } f(x) \in T_x \mathcal{M}$ that satisfies $Df(x)[\eta] = \langle \text{grad } f(x), \eta \rangle_x$ for all $\eta \in T_x \mathcal{M}$, where $Df(x)[\eta]$ is the directional derivative of f in the direction η . When the Riemannian gradient is not available, it can be approximated by the intrinsic finite-difference gradient

$$g_h(x) = \sum_{i=1}^d \frac{f(R_x(he_i(x))) - f(x)}{h} e_i(x), \quad (1)$$

where $\{e_1(x), \dots, e_d(x)\}$ is a basis of $T_x \mathcal{M}$, $h > 0$ controls the finite-difference accuracy, and $R_x(\cdot)$ is defined by a retraction. Retractions allow to make a step on the manifold from an arbitrary point $x \in \mathcal{M}$ along an arbitrary tangent vector $\eta \in T_x \mathcal{M}$. More precisely, a retraction is a smooth mapping $R : T\mathcal{M} \rightarrow \mathcal{M}$, with $(x, \eta) \mapsto R_x(\eta)$, where $T\mathcal{M}$ is the tangent bundle of \mathcal{M} , satisfying $R_x(0) = x$ and $DR_x(0)[\cdot]$ is the identity operator (the so-called local rigidity condition), see, e.g., [3, §3.5.3 and Def. 4.1.1]. In this work, we assume that \mathcal{M} is endowed with a retraction R such that, for all $x \in \mathcal{M}$, $R_x : T_x \mathcal{M} \rightarrow \mathcal{M}$ is defined over the entire $T_x \mathcal{M}$. This assumption holds for several important examples of retraction, it fails for example for some projection-based retractions [4] where the projection is only defined on a neighborhood of $0_x \in T_x \mathcal{M}$; we refer, e.g., to [9] for a complexity analysis of classical Riemannian optimization algorithms that accounts for locally defined retractions, by restricting the stepsize to the domain of definition of the retraction at each iteration. While similar ideas could be used in this work, we relied on our above-mentioned assumption on the retraction to simplify the presentation of our methods and their convergence analyses.

2.2 Problem assumptions and auxiliary results

We consider the optimization problem

$$\min_{x \in \mathcal{M}} f(x), \quad (\text{P})$$

where we make the two following assumptions on the objective function f . The first one is a generalization to manifolds of the classical Lipschitz-smoothness assumption (see [9]), while the second simply requires the objective to be lower-bounded.

A1. $f(\cdot)$ is $(L_{\mathcal{M}}, R)$ -smooth, that is for all $x \in \mathcal{M}$,

$$|f(R_x(\eta)) - f(x) - \langle \text{grad } f(x), \eta \rangle_x| \leq \frac{L_{\mathcal{M}}}{2} \|\eta\|_x^2, \quad \forall \eta \in T_x \mathcal{M}.$$

A2. There exists $f_{\text{low}} \in \mathbb{R}$ such that $f(x) \geq f_{\text{low}}$, for all $x \in \mathcal{M}$.

Since we will consider different finite-difference gradient schemes in Section 3 and Section 4, we encompass here these two schemes in a broader notion of approximate Riemannian gradient, namely, tangent vectors parametrized by some parameter h allowing to control their accuracy with respect to the exact Riemannian gradient.

Definition 2.1. An approximate Riemannian gradient of a function $f : \mathcal{M} \rightarrow \mathbb{R}$ is a mapping $g : [0, +\infty) \times \mathcal{M} \rightarrow T\mathcal{M}$, $(h, x) \mapsto g_h(x)$, such that

$$\|\text{grad } f(x) - g_h(x)\|_x \leq C_f h, \quad (2)$$

for some constant $C_f > 0$ that does not depend on x .

Thereafter, we present two auxiliary lemmas that will be key in the convergence analysis of the methods proposed in this paper. The first lemma shows that the relative error of the approximate Riemannian gradient is bounded by some constant as soon as h is small enough and we have not yet reached ϵ -criticality.

Lemma 2.2 (Adapted from Lemma 2 in [24]). *Given $\epsilon > 0$ and $x \in \mathcal{M}$ such that $\|\text{grad } f(x)\|_x > \epsilon$, let $g_h(x)$ be an approximate Riemannian gradient at x , satisfying Definition 2.1 for some constant $C_f > 0$. If $h \leq \frac{\epsilon}{5C_f}$, then*

$$\|g_h(x)\|_x \geq \frac{4\epsilon}{5},$$

and

$$\|\text{grad } f(x) - g_h(x)\|_x \leq \frac{1}{4} \|g_h(x)\|_x.$$

Proof. Definition 2.1 with $h \leq \frac{\epsilon}{5C_f}$ implies that

$$\|\text{grad } f(x) - g_h(x)\|_x \leq C_f h \leq \frac{\epsilon}{5}.$$

By the triangle inequality,

$$\epsilon < \|\text{grad } f(x)\|_x \leq \|\text{grad } f(x) - g_h(x)\|_x + \|g_h(x)\|_x \leq \frac{\epsilon}{5} + \|g_h(x)\|_x,$$

hence

$$\|g_h(x)\|_x \geq \frac{4\epsilon}{5}$$

and

$$\|\text{grad } f(x) - g_h(x)\|_x \leq \frac{\epsilon}{5} \leq \frac{1}{4} \|g_h(x)\|_x.$$

□

The next lemma uses this bound on the relative error of the gradient approximation to provide a derivative-free Armijo-like decrease of the objective under Lipschitz-smoothness of the objective (i.e., A1).

Lemma 2.3 (Adapted from Lemma 1 in [24]). *Suppose that A1 holds. Let $x \in \mathcal{M}$ and $g_h(x) \in T_x\mathcal{M}$ be such that*

$$\|\text{grad } f(x) - g_h(x)\|_x \leq \frac{1}{4} \|g_h(x)\|_x, \quad (3)$$

and let $x^+ = R_x(-\frac{1}{\sigma}g_h(x))$ for $\sigma \geq L_{\mathcal{M}}$, then

$$f(x) - f(x^+) \geq \frac{1}{4\sigma} \|g_h(x)\|_x^2.$$

Proof. By A1, there holds

$$f(x^+) = f\left(R_x\left(-\frac{1}{\sigma}g_h(x)\right)\right) \quad (4)$$

$$\begin{aligned} &\leq f(x) - \frac{1}{\sigma}\langle \text{grad } f(x), g_h(x) \rangle_x + \frac{L_{\mathcal{M}}}{2\sigma^2} \|g_h(x)\|_x^2 \\ &= f(x) + \frac{1}{\sigma}\langle g_h(x) - \text{grad } f(x), g_h(x) \rangle_x + \frac{L_{\mathcal{M}} - 2\sigma}{2\sigma^2} \|g_h(x)\|_x^2 \\ &\leq f(x) + \frac{1}{\sigma} \|g_h(x) - \text{grad } f(x)\|_x \|g_h(x)\|_x + \frac{L_{\mathcal{M}} - 2\sigma}{2\sigma^2} \|g_h(x)\|_x^2. \end{aligned} \quad (5)$$

Combining this with (3) gives

$$f(x) - f(x^+) \geq \left(\frac{2\sigma - L_{\mathcal{M}}}{2\sigma^2} - \frac{1}{4\sigma}\right) \|g_h(x)\|_x^2 = \frac{3\sigma - 2L_{\mathcal{M}}}{4\sigma^2} \|g_h(x)\|_x^2 \geq \frac{1}{4\sigma} \|g_h(x)\|_x^2.$$

□

3 A novel finite-difference DFRO method

We are now ready to present our main DFRO method, which is novel even in the Euclidean setting. The main specificity of this method is an adaptive scheme for the finite-difference accuracy and stepsize, that exploits information gathered on the objective over past iterations. In particular, our method does not require knowing the smoothness constant $L_{\mathcal{M}}$ in A1, in contrast with existing DFRO methods that use smoothness constants in the stepsize computation [26, 32].

Algorithm 1: Intrinsic Riemannian finite-difference method (Int-RFD)

Step 0. Let $x_0 \in \mathcal{M}$, $\sigma_0 > 0$, $\tau_0 \geq \sigma_0$, $\epsilon > 0$, set $k := 0$. Let d be the dimension of \mathcal{M} .

Step 1.1. Let $h_k = \frac{2\epsilon}{5\sqrt{d}\tau_k}$, compute the intrinsic finite-difference gradient

$$g_{h_k}(x_k) = \sum_{l=1}^d \frac{f(R_{x_k}(h_k e_l(x_k))) - f(x_k)}{h_k} e_l(x_k), \quad (6)$$

where $\{e_1(x_k), \dots, e_d(x_k)\}$ is an orthonormal basis of $T_{x_k} \mathcal{M}$.

Step 1.2. If

$$\|g_{h_k}(x_k)\|_{x_k} < \frac{4\epsilon}{5}, \quad (7)$$

set $x_{k+1} = x_k$, $\sigma_{k+1} = \sigma_k$, $\tau_{k+1} = 2\tau_k$ and $k := k + 1$, go back to Step 1.1.

Step 1.3. If

$$f(x_k) - f\left(R_{x_k}\left(-\frac{1}{\sigma_k}g_{h_k}(x_k)\right)\right) \geq \frac{1}{4\sigma_k} \|g_{h_k}(x_k)\|_{x_k}^2, \quad (8)$$

set $x_{k+1} = R_{x_k}\left(-\frac{1}{\sigma_k}g_{h_k}(x_k)\right)$, $\sigma_{k+1} = \frac{\sigma_k}{2}$, $\tau_{k+1} = \tau_k$ and $k := k + 1$, go back to Step 1.1.

Step 1.4 Set $x_{k+1} = x_k$ and $\sigma_{k+1} = 2\sigma_k$. If $\sigma_{k+1} > \tau_k$ set $\tau_{k+1} = 2\tau_k$ and $k := k + 1$, go back to Step 1.1. Otherwise set $\tau_{k+1} = \tau_k$ and $k := k + 1$, go back to Step 1.3.

Our intrinsic Riemannian finite-difference method, relying on the intrinsic finite-difference scheme (1), is presented in Algorithm 1. The computation of a finite-difference approximation to the Riemannian gradient is done at Step 1.1, and the search for a stepsize that ensures a derivative-free Armijo-type sufficient decrease condition is in Step 1.3. The approximate gradient accuracy h_k and stepsize used in these two steps are determined by two parameters τ_k and σ_k that estimate the constant $L_{\mathcal{M}}$ in A1 during optimization. These two parameters (whose update strategy is described in the next paragraph) differ in their roles: the first estimate τ_k is conservative and ensures sufficiently accurate finite-difference approximations whereas the second estimate σ_k is an optimistic estimate used for the stepsize. When it is strictly smaller than τ_k , it allows the method to perform larger steps. The use of two distinct estimates is the key modification with respect to [24] and was inspired by a similar idea proposed for trust-region derivative-free methods [14].

The two sequences of parameters σ_k and τ_k are updated as follows. When the norm of the estimated gradient is too small at Step 1.2, the parameter τ_k is increased in order to refine the accuracy of the gradient approximation. Indeed, by Lemma 2.2, (7) either implies that x_k is an ϵ -critical point of f (in which case the algorithm succeeded), or that the finite-difference gradient has insufficient accuracy to ensure a guaranteed decrease of the objective by Lemma 2.3, we then increase τ_k . Conversely, when (7) does not hold, we enter Step 1.3, where the sufficient decrease condition (8) is checked for a stepsize $1/\sigma_k$. If the decrease condition is satisfied, we accept the step and decrease σ_k to allow for larger steps in future iterations. If not, the stepsize is decreased (by increasing σ_k) and the condition is checked again with this smaller stepsize. Note that, according to Lemma 2.3 and provided τ_k is large enough, (8) will be satisfied as soon as $\sigma_k \geq L_{\mathcal{M}}$. If the updated value σ_{k+1} becomes larger than τ_k , the failure of the decrease condition is likely due to an inaccurate finite-difference gradient which prevents Lemma 2.3 to hold, and we increase τ_k and recompute an approximate gradient of higher accuracy in Step 1.1.

We next provide a theoretical analysis of the convergence of Algorithm 1. For this, we organize the iterations in Algorithm 1 in four categories, depending on whether they provide a decrease of the objective or not, and if they do not, the reason of failure. We thus define four disjoint sets $\mathcal{U}^{(1)}$, $\mathcal{U}^{(2)}$, $\mathcal{U}^{(3)}$ and \mathcal{S} as follows.

- **Unsuccessful iterations of type I** ($k \in \mathcal{U}^{(1)}$): these iterations do not provide an objective decrease, and the failure comes from an approximate gradient $g_{h_k}(x_k)$ whose norm is too small, satisfying (7). As discussed above, in this case we increase τ_k .
- **Unsuccessful iterations of type II** ($k \in \mathcal{U}^{(2)}$): the norm of the approximate gradient is large enough to enter Step 1.3, but the sufficient decrease (8) does not hold. We then increase σ_k , and get $\sigma_{k+1} > \tau_k$. We thus suspect a lack of accuracy of the finite-difference gradient explaining the failure of the decrease condition, and increase τ_k .
- **Unsuccessful iterations of type III** ($k \in \mathcal{U}^{(3)}$): In Step 1.3, the sufficient decrease (8) does not hold. We then increase σ_k and get $\sigma_{k+1} \leq \tau_k$. In this case, we keep τ_k constant and avoid the computation of a new approximate gradient.
- **Successful iterations** ($k \in \mathcal{S}$): The sufficient decrease condition (8) in Step 1.3 is satisfied and the step is accepted. The stepsize parameter σ_k is decreased.

Let us write the number of iterations to reach ϵ -criticality as

$$T(\epsilon) = \inf\{k \in \mathbb{N} \mid \|\text{grad } f(x_k)\|_{x_k} \leq \epsilon\}.$$

The subsets of successful and unsuccessful iterations before reaching ϵ -criticality are thus:

$$\mathcal{S}_{T(\epsilon)} = \mathcal{S} \cap \{0, 1, \dots, T(\epsilon) - 1\}, \quad \text{and} \quad \mathcal{U}_{T(\epsilon)}^{(i)} = \mathcal{U}^{(i)} \cap \{0, 1, \dots, T(\epsilon) - 1\}, \quad i = 1, 2, 3.$$

We are now ready to derive a worst-case complexity analysis of Algorithm 1 in terms of function evaluations and retractions. It turns out that Algorithm 1 is guaranteed to converge to an ϵ -critical point, assuming that τ_k is bounded from above, as stated by the following result.

Theorem 3.1. *Suppose that A2 holds. Let us assume that there exists $\tau_{\max} > 0$ such that*

$$\tau_k \leq \tau_{\max}, \quad \forall k = 0, 1, \dots, T(\epsilon) - 1,$$

then the number of function and retraction evaluations performed by Algorithm 1 to reach an ϵ -critical point of f (i.e., a point $x^ \in \mathcal{M}$ such that $\|\text{grad } f(x^*)\|_{x^*} \leq \epsilon$), written $FE(\epsilon)$ and $RE(\epsilon)$, satisfy*

$$FE(\epsilon) \leq (FE_g + 2) \left[2 \log_2 \left(\frac{\tau_{\max}}{\tau_0} \right) + \frac{25}{2} \tau_{\max} (f(x_0) - f_{\text{low}}) \epsilon^{-2} \right], \quad (9)$$

$$RE(\epsilon) \leq (RE_g + 1) \left[2 \log_2 \left(\frac{\tau_{\max}}{\tau_0} \right) + \frac{25}{2} \tau_{\max} (f(x_0) - f_{\text{low}}) \epsilon^{-2} \right], \quad (10)$$

where FE_g and RE_g are respectively the number of new function evaluations and retractions to compute one finite-difference gradient $g_h(x)$.

Proof. Let $(x_k)_{k=0}^{T(\epsilon)-1}$ be a sequence of iterates generated by Algorithm 1. We start by finding upper bounds on the cardinality of the sets $\mathcal{S}_{T(\epsilon)}$, $\mathcal{U}_{T(\epsilon)}^{(1)}$, $\mathcal{U}_{T(\epsilon)}^{(2)}$ and $\mathcal{U}_{T(\epsilon)}^{(3)}$. Note that, by A2,

$$f(x_0) - f_{\text{low}} \geq \sum_{k=0}^{T(\epsilon)-1} (f(x_k) - f(x_{k+1})) = \sum_{k \in \mathcal{S}_{T(\epsilon)}} (f(x_k) - f(x_{k+1})).$$

Moreover, each successful iteration $k \in \mathcal{S}_{T(\epsilon)}$ satisfies

$$f(x_k) - f(x_{k+1}) \geq \frac{1}{4\sigma_k} \|g_{h_k}(x_k)\|_{x_k}^2 \geq \frac{1}{4\sigma_k} \left(\frac{4\epsilon}{5} \right)^2.$$

Since $\sigma_k \leq \tau_k$ for all k , it follows from our assumption that

$$f(x_0) - f_{\text{low}} \geq \frac{4\epsilon^2}{25\tau_{\max}} |\mathcal{S}_{T(\epsilon)}|$$

or equivalently,

$$|\mathcal{S}_{T(\epsilon)}| \leq \frac{25}{4} \tau_{\max} (f(x_0) - f_{\text{low}}) \epsilon^{-2}.$$

On the other hand, note that the update rules of τ_k and σ_k imply

$$\tau_{T(\epsilon)-1} = 2^{|\mathcal{U}_{T(\epsilon)}^{(1)}| + |\mathcal{U}_{T(\epsilon)}^{(2)}|} \tau_0 \leq \tau_{\max} \quad \text{and} \quad \sigma_{T(\epsilon)-1} \leq 2^{|\mathcal{U}_{T(\epsilon)}^{(2)}| + |\mathcal{U}_{T(\epsilon)}^{(3)}| - |\mathcal{S}_{T(\epsilon)}|} \tau_0 \leq \tau_{\max}.$$

We deduce that

$$|\mathcal{U}_{T(\epsilon)}^{(1)}| + |\mathcal{U}_{T(\epsilon)}^{(2)}| \leq \log_2 \left(\frac{\tau_{\max}}{\tau_0} \right) \quad \text{and} \quad |\mathcal{U}_{T(\epsilon)}^{(2)}| + |\mathcal{U}_{T(\epsilon)}^{(3)}| \leq |\mathcal{S}_{T(\epsilon)}| + \log_2 \left(\frac{\tau_{\max}}{\tau_0} \right).$$

Thus,

$$\begin{aligned}
T(\epsilon) &= |\mathcal{U}_{T(\epsilon)}^{(1)}| + |\mathcal{U}_{T(\epsilon)}^{(2)}| + |\mathcal{U}_{T(\epsilon)}^{(3)}| + |\mathcal{S}_{T(\epsilon)}| \\
&\leq 2 \log_2 \left(\frac{\tau_{\max}}{\tau_0} \right) + \frac{25}{2} \tau_{\max} (f(x_0) - f_{\text{low}}) \epsilon^{-2}.
\end{aligned} \tag{11}$$

Since each iteration of Algorithm 1 requires at most $(FE_g + 2)$ function evaluations and $(RE_g + 1)$ retractions, it follows from (11), that the bounds (9) and (10) hold. \square

Theorem 3.1 assumes that τ_k is bounded for all k . We next prove that this condition holds under A1. The following lemma shows that the finite-difference gradient defined in (6) is an approximate Riemannian gradient in the sense of Definition 2.1, with $C_f = L_{\mathcal{M}}\sqrt{d}/2$. This allows us to apply Lemma 2.2 and Lemma 2.3 to guarantee a sufficiently large norm of the finite-difference gradient when x_k is not an ϵ -critical point of f , as well as a sufficient decrease of the objective function when τ_k and σ_k are sufficiently large. As a result, the steps are eventually accepted systematically and τ_k does not increase beyond a fixed threshold.

Lemma 3.2. *Assume that A1 holds, and let $g_h(x)$ be the intrinsic finite difference gradient defined in (6). Then*

$$\|\text{grad } f(x) - g_h(x)\|_x \leq \frac{L_{\mathcal{M}}\sqrt{d}}{2} h.$$

Proof. Let $\{e_1(x), \dots, e_d(x)\}$ be an orthonormal basis of $T_x\mathcal{M}$, then any tangent vector $\eta \in T_x\mathcal{M}$ can be written as

$$\eta = \sum_{l=1}^d \langle \eta, e_l(x) \rangle_x e_l(x).$$

In particular,

$$\text{grad } f(x) = \sum_{l=1}^d \langle \text{grad } f(x), e_l(x) \rangle_x e_l(x).$$

There follows:

$$\begin{aligned}
\|g_h(x) - \text{grad } f(x)\|_x^2 &= \left\| \sum_{l=1}^d \frac{f(R_x(he_l(x))) - f(x)}{h} e_l(x) - \text{grad } f(x) \right\|_x^2 \\
&= \frac{1}{h^2} \sum_{l=1}^d (f(R_x(he_l(x))) - f(x) - h \langle \text{grad } f(x), e_l(x) \rangle_x)^2 \\
&\leq d \left(\frac{L_{\mathcal{M}}h}{2} \right)^2,
\end{aligned}$$

where we used the fact that, by A1, for $l = 1, \dots, d$, there holds

$$|f(R_x(he_l(x))) - f(x) - h \langle \text{grad } f(x), e_l(x) \rangle_x| \leq \frac{L_{\mathcal{M}}h^2}{2} \|e_l(x)\|_x^2 = \frac{L_{\mathcal{M}}h^2}{2}.$$

\square

The next result provides an upper bound on τ_k for $k \in \{0, \dots, T(\epsilon) - 1\}$.

Lemma 3.3. *Assume that A1 holds, and let $(x_k)_{k=0}^{T(\epsilon)-1}$ be a sequence generated by Algorithm 1. Then, for all $k \in \{0, \dots, T(\epsilon) - 1\}$, there holds*

$$\tau_k \leq \max\{\tau_0, 4L_{\mathcal{M}}\} =: \tau_{\max}^{\text{Int}}. \quad (12)$$

Proof. We proceed by induction on k . For $k = 0$, (12) holds trivially. Assume that (12) holds for some $k \in \{0, \dots, T(\epsilon) - 2\}$. We show that (12) also holds for $k + 1$ by considering the following three cases.

Case 1: $k \in \mathcal{U}^{(1)}$ (i.e., (7) holds).

In this case, we must have

$$\tau_k < L_{\mathcal{M}}, \quad (13)$$

since otherwise Lemma 3.2 and Lemma 2.2, with $C_f = L_{\mathcal{M}}\sqrt{d}/2$, would imply that $\|g_{h_k}(x_k)\|_x \geq 4\epsilon/5$, contradicting (7). Using (13) and the update rule for τ_k in Step 1.2 of Algorithm 1, we obtain

$$\tau_{k+1} = 2\tau_k < 2L_{\mathcal{M}} \leq \tau_{\max}^{\text{Int}},$$

so that (12) holds for $k + 1$.

Case 2: $k \in \mathcal{U}^{(2)}$ (i.e., (8) fails and $\sigma_{k+1} > \tau_k$).

In this case, we must have

$$\tau_k < 2L_{\mathcal{M}}, \quad (14)$$

since otherwise

$$2\sigma_k = \sigma_{k+1} > \tau_k \geq 2L_{\mathcal{M}},$$

which implies $\tau_k \geq \sigma_k > L_{\mathcal{M}}$. By Lemma 2.3 and Lemma 2.2, with $C_f = L_{\mathcal{M}}\sqrt{d}/2$, this would imply that (8) holds, contradicting the definition of $\mathcal{U}^{(2)}$. Therefore, using (14) and the update rule for τ_k in Step 1.4 of Algorithm 1, we obtain

$$\tau_{k+1} = 2\tau_k < 4L_{\mathcal{M}} \leq \tau_{\max}^{\text{Int}},$$

and thus (12) holds for $k + 1$.

Case 3: $k \in \mathcal{U}^{(3)} \cup \mathcal{S}$.

In this case, the update rule for τ_k in Step 1.4 and the induction hypothesis imply that $\tau_{k+1} = \tau_k \leq \tau_{\max}^{\text{Int}}$. Hence, (12) also holds in this case. \square

Now we can provide a complete complexity result for Algorithm 1.

Corollary 3.4. *Under A1 and A2, Algorithm 1 finds an ϵ -critical point of (P) (i.e., a point $x^* \in \mathcal{M}$ such that $\|\text{grad } f(x^*)\|_{x^*} \leq \epsilon$) with the following worst-case complexity bounds:*

$$\begin{aligned} FE(\epsilon) &\leq (d+2) \left[2 \log_2 \left(\frac{\max\{\tau_0, 4L_{\mathcal{M}}\}}{\tau_0} \right) + \frac{25}{2} \max\{\tau_0, 4L_{\mathcal{M}}\} (f(x_0) - f_{\text{low}}) \epsilon^{-2} \right], \\ RE(\epsilon) &\leq (d+1) \left[2 \log_2 \left(\frac{\max\{\tau_0, 4L_{\mathcal{M}}\}}{\tau_0} \right) + \frac{25}{2} \max\{\tau_0, 4L_{\mathcal{M}}\} (f(x_0) - f_{\text{low}}) \epsilon^{-2} \right], \end{aligned}$$

Proof. It follows directly from Theorem 3.1 and Lemma 3.3, noting that the computation in (6) requires $FE_g = d$ function evaluations in addition to $f(x_k)$, as well as $RE_g = d$ retractions. \square

Comparing Corollary 3.4 with the results from [24], we recover the linear dependency in problem dimension (the dimension d of the manifold in our case). Note also that choosing τ_0 larger typically improves the logarithmic terms in the bounds, which is expected since it corresponds to more accurate Riemannian gradient approximations, and equivalently, fewer unsuccessful iterations of type I and III. On the other hand, increasing σ_0 (and so τ_0) slows down the convergence in general since step size may be too conservative. We therefore recommend to implement the method with τ_0 much larger than σ_0 . Note also that the number of retractions scales similarly to the number of function evaluations, we will show next that it is possible to reduce it drastically when \mathcal{M} is a Riemannian submanifold of \mathbb{R}^n .

4 An extrinsic finite-difference scheme on manifolds

Let us assume now that the objective function f in (P) is defined over an n -dimensional Euclidean space \mathcal{E} , but that its minimization is restricted to a d -dimensional Riemannian submanifold \mathcal{M} of \mathcal{E} . We make the following assumption on f .

A3. The function $f : \mathcal{E} \rightarrow \mathbb{R}$ is $L_{\mathcal{E}}$ -smooth, i.e., for all $x \in \mathcal{E}$,

$$|f(x + \eta) - f(x) - \langle \nabla f(x), \eta \rangle| \leq \frac{L_{\mathcal{E}}}{2} \|\eta\|^2, \quad \forall \eta \in \mathcal{E},$$

where $\langle \cdot, \cdot \rangle$ and $\|\cdot\|$ are the usual Euclidean inner product and norm, respectively.

Remark 4.1. According to [9], when \mathcal{M} is a compact Riemannian submanifold of \mathcal{E} , Assumption A3 implies Assumption A1 for some constant $L_{\mathcal{M}}$ that depends on its Euclidean counterpart, on the diameter of the manifold and on some norm of the differential of the retraction.

The goal of this section is to propose an alternative finite-difference scheme, presented in Algorithm 2, that relies on the embedding of \mathcal{M} in \mathcal{E} to avoid the need of computing retractions in the intrinsic finite-difference scheme used in last section. This is motivated by the observation that the computational cost of retractions can be substantial in practice. For instance, on the orthogonal group, retractions involve matrix factorizations such as the QR or polar decompositions, whose computation requires a cubic number of floating point operations in terms of the size of the matrix, becoming prohibitive in high dimensions (see, e.g., [3, Example 4.1.2]).

The convergence analysis of Algorithm 2 relies on the following result, which shows that the extrinsic finite-difference scheme (15) leads to an approximate Riemannian gradient in the sense of Definition 2.1, with $C_f = L_{\mathcal{E}}\sqrt{d}/2$.

Lemma 4.2. Let $f : \mathcal{E} \rightarrow \mathbb{R}$ and let \mathcal{M} be a d -dimensional Riemannian submanifold of \mathcal{E} in (P). Assume that A3 hold, and let $g_h(x)$ be the extrinsic finite-difference gradient of f defined in (15), then

$$\|\text{grad } f(x) - g_h(x)\|_x \leq \frac{L_{\mathcal{E}}\sqrt{d}}{2} h,$$

where $\|\cdot\|_x$ coincides with the Euclidean norm, since \mathcal{M} is a Riemannian submanifold of \mathcal{E} .

Algorithm 2: Extrinsic Riemannian finite-difference method (Ext-RFD)

Step 0. Let $x_0 \in \mathcal{M}$, $\sigma_0 > 0$, $\tau_0 > \sigma_0$, $\epsilon > 0$, and set $k := 0$. Let d be the dimension of \mathcal{M} .

Step 1.1. Let $h_k = \frac{2\epsilon}{5\sqrt{d}\tau_k}$ and compute the approximate Riemannian gradient

$$g_{h_k}(x_k) = \sum_{l=1}^d \frac{f(x_k + h_k e_l(x_k)) - f(x_k)}{h_k} e_l(x_k), \quad (15)$$

where $\{e_1(x_k), \dots, e_d(x_k)\}$ is an orthonormal basis of $T_{x_k} \mathcal{M}$.

Steps 1.2–1.4: Same as Steps 1.2–1.4 of Algorithm 1.

Proof. Since \mathcal{M} is a Riemannian submanifold of \mathcal{E} , the Riemannian gradient $\text{grad } f(x)$ is the orthogonal projection of its Euclidean counterpart $\nabla f(x)$ onto the tangent space $T_x \mathcal{M}$, i.e., $\text{grad } f(x) = \text{Proj}_{T_x \mathcal{M}} \nabla f(x)$ for all $x \in \mathcal{M}$, see [3, eq. 3.37]. Let $\{e_1(x), \dots, e_d(x)\}$ be an orthonormal basis of $T_x \mathcal{M}$, then any tangent vector $\eta \in T_x \mathcal{M}$ can be written as

$$\eta = \sum_{l=1}^d \langle \eta, e_l(x) \rangle e_l(x).$$

Particularizing this equality to the Riemannian gradient gives

$$\text{grad } f(x) = \sum_{l=1}^d \langle \text{grad } f(x), e_l(x) \rangle e_l(x) = \sum_{l=1}^d \langle \text{Proj}_{T_x \mathcal{M}} \nabla f(x), e_l(x) \rangle e_l(x) = \sum_{l=1}^d \langle \nabla f(x), e_l(x) \rangle e_l(x),$$

by definition of the orthogonal projection. There follows:

$$\begin{aligned} \|g_h(x) - \text{grad } f(x)\|^2 &= \left\| \sum_{l=1}^d \left(\frac{f(x + h e_l(x)) - f(x)}{h} - \langle \nabla f(x), e_l(x) \rangle \right) e_l(x) \right\|^2 \\ &= \frac{1}{h^2} \sum_{l=1}^d (f(x + h e_l(x)) - f(x) - h \langle \nabla f(x), e_l(x) \rangle)^2 \\ &\leq d \left(\frac{L_{\mathcal{E}} h}{2} \right)^2, \end{aligned}$$

where the inequality results from Assumption A3, which implies that, for $l = 1, \dots, d$

$$|f(x + h e_l(x)) - f(x) - h \langle \nabla f(x), e_l(x) \rangle| \leq \frac{L_{\mathcal{E}} h^2}{2} \|e_l(x)\|^2 = \frac{L_{\mathcal{E}} h^2}{2}.$$

□

Similarly as in the previous section, we can use this upper bound on the relative error on gradient approximations to derive upper bounds on the parameters σ_k and τ_k in Algorithm 2, for $k = 0, \dots, T(\epsilon) - 1$. For this, we rely on the four sets of iterations $\mathcal{U}^{(1)}, \mathcal{U}^{(2)}, \mathcal{U}^{(3)}, \mathcal{S}$, defined in Section 3 for Algorithm 1, and whose definition can be identically transferred to Algorithm 2.

Lemma 4.3. *Let $f : \mathcal{E} \rightarrow \mathbb{R}$ and let \mathcal{M} be a d -dimensional Riemannian submanifold of \mathcal{E} . Assume that A1 and A3 hold and let $(x_k)_{k=0}^{T(\epsilon)-1}$ be a sequence generated by Algorithm 2. Then, for all $k \in \{0, \dots, T(\epsilon) - 1\}$,*

$$\tau_k \leq \max\{\tau_0, 4L_{\mathcal{M}}, 4L_{\mathcal{E}}\} =: \tau_{\max}^{\text{Ext}}.$$

Proof. It follows exactly as in the proof of Lemma 3.3, by using Lemma 4.2 with $C_f = L_{\mathcal{E}}\sqrt{d}/2$. \square

Remark 4.4. *Note that, while the sufficient decrease of the objective is only governed by the Lipschitz constant of the gradient along the manifold (written $L_{\mathcal{M}}$, see Lemma 2.3), the minimum accuracy of the extrinsic finite-difference gradient required in Lemma 4.2 and Lemma 2.2 depends on the Lipschitz constant $L_{\mathcal{E}}$ of the gradient of the objective in the Euclidean space, which will also account for variations of the objective in directions normal to the manifold. In other words, while by Remark 4.1, Lipschitz-smoothness of f in \mathcal{E} ensures Lipschitz-smoothness of f on the manifold when the latter is compact, $L_{\mathcal{E}}$ may be much larger than $L_{\mathcal{M}}$. This is accounted for in Algorithm 2 by the fact that repeated failures of the descent condition in Step 1.3 will lead to a progressive increase of σ_k , possibly going beyond $L_{\mathcal{M}}$, triggering in Step 1.4 an associated increase of τ_k , until the extrinsic finite-difference scheme gets sufficiently accurate to allow for a decrease of the objective in Step 1.3, which will occur as soon as $\tau_k \geq L_{\mathcal{E}}$.*

Note that Theorem 3.1 also applies to Algorithm 2. Since in Algorithm 2, each computation of the finite-difference gradient requires $FE_g = d$ function evaluations in addition to $f(x_k)$, and $RE_g = 0$ retractions, we obtain the following complexity result.

Corollary 4.5. *Let $f : \mathcal{E} \rightarrow \mathbb{R}$ and let \mathcal{M} be a d -dimensional Riemannian submanifold of \mathcal{E} . Assume that A1, A2 and A3 hold. Then, Algorithm 2 finds an ϵ -critical point of (P) (i.e., a point $x^* \in \mathcal{M}$ such that $\|\text{grad } f(x^*)\|_{x^*} \leq \epsilon$) with the following worst-case complexity bounds:*

$$\begin{aligned} FE(\epsilon) &\leq (d+2) \left[2 \log_2 \left(\frac{\max\{\tau_0, 4L_{\mathcal{E}}, 4L_{\mathcal{M}}\}}{\tau_0} \right) + \frac{25}{2} \max\{\tau_0, 4L_{\mathcal{E}}, 4L_{\mathcal{M}}\} (f(x_0) - f_{\text{low}})\epsilon^{-2} \right], \\ RE(\epsilon) &\leq 2 \log_2 \left(\frac{\max\{\tau_0, 4L_{\mathcal{E}}, 4L_{\mathcal{M}}\}}{\tau_0} \right) + \frac{25}{2} \max\{\tau_0, 4L_{\mathcal{E}}, 4L_{\mathcal{M}}\} (f(x_0) - f_{\text{low}})\epsilon^{-2}. \end{aligned}$$

Proof. It follows directly from Theorem 3.1 combined with Lemma 4.3. \square

Note that Corollary 4.5 implies that $RE(\epsilon)$ does not depend on the problem dimension, because retractions are no longer used in the extrinsic finite-difference scheme. This feature makes the method particularly well-suited for large-scale problems where the retractions are expensive to compute. On the other hand, this may come at the cost of an increased number of function evaluations. Indeed, the smoothness constant $L_{\mathcal{E}}$ may be larger than $L_{\mathcal{M}}$, which means that more function evaluations may be needed to build sufficiently accurate finite-difference gradients. The worst-case complexity rates of the two methods proposed in this paper are summarized in Table 1.

	Int-RFD	Ext-RFD
$FE(\epsilon)$	$O(L_{\mathcal{M}}d\epsilon^{-2})$	$O(\max\{L_{\mathcal{M}}, L_{\mathcal{E}}\}d\epsilon^{-2})$
$RE(\epsilon)$	$O(L_{\mathcal{M}}d\epsilon^{-2})$	$O(\max\{L_{\mathcal{M}}, L_{\mathcal{E}}\}\epsilon^{-2})$

Table 1: Worst-case complexity bounds to reach an ϵ -critical point in terms of number of function evaluations $FE(\epsilon)$ and number of retractions $RE(\epsilon)$.

5 Numerical results

We compare here numerically our proposed algorithms to existing derivative-free methods. The first part of the section addresses the Euclidean setting, where we compare Algorithm 1 to the finite-difference method proposed in [24]. We then compare Algorithm 1 and Algorithm 2 to the DFRO method proposed in [30] on three common Riemannian optimization problems. The implementation of Algorithm 1 and Algorithm 2 are available on <https://github.com/taminiaut/Extrinsic-Intrinsic-Riemannian-Finite-Difference-Method>.

5.1 Experiments in the Euclidean setting

Our proposed algorithms (Algorithm 1 and Algorithm 2) rely on two sequences of parameters $(\sigma_k)_{k \geq 0}$, $(\tau_k)_{k \geq 0}$ to account for the smoothness of the objective in the stepsize and finite-difference accuracy. In contrast the Derivative-Free Quadratic Regularization Method (DFQRM), proposed in [24], only uses one such sequence of parameters (equivalently, σ_k is constrained to be equal to τ_k at each iteration). In this section, we show numerically that relying on two sequences results in a faster convergence compared to DFQRM. Our comparison is run on the OPM test set proposed in [25] which is a subset of 134 problems the well-known CUTEst collection [21]. Note also that, since, in the Euclidean setting, $\mathcal{M} = \mathbb{R}^n$ in (P), our proposed Int-RFD and Ext-RFD algorithms (respectively, Algorithm 1 and Algorithm 2) are equivalent, hence we only consider here the Int-RFD algorithm. Note that, for the sake of comparison, we do not use Hessian approximations in DFQRM, and choose exactly the same parameters for DFQRM (in particular, $\sigma_0 = 1$ and $\epsilon = 10^{-5}$) as for Algorithm 1 (with $\tau_0 = 100$).

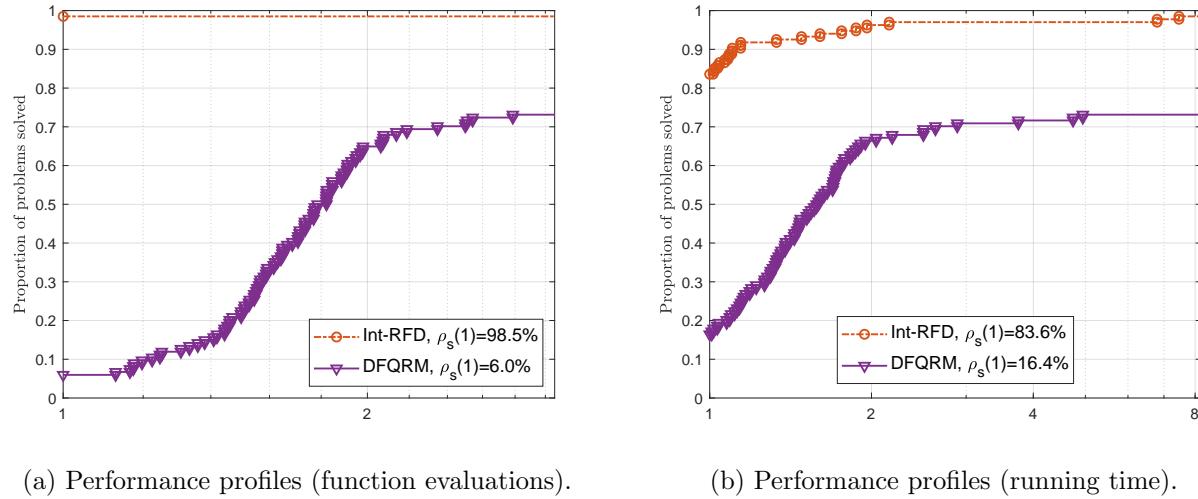
We use the performance profiles introduced in [16] to compare the two methods. Let \mathcal{P} be a set of problems, \mathcal{S} a set of solvers and $\eta \in (0, 1)$ a tolerance, we measure the computational cost $t_{p,s}$ (either in terms of number of function evaluations, or running time) required by any solver $s \in \mathcal{S}$ to find a point x that satisfies the following convergence test, on problem $p \in \mathcal{P}$ and for some initial iterate x_0 :

$$f(x_0) - f(x) \geq (1 - \eta)(f(x_0) - f_{\text{best}}), \quad (16)$$

where f_{best} is the best function value found by any of the solver in \mathcal{S} within a fixed maximum budget of $100(d + 1)$ function evaluations (where d is the dimension of the manifold). The performance profile curve of the solver s is then given by

$$\rho_s(\alpha) = \frac{1}{|\mathcal{P}|} \left| \left\{ p \in \mathcal{P} : \frac{t_{p,s}}{\min_{s' \in \mathcal{S}} t_{p,s'}} \leq \alpha \right\} \right|.$$

Performance profiles allow to compare simultaneously the efficiency and robustness of the methods. The value $\rho_s(1)$ indicates the proportion of problems that were solved, according to (16), by the method s with the smallest budget; more efficient methods are associated with a larger $\rho_s(1)$. On the other hand, more robust methods, i.e., methods that can solve more problems possibly with an increased budget, will have higher $\rho_s(\alpha)$ for α large. In our experiments, we chose $\eta = 10^{-3}$ in (16) and we relied on the implementation of performance profiles from [36]. The resulting performance profiles in terms of function evaluations and running time are presented in Figure 1a and Figure 1b, respectively. These plots show that using two distinct estimates makes the method better on almost all the problems and it is also better in terms of robustness.



(a) Performance profiles (function evaluations).

(b) Performance profiles (running time).

Figure 1: Numerical experiments in the Euclidean case on the OPM test set [25].

5.2 Experiments on Riemannian test sets

Secondly, we compare our methods with the RDSE-SB algorithm proposed in [30], on three smooth Riemannian optimization problems. We constructed our test set based on these problem implementations from [30], considering for each problem several random instances and several problem dimensions.

Algorithm 1 and Algorithm 2 are implemented with $\sigma_0 = 1$, $\tau_0 = 100$ and $\epsilon = 10^{-5}$. As all problems considered are defined on Riemannian submanifolds, orthonormal bases of the tangent space were generated by projecting random vectors of the ambient space into the tangent space, then applying the Gram-Schmidt procedure to these vectors to make them orthonormal.

For the RDSE-SB method we used the implementations provided in [30], while the manifold structures are handled thanks to the MANOPT toolbox [10]. We next describe the three test set problems selected from [30] that we used in this work. For each problem, we generate a set of random problem instances of different dimensions given in Table 2. To give the reader a better view on the associated manifold and ambient space dimensions (respectively d and n), we also include their values in Table 2.

Problem I : Top singular vectors. The problem of finding the m_3 leading left and right singular vectors (associated to the largest singular values) of a matrix $A \in \mathbb{R}^{m_1 \times m_2}$ is formulated as the Riemannian optimization problem

$$\min_{\substack{X \in St(m_1, m_3) \\ Y \in St(m_2, m_3)}} -\text{trace}(X^T A Y),$$

where $St(m, p) = \{X \in \mathbb{R}^{m \times p} \mid X^T X = I\}$ for $m \geq p$, is the Stiefel manifold, seen as a Riemannian submanifold of $\mathbb{R}^{m \times p}$.

Problem II : Dictionary learning. Let $Y \in \mathbb{R}^{m_1 \times m_2}$ be a matrix whose columns are m_1 -dimensional vectors that we would like to represent as well as possible by a sparse linear combination

of the columns of a dictionary $D \in \mathbb{R}^{m_1 \times m_3}$, that is, we search a dictionary $D \in \mathbb{R}^{m_1 \times m_3}$ and a (sparse) matrix of coefficients $C \in \mathbb{R}^{m_3 \times m_2}$ such that

$$\min_{\substack{D \in Ob(m_1, m_3) \\ C \in \mathbb{R}^{m_3 \times m_2}}} \|Y - DC\|_F^2 + \lambda \phi(C)$$

where $Ob(m, p) = \{X = (x_1 \dots x_p) \in \mathbb{R}^{m \times p} \mid \|x_i\|_2 = 1, \text{ for } i = 1, \dots, p\}$ is the oblique manifold, seen as a Riemannian submanifold of $\mathbb{R}^{m \times p}$, and $\phi(C) = \sum_{i=1}^p \sum_{j=1}^n \sqrt{C_{ij}^2 + \delta^2}$ is a smoothed l_1 norm that promotes sparsity. Note that the oblique manifold constraint allows getting rid of the invariance under scaling of the columns of D and rows of C . In our experiments, we set the regularization and smoothing parameters to $\lambda = 0.01$ and $\delta = 0.001$, respectively.

Problem III : Rotation synchronization. The goal here is to estimate m_2 rotation matrices $R_1, \dots, R_{m_2} \in \mathbb{R}^{m_1 \times m_1}$ from noisy measurements $H_{ij} \approx R_i R_j^{-1}$ for all $i, j = 1, \dots, m_2$, leading to the optimization problem

$$\min_{R_1, \dots, R_{m_2} \in SO(m_1)} \sum_{i=1}^{m_2} \sum_{j=1}^{i-1} \|R_i - H_{ij} R_j\|_F^2,$$

where $SO(m) = \{X \in \mathbb{R}^{m \times m} \mid X^T X = I, \det(X) = 1\}$ is the special orthogonal group, seen as a Riemannian submanifold of $\mathbb{R}^{m \times m}$.

Problem I					Problem II					Problem III				
Top singular vectors					Dictionary learning					Rotations synchronization				
m_1	m_2	m_3	n	d	m_1	m_2	m_3	n	d	m_1	m_2	m_3	n	d
2	2	1	4	2	2	3	1	5	4	2	2	\times	8	2
3	3	2	12	16	3	5	2	16	14	2	4	\times	16	4
5	5	2	20	14	4	6	3	30	27	2	6	\times	24	6
10	10	2	40	34	5	7	4	48	44	4	2	\times	32	12
15	15	2	60	54	6	8	5	70	65	4	4	\times	64	24
20	20	2	80	74	8	10	6	108	102	4	6	\times	96	36
30	30	2	120	114	10	12	7	154	147	6	2	\times	72	30
5	5	4	40	20	12	14	8	208	200	6	4	\times	144	60
10	10	4	80	60	14	16	10	300	290	6	6	\times	216	90
20	20	4	160	140	5	20	3	75	72	8	2	\times	128	56
30	30	4	240	220	7	20	5	135	130	8	4	\times	256	112
30	10	6	240	198	12	20	3	96	93	8	6	\times	384	168
30	15	6	270	228	3	20	5	115	110	10	2	\times	200	90
30	10	8	320	248	5	20	7	175	168	10	4	\times	400	180
30	15	8	360	288	3	20	12	276	264	10	6	\times	600	270

Table 2: Values of the parameters (m_1, m_2, m_3) of the problem instances included in our test set, and corresponding dimensions of the ambient space (n) and of the manifold (d).

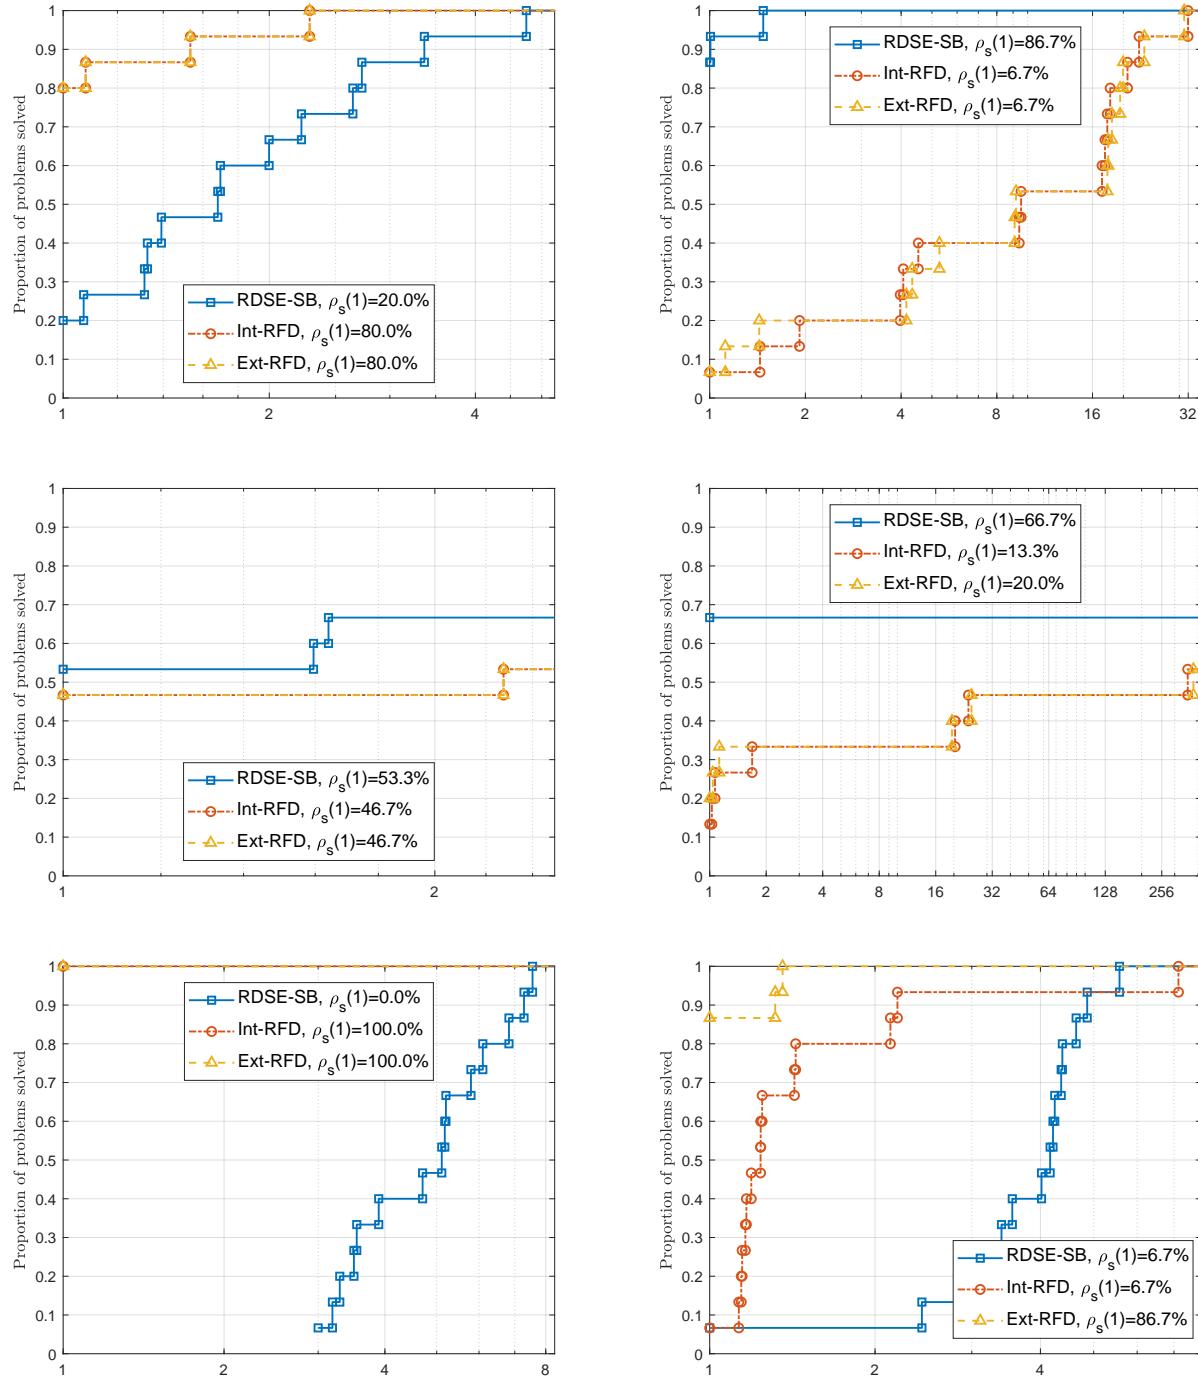


Figure 2: Experiments for the top singular vectors problem (top row), the dictionary learning problem (second row) and the rotation synchronization problem (bottom row), on a set of problem instances described in Table 2.

As can be seen in Figure 2a, the Int-RFD and Ext-RFD methods strongly outperform the RSDE-SB method of [30] in two of the three problems regarding accuracy vs function evaluations. For the dictionary learning problem, higher efficiency and robustness were achieved by the RSDE-SB method, though no method manages to solve all problem instances. We also notice that the Ext-RFD methods remains robust with respect to the Int-RFD method, despite having a worse complexity bound for function evaluations. Considering the additional computational cost of handling the manifold structure, we see in Figure 2b that the comparison between Int-RFD and Ext-RFD in terms of running time depends on the manifold considered. When the manifold dimension is much smaller than that of the ambient space ($d \ll n$), substantial benefits can be achieved by the Ext-RFD method, see the performance profile of the rotation synchronization problem. The comparable running times of Int-RFD and Ext-RFD on other problems is due to the fact that the Ext-RFD method still requires an orthonormal basis of the tangent space, which is computed here numerically using Gram-Schmidt; the computational cost of the later dominates the cost of retractions on some manifolds. Although closed-form expressions for orthonormal bases of the tangent space exist for all manifolds considered in this section, we include their computation in the cost to present numerical results that reflect performance in generic Riemannian optimization problems.

6 Conclusions

In this paper, we introduced a novel Riemannian finite-difference method that generalizes the (Euclidean) derivative-free method from [24], relying on two sequences approximating the smoothness constant of the problem in an optimistic and conservative way, respectively. This strategy, which was inspired from [14], amounts to use different smoothness constant estimates for the stepsize than for the finite-difference gradient accuracy. The resulting algorithm, which is novel even in the Euclidean setting, was showed numerically to converge faster than the original method proposed in [24]. The second contribution of this work is to extend this finite-difference algorithm to the Riemannian setting. For this, we relied on the usual definition of finite differences on Riemannian manifolds, leading to our Int-RFD algorithm presented in Algorithm 1. We derive worst-case complexity bounds for Int-RFD, showing that it reaches an ϵ -critical point after at most $\mathcal{O}(d\epsilon^{-2})$ function evaluations and retractions. Since retractions are often computationally intensive, we propose a variant of our algorithm, which we call Ext-RFD (see Algorithm 2), that relies on an extrinsic finite-difference scheme, for optimization over Riemannian submanifolds of a Euclidean space \mathcal{E} . We prove that this second method reaches an ϵ -critical point after at most $\mathcal{O}(d\epsilon^{-2})$ function evaluations, while the number of retractions is now independent of problem dimension. The numerical experiments indicate an improvement over the derivative-free method proposed in [24] in the Euclidean setting and that the extrinsic method significantly outperforms the intrinsic approach when the manifold dimension is much smaller than the dimension of the embedding space. Nevertheless, one should be careful that the worst-case complexity of the low cost method involves the Euclidean smoothness constant $L_{\mathcal{E}}$ which can be substantially larger than its Riemannian counterpart when the objective varies strongly in directions normal to the manifold in the embedding space.

References

[1] P. Ablin and G. Peyré. Fast and accurate optimization on the orthogonal manifold without retraction. In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera, editors, *Proceedings*

of The 25th International Conference on Artificial Intelligence and Statistics, volume 151 of *Proceedings of Machine Learning Research*, pages 5636–5657. PMLR, 2022.

- [2] P. Ablin, S. Vary, B. Gao, and P.-A. Absil. Infeasible Deterministic, Stochastic, and Variance-Reduction Algorithms for Optimization under Orthogonality Constraints. *Journal of Machine Learning Research*, 25(389):1–38, 2024.
- [3] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2008.
- [4] P.-A. Absil and J. Malick. Projection-like Retractions on Matrix Manifolds. *SIAM Journal on Optimization*, 22(1):135–158, 2012.
- [5] P.-A. Absil and I. V. Oseledets. Low-rank retractions: a survey and new results. *Computational Optimization and Applications*, 62(1):5–29, 2014.
- [6] C. Audet and W. Hare. *Derivative-Free and Blackbox Optimization*. Springer Series in Operations Research and Financial Engineering. Springer International Publishing, Cham, 2017.
- [7] N. Boumal. Riemannian Trust Regions with Finite-Difference Hessian Approximations are Globally Convergent. In *Geometric Science of Information*, pages 467–475. Springer International Publishing, 2015.
- [8] N. Boumal. *An Introduction to Optimization on Smooth Manifolds*. Cambridge University Press, 2023.
- [9] N. Boumal, P.-A. Absil, and C. Cartis. Global rates of convergence for nonconvex optimization on manifolds. *IMA Journal of Numerical Analysis*, 39(1):1–33, 2019.
- [10] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre. Manopt, a Matlab Toolbox for Optimization on Manifolds. *Journal of Machine Learning Research*, 15:1455–1459, 2014.
- [11] B. Cavarretta, F. Goyens, C. W. Royer, and F. Yger. Complexity guarantees and polling strategies for Riemannian direct-search methods. *arXiv e-prints*, page arXiv:2511.15360, 2025.
- [12] R. Chakraborty, J. Bouza, J. H. Manton, and B. C. Vemuri. Manifoldnet: A deep neural network for manifold-valued data with applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(2), 2022.
- [13] A. R. Conn, K. Scheinberg, and L. N. Vicente. *Introduction to Derivative-Free Optimization*. Society for Industrial and Applied Mathematics, USA, 2009.
- [14] D. Davar and G. N. Grapiglia. TRFD: A Derivative-Free Trust-Region Method Based on Finite Differences for Composite Nonsmooth Optimization. *SIAM Journal on Optimization*, 35(3):1792–1821, 2025.
- [15] M. Dodangeh, L. N. Vicente, and Z. Zhang. On the optimal order of worst case complexity of direct search. *Optimization Letters*, 10(4):699–708, 2015.
- [16] E. D. Dolan and J. J. Moré. Benchmarking optimization software with performance profiles. *Mathematical Programming*, 91(2):201–213, 2002.

- [17] D. W. Dreisigmeyer. Direct Search Algorithms over Riemannian Manifolds, 2007. Technical Report LA-UR-06-7416, Los Alamos National Laboratory.
- [18] D. W. Dreisigmeyer. Direct Search Methods on Reductive Homogeneous Spaces. *Journal of Optimization Theory and Applications*, 176(3):585–604, 2018.
- [19] R. Simon Fong and P. Tino. *Population-Based Optimization on Riemannian Manifolds*. Springer International Publishing, 2022.
- [20] B. Gao, G. Hu, Y. Kuang, and X. Liu. An Orthogonalization-Free Parallelizable Framework for All-Electron Calculations in Density Functional Theory. *SIAM Journal on Scientific Computing*, 44(3):B723–B745, 2022.
- [21] N. I. M. Gould, D. Orban, and P. L. Toint. CUTEST: a Constrained and Unconstrained Testing Environment with safe threads for mathematical optimization. *Computational Optimization and Applications*, 60(3):545–557, 2015.
- [22] F. Goyens, P.-A. Absil, and F. Feppon. Geometric Design of the Tangent Term in Landing Algorithms for Orthogonality Constraints. In *Geometric Science of Information*, pages 133–141. Springer Nature Switzerland, 2025.
- [23] G. N. Grapiglia. Quadratic regularization methods with finite-difference gradient approximations. *Computational Optimization and Applications*, 85(3):683–703, 2022.
- [24] G. N. Grapiglia. Worst-case evaluation complexity of a derivative-free quadratic regularization method. *Optimization Letters*, 18:1–19, 2023.
- [25] S. Gratton and P. L. Toint. OPM, a collection of Optimization Problems in Matlab. *arXiv e-prints*, page arXiv:2112.05636, 2021.
- [26] C. He, Z. Pan, X. Wang, and B. Jiang. Riemannian accelerated zeroth-order algorithm: improved robustness and lower query complexity. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *ICML’24*, pages 17972–18009, 2024.
- [27] Z. Huang and L. Van Gool. A Riemannian Network for SPD Matrix Learning. *arXiv e-prints*, page arXiv:1608.04233v2, 2016.
- [28] N. Jaquier, L. Rozo, S. Calinon, and M. Bürger. Bayesian Optimization Meets Riemannian Manifolds in Robot Learning. *Proceedings of the Conference on Robot Learning*, PMLR 100:233–246, 2020.
- [29] H. R. Karbasian and B. C. Vermeire. Gradient-free aerodynamic shape optimization using large eddy simulation. *Computers and Fluids*, 232:105185, 2022.
- [30] V. Kungurtsev, F. Rinaldi, and D. Zeffiro. Retraction-Based Direct Search Methods for Derivative Free Riemannian Optimization. *Journal of Optimization Theory and Applications*, 203(2):1710–1735, 2023.
- [31] J. Larson, M. Menickelly, and S. M. Wild. Derivative-free optimization methods. *Acta Numerica*, 28:287–404, 2019.

- [32] J. Li, K. Balasubramanian, and S. Ma. Stochastic Zeroth-Order Riemannian Derivative Estimation and Optimization. *Mathematics of Operations Research*, 48(2):1183–1211, 2023.
- [33] X. Liu, N. Xiao, and Y.-x. Yuan. A Penalty-Free Infeasible Approach for a Class of Nonsmooth Optimization Problems Over the Stiefel Manifold. *Journal of Scientific Computing*, 99(2), 2024.
- [34] M. S. Louzeiro, G. N. Silva, J. Yuan, and D. Zhang. An Adaptive Cubic Regularization quasi-Newton Method on Riemannian Manifolds. *arXiv e-prints*, page arXiv:2402.12464, 2024.
- [35] A. I. Maass, C. Manzie, D. Nešić, J. H. Manton, and I. Shames. Tracking and Regret Bounds for Online Zeroth-Order Euclidean and Riemannian Optimization. *SIAM Journal on Optimization*, 32(2):445–469, 2022.
- [36] J. Moré and S. Wild. Benchmarking Derivative-Free Optimization Algorithms. *SIAM Journal on Optimization*, 20:172–191, 2009.
- [37] S. Najafi and M. Hajarian. Derivative-Free Optimization on Riemannian Manifolds Using Simplex Gradient Approximations. *Journal of Optimization Theory and Applications*, 208(1):1–22, 2025.
- [38] H. Oviedo and R. Herrera. A collection of efficient retractions for the symplectic Stiefel manifold. *Computational and Applied Mathematics*, 42(4), 2023.
- [39] H. Sato. *Riemannian Optimization and Its Applications*. Springer, 2021.
- [40] H. Sato and K. Aihara. Cholesky QR-based retraction on the generalized Stiefel manifold. *Computational Optimization and Applications*, 72(2):293–308, 2018.
- [41] S. E. Selvan, P. B. Borckmans, A. Chattopadhyay, and P.-A. Absil. Spherical Mesh Adaptive Direct Search for Separating Quasi-Uncorrelated Sources by Range-Based Independent Component Analysis. *Neural Computation*, 25(9):2486–2522, 2013.
- [42] G. Ughi, V. Abrol, and J. Tanner. An empirical study of derivative-free-optimization algorithms for targeted black-box attacks in deep neural networks. *Optimization and Engineering*, 23(3):1319–1346, 2021.
- [43] S. Vary, P. Ablin, B. Gao, and P.-A. Absil. Optimization without Retraction on the Random Generalized Stiefel Manifold. In *Proceedings of the 41st International Conference on Machine Learning (ICML 2024)*, 2024.