

Open Problem: Polynomial linearly-convergent method for geodesically convex optimization?

Christopher Criscitiello

Ecole Polytechnique Fédérale de Lausanne (EPFL), Institute of Mathematics

CHRISTOPHER.CRISCITIELLO@EPFL.CH

David Martínez-Rubio

Zuse Institute Berlin and Technische Universität

MARTINEZ-RUBIO@ZIB.DE

Nicolas Boumal

Ecole Polytechnique Fédérale de Lausanne (EPFL), Institute of Mathematics

NICOLAS.BOUMAL@EPFL.CH

Editors: Gergely Neu and Lorenzo Rosasco

Abstract

Let $f: \mathcal{M} \rightarrow \mathbb{R}$ be a Lipschitz and geodesically convex function defined on a d -dimensional Riemannian manifold \mathcal{M} . Does there exist a first-order deterministic algorithm which (a) uses at most $O(\text{poly}(d) \log(\epsilon^{-1}))$ subgradient queries to find a point with target accuracy ϵ , and (b) requires only $O(\text{poly}(d))$ arithmetic operations per query? In convex optimization, the classical ellipsoid method achieves this. After detailing related work, we provide an ellipsoid-like algorithm with query complexity $O(d^2 \log^2(\epsilon^{-1}))$ and per-query complexity $O(d^2)$ for the limited case where \mathcal{M} has constant curvature (hemisphere or hyperbolic space). We then detail possible approaches and corresponding obstacles for designing an ellipsoid-like method for general Riemannian manifolds.

Keywords: ellipsoid method; geodesic convexity; Riemannian optimization; hyperbolic space

1 Introduction and background

The ellipsoid method (developed by [Shor \(1977\)](#) and [Yudin and Nemirovskii \(1976a,b\)](#)) is a fundamental algorithm in convex optimization and computational complexity.¹ The first proof that linear programs can be solved in polynomial time used the ellipsoid method ([Khachiyan, 1979](#)).

Let us briefly review how the method works. Suppose we seek to minimize a Lipschitz convex function $f: \mathbb{R}^d \rightarrow \mathbb{R}$ constrained to a closed ball $\bar{B}(x_{\text{ref}}, r)$ of radius r and center x_{ref} . At each step of the ellipsoid method, there is an ellipsoid $E_k \subset \mathbb{R}^d$ which by construction contains the minimizers of $\min_{x \in \bar{B}(x_{\text{ref}}, r)} f(x)$. The ellipsoid method queries the center of the ellipsoid and receives a subgradient, which in turn determines a halfspace containing the minimizers. The next ellipsoid E_{k+1} is the minimum-volume ellipsoid containing the intersection of the halfspace and E_k . The volume of the ellipsoids E_k decreases at a linear rate, leading to the linear rate of convergence for the method. The ellipsoid method has two key properties. First, given a target accuracy ϵ , it finds an ϵ -approximate solution in at most $O(d^2 \log(\epsilon^{-1}))$ subgradient queries. Second, each iteration requires only $O(d^2)$ arithmetic operations to determine the next query.

We ask whether there is an algorithm for geodesically convex (g-convex) optimization on a Riemannian manifold which has linear convergence and each query can be computed efficiently. This problem has been stated informally before by [Allen-Zhu et al. \(2018, Sec. 2.2\)](#) and [Rusciano \(2019, Sec. 1.1,1.3\)](#). Our contributions are to point out a partial solution for the case of constant curvature, to suggest possible avenues of attack, and to bring the problem to a wider audience.

1. For a description and analysis of the ellipsoid method, see ([Nemirovski, 1994, Sec. 3.2](#)) or ([Nesterov, 2004, Sec. 3.2](#)).

Let us recall the relevant definitions. For references on g -convex optimization, see (Udriște, 1994) or (Boumal, 2023, Ch. 11). Throughout, \mathcal{M} denotes a complete d -dimensional Riemannian manifold with Riemannian metric $\langle \cdot, \cdot \rangle$, distance dist , and exponential and logarithm maps \exp and \log . A subset D of \mathcal{M} is g -convex if for all $x, y \in D$ there is a unique minimizing geodesic segment γ in \mathcal{M} connecting x, y , and γ is contained in D (this is sometimes called strongly g -convex (Boumal, 2023, Sec. 11.3)). A function $f: D \rightarrow \mathbb{R}$ is g -convex in D if $f \circ \gamma$ is convex for all geodesic segments γ contained in D . A function f is M -Lipschitz in D if $|f(x) - f(y)| \leq M \text{dist}(x, y)$ for all $x, y \in D$. A tangent vector g in the tangent space at x , denoted $\mathbb{T}_x \mathcal{M}$, is a *subgradient* of $f: D \rightarrow \mathbb{R}$ at x if $f(y) \geq f(x) + \langle g, \log_x(y) \rangle$ for all $y \in D$. A *halfspace* on \mathcal{M} is a set $\{y \in \mathcal{M} : \langle g, \log_x(y) \rangle \leq 0\}$ for some (x, g) in the tangent bundle $\mathbb{T}\mathcal{M}$.

We adopt the black-box model of optimization (Nesterov, 2004, Sec. 1.1). A first-order algorithm can access the function f to be minimized through oracle queries $x_k \in \mathcal{M}$. After query x_k , the algorithm is given $f(x_k)$ and a subgradient g at x_k . We can now state the open question:

Open Question 1 *Assume \mathcal{M} has sectional curvatures in the interval $[-K, K]$. Assume $\bar{B}(x_{\text{ref}}, r)$ is g -convex. Let $f: \bar{B}(x_{\text{ref}}, r) \rightarrow \mathbb{R}$ be g -convex and M -Lipschitz in $\bar{B}(x_{\text{ref}}, r)$. Define $\zeta_{r\sqrt{K}} = \frac{r\sqrt{K}}{\tanh(r\sqrt{K})}$. Is there a deterministic first-order algorithm \mathcal{A} with the following properties?*

- (a) *For every $\epsilon \in (0, 1)$, algorithm \mathcal{A} finds a point x such that $f(x) - f^* \leq \epsilon \cdot Mr$ in at most $O(\text{poly}(\zeta_{r\sqrt{K}}, d) \log(\epsilon^{-1}))$ oracle queries, where $f^* = \min_{x \in \bar{B}(x_{\text{ref}}, r)} f(x)$.*
- (b) *Each iteration requires only a polynomial number of arithmetic operations, $\text{poly}(\zeta_{r\sqrt{K}}, d)$, to determine the next query (in addition to the subgradient queries).*

Unlike in the Euclidean case, we permit the complexity of the method to depend on the curvature through $\zeta_{r\sqrt{K}} = \Theta(1 + r\sqrt{K})$ as such dependence is unavoidable (Criscitiello and Boumal, 2023).

A method satisfying properties (a) and (b) in Open Question 1 is of interest for several reasons. First, the ellipsoid method is of fundamental theoretical importance in convex optimization. As geodesic convexity is a generalization of convexity, it is natural to ask for such a generalization. Second, a method solving Open Question 1 applies to nonsmooth g -convex optimization problems, for example computing the geometric median such as for computational anatomy (Fletcher et al., 2009) or phylogenetics (Bacák, 2014, Ch. 8). Third, there are applications where the cost function is g -convex, but not strongly g -convex, and one seeks a linear convergence rate. A notable example is *operator scaling* (Bürgisser et al., 2019), where it is also important that the method is deterministic.

The operator scaling problem encompasses several statistical problems including robust covariance estimation (Wiesel and Zhang, 2014; Sra and Hosseini, 2015; Franks and Moitra, 2020) and estimation for matrix normal models (Tang and Allen, 2021; Franks et al., 2021). It also encompasses several questions in theoretical computer science, including a variant on polynomial identity testing (see references in (Bürgisser et al., 2019)). For operator scaling, Allen-Zhu et al. (2018) propose a linearly-convergent box-constrained Newton method. This method does not solve Open Question 1 as it assumes additional properties about the objective (e.g., second-order robustness).

There are other methods for convex optimization which have similar properties as the ellipsoid method (Nesterov, 2004, pg. 156). We focus on generalizing the ellipsoid method because it arguably has the simplest analysis among such methods, and similar obstacles appear in generalizing other methods. In Section 2, using the tools of geodesic maps introduced by Martínez-Rubio (2023), we describe an ellipsoid-like method for spaces of constant curvature with query com-

plexity $O(\zeta_{r\sqrt{K}} d^2 \log^2(\epsilon^{-1}))$ and per-query complexity $O(d^2)$. In Section 3, we discuss possible approaches and obstacles to solving Open Question 1 for general Riemannian manifolds.

Related work There are algorithms which satisfy properties (a) and (b) of Open Question 1 individually, but not concurrently. Each iteration of the subgradient method (Zhang and Sra, 2016) can be computed efficiently. However, without additional assumptions, this method does not have linear convergence. The center of gravity method due to Rusciano (2019) satisfies property (a). However, each iteration requires computing the center of gravity, which we do not know how to compute efficiently. Lai and Yoshise (2022) and Hirai et al. (2023) give Riemannian interior-point methods, which can converge quickly, but only under additional assumptions on the cost function.

Intermediate milestones Open Question 1 asks for a deterministic algorithm, but one may initially permit the use of random queries. For example, randomness can be helpful for computing the center of gravity in \mathbb{R}^d (Nemirovski, 1994, Sec. 3.3). One may also focus on the Hadamard manifold of $n \times n$ positive definite matrices endowed with the affine-invariant metric (Boumal, 2023, Sec. 11.7). This is the most common manifold occurring in applications of g-convexity and possesses a well-studied structure (Bridson and Haefliger, 1999; Dolcetti and Pertici, 2018).

2 An ellipsoid-like method for spaces of constant curvature

Consider the setting described in Open Question 1. We first observe that it is essentially enough to solve Open Question 1 with radius $R = 1/\sqrt{K}$. Let $D = \bar{B}(x_{\text{ref}}, r)$ and assume $r > R$. (The case $r \leq R$ is easier to handle: we only need to pull f back by a geodesic map once.) Consider the following algorithm. Initialize $x_0 = x_{\text{ref}}$. Given x_k , approximately solve the subproblem $\min_{\bar{B}(x_k, R) \cap D} f$ for a point x_{k+1} such that $f(x_{k+1}) - \min_{\bar{B}(x_k, R) \cap D} f \leq \frac{\epsilon}{4} MR$. Repeat the process. Using only g-convexity of f and D , one can show that after repeating this procedure $T = \lceil \frac{2r}{R} \log(\frac{2}{\epsilon}) \rceil$ times, we have $f(x_T) - f^* \leq \epsilon$ (see the proof of Theorem 7 in (Martínez-Rubio, 2023, App. A) for details). Therefore, if we can solve each subproblem with $O(d^2 \log(\epsilon^{-1}))$ subgradient queries, then we can solve the original problem in $O(\zeta_{r\sqrt{K}} d^2 \log^2(\epsilon^{-1}))$ queries.

Now further assume \mathcal{M} is a d -dimensional hyperbolic space \mathbb{H}^d (curvature equals -1 without loss of generality), and let us show how to solve each subproblem $\min_{\bar{B}(x_k, R) \cap D} f$. The analysis for a sphere is similar, so we omit it. The key tool we need is geodesic maps. A geodesic map with base point $x \in \mathcal{M}$ is a diffeomorphism from \mathbb{H}^d to the open unit ball $B(0, 1)$ of \mathbb{R}^d which maps x to the origin and which maps geodesics of \mathbb{H}^d to straight lines of \mathbb{R}^d intersected with $B(0, 1)$. Geodesic maps of \mathbb{H}^d have an explicit formula given by the Beltrami-Klein model, see (Martínez-Rubio, 2023, App. C). Our strategy to approximately solve $\min_{\bar{B}(x_k, R) \cap D} f$ is simple: we first pull f back to a Euclidean space via a geodesic map $\phi_k: \mathbb{H}^d \rightarrow B(0, 1)$ based at x_k , and then solve the resulting Euclidean problem $\min_{\bar{B}(0, \tilde{R}) \cap \phi_k^{-1}(D)} f \circ \phi_k^{-1}$ using the usual ellipsoid method. We initialize that method with the ball $\bar{B}(0, \tilde{R})$. Here \tilde{R} equals R times an absolute constant, i.e., $\tilde{R} = \Theta(1/\sqrt{K})$.

Why does our strategy for solving the subproblem work? First, geodesic maps have the key property that they map halfspaces in \mathbb{H}^d to halfspaces of \mathbb{R}^d intersected with $B(0, 1)$ (this fact can be verified by inspecting the explicit formula of a geodesic map). In particular, given $y_k \in \bar{B}(0, \tilde{R}) \cap \phi_k^{-1}(D)$ (e.g., a query from the ellipsoid method), we can compute a Euclidean halfspace containing the sublevel set $\{y \in \phi_k^{-1}(D) \cap B(0, 1) : f(\phi_k^{-1}(y)) \leq f(\phi_k^{-1}(y_k))\}$ by pulling back the corresponding Riemannian halfspace $\{x \in D : f(x) \leq f(\phi_k^{-1}(y_k))\}$ with ϕ_k .² Likewise if

2. This is alternatively quantified with the notion of *tilted convexity* introduced by Martínez-Rubio (2023, Lem. 3).

$y_k \notin \bar{B}(0, \tilde{R}) \cap \phi_k^{-1}(D)$, one can compute a Euclidean separating halfspace for $\bar{B}(0, \tilde{R}) \cap \phi_k^{-1}(D)$ by pulling back a Riemannian separating halfspace for $\bar{B}(x_k, R) \cap D$. Second, since we have chosen R to be small, the metric distortion due to the geodesic map is small: in particular, $f \circ \phi_k^{-1}$ is still $O(M)$ -Lipschitz (Martínez-Rubio, 2023, Lem. 2.c). Lastly, note that the convexity of f is not needed for the classical ellipsoid method to work: we just need that f is Lipschitz and at each query point we are provided with a halfspace containing the minimizers of the problem (as guaranteed by our first observation). One can verify that the Euclidean volume of $\bar{B}(0, \tilde{R}) \cap \phi_k^{-1}(D)$ is at least the volume of $B(0, c\tilde{R})$ for some absolute constant $c \in (0, 1)$. This allows us to conclude the ellipsoid method finds an $\frac{\epsilon}{4}MR$ -approximate solution to each subproblem in $O(d^2 \log(\epsilon^{-1}))$ queries.

Can we extend this technique beyond hyperbolic spaces? Unfortunately, Beltrami’s theorem states that the only Riemannian manifolds which admit geodesic maps to Euclidean space are those of constant curvature (Matveev, 2006). One could replace a geodesic map with the exponential map, i.e., solve subproblems of the form $\min_{T_x \mathcal{M}} f \circ \exp_x$, and use comparison theorems (Cheeger and Ebin, 2008) to analyze this method; but it is unclear how to make this approach work.

3 Possible approaches and obstacles

There is no convenient notion of ellipsoid on manifolds. We consider two ways to handle this issue.

Ellipsoids in tangent spaces One could maintain ellipsoids in the tangent spaces of \mathcal{M} . Consider the following algorithm. Let $x_0 = x_{\text{ref}}$ and $x^* \in \arg \min_{x \in \bar{B}(x_{\text{ref}}, r)} f(x)$. We initially know $\log_{x_0}(x^*)$ is contained in the ellipsoid $E_0 := \{v \in T_{x_0} \mathcal{M} : \|v\| \leq r\}$. Querying at x_0 to get a subgradient g_0 , we know $\log_{x_0}(x^*)$ is contained in the intersection $E_0 \cap \{v \in T_{x_0} \mathcal{M} : \langle v, g_0 \rangle \leq 0\}$. Take an ellipsoid $\tilde{E}_0 \subset T_{x_0} \mathcal{M}$ with center \tilde{c}_0 which contains this intersection and has sufficiently small volume. Define $x_1 = \exp_{x_0}(\tilde{c}_1)$. To repeat this process, we need a method for *transferring* the ellipsoid \tilde{E}_0 in $T_{x_0} \mathcal{M}$ to an ellipsoid E_1 in $T_{x_1} \mathcal{M}$ whose center is the origin, and such that $\exp_{x_1}(E_1) \supseteq \exp_{x_0}(\tilde{E}_0)$. Kim and Yang (2022, Lem. 1,2) implicitly introduced a way of transferring *balls* between tangent spaces while controlling the distortion of this transfer in order to obtain an accelerated method on manifolds. It is not clear how to generalize their results to ellipsoids.

How should we keep track of the size of these ellipsoids? We could track their Euclidean volumes. However, even if the Euclidean volume of an ellipsoid $E \subset T_x \mathcal{M}$ is small, the Riemannian volume of $\exp_x(E)$ is not necessarily small, unless the ellipsoid is *bounded*. A bound on the Riemannian volume is important because it provides a bound on the optimality gap (Rusciano, 2019).

Halfspace- and volume-preserving maps A key ingredient in the analysis of the ellipsoid method is that there is a nontrivial family of *halfspace-preserving* and *volume-preserving* diffeomorphisms from \mathbb{R}^d to \mathbb{R}^d : the invertible affine maps $x \mapsto Ax + b$ with $\det(A) = 1$. The explicit formula for the minimum-volume ellipsoid containing the intersection of an initial ellipsoid and a halfspace is derived by using a volume-preserving affine map to transform the ellipsoid to a ball (and the halfspace to another halfspace), and then computing the minimum-volume ellipsoid containing the intersection of a halfspace and a ball (a much easier calculation). Let \mathcal{G} be the group of halfspace-preserving and volume-preserving diffeomorphisms from \mathcal{M} to \mathcal{M} . Define the set of “ellipsoids” \mathcal{E} on \mathcal{M} as the result of applying all elements of \mathcal{G} to every geodesic ball. If \mathcal{M} is a Euclidean space, \mathcal{E} contains all ellipsoids in the usual sense. Unfortunately, even for highly symmetric spaces like hyperbolic space, it seems \mathcal{G} is the set of isometries of \mathcal{M} (McKay, 2019; Kajelad, 2020). In this case, \mathcal{E} is the set of geodesic balls, which is not a sufficiently flexible family of subsets to be useful.

References

- Z. Allen-Zhu, A. Garg, Y. Li, R. Oliveira, and A. Wigderson. Operator scaling via geodesically convex optimization, invariant theory and polynomial identity testing. In *Proceedings of the 50th Annual ACM SIGACT Symposium on the Theory of Computing (STOC 2018)*, 2018. doi: 10.1145/3188745.3188942.
- M. Bacák. *Convex analysis and optimization in Hadamard spaces*, volume 22 of *De Gruyter Series in Nonlinear Analysis and Applications*. Walter de Gruyter GmbH & Co KG, 2014.
- N. Boumal. *An introduction to optimization on smooth manifolds*. Cambridge University Press, 2023. doi: 10.1017/9781009166164. URL <https://www.nicolasboumal.net/book>.
- M. R. Bridson and A. Haefliger. *Metric Spaces of Non-Positive Curvature*. Springer-Verlag Berlin Heidelberg, 1999. doi: 10.1007/978-3-662-12494-9.
- P. Bürgisser, C. Franks, A. Garg, R. Oliveira, M. Walter, and A. Wigderson. Towards a theory of non-commutative optimization: Geodesic 1st and 2nd order methods for moment maps and polytopes. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 845–861, 2019. doi: 10.1109/FOCS.2019.00055.
- J. Cheeger and D. G. Ebin. *Comparison Theorems in Riemannian Geometry*. AMS Chelsea Publishing, 2008. ISBN 978-0-8218-4417-5.
- Christopher Criscitiello and Nicolas Boumal. Curvature and complexity: Better lower bounds for geodesically convex optimization, 2023. URL <https://arxiv.org/abs/2306.02959>.
- A. Dolcetti and D. Pertici. Differential properties of spaces of symmetric real matrices. *arXiv: 1807.01113*, 2018.
- P. T. Fletcher, S. Venkatasubramanian, and S. C. Joshi. The geometric median on Riemannian manifolds with application to robust atlas estimation. *NeuroImage*, 45:S143–S152, 2009.
- C. Franks and A. Moitra. Rigorous guarantees for tyler’s m-estimator via quantum expansion. In *33rd Annual Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pages 1–32. PMLR, 2020. URL <http://proceedings.mlr.press/v125/franks20a/franks20a.pdf>.
- C. Franks, R. Oliveira, A. Ramachandran, and M. Walter. Near optimal sample complexity for matrix and tensor normal models via geodesic convexity. *arXiv: 2110.07583*, 2021.
- Hiroshi Hirai, Harold Nieuwboer, and Michael Walter. Interior-point methods on manifolds: theory and applications, 2023. URL <https://arxiv.org/abs/2303.04771>.
- Kajelad. Mathstackexchange: Is every geodesic-preserving diffeomorphism an isometry? <https://math.stackexchange.com/questions/3896934/is-every-geodesic-preserving-diffeomorphism-an-isometry>, 2020.
- L. G. Khachiyan. A polynomial algorithm in linear programming. *Dokl. Akad. Nauk SSSR*, 244(5): 1093–1096, 1979. URL <https://www.mathnet.ru/dan42319>.

- Jungbin Kim and Insoon Yang. Accelerated gradient methods for geodesically convex optimization: Tractable algorithms and convergence analysis. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 11255–11282. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/kim22k.html>.
- Zhijian Lai and Akiko Yoshise. Riemannian interior point methods for constrained optimization on manifolds, 2022. URL <https://arxiv.org/abs/2203.09762>.
- David Martínez-Rubio. Global riemannian acceleration in hyperbolic and spherical spaces, 2023. URL <https://arxiv.org/abs/2012.03618>.
- Vladimir S. Matveev. Geometric explanation of the beltrami theorem. *International Journal of Geometric Methods in Modern Physics*, 03(03):623–629, 05 2006. ISSN 0219-8878. doi: 10.1142/s0219887806001296. URL <https://cir.nii.ac.jp/crid/1361418521514309888>.
- Ben McKay. Mathoverflow: Geodesic preserving diffeomorphisms of constant curvature spaces. <https://mathoverflow.net/questions/338075/geodesic-preserving-diffeomorphisms-of-constant-curvature-spaces>, 2019.
- A. Nemirovski. *Information-based complexity of convex programming*. 1994. URL https://www2.isye.gatech.edu/~nemirovs/Lec_EMCO.pdf.
- Y. Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87 of *Applied optimization*. Springer, 2004. ISBN 978-1-4020-7553-7.
- A. Rusciano. A riemannian corollary of helly’s theorem. *arXiv: 1804.10738*, 2019.
- Naum Z Shor. Cut-off method with space extension in convex programming problems. *Cybernetics*, 13(1):94–96, 1977.
- S. Sra and R. Hosseini. Conic geometric optimization on the manifold of positive definite matrices. *SIAM Journal on Optimization*, 25(1):713–739, 2015. doi: 10.1137/140978168.
- T. M. Tang and G. I. Allen. Integrated principal components analysis. *arXiv: 1810.00832*, 2021.
- C. Udriște. *Convex functions and optimization methods on Riemannian manifolds*, volume 297 of *Mathematics and its applications*. Kluwer Academic Publishers, 1994. doi: 10.1007/978-94-015-8390-9.
- A. Wiesel and T. Zhang. Structured robust covariance estimation. *Foundations and Trends in Signal Processing*, 8(3):127–216, 2014. doi: 10.1561/20000000053.
- David B Yudin and Arkadi S Nemirovskii. Evaluation of the informational complexity of mathematical programming problems. *Matekon*, 13(2):3–25, 1976a.
- David B Yudin and Arkadi S Nemirovskii. Informational complexity and efficient methods for the solution of convex extremal problems. *Matekon*, 13(3):25–45, 1976b.

H. Zhang and S. Sra. First-order methods for geodesically convex optimization. In *Conference on Learning Theory*, pages 1617–1638, 2016.