

# Geometry-Aware Sampling-Based Motion Planning on Riemannian Manifolds

Phone Thiha Kyaw<sup>✉</sup> and Jonathan Kelly<sup>✉</sup>

Institute for Aerospace Studies, University of Toronto, Toronto, Canada  
{phone.thiha,jonathan.kelly}@robotics.utias.utoronto.ca

**Abstract.** In many robot motion planning problems, task objectives and physical constraints induce non-Euclidean geometry on the configuration space, yet many planners operate using Euclidean distances that ignore this structure. We address the problem of planning collision-free motions that minimize length under configuration-dependent Riemannian metrics, corresponding to geodesics on the configuration manifold. Conventional numerical methods for computing such paths do not scale well to high-dimensional systems, while sampling-based planners trade scalability for geometric fidelity. To bridge this gap, we propose a sampling-based motion planning framework that operates directly on Riemannian manifolds. We introduce a computationally efficient midpoint-based approximation of the Riemannian geodesic distance and prove that it matches the true Riemannian distance with third-order accuracy. Building on this approximation, we design a local planner that traces the manifold using first-order retractions guided by Riemannian natural gradients. Experiments on a two-link planar arm and a 7-DoF Franka manipulator under a kinetic-energy metric, as well as on rigid-body planning in  $SE(2)$  with non-holonomic motion constraints, demonstrate that our approach consistently produces lower-cost trajectories than Euclidean-based planners and classical numerical geodesic-solver baselines.

**Keywords:** Motion and Path Planning · Sampling-Based Algorithms · Riemannian Geometry · Geodesics · Nonholonomic Planning

## 1 Introduction

Robotic motion planning is often posed as a search for collision-free paths through a configuration space. For many robotic systems, the configuration space is a non-Euclidean manifold. For example, rigid-body poses live on Lie groups such as  $SE(2)$  and  $SE(3)$ , while articulated manipulators live on products of circles (tori). More generally, closed-chain or task constraints induce implicit manifolds embedded in a higher-dimensional ambient space [32]. In these settings, planning feasibility and optimality have clear geometric interpretations: feasibility is governed by the intrinsic manifold structure and constraints, while optimality depends on how we measure the cost of motion along the manifold.

A natural way to encode such costs is via a Riemannian metric on the appropriate configuration manifold. This metric induces a notion of distance




Fig. 1: Geodesic motion planning for a 7-DoF Franka manipulator in a cluttered environment. Conventional sampling-based methods compute joint-space shortest paths under the ambient Euclidean metric (red), often ignoring the intrinsic geometry of the configuration manifold (left). Our geodesic formulation instead recovers minimum-energy trajectories under the kinetic-energy Riemannian metric (green), with path thickness indicating relative energy consumption (right).

on the manifold, and the resulting shortest paths, called geodesics, generalize straight lines in Euclidean spaces to curved manifolds [35]. This abstraction recovers classical shortest-path planning in Euclidean configuration spaces as a special case, while also capturing costs that vary smoothly with configuration. Examples of configuration-dependent metrics include the kinetic-energy metric commonly used in manipulation [7, 21, 36], as well as group-invariant metrics on Lie groups [18, 40]. In these settings, optimal motion planning becomes the problem of finding collision-free paths that minimize Riemannian arc length under either a constant or smoothly varying metric.

We study the problem of computing minimum-length geodesics on Riemannian manifolds arising in robotic motion planning. Finding such paths in high dimensions is computationally expensive. Classical approaches either directly solve the geodesic ordinary differential equations as a boundary value problem or minimize a variational energy functional [42]. While effective in low dimensions, these methods scale poorly to high-dimensional robotic systems, struggling to satisfy feasibility and optimality while respecting constraints like joint limits and obstacles in real-time. Sampling-based motion planning offers a scalable alternative. Algorithms such as the Rapidly-exploring Random Tree (RRT) and its anytime variants have been widely adopted in robotic motion planning and shown to scale well in high dimensions [23, 32, 33]. Most existing sampling-based approaches measure distances and interpolate motions using an ambient Euclidean metric. Consequently, they may ignore the intrinsic geometry of the configuration manifold, leading to motions that violate manifold constraints or are infeasible. Constrained motion planning techniques combine numerical continuation with sampling-based planning to address these feasibility problems by searching directly on the manifold [19, 20, 25]. However, optimality is usually expressed through a fixed Euclidean metric, leaving a gap between manifold-aware feasibility and metric-aware optimality.

In this work, we develop geometry-aware subroutines that allow anytime sampling-based planners to optimize Riemannian path length directly on configuration manifolds. This generalizes prior approaches by supporting both constant

and smoothly varying Riemannian metrics (Figure 1). Our contributions are summarized as follows.

- We propose a midpoint-based approximation of the Riemannian geodesic distance and prove that its approximation error vanishes asymptotically with third-order accuracy.
- We design a geometry-aware local planner that traces the manifold using retraction steps along the Riemannian natural gradient under a configuration-dependent metric.
- We empirically validate our approach on energy-minimizing motion planning problems for serial manipulators and on non-holonomic planning tasks on  $SE(2)$ , demonstrating consistently lower-cost paths compared to baselines.

## 2 Related Work

Robot motion generation algorithms often optimize smoothness criteria, such as minimum jerk or acceleration, inspired by human arm movements (see [55] for a review). While effective in certain contexts, these models typically neglect the nonlinear coupling and dynamics that are intrinsic to articulated systems like humans and (most) robots. As a result, linear methods based on Euclidean geometry often fail to capture the true system behaviour, leading to trajectories that are dynamically inconsistent or infeasible. In contrast, geometric methods provide an alternative formulation that more faithfully reflects the underlying nonlinear dynamics, for example by leveraging tools from differential geometry to analyze the mechanics of motor control [15–17]. Building on this perspective, Bullo and Lewis [7] model the configuration space of a multi-linked system as a Riemannian manifold whose metric encodes the system’s structural and dynamic properties. Natural motions then follow geodesics that minimize intrinsic costs such as muscular effort [5, 49] or variations in joint torque [4]. Riemannian metrics have also been proposed to optimize additional criteria, including manipulability [22], joint stiffness [48], and distance to kinematic singularities [39]. Motivated by these geometric insights, our approach frames the robot motion planning problem as a search for collision-free, minimum-length geodesics on Riemannian manifolds defined over high-dimensional configuration spaces.

To compute such minimum-length geodesics, one classical approach is to minimize an energy functional between two joint configurations, resulting in a boundary value problem defined by the Euler–Lagrange equations [42, 51]. Although this yields accurate solutions, it is computationally expensive and often impractical in high-dimensional configuration spaces. As a result, several approximate methods have been proposed to estimate geodesic distances on manifolds. One representative approach solves the Eikonal equation using the Fast Marching Method, which estimates geodesic distances by propagating wavefronts over the manifold geometry [41, 50]. Other techniques include the heat method [12] and shortest-path computations on discretized meshes [54]. Though less expensive than direct geodesic solvers, these methods rely on explicit manifold discretization or mesh-based representations, limiting their applicability to general high-dimensional robotic configuration spaces.

In contrast to techniques that rely on explicit discretization of the configuration space, sampling-based planners such as RRT [33] and RRT\* [23] scale well to high-dimensional spaces. However, standard implementations typically rely on Euclidean metrics for distance computations, often ignoring the underlying geometry of the configuration space. While constrained motion planning methods ensure feasibility by employing projection or retraction operators to maintain samples on implicitly defined manifolds [19, 20, 24–26, 52], they generally still use the ambient metric from the configuration space (e.g., Euclidean distance) to measure the distances between samples. Some methods approximate geodesic paths through sampling [47], while others use Riemannian metrics to guide sampling [56]. More recent approaches, such as RRT\*-R, incorporate Riemannian metrics into planning on low-dimensional submanifolds [57]. Our method generalizes these approaches by integrating geometry-aware subroutines into sampling-based planners, ensuring that resulting paths are not only feasible with respect to manifold constraints but also asymptotically optimal with respect to the intrinsic Riemannian metric.

Several methods exploit geometric structure by designing or reshaping Riemannian metrics to encode constraints or task-relevant features. Under this formulation, the cost of moving between configurations is expressed through a Riemannian metric on joint space, capturing the geometry of the configuration manifold [10, 30, 31]. Metrics can also blend information from both joint and task spaces. This is especially useful when dealing with task-space constraints such as obstacle avoidance. Rather than treating these as cost terms or hard constraints, one can define metrics in task space that reflect obstacle geometry and pull them back to configuration space, yielding geodesics that naturally curve around obstacles [1, 37, 38, 44]. This geometric reformulation has been shown to improve convergence in motion optimization methods such as CHOMP [45].

Beyond reshaping metrics for single objectives, Riemannian geometry also provides a natural mechanism for blending several distinct motion objectives or constraints. In [11, 46], the authors propose combining multiple motion policies, each representing a different robot behaviour, into a single, geometrically consistent policy using Riemannian pullback operations. Similar pullback operations have been used to reshape the configuration space metric based on barrier functions, allowing constraints like joint limits, self-collisions, and obstacles to be naturally encoded [2, 27]. In contrast to these approaches, we do not reshape the metric. Instead, we work directly with the intrinsic Riemannian metric and minimize the curve length based on it, rather than minimizing the energy functional. Constraints such as obstacle avoidance are handled implicitly through the sampling-based planner, without being explicitly encoded in the metric. This approach allows us to leverage all the benefits of sampling-based planning while producing collision-free, minimum-length motions without requiring any handcrafting in the metric design.

### 3 Preliminaries

This section presents a minimal set of tools from differential geometry required to follow the development in this work. For a more in-depth treatment of these concepts, we refer the reader to the following textbooks [6, 34, 35].

### 3.1 Riemannian Metrics and Manifolds

Let  $\mathcal{M}$  be an  $n$ -dimensional manifold embedded in an ambient linear space  $\mathcal{E}$  (e.g.,  $\mathcal{M} \subseteq \mathbb{R}^d$  with  $d \geq n$ ). By definition,  $\mathcal{M}$  is a topological space that is locally Euclidean, meaning that each point  $q \in \mathcal{M}$  has a neighborhood homeomorphic to an open subset of  $\mathbb{R}^n$ . The tangent space  $\mathcal{T}_q\mathcal{M}$  at a point  $q \in \mathcal{M}$  consists of the velocity vectors of all smooth curves on  $\mathcal{M}$  passing through  $q$ . When  $\mathcal{M}$  is embedded in  $\mathcal{E}$ , the tangent space  $\mathcal{T}_q\mathcal{M}$  may be identified with a linear subspace of  $\mathcal{E}$ . The collection of all tangent spaces forms the tangent bundle, defined as the disjoint union

$$\mathcal{TM} = \{(q, v) : q \in \mathcal{M} \text{ and } v \in \mathcal{T}_q\mathcal{M}\}.$$

Since each tangent space is a vector space, we can define an inner product on  $\mathcal{T}_q\mathcal{M}$ , given by a bilinear, symmetric and positive definite map  $\langle \cdot, \cdot \rangle_q : \mathcal{T}_q\mathcal{M} \times \mathcal{T}_q\mathcal{M} \rightarrow \mathbb{R}$ . This inner product induces a norm  $\|v\|_q = \langle v, v \rangle_q^{1/2}$  on tangent vectors. A Riemannian metric  $G_q$  on  $\mathcal{M}$  is a smoothly varying choice of such inner products for each  $q \in \mathcal{M}$  acting on  $\mathcal{TM}$ . In local coordinates, the Riemannian inner product between two tangent vectors  $u, v \in \mathcal{T}_q\mathcal{M}$  can be written as

$$\langle u, v \rangle_q = u^\top G_q v,$$

where  $G_q$  is a symmetric positive definite matrix representing the metric at  $q$ . A manifold  $\mathcal{M}$  equipped with a Riemannian metric is called a Riemannian manifold.

### 3.2 Distances and Geodesics

Given a Riemannian manifold  $\mathcal{M}$ , the metric induces a natural notion of length for smooth curves on  $\mathcal{M}$ , and consequently defines a distance function, making  $\mathcal{M}$  into a metric space. For a piecewise smooth curve  $\pi : [0, 1] \rightarrow \mathcal{M}$ , its length is defined as the integral of its speed under the Riemannian metric,

$$L(\pi) = \int_0^1 \|\dot{\pi}(t)\|_{\pi(t)} dt = \int_0^1 \sqrt{\dot{\pi}(t)^\top G_{\pi(t)} \dot{\pi}(t)} dt.$$

This induces the *Riemannian distance* between two points on the manifold

$$d_{\mathcal{M}}(q_x, q_y) = \inf_{\pi} L(\pi), \quad (1)$$

where the infimum is taken over all piecewise smooth curves  $\pi$  such that  $\pi(0) = q_x$  and  $\pi(1) = q_y$ . Geodesics are curves that locally minimize this distance, generalizing the notion of straight lines in Euclidean space. Equivalently, geodesics between two fixed endpoints minimize the curve energy

$$E(\pi) = \frac{1}{2} \int_0^1 \dot{\pi}(t)^\top G_{\pi(t)} \dot{\pi}(t) dt. \quad (2)$$

They satisfy the geodesic ordinary differential equation (ODE)

$$\ddot{q}^k(t) + \Gamma_{ij}^k(q(t)) \dot{q}^i(t) \dot{q}^j(t) = 0, \quad i, j, k \in \{1, \dots, n\}, \quad (3)$$

where  $q(t) = (q^1(t), \dots, q^n(t))$  denotes the local coordinate representation of  $\pi(t)$ .<sup>1</sup> Here,  $\Gamma_{ij}^k$  are the Christoffel symbols of the second kind, computed directly

<sup>1</sup> We use the Einstein summation convention here.

from the Riemannian metric and given by

$$\Gamma_{ij}^k = \frac{1}{2} G^{kl} \left( \frac{\partial G_{il}}{\partial q^j} + \frac{\partial G_{jl}}{\partial q^i} - \frac{\partial G_{ij}}{\partial q^l} \right), \quad (4)$$

which describe how the coordinate basis varies across the manifold. Computing geodesics from (3) requires solving a boundary value problem, which becomes computationally challenging in high-dimensional spaces. For this reason, rather than solving the geodesic ODEs directly, we consider minimizing the curve length in (1) under the intrinsic Riemannian metric (see Section 3.3).

### 3.3 Motion Planning Problem Formulation

Let  $\mathcal{Q} \subseteq \mathcal{M}$  denote the configuration manifold (for example, a constrained lower-dimensional submanifold), with each  $q \in \mathcal{Q}$  representing a configuration of the system. In this work, we restrict attention to the unconstrained case  $\mathcal{Q} = \mathcal{M}$ , while retaining the notation  $\mathcal{Q}$  for generality. Let  $\mathcal{Q}_{\text{obs}} \subsetneq \mathcal{Q}$  denote the set of configurations in collision with obstacles. The obstacle-free configuration space is defined as  $\mathcal{Q}_{\text{free}} = \mathcal{Q} \setminus \mathcal{Q}_{\text{obs}}$ . Given an initial configuration  $q_{\text{start}}$  and a goal configuration  $q_{\text{goal}}$ , the objective is to find a feasible path  $\pi^*$  in  $\mathcal{Q}_{\text{free}}$  that minimizes the Riemannian length objective given in (1):

$$\pi^* = \underset{\pi \in \Sigma}{\operatorname{argmin}} \left\{ L(\pi) \mid \pi(0) = q_{\text{start}}, \pi(1) = q_{\text{goal}}, \pi(t) \in \mathcal{Q}_{\text{free}}, \forall t \in [0, 1] \right\}, \quad (5)$$

where  $\Sigma$  denotes the set of all piecewise smooth feasible paths.

In this work, we address the planning problem in (5) by adopting sampling-based motion planning methods. Specifically, we search for solutions by incrementally constructing a rapidly-exploring random tree through random sampling, and refining it via incremental rewiring to asymptotically approach shortest paths in  $\mathcal{Q}_{\text{free}}$  [23, 33]. This strategy enables scalability to high-dimensional configuration spaces while respecting the intrinsic geometry of the manifold. Because the Riemannian metric generally distorts the space, shortest paths are no longer straight lines in the ambient space  $\mathcal{E}$  (e.g.,  $\mathbb{R}^d$ ), and it is therefore critical that distance computations and interpolations within the planner remain consistent with the underlying manifold structure. We address geometric consistency in Section 4.

## 4 Geodesic Motion Planning on Riemannian Manifolds

This section develops the geometry-aware subroutines that enable sampling-based motion planning algorithms to find geodesic paths on Riemannian manifolds. We first introduce a computationally efficient approximation of geodesic distance between configurations (Section 4.1). We then describe a gradient-based interpolation procedure that uses this distance to extend the search tree while respecting the underlying geometry of the manifold (Section 4.2).

### 4.1 Distance Between Configurations

The distance between two configurations on a Riemannian manifold is defined as the length of the shortest connecting curve under the intrinsic metric (1). Exact




Fig. 2: Midpoint-based geodesic distance between configurations  $q_x$  and  $q_y$  on the manifold  $\mathcal{M}$ . The geodesic midpoint  $q_{\text{mid}}$  is constructed by interpolating halfway along the tangent vector connecting  $q_x$  and  $q_y$  in either of their tangent spaces and mapping the result back to the manifold (left). The distance is then computed in the tangent space at  $q_{\text{mid}}$  as the Riemannian norm of the difference between corresponding tangent vectors (right).

evaluation of this distance requires solving a geodesic boundary value problem, which is impractical for online use due to the repeated calls required by nearest-neighbour queries and tree rewiring in sampling-based planners. Accordingly, we approximate the geodesic distance by evaluating the Riemannian metric at the midpoint of the two configurations and computing the length of a piecewise path in the midpoint tangent space. This approximation is computationally ‘cheap’ since it requires only a single metric evaluation per subroutine call. We later show that this midpoint-based approximation converges to the true geodesic distance as the configurations become arbitrarily close to each other (Theorem 1).

We begin by analyzing the geometry of the geodesic midpoint and establish that this formulation yields an exact distance identity.

**Definition 1 (Geodesic Midpoint).** *Let  $\mathcal{M}$  be a Riemannian manifold and let  $q_x, q_y \in \mathcal{M}$  denote two configurations (points) on the manifold. Assuming  $q_x$  and  $q_y$  lie within a geodesically convex neighbourhood, the geodesic midpoint  $q_{\text{mid}}$  is constructed by mapping half the geodesic distance from  $q_x$  towards  $q_y$*

$$q_{\text{mid}} = \text{exp}_{q_x} \left( \frac{1}{2} \log_{q_x}(q_y) \right).$$

By [35, Theorem 6.17],  $q_{\text{mid}}$  lies on the unique minimizing geodesic  $\pi : [0, 1] \rightarrow \mathcal{M}$  from  $q_x$  to  $q_y$  such that  $\pi(0) = q_x$ ,  $\pi(1) = q_y$ , and  $\pi(1/2) = q_{\text{mid}}$ .

This construction is particularly useful since it provides a symmetric reference frame (Figure 2). As the following lemma shows, projecting both configurations into the tangent space of the midpoint yields a symmetric distance relation.

**Lemma 1.** *The Riemannian distance between  $q_x$  and  $q_y$  satisfies the identity*

$$d_{\mathcal{M}}(q_x, q_y) = \left\| \log_{q_{\text{mid}}}(q_y) - \log_{q_{\text{mid}}}(q_x) \right\|_{q_{\text{mid}}}, \quad (6)$$

where  $\log_{q_{\text{mid}}}(\cdot)$  denotes the Riemannian logarithmic map on the tangent space  $\mathcal{T}_{q_{\text{mid}}}\mathcal{M}$  and  $\|\cdot\|_{q_{\text{mid}}}$  denotes the norm induced by the metric at the midpoint.

*Proof.* By Definition 1,  $\pi$  is the unique minimizing geodesic connecting  $q_x$  to  $q_y$  within a geodesically convex neighbourhood on  $\mathcal{M}$ . Since  $\pi$  is a geodesic, the curve segment connecting  $q_{\text{mid}}$  to  $q_y$  is also a geodesic, given by the reparameterization

$\pi_y(t) = \pi(\frac{1}{2} + \frac{t}{2})$  for all  $t \in [0, 1]$ . The value of the Riemannian logarithmic map at  $q_{\text{mid}}$  is precisely the initial velocity of this segment

$$\log_{q_{\text{mid}}}(q_y) = \dot{\pi}_y(0) = \frac{1}{2} \dot{\pi}(1/2).$$

Similarly, the segment connecting  $q_{\text{mid}}$  to  $q_x$  is given by  $\pi_x(t) = \pi(\frac{1}{2} - \frac{t}{2})$ , yielding

$$\log_{q_{\text{mid}}}(q_x) = \dot{\pi}_x(0) = -\frac{1}{2} \dot{\pi}(1/2).$$

Subtracting the two tangent vectors yields

$$\log_{q_{\text{mid}}}(q_y) - \log_{q_{\text{mid}}}(q_x) = \frac{1}{2} \dot{\pi}(1/2) - \left(-\frac{1}{2} \dot{\pi}(1/2)\right) = \dot{\pi}(1/2).$$

Taking the norm with respect to the Riemannian metric at  $q_{\text{mid}}$ , we obtain

$$\left\| \log_{q_{\text{mid}}}(q_y) - \log_{q_{\text{mid}}}(q_x) \right\|_{q_{\text{mid}}} = \|\dot{\pi}(1/2)\|_{q_{\text{mid}}}.$$

Since  $\pi$  is a minimizing geodesic, it has constant speed, implying that  $\|\dot{\pi}(t)\|_{\pi(t)}$  is constant for all  $t$ . Consequently,

$$\|\dot{\pi}(1/2)\|_{q_{\text{mid}}} = \int_0^1 \|\dot{\pi}(t)\|_{\pi(t)} dt = d_{\mathcal{M}}(q_x, q_y) \quad \square$$

In practice, computing exponential and logarithmic maps requires solving a geodesic boundary value problem, which is typically too computationally expensive for use in sampling-based motion planners. In this work, we propose replacing these operations with *retractions*, which are first-order approximations of the exponential map, and their local inverses, which serve as computationally less expensive approximations of the logarithmic map.

**Definition 2 (Retraction).** *A retraction on a Riemannian manifold  $\mathcal{M}$  is a smooth map  $\mathcal{R} : \mathcal{T}\mathcal{M} \rightarrow \mathcal{M}; (q, v) \mapsto \mathcal{R}_q(v)$  from the tangent bundle  $\mathcal{T}\mathcal{M}$  onto  $\mathcal{M}$  such that its restriction  $\mathcal{R}_q : \mathcal{T}_q\mathcal{M} \rightarrow \mathcal{M}$  to the tangent space at  $q$  satisfies  $\mathcal{R}_q(0) = q$  and the identity map  $D\mathcal{R}_q(0)[v] = v$ .*

Since the differential  $D\mathcal{R}_q(0)$  is nonsingular, the inverse function theorem guarantees that  $\mathcal{R}_q$  is a local diffeomorphism around the origin. Consequently, it admits a local inverse  $\mathcal{R}_q^{-1}$  defined in a neighbourhood of  $q$ . Leveraging these operators, we formulate an approximation of the midpoint distance in (6) by substituting the exp and log maps with their retraction counterparts,

$$\hat{d}_{\mathcal{M}}(q_x, q_y) = \left\| \mathcal{R}_{\hat{q}_{\text{mid}}}^{-1}(q_y) - \mathcal{R}_{\hat{q}_{\text{mid}}}^{-1}(q_x) \right\|_{\hat{q}_{\text{mid}}}, \quad (7)$$

where the retraction midpoint  $\hat{q}_{\text{mid}}$  can be constructed as

$$\hat{q}_{\text{mid}} = \mathcal{R}_{q_x} \left( \frac{1}{2} \mathcal{R}_{q_x}^{-1}(q_y) \right). \quad (8)$$

The primary advantage of the construction in (7) is its higher-order local accuracy relative to endpoint-based retraction distances, analogous to the improved accuracy of central finite differences over forward or backward differences

in numerical analysis. Although the retraction midpoint only approximates the true geodesic midpoint, it is accurate enough that the first- and second-order distortion terms introduced by first-order retractions cancel in the final difference, effectively yielding an extra order of accuracy for free. Theorem 1 formalizes this argument, establishing that the leading term in the Taylor expansion of the approximation error is third order in the separation between configurations.

**Theorem 1.** *Let  $\mathcal{M}$  be a smooth Riemannian manifold and  $U \subseteq \mathcal{M}$  be a compact subset. For any configurations  $q_x, q_y \in U$ , the midpoint retraction distance approximates the true Riemannian distance with cubic accuracy, that is, the approximation error satisfies*

$$\left| \hat{d}_{\mathcal{M}}(q_x, q_y) - d_{\mathcal{M}}(q_x, q_y) \right| = \mathcal{O}(d_{\mathcal{M}}(q_x, q_y)^3).$$

*Proof.* The result follows from the Taylor expansion of the metric tensor in Riemann normal coordinates centred at the midpoint; see Appendix A.  $\square$

## 4.2 Vertex Expansion

Sampling-based motion planning methods typically require a fast and efficient procedure for connecting two configurations, a process often referred to as local planning. In the case of RRTs, for example, this operation attempts to connect the nearest neighbour to a randomly sampled configuration using an edge of bounded length, thereby generating a new candidate configuration. In Euclidean spaces, such connections are typically implemented using straight-line interpolation. On curved configuration spaces, however, this task requires finding geodesics that locally minimize distance while respecting the intrinsic geometry of the manifold.

At a high level, our method performs this interpolation by constructing a discrete approximation to a geodesic that follows the natural gradient of the squared Riemannian distance potential. Rather than solving for a continuous curve in closed form, we generate a sequence of configurations by iteratively descending along this gradient. Such descent requires moving around the manifold along a specified direction, for which we make use of retractions (Definition 2). A retraction maps a configuration  $q \in \mathcal{M}$  and a tangent vector  $v \in \mathcal{T}_q\mathcal{M}$  to a new point on the manifold, thereby allowing local movement while remaining on  $\mathcal{M}$ .<sup>2</sup> We next formalize the notion of gradients on a Riemannian manifold (Definition 3).

**Definition 3 (Riemannian gradient).** *Let  $\phi : \mathcal{M} \rightarrow \mathbb{R}$  be a smooth function on a Riemannian manifold  $\mathcal{M}$ . The Riemannian gradient of  $\phi$  is the vector field  $\text{grad } \phi$  on  $\mathcal{M}$  uniquely defined by*

$$D\phi(q)[v] = \langle v, \text{grad } \phi(q) \rangle_q \quad \forall v \in \mathcal{T}_q\mathcal{M},$$

where  $\langle \cdot, \cdot \rangle_q$  is the inner product induced by the Riemannian metric  $G$  on  $\mathcal{T}_q\mathcal{M}$  and  $D\phi(q)[v]$  is the differential of  $\phi$  at  $q$  along  $v$ .

<sup>2</sup> For example, on a linear manifold, the mapping  $\mathcal{R}_q(v) = q + v$  is a valid retraction.

**Algorithm 1:** Expansion ( $q_{\text{near}}, q_{\text{rand}} \in \mathcal{M}$ )

---

```

1  $q \leftarrow q_{\text{near}}, d \leftarrow 0;$ 
2 repeat
3    $u \leftarrow \mathcal{R}_q^{-1}(q_{\text{rand}});$ 
4    $v \leftarrow G(q)^{-1} \nabla_u (\phi \circ \mathcal{R}_q)(0);$ 
5    $q_{\text{next}} \leftarrow \mathcal{R}_q(-s \hat{v});$ 
6   if  $\hat{d}_{\mathcal{M}}(q, q_{\text{next}}) > \lambda s$  then
7      $s \leftarrow \frac{1}{2} s;$ 
8     if  $s < s_{\text{min}}$  then
9       break
10    continue
11    $d \leftarrow d + \hat{d}_{\mathcal{M}}(q, q_{\text{next}});$ 
12   if  $d > d_{\text{max}}$  then
13     break
14    $q \leftarrow q_{\text{next}};$ 
15 until  $\hat{d}_{\mathcal{M}}(q, q_{\text{rand}}) \leq s;$ 
16 return  $q;$ 

```

---

In practice, the Riemannian gradient of  $\phi$  can be computed by evaluating a classical gradient at the origin of the linear tangent model provided by the retraction. Specifically, for any retraction  $\mathcal{R}$  on  $\mathcal{M}$ ,

$$\text{grad } \phi(q) = \text{grad}(\phi \circ \mathcal{R}_q)(0),$$

where  $\phi \circ \mathcal{R}_q : \mathcal{T}_q \mathcal{M} \rightarrow \mathbb{R}$  is a smooth function defined on the linear space  $\mathcal{T}_q \mathcal{M}$  endowed with the inner product  $\langle \cdot, \cdot \rangle_q$  [6]. Equivalently, in local coordinates,

$$\text{grad } \phi(q) = G(q)^{-1} \nabla_u (\phi \circ \mathcal{R}_q)(0) \quad (9)$$

where  $u \in \mathbb{R}^n$  are coordinates on  $\mathcal{T}_q \mathcal{M}$  and  $\nabla_u$  denotes the usual Euclidean gradient on  $\mathcal{T}_q \mathcal{M}$  expressed in any basis. Using (9), the *Riemannian natural gradient* direction at iteration  $k$  is  $v_k = \text{grad } \phi(q_k)$ . An update rule for taking a single discrete step along the direction of steepest descent on  $\mathcal{M}$  can then be written as

$$q_{k+1} = \mathcal{R}_{q_k}(-s_k \hat{v}_k), \quad \hat{v}_k = \frac{v_k}{\|v_k\|_{q_k}}$$

with a step length  $s_k > 0$ , which may be chosen by any backtracking test. Here, we normalize  $v_k$  to ensure the step size  $s_k$  approximates the arc length along the retraction curve, and use the distance function in (7) as our squared cost potential. Specifically, for any fixed configuration, say  $q^\dagger$ , we define the function

$$\phi(q) = \frac{1}{2} \hat{d}_{\mathcal{M}}(q, q^\dagger)^2 \quad (10)$$

and compute the gradient of (10) in local coordinates using (9).

The complete vertex expansion procedure is presented in Algorithm 1. The algorithm performs the expansion process by repeatedly taking small retraction steps on the manifold starting from the nearest configuration  $q_{\text{near}}$  until the extended branch is close to the random configuration  $q_{\text{rand}}$ . At each iteration, it maps  $q_{\text{rand}}$  to the current tangent space at  $q$  via the inverse retraction, yielding

a tangent vector  $u \in \mathcal{T}_q\mathcal{M}$  expressed in local coordinates (Algorithm 1, Line 3). It then evaluates the gradient of (10) in the current tangent space and computes the negative natural-gradient direction  $v$  with respect to the Riemannian metric (Algorithm 1, Line 4). Taking a step size  $s$  along the normalized direction  $\hat{v}$  using retraction yields a new configuration  $q_{\text{next}}$  (Algorithm 1, Line 5). If the resulting displacement exceeds a threshold proportional to the step size, specifically a factor  $\lambda s$ , the step length is halved (Algorithm 1, Lines 6–10). We also track the cumulative distance traveled, denoted  $d$ , and enforce an upper bound  $d_{\text{max}}$  to avoid unnecessarily long expansions, which may otherwise arise due to manifold curvature or variations in the Riemannian metric (Algorithm 1, Lines 11–13). If none of the stopping conditions is triggered, the step is accepted and the current configuration is updated accordingly (Algorithm 1, Line 14).

## 5 Experiments

We validate our proposed approach through two use cases to demonstrate the versatility of geometry-aware sampling-based planners for robot motion planning problems. First, we evaluate the planner’s ability to find minimum-energy motions for serial-link manipulators (Section 5.1). Second, we apply the method to rigid-body planning on SE(2) under nonholonomic constraints (Section 5.2). We compare against standard numerical optimization baselines, including a boundary value problem (BVP) solver and a variational energy minimization method, as well as a sampling-based planner using a Euclidean metric. To enforce obstacle avoidance in the optimization-based methods, we reshape the metric using standard exponential barrier functions and tune the corresponding parameters to achieve the best performance possible. To account for the sensitivity of variational solvers to initial conditions and the stochasticity of sampling-based methods, we conduct all experiments over 50 trials with start and goal configurations perturbed by Gaussian noise, running each sampling-based planner 10 times per trial. In all experiments, we use the Open Motion Planning Library (OMPL) [53] implementation of RRT\* as the underlying sampling-based planner.<sup>3</sup> For geodesic computation with the variational solvers, we use the StochMan Python library [13] to represent geodesics as cubic splines with a fixed number of control points and optimize them to minimize the Riemannian energy functional, following [2, 27].

Throughout our analysis, we evaluate path quality using *geodesic length* and *energy*, both measured under the Riemannian metric specific to each experiment. To ensure consistent reporting, we reparameterize all solution paths to unit-speed curves prior to evaluation and then report the geometric length (which is invariant to parameterization) and the Dirichlet energy functional defined in (2), which quantifies the smoothness of the geodesic.

### 5.1 Serial Link Manipulators

This section presents experimental results assessing the effectiveness of our approach for motion planning with configuration-dependent Riemannian metrics

<sup>3</sup> Code is publicly available at <https://github.com/utiasSTARS/geodex>.

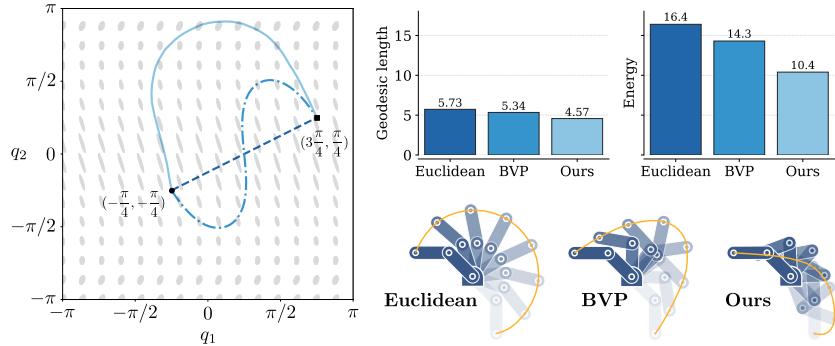


Fig. 3: Geodesics found by various motion planning methods for the 2-DoF planar manipulator experiment in Section 5.1. Configuration space paths from start ( $\bullet$ ) and goal ( $\blacksquare$ ) are shown (left), with shaded regions indicating kinetic-energy ellipsoids at different configurations. Corresponding task space motions with end-effector trajectories are shown in yellow (right). The Euclidean approach minimizes joint space shortest distance, yielding straight-line geodesics ( $- -$ ). Numerical methods optimize the energy functional, producing lower-cost paths ( $- \cdot -$ ) but often converge to local minima. Our geometry-aware sampling-based approach recovers globally optimal geodesics under the intrinsic kinetic-energy Riemannian metric ( $-$ ), achieving lower geodesic length and energy.

for serial manipulators. In all experiments, we define the Riemannian metric using the manipulator’s mass-inertia matrix [21], computed via the composite rigid body algorithm implemented in the Pinocchio library [8]. Under this metric, geodesics correspond to motions that minimize kinetic energy for a given traversal time.

**Two-Link Planar Arm** We first consider a simple 2-DoF planar manipulator in an obstacle-free environment. The arm consists of two identical links, each with a length of 1.0 m and a mass of 1.0 kg<sup>4</sup>. The task is to plan a motion from an initial configuration of  $q_{\text{start}} = [-\pi/4, -\pi/4]^T$  to a goal configuration of  $q_{\text{goal}} = [3\pi/4, 3\pi/4]^T$ . Unlike the Euclidean baseline, which yields a straight-line path in configuration space, our approach correctly recovers the curved Riemannian geodesic (Figure 3). Even with identical link masses, the effective inertia at the base joint is higher than at the elbow joint; our planner naturally captures this property, generating trajectories that minimize the mechanical work. We also observe that classical numerical baselines (both BVP and variational solvers) are highly sensitive to the initial guess; when initialized with a standard straight-line path, they frequently converge to suboptimal solutions. In contrast, our geometry-aware sampling-based planner consistently recovers the globally optimal geodesic without requiring any prior knowledge of the solution geometry, successfully navigating the nonlinear energy landscape where optimization methods struggle. This behavior highlights the strong dependence of classical numerical methods on initialization in highly nonconvex energy landscapes. By contrast, the sampling-based formulation instead leverages geometric structure to explore multiple homotopy classes, enabling reliable convergence to the global geodesic without requiring problem-specific initialization.

<sup>4</sup> Links are modelled as uniform slender rods with centres of mass at the midpoints.

Table 1: Performance statistics over 50 trials for the 7-DoF Franka arm and anisotropic SE(2) planning tasks described in Sections 5.1 and 5.2. We report median values for geodesic length and energy measured under the corresponding Riemannian metrics; success denotes the fraction of trials that produced a collision-free path. While timing information is not shown here, achieving the indicated performance requires at least 2 to 3 minutes for variational methods, whereas our approach is limited to just one minute.

		Length ( $\downarrow$ )	Energy ( $\downarrow$ )	Success ( $\uparrow$ )
<b>Franka</b>	Variational	$2.5 \pm 0.6$	$3.1 \pm 1.5$	<b>96%</b>
	Sampling (Euclidean)	$2.6 \pm 0.5$	$3.5 \pm 1.5$	85%
	Sampling (Ours)	<b><math>2.1 \pm 0.2</math></b>	<b><math>2.3 \pm 0.4</math></b>	90%
<b>SE(2) Doorway</b>	Variational	$24.9 \pm 0.6$	$310.4 \pm 14.0$	86%
	Sampling (Euclidean)	$43.7 \pm 3.7$	$954.4 \pm 174.7$	<b>100%</b>
	Sampling (Ours)	<b><math>23.2 \pm 0.5</math></b>	<b><math>269.1 \pm 11.7</math></b>	<b>100%</b>
<b>SE(2) Corridor</b>	Variational	$\infty$	$\infty$	8%
	Sampling (Euclidean)	$95.6 \pm 9.2$	$4571.2 \pm 873.0$	<b>100%</b>
	Sampling (Ours)	<b><math>43.0 \pm 0.5</math></b>	<b><math>925.9 \pm 22.5</math></b>	<b>100%</b>

**7-DoF Franka** To evaluate scalability in high-dimensional spaces, we apply our planner to a 7-DoF Franka Emika robot operating in a cluttered environment. We utilize the *table pick* environment from the MotionBenchMaker dataset [9], where the robot must navigate from a start configuration to a grasp pose while avoiding collisions with the table and surrounding obstacles (Figure 1). Because of the dimensionality and complexity of this problem, we exclude the BVP solver from the set of baselines due to its poor scaling behavior. Table 1 summarizes the results over 50 trials. From the data, we observe that all methods achieved high success rates in this setting. Since we employ a single-tree RRT\* as the underlying sampling-based planner, some trials fail to reach the goal due to sampling stochasticity; this limitation could be mitigated by adopting a bidirectional search strategy [28], for example. While the Euclidean approach reliably finds feasible paths, it does not account for the robot’s configuration-dependent inertia, often producing high-energy motions that unnecessarily excite the heavy base joints. Conversely, although the variational solver explicitly minimizes energy, the reshaping of the metric induces a complex, non-convex landscape, making the method sensitive to local minima and requiring careful tuning of barrier parameters. In contrast, our geometry-aware planner consistently produces lower kinetic-energy paths than both baselines, effectively handling obstacle avoidance constraints implicitly without the need for explicit metric design.

## 5.2 Planning on SE(2)

We next apply our framework to rigid-body planning on the special Euclidean group SE(2) to demonstrate how metric design can enforce specific kinematic behaviors without explicitly encoding them as nonholonomic constraints, as is common in kinodynamic motion planners [14, 29]. We evaluate our approach on a large-scale navigation environment from the literature, the Willow Garage map [43], con-




Fig. 4: Comparison of collision-free geodesics produced by the benchmarked methods in the Willow Garage environment for the SE(2) planning experiment described in Section 5.2. The proposed method (—) is compared against the variational method (--) and a sampling-based planner using a Euclidean metric (····). Arrows along each path represent the SE(2) pose at discrete intervals; for visual clarity, only the body-frame  $x$ -axis is shown to illustrate orientation along the trajectory. Start and goal configurations are indicated by green and red dots, respectively, with their full orientation frames.

sidering two distinct scenarios, *Doorway* and *Corridor* (Figure 4). To examine the effect of metric shaping, we employ left-invariant Riemannian metrics defined by a diagonal weight matrix in the body frame,  $G = \text{diag}(w_x, w_y, w_\theta)$ , which assigns independent costs to longitudinal translation, lateral translation, and rotation [3]. In this experiment, we penalize lateral translation by making the metric highly anisotropic (i.e.,  $w_y \gg w_x$ ), effectively creating a ‘soft’ nonholonomic constraint.

As reported in Table 1, both the Euclidean baseline and our proposed method achieve a 100% success rate in finding collision-free paths within the available planning time. The variational method, however, often struggles with the complex geometry of the environment, particularly in the *Corridor* scenario, where it attains only an 8% success rate. This behavior reflects the sensitivity of variational solvers to initialization and the difficulty of tuning barrier-function weights to navigate narrow passages without becoming trapped in local minima or violating collision constraints. Although the Euclidean planner reliably finds feasible paths, the quality of its solutions is poor, incurring substantially higher geodesic length and energy costs compared to our method. As shown in Figure 4, the geodesics found by the Euclidean planner ignore the anisotropic nature of the underlying metric, resulting in unnatural skidding or screw motions where the rigid body translates laterally. Conversely, our geometry-aware planner naturally recovers the soft nonholonomic behavior implied by the intrinsic Riemannian metric, aligning the orientation with the direction of travel to minimize energy and producing significantly shorter geodesics than the baselines.

## 6 Conclusion

This work presents a geometry-aware sampling-based planning framework for robot configuration spaces equipped with a Riemannian metric. Our main contribution is a midpoint-based distance approximation on Riemannian manifolds that

can be evaluated using only retractions and local metric information, yet matches the true Riemannian distance with third-order accuracy. This approximation makes the distance-based subroutines of sampling-based planners more consistent with the intrinsic motion costs than those of Euclidean baselines, while avoiding the expense of solving geodesic boundary-value problems. We further show how the same ingredients enable a geometry-aware local interpolation approach based on discrete retraction steps and Riemannian natural gradients. Across manipulation and SE(2) problems under anisotropic metrics, our method consistently produces higher-quality solutions under the target metric, especially in settings where Euclidean or isotropic assumptions are misleading. Future work will explore the full implications of our approach. In particular, we plan to investigate the design of heuristic functions in curved spaces to focus search on promising regions of the configuration space. We are also interested in studying tighter theoretical guarantees when our proposed geometry-aware subroutines are used within asymptotically optimal planners.

## Appendix A Proofs

We first define the notation and assumptions used throughout. Let  $q_x, q_y \in \mathcal{M}$  be two configurations within a geodesically convex neighbourhood. Then, there exists a unique minimizing geodesic  $\pi : [0, 1] \rightarrow \mathcal{M}$  with  $q_x = \pi(0)$  and  $q_y = \pi(1)$ . We denote the geodesic midpoint by  $q_{\text{mid}} = \pi(1/2)$  and the distance by  $h = d_{\mathcal{M}}(q_x, q_y)$ . We define  $u = \log_{q_{\text{mid}}}(q_y) \in \mathcal{T}_{q_{\text{mid}}}\mathcal{M}$  as the tangent vector at  $q_{\text{mid}}$  pointing toward  $q_y$ . Consequently,  $\|u\|_{q_{\text{mid}}} = d_{\mathcal{M}}(q_{\text{mid}}, q_y) = h/2$ . By the symmetry of the minimizing geodesic, we have  $\log_{q_{\text{mid}}}(q_x) = -u$ . For convenience, we utilize Riemann normal coordinates centred at  $q_{\text{mid}}$ . In these coordinates, the metric tensor at the origin is the identity matrix  $I$ , and its first-order partial derivatives vanish. Accordingly, the coordinates of  $q_x$  and  $q_y$  are given by  $-u$  and  $u$ , respectively.

**Lemma 2.** *Let  $\delta \in \mathbb{R}^n$  be the coordinate representation of the retraction midpoint  $\hat{q}_{\text{mid}}$ . Then  $\delta$  approximates the origin with second-order accuracy.*

*Proof.* Since  $\mathcal{R}_q(0) = q$  and  $D\mathcal{R}_q(0) = I$ , the coordinate expansion is

$$\mathcal{R}_q(v) = q + v + \mathcal{O}(\|v\|^2). \quad (11)$$

Consequently, the inverse retraction satisfies

$$\mathcal{R}_q^{-1}(p) = (p - q) + \mathcal{O}(\|p - q\|^2), \quad (12)$$

for any  $p \in \mathcal{M}$  sufficiently close to  $q$ . Let  $v = \mathcal{R}_{q_x}^{-1}(q_y) \in \mathcal{T}_{q_x}\mathcal{M}$ . Substituting the coordinate representations  $q_x = -u$  and  $q_y = u$  into (12), we expand  $v$  as

$$v = \mathcal{R}_{q_x}^{-1}(q_y) = (q_y - q_x) + \mathcal{O}(\|q_y - q_x\|^2) = 2u + \mathcal{O}(\|u\|^2).$$

Recall from (8) that the retraction midpoint is defined as  $\hat{q}_{\text{mid}} = \mathcal{R}_{q_x}(\frac{1}{2}v)$ . Substituting  $q_x = -u$  and  $v = 2u + \mathcal{O}(\|u\|^2)$ , we expand  $\hat{q}_{\text{mid}}$  using (11)

$$\hat{q}_{\text{mid}} = \mathcal{R}_{q_x}(\frac{1}{2}v) = -u + \left(u + \mathcal{O}(\|u\|^2)\right) + \mathcal{O}(\|u\|^2) = \mathcal{O}(\|u\|^2).$$

Since  $\|\delta\| = \mathcal{O}(\|u\|^2)$ , substituting  $\|u\| = h/2$  gives  $\|\delta\| = \mathcal{O}(h^2)$ .  $\square$

**Lemma 3.** *The difference of inverse retractions at  $\hat{q}_{\text{mid}}$  satisfies*

$$\mathcal{R}_{\hat{q}_{\text{mid}}}^{-1}(q_y) - \mathcal{R}_{\hat{q}_{\text{mid}}}^{-1}(q_x) = 2u + \mathcal{O}(\|u\|^3).$$

*Proof.* Let  $\mathcal{R}(z, \zeta)$  denote the coordinate representation of the retraction, that is,  $\mathcal{R}(z, \zeta)$  is the coordinate of  $\mathcal{R}_{\text{exp}_q(z)}(\zeta)$  in the same chart. Since  $\mathcal{R}(0, 0) = 0$  and  $D_\zeta \mathcal{R}(0, 0) = I$ , a second-order Taylor expansion yields

$$\mathcal{R}(z, \zeta) = z + \zeta + \mathcal{Q}(\zeta, \zeta) + \mathcal{B}(z, \zeta) + \mathcal{O}(\|(z, \zeta)\|^3), \quad (13)$$

where  $\mathcal{Q}$  is quadratic in  $\zeta$  and  $\mathcal{B}$  is bilinear in  $(z, \zeta)$ . Because  $\mathcal{R}_{\hat{q}_{\text{mid}}}$  is a local diffeomorphism near  $0 \in \mathcal{T}_{\hat{q}_{\text{mid}}}\mathcal{M}$ , the inverse-retraction vectors  $w_y = \mathcal{R}_{\hat{q}_{\text{mid}}}^{-1}(q_y)$  and  $w_x = \mathcal{R}_{\hat{q}_{\text{mid}}}^{-1}(q_x)$  are well-defined (for  $\|u\|$  small) and satisfy

$$\mathcal{R}(\delta, w_y) = u, \quad \mathcal{R}(\delta, w_x) = -u.$$

We solve for  $w_y$  and  $w_x$  given base point  $z = \delta$  and targets  $\pm u$ . Let  $w_y = u + e_y$  and  $w_x = -u + e_x$ , where  $e_y, e_x$  are correction terms. Substituting into (13)

$$\begin{aligned} u &= \mathcal{R}(\delta, u + e_y) = \delta + (u + e_y) + \mathcal{Q}(u, u) + \mathcal{B}(\delta, u) + \mathcal{O}(\|u\|^3), \\ -u &= \mathcal{R}(\delta, -u + e_x) = \delta + (-u + e_x) + \mathcal{Q}(-u, -u) + \mathcal{B}(\delta, -u) + \mathcal{O}(\|u\|^3). \end{aligned}$$

Solving for the errors,

$$\begin{aligned} e_y &= -\delta - \mathcal{Q}(u, u) - \mathcal{B}(\delta, u) + \mathcal{O}(\|u\|^3), \\ e_x &= -\delta - \mathcal{Q}(-u, -u) - \mathcal{B}(\delta, -u) + \mathcal{O}(\|u\|^3). \end{aligned}$$

Using the symmetry  $\mathcal{Q}(-u, -u) = \mathcal{Q}(u, u)$  and bilinearity  $\mathcal{B}(\delta, -u) = -\mathcal{B}(\delta, u)$ , subtracting the two expressions cancels out the even terms  $-\delta$  and  $-\mathcal{Q}$

$$w_y - w_x = 2u + (e_y - e_x) = 2u + (-2\mathcal{B}(\delta, u) + \mathcal{O}(\|u\|^3)).$$

Since  $\mathcal{B}$  is bilinear, and by Lemma 2,  $\|\delta\| = \mathcal{O}(\|u\|^2)$ , it follows that  $\mathcal{B}(\delta, u) = \mathcal{O}(\|\delta\|\|u\|) = \mathcal{O}(\|u\|^3)$ . Therefore, we have

$$w_y - w_x = \mathcal{R}_{\hat{q}_{\text{mid}}}^{-1}(q_y) - \mathcal{R}_{\hat{q}_{\text{mid}}}^{-1}(q_x) = 2u + \mathcal{O}(\|u\|^3). \quad \square$$

*Proof (Theorem 1).* In normal coordinates, the metric tensor at  $\hat{q}_{\text{mid}}$  expands as

$$G(\delta) = I + \mathcal{O}(\|\delta\|^2).$$

Substituting  $\|\delta\| = \mathcal{O}(\|u\|^2)$  from Lemma 2, we get  $G(\delta) = I + \mathcal{O}(\|u\|^2)$ . Let  $\Delta w = w_y - w_x$ . Using Lemma 3, we write  $\Delta w = 2u + r$ , where  $r = \mathcal{O}(\|u\|^3)$ . The midpoint retraction distance in (7) is the norm of  $\Delta w$  under  $G(\delta)$ , which relates to the Euclidean norm  $G(0)$  by the metric expansion

$$\begin{aligned} \hat{d}_{\mathcal{M}}(q_x, q_y) &= \|2u + r\|_{G(\delta)} \\ &= \|2u + r\|_{G(0)} \left(1 + \mathcal{O}(\|u\|^2)\right) \\ &= \left(2\|u\| + \mathcal{O}(\|u\|^3)\right) \left(1 + \mathcal{O}(\|u\|^2)\right) \\ &= 2\|u\| + \mathcal{O}(\|u\|^3). \end{aligned}$$

Since  $d_{\mathcal{M}}(q_x, q_y) = 2\|u\| = h$ , we have

$$\left| \hat{d}_{\mathcal{M}}(q_x, q_y) - d_{\mathcal{M}}(q_x, q_y) \right| = \mathcal{O}(h^3). \quad \square$$

## References

1. Bao, R., Wang, J., Wang, S.: Geodesic-based path planning for port transfer robots on Riemannian manifolds. *Expert Systems with Applications* p. 129706 (2025)
2. Beik-Mohammadi, H., Hauberg, S., Arvanitidis, G., Neumann, G., Rozo, L.: Reactive motion generation on learned Riemannian manifolds. *The International Journal of Robotics Research* **42**(10), 729–754 (2023)
3. Belta, C., Kumar, V.: Euclidean metrics for motion generation on  $SE(3)$ . *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science* **216**(1), 47–60 (2002)
4. Biess, A., Flash, T., Liebermann, D.G.: Riemannian geometric approach to human arm dynamics, movement optimization, and invariance. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics* **83**(3), 031927 (2011)
5. Biess, A., Liebermann, D.G., Flash, T.: A computational model for redundant human three-dimensional pointing movements: integration of independent spatial and temporal motor plans simplifies movement dynamics. *Journal of Neuroscience* **27**(48), 13045–13064 (2007)
6. Boumal, N.: *An Introduction to Optimization on Smooth Manifolds*. Cambridge University Press (2023)
7. Bullo, F., Lewis, A.D.: *Geometric control of mechanical systems: modeling, analysis, and design for simple mechanical control systems*, vol. 49. Springer (2019)
8. Carpentier, J., Saurel, G., Buondonno, G., Mirabel, J., Lamiroux, F., Stasse, O., Mansard, N.: The Pinocchio C++ library: A fast and flexible implementation of rigid body dynamics algorithms and their analytical derivatives. In: *2019 IEEE/SICE International Symposium on System Integration (SII)*. pp. 614–619 (2019)
9. Chamzas, C., Quintero-Pena, C., Kingston, Z., Orthey, A., Rakita, D., Gleicher, M., Toussaint, M., Kavraki, L.E.: MotionBenchMaker: A tool to generate and benchmark motion planning datasets. *IEEE Robotics and Automation Letters* **7**(2), 882–889 (2021)
10. Chen, Y., Li, L., Tang, W.: An improved geodesic algorithm for trajectory planning of multi-joint robots. *International Journal of Advanced Robotic Systems* **13**(5), 1729881416657742 (2016)
11. Cheng, C.A., Mukadam, M., Issac, J., Birchfield, S., Fox, D., Boots, B., Ratliff, N.: RMPflow: A geometric framework for generation of multitask motion policies. *IEEE Transactions on Automation Science and Engineering* **18**(3), 968–987 (2021)
12. Crane, K., Weischedel, C., Wardetzky, M.: Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Transactions on Graphics* **32**(5), 1–11 (2013)
13. Detlefsen, N.S., Pouplin, A., Feldager, C.W., Geng, C., Kalatzis, D., Hauschultz, H., González-Duque, M., Warburg, F., Miani, M., Hauberg, S.: *StochMan*. GitHub. Note: <https://github.com/MachineLearningLifeScience/stochman/> (2021)
14. Donald, B., Xavier, P., Canny, J., Reif, J.: Kinodynamic motion planning. *Journal of the ACM (JACM)* **40**(5), 1048–1066 (1993)
15. Flash, T., Handzel, A.A.: Affine differential geometry analysis of human arm movements. *Biological Cybernetics* **96**(6), 577–601 (2007)
16. Flash, T., Karklinsky, M., Fuchs, R., Berthoz, A., Bennequin, D., Meirovitch, Y.: Motor compositionality and timing: combined geometrical and optimization approaches. In: *Biomechanics of Anthropomorphic Systems*, pp. 155–184. Springer (2018)

17. Handzel, A.A., Flash, T.: Geometric methods in the study of human motor control. *Cognitive Studies: Bulletin of the Japanese Cognitive Science Society* **6**(3), 309–321 (1999)
18. Helgason, S.: *Differential Geometry, Lie Groups, and Symmetric Spaces*, vol. 80. Academic Press (1979)
19. Henderson, M.E.: Multiple parameter continuation: Computing implicitly defined k-manifolds. *International Journal of Bifurcation and Chaos* **12**(03), 451–476 (2002)
20. Jaillet, L., Porta, J.M.: Path planning under kinematic constraints by rapidly exploring manifolds. *IEEE Transactions on Robotics* **29**(1), 105–117 (2012)
21. Jaquier, N., Asfour, T.: Riemannian geometry as a unifying theory for robot motion learning and control. In: *The International Symposium of Robotics Research (ISRR)*. pp. 395–403. Springer (2022)
22. Jaquier, N., Rozo, L., Caldwell, D.G., Calinon, S.: Geometry-aware manipulability learning, tracking, and transfer. *The International Journal of Robotics Research* **40**(2-3), 624–650 (2021)
23. Karaman, S., Frazzoli, E.: Sampling-based algorithms for optimal motion planning. *The International Journal of Robotics Research* **30**(7), 846–894 (2011)
24. Kim, B., Um, T.T., Suh, C., Park, F.C.: Tangent bundle RRT: A randomized algorithm for constrained motion planning. *Robotica* **34**(1), 202–225 (2016)
25. Kingston, Z., Moll, M., Kavraki, L.E.: Sampling-based methods for motion planning with constraints. *Annual Review of Control, Robotics, and Autonomous Systems* **1**(1), 159–185 (2018)
26. Kingston, Z., Moll, M., Kavraki, L.E.: Exploring implicit spaces for constrained sampling-based planning. *The International Journal of Robotics Research* **38**(10-11), 1151–1178 (2019)
27. Klein, H., Jaquier, N., Meixner, A., Asfour, T.: On the design of region-avoiding metrics for collision-safe motion generation on Riemannian manifolds. In: *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 2346–2353 (2023)
28. Kuffner, J.J., LaValle, S.M.: RRT-Connect: An efficient approach to single-query path planning. In: *Proceedings 2000 ICRA. Millennium conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*. vol. 2, pp. 995–1001. IEEE (2000)
29. Laumond, J.P., Jacobs, P.E., Taix, M., Murray, R.M.: A motion planner for non-holonomic mobile robots. *IEEE Transactions on Robotics and Automation* **10**(5), 577–593 (2002)
30. Laux, M., Zell, A.: Robot arm motion planning based on geodesics. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 7585–7591 (2021)
31. Laux, M., Zell, A.: Boundary conditions in geodesic motion planning for manipulators. In: *2023 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 1558–1564 (2023)
32. LaValle, S.M.: *Planning Algorithms*. Cambridge University Press (2006)
33. LaValle, S.M., Kuffner Jr, J.J.: Randomized kinodynamic planning. *The International Journal of Robotics Research* **20**(5), 378–400 (2001)
34. Lee, J.M.: *Introduction to Smooth Manifolds*. Springer, 2 edn. (2012)
35. Lee, J.M.: *Introduction to Riemannian Manifolds*. Springer, 2 edn. (2018)
36. Li, Y., Qiu, J., Calinon, S.: A Riemannian take on distance fields and geodesic flows in robotics. *The International Journal of Robotics Research* p. 02783649261420233 (2024)

37. Mainprice, J., Ratliff, N., Schaal, S.: Warping the workspace geometry with electric potentials for motion optimization of manipulation tasks. In: 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 3156–3163 (2016)
38. Mainprice, J., Ratliff, N., Toussaint, M., Schaal, S.: An interior point method solving motion planning problems with narrow passages. In: 2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN). pp. 547–552 (2020)
39. Marić, F., Petrović, L., Guberina, M., Kelly, J., Petrović, I.: A Riemannian metric for geometry-aware singularity avoidance by articulated robots. *Robotics and Autonomous Systems* **145**, 103865 (2021)
40. Milnor, J.: Curvatures of left invariant metrics on Lie groups. *Advances in Mathematics* **21**(3), 293–329 (1976)
41. Mirebeau, J.M.: Anisotropic fast-marching on cartesian grids using Voronoi’s first reduction of quadratic forms. HAL preprint (2017)
42. Peyré, G., Péchaud, M., Keriven, R., Cohen, L.D., et al.: Geodesic methods in computer vision and graphics. *Foundations and Trends® in Computer Graphics and Vision* **5**(3–4), 197–397 (2010)
43. Quigley, M., Conley, K., Gerkey, B., Faust, J., Foote, T., Leibs, J., Wheeler, R., Ng, A.Y., et al.: ROS: An open-source robot operating system. In: ICRA Workshop on Open Source Software. vol. 3, p. 5. Kobe (2009)
44. Ratliff, N., Toussaint, M., Schaal, S.: Understanding the geometry of workspace obstacles in motion optimization. In: 2015 IEEE International Conference on Robotics and Automation (ICRA). pp. 4202–4209 (2015)
45. Ratliff, N., Zucker, M., Bagnell, J.A., Srinivasa, S.: CHOMP: Gradient optimization techniques for efficient motion planning. In: 2009 IEEE International Conference on Robotics and Automation (ICRA). pp. 489–494 (2009)
46. Ratliff, N.D., Issac, J., Kappler, D., Birchfield, S., Fox, D.: Riemannian motion policies. arXiv preprint arXiv:1801.02854 (2018)
47. Salzman, O., Hemmer, M., Raveh, B., Halperin, D.: Motion planning via manifold samples. *Algorithmica* **67**(4), 547–565 (2013)
48. Saveriano, M., Abu-Dakka, F.J., Kyrki, V.: Learning stable robotic skills on Riemannian manifolds. *Robotics and Autonomous Systems* **169**, 104510 (2023)
49. Sekimoto, M., Arimoto, S., Prilutsky, B.I., Isaka, T., Kawamura, S.: Observation of human multi-joint arm movement from the viewpoint of a Riemannian distance. In: 2009 ICCAS-SICE. pp. 2664–2669 (2009)
50. Sethian, J.A.: A fast marching level set method for monotonically advancing fronts. *Proceedings of the National Academy of Sciences* **93**(4), 1591–1595 (1996)
51. Shao, H., Kumar, A., Thomas Fletcher, P.: The Riemannian geometry of deep generative models. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 315–323 (2018)
52. Sucas, I.A., Chitta, S.: Motion planning with constraints using configuration space approximations. In: 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 1904–1910 (2012)
53. Sucas, I.A., Moll, M., Kavraki, L.E.: The open motion planning library. *IEEE Robotics & Automation Magazine* **19**(4), 72–82 (2012)
54. Surazhsky, V., Surazhsky, T., Kirsanov, D., Gortler, S.J., Hoppe, H.: Fast exact and approximate geodesics on meshes. *ACM Transactions on Graphics* **24**(3), 553–560 (2005)
55. Todorov, E.: Optimality principles in sensorimotor control. *Nature Neuroscience* **7**(9), 907–915 (2004)

56. Wilmarth, S.A., Amato, N.M., Stiller, P.F.: Motion planning for a rigid body using random networks on the medial axis of the free space. In: Proceedings of the Fifteenth Annual Symposium on Computational Geometry. pp. 173–180 (1999)
57. Zhang, Y., Zhou, Q., Yang, X.S.: An RRT\* algorithm based on Riemannian metric model for optimal path planning. arXiv preprint arXiv:2507.01697 (2025)