Manifold-Guided Motion Planning for Tight Assemblies

Dror Livnat, Michael M. Bilevich, Michal Kleinbort, Dan Halperin

Blavatnik School of Computer Science and AI, Tel Aviv University
drorlivnat@mail.tau.ac.il

Abstract. Motion planning for rigid-body assembly poses a fundamental challenge in robotics due to tight geometric constraints. In such scenarios, feasible motions often require passing through (near-)zero clearance configurations in which the parts are tightly constrained by contact. In this work, we introduce Critical-Manifold Guided RRT (CMG-RRT), a sampling-based planner designed specifically for tight assembly problems. Our key observation is that in tight assemblies, valid solution paths lie on or near a critical manifold: the subset of configuration space consisting of poses with at least one contact point between parts. CMG-RRT guides exploration by adaptively biasing sampling toward neighborhoods of the critical manifold using a hierarchical subdivision of the configuration space. We prove that CMG-RRT is probabilistically complete under standard clearance assumptions. Empirical evaluation on challenging rotational assembly benchmarks demonstrates a 100% success rate across all tested instances, including, to the best of our knowledge, the first fully automatic solution of the *Elk* disentanglement puzzle. Our open source software is available through our project page: https://www.cgl.cs.tau.ac.il/projects/tight-assembly-planning

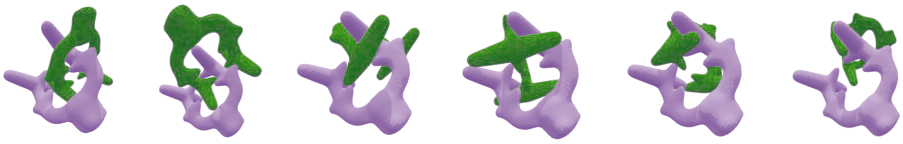

Fig. 1: Snapshots of the *Elk* disentanglement puzzle, previously unsolved (to the best of our knowledge), showing a complete solution generated by our algorithm. From left to right: the start configuration, traversal through four narrow C-space tunnels interleaved with wide free-space regions, and the goal configuration.

1 Introduction

Sampling-based motion planning is a fundamental paradigm in robotics, used in applications such as autonomous navigation, manipulation, assembly, and numerous others [38]. Methods such as Probabilistic Roadmaps (PRM) and Rapidly-Exploring Random Trees (RRT) have become standard due to their scalability

to high-dimensional configuration spaces and their probabilistic completeness guarantees [28,37,39,64]. When feasible motions have sufficient clearance from obstacles, uniform random sampling in configuration space enables these planners to efficiently discover collision-free paths.

A longstanding difficulty arises in *tight* motion-planning problems, where any valid solution must pass through regions of extremely small measure in configuration space, commonly referred to as *narrow passages* or *tunnels* [22,29,50]. Such situations are common in robotic assembly, particularly for rigid parts that require tightly coupled translation and rotation. In these settings, uniform sampling becomes highly inefficient: the overwhelming majority of samples lie in regions irrelevant to the solution, while the critical regions that enable progress are sampled with vanishing probability. This failure mode is a canonical limitation of sampling-based planners and is often illustrated through "bug-trap" examples [38].

To address narrow passages, a variety of nonuniform sampling strategies have been proposed, including bridge tests, obstacle-based sampling, spatial subdivisions,and Gaussian or medial-axis inspired heuristics [21,60,65,69]. While these methods can improve performance on certain instances, they typically rely on geometric heuristics that are unreliable in high-dimensional spaces and offer no general guarantees for complex, contact-rich motions. The Soft Subdivision Search (SSS) framework for motion planning [67,68,71] has guarantees, but applying it to sufficiently complex assemblies or robots seems (as of now) a prohibitively hard task.

Many tight assemblies may require sliding or rolling along obstacle boundaries, rather than maintaining clearance throughout the motion. This observation has motivated work on compliant and contact-aware motion planning [15,41,51], as well as physics- and simulation-based planners that can exploit contact and near-contact interactions to solve tight disentanglement tasks [61]. Feature-driven tunnel discovery methods similarly leverage local geometric structure to guide search toward narrow passages without assuming continuous contact [70].

While these approaches can be effective for problems dominated by tight interactions or exhibiting specific structural cues, they are less suited to mixed settings in which wide free-space motion is interleaved with short but critical tight transitions whose locations are not known in advance.

Recent work has explored learning-based and data-driven techniques to guide sampling toward narrow passages or critical connectivity regions [40,42] Another influential line of work studies motion planning on *constraint manifolds*. Projection-based planners such as CBiRRT [5] and atlas-based methods, including AtlasRRT and its variants [24,62], explore lower-dimensional manifolds defined implicitly by kinematic or task constraints. Comprehensive surveys [30,31] have formalized this perspective and shown its effectiveness for problems with fixed, known constraints such as closed kinematic chains. In contrast, contact constraints in assembly planning are neither fixed nor known a priori: the set of active contacts changes along the motion, and the corresponding manifolds appear and disappear as the robot moves.

In prior work, we introduced TouchRoll-RRT (`TR-RRT`) [45], a sampling-based planner designed for mixed wide-tight settings. `TR-RRT` combines standard RRT exploration in free space with a contact-aware extension procedure guided by a

signed distance function (SDF). When the search reaches the vicinity of obstacles, `TR-RRT` maintains multiple contact points and advances by projecting steering directions onto the tangent space of a local *contact critical manifold*, followed by a retraction step. This reduces the effective degrees of freedom from six to five, four, and at times even three or two, and enables traversal of narrow passages that outperforms standard RRT.

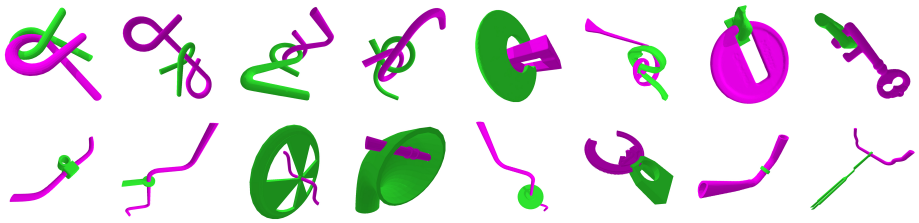

Fig. 2: Dataset provided by Tian et al. [61]. Excerpts from the more challenging *Rotational* benchmarks: `Puzzles` in the first row and `Others` in the second row.

However, `TR-RRT` retains a fundamental mismatch between sampling and exploration. While its extension step explicitly follows the critical manifold, its sampling step remains uniform in the six-dimensional space SE(3). As a result, only samples that happen to induce extensions near the manifold benefit from the contact-aware machinery. Many samples are generated in regions that do not contribute to progress towards the goal causing the tree to repeatedly grow into misleading dead ends—an effect similar to bug-trap behavior [38]. This sampling bias limits performance and prevents `TR-RRT` from solving particularly challenging instances.

In this work, we present *Critical-Manifold Guided RRT* (`CMG-RRT`), which resolves this limitation by explicitly biasing sampling toward the critical manifolds. We consider the configuration space $\mathcal{C} = \mathrm{SE}(3)$, the six-dimensional manifold of rigid body transformations in $\mathbb{R}^3$. We maintain an adaptive subdivision of the configuration space into axis-aligned boxes and discard boxes that are provably far from the contact manifold, using an SDF-based distance oracle. Sampling is restricted to the remaining boxes, whose union is contained in an increasingly small D-neighborhood of the critical manifold as the box diagonal D decreases. In addition, we occasionally apply random rotations to the entire system, an idea already in use for other purposes in the field, such as efficiently finding pairs of nearest neighbors [2,32]. In this work it helps avoiding gimbal locks [19,38], as well as enables the completeness proof.

Search and refinement are interleaved adaptively, so that easy problems incur little overhead, while harder problems trigger additional refinement that sharpens the sampling distribution.

As we prove below, this novel manifold-guided sampling scheme allows `CMG-RRT` to be *probabilistically complete*; that is, under standard assumptions of positive clearance, the probability that `CMG-RRT` finds a solution approaches one as the number of iterations grows. Moreover, `CMG-RRT` accelerates planning on problems already solvable by `TR-RRT` and, crucially, solves significantly harder instances.

A notable example is the *Elk* disentanglement puzzle[1], previously identified as particularly challenging due to the absence of exploitable geometric features and the presence of multiple dead-end tunnels [45,70]. Using `CMG-RRT`, we demonstrate—to the best of our knowledge—the first fully automatic solution of this puzzle using a general-purpose motion planner without handcrafted features.

Our Contribution Our contributions are as follows: (i) We introduce `CMG-RRT`, a sampling-based planner that adaptively concentrates samples near contact critical-manifolds. (ii) We extend previous probabilistic completeness proofs to an SE(3) parametrization with non-Euclidean topology, and show how choices in our algorithm directly benefit the completeness proof. (iii) We present experimental results, demonstrating the effectiveness of our method over prior methods. (iv) Our software is open-source and publicly available online. [2]

In particular, we show that our method solves challenging assembly problems, including the Elk puzzle. In [70], the authors mention the Elk puzzle as a tough challenge that their algorithm fails to solve. In [45], we tried applying TR-RRT to the Elk puzzle, and although it was able to go through individual c-space tunnels, it failed to solve the entire puzzle due to the c-space uniform sampling. This is the first time that an algorithm solving the Elk puzzle is presented, as far as we know.

Organization The remainder of the paper is organized as follows. Section 2 reviews the background and formalizes the problem. Section 3 presents the `CMG-RRT` algorithm. Section 4 provides a probabilistic completeness proof. Experimental results are reported in Section 5, followed by discussion and future directions in Section 6.

2 Preliminaries and Problem Statement

2.1 The SO(3) and SE(3) Groups and the Configuration Space

When dealing with rigid body transformations, there are two groups that naturally arise [18,53]—the special orthogonal group SO(3) and the special Euclidean group SE(3). Both are Lie groups and are extensively researched in mathematics, physics, and robotics.

As previously mentioned, in this work, the configuration space [38] is $\mathcal{C} = \mathrm{SE}(3)$, as we solve for the assembly of a part among other static parts.

We note that SO(3) is a three-dimensional manifold. However, there are many common methods of representing orientations, and each has its advantages and shortcomings. Examples are unit quaternions, angle axis, and Euler angles, to name a few [11]. In this work, we use only the Euler angle representation, and in particular, the *rpy* (roll-pitch-yaw) representation.

Recall the notion of the Haar measure on a group [14], which is a generalization of the Lebesgue measure to locally compact groups. The Haar measure is a measure that is invariant under group multiplication and is unique up to a constant scaling

[1] A commercially available cast-metal disentanglement puzzle designed by Nobuyuki Yoshigahara and marketed by Hanayama of Japan.

[2] Project page: `www.cgl.cs.tau.ac.il/projects/tight-assembly-planning`

factor. For example, the Lebesgue measure on $\mathbb{R}^n$ is also a Haar measure when $\mathbb{R}^n$ is viewed as an additive group. Since SO(3) is a *compact* Lie group, its Haar measure can be normalized to a probability measure, providing a notion of uniform sampling over rotations.

We focus on the *rpy* representation, let $\phi,\psi \in [-\pi,\pi)$ be the roll and yaw, respectively, and $\theta \in [-\pi/2,\pi/2]$ be the pitch. Then the volume element of the Haar measure using this representation [14] is

$$\frac{1}{8\pi^2}\cos\theta \; d\phi \; d\theta \; d\psi \,. \tag{1}$$

Finally, we note that due to this parametrization of SO(3), we regard our configuration space as $\mathcal{C} = \mathrm{SE}(3) \simeq \mathbb{R}^3 \times [-\pi,\pi) \times [-\pi/2,\pi/2] \times [-\pi,\pi) \subset \mathbb{R}^6$. Hence, in this work, we regard points in the Lie group SE(3) as their six-dimensional (translation + *rpy*) representation.

2.2 Pseudo-metric Spaces

Metric spaces [44,52] arise in many different applications. A metric space (X,d) is some set X equipped with a function $d: X \times X \to \mathbb{R}$ called *metric* that satisfies three axioms: (i) (*symmetry*) $d(x,y) = d(y,x)$, (ii) (*positivity*) $d(x,y) \geq 0$ and equality if and only if $x = y$, and (iii) for any $z \in X$, (*triangle inequality*) $d(x,y) \leq d(x,z) + d(z,y)$.

Similarly, we can define a *pseudo-metric* space [20], by replacing axiom (ii) with a softer requirement: for any $x,y \in X$, $d(x,y) \geq 0$ and $d(x,x) = 0$. Hence, we possibly allow for different points $x \neq y$ to have $d(x,y) = 0$.

One example is $d_{\mathbb{R}^3} : \mathbb{R}^6 \times \mathbb{R}^6 \to \mathbb{R}$, defined by $d_{\mathbb{R}^3}(x,y) = \sqrt{\sum_{i=1}^{3}(x_i - y_i)^2}$. Clearly, it is symmetric and has the triangle inequality (since the Euclidean norm on $\mathbb{R}^3$ has the triangle inequality). However, if $e_1,\ldots,e_6$ are the standard basis of $\mathbb{R}^6$, then $d_{\mathbb{R}^3}(e_4,e_5) = 0$, meaning the $d_{\mathbb{R}^3}$ is not a metric, but is a pseudo-metric.

Let us define the *great-circle metric* on $\mathbb{S}^1$, which we parametrize as $\mathbb{S}^1 \simeq [-\pi,\pi)$. Take $d_{\mathbb{S}^1} : \mathbb{S}^1 \times \mathbb{S}^1 \to \mathbb{R}$,

$$d_{\mathbb{S}^1}(x,y) = \min(|x-y|, 2\pi - |x-y|) \,. \tag{2}$$

Lemma 1. *The function $d_{\mathbb{S}^1}$ is a metric on $\mathbb{S}^1 \simeq [-\pi,\pi)$.*

Proof. Notice that this function is the distance of the shorter arc length (in radians) between x and y. □

We can extend the metric $d_{\mathbb{S}^1}$ to a pseudo-metric on $\mathcal{C}$: choose some index $i \in \{1,\ldots,6\}$. Then $d_{\mathbb{S}^1[i]} : \mathcal{C} \times \mathcal{C} \to \mathbb{R}$, defined as $d_{\mathbb{S}^1[i]}(x,y) = d_{\mathbb{S}^1}(x_i,y_i)$ is indeed a pseudo-metric on $\mathcal{C}$.

Notice the following lemma on pseudo-metrics:

Lemma 2. *Let $d_1,d_2 : X \times X \to \mathbb{R}$ be pseudo-metrics. Then $d_3 : X \times X \to \mathbb{R}$, defined by*

$$d_3(x,y) \coloneqq \sqrt{(d_1(x,y))^2 + (d_2(x,y))^2} \tag{3}$$

is also a pseudo-metric.

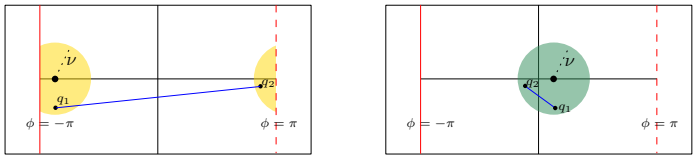

Fig. 3: Illustrations of projections of $\mathcal{B}_\nu^{\mathcal{C}}(\cdot)$ balls onto $[-\pi,\pi)\times\mathbb{R}$. The center of the yellow ball is located in a subset of $[-\pi,\pi)\times\mathbb{R}$ such that the ball is not convex. Therefore, there exist points q_1, q_2 in the ball such that the straight line segment (in blue) connecting them does not lie in the ball. With high probability (specified in Lemma 3), a random rotation will shift the center of the ball to a new position (see the green ball) such that the ball is convex.

Proof. The proof is straightforward by using the Cauchy-Schwarz inequality. □

Finally, we combine $d_{\mathbb{R}^3}$ and $d_{\mathbb{S}^1}$ into a pseudo-metric on our configuration space $\mathcal{C} = \mathrm{SE}(3) \subset \mathbb{R}^6$. Define $d_{\mathcal{C}} : \mathcal{C}\times\mathcal{C} \to \mathbb{R}$ as

$$d_{\mathcal{C}}(x,y) = \sqrt{(d_{\mathbb{R}^3}(x,y))^2 + \sum_{i=4}^{6} (d_{\mathbb{S}^1[i]}(x,y))^2} \tag{4}$$

Proof. Follows immediately by applying Lemma 2 four times. □

We choose this $d_{\mathcal{C}}$ to be our distance function as it is very simple to compute, benefits the manifold search method described in Section 3, and it is a sufficient choice to show that our proposed method is probabilistically complete.

Using the distance function $d_{\mathcal{C}}$, we can define an open ball of radius r centered at a point q, $\mathcal{B}_r^{\mathcal{C}}(q)$, as the set of all points whose distance $d_{\mathcal{C}}$ from q is less than r. Note that such balls are not necessarily path-connected in the natural Euclidean topology. See Figure 3 for an example. As a consequence, in the topology induced by $d_{\mathcal{C}}$, balls $\mathcal{B}_r^{\mathcal{C}}(q)$ are not necessarily convex.

2.3 Signed Distance Functions (SDF)

Signed distance functions (sometimes also referred to as signed distance fields, or SDFs) are a common approach for representing geometry in computer graphics [26]. The geometry is represented as the zero-level set of its SDF. Formally, assume that we have some compact manifold[3] $M \subseteq \mathbb{R}^3$. In this work, we assume that this manifold M is given as a triangle soup. The signed distance function (SDF) of M is a function $F_M : \mathbb{R}^3 \to \mathbb{R}$ that maps a point $p \in \mathbb{R}^3$ to the distance to its closest point on the manifold M. The distance is signed positive if p is outside M, zero if $p \in M$ and negative if p is inside M.

A straightforward approach for computing the SDF is to directly query the underlying geometry for the distance. For example, one may compare the query

[3] We assume that the manifold is compact and without a boundary, hence, water-tight [59].

point against all triangles in the triangle soup representation of M to find the nearest distance, and use the parity of the number of intersections of an arbitrary ray with M to determine the sign of the SDF for that query point [8,35]. While there are data structures and algorithms that significantly speed up this query [12,35,55], a common approach is to simplify this SDF by evaluating it on a grid, and use that grid as an approximation for the SDF value when querying new points [61,72]. In recent years, there has also been an emergence of deep-learning approaches for approximating the SDF of a given object [10,13,23,43,49,54,66].

One straightforward application of the SDF is as a tool for collision detection. Collision detection is the task of determining whether two (or more) objects overlap [25]. We can represent one object as a collection of points, and evaluate each point in the SDF of the second body [4,6,23,34,47,63]. If there are any points with negative distance, then they penetrate the second object and thus a collision occurs. Conversely, if all points have non-negative signed distance, we assume that no collision occurs if the point sample set is sufficiently dense.

2.4 Problem Statement

Finally, we formally define the problem we address in this work. Assume that $M_1, M_2 \subseteq \mathbb{R}^3$ are two rigid bodies, which are compact sub-manifolds, with or without a boundary. In this work, we assume that both bodies are scaled down such that they fit in the unit sphere in $\mathbb{R}^3$. We assume we have some digital representation of those M_1 and M_2 for which we can sample random points on the bodies' boundary, and for which we can efficiently and effectively evaluate the SDF $F_{M_1} : \mathbb{R}^3 \to \mathbb{R}$ that is defined in Section 2.3. We assume that M_1 is static, and is referred to as the *obstacle(s)*, and M_2 can freely translate and rotate in the workspace, and is referred to as the *robot*. We also define the following function $F_{M_2 \to M_1} : \mathrm{SE}(3) \to \mathbb{R}$, which is the smallest signed distance of a point on the manifold M_2 transformed by a configuration $q \in \mathrm{SE}(3)$ from the static obstacle M_1. Formally:

$$F_{M_2 \to M_1}(q) = \min_{p \in M_2} F_{M_1}(q \cdot p), \tag{5}$$

where $q \cdot p$ is the application of a rigid body transformation $q \in \mathrm{SE}(3)$ on a three-dimensional point $p \in \mathbb{R}^3$.

The *δ-free space*, denoted by $\mathcal{F}_\delta \subseteq \mathcal{C}$ is defined as

$$\mathcal{F}_\delta := \{q \in \mathcal{C} \mid F_{M_2 \to M_1}(q) > -\delta\}, \tag{6}$$

i.e, the set of all configurations for which the robot's penetration into obstacles is at most δ. We refer to this δ as the *allowance*, or the allowed penetration.

A *motion-planning problem* is implicitly defined by the triplet $(\mathcal{F}_\delta, q_{\text{start}}, q_{\text{goal}})$, with $q_{\text{start}}, q_{\text{goal}} \in \mathcal{F}_\delta$. A solution to such a problem is a continuous path that moves the robot from the initial configuration q_{start} to the goal q_{goal} while avoiding

collision with obstacles. Formally, a valid path is a continuous[4] map $\gamma : [0,1] \to \mathcal{F}_\delta$, such that $\gamma(0) = q_{\text{start}}$ and $\gamma(1) = q_{\text{goal}}$.

3 Critical-Manifold Guided RRT (CMG-RRT)

In this section, we introduce our proposed algorithm. We begin with the motivation in Subsection 3.1, identifying the challenges inherent in tight assembly planning and motivating the need for a more principled solution. We then summarize the TR-RRT algorithm [45], which addresses these challenges using contact-aware exploration. Finally, in Subsection 3.2, we present our new method, *Critical-Manifold Guided RRT* (CMG-RRT), and describe its core principles and implementation details.

3.1 Motivation

The introduction motivates tight assembly planning at a high level. We revisit the motivation here to make explicit the algorithmic bottleneck that CMG-RRT is designed to resolve: in mixed wide-tight settings, the difficulty is not only in generating contact-rich motions once near obstacles, but in *reaching* the relevant near-contact regions efficiently.

Sampling-based planners such as RRT [37] and PRM [28] perform well when feasible paths have wide clearance. In tight phases, simulation-based [61] and contact-aware methods [41] can exploit local geometric constraints. Still, many practical assembly tasks interleave wide-clearance motion with short, critical near-contact transitions whose locations are unknown a priori. Uniform sampling in SE(3) is therefore inefficient, as most samples fall in regions that do not contribute to progress.

Several methods bias exploration toward narrow passages, e.g., by detecting geometric cues for tunnels [70]. While effective on certain families, such cues may be weak or absent in puzzles like the Elk (as mentioned in [70]), motivating a bias that depends only on a general distance-to-contact oracle rather than problem-specific features.

Our objective is thus to leverage the structure of the *critical manifold* to concentrate sampling in the small subset of SE(3) that enables tight transitions, while still allowing efficient exploration of wide regions. This motivates the subdivision-based sampling scheme introduced below.

3.2 Critical-Manifold Guided RRT (CMG-RRT)

As a brief reminder, TouchRoll-RRT (TR-RRT) [45] augments a standard RRT in SE(3) with a contact-aware extension procedure based on signed distance function (SDF) queries. During tree expansion, TR-RRT steers toward a random sample as usual when the motion has clearance; however, when an extension reaches the δ-vicinity of obstacles, it identifies contact points via SDF values and uses the

[4] Note that γ is not necessarily continuous under the natural Euclidean topology, but rather in the topology induced by the pseudo-metric d_C.

corresponding SDF gradients to maintain sliding/rolling motion along a local contact critical manifold. In effect, this reduces the number of degrees of freedom during tight phases and enables traversal through narrow C-space tunnels that defeat a purely free-space RRT.

Despite this manifold-aware *extension* step, TR-RRT still *samples* uniformly in SE(3), and therefore benefits from its contact-aware machinery only when random samples happen to induce extensions near the critical manifold. As a result, many samples do not contribute to progress and may repeatedly grow the tree into misleading dead ends. We address this mismatch by modifying the sampling process itself: CMG-RRT adaptively concentrates sampling near the critical manifold using a hierarchical subdivision scheme.

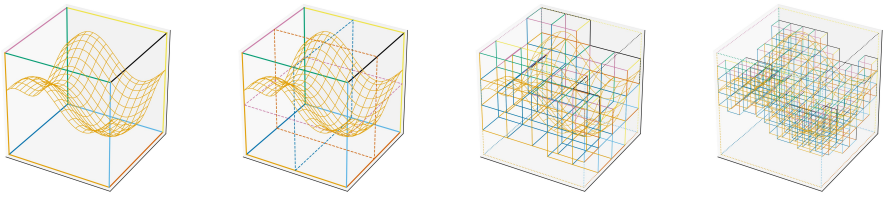

Fig. 4: Illustration of the REFINE (subdivision) procedure in $3D$. In CMG-RRT, this refinement takes place in $6D$.

Figure 4 illustrates the central mechanism in CMG-RRT: an adaptive subdivision that progressively filters the configuration space to retain only regions that may lie near the critical manifold.[5] We next formalize the resulting planner and its refinement search schedule.

At a high level, CMG-RRT replaces uniform sampling by a subdivision-based sampler [7] that concentrates samples near the critical manifold. The idea is to recursively subdivide SE(3) into shrinking axis-aligned boxes, discard boxes whose configurations are guaranteed not to intersect the critical manifold, and sample only from the remaining boxes (see Figure 4). A more formal description follows.

SDF and distance computation In the original TR-RRT [45], the SDF F_{M_1} was approximated by a neural network, evaluated via GPU forward/backward passes [10,49]. In CMG-RRT we instead use a three-dimensional grid with trilinear interpolation [16]; gradients are computed by centered finite differences. This yields faster query times and allows for efficient parallelization on CPU cores.

We now provide a high-level overview of the algorithmic flow, highlighting how refinement, sampling, and tree expansion interact, and referring to Algorithms 1-3 for precise pseudo-code.

Subdivision module The configuration space is six-dimensional with ranges $x, y, z \in [-1, 1]$, $\phi, \psi \in [-\pi, \pi)$, and $\theta \in [-\pi/2, \pi/2]$. The SUBDIVISION module

[5] We note that the number of boxes maintained depends on the Hausdorff measure of the critical manifold and is $O(1/\delta^k)$ with $k \leq 5$ [7], which is asymptotically better than naively splitting the configuration space into $O(1/\delta^6)$ boxes.

Algorithm 1 Critical-Manifold Guided RRT

Require: $\eta > 0, \delta > 0$
Require: $q_{\text{start}}, q_{\text{goal}} \in \mathcal{F}_\delta$
$V \leftarrow \{q_{\text{start}}\}$, $E \leftarrow \emptyset$
$T_{\text{expand}} \leftarrow 0, T_{\text{refine}} \leftarrow 0, \text{factor}_{\text{refine}} \leftarrow 10$
$T \leftarrow 0$
while $T <$ LIMIT **do**
  if $T_{\text{refine}} \cdot \text{factor}_{\text{refine}} < T_{\text{expand}}$ **then**
    REFINE()
    UPDATE(T_{refine})
  if $i \bmod N_{\text{rot}} = 0$ **then**
    RANDOMROTATE()
  $s \leftarrow$ SAMPLE()
  $q_{\text{near}} \leftarrow \text{argmin}_{q \in V}(d_{\mathcal{C}}(q, s))$
  $v \leftarrow$ DIRECTION(q_{near}, s)
  $q_{\text{new}} \leftarrow$ EXTEND(q_{near}, v, η)
  if ISEDGEVALID($q_{\text{new}}, q_{\text{goal}}$) **then**
    return GETPATH($V, E, q_{\text{start}}, q_{\text{goal}}$)
  $V \leftarrow V \cup \{q_{\text{new}}\}$
  $E \leftarrow E \cup \{(q_{\text{near}}, q_{\text{new}})\}$
  UPDATE(T_{search})
  $T \leftarrow T_{\text{search}} + T_{\text{refine}}$
return NONE

Algorithm 2 REFINE(SUBDIVISION)

Require: Box set $\mathbb{B}$; distance oracle $\mathsf{dist}(\cdot)$
$a \leftarrow \text{argmax}_{\text{axis} \in \{x, y, z, R_x, R_y, R_z\}}$(side length of boxes in $\mathbb{B}$) ▷ largest dimension
$\mathbb{B}_{\text{new}} \leftarrow \emptyset$
for each $b \in \mathbb{B}$ **do**
  Split b into two children b^-, b^+ by bisecting along axis a
  for each $b' \in \{b^-, b^+\}$ **do**
    $q_c \leftarrow center(b')$; $D \leftarrow Diagonal(b')$
    $d \leftarrow F_{M_2 \to M_1}(q_c)$
    if $d < D/2$ **then**
      $\mathbb{B}_{\text{new}} \leftarrow \mathbb{B}_{\text{new}} \cup \{b'\}$ ▷ keep potentially near contact
$\mathbb{B} \leftarrow \mathbb{B}_{\text{new}}$

maintains a finite set of axis-aligned boxes $\mathbb{B}$; each box $b \in \mathbb{B}$ is represented by its center and shares a common diagonal length D. During each call to REFINE (Algorithm 2), the boxes are bisected along their longest dimension, and each child box is tested for potential intersection with the contact manifold. Boxes for which $|F_{M_2 \to M_1}(\text{Center}(b))| < D/2$ are retained. SAMPLE (Algorithm 3) selects a box uniformly from $\mathbb{B}$ and then selects a configuration uniformly from within that box.

Algorithm description `CMG-RRT` begins with several iterations of REFINE. Then, during search, whenever the relative computational budget spent on refinement falls below a threshold, another refinement step is invoked. This balances refinement cost with search cost: easy problems require few refinements, while hard

Algorithm 3 SAMPLE(SUBDIVISION)

Require: Nonempty box set $\mathbb{B}$
1: $b \leftarrow$ **uniformly at random** select a box from $\mathbb{B}$
2: $q \leftarrow$ **uniformly at random** sample a configuration from b ▷ independent uniform over each coordinate range of b
3: **return** q

problems trigger more refinement steps, effectively shrinking boxes toward the contact manifold and yielding highly informative samples.

This refinement-driven sampling accelerates performance on previously solved problems (e.g., the Alpha puzzle) and, crucially, enables solving significantly harder instances such as the Elk puzzle (Fig. 1). Zhang et al. [70] identified Elk as a particularly challenging case due to the absence of identifiable per-piece geometric features and the presence of multiple tunnels involving disjoint contact points from both pieces simultaneously. Our manifold-guided sampling circumvents these limitations by exploring the C-space directly along its critical manifold, rather than relying on local structural cues.

4 Probabilistic Completeness

We derive a probabilistic completeness proof for `CMG-RRT`. We assume that there exists a valid path $\gamma:[0,1]\to\mathcal{F}_\delta$, where $\mathcal{F}_\delta$ is the set of all configurations for which the robot's penetration into the obstacles is at most $\delta>0$.

We follow our general approach as described in [33,56], but we adapt it to handle rotations explicitly, and we connect it to the boxes that arise in the subdivision process. We define a sequence of balls covering the path and show that with high probability, `CMG-RRT` will generate a path that goes through the union of these balls in the order of the sequence. We show that this probability converges to one as the number of samples tends to infinity.

Denote by L the length of the path γ under the pseudo-metric $d_{\mathcal{C}}$. Let $m=\frac{5L}{\nu}$, where $\nu=\min(\delta,\eta)$, and η is the maximal step size used by the algorithm. Then, define a sequence of $m+1$ points $q_0=\gamma(0),\ldots,q_m=\gamma(1)$ along γ, such that the $d_{\mathcal{C}}$ length of the sub-path between every two consecutive points is $\nu/5$. Therefore, $d_{\mathcal{C}}(q_i,q_{i+1})\le\nu/5$ for every $0\le i<m$. Next, we define a set of $m+1$ balls of radius $\nu/5$, centered at these points. We now prove that with high probability, `CMG-RRT` will generate a path that goes through these balls.

First, we claim that for every configuration $q\in\mathcal{C}$, we can sample a random rotation R such that with high probability the ball $\mathcal{B}^{\mathcal{C}}_{\nu}(q')$ centered at $q'=R\cdot q$ is convex. Denote this probability by p_{rotate}.

Lemma 3. *Fix some $q\in\mathcal{C}$. Let $R\in\mathrm{SO}(3)$ which is sampled uniformly, i.e., $R\sim\mathcal{U}(\mathrm{SO}(3))$. Then, with probability at least $1-\frac{2}{\pi}\nu$, the ball $\mathcal{B}^{\mathcal{C}}_{\nu}(R\cdot q)$ is convex.*

Proof. Throughout the proof, with a slight abuse of notation, R denotes both a rotation and the corresponding rigid-body transformation, by embedding R in $\mathrm{SE}(3)$ with zero translation. Let $\pi_{\mathrm{SO}(3)}:\mathrm{SE}(3)\to\mathrm{SO}(3)$ be the projection from

SE(3) to its orientation subgroup SO(3). We first notice that sampling R uniformly from SO(3) is equally likely to sampling $R' = R \cdot \pi_{\mathrm{SO}(3)}(q^{-1})$. We also notice that the ball's $\mathcal{B}^{\mathcal{C}}_{\nu}$ convexity does not depend on the translation (which is Euclidean) but only on the orientation. Hence, it suffices to prove that the ball $\mathcal{B}^{\mathcal{C}}_{\nu}(R)$ is convex.

Let $\pi_\phi, \pi_\psi : \mathrm{SO}(3) \to [-\pi, \pi)$ be the projection of the orientation into the roll and yaw, respectively. Note that when both angles have $\pi_\phi(R), \pi_\psi(R) \in (-\pi+\nu, \pi-\nu)$, by definition of d_c, for any two points $x, y \in \mathcal{B}^{\mathcal{C}}_{\nu}(R)$, the pseudo-metric d_c coincides with the Euclidean metric on $\mathbb{R}^6$. Thus, the ball $\mathcal{B}^{\mathcal{C}}_{\nu}(R)$ coincides with the Euclidean ball, and as such, it is convex.

Using the volume element defined in Section 2.1, we can bound the probability for either $\pi_\phi(R), \pi_\psi(R) \notin (-\pi+\nu, \pi-\nu)$:

$$
\begin{aligned}
&\mathbb{P}[\text{either } \pi_\phi(R), \pi_\psi(R) \notin (-\pi+\nu, \pi-\nu)] \\
&\quad \leq \mathbb{P}[\pi_\phi(R) \notin (-\pi+\nu, \pi-\nu)] + \mathbb{P}[\pi_\psi(R) \notin (-\pi+\nu, \pi-\nu)] \\
&\quad = \mathbb{P}[\pi_\phi(R) \in [-\pi, -\pi+\nu]] + \mathbb{P}[\pi_\phi(R) \in [\pi-\nu, \pi)] \\
&\qquad + \mathbb{P}[\pi_\psi(R) \in [-\pi, -\pi+\nu]] + \mathbb{P}[\pi_\psi(R) \in [\pi-\nu, \pi)] \\
&\quad = 4 \cdot \int_{-\pi}^{-\pi+\nu} \int_{-\pi/2}^{\pi/2} \int_{-\pi}^{\pi} \frac{1}{8\pi^2} \cos\theta \, d\phi \, d\theta \, d\psi = 4 \cdot \frac{1}{8\pi^2} \cdot 2 \cdot 2\pi \cdot \nu = \frac{2}{\pi} \nu .
\end{aligned}
$$

Notice that due to symmetry, the four probability computations amount to the same triple integral. Hence, the probability that the ball $\mathcal{B}^{\mathcal{C}}_{\nu}(R)$ is Euclidean (and thus convex) is at least $1 - \frac{2}{\pi}\nu$. □

Next, we prove that if `CMG-RRT` has reached a certain ball, with probability greater than zero, it will reach the next consecutive ball.

Lemma 4. *Suppose that `CMG-RRT` has reached $\mathcal{B}^{\mathcal{C}}_{\nu/5}(q_i)$, that is, T contains a vertex q'_i such that $q'_i \in \mathcal{B}^{\mathcal{C}}_{\nu/5}(q_i)$. With high probability `CMG-RRT` will reach $\mathcal{B}^{\mathcal{C}}_{\nu/5}(q_{i+1})$.*

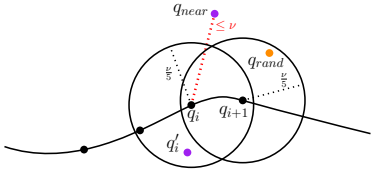

Fig. 5: Illustration of Lemma 4. Note that the balls are not necessarily convex. However, with probability p_{rotate}, we will choose a random rotation such that, in each transition step, the relevant balls are convex.

Proof. Suppose that q_{rand} is drawn such that $q_{\mathrm{rand}} \in \mathcal{B}^{\mathcal{C}}_{\nu/5}(q_{i+1})$. Denote by q_{near} the nearest neighbor of q_{rand} among the `CMG-RRT` vertices. See Fig. 5 for an illustration. Then, from the definition of q_{near}, it follows that $d_c(q_{\mathrm{near}}, q_{\mathrm{rand}}) \leq d_c(q'_i, q_{\mathrm{rand}})$.

From $d_c(q_{\mathrm{near}}, q_{\mathrm{rand}}) \leq d_c(q'_i, q_{\mathrm{rand}})$ and the triangle inequality, we have:

$$
\begin{aligned}
d_c(q_{\mathrm{near}}, q_i) &\leq d_c(q_{\mathrm{near}}, q_{\mathrm{rand}}) + d_c(q_{\mathrm{rand}}, q_i) \\
&\leq d_c(q'_i, q_{\mathrm{rand}}) + d_c(q_{\mathrm{rand}}, q_i).
\end{aligned}
$$

From the triangle inequality, we have that

$$d_c(q_{\text{rand}},q_i) \leq d_c(q_{\text{rand}},q_{i+1}) + d_c(q_{i+1},q_i),$$

$$d_c(q_{\text{rand}},q'_i) \leq d_c(q_{\text{rand}},q_i) + d_c(q_i,q'_i)).$$

Therefore, and since $d_c(\cdot,\cdot)$ is symmetric:

$$d_c(q_{\text{near}},q_i) \leq d_c(q'_i,q_i) + 2\cdot(d_c(q_{i+1},q_{\text{rand}}) + d_c(q_{i+1},q_i)) \leq 5\frac{\nu}{5} = \nu.$$

Note that by applying a random rotation R, we split and reform the ball $\mathcal{B}^{\mathcal{C}}_{\nu}(q_i)$. From Lemma 3, the probability that the random rotation chosen will cause $\mathcal{B}^{\mathcal{C}}_{\nu}(q_i)$ to be convex is p_{rotate}.

If $\mathcal{B}^{\mathcal{C}}_{\nu}(q_i)$ is convex then since $q_{\text{near}} \in \mathcal{B}^{\mathcal{C}}_{\nu}(q_i) \subseteq \mathcal{F}_\delta$ and since $q_{\text{rand}} \in \mathcal{B}^{\mathcal{C}}_{\nu}(q_i) \subseteq \mathcal{F}_\delta$ we obtain $\overline{q_{\text{near}}q_{\text{rand}}} \subseteq \mathcal{B}^{\mathcal{C}}_{\nu}(q_i) \subseteq \mathcal{F}_\delta$.

Also, the distance between q_{near} and q_{rand} is at most η since: $d_c(q_{\text{rand}},q_{\text{near}}) \leq d_c(q_{\text{rand}},q'_i) \leq d_c(q'_i,q_i) + d_c(q_i,q_{i+1}) + d_c(q_{i+1},q_{\text{rand}}) \leq 3\cdot\frac{\nu}{5} < \nu \leq \eta$. The fact that $d_c(q_{\text{near}},q_{\text{rand}}) \leq \eta$, means that $q_{\text{new}} = q_{\text{rand}}$.

Finally, we bound the probability to sample q_{rand} such that it lies in $\mathcal{B}^{\mathcal{C}}_{\nu/5}(q_{i+1})$. After a finite number of refinement steps, the diagonal D of a box will be at most $\nu/5$. When $D \leq \nu/5$, and assuming that $\mathcal{B}^{\mathcal{C}}_{\nu}(q_i)$ is convex, then each such ball will contain at least one box $b \in \mathbb{B}$. Since we sample uniformly from the boxes in the box set $\mathbb{B}$, the probability p_{sample} to sample a point within $\mathcal{B}^{\mathcal{C}}_{\nu/5}(q_{i+1})$ given that $\mathcal{B}^{\mathcal{C}}_{\nu}(q_i)$ is convex is at least $|b|/|\mathbb{B}|$, where $|b|$ is the Lebesgue measure[6] of a box $b \in \mathbb{B}$ and $|\mathbb{B}|$ is the Lebesgue measure of the box set $\mathbb{B}$.

Thus, the probability p that the straight line segment between a random sample q_{rand} and its nearest neighbor q_{near} in T lies entirely in $\mathcal{F}_\delta$ is $p = p_{\text{rotate}} \cdot p_{sample} > 0$. □

We now prove our main theorem.

Theorem 1. *The probability that* `CMG-RRT` *fails to reach q_{goal} from q_{init} after k iterations is at most ae^{-bk}, for some constants $a,b \in \mathbb{R}_{>0}$.*

Proof. Assume that $\mathcal{B}^{\mathcal{C}}_{\nu/5}(q_i)$ already contains a `CMG-RRT` vertex. From Lemma 4, with probability $p > 0$ in the next iteration a `CMG-RRT` vertex will be added to $\mathcal{B}^{\mathcal{C}}_{\nu/5}(q_{i+1})$.

Reaching q_{goal} from q_{init} with `CMG-RRT` requires repeating this step m times, transitioning from q_i to q_{i+1} for $0 \leq i < m$. With probability $p_{\text{rotate}} \geq p$ the mth ball $\mathcal{B}^{\mathcal{C}}_{\nu/5}(q_{\text{goal}})$ is convex, and then any line segment from a point in the ball to its center q_{goal} is in $\mathcal{F}_\delta$. Therefore, this process can be described as k Bernoulli trials with success probability p, where we would like to bound the probability to obtain $m+1$ successful outcomes (m for reaching $\mathcal{B}^{\mathcal{C}}_{\nu/5}(q_{\text{goal}})$ from $\mathcal{B}^{\mathcal{C}}_{\nu/5}(q_{\text{start}})$ and another successful outcome in choosing a random rotation such that $\mathcal{B}^{\mathcal{C}}_{\nu/5}(q_{\text{goal}})$ is convex and thus any straight line of points in it is in $\mathcal{F}_\delta$, as described before). By defining

[6] Note that we use the Lebesgue measure here since we sample points in $\mathbb{R}^6$.

success to be $m+1$ successful outcomes, we obtain an upper bound on the probability of failure, as the process may end after less than $m+1$ successful outcomes.

As in the analysis of [33], we can bound the probability of failure. That is, the probability that the process does not reach state $(m+1)$, after k steps. Let X_k denote the number of successes in k trials, then

$$\begin{aligned}
\Pr[X_k < (m+1)] &= \sum_{i=0}^{m} \binom{k}{i} p^i (1-p)^{k-i} \\
&\leq \sum_{i=0}^{m} \binom{k}{m} p^i (1-p)^{k-i} \leq \binom{k}{m} \sum_{i=0}^{m} (1-p)^k \\
&\leq \binom{k}{m} \sum_{i=0}^{m} (e^{-p})^k = \binom{k}{m} (m+1) e^{-pk} \\
&= \frac{\prod_{i=k-m-1}^{k} i}{m!} (m+1) e^{-pk} \leq \frac{1}{m!} k^{m+1} (m+1) e^{-pk},
\end{aligned}$$

where the transitions rely on (i) $m \ll k$, (ii) $p < \frac{1}{2}$, and (iii) $(1-p) \leq e^{-p}$.

As p,m are fixed and independent of k, the expression $\frac{1}{m!} k^{m+1} (m+1) e^{-pk}$ decays to zero exponentially with k. Therefore, `CMG-RRT` is probabilistically complete. □

5 Experiments and Results

In this section, we present comparisons of `CMG-RRT` with state-of-the-art benchmarks and results. Specifically, we refer to a subset of 16 instances from the benchmark presented in [61] which are challenging because they require a non-trivial combination of translation and rotation. In addition, we include the Elk puzzle, mentioned in [45,70] which is notably difficult.

5.1 Implementation Details

Our open-source software is written in Python and is available online (see project page, footnote 2). The code and all evaluations were run on a Linux machine with an Intel Core i7-12700 CPU. The tree exploration is parallelized among 16 cores. All models are scaled beforehand such that they fit inside the three-dimensional unit sphere $\mathbb{S}^2 \subset \mathbb{R}^3$.

Collision Detection To perform collision detection and to find contact points, we do the following. At the beginning of the algorithm, we sample 10,000 points on the boundary of the robot M_2. Then, for a configuration $q \in \mathrm{SE}(3)$, we transform those sampled points by q and evaluate the SDF of M_1. That SDF is implemented as a 240^3 grid. The SDF construction takes 7.4 seconds on average and at most 12.9 seconds for all tested models. We evaluate grid points in parallel using OpenMP [9], and perform distance queries with CGAL's Axis Aligned Bounding Box (AABB) tree [3].

Table 1: Results on 16 rotational assemblies from [61] under *puzzle* and *other*.

Algorithm	Success (%) ↑	AST (min) ↓
PRM* [27]	0	-
BIT* [17]	0	-
BKPIECE [57]	6.9	9.6
RRTConnect [36]	12.5	5.2
BK-RRT [73]	75.0	10.9
Tian et al. [61]	88.0	8.7
Zhang et al.† [70]	100.0	37.6
TR-RRT [45]	100.0	61.9
`CMG-RRT`	**100.0**	**3.0**

† Results reported from the original paper on *puzzles* only, running on a HTCondor cluster. This method, by design, can be applied only for puzzles.

Points that have an absolute signed distance less than a threshold $\delta = 0.005$ are considered contact points. Points with signed distance less than $-\delta$ are considered in penetration, and we report a collision.

Identifying too many contact points can artificially eliminate all the robot's degrees of freedom, particularly when nearby points induce nearly identical gradients. Therefore, after identifying all potential contact points, we cluster them by proximity of their gradients using the K-means algorithm [1,48] and achieve representative contact points for calculating of the tangent or retract direction.

5.2 Evaluation Against Baseline Methods

We compare our results against four representative sampling-based methods: PRM* [27], BIT* [17], BKPIECE [57], and RRTConnect [36], which are all implemented in the OMPL library [58], as well as four methods dedicated to tight assemblies and puzzles: Tian et al [61], BK-RRT [73], Zhang et al. [70], and the original TR-RRT [45]. We evaluated our method on the dataset provided by Tian et al. [61] (see Figure 2). Specifically, we deal with their *rotational assemblies*, which are problems that require a combination of simultaneous translation and rotation to disassemble the parts. This dataset is comprised of overall 24 instances, divided into three categories (of eight instances each): `screws`, `puzzles`, and `others`. Since our method deals with cases where there is only a discrete number of contact points at all times, we have not evaluated it on the screws category, which have contact surfaces (i.e., the set of all contact points in the workspace is a 2-manifold). Finally, we have tested on the Elk puzzle, which was first presented as an unsolved benchmark in [70], and later partially solved in [45], but which, to the best of our knowledge, has never been completely solved automatically. Our method, `CMG-RRT`, solves this puzzle with 100% success rate and with an average time of 156 minutes.

We ran each method on each instance 10 times, with a timeout of 4 hours. If, before that timeout, a valid path of motion was found, then we consider this a success.

In Table 1, we show the success rates and average success time (AST) of the methods. Running time includes the SDF grid construction. `CMG-RRT`'s success rate outperforms all but the original TR-RRT, consistently solving all evaluated instances, and runs faster than all tested methods, when considering succefull instances.

6 Discussion

This work targets a fundamental challenge of sampling-based planners in tight assemblies: although feasible motions often concentrate near contact, uniform sampling in SE(3) spends most effort in regions that do not contribute to progress. `CMG-RRT` addresses this mismatch by adaptively restricting sampling to a shrinking set of boxes that are provably near contact according to an SDF-based oracle, while preserving probabilistic completeness. Empirically, this bias toward the critical manifold substantially improves performance on challenging rotational assembly benchmarks, solving instances that prior methods have not, including the Elk puzzle.

Fig. 6: Execution of a `CMG-RRT` planned trajectory for the `az` puzzle using two UR5e robotic arms. Arms trajectories are based on a separate continuous IK derived from the exact `CMG-RRT` trajectory for the free-flying objects.

Figure 6 provides a complementary, qualitative validation of the trajectories produced by `CMG-RRT`. While `CMG-RRT` plans for free-flying rigid bodies in SE(3), we executed one planned solution using two UR5e robotic arms by converting the free-flying pose trajectory to continuous joint-space motion using a separate IK-based pipeline [46]. The resulting motion illustrates that the planned path is smooth and physically plausible and that the contact-rich portion of the motion can be tracked without requiring manual modification of the underlying trajectory. We emphasize that this experiment is intended as a qualitative demonstration of executability rather than a full manipulation-planning evaluation.

`CMG-RRT` has a couple of noticeable limitations. First, our current implementation assumes a grid-based SDF with finite-difference gradients; while this yields predictable query times, the approximation quality depends on grid resolution and scaling. Second, our method is designed for rigid-body assemblies with point contacts; extending it to settings with sustained surface contact (e.g., screw-like motions) may require additional progress.

A natural next step is to extend `CMG-RRT` from two-part to multi-part assemblies, where progress requires simultaneous and mutually consistent contacts among three or more parts. We expect the subdivision-and-prune approach to be particularly valuable in this setting, as coordinated contacts implicitly restrict feasible motion to a small, structured subset of the full configuration space, allowing sampling to focus on regions where progress is possible despite the increased dimensionality.

References

1. Ahmed, M., Seraj, R., Islam, S.M.S.: The K-Means algorithm: A comprehensive survey and performance evaluation. Electronics **9**(8), 1295 (2020)
2. Aiger, D., Kaplan, H., Sharir, M.: Reporting neighbors in high-dimensional euclidean space. SIAM Journal on Computing **43**(4), 1363–1395 (2014)
3. Alliez, P., Tayeb, S., Wormser, C.: 2D and 3D fast intersection and distance computation. In: CGAL User and Reference Manual. CGAL Editorial Board, 6.1.1 edn. (2026), `https://doc.cgal.org/6.1.1/Manual/packages.html#PkgAABBTree`
4. Bender, J., Duriez, C., Jaillet, F., Zachmann, G.: Continuous collision detection between points and signed distance fields. In: Workshop on virtual reality interaction and physical simulation. vol. 8 (2014)
5. Berenson, D., Srinivasa, S., Ferguson, D., Kuffner, J.: Manipulation planning on constraint manifolds. In: IEEE International Conference on Robotics and Automation (ICRA). pp. 625–632 (2009)
6. Bertiche, H., Madadi, M., Escalera, S.: Neural implicit surfaces for efficient and accurate collisions in physically based simulations. arXiv preprint arXiv:2110.01614 (2021)
7. Bilevich, M.M., Halperin, D.: A note on the time complexity of using subdivision methods for the approximation of fibers. arXiv preprint arXiv:2503.01626 (2025)
8. Chang, B., Cha, D., Ihm, I.: Computing local signed distance fields for large polygonal models. In: Computer Graphics Forum. vol. 27, pp. 799–806. Wiley Online Library (2008)
9. Dagum, L., Menon, R.: Openmp: an industry standard api for shared-memory programming. IEEE computational science and engineering **5**(1), 46–55 (1998)
10. Davies, T., Nowrouzezahrai, D., Jacobson, A.: On the effectiveness of eeight-encoded neural implicit 3D shapes. arXiv preprint arXiv:2009.09808 (2020)
11. Diebel, J., et al.: Representing attitude: Euler angles, unit quaternions, and rotation vectors. Matrix **58**(15-16), 1–35 (2006)
12. Dong, W., Wang, Q., Wang, X., Zha, H.: PSDF fusion: Probabilistic signed distance function for on-the-fly 3D data fusion and scene reconstruction. In: Proceedings of the European conference on computer vision (ECCV). pp. 701–717 (2018)
13. Driess, D., Ha, J.S., Toussaint, M., Tedrake, R.: Learning models as functionals of signed-distance fields for manipulation planning. In: Conference on robot learning. pp. 245–255. PMLR (2022)
14. Ecker, C., Kolev, B.: The Haar Measure in Solid Mechanics. arXiv preprint arXiv:2410.03371 (2024)
15. Friedman, Z., Paul, R.: Compliant manipulation planning. IEEE Transactions on Robotics and Automation **12**(4), 554–564 (1996)
16. Frisken, S.F., Perry, R.N., Rockwood, A.P., Jones, T.R.: Adaptively sampled distance fields: A general representation of shape for computer graphics. In: Proceedings of SIGGRAPH. pp. 249–254 (2000)
17. Gammell, J.D., Barfoot, T.D., Srinivasa, S.S.: Batch informed trees (bit*): Informed asymptotically optimal anytime search. The International Journal of Robotics Research **39**(5), 543–567 (2020)
18. Hashim, H.A.: Special orthogonal group SO(3), Euler angles, angle-axis, Rodriguez vector and unit-quaternion: Overview, mapping and challenges. arXiv preprint arXiv:1909.06669 (2019)
19. Hemingway, E.G., O'Reilly, O.M.: Perspectives on Euler angle singularities, gimbal lock, and the orthogonality of applied forces and applied moments. Multibody system dynamics **44**(1), 31–56 (2018)

20. Howes, N.R.: Modern Analysis and Topology. Springer Science & Business Media (1995)
21. Hsu, D., Jiang, T., Reif, J., Sun, Z.: The bridge test for sampling narrow passages with probabilistic roadmap planners. In: IEEE International Conference on Robotics and Automation (ICRA). pp. 4420–4426 (2003)
22. Hsu, D., Latombe, J.C., Motwani, R.: Path planning in expansive configuration spaces. The International Journal of Robotics Research **18**(5), 495–512 (1999)
23. Jacquet, M., Harms, M., Alexis, K.: Neural NMPC through signed distance field encoding for collision avoidance. The International Journal of Robotics Research p. 02783649251401223 (2025)
24. Jaillet, L., Porta, J.: Path planning under kinematic constraints by Rapidly-Exploring Random Trees. IEEE Transactions on Robotics **29**(1), 105–117 (2013)
25. Jiménez, P., Thomas, F., Torras, C.: 3D collision detection: A survey. Computers & Graphics **25**(2), 269–285 (2001)
26. Jones, M.W., Baerentzen, J.A., Sramek, M.: 3d distance fields: A survey of techniques and applications. IEEE Transactions on visualization and Computer Graphics **12**(4), 581–599 (2006)
27. Karaman, S., Frazzoli, E.: Sampling-based algorithms for optimal motion planning. The international journal of robotics research **30**(7), 846–894 (2011)
28. Kavraki, L.E., Švestka, P., Latombe, J.C., Overmars, M.H.: Probabilistic Roadmaps for path planning in high-dimensional configuration spaces. IEEE Transactions on Robotics and Automation **12**(4), 566–580 (1996)
29. Kavraki, L.E., Švestka, P., Latombe, J.C., Overmars, M.H.: Analysis of probabilistic roadmaps for path planning. IEEE Transactions on Robotics and Automation **14**(1), 166–171 (1998)
30. Kingston, Z., Moll, M.: Exploring implicit configuration spaces. The International Journal of Robotics Research **38**(8), 998–1012 (2019)
31. Kingston, Z., Moll, M., Kavraki, L.E.: Sampling-based methods for motion planning with constraints. Annual Review of Control, Robotics, and Autonomous Systems **1**, 159–185 (2018)
32. Kleinbort, M., Salzman, O., Halperin, D.: Efficient high-quality motion planning by fast all-pairs r-nearest-neighbors. In: 2015 IEEE International Conference on Robotics and Automation (ICRA). pp. 2985–2990. IEEE (2015)
33. Kleinbort, M., Solovey, K., Littlefield, Z., Bekris, K.E., Halperin, D.: Probabilistic completeness of RRT for geometric and kinodynamic planning with forward propagation. IEEE Robotics and Automation Letters **4**(2), i–vii (2018)
34. Koschier, D., Deul, C., Bender, J.: Hierarchical HP-Adaptive signed distance fields. In: Symposium on Computer Animation. pp. 189–198 (2016)
35. Krayer, B., Müller, S.: Generating signed distance fields on the GPU with ray maps. The Visual Computer **35**(6), 961–971 (2019)
36. 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)
37. LaValle, S.: Rapidly-Exploring Random Trees: a new tool for path planning. Research Report 9811 (1998)
38. LaValle, S.M.: Planning Algorithms. Cambridge University Press (2006)
39. LaValle, S.M., Kuffner, J.J.: Rapidly-exploring random trees: Progress and prospects: Steven m. lavalle, iowa state university, a james j. kuffner, jr., university of tokyo, tokyo, japan. Algorithmic and computational robotics pp. 303–307 (2001)

40. Lee, Y., Chamzas, C., Kavraki, L.E.: Adaptive experience sampling for motion planning using the generator–critic framework. In: IEEE International Conference on Robotics and Automation (ICRA). pp. 7104–7110 (2022)
41. Lefebvre, T., Bruyninckx, H., Schutter, J.D.: Active compliant motion: A survey. Advanced Robotics **19**(5), 479–499 (2005)
42. Li, S., Dantam, N.T.: Sample-driven connectivity learning for motion planning in narrow passages. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA). pp. 5681–5687 (2023)
43. Li, Y., Zhang, Y., Razmjoo, A., Calinon, S.: Representing robot geometry as distance fields: Applications to whole-body manipulation. In: 2024 IEEE International Conference on Robotics and Automation (ICRA). pp. 15351–15357. IEEE (2024)
44. Linial, N.: Finite metric spaces–combinatorics, geometry and algorithms. arXiv preprint math/0304466 (2003)
45. Livnat, D., Bilevich, M.M., Halperin, D.: Tight Motion Planning by Riemannian Optimization for Sliding and Rolling with Finite Number of Contact Points. In: 2024 IEEE International Conference on Robotics and Automation (ICRA). pp. 14333–14340. IEEE (2024)
46. Livnat, D., Lavi, Y., Halperin, D.: A full-cycle assembly operation: From digital planning to trajectory execution using a robotic arm. In: 2025 IEEE International Conference on Robotics and Automation (ICRA). pp. 9184–9191. IEEE (2025)
47. Macklin, M., Erleben, K., Müller, M., Chentanez, N., Jeschke, S., Corse, Z.: Local optimization for robust signed distance field collision. Proceedings of the ACM on Computer Graphics and Interactive Techniques **3**(1), 1–17 (2020)
48. MacQueen, J.: Some methods for classification and analysis of multivariate observations. In: Le Cam, L.M., Neyman, J. (eds.) Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability. pp. 281–297. No. 14, Oakland, CA, USA (1967)
49. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 165–174 (2019)
50. Salzman, O., Hemmer, M., Halperin, D.: On the power of manifold samples in exploring configuration spaces and the dimensionality of narrow passages. IEEE Trans Autom. Sci. Eng. **12**(2), 529–538 (2015)
51. Schutter, J.D., Brussel, J.V.: Compliant robot motion i: A formalism for specifying compliant motion tasks. The International Journal of Robotics Research **7**(4), 3–17 (1988)
52. Searcoid, M.O.: Metric Spaces. Springer (2007)
53. Selig, J.M.: Geometrical Methods in Robotics. Springer Science & Business Media (2013)
54. Shim, J., Kang, C., Joo, K.: Diffusion-based signed distance fields for 3D shape generation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 20887–20897 (2023)
55. Sigg, C., Peikert, R., Gross, M.: Signed distance transform using graphics hardware. In: IEEE Visualization, 2003. VIS 2003. pp. 83–90. IEEE (2003)
56. Solovey, K., Janson, L., Schmerling, E., Frazzoli, E., Pavone, M.: Revisiting the asymptotic optimality of RRT. In: 2020 IEEE international conference on robotics and automation (ICRA). pp. 2189–2195. IEEE (2020)
57. Şucan, I.A., Kavraki, L.E.: Kinodynamic motion planning by interior-exterior cell exploration. In: Algorithmic foundation of robotics VIII: selected contributions

of the eight international workshop on the algorithmic foundations of robotics. pp. 449–464. Springer (2009)
58. Şucan, I.A., Moll, M., Kavraki, L.E.: The Open Motion Planning Library. IEEE Robotics & Automation Magazine **19**(4), 72–82 (December 2012). https://doi.org/10.1109/MRA.2012.2205651, https://ompl.kavrakilab.org
59. Sulzer, R., Marlet, R., Vallet, B., Landrieu, L.: A survey and benchmark of automatic surface reconstruction from point clouds. IEEE Transactions on Pattern Analysis and Machine Intelligence (2024)
60. Sun, Z., Hsu, D., Jiang, T., Reif, J.: Narrow passage sampling for probabilistic roadmap planners. In: IEEE International Conference on Robotics and Automation (ICRA). pp. 4423–4428 (2005)
61. Tian, Y., Xu, J., Li, Y., Luo, J., Sueda, S., Li, H., Willis, K.D., Matusik, W.: Assemble Them All: Physics-Based Planning for Generalizable Assembly By Disassembly. ACM Transactions on Graphics (TOG) **41**(6), 1–11 (2022)
62. Voss, C., Moll, M., Kavraki, L.E.: A general method for sampling on implicit manifolds. In: IEEE International Conference on Robotics and Automation (ICRA). pp. 2420–2426 (2017)
63. Xu, H., Barbič, J.: 6-DOF Haptic rendering using continuous collision detection between points and signed distance fields. IEEE transactions on haptics **10**(2), 151–161 (2016)
64. Xu, T.: Recent advances in Rapidly-Exploring Random Tree: A review. Heliyon **10**(11) (2024)
65. Yan, F., Liu, Y.S., Xiao, J.Z.: Path planning in complex 3d environments using a probabilistic roadmap method. International Journal of Automation and computing **10**(6), 525–533 (2013)
66. Yang, W., Jin, W.: ContactSDF: Signed Distance Functions as multi-contact models for dexterous manipulation. IEEE Robotics and Automation Letters (2025)
67. Yap, C.K.: Soft subdivision search in motion planning. In: Proceedings, 1st Workshop on Robotics Challenge and Vision (RCV 2013) (2013)
68. Yap, C.K.: Soft subdivision search in motion planning, ii: axiomatics. In: International Workshop on Frontiers in Algorithmics. pp. 7–22. Springer (2015)
69. Zhang, L., Kim, Y.J., Manocha, D.: A hybrid approach for complete motion planning. In: 2007 IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 7–14. IEEE (2007)
70. Zhang, X., Belfer, R., Kry, P.G., Vouga, E.: C-space tunnel discovery for puzzle path planning. ACM Trans. Graph. **39**(4) (Aug 2020)
71. Zhang, Z., Chiang, Y.J., Yap, C.: Theory and explicit design of a path planner for an se (3) robot. arXiv preprint arXiv:2407.05135 (2024)
72. Zhao, H.: A fast sweeping method for Eikonal equations. Mathematics of computation **74**(250), 603–627 (2005)
73. Zickler, S., Veloso, M.M.: Efficient physics-based planning: Sampling search via non-deterministic tactics and skills. In: AAMAS (1). pp. 27–33 (2009)