RMRRT: Riemannian Barrier Metric RRT for Inequality-Aware Steering on Equality Manifolds

***Abstract*—This paper presents a motion planning framework that unifies equality and inequality constraints within a single geometric formulation for sampling-based planning in high-dimensional robotic systems. In conventional sampling-based planners, equality constraints are typically enforced through projection, whereas inequality constraints are handled separately through binary validity checks such as collision testing, often leading to inefficient exploration. To address this limitation, we propose Riemannian Barrier Metric RRT (RMRRT), which constructs a unified local geometry for planning on equality-constrained manifolds. RMRRT first builds an ambient barrier metric from inequality-sensitive barrier terms and then induces a tangent-space metric via a G-orthogonal projection associated with the equality constraints. The resulting tangent-space metric is used consistently in both steering and nearest-neighbor selection, biasing exploration away from nearby inequality boundaries while preserving first-order equality consistency. Experimental results show that RMRRT achieves a 100% success rate across diverse constrained manipulation tasks in both simulation and real-world settings, while reducing planning time relative to representative constrained planning baselines. Ablation studies further demonstrate that the proposed metric improves exploration quality by reducing rejected samples and shortening path length. Experiment videos and source code are available at: https://rmrrt-anonymous.github.io**

***Index Terms*—Constrained motion planning, Sampling-based planning, Riemannian metric, Barrier function, Multi-arm manipulation.**

I. Introduction

SAMPLING-based motion planning becomes particularly challenging when robotic systems are subject to simultaneous equality and inequality constraints [1]–[3]. Equality constraints commonly arise from closed-chain kinematics [4]–[6], grasp maintenance [7]–[9], and cooperative multi-arm manipulation [10], [11], restricting feasible configurations to an implicit manifold embedded in the ambient configuration space [12]–[14]. Inequality constraints, such as obstacle avoidance [15], [16], self-collision [17]–[19], and joint limits [2], [20], further restrict admissible regions through state-dependent boundaries. Under such conditions, feasible configurations typically occupy only a small fraction of the ambient space, making exploration efficiency and steering quality central to planning performance.

In conventional sampling-based planners, equality constraints are typically enforced through projection onto the constraint manifold, whereas inequality constraints are handled separately through binary validity checks such as collision testing [12]–[14], [21]. This decoupled treatment often leads to inefficient sample rejection and unstable steering, particularly in highly constrained environments where admissible regions are narrow or fragmented. Consequently, integrating equality and inequality constraints into a unified exploration framework remains a fundamental challenge in sampling-based motion planning. In this work, we address this challenge by proposing a unified geometric formulation that directly incorporates inequality-induced boundary geometry into exploration on equality-constrained manifolds.

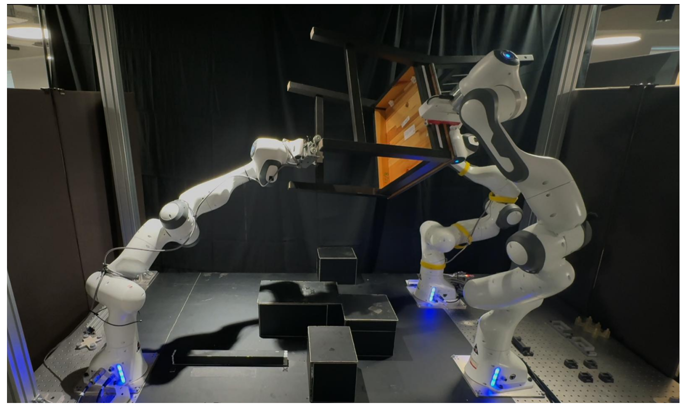

Fig. 1: Real-world demonstration of RMRRT with three manipulators cooperatively transporting a chair under simultaneous equality and inequality constraints for closed-chain grasp maintenance and collision avoidance.

A. Related Work

Sampling-based robot motion planning under constraints has been studied from multiple perspectives, largely distinguished by how equality and inequality constraints are represented and integrated into exploration. Prior work can be broadly grouped into three categories: (i) manifold-based approaches that plan directly on equality-constrained manifolds, (ii) metric-based approaches that encode inequality-induced geometry into distances or local metrics, and (iii) data-driven approaches that accelerate exploration using learned or precomputed priors.

1) Manifold-Based Motion Planning: Manifold-based motion planning explores directly on implicitly defined equality-constrained manifolds, aiming to preserve equality feasibility throughout planning [5], [12]–[14]. Atlas-based planners approximate the manifold using collections of local charts and transition across neighboring charts during expansion, improving robustness in highly curved regions [13], [22]. Tangent-space methods instead generate local motions in the Jacobian-defined tangent space and apply manifold-consistent correction, such as local projection, to obtain equality-feasible extensions [12], [14]. However, in most such approaches, tangent-space information is used primarily to maintain equality feasibility, while nearest-neighbor selection and steering are not consistently informed by nearby inequality boundaries. As a result, inequality constraints often affect exploration only through post-hoc rejection by external validity checks such as collision testing [23], which can lead to many unproductive extension attempts in narrow or cluttered environments [24], [25].

2) Metric-Based Motion Planning: Metric-based approaches incorporate inequality constraints into costs, distance measures, or state-dependent metrics to discourage motions toward constraint boundaries and bias exploration toward admissible regions. Classical examples include artificial potential fields and continuous obstacle-based cost formulations [26], [27], while more recent Riemannian approaches construct non-Euclidean metrics from curvature or Hessian information associated with inequality constraints [28], [29]. Related reactive frameworks such as RMPflow also use task-space Riemannian structures to generate local motion policies under multiple objectives [30]. While effective at encoding inequality-aware local geometry, these approaches are typically formulated in the ambient configuration space and do not directly provide equality-consistent motions on implicitly defined manifolds. In practice, additional projection or correction mechanisms are therefore required to recover manifold consistency during steering and tree expansion.

3) Data-Driven Motion Planning: Data-driven approaches accelerate constrained planning by exploiting learned or pre-computed structure in the feasible space. Representative directions include learning low-dimensional latent representations of equality-constrained manifolds [6], [31], learned samplers or policies that bias expansion toward promising regions, and offline roadmap-based methods that reduce online search burden through precomputed feasible-state connectivity [5], [32]. These methods can significantly improve efficiency when the learned prior or precomputed structure matches the target problem well. However, their performance may depend on training coverage, prior quality, or the suitability of the offline structure under changed environments or constraints.

4) Structural Gap Across Existing Approaches: Despite their respective strengths, existing approaches share a deeper structural limitation: they do not provide a single local geometry that is simultaneously equality-consistent and inequality-aware. In manifold-based planners, equality constraints define admissible first-order motion through tangent spaces and projection operators, but the local geometry used for exploration remains largely insensitive to nearby inequality boundaries. In metric-based planners, inequality constraints reshape local distances or energies, but the resulting geometry is typically defined in the ambient space and does not directly yield equality-consistent motions. Data-driven approaches can bias exploration toward favorable regions, but they generally do so through learned or precomputed priors rather than through an explicit geometric object coupling both constraint classes.

This separation reflects a genuine geometric mismatch rather than a mere modeling choice. A natural way to encode inequality constraints as a continuous local geometry is through a barrier-induced metric, which increases near active constraint boundaries. However, such geometry is naturally defined in the ambient configuration space, whereas admissible motions under equality constraints are restricted to a Jacobian-defined tangent space. Accordingly, simply combining a barrier metric with a standard Euclidean null-space projector is insufficient; the projector itself must be compatible with the metric defining the local geometry. This observation motivates the central construction of this work: a metric-compatible tangent projection and the induced tangent-space metric, used consistently across nearest-neighbor selection, local steering, and manifold correction to provide planner-wide geometric consistency.

B. Overview of Our Approach

As discussed above, existing sampling-based planners do not provide a single local geometry that is simultaneously equality-consistent and inequality-aware. This geometric mismatch requires a projection that is compatible with the metric defining the local geometry, which leads to the G-orthogonal projector and the induced tangent-space metric used in this work.

Motivated by this observation, we propose Riemannian Barrier Metric RRT (RMRRT), a sampling-based planner that redefines local planning geometry on equality-constrained manifolds. The key idea is to endow equality-consistent motions with a barrier-informed Riemannian structure that captures the local influence of inequality constraints such as obstacle proximity, self-collision, and joint limits. Concretely, we first construct a state-dependent ambient barrier metric from inequality-sensitive barrier terms, and then induce a metric on the tangent space through a barrier-compatible projection associated with the equality-constraint Jacobian. The resulting geometry not only preserves equality consistency, but also reshapes local exploration according to nearby inequality boundaries.

Based on this induced geometry, RMRRT performs local steering by selecting tangent directions that make first-order progress toward the target while penalizing motions toward nearby inequality boundaries. The same induced tangent-space metric is also used in nearest-neighbor selection, yielding a notion of local connectability that is consistent with equality-constrained expansion. As a result, inequality constraints influence the exploration mechanism itself rather than entering only through post-hoc validity checks.

The main contributions of this paper are summarized as follows:

- We introduce an ambient barrier metric and a G-orthogonal projection that together induce a tangent-space metric for sampling-based planning on equality-constrained manifolds. This construction captures both manifold structure and inequality-boundary sensitivity within a single local geometry.
- Building on this construction, we develop a bidirectional RRT-based planner for inequality-aware exploration on equality-constrained manifolds. The proposed planner incorporates the induced tangent-space metric into the expansion process so as to bias tree growth away from nearby inequality boundaries while preserving equality consistency.
- We use the same induced tangent-space metric consistently in both steering and nearest-neighbor selection, establishing planner-wide geometric consistency during exploration. This aligns node selection with the local directions favored

during expansion, in contrast to planners that rely on Euclidean proximity for selection and enforce equality consistency only during steering.

- We validate the proposed method through simulation and real-world experiments on constrained manipulation tasks involving multi-arm and humanoid systems. The results show improved planning efficiency and reliability over representative constrained planning baselines in challenging environments with coupled equality and inequality constraints [6], [13], [14], [21], [32].

The remainder of this paper is organized as follows. Section II introduces the necessary preliminaries, including equality-constrained manifolds and barrier-based metrics. Section III presents the proposed RMRRT methodology in detail. Section IV reports experimental results on a variety of constrained manipulation scenarios. Finally, Section V concludes the paper and discusses directions for future work.

II. Preliminaries

This section introduces the notation and geometric preliminaries used throughout the paper. Specifically, we formulate the motion planning problem under equality and inequality constraints, review equality-constrained manifolds and their tangent spaces, and summarize barrier-based representations of inequality constraints used in the metric construction of Section III.

Notation. Throughout the paper, n denotes the configuration-space dimension, m_{eq} the number of equality constraints, k the number of inequality constraints, and m_{proxy} the number of geometric proxy points used in the barrier metric. The scalar $\lambda > 0$ denotes the metric regularization parameter, and Lagrange multipliers are written as ν or μ. The equality-constraint Jacobian is denoted by $J(q) \in \mathbb{R}^{m_{\mathrm{eq}} \times n}$, while proxy-point Jacobians are written as $J_j(q) \in \mathbb{R}^{3 \times n}$.

A. Problem Formulation

We consider motion planning problems for robotic systems operating in a continuous configuration space $\mathcal{Q} \subset \mathbb{R}^n$, where a configuration is denoted by $q \in \mathcal{Q}$. The system is subject to both equality and inequality constraints, which jointly define the feasible configuration space. Equality constraints are expressed as

$$g(q) = 0, \tag{1}$$

where $g : \mathbb{R}^n \to \mathbb{R}^{m_{\mathrm{eq}}}$ is a differentiable constraint function. These constraints restrict feasible configurations to an implicitly defined manifold embedded in the ambient configuration space.

In addition, inequality constraints are given by

$$h_i(q) > 0, \quad i = 1, \ldots, k \tag{2}$$

where each $h_i : \mathbb{R}^n \to \mathbb{R}$ represents a state-dependent constraint. Together, the inequality constraints define admissible regions within the equality-constrained manifold. The overall feasible configuration set is therefore defined as

$$\mathcal{F} = \{q \in \mathcal{Q} \mid g(q) = 0,\ h_i(q) > 0,\ \forall i\}. \tag{3}$$

We assume that valid start and goal configurations lie within $\mathcal{F}$, and focus on planning collision-free motions that remain feasible throughout the planning process. In sampling-based motion planning, this requires generating tree expansions that preserve equality constraints while avoiding violations of inequality constraints.

B. Equality-Constrained Manifolds and Tangent Spaces

Equality constraints of the form $g(q) = 0$ define an implicit manifold embedded in the ambient configuration space $\mathcal{Q}$. Specifically, the equality-constrained manifold is given by

$$\mathcal{M} = \{\, q \in \mathcal{Q} \mid g(q) = 0 \,\}. \tag{4}$$

Let $J(q)$ denote the Jacobian of the constraint function,

$$J(q) = \frac{\partial g(q)}{\partial q} \in \mathbb{R}^{m_{\mathrm{eq}} \times n}. \tag{5}$$

If $J(q)$ has full row rank m_{eq} for all $q \in \mathcal{M}$, then $\mathcal{M}$ is a smooth manifold of dimension $n - m_{\mathrm{eq}}$[1].

At a configuration $q \in \mathcal{M}$, the tangent space of $\mathcal{M}$ is defined as

$$T_q\mathcal{M} = \{\, v \in \mathbb{R}^n \mid J(q)v = 0 \,\}. \tag{6}$$

This space characterizes the set of infinitesimal velocities that preserve the equality constraints to first order. In sampling-based motion planning, tangent directions are used to construct local expansions, and a retraction is applied when needed to keep the expanded state on the manifold $\mathcal{M}$. Moreover, when $J(q)$ has full row rank m_{eq}, the matrix $J(q)J(q)^\top$ is invertible, which allows defining the orthogonal projector onto $T_q\mathcal{M}$ as

$$Q(q) = I - J(q)^\top \big(J(q)J(q)^\top\big)^{-1} J(q), \tag{7}$$

where I denotes the $n \times n$ identity matrix. The projector $Q(q)$ satisfies $J(q)Q(q) = 0$ and maps arbitrary vectors in $\mathbb{R}^n$ onto the tangent space $T_q\mathcal{M}$. This tangent-space formulation provides a fundamental tool for constructing constraint-consistent steering and expansion operators, which will be exploited in the proposed approach.

C. Inequality Constraints and Log-Barrier Functions

As introduced in (2), the inequality constraints induce state-dependent boundaries that restrict admissible configurations within the equality-constrained manifold. In sampling-based planning, such constraints are typically handled via binary feasibility checks. To obtain a continuous representation of inequality constraints, we adopt a log-barrier formulation. Specifically, given a set of k inequality constraints $\{h_i(q)\}_{i=1}^k$, we define the logarithmic barrier function

$$\phi(q) = -\sum_{i=1}^{k} \log h_i(q). \tag{8}$$

This function grows rapidly as a configuration q approaches the boundary of the admissible region, thereby smoothly encoding proximity to inequality constraints. Moreover, the Hessian of the barrier function captures second-order information associated with the curvature of the constraint boundaries, providing a principled means to quantify how inequality constraints locally distort the geometry of the configuration space. This property makes log-barrier functions well-suited for defining geometry-aware structures that can be leveraged during planning, as will be described in Section III.

[1] This condition corresponds to the standard regularity assumption of the implicit function theorem.

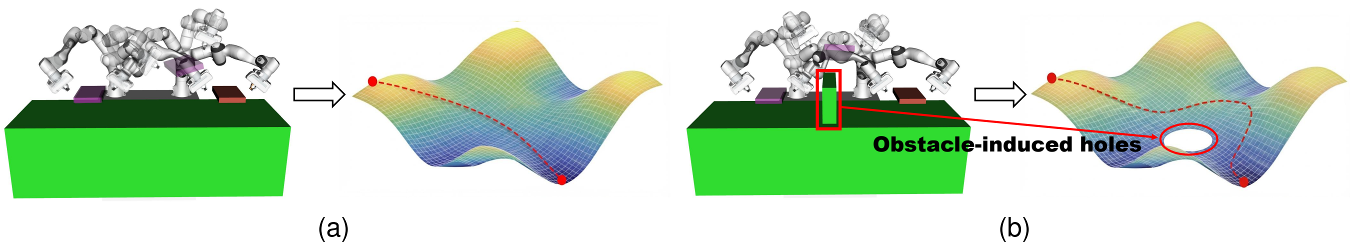

Fig. 2: Conceptual illustration of inequality-induced geometry on an equality-constrained manifold. (a) Equality-constrained manifold under the implicit assumption of full admissibility, as commonly treated in tangent-space and projection-based planners. (b) The same manifold after removing infeasible regions induced by inequality constraints, which form holes in the feasible subset. Although the equality constraint is unchanged, the admissible region exhibits nontrivial geometry that is not captured by conventional steering. RMRRT accounts for this geometry through a Riemannian metric defined on the tangent space.

III. Riemannian Barrier Metric RRT

In this section, we present Riemannian Barrier Metric RRT (RMRRT), a sampling-based planner that combines equality-manifold admissibility with inequality-aware local geometry. We first construct an ambient barrier metric and the induced tangent-space metric, and then show how the resulting geometry is used consistently in steering, nearest-neighbor selection, and manifold correction.

A. Proposed Framework

RMRRT builds on a bidirectional RRT-Connect backbone, but replaces purely Euclidean local expansion with a geometry that reflects both equality and inequality constraints. The motivation is illustrated in Fig. 2a–2b. In the absence of active inequality boundaries, tangent-space steering on the equality-constrained manifold provides a natural local motion model and is widely used in projection-based constrained planners [5], [12], [13]. However, when inequality constraints carve out infeasible regions on the same manifold, the feasible subset becomes geometrically nontrivial. In such settings, steering based only on tangent feasibility repeatedly proposes motions toward nearby invalid regions, leading to frequent rejection and unstable exploration.

The key idea of RMRRT is therefore to replace uniform tangent-space geometry with an inequality-aware local geometry on admissible motions. Equality constraints define the tangent space of feasible first-order motion, while inequality information reshapes the local exploration bias on that space through a barrier-informed metric. This yields local expansions that remain equality-consistent to first order while being biased away from nearby inequality boundaries. The next subsection formalizes this construction through an ambient barrier metric, a metric-compatible projector, and the induced tangent-space metric.

B. Riemannian Barrier Metric Construction

We first construct an ambient barrier metric that captures how proximity to inequality boundaries reshapes local geometry. In principle, this metric could be formed from the exact Hessians of the inequality barrier terms:

$$G(q) = I + \sum_{i=1}^{k} \alpha_i \nabla^2 \phi_i(q), \tag{9}$$

where ϕ_i is the log-barrier associated with the i-th inequality constraint and $\alpha_i > 0$ is a weighting factor. This ideal form directly reflects boundary-induced curvature, assigning larger metric cost to motions that reduce constraint margins.

In practice, however, exact barrier Hessians are expensive and numerically fragile in cluttered scenes, since signed-distance fields may become nonsmooth under closest-feature changes and contact switching. We therefore adopt a Gauss–Newton (GN) surrogate that retains the dominant positive semidefinite contribution of the barrier curvature. For a distance-based inequality constraint with signed distance $d(q)$ and barrier $\phi(q) = \phi(d(q))$, the chain rule gives

$$\nabla^2 \phi(q) = \phi''(d)\, \nabla d(q) \nabla d(q)^\top + \phi'(d)\, \nabla^2 d(q). \tag{10}$$

The second term depends on the curvature of the signed-distance field and may be ill-conditioned or undefined near feature switching. We therefore retain only the dominant positive semidefinite term and use the approximation

$$\nabla^2 \phi(q) \approx \phi''(d)\, \nabla d(q) \nabla d(q)^\top. \tag{11}$$

Now let $p(q) \in \mathbb{R}^3$ be a geometric proxy point with Jacobian $J_p(q) = \partial p(q)/\partial q \in \mathbb{R}^{3\times n}$, and let $n(q) = \partial d/\partial p \in \mathbb{R}^3$

Algorithm 1 $\textsc{BuildMetric}(q)$: Riemannian barrier metric

Require: Configuration q
Ensure: Ambient barrier metric $G(q)$, tangent projector $P(q)$, induced tangent-space metric $M(q)$

- **for** $j = 1, \ldots, m_{\text{proxy}}$ **do**
  - Compute proxy point $p_j(q)$ and signed distance $d_j(q)$
  - Compute point Jacobian $J_j(q) = \partial p_j(q)/\partial q$
- **end for**
- $G(q) \leftarrow \lambda I + \sum_{j=1}^{m_{\text{proxy}}} w_j(q)\, J_j^\top(q) J_j(q)$
- Compute equality Jacobian $J(q)$
- $A \leftarrow J(q)\, G^{-1}(q)\, J^\top(q)$
- $P(q) \leftarrow I - G^{-1}(q) J^\top(q)\, A^{-1}\, J(q)$
- $M(q) \leftarrow P^\top(q)\, G(q)\, P(q)$
- **return** $G(q), P(q), M(q)$

---

denote the local distance normal. Then

$$\nabla d(q) = n(q)^\top J_p(q), \tag{12}$$

which yields

$$\nabla^2 \phi(q) \approx \phi''(d)\, J_p(q)^\top \big(n(q)n(q)^\top\big) J_p(q). \tag{13}$$

Since $n(q)n(q)^\top \preceq I$, we obtain the upper bound

$$\nabla^2 \phi(q) \preceq \phi''(d)\, J_p(q)^\top J_p(q). \tag{14}$$

Motivated by this bound, we replace the exact directional term with a stable positive semidefinite surrogate and absorb the barrier curvature into a nonnegative distance-dependent activation weight. Summing over m_{proxy} proxy points leads to the practical ambient barrier metric

$$G(q) = \lambda I + \sum_{j=1}^{m_{\text{proxy}}} w_j(q)\, J_j(q)^\top J_j(q), \tag{15}$$

where $\lambda > 0$ is a regularizer, $d_j(q)$ is the signed distance associated with the j-th proxy point, $J_j(q) \in \mathbb{R}^{3\times n}$ is the corresponding point Jacobian, and $w_j(q)$ is a distance-based activation weight.

Definition 1 (Barrier Activation Weight)**.** For the j-th proxy point with signed distance $d_j(q)$ to the nearest obstacle, define

$$w_j(q) = \alpha_j\, \psi\big(\bar{d}_j(q)\big)\, \mathbf{1}[d_j(q) < \delta_{\text{act}}], \tag{16}$$

where

$$\bar{d}_j(q) = \max\big(d_j(q) - \delta_{\text{safe}},\, 0\big), \tag{17}$$

$\alpha_j > 0$ is a category-dependent scaling factor, $\delta_{\text{safe}} \geq 0$ is a safety margin, $\delta_{\text{act}} > 0$ is the activation radius, and

$$\psi(d) = \min\left(\frac{1}{d^2 + \epsilon_d^2},\, \psi_{\max}\right). \tag{18}$$

Here, $\epsilon_d > 0$ avoids singular behavior near contact, and $\psi_{\max}$ caps the metric contribution for numerical stability.

Remark 1 (Barrier-Kernel Interpretation)**.** The kernel $\psi(d)$ is a regularized and saturated approximation of the dominant $1/d^2$ behavior of the barrier curvature. As $d \to \infty$, $\psi(d) \to 0$ and the metric approaches λI, recovering near-isotropic exploration. As $d \to 0$, $\psi(d)$ saturates at $\psi_{\max}$, ensuring bounded metric contribution near contact. Thus, the metric smoothly transitions from near-isotropic exploration far from obstacles to anisotropic penalization near active inequality boundaries.

To make this geometry compatible with equality constraints, we restrict the ambient barrier metric to the tangent space of the equality-constrained manifold $\mathcal{M}$. Recall that

$$T_q\mathcal{M} = \{v \in \mathbb{R}^n \mid J(q)v = 0\}. \tag{19}$$

Given any ambient vector $w \in \mathbb{R}^n$, its G-orthogonal projection onto $T_q\mathcal{M}$ is defined by

$$\min_{v\in\mathbb{R}^n} \frac{1}{2}(v-w)^\top G(q)(v-w) \quad \text{s.t.} \quad J(q)v = 0. \tag{20}$$

The corresponding KKT conditions are

$$G(q)(v-w) + J(q)^\top \nu = 0, \qquad J(q)v = 0, \tag{21}$$

which yield

$$v = P(q)w, \tag{22}$$

with

$$P(q) = I - G^{-1}(q)J(q)^\top \big(J(q)G^{-1}(q)J(q)^\top\big)^{-1} J(q). \tag{23}$$

The induced tangent-space metric is then defined as

$$M(q) = P(q)^\top G(q) P(q). \tag{24}$$

This induced tangent-space metric is the metric used in all subsequent planner operations.

Lemma 1 (Tangent-Space Identity of the G-Orthogonal Projector)**.** *Let $P(q)$ be defined as in* (23) *and suppose $J(q)$ has full row rank. Then, for any $v \in T_q\mathcal{M}$,*

$$P(q)v = v. \tag{25}$$

Proof. Expanding (23) gives

$$P(q)v = v - G^{-1}(q)J(q)^\top \big(J(q)G^{-1}(q)J(q)^\top\big)^{-1} J(q)v. \tag{26}$$

Since $v \in \text{Null}(J(q))$, we have $J(q)v = 0$, and therefore $P(q)v = v$. □

Lemma 2 (Metric Equivalence on the Tangent Space)**.** *For any tangent-projected displacement $\tilde{r} = P(q)(q_{\text{rand}} - q) \in T_q\mathcal{M}$, the quadratic form induced by $M(q)$ satisfies*

$$\tilde{r}^\top M(q)\tilde{r} = \tilde{r}^\top G(q)\tilde{r}. \tag{27}$$

Proof. Since $M(q) = P(q)^\top G(q) P(q)$ and $\tilde{r} \in T_q\mathcal{M}$, Lemma 1 implies $P(q)\tilde{r} = \tilde{r}$. Therefore,

$$\tilde{r}^\top M(q)\tilde{r} = \tilde{r}^\top P(q)^\top G(q) P(q)\tilde{r} = \tilde{r}^\top G(q)\tilde{r}.$$

□

Remark 2 (Metric Evaluation after Projection)**.** Lemma 2 shows that once a displacement has been projected into $T_q\mathcal{M}$, metric-aware distance evaluation can be carried out using $G(q)$, without explicitly forming $P(q)^\top G(q) P(q)$ for each candidate.

Proposition 1 (Adaptive Metric Distortion Bound)**.** *For any feasible configuration $q \in \mathcal{F}$ with clearance*

$$\delta_{\min} = \min_j d_j(q) > 0, \tag{28}$$

the condition number of the induced tangent-space metric restricted to $T_q\mathcal{M}$ satisfies

$$\kappa\Big(M(q)\big|_{T_q\mathcal{M}}\Big) \leq 1 + \frac{m_{\text{proxy}}\, \alpha_{\max}\, \psi(\bar{\delta}_{\min})\, C_J}{\lambda}, \tag{29}$$

where

$$C_J = \max_j \|J_j(q)^\top J_j(q)\|_2, \qquad \alpha_{\max} = \max_j \alpha_j, \tag{30}$$

and

$$\bar{\delta}_{\min} = \max\big(\delta_{\min} - \delta_{\text{safe}},\, 0\big). \tag{31}$$

Proof. Let $v \in T_q\mathcal{M}$ with $\|v\| = 1$. By Lemma 1, $P(q)v = v$, and by Lemma 2, $v^\top M(q)v = v^\top G(q)v$. Using (15) and Definition 1,

$$v^\top G(q)v = \lambda + \sum_{j=1}^{m_{\text{proxy}}} w_j(q)\, \|J_j(q)v\|^2. \tag{32}$$

Since $d_j(q) \geq \delta_{\min}$ for all active proxy points and ψ is non-increasing,

$$\lambda \leq v^\top G(q)v \leq \lambda + m_{\text{proxy}}\, \alpha_{\max}\, \psi(\bar{\delta}_{\min})\, C_J. \tag{33}$$

The claimed bound follows immediately. □

The construction above yields an induced tangent-space metric $M(q)$ that preserves admissible first-order motions while encoding proximity to active inequality boundaries. In particular, Proposition 1 shows that the resulting metric distortion remains controlled by obstacle clearance and regularization. The next two subsections use this same metric for local steering and nearest-neighbor selection.

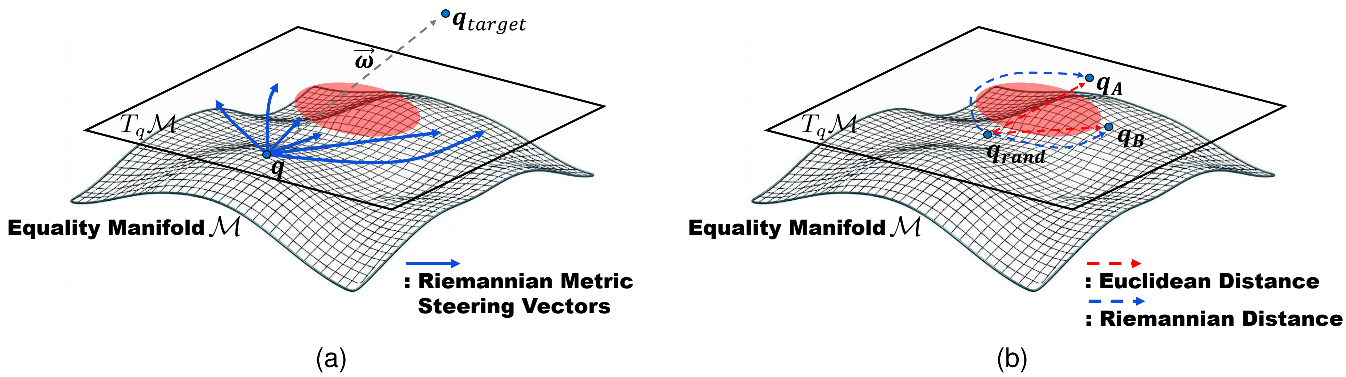

Fig. 3: Conceptual illustration of the metric-aware exploration mechanism in RMRRT. (a) Metric-induced tangent-space steering under the induced tangent-space metric. (b) Metric-aware nearest-neighbor selection under local inequality-dependent anisotropy.

Algorithm 2 METRICSTEER$(q, q_{\text{target}}, \eta)$: metric-induced tangent-space steering

Require: Feasible node $q \in \mathcal{M}$, target sample q_{target}, step size $\eta > 0$
Ensure: New configuration q_{new}

```
w ← q_target − q
(G, P, M) ← BUILDMETRIC(q)
w̃ ← P w
if ‖w̃‖ < ε then
    return q
end if
Solve M v_tmp = w̃
v_0 ← P v_tmp
s ← sqrt(v_0^T M v_0)
if s < ε then
    return q
end if
v̂ ← v_0 / s
q_new ← q + η v̂
return q_new
```

C. Metric-Induced Tangent-Space Steering

Given the induced tangent-space metric $M(q)$, we define a local steering rule that selects a tangent direction making first-order progress toward the target while minimizing local metric energy. As illustrated in Fig. 3a, the resulting direction remains equality-consistent to first order while being biased away from nearby active inequality boundaries.

For a tree node $q \in \mathcal{M}$ and target sample q_{target}, define the ambient displacement

$$w = q_{\text{target}} - q. \tag{34}$$

Its tangent-space component is

$$\tilde{w} = P(q)w \in T_q\mathcal{M}. \tag{35}$$

Proposition 2 (Well-Posedness of Metric-Induced Steering)**.** *Assume that $J(q)$ has full row rank and that $G(q) \succ 0$. Then the projector $P(q)$ is well-defined, the induced tangent-space metric $M(q)$ is positive definite on $T_q\mathcal{M}$, and for any nonzero $\tilde{w} \in T_q\mathcal{M}$, the optimization problem*

$$v^\star = \arg\min_{v \in T_q\mathcal{M}} \frac{1}{2} v^\top M(q) v \quad \textit{s.t.} \quad v^\top \tilde{w} = 1 \tag{36}$$

admits a unique minimizer.

Proof. Since $G(q) \succ 0$, the Schur complement

$$S(q) = J(q)G^{-1}(q)J(q)^\top \tag{37}$$

is positive definite under the full-row-rank assumption on $J(q)$, so $P(q)$ is well-defined. Moreover, for any nonzero $v \in T_q\mathcal{M}$, Lemma 1 implies $P(q)v = v$, hence

$$v^\top M(q) v = v^\top G(q) v > 0.$$

Therefore $M(q)$ is positive definite on $T_q\mathcal{M}$. The objective in (36) is strictly convex on the tangent space, and the affine constraint is nonempty whenever $\tilde{w} \neq 0$, so the minimizer is unique. □

Introducing a scalar Lagrange multiplier ν, define

$$L(v, \nu) = \frac{1}{2} v^\top M(q) v + \nu(1 - v^\top \tilde{w}). \tag{38}$$

The first-order optimality conditions yield

$$v^\star = \frac{M^{-1}(q)\tilde{w}}{\tilde{w}^\top M^{-1}(q)\tilde{w}}. \tag{39}$$

Equation (39) shows that the projected displacement is reweighted by the inverse induced tangent-space metric, amplifying inexpensive directions while suppressing directions toward nearby active boundaries.

For numerical robustness, we optionally reapply the tangent projector and define

$$v_0 = P(q)M^{-1}(q)\tilde{w}. \tag{40}$$

We then normalize this direction with respect to $M(q)$ and take a fixed-length step:

$$\hat{v} = \frac{v_0}{\sqrt{v_0^\top M(q) v_0}}, \qquad q_{\text{new}} = q + \eta\hat{v}, \tag{41}$$

where $\eta > 0$ is the step size. The complete procedure is summarized in Algorithm 2.

D. Metric-Aware Nearest-Neighbor Selection

We next use the same induced tangent-space metric $M(q)$ to define a geometry-consistent nearest-neighbor rule. Rather

Algorithm 3 $\textsc{MetricNearest}(q_{\mathrm{rand}}, \mathcal{V}, K)$: metric-aware nearest-neighbor selection

Require: Sample q_{rand}, tree nodes $\mathcal{V} \subset \mathcal{M}$, candidate size K
Ensure: Nearest node q_{near}

$\mathcal{V}_K \leftarrow \textsc{EuclideanKNN}(q_{\mathrm{rand}}, \mathcal{V}, K)$
$q_{\mathrm{near}} \leftarrow \emptyset,\ d_{\min} \leftarrow +\infty$
for all $q \in \mathcal{V}_K$ **do**
  $(G, P, M) \leftarrow \textsc{BuildMetric}(q)$
  $r \leftarrow q_{\mathrm{rand}} - q$
  $\tilde{r} \leftarrow P\, r$
  $d \leftarrow \sqrt{\tilde{r}^\top M\, \tilde{r}}$
  if $d < d_{\min}$ **then**
    $d_{\min} \leftarrow d, \quad q_{\mathrm{near}} \leftarrow q$
  end if
end for
return q_{near}

than ranking tree nodes solely by ambient Euclidean proximity, the planner evaluates the local cost of the tangent-projected displacement under $M(q)$.

Given a sample q_{rand} and a set of tree nodes $\mathcal{V} \subset \mathcal{M}$, define

$$d_M(q, q_{\mathrm{rand}}) = \left\| P(q)(q_{\mathrm{rand}} - q) \right\|_{M(q)}. \tag{42}$$

The nearest node is then selected as

$$q_{\mathrm{near}} = \arg\min_{q \in \mathcal{V}} \left\| P(q)(q_{\mathrm{rand}} - q) \right\|_{M(q)}. \tag{43}$$

Because the displacement is first projected onto $T_q\mathcal{M}$, the comparison depends only on equality-consistent first-order motions. Moreover, by Lemma 2, the quadratic form can be evaluated using the ambient barrier metric $G(q)$ once the projection has been applied.

A full metric-aware search over all nodes would be unnecessarily expensive. We therefore use Euclidean preselection followed by metric reranking. Let $\mathcal{V}_K \subset \mathcal{V}$ denote the set of the K Euclidean-nearest candidates to q_{rand}, and perform reranking only over this subset:

$$q_{\mathrm{near}} = \arg\min_{q \in \mathcal{V}_K} \left\| P(q)(q_{\mathrm{rand}} - q) \right\|_{M(q)}. \tag{44}$$

Remark 3 (Consistency of Euclidean Preselection)**.** The two-stage strategy in Algorithm 3 is consistent with metric-aware selection when K is sufficiently large, because Proposition 1 implies norm equivalence between the Euclidean norm and the induced tangent-space metric on the tangent space. Specifically, for any feasible node $q \in \mathcal{F}$ with clearance $\delta_{\min} > 0$ and any tangent-projected displacement $\tilde{r} = P(q)(q_{\mathrm{rand}} - q)$,

$$\sqrt{\lambda}\, \|\tilde{r}\| \ \le\ \|\tilde{r}\|_{M(q)} \ \le\ \sqrt{\Lambda}\, \|\tilde{r}\|, \tag{45}$$

where λ and Λ are the spectral bounds implied by Proposition 1.

Remark 4 (Asymmetry of the Metric-Aware Distance)**.** The distance $d_M(q, q_{\mathrm{rand}})$ is evaluated using the local metric at the tree node q, not at the sampled configuration q_{rand}. This is intentional: the metric at q encodes the local geometry of the region from which expansion occurs, ensuring consistency between nearest-neighbor selection and subsequent steering.

In our implementation, the candidate size is chosen adaptively as

$$K = \min\!\Big(K_{\max},\ \max(4,\ \lfloor 2\log_2(|\mathcal{V}| + 1) \rfloor)\Big), \tag{46}$$

Algorithm 4 Overall RMRRT Framework

Require: Feasible start/goal configurations $q_{\mathrm{start}}, q_{\mathrm{goal}} \in \mathcal{M}$, step size η, candidate size K, maximum iterations $N_{\max}$
Ensure: Feasible path σ or $\textsc{Failure}$

Initialize two trees $\mathcal{T}_a \leftarrow \{q_{\mathrm{start}}\}$, $\mathcal{T}_b \leftarrow \{q_{\mathrm{goal}}\}$
for $n = 1$ to $N_{\max}$ **do**
  $q_{\mathrm{rand}} \sim \textsc{UniformSample}(\mathbb{R}^n)$
  $q_{\mathrm{near}} \leftarrow \textsc{MetricNearest}(q_{\mathrm{rand}}, \mathcal{V}(\mathcal{T}_a), K)$
  $q_{\mathrm{new}} \leftarrow \textsc{MetricSteer}(q_{\mathrm{near}}, q_{\mathrm{rand}}, \eta)$
  $isStateValid \leftarrow \textsc{StateValid}(q_{\mathrm{new}})$
  $isEdgeValid \leftarrow \textsc{EdgeValid}(q_{\mathrm{near}}, q_{\mathrm{new}})$
  if not $isStateValid$ **or not** $isEdgeValid$ **then**
    continue
  end if
  Add q_{new} to $\mathcal{T}_a$ with parent q_{near}
  $q_{\mathrm{curr}} \leftarrow \textsc{MetricNearest}(q_{\mathrm{new}}, \mathcal{V}(\mathcal{T}_b), K)$
  repeat
    $q_{\mathrm{next}} \leftarrow \textsc{MetricSteer}(q_{\mathrm{curr}}, q_{\mathrm{new}}, \eta)$
    $isStateValid \leftarrow \textsc{StateValid}(q_{\mathrm{next}})$
    $isEdgeValid \leftarrow \textsc{EdgeValid}(q_{\mathrm{curr}}, q_{\mathrm{next}})$
    if not $isStateValid$ **or not** $isEdgeValid$ **then**
      break
    end if
    Add q_{next} to $\mathcal{T}_b$ with parent q_{curr}
    $q_{\mathrm{curr}} \leftarrow q_{\mathrm{next}}$
  until $\|q_{\mathrm{curr}} - q_{\mathrm{new}}\| < \epsilon_{\mathrm{conn}}$
  if $\|q_{\mathrm{curr}} - q_{\mathrm{new}}\| < \epsilon_{\mathrm{conn}}$ **then**
    $\sigma \leftarrow \textsc{ExtractPath}(\mathcal{T}_a, q_{\mathrm{new}}, \mathcal{T}_b, q_{\mathrm{curr}})$
    return σ
  end if
  Swap $(\mathcal{T}_a, \mathcal{T}_b)$
end for
return $\textsc{Failure}$

which serves as a practical candidate-pool heuristic prior to metric reranking. The complete nearest-neighbor procedure is summarized in Algorithm 3.

E. *Overall RMRRT Algorithm*

Algorithm 4 summarizes the overall planner. At each iteration, RMRRT samples a random configuration in the ambient space, selects a node from the active tree using the metric-aware nearest-neighbor rule, and performs a local expansion using the tangent-space steering rule. If the resulting state and edge are valid, the new node is inserted into the tree, after which the opposite tree is extended toward it in a connect phase.

When finite-step integration introduces noticeable equality drift, a lightweight metric-aware correction step is applied to return the configuration toward the equality-constrained manifold. This correction serves only as a local numerical recovery mechanism; admissible motion is already determined by the tangent-space geometry used in nearest-neighbor selection and steering.

Lemma 3 (Riemannian Newton Projection)**.** *Given a configuration q with equality violation $\|g(q)\| > \tau$, the iterative* *update*

$$q_{k+1} = q_k - G^{-1}(q_k)J(q_k)^\top \left(J(q_k)G^{-1}(q_k)J(q_k)^\top\right)^{-1} g(q_k) \tag{47}$$

converges locally to the equality-constrained manifold $\mathcal{M}$, provided $G(q)$ is continuous and $J(q)$ has full row rank in a neighborhood of $\mathcal{M}$.

Proof. Each iteration solves

$$\Delta q^* = \arg\min_{\Delta q} \frac{1}{2}\Delta q^\top G(q_k)\Delta q \quad \text{s.t.} \quad g(q_k)+J(q_k)\Delta q = 0. \tag{48}$$

The corresponding KKT system yields

$$\Delta q^* = -G^{-1}(q_k)J(q_k)^\top \left(J(q_k)G^{-1}(q_k)J(q_k)^\top\right)^{-1} g(q_k), \tag{49}$$

which recovers (47). The Schur complement remains positive definite under the same conditions used above, so the step is well-defined, and local convergence follows from standard equality-constrained Newton arguments. □

Remark 5 (Minimum-Norm Manifold Correction)**.** The correction step is the minimum-G-norm update that removes the first-order equality violation. Because $G(q)$ already reflects nearby inequality geometry, the correction remains consistent with the same local metric used in steering and nearest-neighbor selection.

Under standard regularity assumptions, the same geometric ingredients remain compatible with the incremental path-covering arguments used in probabilistic completeness analyses for constrained sampling-based planners.

Proposition 3 (Compatibility with Standard Completeness Arguments)**.** *Let $\mathcal{F} \subset \mathcal{M}$ be the feasible subset of the equality-constrained manifold. Suppose there exists a feasible path*

$$\sigma : [0,1] \to \mathcal{F} \tag{50}$$

connecting q_{start} to q_{goal}, and that this path has clearance $\delta > 0$ from all inequality boundaries. Assume further that the uniform sampling fraction satisfies $\rho_u > 0$, the equality Jacobian $J(q)$ has full row rank for all $q \in \mathcal{M}$, and the local metric satisfies the bounded-distortion conditions stated in Proposition 1 along $\sigma([0,1])$. Then the metric-aware nearest-neighbor selection, tangent-space steering, and correction steps of RMRRT remain compatible with the standard path-covering arguments used in probabilistic completeness analyses for constrained sampling-based planning.

Proof. The argument proceeds in three steps.

Step 1 (Uniform metric bounds along the feasible path). Because σ has positive clearance $\delta > 0$, the barrier activation weights remain uniformly bounded along $\sigma([0,1])$. Together with Proposition 1, this implies the existence of constants $0 < \lambda \le \Lambda < \infty$ such that, for every $q \in \sigma([0,1])$ and every $v \in \mathbb{R}^n$,

$$\lambda\|v\|^2 \le v^\top G(q)v \le \Lambda\|v\|^2. \tag{51}$$

Step 2 (Uniform norm equivalence on admissible motions). For every tangent vector $v \in T_q\mathcal{M}$, Lemma 2 gives

$$v^\top M(q)v = v^\top G(q)v.$$

Hence

$$\lambda\|v\|^2 \le v^\top M(q)v \le \Lambda\|v\|^2, \qquad \forall q \in \sigma([0,1]),\ \forall v \in T_q\mathcal{M}. \tag{52}$$

Step 3 (Existence of admissible local progress). Let $\hat{v} \in T_q\mathcal{M}$ denote the unit-$M(q)$-norm steering direction returned by Algorithm 2. By (52), its Euclidean norm is uniformly bounded above and below, so the steering rule generates a nondegenerate tangent displacement. Now cover the feasible path σ by a finite sequence of overlapping balls of radius $r = \eta/(2\sqrt{\Lambda})$. Because each ball has positive measure and the uniform sampling fraction satisfies $\rho_u > 0$, these neighborhoods are sampled infinitely often with probability one. Whenever a tree node lies in one ball and a subsequent sample falls in the next overlapping ball, the local norm-equivalence bounds imply that the metric-aware steering rule admits progress of controlled size toward the next neighborhood. Under the additional local-neighborhood coverage assumption that the Euclidean preselection stage retains access to a node in the relevant neighborhood, the subsequent metric-aware reranking and steering preserve this local progress property. Feasibility is enforced separately through state validity, edge validity, and the correction step. Since the feasible path has positive clearance, sufficiently small steps remain inside the admissible region and can be returned to $\mathcal{M}$ with bounded correction. Therefore, under the stated regularity, clearance, and local-neighborhood coverage assumptions, the metric-aware components of RMRRT remain compatible with the standard path-covering mechanism used in probabilistic completeness arguments for constrained sampling-based planning. □

Remark 6 (Barrier Metric as Geometric Bias)**.** The barrier metric is introduced to bias exploration geometrically, not to suppress it. Equality constraints define the admissible tangent structure, while inequality information reshapes the local geometry on that space through the induced tangent-space metric. The correction step is used only to remove finite-step drift from the manifold, whereas hard feasibility is enforced separately through state and edge validity checks.

F. Computational Complexity and Practical Overhead

We now examine the additional computational overhead introduced by the metric-aware operations. The dominant added cost arises from constructing the ambient barrier metric and reranking a small set of Euclidean-preselected candidates during nearest-neighbor search.

Proposition 4 (Per-Expansion Complexity)**.** *One metric-aware expansion of RMRRT, including local metric construction, tangent projection, and nearest-neighbor reranking over K Euclidean-preselected candidates, has computational complexity*

$$O\big(K\left(n^2 m_{\text{eq}} + n^2 m_{\text{proxy}}\right)\big), \tag{53}$$

where n is the configuration dimension, m_{eq} is the number of equality constraints, m_{proxy} is the number of proxy points, and K is the candidate size used in Algorithm 3.

Proof. We focus on the additional metric-aware cost beyond standard validity checking. For each candidate node, forming

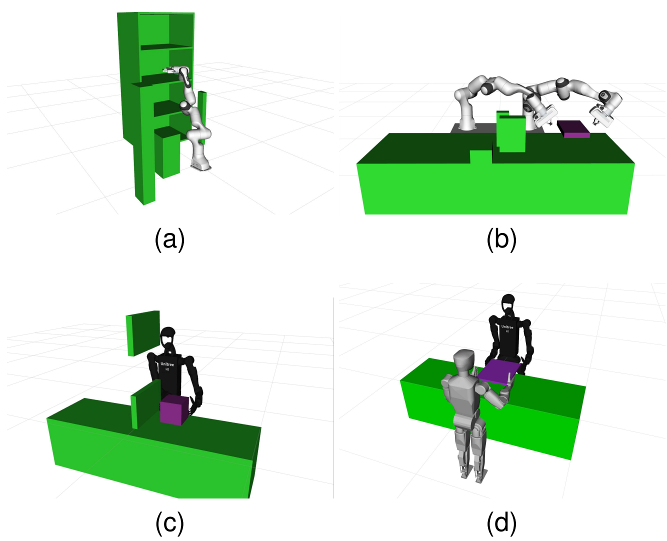

Fig. 4: Representative examples demonstrating the applicability of RMRRT to multiple robotic platforms. (a) single arm manipulator, (b) multi arm manipulators, (c) single humanoid, (d) dual humanoids.

the ambient barrier metric in (15) requires accumulating m_{proxy} rank-3 outer products, with cost $O(n^2 m_{\text{proxy}})$. Constructing the tangent projector requires solving for $G^{-1}(q)J(q)^\top$ and forming the associated Schur complement, whose dominant cost is $O(n^2 m_{\text{eq}})$ when $m_{\text{eq}} \ll n$. Since Algorithm 3 reranks K candidates, the total added cost is

$$O\big(K\left(n^2 m_{\text{eq}} + n^2 m_{\text{proxy}}\right)\big)\,.$$

Thus, when $m_{\text{eq}} \ll n$, the metric-aware overhead is dominated by local metric construction and scales linearly with the number of proxy points. □

Remark 7 (Low-Rank Structure and Efficient Inversion)**.** The ambient barrier metric

$$G(q) = \lambda I + \sum_{j=1}^{m_{\text{proxy}}} w_j(q)\, J_j(q)^\top J_j(q)$$

has a low-rank-plus-diagonal structure. Let

$$\bar{J} = \left[\sqrt{w_1}J_1; \ldots ; \sqrt{w_{m_{\text{proxy}}}}J_{m_{\text{proxy}}}\right] \in \mathbb{R}^{3m_{\text{proxy}}\times n}.$$

Then Woodbury's identity gives

$$G^{-1} = \frac{1}{\lambda}\left(I - \bar{J}^\top(\lambda I_{3m_{\text{proxy}}} + \bar{J}\bar{J}^\top)^{-1}\bar{J}\right), \tag{54}$$

which reduces the cost of applying G^{-1} when $m_{\text{proxy}} \ll n$.

Remark 8 (Practical Overhead Relative to Validity Checking)**.** In constrained sampling-based planning, collision and validity checking often dominate runtime. The metric-aware computations reuse the same signed distances and point Jacobians already computed or cached for validity evaluation, so their marginal overhead is largely algebraic. Moreover, because the barrier metric biases expansions away from nearby invalid regions, it can partially offset its own cost by reducing rejection-driven validity-check cycles in cluttered environments.

IV. Experiments

This section evaluates RMRRT under coupled equality and inequality constraints to assess its reliability, efficiency, and practical applicability. As shown in Fig. 4 and summarized in Table I, the proposed method is validated across multiple robotic platforms, including single-arm and multi-arm manipulators as well as single and dual humanoid systems. Simulation benchmarks are reported on representative constrained manipulation scenarios of increasing geometric and kinematic complexity. The experiments examine planning success rate, runtime, and computational overhead in broad benchmarks, isolate the contribution of the proposed metric-aware components through ablation, and further validate real-world performance in physical deployment. Additional videos are available on the project page: https://rmrrt-anonymous.github.io.

TABLE I: Problem scale of the benchmark scenarios.

Scenario	**Robot**	n	m_{eq}	m_{proxy}	Obs.
Scenario 1	Single arm	7	2	6	4
Scenario 2	Dual arm	14	6	11	5
Scenario 3	Dual arm	14	8	11	5
Scenario 4	Triple arm	21	12	16	7
Scenario 5	Humanoid H1-2	14	6	57	4

A. Simulation Benchmarks

For the simulation benchmarks, we constructed constrained manipulation scenarios spanning single-arm to triple-arm systems. All simulations were conducted on the same computing platform, equipped with an Intel Core i7-13700F processor (2.1 GHz), 32 GB of RAM, and an NVIDIA GeForce RTX 4060 GPU. For each scenario, 100 independent trials were performed using the same start state q_{start} and goal state q_{goal}, and a trial was counted as successful only if a valid solution was found within 300 s. All methods were evaluated under the same experimental conditions, while each planner was configured individually according to its own internal parameterization. Whenever possible, the parameters were tuned fairly based on quantities with analogous functional roles across methods, including expansion length, local extension granularity, and manifold-feasibility tolerance.

The baselines considered in this benchmark were restricted to methods that either operate natively on implicitly defined equality-constrained manifolds or provide a directly comparable constrained planning formulation. As discussed in Section I, purely metric-based approaches typically define local geometry in the ambient configuration space and do not directly provide equality-consistent motion generation on an implicitly defined equality-constrained manifold. Including such methods in the present benchmark would require additional design choices beyond their original formulation, which would make a direct apples-to-apples comparison less clear. For this reason, we focus on planners whose treatment of equality constraints is directly compatible with the problem class considered here. This choice is also aligned with the central motivation of the present work: RMRRT aims to unify the inequality-aware local geometry of metric-based approaches with the equality consistency of manifold-based planning within a single geometric object.

Table I summarizes the scale of each benchmark scenario in terms of configuration-space dimension, equality-constraint dimension, number of proxy points, and obstacle complexity,

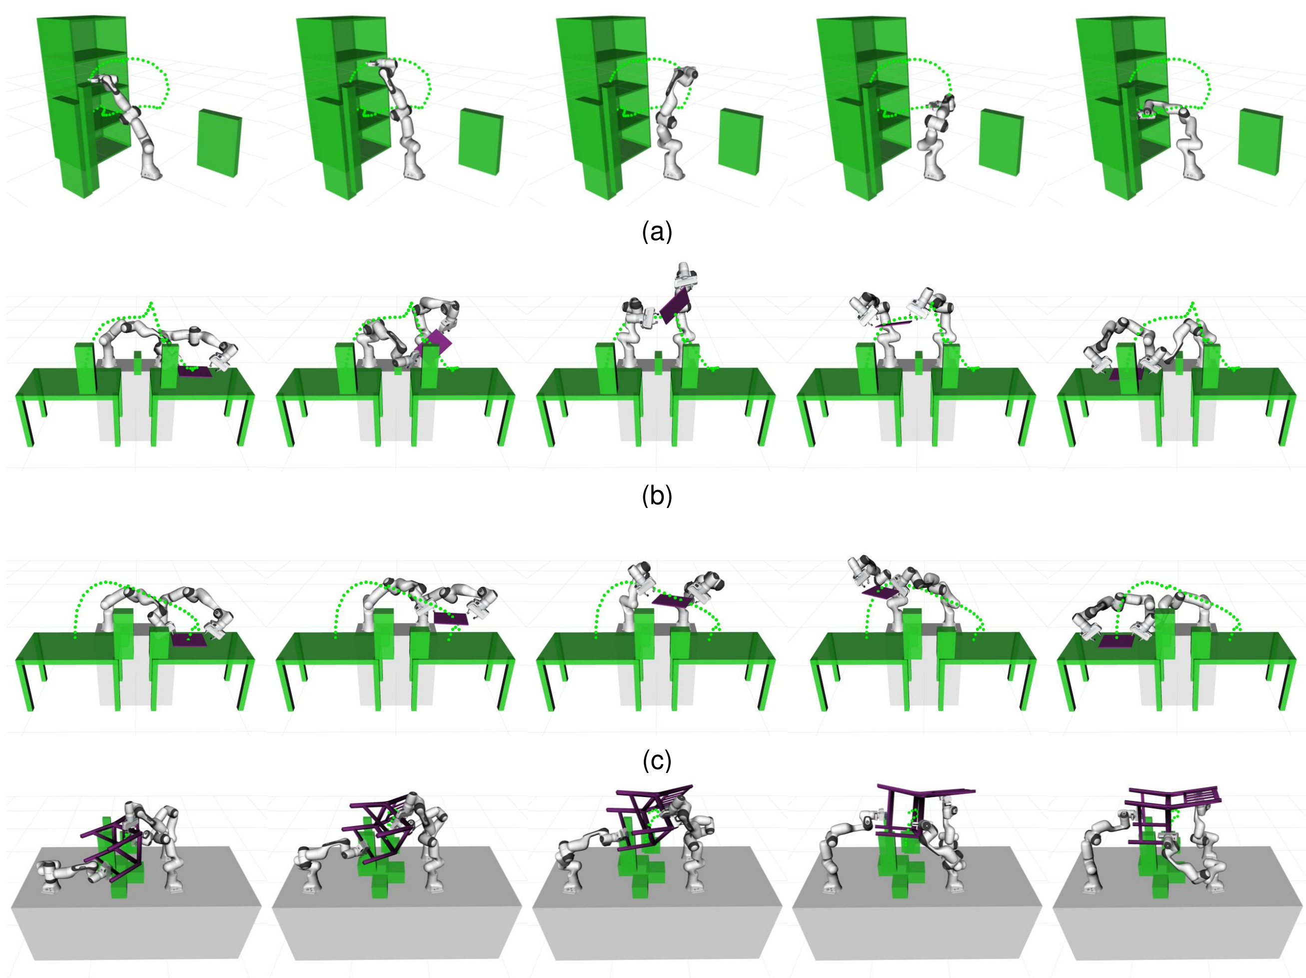
(a)

(b)

(c)

(d)

Fig. 5: Planned trajectories in four representative manipulation scenarios. The dotted curve indicates the end-effector (or object) trajectory, and snapshots depict key configurations along each path. (a) Single-arm cabinet reaching with fixed orientation constraint. (b) Dual-arm constrained transfer. (c) Dual-arm transfer with fixed orientation constraint. (d) Triple-arm constrained manipulation.

while Fig. 5 visualizes the representative tasks and planned trajectories. Table II reports planning performance across four simulation scenarios of increasing geometric and kinematic complexity. Across all scenarios, RMRRT achieves a 100% success rate while maintaining consistently low planning time and low inter-trial variance.

Manifold-based planners such as CBiRRT [21], Atlas-RRT [13], and Tangent Bundle RRT [13], [14] operate on the equality-constrained manifold by constructing local tangent spaces and correcting drift via projection. However, in these methods, tangent-space expansion is generated without explicitly reflecting nearby active inequality boundaries. As a result, many candidate motions may later be rejected by collision or validity checks, which can lead to repeated steering-rejection cycles in cluttered environments. This limitation can become more pronounced during nearest-neighbor selection. Because the nearest node is typically selected according to Euclidean distance in the ambient space, local connectability on the constraint manifold is not fully reflected, which can make the selected node less aligned with the locally admissible direction of progress.

By contrast, RMRRT uses the same induced tangent-space metric consistently in both nearest-neighbor selection and steering. This coupling prioritizes nodes that admit effective equality-consistent motion toward the target while suppressing locally unproductive motions at an earlier stage, thereby reducing rejection-driven exploration and improving both planning efficiency and path quality. This contrast is most pronounced in Scenario 4. In this setting, a 21-DoF triple-arm system must pass through a narrow region under 12 equality constraints. As shown in Table II, the manifold-based baselines exhibit substantially larger runtime variance and, in some cases, reduced success rates. These results are consistent with the interpretation that narrow admissible regions and dense active inequality boundaries increase the frequency of locally unproductive extension attempts. In contrast, RMRRT maintains a 100% success rate together with low runtime and low inter-trial variance in the same scenario. The per-iteration runtime

TABLE II: Planning performance across four experimental scenarios.

Planner	Scenario 1			Scenario 2		
	Time (s)	Succ. (%)	Path Len.	Time (s)	Succ. (%)	Path Len.
CBiRRT [21]	0.36 ± 0.54	100	8.54 ± 3.68	2.97 ± 4.48	100	9.07 ± 2.75
Atlas RRT [13]	1.04 ± 2.65	100	7.45 ± 3.08	3.87 ± 6.34	100	9.87 ± 3.22
Tangent Bundle RRT [14]	0.87 ± 3.34	100	9.12 ± 3.05	3.03 ± 4.32	100	9.61 ± 3.83
Precomputed Graph RRT [32]	0.42 ± 0.69	100	8.89 ± 3.79	2.72 ± 3.67	100	9.06 ± 2.62
LJCMP [6]	1.17 ± 3.30	100	12.45 ± 4.71	0.28 ± 0.65	100	11.97 ± 4.45
RMRRT (ours)	$\mathbf{0.09 \pm 0.14}$	100	8.05 ± 3.33	$\mathbf{0.23 \pm 0.31}$	100	10.92 ± 4.86
Planner	**Scenario 3**			**Scenario 4**		
	Time (s)	**Succ. (%)**	**Path Len.**	**Time (s)**	**Succ. (%)**	**Path Len.**
CBiRRT [21]	4.58 ± 9.20	100	9.53 ± 2.78	74.75 ± 118.59	90	10.67 ± 2.49
Atlas RRT [13]	5.71 ± 9.27	100	9.56 ± 3.23	42.11 ± 105.49	92	10.54 ± 3.64
Tangent Bundle RRT [14]	4.94 ± 10.11	100	9.80 ± 7.35	15.80 ± 57.35	87	11.01 ± 3.89
Precomputed Graph RRT [32]	2.07 ± 3.29	100	9.51 ± 2.80	13.34 ± 35.11	100	11.41 ± 35.11
LJCMP [6]	0.48 ± 1.56	100	12.31 ± 3.34	0.55 ± 1.69	100	16.44 ± 5.65
RMRRT (ours)	$\mathbf{0.25 \pm 0.26}$	100	12.63 ± 8.17	$\mathbf{0.31 \pm 0.39}$	100	11.09 ± 3.42

implications of this behavior are analyzed in the following subsection together with Table III.

We also compare against Precomputed Graph RRT [32] and LJCMP [6]. Precomputed Graph RRT exhibits relatively stable performance in simpler settings, but in Scenario 4 both its mean planning time and variance increase substantially. This suggests that a fixed offline roadmap structure may not fully capture local geometric variation in higher-dimensional constrained spaces. In contrast, LJCMP performs expansion on a learned latent embedding and achieves low average planning time across all scenarios; however, as shown in Table II, its path lengths are generally longer. One possible explanation is that the learned embedding is optimized primarily for feasibility-related structure rather than for preserving Euclidean distance in the original configuration space. Consequently, trajectories that are natural in the latent space may appear as detoured motions after decoding back to the configuration space.

RMRRT shares with LJCMP the fact that it does not explicitly minimize Euclidean path cost. However, there is an important structural distinction between the two approaches. The induced tangent-space metric of RMRRT is explicitly norm-equivalent to the Euclidean norm, and its distortion is quantitatively bounded by clearance and regularization. In other words, although the metric is locally reshaped according to nearby inequality geometry, its deviation from Euclidean distance remains mathematically controlled. This is not generally a property that can be directly assumed for a learned latent embedding. The path-length results in Table II are consistent with this distinction. In particular, RMRRT achieves substantially lower average planning time across all scenarios, while maintaining path lengths that remain competitive with the baselines. These results indicate that the proposed metric improves exploration efficiency without causing severe path-length degradation. These improvements are obtained without relying on learning or precomputed roadmaps, thereby providing a purely geometric and portable alternative for constrained motion planning.

TABLE III: Per-iteration runtime breakdown (ms) in Scenario 4.

Component	RMRRT	CBiRRT	LJCMP
Metric build	0.31	N/A	N/A
Nearest-neighbor	0.01	0.04	1.65
Steering	17.56	999.35	1.48
Collision / validity	0.95	4.83	1.72
Total per iteration	18.83	1016.66	4.86

B. Computational Overhead

To complement the theoretical complexity discussion in Section III, we report an empirical runtime breakdown of the main per-iteration components in Table III. Specifically, we separate the costs of metric construction, nearest-neighbor evaluation, steering, and collision/validity checking. This breakdown is particularly informative in higher-dimensional scenarios such as Scenario 4, where the ambient configuration dimension reaches $n = 21$. The purpose of this analysis is to clarify the practical overhead introduced by the proposed metric-aware components beyond standard validity checking. In particular, it quantifies whether the additional cost of constructing and using the ambient barrier metric remains manageable relative to the overall planning process, and whether improved exploration efficiency can offset part of this overhead in constrained environments.

As shown in Table III, the additional cost of metric construction in RMRRT remains modest on a per-iteration basis. The metric-build step requires only 0.31 ms per iteration, accounting for a small fraction of the total runtime. Nearest-neighbor evaluation also incurs only a minor overhead due to Euclidean preselection followed by metric reranking over a limited candidate set. The dominant per-iteration cost in RMRRT arises from steering, which includes metric-aware tangent projection and local extension under the induced tangent-space metric. Nevertheless, the total per-iteration runtime of RMRRT remains far smaller than that of CBiRRT and remains practically manageable in the benchmark setting considered here.

Compared with CBiRRT, the largest difference appears in the steering component. While both methods operate on

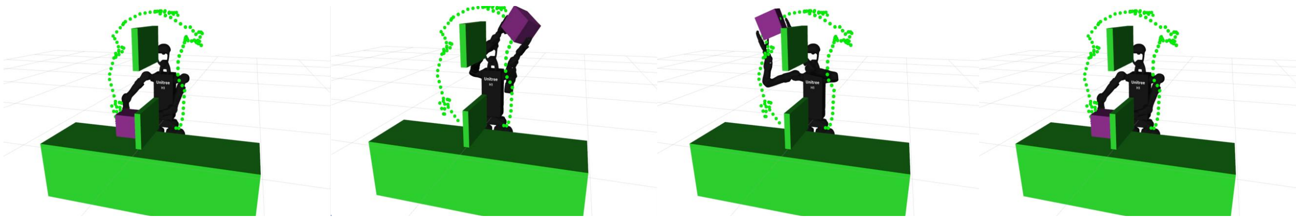

(a)

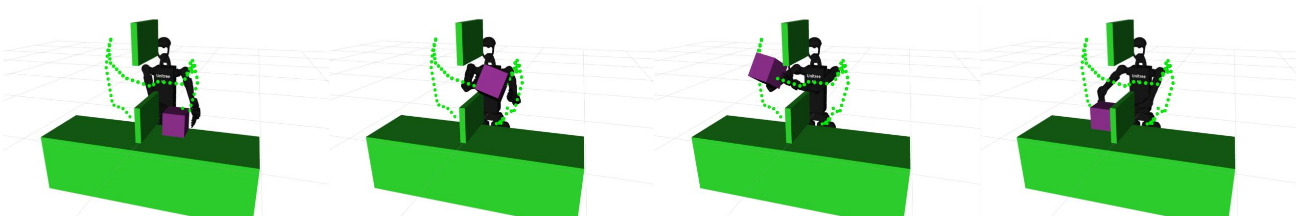

(b)

Fig. 6: Ablation study on the narrow-passage constrained transfer task in Scenario 5 using a Unitree H1-2 humanoid. (a) No barrier variant. (b) Full RMRRT. The no-barrier variant exhibits a longer detour around the restrictive region, whereas Full RMRRT generates a more direct trajectory through the passage.

TABLE IV: Ablation study on the contribution of the metric-aware components (Scenario 5).

	Outcome		**Exploration Quality**		**Cost Breakdown (ms/iter)**		
Method	Time (s)	Path Len. (m)	Valid Ext. Ratio	Path Len. Std	Check	Steer	Near
No barrier	21.26 ± 7.34	35.39 ± 9.84	0.333	9.84	12349	1193	5.4
Steering-only barrier	17.46 ± 7.42	28.13 ± 7.38	**0.349**	7.38	**8174**	1442	3.4
Nearest-only barrier	18.75 ± 8.97	25.66 ± 6.36	0.329	6.36	10987	1167	5.1
Full RMRRT	19.54 ± 11.08	$\mathbf{25.21 \pm 2.82}$	**0.349**	**2.82**	9243	1541	7.3

equality-constrained manifolds, CBiRRT exhibits a substantially higher steering cost, reflecting the heavier computational burden associated with repeated projection-based local expansion. By contrast, RMRRT performs steering through the induced tangent-space metric and reduces the total per-iteration runtime from $1016.66\,\text{ms}$ to $18.83\,\text{ms}$. Relative to LJCMP, RMRRT introduces additional geometric computation, particularly in metric construction and steering, resulting in a higher per-iteration cost ($18.83\,\text{ms}$ versus $4.86\,\text{ms}$). However, this additional computation does not translate into slower overall planning.

To better interpret this discrepancy, we also measured the number of planning iterations required to obtain a valid solution. Although LJCMP is cheaper per iteration, it requires substantially more iterations on average before finding a solution, whereas RMRRT converges in far fewer iterations. In our measurements, LJCMP required approximately 113 planning iterations on average, while RMRRT required only about 17. This suggests that the metric-aware expansion mechanism of RMRRT makes more effective progress per iteration, so that the added geometric cost is offset by a substantial reduction in the number of expansion cycles needed to reach a feasible solution. A similar contrast is observed with CBiRRT, whose very large per-iteration steering cost leads to both high iteration cost and large overall runtime.

Overall, these results support the practical viability of the proposed metric-aware design. Although RMRRT introduces additional geometric computation beyond standard constrained RRT expansion, the marginal overhead remains controlled and is accompanied by substantially improved search efficiency. In cluttered environments, this overhead can be offset by requiring fewer expansion attempts to reach a valid solution, yielding a favorable trade-off between richer local geometry and overall planning efficiency.

C. Ablation Study

To isolate the contribution of the proposed metric-aware components, we conduct an ablation study on the narrow-passage constrained transfer task in Scenario 5 using a Unitree H1-2 single-humanoid dual-arm system. In this setting, densely arranged obstacles induce strong inequality constraints that sharply restrict locally admissible directions and passable regions on the equality-constrained manifold. Scenario 5 is therefore particularly suitable for examining how each metric-aware component affects exploration under restrictive local geometry.

We compare four variants:

- **Full RMRRT**: metric-aware nearest-neighbor selection and metric-induced steering.
- **Steering-only barrier**: barrier-informed metric used only in steering, with Euclidean nearest-neighbor selection.
- **Nearest-only barrier**: barrier-informed metric used only in nearest-neighbor selection, with standard tangent-space steering.

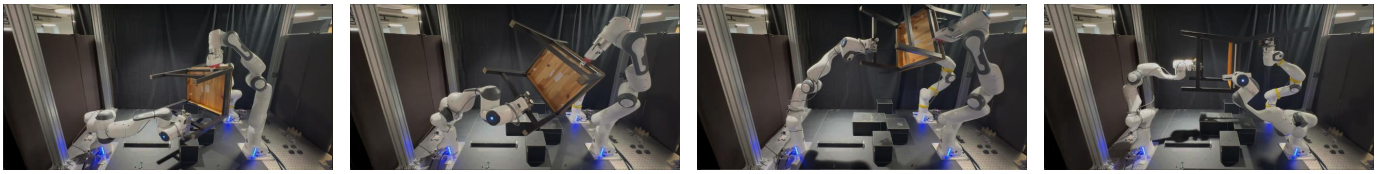

Fig. 7: Real-world validation of RMRRT on a triple-arm constrained manipulation task. The robots cooperatively manipulate a table under simultaneous equality and inequality constraints, including grasp-maintenance and collision-avoidance requirements.

- **No barrier**: neither nearest-neighbor selection nor steering uses the barrier-informed metric.

All other components, including the sampling strategy, validity checking, termination conditions, and major hyperparameters, are kept identical across the four variants. Table IV reports not only outcome metrics such as planning time and path length, but also exploration-quality metrics and per-iteration cost breakdowns. This decomposition allows us to distinguish whether the observed gains arise primarily from steering, from nearest-neighbor selection, or from their consistent coupling under a shared induced tangent-space metric.

As illustrated in Fig. 6, the barrier-informed metric does not alter equality feasibility itself. Rather, it changes the quality of local exploration under strong inequality-induced geometric constraints. In the narrow-passage setting, removing the barrier metric increases locally unproductive expansions, lengthens the resulting paths, and reduces extension efficiency. By contrast, Full RMRRT combines metric-aware nearest-neighbor selection and metric-induced steering under the same induced tangent-space metric, yielding more direct and more consistent traversal through the restrictive passage.

The results reveal a clear division of roles between the two metric-aware components. Applying the barrier-informed metric only in steering primarily improves local extension quality. In particular, the Steering-only barrier variant attains the highest valid extension ratio among the four variants and substantially reduces per-iteration validity-check-related cost relative to the No barrier baseline. This indicates that metric-induced steering suppresses locally unproductive motions toward nearby inequality boundaries before they propagate into repeated rejection cycles. This behavior is also consistent with the interpretation in Remark 8: the additional geometric computation introduced in steering can be partially offset by a reduction in rejection-driven validity-check overhead.

By contrast, applying the metric only in nearest-neighbor selection contributes more strongly to parent-node quality and path formation. The Nearest-only barrier variant significantly reduces the average path length relative to the No barrier baseline, suggesting that metric-aware parent selection favors nodes with better local connectability under the induced geometry. However, its valid extension ratio remains comparable to that of the baseline. This indicates that improved parent selection alone is insufficient to improve local extension quality when the extension rule itself remains inequality-unaware. In other words, even if the planner starts from a better parent node, the resulting extension can still proceed in a locally unproductive direction if steering does not reflect the same geometry.

Full RMRRT combines both components under a shared induced tangent-space metric. Among the four variants, it achieves the shortest average path length and, more importantly, the smallest path-length standard deviation. This reduction in inter-trial variability is especially important because it indicates that the proposed geometry improves not only mean trajectory quality but also the consistency of the generated trajectories across runs. That is, Full RMRRT does not merely produce shorter paths on average; it produces them more reliably. This provides direct experimental support for the central claim of the paper: planner-wide geometric consistency emerges when the same local geometry is used consistently in both nearest-neighbor selection and steering.

TABLE V: Quantitative real-world results.

Metric	Scenario 4
# Trials	100
Planning Time (s)	0.53 ± 0.12
Execution Success (%)	100

Taken together, these results show that the two metric-aware components play complementary but distinct roles. Metric-induced steering directly improves local extension quality, whereas metric-aware nearest-neighbor selection improves parent-node quality and path formation. Each component provides a meaningful benefit, but neither alone fully realizes the geometric intent of the proposed framework. Only when both are used consistently under the same induced tangent-space metric do parent selection and local expansion become aligned under a common geometric criterion, leading to the shortest and most consistent path quality in Scenario 5. These findings experimentally support the main thesis of this work: consistently applying the same local geometry in nearest-neighbor selection and steering improves exploration quality in constrained sampling-based planning.

D. Real-World Quantitative Evaluation

To assess practical applicability beyond simulation, we also report quantitative results from real-world execution on Scenario 4, which involves the highest-dimensional configuration space and the largest number of equality constraints among the benchmark scenarios. These experiments evaluate planning latency and execution reliability on hardware using the same equality- and inequality-constrained task family illustrated in Fig. 7. Table V summarizes the number of hardware trials, planning time, and execution success rate for Scenario 4.

In the real-world experiments, RMRRT was evaluated over 100 hardware trials on the triple-arm constrained manipulation task. The planner achieved an average planning time of $0.53 \pm 0.12\,\mathrm{s}$ and an execution success rate of 100%. These results indicate that the proposed planning framework remains reliable under real sensing, actuation, and model mismatch, while retaining planning latency at a practically usable level. Additional qualitative demonstrations and videos are available on the project page.

V. Conclusion

In this paper, we proposed RMRRT, a sampling-based motion planning framework for equality-constrained manifolds with inequality-aware local exploration. By constructing an ambient barrier metric and the induced tangent-space metric, RMRRT incorporates inequality-induced geometry directly into both steering and nearest-node selection, guiding tree expansion away from nearby infeasible regions without relying solely on collision-based sample rejection. The proposed framework was validated through a wide range of manipulation scenarios including multi-arm and humanoid systems, with detailed benchmark comparisons and ablation studies demonstrating improved planning efficiency and reliability over representative constrained planning baselines. Additional real-world experiments and demonstration videos are available on https://rmrrt-anonymous.github.io.

Future work will explore the integration of the proposed geometric construction with learning-based motion planning approaches. In particular, the induced tangent-space metric may serve as a complementary local planning component when combined with learned sampling strategies, heuristics, or latent-space exploration.

References

[1] C. Mu, Z. Jiang, J. Zhang, K. Wang, X. Xu, and J. Yi, "Efficient exploration in large state-action space through structured action space for learning multirobots motion planning," *IEEE Transactions on Control Systems Technology*, 2025.

[2] M. Stilman, "Task constrained motion planning in robot joint space," in *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2007, pp. 3074–3081.

[3] Z. Xian, P. Lertkultanon, and Q.-C. Pham, "Closed-chain manipulation of large objects by multi-arm robotic systems," *IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 1832–1839, 2017.

[4] R. Bordalba, L. Ros, and J. M. Porta, "A randomized kinodynamic planner for closed-chain robotic systems," *IEEE Transactions on Robotics*, vol. 37, no. 1, pp. 99–115, 2020.

[5] K. Jang, J. Baek, S. Park, and J. Park, "Motion planning for closed-chain constraints based on probabilistic roadmap with improved connectivity," *IEEE/ASME Transactions on Mechatronics*, vol. 27, no. 4, pp. 2035–2043, 2022.

[6] S. Park, S. Jeon, and J. Park, "A constrained motion planning method exploiting learned latent space for high-dimensional state and constraint spaces," *IEEE/ASME Transactions on Mechatronics*, vol. 29, no. 4, pp. 3001–3009, 2024.

[7] W. Liu, M. Ren, K. Song, M. Y. Wang, and Z. Xiong, "A planning framework for complex flipping manipulation of multiple mobile manipulators," *IEEE Robotics and Automation Letters*, 2025.

[8] M. Dogar, A. Spielberg, S. Baker, and D. Rus, "Multi-robot grasp planning for sequential assembly operations," *Autonomous Robots*, vol. 43, no. 3, pp. 649–664, 2019.

[9] S. Zimmermann, G. Hakimifard, M. Zamora, R. Poranne, and S. Coros, "A multi-level optimization framework for simultaneous grasping and motion planning," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2966–2972, 2020.

[10] K. Shi, J. Liu, J. Chang, and G. Hu, "Distributed quadratic programming for safe human multirobot interaction under task constraints," *IEEE/ASME Transactions on Mechatronics*, 2025.

[11] J. Alonso-Mora, R. Knepper, R. Siegwart, and D. Rus, "Local motion planning for collaborative multi-robot manipulation of deformable objects," in *2015 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2015, pp. 5495–5502.

[12] C. Suh, T. T. Um, B. Kim, H. Noh, M. Kim, and F. C. Park, "Tangent space rrt: A randomized planning algorithm on constraint manifolds," in *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 4968–4973.

[13] L. Jaillet and J. M. Porta, "Path planning under kinematic constraints by rapidly exploring manifolds," *IEEE Transactions on Robotics*, vol. 29, no. 1, pp. 105–117, 2012.

[14] B. Kim, T. T. Um, C. Suh, and F. C. Park, "Tangent bundle rrt: A randomized algorithm for constrained motion planning," *Robotica*, vol. 34, no. 1, pp. 202–225, 2016.

[15] J. J. Quiroz-Omaña and B. V. Adorno, "Whole-body control with (self) collision avoidance using vector field inequalities," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 4048–4053, 2019.

[16] D. Guo and Y. Zhang, "A new inequality-based obstacle-avoidance mvn scheme and its application to redundant robot manipulators," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 42, no. 6, pp. 1326–1340, 2012.

[17] R. Grassmann, L. Johannsmeier, and S. Haddadin, "Smooth point-to-point trajectory planning in $se(3)$ with self-collision and joint constraints avoidance," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1–9.

[18] P. Hu, A. Huang, J. Dolan, D. Held, and D. Ramanan, "Safe local motion planning with self-supervised freespace forecasting," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 12 732–12 741.

[19] K. Cai, L. Zhang, X. Su, K. Chen, C. Wang, S. Haddadin, A. Knoll, A. Ajoudani, and L. Figueredo, "Just-in-time informed trees: Manipulability-aware asymptotically optimized motion planning," *IEEE/ASME Transactions on Mechatronics*, 2025.

[20] W. Li, Z. Yi, R. Rao, D. Guo, Y. Pan, and K. Huang, "A quadratic programming solution to position-level repetitive motion planning and obstacle avoidance of joint-constrained robots," *IEEE Transactions On Industrial Electronics*, 2024.

[21] D. Berenson, S. S. Srinivasa, D. Ferguson, and J. J. Kuffner, "Manipulation planning on constraint manifolds," in *2009 IEEE international conference on robotics and automation*. IEEE, 2009, pp. 625–632.

[22] L. Jaillet and J. M. Porta, "Path planning with loop closure constraints using an atlas-based rrt," in *Robotics Research: The 15th International Symposium ISRR*. Springer, 2016, pp. 345–362.

[23] Z. Kingston, M. Moll, and L. E. Kavraki, "Sampling-based methods for motion planning with constraints," *Annual review of control, robotics, and autonomous systems*, vol. 1, no. 1, pp. 159–185, 2018.

[24] D. Hsu, T. Jiang, J. Reif, and Z. Sun, "The bridge test for sampling narrow passages with probabilistic roadmap planners," in *2003 IEEE international conference on robotics and automation (cat. no. 03CH37422)*, vol. 3. IEEE, 2003, pp. 4420–4426.

[25] J. Pan and D. Manocha, "Fast probabilistic collision checking for sampling-based motion planning using locality-sensitive hashing," *The International Journal of Robotics Research*, vol. 35, no. 12, pp. 1477–1496, 2016.

[26] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," *The international journal of robotics research*, vol. 5, no. 1, pp. 90–98, 1986.

[27] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, "Chomp: Gradient optimization techniques for efficient motion planning," in *2009 IEEE international conference on robotics and automation*. IEEE, 2009, pp. 489–494.

[28] A. Manjunath and Q. Nguyen, "Safe and robust motion planning for dynamic robotics via control barrier functions," in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 2122–2128.

[29] G. Yang, B. Vang, Z. Serlin, C. Belta, and R. Tron, "Sampling-based motion planning via control barrier functions," in *Proceedings of the 2019 3rd International Conference on Automation, Control and Robots*, 2019, pp. 22–29.

[30] C.-A. Cheng, M. Mukadam, J. Issac, S. Birchfield, D. Fox, B. Boots, and N. Ratliff, "Rmpflow: A geometric framework for generation of multitask motion policies," *IEEE Transactions on Automation Science and Engineering*, vol. 18, no. 3, pp. 968–987, 2021.

[31] J. Zhang, C. Bai, W. Pan, T. Liu, and J. Guo, "Local path optimization in the latent space using learned distance gradient," in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2025, pp. 15 940–15 946.

[32] I. A. Şucan and S. Chitta, "Motion planning with constraints using configuration space approximations," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 1904–1910.