

Implicit Null-space Manifold Generation for Redundant Robotic Systems

Taiki Ishigaki

Tokyo University of Science, Japan
taiki.ishigaki@rs.tus.ac.jp

Ko Ayusawa

National Institute of Advanced Industrial Science and Technology, Japan

Teresa Vidal-Calleja

University of Technology Sydney, Australia

Eiichi Yoshida

Tokyo University of Science, Japan

Abstract—Robotic systems with redundant degrees of freedom can achieve the same task outcome using multiple configurations, resulting in solution sets that form manifolds in the configuration space. Existing approaches typically exploit such redundancy locally through Jacobian-based techniques to compute individual solutions or trajectories. While effective for solution computation, these methods do not retain a representation of the geometry of the solution set itself. In this work, we adopt a representation-centric approach to estimate the geometric structure of the solution space. We consider solution manifolds induced by general task-defining maps and construct an implicit scalar field over the configuration space, whose zero-level set corresponds to the solution manifold. To this end, we generate samples in the neighborhood of the solution manifold using a Jacobian-guided exploration strategy, which efficiently captures its local and global structure. The resulting implicit representation is defined over the configuration space and naturally induces a continuous, distance field that encodes proximity to the solution manifold. Experiments on a planar three-link robot and a seven-degree-of-freedom Franka manipulator demonstrate the effectiveness of the proposed representation. Furthermore, the framework enables consistent modeling of solution spaces across families of tasks with continuous variation.

I. INTRODUCTION

Robotic systems often exhibit redundancy, where multiple configurations yield the same task outcome. Such redundancy naturally induces solution sets in the configuration space and plays a fundamental role in inverse kinematics [11, 12], redundancy resolution, and motion generation. These problems are typically addressed in a query-driven manner, where feasible configurations or trajectories are computed for individual task instances.

Existing methods, in both kinematic [10] and dynamic [5] settings, primarily exploit redundancy via local Jacobian-based techniques to compute solutions or motions. While effective for obtaining feasible instances, these approaches do not provide an explicit representation of the underlying solution space. As a result, the geometric structure of the null-space is discarded once a solution is obtained, limiting reuse across tasks and subsequent queries. Our approach is motivated by classical results on self-motion manifolds [2, 8], which establish that redundancy-induced solution sets form smooth manifolds in the configuration space.

In this work, we adopt a representation-centric perspective on redundancy, focusing on modeling the geometric structure




Fig. 1: Overview of the proposed framework. From a task-defining map $p = f(q)$, Jacobian-guided sampling produces samples in the neighborhood of the solution manifold in configuration space. An implicit representation is then constructed using a Gaussian process, providing a configuration-space model whose zero-level set represents the solution manifold.

of solution spaces rather than computing individual solutions. We consider tasks defined by a differentiable mapping

$$p = f(q), \quad (1)$$

where $q \in \mathbb{R}^n$ denotes configuration variables and $p \in \mathbb{R}^m$ denotes task variables. Importantly, the mapping f is not restricted to forward kinematics; it may represent any task-defining function for which function and Jacobian evaluations are available.

For a given target task value p^* , the set of configurations satisfying $f(q) = p^*$ forms an implicitly defined solution manifold,

$$\mathcal{M}(p^*) = \{q \in \mathbb{R}^n \mid f(q) = p^*\}. \quad (2)$$

It has been shown that, under redundancy, inverse kinematic solution sets form smooth self-motion manifolds in the configuration space [2]. Recent progress has also been made toward directly computing such manifolds as global solution traces, for example via ODE-based integration in the Jacobian null space together with mechanisms for handling sign ambiguities and searching disconnected components [4]. However, trace-based approaches primarily yield instance-specific point sets and do not produce a reusable configuration-space representation that supports geometric queries or downstream planning.

We propose a representation-centric framework for modeling these solution manifolds. Specifically, we construct a probabilistic implicit representation whose zero-level set corresponds to $\mathcal{M}(\mathbf{p}^*)$. The implicit function $\phi(\mathbf{q})$ is modeled as a Gaussian process (GP) and is constructed from configuration samples generated on and near the manifold. Sampling is guided by the Jacobian of the task-defining map, which encodes local null-space geometry, enabling informative sampling without exhaustive exploration of the configuration space or distance supervision. Once constructed, the resulting implicit representation can be reused for a fixed task instance to support proximity evaluation and motion generation, and it induces a continuous, distance field over the configuration space that reflects proximity to the solution manifold.

Although inverse kinematics provides a convenient testbed in our experiments, the proposed framework does not aim to solve inverse kinematics problems. Instead, it targets reusable representation of solution spaces induced by general task-defining maps, enabling proximity queries and task-aware sampling beyond inverse kinematics.

II. RELATED WORK

A. Inverse Kinematics and Redundancy Resolution

Inverse kinematics (IK) [11, 12] and redundancy resolution have been extensively studied as means of computing configurations or trajectories that satisfy task constraints. Classical approaches exploit the Jacobian null-space to locally resolve redundancy and generate feasible motions. These methods are inherently query-driven, focusing on computing feasible solutions for individual task instances.

B. Self-Motion Manifolds and Geometric Characterization

The geometric structure of solution sets induced by redundancy has been studied from differential geometric and topological perspectives. Burdick [2] characterized inverse kinematic solution sets of redundant manipulators as self-motion manifolds, analyzing their dimension and topology. More recently, self-motion manifolds have been explicitly computed as global solution traces via ODE-based integration of Jacobian null-space directions [4]. These approaches provide accurate manifold traces for specific task instances.

C. Learning-Based Inverse Models

Learning-based approaches have been proposed to approximate inverse mappings from task space to configuration space and to address multi-solution IK problems [1]. These methods typically rely on dense training data and focus on predicting configurations directly. Moreover, they do not explicitly exploit Jacobian-based geometric structure.

D. Implicit Representations and Distance Fields

Implicit surface representations and distance fields provide a framework for encoding geometry as scalar-valued functions over a space. Gaussian process implicit surfaces (GPIS) have been used to model shapes [13, 14, 9], and distance fields [6] without explicit parametrization. More recently, configuration




Fig. 2: Jacobian-guided sampling for traversing a solution manifold.

space distance fields (CDFs) have been proposed to encode task-relevant geometry directly in joint space [7]. These representations are primarily designed for distance queries and motion planning.

E. Position of This Work

In contrast to the above approaches, this work focuses on representing the geometry of solution manifolds induced by general task-defining maps. Rather than computing individual inverse solutions or learning direct inverse mappings, we construct a probabilistic implicit representation of the solution space itself. By exploiting Jacobian-guided local geometric information, the proposed method bridges classical self-motion manifold theory and implicit representations, enabling reusable modeling of null-space geometry beyond inverse kinematics.

III. METHOD

Given a differentiable task-defining map $\mathbf{p} = \mathbf{f}(\mathbf{q})$ with $\mathbf{q} \in \mathbb{R}^n$ and $\mathbf{p} \in \mathbb{R}^m$, we fix a target $\mathbf{p}^*$ and define the residual

$$\mathbf{r}(\mathbf{q}) := \mathbf{f}(\mathbf{q}) - \mathbf{p}^*. \quad (3)$$

The solution manifold is

$$\mathcal{M}(\mathbf{p}^*) = \{\mathbf{q} \in \mathbb{R}^n \mid \mathbf{r}(\mathbf{q}) = \mathbf{0}\}. \quad (4)$$

Although $\mathbf{p}$ may lie in $\text{SE}(3)$, we represent it as $\mathbf{p} \in \mathbb{R}^p$ to explicitly encode the dimensionality of the task, while the task error is computed in the corresponding tangent space (e.g., via $\log(H^{-1}H')$ for $\text{SE}(3)$). We assume access only to $\mathbf{r}(\mathbf{q})$ and its Jacobian

$$\mathbf{J}(\mathbf{q}) := \frac{\partial \mathbf{f}}{\partial \mathbf{q}} = \frac{\partial \mathbf{r}}{\partial \mathbf{q}}. \quad (5)$$

Since $\Delta \mathbf{p} = \mathbf{J} \Delta \mathbf{q}$, displacements in $\text{Null}(\mathbf{J})$ do not affect the task to first order and thus span the local tangent space of $\mathcal{M}(\mathbf{p}^*)$.

Our framework has two components: (i) Jacobian-guided sampling of points on/near $\mathcal{M}(\mathbf{p}^*)$, and (ii) GP-based implicit modeling to obtain a distance-like field in configuration space.

A. Jacobian-Guided Sampling

Fig. 2 illustrates the proposed Jacobian-guided sampling strategy for traversing the solution manifold $\mathcal{M}(\mathbf{p}^*)$ in configuration space. Starting from a configuration that satisfies the task constraint, sampling proceeds by alternately (1) advancing along directions tangent to the manifold and (2) restoring task consistency through local correction. The tangent directions are obtained from the null space of the Jacobian, which locally approximates the manifold geometry. Two variants of this traversal are shown in Fig. 2: a Newton-based method that tightly follows the manifold via frequent projection, and a zigzag method that relaxes projection accuracy and alternates across the manifold to accelerate exploration. Since joint angle limits can be straightforwardly represented in the configuration space, we consider a configuration space in which each joint angle is constrained to lie within the interval $[-\pi, \pi]$.

1) *Initial Point Search*: We first obtain an initial configuration $\mathbf{q}_0$ satisfying $\mathbf{r}(\mathbf{q}_0) = \mathbf{0}$ by iterating a Gauss–Newton update:

$$\Delta \mathbf{q} = \underset{\Delta \mathbf{q}}{\operatorname{argmin}} \|\mathbf{r}(\mathbf{q}) + \mathbf{J}(\mathbf{q})\Delta \mathbf{q}\|^2, \quad \mathbf{q} \leftarrow \mathbf{q} + \Delta \mathbf{q}, \quad (6)$$

until $\|\mathbf{r}(\mathbf{q})\| \leq \varepsilon$.

2) *Jacobian-Guided Manifold Traversal*: Given $\mathbf{q}_k$, we advance along a null-space direction and then restore constraint satisfaction:

$$\tilde{\mathbf{q}}_{k+1} = \mathbf{q}_k + \beta \mathbf{v}_k, \quad \mathbf{v}_k \in \operatorname{Null}(\mathbf{J}(\mathbf{q}_k)), \quad (7)$$

$$\mathbf{q}_{k+1} = \Pi(\tilde{\mathbf{q}}_{k+1}), \quad (8)$$

where $\Pi(\cdot)$ denotes an operator that maps a configuration to a neighborhood of the manifold. The vector $\mathbf{v}_k$ is a nontrivial solution satisfying $\mathbf{J}(\mathbf{q}_k)\mathbf{v}_k = \mathbf{0}$. In the following, we introduce two different realizations of this operator. The parameter β represents the sampling interval: larger values result in coarser sampling. However, for manifolds with strong nonlinearity, excessively large β may cause the tracing process to deviate from the manifold.

Joint angles are wrapped to the interval $[-\pi, \pi]$ to account for angular periodicity during exploration. Equation (7) is specialized for one-dimensional manifolds, where a single tangent direction is obtained and the termination condition for manifold exploration can be defined in a straightforward manner. Accordingly, in our implementation, the exploration on a given manifold is terminated when a previously sampled point is found within a β -neighborhood of the current sampling point. While this work focuses on the case of a single redundant degree of freedom, tangent directions of a manifold can in general be obtained from the null space of the Jacobian. By developing appropriate exploration strategies and termination conditions, the proposed framework can be extended to exploration on higher-dimensional manifolds.

3) *Newton-Based Traversal*: In the Newton-based variant, $\mathbf{v}_k$ is recomputed at each step from the current Jacobian, yielding stable manifold-following at the cost of more frequent direction/projection updates.

4) *Zigzag Traversal*: To reduce projection effort, we propose a zigzag strategy that relaxes strict projection and instead maintains successive samples that approximately lie on opposite sides of the manifold.

We define a correction direction toward the manifold by the minimum-norm linearized solution of $\mathbf{J}(\mathbf{q})\Delta \mathbf{q} = -\mathbf{r}(\mathbf{q})$:

$$\mathbf{d}(\mathbf{q}) = -\mathbf{J}(\mathbf{q})^\# \mathbf{r}(\mathbf{q}), \quad (9)$$

where $(\cdot)^\#$ is the pseudo-inverse. Starting from $\tilde{\mathbf{q}}_{k+1}$, we perform a small number of correction steps

$$\mathbf{q}_\ell = \mathbf{q}_{\ell-1} + \gamma \mathbf{d}(\mathbf{q}_{\ell-1}), \quad (10)$$

and terminate when the direction to the manifold flips:

$$\mathbf{d}(\mathbf{q}_{\ell-1}) \cdot \mathbf{d}(\mathbf{q}_\ell) < 0. \quad (11)$$

Equation (11) checks the relative orientation of the correction directions toward the manifold at the two consecutive iterations. A negative inner product means that the angle between these vectors is obtuse, i.e., the two directions point to opposite sides. This is a simple indicator that the manifold has been crossed and lies between $\mathbf{q}_{\ell-1}$ and $\mathbf{q}_\ell$ (up to the local linearization used to compute $\mathbf{d}$), so we stop without enforcing exact convergence. We set $\gamma > 1$ in (10) to intentionally overshoot and promote such crossings.

B. Probabilistic Implicit Function Representation

GP regression is a nonparametric method that has been widely used for nonlinear regression problems. It enables probabilistic estimation of unknown function values at arbitrary query points from a limited number of samples that may include noise and variability.

GPIS methods [13, 14, 9] employ GP to estimate an implicit surface from a set of noisy sampling points. Using GPIS, it is also possible to construct a distance field from the implicit surface for arbitrary query points [6].

In this study, we utilize manifold neighborhood samples generated by the method described in Section III-A and represent, using a shifted GPIS, a field that simultaneously encodes the manifold structure and a customizable notion of distance from the manifold. Unlike the standard GPIS method, which uses a constant zero-mean function, we shift the mean to one. This field can be interpreted as a smooth occupancy field, or equivalently, as an implicit surface on the 1-level set. While the surface can be modelled directly using this occupancy field, [6] proposed applying an inverse mapping to recover a distance field far away from the surface. This can be understood as reversing the effect of the kernel function, which essentially represents the distance between two points when only one surface point and one query point are involved.

Let us consider a GP with one mean:

$$\phi(\mathbf{q}) \sim \mathcal{GP}(1, k(\mathbf{q}, \mathbf{q}')), \quad (12)$$

where $k(\mathbf{q}, \mathbf{q}')$ is a covariance function. As the covariance function, we employ the squared exponential (SE) kernel

$$k(\mathbf{q}, \mathbf{q}') = \exp\left(-\frac{\|\mathbf{q} - \mathbf{q}'\|^2}{2\ell^2}\right). \quad (13)$$

TABLE I: Comparison of manifold sampling methods

Method	β	Sampling Time [ms]	Coverage Volume $V[\text{m}^3]$	Mean of Norm of Residual [m]	Sampling Size $n[-]$
Newton-based	1.5	1.3	0.0056	4.6×10^{-5}	11
	1.0	1.4	0.0076	2.4×10^{-5}	15
	0.5	2.1	0.0148	1.7×10^{-5}	29
	0.1	7.7	0.0734	4.0×10^{-6}	140
Zigzag	1.5	1.1	0.0052	9.9×10^{-4}	11
	1.0	1.3	0.0079	8.0×10^{-4}	15
	0.5	2.3	0.0146	1.4×10^{-4}	29
	0.1	10.1	0.0735	4.0×10^{-6}	140
Random IK		3.2	0.0080	1.0×10^{-6}	15
		6.5	0.0143	1.0×10^{-6}	30
		31.3	0.0525	1.0×10^{-6}	150
		63.3	0.0751	1.0×10^{-6}	300

where ℓ is the length-scale. Assuming observation noise, by using the observation value $\hat{\mathbf{q}}$, the covariance matrix is defined as

$$\mathbf{K} = k(\hat{\mathbf{q}}, \hat{\mathbf{q}}) + \sigma^2 \mathbf{E}. \quad (14)$$

Function values at the sampling points are set to $\mathbf{y} = \mathbf{1}$. The coefficient vector $\boldsymbol{\alpha}$ in the posterior mean, requirement to construct the GP-based implicit representation, is computed as

$$\boldsymbol{\alpha} = \mathbf{K}^{-1} \mathbf{y}. \quad (15)$$

Using $\boldsymbol{\alpha}$, the GPIS value at an arbitrary query point $\mathbf{q}$ is evaluated as

$$\phi(\mathbf{q}) = \mathbf{k}(\hat{\mathbf{q}}, \mathbf{q})^T \boldsymbol{\alpha}. \quad (16)$$

Since the GP is constructed from samples on the null-space manifold, we introduce a threshold $\beta \sim 1$ and regard configurations $\mathbf{q}$ satisfying $\phi(\mathbf{q}) > \beta$ as belonging to the null-space manifold.

The gradient of $\phi(\mathbf{q})$ can be computed analytically as

$$\nabla \phi(\mathbf{q}) = \sum_i -\alpha_i k(\hat{\mathbf{q}}_i, \mathbf{q}) \frac{(\hat{\mathbf{q}}_i - \mathbf{q})}{\ell^2} \quad (17)$$

where $\boldsymbol{\alpha} = [\alpha_1, \dots, \alpha_N]^T$ and $\hat{\mathbf{q}} = [\hat{\mathbf{q}}_1, \dots, \hat{\mathbf{q}}_N]^T$. Finally, following the kernel inversion in [6], we can estimate a pseudo-Euclidean distance field in configuration space induced by the implicit function as

$$d(\mathbf{q}) = \sqrt{\max(0, -2\ell^2 \log \phi(\mathbf{q}))}. \quad (18)$$

IV. EXPERIMENTS

A. Comparison of Jacobian-guided Sampling Methods

We compare the Jacobian-guided sampling method proposed in III-A2 with a baseline approach that randomly samples $\mathcal{M}(\mathbf{p}^*)$ based on inverse kinematics computations. As a benchmark problem, we consider a 3-DoF planar manipulator and sample points in the neighborhood of the solution manifold corresponding to a target point in a two-dimensional plane. Since the system has one redundant degree of freedom, the solution set $\mathcal{M}(\mathbf{p}^*)$ forms a one-dimensional manifold.

In the Jacobian-guided sampling method, the parameter controlling the step size along the tangent direction in (7) is




Fig. 3: Comparison of sampling results in configuration space. Left: Jacobian-guided methods ($\beta = 0.1$) vs. Random IK ($n = 150$). Right: Jacobian-guided methods ($\beta = 1.0$) vs. Random IK ($n = 15$).

set to $\beta = 1.5, 1.0, 0.5, 0.1$. For the inverse-kinematics-based random sampling method, we choose the number of samples as $n = 15, 30, 150$, and 300 , which are close to and comparable with those obtained by the Jacobian-guided method. Fig. 3 shows the sampled configurations in the joint-angle space. The upper figure corresponds to the Random IK method with $n = 150$ and the Jacobian-guided methods (Newton-based and Zigzag) with $\beta = 0.1$. The lower figure corresponds to the Random IK method with $n = 15$ and the Jacobian-guided methods with $\beta = 1.0$. As illustrated in Fig. 3, both methods successfully track the one-dimensional solution manifold.

For each method, we report the computation time required for sampling, the occupied volume of the sampled points, and the mean distance from the sampled points to the manifold evaluated using the residual function, as summarized in Table I. The coverage volume is approximated by discretizing the configuration space into a uniform grid with spacing s and counting the number of grid points that lie within an ϵ -neighborhood of the sampled points, multiplied by the volume of a single grid cell (in this case $s = 0.05$, $\epsilon = 0.5$). A larger coverage volume indicates that the solution manifold is covered more broadly.

From the experimental results, we observe that for large values of β , both the Newton-based and Zigzag methods sample the manifold coarsely, resulting in a small number of samples and reduced computation time. As β decreases, the sampling proceeds with smaller steps along the manifold, increasing the number of sampled points and the computational cost. However, this also leads to a larger covered volume and a smaller average distance from the manifold, indicating improved sampling accuracy.

The Zigzag method is expected to be computationally efficient because it relaxes the manifold projection condition, thereby reducing the cost required for convergence compared to the Newton-based method. As shown in Table I, under coarse sampling conditions ($\beta = 1.5, 1.0$), the Zigzag method achieves shorter computation times. In contrast, as β becomes smaller, the Newton-based method becomes faster than the Zigzag method. This behavior can be explained by the number of iterations required for convergence. When $\beta \geq 1.0$, the update in (7) induces a large movement along the tangent

TABLE II: Computation time for manifold neighborhood sampling and implicit GP construction across two robots and tasks

Robot	Task	Redundant DOF	Samples	Sampling Time [ms]	GP Construct Time [ms]
3 DOF arm	position (2D)	1	29	2.91	0.14
7 DOF arm	position and rotation (6D)	1	34	8.43	0.14

space, causing the iterate to deviate from the manifold. In this case, the Newton-based method typically requires three to four iterations to project back onto the manifold, whereas the Zigzag method converges within approximately one step. On the other hand, when $\beta < 1.0$, the tangent-space update remains within the vicinity of the manifold, and even the Newton-based method requires zero or one iteration for projection. Under these conditions, the additional computations in the Zigzag method, such as inner product evaluations at each step, dominate the overall cost, making the Newton-based method faster for smaller values of β . Regarding accuracy, the Zigzag method consistently exhibits larger errors measured by the distance to the manifold, which is expected since it does not enforce strict manifold projection at each iteration.

In the random inverse-kinematics (Random IK) sampling method, inverse kinematics is solved n times to find joint configurations $\mathbf{q}$ satisfying $\mathbf{p}^* = \mathbf{f}(\mathbf{q})$. Compared with the Jacobian-guided sampling methods that produce a similar number of samples, the Random IK method requires two to three times more computation time and results in a smaller coverage volume. These results indicate that the proposed Jacobian-guided sampling methods can sample points in the vicinity of the manifold more efficiently. In Fig. 3, the sampled points obtained by the Newton-based and Zigzag methods overlap visually; however, both methods sample the manifold in a nearly uniform manner. In contrast, the Random IK method produces noticeable gaps between sampled points, leading to non-uniform coverage of the manifold. This behavior is also reflected in the coverage-volume metric.

Based on the experimental results, we adopt the Newton-based Jacobian-guided sampling method in the following sections.

B. Generation of Implicit Null-space Manifolds

In this section, we evaluate whether a representation of the null-space manifold can be obtained from the sampled points generated by the Jacobian-guided sampling method using a Gaussian Process (GP). Table II summarizes the computation times required for Newton-based Jacobian-guided sampling of points in the vicinity of the manifold and for constructing a GP-based manifold representation, for a 3-DoF planar robot arm and a 7-DoF robot arm. The step size for the tangent update in the Newton-based method is set to $\beta = 0.5$.

The computation times required for sampling the manifold neighborhood were 2.91 ms for the 3-DoF planar arm and 8.41 ms for the 7-DoF arm, indicating that the 7-DoF case required approximately 2.89 times more computation time. In contrast, the time required for GP construction was ap-

proximately 0.14 ms for both models. This difference arises because the computational cost of sampling depends on the dimensionality of the configuration space, i.e., the robot's degrees of freedom, whereas the cost of GP construction mainly depends on the number of sampled points. Since the numbers of sampled points for the two models were similar (29 and 34, respectively), the GP construction times remained nearly identical.

We verify whether the constructed null-space manifold is properly represented. The configuration space is discretized, and inference is performed at each grid point using (16). Since the GP is constructed such that the output value at the sampled points is 1.0, grid points with inference values close to 1 can be regarded as lying on the manifold. Fig. 4 (a) shows the inference results obtained by discretizing the configuration space with a resolution of 0.05 for a 3-DoF planar robot arm, where a two-dimensional planar position is specified as the task. The manifold reconstructed from grid points whose GP inference values exceed 0.995 is shown in Fig. 4 (a). As illustrated in Fig. 4 (b), the one-dimensional null-space manifold is successfully captured. In addition, Fig. 4 (c) visualizes the distance field from the manifold computed based on (18). To further assess the accuracy of this field, we compare the inferred distances with ground-truth distances obtained via dense sampling for the 1D null-space manifold of the 3-DoF planar arm corresponding to Fig. 4 (c). The root-mean-square error (RMSE) is 7.67×10^{-3} , indicating that the proposed method provides an accurate local approximation of the distance field. The null-space manifold can also be constructed in the case of a 2D redundant manifold, as illustrated in Fig. 5, for a 3-link manipulator.

Fig. 7 visualizes the inferred manifold, distance field, and GP values for a 7-DoF robot arm. Since the configuration space is seven-dimensional, three joint axes are selected for visualization, while the remaining dimensions are marginalized by taking the maximum value. Using the constructed null-space manifold, it is also possible to compute solutions that satisfy the end-effector task, as illustrated in Fig. 6.

C. Global Proximity Queries using Implicit Null-space Manifold

By using (18), the distance from null-space the manifold can be computed, which makes it possible to obtain the shortest path from an arbitrary pose to the manifold. Moreover, the direction toward the manifold is given by (17), and moving along this direction by the distance computed from (18) allows the configuration to reach the manifold approximately in a single step. In Figs. 7 (a), (b), and (c), the red lines indicate trajectories in the configuration space, while in Fig. 7 (d), the red point represents the target, together with snapshots of the motion trajectory of an actual 7-DoF robot. Thus, the proposed method enables efficient computation of a path to the manifold. In practice, this projection step can be computed efficiently, requiring approximately 0.04 ms per evaluation including (17) and (18).




Fig. 4: Representation of the solution manifold and its induced fields. (a) Manifold field. (b) Null-space manifold. (c) Distance field.




Fig. 5: 2D null-space manifold and distance field (3D planar robot arm with 1 dof task ($x = 0.05$)).




Fig. 6: Get poses from null-space manifold.

D. Reconstruction of Spatially Distributed Tasks in the Configuration Space

So far, we have focused on tasks defined at a single point in the Euclidean space. This formulation can be extended to spatially distributed tasks by interpreting $\mathbf{p}$ as a field defined over the Euclidean space and reconstructing it in the configuration space within the proposed framework.


We consider two types of spatially distributed tasks: a one-dimensional task defined along a line and a two-dimensional task defined over a rectangular region. The line task corresponds to finding candidate joint-space trajectories that satisfy a given end-effector trajectory, while the rectangular task represents motions over a planar surface, such as a table-wiping motion.

For the line task, the trajectory is defined as $y = 0.6$, $-0.5 \leq x \leq 0.5$, $z = 0.3$, and the end-effector of a 7-DoF Franka robot is constrained to maintain a vertically downward orientation (Fig. 8 (a)). The null-space manifolds corresponding to the discretized task trajectory (30 points) are visualized simultaneously in Fig. 9 (b), where the color indicates the position along the task trajectory. As a result, a two-dimensional manifold is formed by combining the one-dimensional task structure with the single redundant degree of freedom at each task point.


Figure 9 (b) illustrates the null-space manifold associated with a two-dimensional rectangular task, in which the end-effector maintains a vertically downward orientation while moving over the region $0.4 \leq x \leq 0.5$, $-0.4 \leq y \leq 0.4$, $z = 0.3$ (Fig. 8 (b)). The corresponding joint configurations form a three-dimensional manifold, as shown in Fig. 9 (b), resulting from the combination of the two-dimensional task structure and one redundant degree of freedom. The apparent overlap of differently colored points is due to the projection of the seven-dimensional configuration-space manifold onto a three-dimensional space for visualization. Since the color encodes a field defined on the manifold, this representation suggests that, in the future, it may be possible to represent not only the distance field from the manifold in the configuration space but also fields defined directly on the manifold.

V. DISCUSSION


The proposed framework builds upon classical results on self-motion manifolds, which establish that solution sets induced by redundancy form smooth manifolds in the configu-




(a) Manifold field



(b) Null-space manifold



(c) Distance field



(d) The robot executing planned motion. The induced distance field guides the robot from a distant configuration to the solution manifold.


Fig. 7: Motion generation using the induced distance field. (a)-(c) The visualized distance-like fields and generated trajectories (red lines) in the configuration space. (d) The robot executing the planned motion to reach the target (red dot).

TABLE III: Quantitative Evaluation of Implicit Manifold Representation on 7-DOF Franka Manipulator

Task Type	Samples	Sampling Time [s]	GP Construct Time [s]	RMSE Error [m]
Line Task	3532	0.67	0.52	2.56×10^{-5}
Rectangle Task	8738	1.59	6.41	1.96×10^{-5}

ration space [2]. While these studies primarily focus on the analytical characterization of manifold properties, our work addresses a complementary and practical question: how to construct a reusable representation of such manifolds that can support repeated queries efficiently.


In conventional inverse kinematics, exploration of the configuration space is performed locally and repeatedly for each query in order to obtain a feasible solution. In contrast, the




(a) 1D line task

(b) 2D planar task

Fig. 8: Task definitions with continuous variation. (a) A line task where the end-effector follows a line while maintaining orientation. (b) A planar task over a rectangular region.



(a) 1D task: null-space manifold (colored by task-space position)



(b) 2D task: corresponding 3D null-space manifold

Fig. 9: Reconstructed manifolds in configuration space.

proposed approach performs a task-specific exploration of the configuration space in advance and encodes the resulting manifold structure into an implicit representation. This representation allows subsequent solution queries—such as distance evaluation, gradient computation, and projection onto the solution manifold—to be carried out efficiently, thereby amortizing the cost of exploration across multiple queries.

From this viewpoint, it is also instructive to compare the proposed approach with existing methods for exploring or representing solution manifolds. In contrast to ODE-based methods [4], which trace the manifold along a single tangent direction and are designed for one-dimensional null-space manifolds, the proposed approach incorporates an explicit projection step onto the manifold. While ODE-based methods rely on fixed step-size integration (e.g., 0.05 in [4]) to maintain constraint accuracy, our approach corrects deviations via projection and remains stable across a wider range of step sizes (e.g., 1.0 to 0.1). Moreover, extending ODE-based approaches

to higher-dimensional null spaces requires handling multiple tangent directions, whereas the proposed framework naturally accommodates such cases. The method further constructs a reusable implicit representation via sampling and GP-based modeling, enabling efficient exploration and fast proximity queries.

Here, reusability refers to the ability to retain the result of this precomputation as a single implicit field and to reuse it across multiple functionalities for a fixed task definition. Once constructed, the representation supports proximity queries, gradient-based updates, and solution generation without resolving inverse kinematics from scratch. In our implementation, such query operations are executed in the order of tens of microseconds, enabling fast online access to solution-space information.

From this perspective, it is instructive to compare the proposed framework with learning-based implicit representations, such as configuration-space distance field methods [7] and related neural-network-based models. These approaches construct a scalar field over the configuration space that encodes distances to constraint sets, and use gradient-based projection for inverse kinematics and planning. While both the proposed method and these approaches employ implicit representations, they differ in objective and construction methodology. Learning-based distance field methods focus on approximating a distance function through offline training, whereas the proposed method directly models the structure of the solution manifold, providing structured access to the solution set for proximity queries, gradient-based updates, and sampling. Moreover, such learning-based approaches typically rely on computationally expensive offline training procedures, requiring large datasets and training times ranging from hours to days. In contrast, the proposed method constructs the representation directly from local geometric information obtained via sampling, without requiring iterative training on large datasets. As a result, the construction time is on the order of milliseconds, substantially reducing the time-to-deployment, particularly in scenarios where task definitions or robot configurations change frequently.

The GPIS formulation adopted in this work plays a key role in enabling this efficient precomputation. GPIS allows meaningful manifold representations to be constructed from relatively small numbers of samples, which is particularly advantageous when sampling on solution manifolds is costly. Moreover, GPIS explicitly models uncertainty through predictive variance, providing confidence-aware information about the represented solution space. The resulting implicit function naturally induces a continuous distance field over the configuration space, from which gradients and projections can be computed directly, supporting efficient query-time operations.

At the same time, the present method is subject to practical limitations that primarily arise from the current strategy for exploring solution manifolds in the configuration space. Since the proposed framework relies on sampling-based exploration to construct the implicit representation, its scalability is inherently constrained by the curse of dimensionality, and efficient

exploration becomes increasingly challenging as the system dimension grows.

Topological transitions may occur near singular configurations, where the tangent space becomes ill-defined, affecting traversal, while projection still enforces approximate constraint satisfaction. In our 3-DoF planar robot example with a 2D task, the resulting null-space manifold forms a circular structure, and a singular configuration lies at the center of this circle, illustrating such topological behavior.

Joint limits can be incorporated by restricting the configuration domain, while self-collisions are not considered in the current formulation and remain future work. `macro`:

Topological transitions may occur near singular configurations, where the tangent space becomes ill-defined, affecting traversal, while projection still enforces approximate constraint satisfaction. In our 3-DoF planar robot example with a 2D task, the resulting null-space manifold forms a circular structure, and a singular configuration lies at the center of this circle, illustrating such topological behavior.

Joint limits can be incorporated by restricting the configuration domain, while self-collisions are not considered in the current formulation and remain future work. `macro`:

In addition, multiple disconnected solution components may arise in general. The proposed method explores a single connected component from a given seed using tangent-based traversal and projection. Additional components can be covered by reinitializing from other feasible configurations, for example obtained via inverse kinematics or random sampling.

Topological transitions may occur near singular configurations, where the tangent space becomes ill-defined, affecting traversal, while projection still enforces approximate constraint satisfaction. In our 3-DoF planar robot example with a 2D task, the resulting null-space manifold forms a circular structure, and a singular configuration lies at the center of this circle, illustrating such topological behavior.

Joint limits can be incorporated by restricting the configuration domain, while self-collisions are not considered in the current formulation and remain future work.

In contrast, the GPIS representation itself does not constitute a fundamental limitation of the framework. The main computational challenge associated with GPIS lies in the growth of inference cost as the number of samples increases, which is a well-known issue in GP models. This limitation can be addressed through established techniques such as sparse GP methods based on structured kernel interpolation [3] without altering the underlying representation-centric formulation.

Finally, while the current formulation targets kinematic task mappings, the underlying viewpoint—precomputing and reusing configuration-space structure through an implicit representation—is not restricted to them. Extending the framework to incorporate additional constraints, such as obstacle avoidance or interaction conditions, as well as improving scalability through sparse or approximate probabilistic models, constitutes an important direction for future work.

VI. CONCLUSION

In this paper, we proposed a representation-centric framework for explicitly modeling solution manifolds induced by redundancy in robotic systems. By combining Jacobian-guided sampling with GPIS, the proposed method constructs a probabilistic implicit representation whose zero-level set corresponds to the solution manifold in the configuration space.

The main contribution of this work is a shift in emphasis from conventional point-wise solution computation, such as inverse kinematics, to reusable spatial representations of solution sets. The constructed implicit representation is defined over the configuration space and naturally induces a continuous, distance field, which facilitates proximity queries, task-aware sampling, and reuse across related tasks.

While the proposed framework demonstrates promising capabilities, it also presents several limitations and directions for future work. In particular, the computational cost of GP inference poses scalability challenges for high-dimensional systems and large sample sets. Future research will investigate scalable probabilistic models and sparse representations. Moreover, extending the framework to incorporate more complex task definitions, such as obstacle avoidance and interaction constraints, remains an important direction toward broader applicability.

ACKNOWLEDGMENTS

This work was supported by JSPS KAKENHI Grant Numbers JP22H05002, JP25K21310. The authors would like to thank Prof. Wael Suleiman for the valuable discussions and insightful comments.

REFERENCES

- [1] Barrett Ames, Jeremy Morgan, and George Konidaris. Ikflow: Generating diverse inverse kinematics solutions. *IEEE Robotics and Automation Letters*, 7(3):7177–7184, 2022.
- [2] Joel W Burdick. On the inverse kinematics of redundant manipulators: Characterization of the self-motion manifolds. In *Advanced Robotics: 1989: Proceedings of the 4th International Conference on Advanced Robotics Columbus, Ohio, June 13–15, 1989*, pages 25–34. Springer, 1989.
- [3] David Eriksson, Kun Dong, Eric Hans Lee, David Bindel, and Andrew Gordon Wilson. Scaling gaussian process regression with derivatives. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS’18*, page 6868–6878, Red Hook, NY, USA, 2018. Curran Associates Inc.
- [4] Dominic Guri and George Kantor. Ode methods for computing one-dimensional self-motion manifolds. *arXiv preprint arXiv:2507.21957*, 2025.
- [5] Oussama Khatib. A unified approach for motion and force control of robot manipulators: The operational space formulation. *IEEE Journal on Robotics and Automation*, 3(1):43–53, 2003.
- [6] Cedric Le Gentil, Lan Wu, and Teresa Vidal-Calleja. Accurate gaussian process distance fields with applications to echolocation and mapping. *IEEE Robotics and Automation Letters*, 9(2):1365–1372, 2024.
- [7] Yiming Li, Xuemin Chi, Amirreza Razmjoo, and Sylvain Calinon. Configuration space distance fields for manipulation planning. *Robotics: Science and Systems (RSS) 2024*, 2024.
- [8] Carlos L Luck and Sukhan Lee. Redundant manipulator self-motion topology under joint limits with an 8-dof case study. In *Proceedings of 1993 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS’93)*, volume 2, pages 848–855. IEEE, 1993.
- [9] Wolfram Martens, Yannick Poffet, Pablo Ramón Soria, Robert Fitch, and Salah Sukkarieh. Geometric priors for gaussian process implicit surfaces. *IEEE Robotics and Automation Letters*, 2(2):373–380, 2016.
- [10] Yoshihiko Nakamura. *Advanced robotics: redundancy and optimization*. Addison-Wesley Longman Publishing Co., Inc., 1990.
- [11] Yoshihiko Nakamura and Hideo Hanafusa. Inverse kinematic solutions with singularity robustness for robot manipulator control. *Journal of Dynamic Systems, Measurement, and Control*, 1986.
- [12] Tomomichi Sugihara. Solvability-unconcerned inverse kinematics by the levenberg–marquardt method. *IEEE transactions on robotics*, 27(5):984–991, 2011.
- [13] Oliver Williams and Andrew Fitzgibbon. Gaussian process implicit surfaces. In *Gaussian Processes in Practice*, 2006.
- [14] Lan Wu, Raphael Falque, Victor Perez-Puchalt, Liyang Liu, Nico Pietroni, and Teresa Vidal-Calleja. Skeleton-based conditionally independent gaussian process implicit surfaces for fusion in sparse to dense 3d reconstruction. *IEEE Robotics and Automation Letters*, 5(2):1532–1539, 2020.