

# Connectivity-Aware Representations for Constrained Motion Planning via Multi-Scale Contrastive Learning

Suhyun Jeon<sup>1</sup>, Yumin Lim<sup>1</sup>, Woo-Jeong Baek<sup>2</sup>, Hyeonsoo Kim<sup>3</sup>, Suhan Park<sup>4,†</sup>, and Jaeheung Park<sup>1,5,†</sup>

**Abstract**—The objective of constrained motion planning is to connect start and goal configurations while satisfying task-specific constraints. Motion planning becomes inefficient or infeasible when the configurations lie in disconnected regions, known as essentially mutually disconnected (EMD) components. Constraints further restrict feasible space to a lower-dimensional submanifold, while redundancy introduces additional complexity because a single end-effector pose admits infinitely many inverse kinematic solutions that may form discrete self-motion manifolds. This paper addresses these challenges by learning a connectivity-aware representation for selecting start and goal configurations prior to planning. Joint configurations are embedded into a latent space through multi-scale manifold learning across neighborhood ranges from local to global, and clustering generates pseudo-labels that supervise a contrastive learning framework. The proposed framework provides a connectivity-aware measure that biases the selection of start and goal configurations in connected regions, avoiding EMDs and yielding higher success rates with reduced planning time. Experiments on various manipulation tasks showed that our method achieves 1.9 times higher success rates and reduces the planning time by a factor of 0.43 compared to baselines.

## I. INTRODUCTION

Constrained motion planning refers to the problem of computing a feasible path between a start and a goal configuration, where the entire path lies on the constraint manifold [1]. Such constraints commonly arise in robotic manipulation, for example, in bimanual or contact-rich tasks that require closed-chain kinematics, contact maintenance, and other physical or geometric conditions [2], [3], [4], [5]. These conditions restrict feasible configurations to a low-dimensional manifold embedded in the high-dimensional joint space. While this causes difficulties for conventional planning techniques like sampling-based planners in [6], addressing these challenges is essential for enabling robots to perform complex manipulation tasks in real-world environments.

\*This work was supported by the National Research Foundation of Korea (NRF) and the Institute of Information & Communications Technology Planning & Evaluation (IITP), both funded by the Korea government (MSIT) (Nos. RS-2024-00461583, RS-2024-00459435).

<sup>1</sup>Suhyun Jeon, Yumin Lim and Jaeheung Park are with the Department of Intelligence and Information, Graduate School of Convergence Science and Technology, Seoul National University, Republic of Korea. [suhyun0606, ckrqksakdmac, park731@snu.ac.kr

<sup>2</sup>Woo-Jeong Baek is with the Artificial Intelligence Institute (AIIS), Seoul National University, Republic of Korea. wjbaek@snu.ac.kr

<sup>3</sup>Hyeonsoo Kim is with the Department of Mechanical Engineering, Seoul National University, Republic of Korea. smtamh@snu.ac.kr

<sup>4</sup>Suhan Park is with the Department of Robotics, Kwangwoon University, Republic of Korea. park94@kw.ac.kr

<sup>5</sup>Jaeheung Park is also with ASRI, AIIS, Seoul National University and the Advanced Institutes of Convergence Technology, Republic of Korea.

† Equally advising

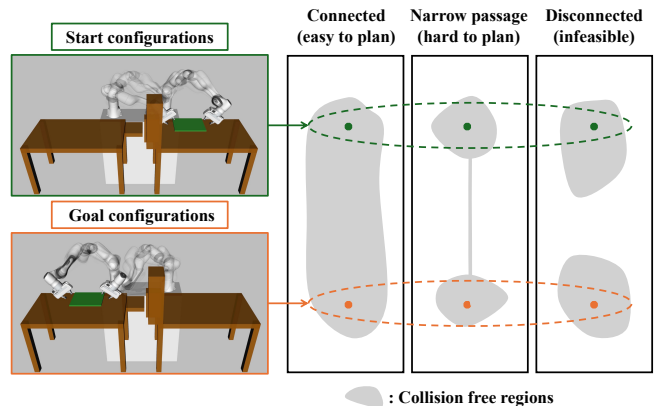


Fig. 1: Illustration of how different configuration selections in constrained motion planning result in varying levels of planning difficulty. (Left) Multiple inverse kinematic solutions for the start and goal pose. (Right) Connected configurations are easy to plan, configurations in narrow passages are difficult to plan, and disconnected configurations are infeasible.

A central difficulty in constrained motion planning is that the exact structure of the constraint-induced manifold is unknown. Traditional approaches have focused on extending sampling-based planners to handle such constraints [7], [8], [9], [10] through manifold projection, geodesic distances, or local planners designed to respect constraints. More recently, learning-based approaches have been proposed to approximate constrained manifolds directly, allowing motion planning to be guided by latent representations [11], [12]. Although these efforts have advanced constrained motion planning, they have focused mainly on improving sampling efficiency and designing local planners.

In redundant manipulators, a single end-effector pose admits infinitely many inverse kinematic solutions that may form continuous self-motion manifolds or fragment into discrete components [2], [13]. Given a start pose and a goal pose, each admits multiple inverse kinematics solutions, yielding many start–goal configuration pairs whose connectivity is unknown. As illustrated in Fig. 1, different choices of start and goal configurations lead to different planning performances: (i) connected configuration pairs can be solved quickly, (ii) configuration pairs separated by planner-hard environments, such as narrow passages, require significantly longer search time, and (iii) truly disconnected configuration pairs lead to planning failure. Following Xian et al. [14],

cases (ii) and (iii) are instances where the configurations lie in essentially mutually disconnected (EMD) components. EMD components are defined as regions of the configuration space that are either truly disconnected or practically unreachable, such that no feasible path can be found within a reasonable planning budget.

This work introduces a framework that learns connectivity-aware representations of the configuration space through multi-scale contrastive learning. The objective is to construct a feature-space where shorter distances between points correspond to a higher probability of connectivity. To achieve this objective, the proposed method proceeds in two stages. First, manifold learning is applied at multiple neighborhood scales to obtain embeddings, and each embedding is clustered to generate pseudo-labels. Second, these pseudo-labels supervise a neural network via multi-scale contrastive learning, where pseudo-labels from different scales are jointly incorporated so that the feature-space reflects connectivity information across multiple levels of granularity. The resulting representation is used in the pre-planning stage of constrained motion planning: multiple start–goal inverse kinematic candidates are embedded, and the closest pair in the learned feature-space is selected. This process leads to a higher likelihood of generating feasible plans, improving success rates and reducing planning times.

The main contributions of this paper are as follows:

- This paper proposes a simple and efficient method for generating multi-scale pseudo-labels for constrained motion planning, utilizing manifold learning across neighborhood scales and clustering to capture the structure of the configuration space.
- The proposed method learns connectivity-aware representations that guide the pre-planning stage to select start–goal configurations more likely to find feasible paths efficiently.
- Comprehensive experiments on diverse manipulation tasks validate the effectiveness of the proposed method, demonstrating higher success rates and reduced planning times compared to baseline methods.

The remainder of this paper is organized as follows. Section II reviews related works. Section III formulates the problem statement. Section IV presents the proposed method in detail. Section V reports the experimental setup and comparative results. Finally, Section VI concludes the paper.

## II. RELATED WORKS

### A. Sampling-Based Constrained Motion Planning

Sampling-based planners such as RRT [15], [16] and PRM [17] have been widely extended to constrained settings. Typical strategies include enforcing loop-closure constraints through projection or biasing sampling toward configurations that satisfy task-specific constraints [7], [8], [9], [10], [18], [19]. These approaches improve the exploration of low-dimensional restricted configuration spaces. However, a fundamental limitation remains that sampling-based planners cannot certify that no feasible solution exists [6].

This limitation becomes more severe in closed-chain problems [20], [21], where loop-closure constraints can fragment the feasible space into disconnected components [13], [14]. Although strategies such as regrasping or contact reconfiguration may restore feasibility [14], [22], they increase planning complexity and are not always applicable. Alternative approaches attempt to mitigate these issues by diversifying the set of candidate goal states, for example, through multi-goal pose selection [23] or iterative addition of inverse kinematic solutions [24]. Nevertheless, these methods do not explicitly consider the connectivity of the constraint manifold. In particular, feasibility depends on whether the start and goal states lie in the same connected region. This gap highlights the need for connectivity-aware representations.

### B. Manifold Approximation and Representation Learning

Data-driven approaches have recently emerged to approximate constrained manifolds directly, enabling planners to generalize feasible regions without explicit reliance on geometric models [12], [11], [25], [26], [27]. These methods improve sampling efficiency and feasibility prediction, but they primarily capture local validity and do not explicitly address whether configurations are connected or not. These limitations motivate the need for representation learning methods that are connectivity-aware and capable of preserving global topological properties of the configuration space.

In parallel, representation learning has gained increasing traction in robotics [28], inspired by its success in vision and natural language processing. Contrastive learning, in particular, has proven highly effective at learning structured embeddings from images, point clouds, and videos [29], [30], [31]. However, its application to robot configuration spaces remains limited because explicit connectivity labels are unavailable. As a result, most prior efforts in robotics have been restricted to perception-oriented tasks, and no existing work has leveraged contrastive learning to explicitly encode connectivity for motion planning. To address this gap, we generate pseudo-labels via clustering to supervise contrastive learning over configuration embeddings.

## III. CONSTRAINED MOTION PLANNING: CLOSED-CHAIN CONSTRAINTS

Multi-arm manipulation tasks with closed-chain constraints can be expressed in terms of the relative pose between end-effectors. Let  $T_i(q) \in SE(3)$  denote the end-effector pose of the manipulator  $i$  at configuration  $q \in \mathcal{C}_{\text{robot}}$ . The relative transform between manipulators  $i$  and  $j$  is

$$T_{ij}(q) = T_i(q)^{-1}T_j(q).$$

Let  $\mathcal{C}_{\text{free}} \subset \mathbb{R}^n$  denote the collision-free configuration space of the multi-arm system. The closed-chain constraint can be expressed as

$$T_{ij}(q) = T_i(q)^{-1}T_j(q) = T_{ij}^d, \quad \forall (i, j) \in \mathcal{G},$$

where  $\mathcal{G}$  is the set of manipulator pairs grasping the object and  $T_{ij}^d$  is the desired relative pose. The corresponding

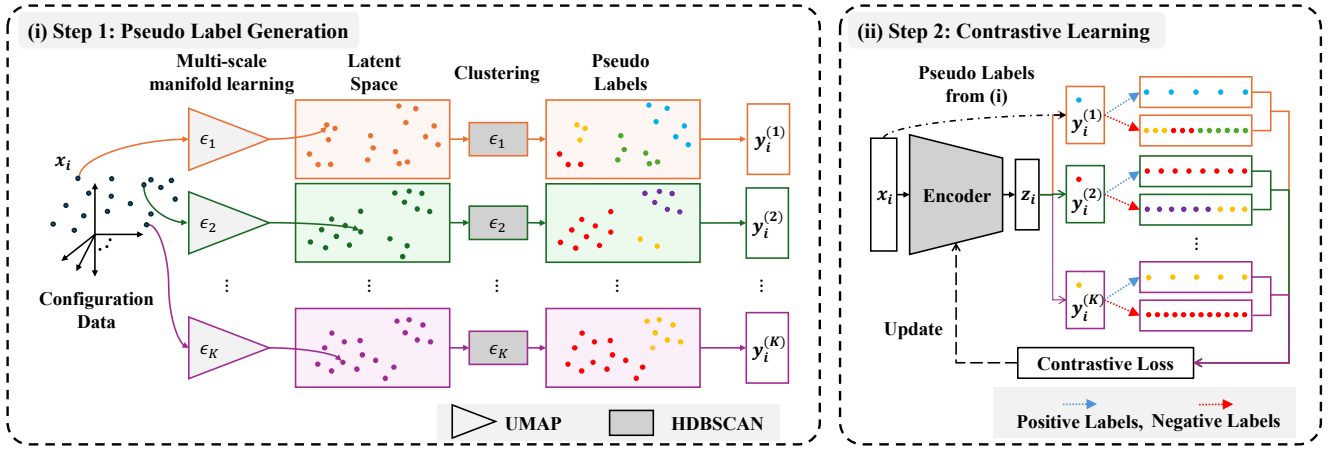


Fig. 2: Overview of the proposed method. The pipeline consists of two stages: (i) multi-scale pseudo-label generation via manifold learning and clustering at different neighborhood scales, and (ii) multi-scale contrastive learning guided by these labels, where configurations sharing the same pseudo-label are pulled closer in the feature-space and those with different labels are pushed apart, resulting in a connectivity-aware representation.

feasible manifold is

$$\mathcal{M} = \{q \in \mathcal{C}_{\text{free}} \mid T_{ij}(q) = T_{ij}^d, \forall (i, j) \in \mathcal{G}\}.$$

Given object-scale start and goal poses  $x_{\text{start}}, x_{\text{goal}} \in SE(3)$ , the inverse kinematics solution sets on the manifold are

$$Q_{\text{start}} = \{q \in \mathcal{M} \mid T_{\text{obj}}(q) = x_{\text{start}}\}$$

$$Q_{\text{goal}} = \{q \in \mathcal{M} \mid T_{\text{obj}}(q) = x_{\text{goal}}\}.$$

The constrained motion planning problem is defined as finding a continuous path

$$\gamma : [0, 1] \rightarrow \mathcal{M}, \quad \gamma(0) \in Q_{\text{start}}, \gamma(1) \in Q_{\text{goal}},$$

that moves the object from  $x_{\text{start}}$  to  $x_{\text{goal}}$  while satisfying the closed-chain constraints throughout.

#### IV. LEARNING CONNECTIVITY-AWARE REPRESENTATIONS FOR CONSTRAINED MOTION PLANNING

Robot configurations subject to task-specific constraints, such as closed-chain kinematics, form a low-dimensional manifold within the high-dimensional joint space. Feasible regions on this manifold are highly complex and often connected only through narrow passages, making direct analysis and planning difficult. While connectivity is formally defined as the existence of a feasible path between two configurations, in practice, it is more effective to reason about connectivity awareness, namely estimating which configurations are likely to be mutually reachable. Narrow passages, while mathematically continuous, are rarely sampled by planners and therefore behave as EMD components, highlighting the need for representations that capture connectivity structure and guide the choice of start–goal configurations.

To address this challenge, we construct a connectivity-aware representation of the configuration space (Fig. 2). The approach consists of two stages: (i) applying UMAP with multiple parameter settings and clustering the embeddings

to generate multi-scale pseudo-labels of local connectivity, and (ii) training a neural network model with multi-scale contrastive learning using these pseudo-labels as supervision. The resulting latent space embeds mutually reachable configurations close together while separating disconnected ones, enabling efficient and reliable constrained motion planning.

##### A. Generating Pseudo-Labels via Manifold Learning and Clustering

To enable contrastive learning on robot configuration data, pseudo-labels are generated by embedding high-dimensional configurations into a latent space with manifold learning followed by clustering. Task-specific constraints restrict robot configurations to a low-dimensional manifold. This property motivates the use of manifold learning to capture the intrinsic low-dimensional structure of the feasible configuration manifold while preserving local connectivity.

Firstly, UMAP [32] is applied to map the original joint configurations to a lower-dimensional latent space, preserving local neighborhood relationships and reflecting the connectivity of the constraint manifold. UMAP is chosen over alternatives such as VAE [33] or t-SNE [34] because it effectively captures both local and global manifold structure, handles disconnected regions robustly, and allows efficient mapping of new data points. To capture different scales of local structure, UMAP is applied with multiple parameter settings by varying the number of nearest neighbors and the minimum distance values. According to [32], larger neighbor sizes emphasize global structure while smaller ones highlight local neighborhoods, and the minimum distance parameter controls how tightly points are clustered in the latent space. Prior work has further shown that UMAP can improve clustering performance [35].

After obtaining latent embeddings at multiple scales, density-based clustering is performed using HDBSCAN [36], an extension of the classical DBSCAN [37] algorithm that automatically determines the number of clusters and

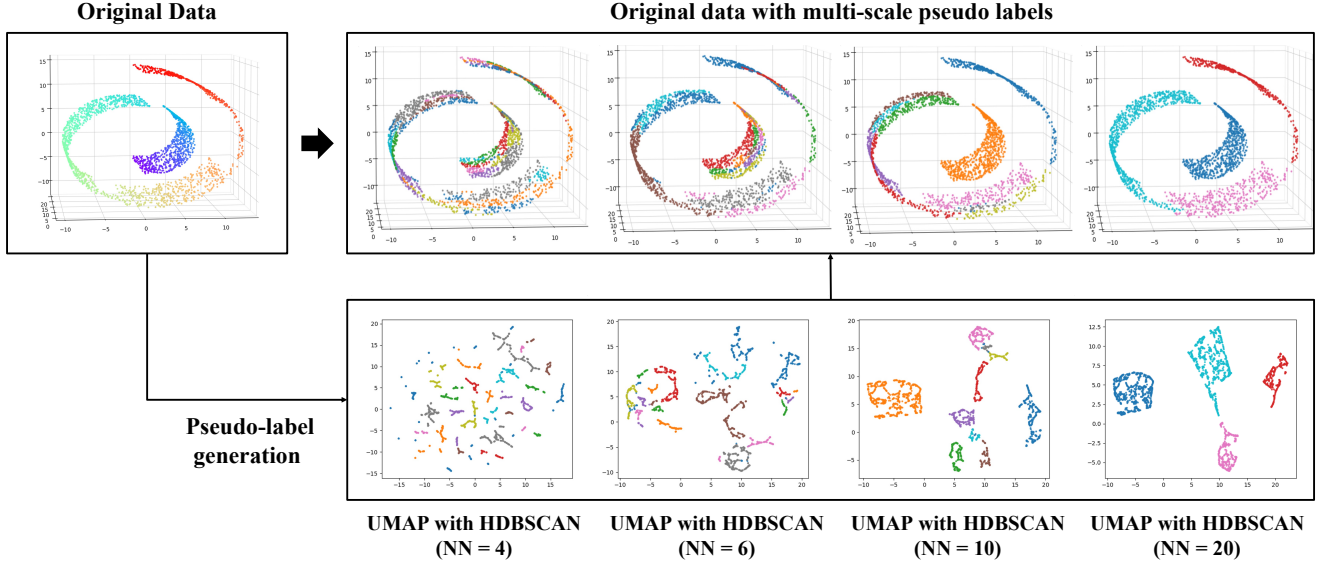


Fig. 3: Multi-scale pseudo-label generation on a modified Swiss Roll dataset emulating constrained regions in robotic configuration spaces. (Top) pseudo-labels projected onto the original data. (Bottom) latent embeddings with clustering. Left to right: increasing nearest neighborhood (NN) scales from local to global connectivity.

identifies outliers. Unlike k-means[38], which assumes uniform cluster density and requires predefining the number of clusters  $k$ , HDBSCAN can naturally detect clusters of varying densities and separate disconnected regions. This is particularly important in constrained configuration spaces, where feasible regions can be disconnected or have narrow bottlenecks. HDBSCAN is applied with tuned parameters, adjusting minimum cluster size and samples to generate robust pseudo-labels that capture the structural connectivity of the configuration manifold.

Finally, clustering results from multiple latent embedding and clustering procedures are aggregated to generate multi-scale pseudo-labels, which serve as supervisory information for the next contrastive learning step. These pseudo-labels encapsulate structural information about the configuration space at various scales, enabling the latent representation to preserve connectivity and support effective constrained motion planning. The overall process, from manifold learning to clustering and pseudo-label aggregation, is illustrated in Fig. 3. The modified Swiss Roll dataset emulates challenging structures of constrained configuration spaces, including narrow passages and disconnected regions. Multi-scale pseudo-labels are obtained by applying UMAP with clustering in the latent space (bottom) and then projected back onto the original data to highlight local connectivity (top).

### B. Contrastive Learning with Multi-scale Pseudo-Labels

Contrastive learning aims to structure a feature-space by pulling together samples that are similar and pushing apart those that are dissimilar. Building on the pseudo-labels obtained in the previous stage a multi-scale contrastive learning strategy is proposed to exploit hierarchical connectivity information. At each scale, configurations assigned to the same cluster are treated as positives, while those

from different clusters are treated as negatives. By jointly optimizing across scales, the learned representation preserves both global connectivity and local geometric structure of the feasible configuration space. This connectivity-aware representation provides a measure for selecting start and goal configurations likely to lie within the same connected region or favorable region, ultimately improving the feasibility and efficiency of constrained motion planning.

Formally, the framework trains a feature  $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^d$ , parameterized by neural network weights  $\theta$ . Let  $\mathcal{X} = \{x_1, x_2, \dots, x_N\}$  denote the dataset of robot configurations, and  $\mathcal{E} = \{\varepsilon_1, \dots, \varepsilon_K\}$  the set of parameters used for multi-scale manifold learning and clustering. For each  $\varepsilon_k$ , cluster assignments  $\{y_i^{(k)}\}_{i=1}^N$  are obtained for the data points  $x_i$ .

At scale  $k$ , the set of positive pairs for configuration  $x_i$  is defined as

$$\mathcal{P}_i^{(k)} = \{j \mid y_j^{(k)} = y_i^{(k)}, j \neq i\},$$

while all remaining configurations are treated as negatives. Following SimCLR [29], null space augmentation [12] is applied by generating perturbed samples  $\tilde{x}_i$  that vary within the null space of the Jacobian while preserving task constraints. These augmented samples are always considered positive, promoting invariance to constraint-preserving variations.

The contrastive loss at scale  $k$  is defined using the InfoNCE objective, which is widely used in contrastive learning:

$$\mathcal{L}_{\text{NCE}}^{(k)} = - \sum_{i \in \mathcal{I}} \log \frac{\sum_{j \in \mathcal{P}_i^{(k)} \cup \{\tilde{x}_i\}} \exp\left(\frac{\text{sim}(\phi(x_i), \phi(x_j))}{\tau}\right)}{\sum_{j \in \mathcal{I} \setminus \{i\}} \exp\left(\frac{\text{sim}(\phi(x_i), \phi(x_j))}{\tau}\right)}, \quad (1)$$

where  $\text{sim}(u, v) = \frac{u^\top v}{\|u\| \|v\|}$  denotes cosine similarity and  $\tau$  is a temperature hyperparameter.

The overall multi-scale contrastive loss is defined as a weighted sum across scales:

$$\mathcal{L}_{\text{MLCL}} = \sum_{k=1}^K \lambda_k \mathcal{L}_{\text{NCE}}^{(k)}, \quad \sum_{k=1}^K \lambda_k = 1, \quad \lambda_k \geq 0. \quad (2)$$

After training, the learned embedding provides a continuous measure of connectivity between two configurations. Formally, given a mapping  $\phi : \mathcal{C}_{\text{free}} \rightarrow \mathbb{R}^d$ , the representation is designed to capture connectivity-awareness:

- 1) Configurations in connected and easily accessible regions are embedded close together.
- 2) Configurations that are disconnected or lie in constrained bottlenecks are embedded farther apart.

The connectivity measure between two configurations  $\mathbf{q}_s$  and  $\mathbf{q}_g$  is modeled as

$$P_\phi(\mathbf{q}_s, \mathbf{q}_g) = \exp(-\|\phi(\mathbf{q}_s) - \phi(\mathbf{q}_g)\|_2). \quad (3)$$

which generalizes beyond sampled configurations. This allows connectivity-aware reasoning, particularly for selecting start and goal IK solutions that are more likely to be connected under constrained motion planning tasks.

In the proposed method, the learned representation serves as a connectivity-aware measure prior to motion planning. Rather than randomly selecting inverse kinematics solutions, start–goal pairs are evaluated in the feature space, and candidates with smaller distances are prioritized. This pre-selection step avoids wasted search in EMD components, improves planning success rate, and reduces planning time by providing favorable initial conditions.

## V. EXPERIMENTS

### A. Constrained Manipulation Planning Experiments

The proposed algorithm is empirically evaluated on constrained manipulation planning tasks across multiple robotic platforms and diverse task constraints. Primary experiments are conducted on the Franka Panda arm, while additional evaluations on the Tocabi humanoid robot [39] assess generality in higher-dimensional settings.

1) *Dataset Generation*: Directly sampling the constrained joint space is inefficient and misaligned with task-relevant regions. Instead, poses are sampled in the general task space to ensure comprehensive coverage of manipulation scenarios [24]. Dataset generation is performed in obstacle-free settings. For each sampled pose, inverse kinematics is solved using TRAC-IK [40] to generate joint configurations, and only self-collision-free solutions are retained to construct the dataset. The associated null-space Jacobians are recorded for each valid configuration. This procedure is applied consistently across single-arm, dual-arm, and triple-arm tasks, and the resulting datasets are subsequently used for learning latent embeddings and null-space data augmentation.

2) *Experimental Setup*: The proposed method is evaluated on constrained manipulation tasks where end-effector poses are specified and feasible motions must be generated under kinematic and environmental constraints. A critical challenge in this setting is that multiple IK solutions may correspond

to the same end-effector pose, yet not all start–goal pairs are mutually reachable. Consequently, selecting appropriate start and goal configurations is crucial for planning success.

To comprehensively assess the approach, experiments are primarily conducted on the Franka Panda arm across four representative scenarios (Fig. 4): a single-arm cup transfer, a dual-arm cooperative tray transfer, a dual-arm tray transfer with fixed orientation, and a triple-arm chair rotation. Each Panda arm has 7 degrees of freedom (DoF). In the dual-arm setting, closed-chain kinematic constraints reduce the dimensionality by 6 DoF, resulting in a 14–6=8 DoF constrained system. Adding a fixed orientation constraint further removes 2 DoF, leaving only 6 DoF. In the triple-arm setting, the combined 21 DoF are reduced by 12 through closed-chain constraints, yielding a 9 DoF constrained system. These reductions illustrate how constraint complexity increases planning difficulty across tasks. For each scenario, planning is performed with both a traditional projection-based planner, CBiRRT [7], and a learning-based latent motion planner [12], enabling assessment of whether the proposed method consistently improves performance across different planning frameworks. Each experiment is conducted on 100 randomized scenes, with three obstacles of varying size and position and a randomly sampled start–goal pose, under a planning time limit of 100 seconds. Note that these obstacles are not included during representation learning; the learned embedding is trained on obstacle-free configuration data and evaluated in environments containing previously unseen obstacles.


Additional experiments on the Tocabi humanoid with 8-DoF arms are conducted to evaluate generality. In the dual-arm setting, 16 DoF are reduced to 10 by closed-chain constraints and further to 8 with fixed orientation constraints. As shown in Fig. 5, two experiments on dual-arm manipulation and dual-arm manipulation with fixed orientation are conducted using CBiRRT under a 100-second time limit. This setup demonstrates that the proposed method scales to higher-dimensional arms and more complex kinematics while effectively handling redundancy.

Three strategies are compared for start–goal selection:


- **Random**: Random selection of start and goal IK solutions.
- **Joint-Space**: Selection of IK pairs minimizing Euclidean distance in joint space.
- **Feature-Space (Proposed)**: Selection of IK pairs minimizing distance in the learned feature space.

### B. Results and Discussion


Table I summarizes the results of 100 constrained manipulation tasks using the Franka Panda robot with randomized obstacle placements and start–goal poses. Planning time is averaged only over successful trials. The method consistently improves both success rates and planning times across all tasks and environments. Compared to the joint-space distance baseline, the proposed method achieves up to 1.9× higher success rates and reduces the average planning time to 0.43×. These gains arise from the connectivity-aware




(a) Single-arm manipulation: cup transfer task subject to fixed orientation constraints.



(b) Dual-arm manipulation: tray transfer task performed under closed-chain constraints.



(c) Dual-arm manipulation: tray transfer task performed under closed-chain constraints with fixed orientation.




(d) Triple-arm manipulation: chair rotation task under closed-chain constraints.


Fig. 4: Experimental setup in simulation using Franka Panda robots across diverse manipulation tasks.

TABLE I: Comparison of start–goal selection method on Franka Panda across diverse manipulation tasks (100 scenes, 100 second time limit). Metrics are the average planning time (successful trials only) and the success rates.

Task	Method	CBiRRT / 100 scenes / 100 sec		Latent Motion / 100 scenes / 100 sec	
		Planning time (sec)	Success rate (%)	Planning time (sec)	Success rate (%)
Fixed Orientation	Random	6.599 ± 11.533	92	2.944 ± 5.996	94
	Joint-Space	6.986 ± 8.744	<b>100</b>	8.230 ± 16.891	99
	Feature-Space (Ours)	<b>5.265 ± 5.221</b>	<b>100</b>	<b>1.808 ± 3.702</b>	<b>100</b>
Dual Manipulation	Random	37.403 ± 34.414	31	4.700 ± 7.904	90
	Joint-Space	47.211 ± 26.268	41	2.510 ± 5.554	<b>100</b>
	Feature-Space (Ours)	<b>30.865 ± 28.612</b>	<b>60</b>	<b>1.805 ± 3.156</b>	<b>100</b>
Dual Manipulation with Fixed Orientation	Random	33.988 ± 31.297	27	1.982 ± 2.261	46
	Joint-Space	56.163 ± 25.422	39	2.092 ± 2.510	99
	Feature-Space (Ours)	<b>32.552 ± 29.684</b>	<b>48</b>	<b>1.672 ± 1.615</b>	<b>100</b>
Triple Manipulation	Random	75.350 ± 0.000	1	2.762 ± 3.834	30
	Joint-Space	56.031 ± 29.892	6	1.906 ± 2.533	76
	Feature-Space (Ours)	<b>32.370 ± 23.106</b>	<b>10</b>	<b>1.500 ± 1.371</b>	<b>84</b>



(a) Dual-arm manipulation: tray transfer task performed under closed-chain constraints.



(b) Dual-arm manipulation: tray transfer task performed under closed-chain constraints with fixed orientation.

Fig. 5: Experimental setup in simulation using the Tocabi humanoid for dual-arm manipulation tasks.

TABLE II: Comparison of start–goal selection method on the humanoid robot Tocabi using CBiRRT (100 scenes, 100 second time limit). Metrics are the average planning time (successful trials only) and the success rates.

Method	Dual Manipulation		Dual Manipulation with Fixed Orientation	
	Planning time (s)	Success rate (%)	Planning time (s)	Success rate (%)
Random	7.351 ± 9.810	34	<b>25.963 ± 25.951</b>	18
Joint-Space	8.372 ± 13.162	<b>98</b>	35.778 ± 25.378	65
Feature-Space (Ours)	<b>6.909 ± 8.584</b>	<b>98</b>	38.341 ± 26.326	<b>80</b>

latent representation, which leverages multi-scale pseudo-labels and contrastive learning to avoid wasted computation in EMD components. Importantly, these benefits are observed not only when integrated with the CBiRRT but also when combined with Latent Motion planner, demonstrating improvements for both data-driven and traditional planner.

On the Tocabi humanoid, Table II shows that planning is generally easier due to the higher redundancy of its 8-DoF arms. As noted in [13], increasing the number of degrees of freedom tends to reduce the number of disconnected self-motion manifolds, which further improves configuration-space connectivity. In the dual-arm manipulation tasks, planning time decreases compared to both baselines, while success rates improve over the random baseline and remain comparable to the joint-space method. In contrast, in the dual-arm manipulation with fixed orientation tasks, planning time increases slightly, but success rates improve further, highlighting the benefit of connectivity-aware selection under stricter constraints.

Although some failures remain, they appear to stem primarily from the inherent difficulty of exploring constrained manifolds within a limited time budget. Notably, success rates improve when a more capable planner is used with the same representation, suggesting that planner limitations, rather than representation deficiencies, contribute to the remaining failures. Nonetheless, learning accurate global connectivity remains challenging, particularly in high-dimensional configuration spaces with complex constraints.

### C. Implementation Details and Computational Efficiency

The model consists of two main components: a four-layer MLP feature encoder with ReLU activations, and a two-layer MLP projection head with ReLU activations. Training is performed using the Adam optimizer with a learning rate of  $10^{-4}$ , batch size 1024, and 5000 epochs. For multi-scale pseudo-label generation, UMAP is applied with neighborhood sizes in the range  $3 \sim 50$  and minimum distance values 0.1, 0.2, followed by HDBSCAN clustering with minimum cluster size 20 and minimum samples 10. All experiments were run on a workstation with an Intel Xeon CPU and an NVIDIA RTX 4080 GPU.

In terms of efficiency, the preprocessing stage incurs a one-time cost of approximately 2 minutes for a dataset of 50,000 configurations in a dual-arm system with 14 DoF. Although the exact cost may vary depending on dataset size, robot configuration, and parameter settings, the overall preprocessing time remains modest. The inference time required to embed a configuration into the latent space is only 0.783 ms, which is negligible compared to the overall planning time. The model size also remains compact, at approximately 600 KB for the MLP model.

## VI. CONCLUSION

This paper presented a connectivity-aware representation learning framework for constrained motion planning. By integrating manifold learning, clustering, and contrastive learning, the proposed method constructs embeddings that pre-

serve both local and global connectivity, enabling improved start–goal configuration selection. Experiments on multi-arm manipulation tasks demonstrated consistently higher success rates and reduced planning times compared to baseline strategies. The current representation is learned in a task-specific manner and requires retraining when the underlying constraint manifold changes. Future work could focus on improving cross-task adaptability by conditioning on task or constraint parameters, learning shared structures across related manifolds, and developing hierarchical representations.

## REFERENCES

- [1] 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.
- [2] T. Cohn, S. Shaw, M. Simchowitz, and R. Tedrake, “Constrained bimanual planning with analytic inverse kinematics,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 6935–6942.
- [3] P. Liu, D. Tateo, H. B. Ammar, and J. Peters, “Robot reinforcement learning on the constraint manifold,” in *Conference on Robot Learning*. PMLR, 2022, pp. 1357–1366.
- [4] A. Pettinger, J. Panthi, F. Alambeigi, and M. Pryor, “Efficient constrained motion planning using direct sampling of screw-constraint manifolds,” *IEEE Transactions on Automation Science and Engineering*, 2024.
- [5] F. Krebs and T. Asfour, “A bimanual manipulation taxonomy,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11031–11038, 2022.
- [6] A. Orthey, C. Chamzas, and L. E. Kavraki, “Sampling-based motion planning: A comparative review,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 7, 2023.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] A. H. Qureshi, J. Dong, A. Choe, and M. C. Yip, “Neural manipulation planning on constraint manifolds,” *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6089–6096, 2020.
- [12] 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.
- [13] J. 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*. Springer, 1989, pp. 25–34.
- [14] 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.
- [15] S. LaValle, “Rapidly-exploring random trees: A new tool for path planning,” *Research Report 9811*, 1998.
- [16] J. J. Kuffner and S. M. LaValle, “Rrt-connect: An efficient approach to single-query path planning,” in *Proceedings 2000 ICRA. Millennium conference. IEEE international conference on robotics and automation. Symposia proceedings (Cat. No. 00CH37065)*, vol. 2. IEEE, 2000, pp. 995–1001.
- [17] L. E. Kavraki, M. N. Kolountzakis, and J.-C. Latombe, “Analysis of probabilistic roadmaps for path planning,” *IEEE Transactions on Robotics and automation*, vol. 14, no. 1, pp. 166–171, 1998.
- [18] M. Stilman, “Global manipulation planning in robot joint space with task constraints,” *IEEE Transactions on Robotics*, vol. 26, no. 3, pp. 576–584, 2010.
- [19] D. Berenson, S. Srinivasa, and J. Kuffner, “Task space regions: A framework for pose-constrained manipulation planning,” *The International Journal of Robotics Research*, vol. 30, no. 12, pp. 1435–1460, 2011.
- [20] J. H. Yakey, S. M. LaValle, and L. E. Kavraki, “Randomized path planning for linkages with closed kinematic chains,” *IEEE Transactions on Robotics and Automation*, vol. 17, no. 6, pp. 951–958, 2002.
- [21] J. Cortés, T. Siméon, and J.-P. Laumond, “A random loop generator for planning the motions of closed kinematic chains using prm methods,” in *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No. 02CH37292)*, vol. 2. IEEE, 2002, pp. 2141–2146.
- [22] P. Lertkultanon and Q.-C. Pham, “A certified-complete bimanual manipulation planner,” *IEEE Transactions on Automation Science and Engineering*, vol. 15, no. 3, pp. 1355–1368, 2018.
- [23] J. D. Hernández, M. Moll, and L. E. Kavraki, “Lazy evaluation of goal specifications guided by motion planning,” in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 944–950.
- [24] 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.
- [25] A. H. Qureshi, J. Dong, A. Baig, and M. C. Yip, “Constrained motion planning networks x,” *IEEE Transactions on Robotics*, vol. 38, no. 2, pp. 868–886, 2021.
- [26] C. Acar and K. P. Tee, “Approximating constraint manifolds using generative models for sampling-based constrained motion planning,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 8451–8457.
- [27] G. Sutanto, I. R. Fernández, P. Englert, R. K. Ramachandran, and G. Sukhatme, “Learning equality constraints for motion planning on manifolds,” in *Conference on Robot Learning*. PMLR, 2021, pp. 2292–2305.
- [28] Y. Cho, J. Han, Y. Cho, and B. Kim, “Corn: Contact-based object representation for nonprehensile manipulation of general unseen objects,” *arXiv preprint arXiv:2403.10760*, 2024.
- [29] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A simple framework for contrastive learning of visual representations,” in *International conference on machine learning*. PMLR, 2020, pp. 1597–1607.
- [30] P. Khosla, P. Teterwak, C. Wang, A. Sarna, Y. Tian, P. Isola, A. Maschinot, C. Liu, and D. Krishnan, “Supervised contrastive learning,” *Advances in neural information processing systems*, vol. 33, pp. 18661–18673, 2020.
- [31] Y. Li, P. Hu, Z. Liu, D. Peng, J. T. Zhou, and X. Peng, “Contrastive clustering,” in *Proceedings of the AAAI conference on artificial intelligence*, vol. 35, no. 10, 2021, pp. 8547–8555.
- [32] L. McInnes, J. Healy, and J. Melville, “Umap: Uniform manifold approximation and projection for dimension reduction,” *arXiv preprint arXiv:1802.03426*, 2018.
- [33] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*, 2013.
- [34] L. v. d. Maaten and G. Hinton, “Visualizing data using t-sne,” *Journal of machine learning research*, vol. 9, no. Nov, pp. 2579–2605, 2008.
- [35] M. Allaoui, M. L. Kherfi, and A. Cheriet, “Considerably improving clustering algorithms using umap dimensionality reduction technique: a comparative study,” in *International conference on image and signal processing*. Springer, 2020, pp. 317–325.
- [36] L. McInnes, J. Healy, S. Astels *et al.*, “hdbscan: Hierarchical density based clustering,” *J. Open Source Softw.*, vol. 2, no. 11, p. 205, 2017.
- [37] M. Ester, H.-P. Kriegel, J. Sander, X. Xu *et al.*, “A density-based algorithm for discovering clusters in large spatial databases with noise,” in *kdd*, vol. 96, no. 34, 1996, pp. 226–231.
- [38] T. Cover and P. Hart, “Nearest neighbor pattern classification,” *IEEE transactions on information theory*, vol. 13, no. 1, pp. 21–27, 1967.
- [39] M. Schwartz, J. Sim, J. Ahn, S. Hwang, Y. Lee, and J. Park, “Design of the humanoid robot tocabi,” in *2022 IEEE-RAS 21st International Conference on Humanoid Robots (Humanoids)*. IEEE, 2022, pp. 322–329.
- [40] P. Beeson and B. Ames, “Trac-ik: An open-source library for improved solving of generic inverse kinematics,” in *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2015, pp. 928–935.