

# Towards Feasible Dynamic Grasping: Leveraging Gaussian Process Distance Field, SE(3) Equivariance, and Riemannian Mixture Models

Ho Jin Choi<sup>†</sup> and Nadia Figueroa<sup>†</sup>

**Abstract**—This paper introduces a novel approach to improve robotic grasping in dynamic environments by integrating Gaussian Process Distance Fields (GPDF), SE(3) equivariant networks, and Riemannian Mixture Models. The aim is to enable robots to grasp moving objects effectively. Our approach comprises three main components: object shape reconstruction, grasp sampling, and implicit grasp pose selection. GPDF accurately models the shape of objects, which is essential for precise grasp planning. SE(3) equivariance ensures that the sampled grasp poses are equivariant to the object’s pose changes, enhancing robustness in dynamic scenarios. Riemannian Gaussian Mixture Models are employed to assess reachability, providing a feasible and adaptable grasping strategies. Feasible grasp poses are targeted by novel task or joint space reactive controllers formulated using Gaussian Mixture Models and Gaussian Processes. This method resolves the challenge of discrete grasp pose selection, enabling smoother grasping execution. Experimental validation confirms the effectiveness of our approach in generating feasible grasp poses and achieving successful grasps in dynamic environments. By integrating these advanced techniques, we present a promising solution for enhancing robotic grasping capabilities in real-world scenarios.

## I. INTRODUCTION

Grasping is a fundamental action enabling robots to execute manipulation tasks. The difficulty in grasping relies on finding an optimal grasp pose for the robot to perform, i.e., grasp sampling. To accomplish effective grasp sampling, consideration of three major components is required: (a) the shape and physical properties of the object; (b) the robot’s kinematics; and (c) potential collisions with the surroundings. Therefore, research in this field focuses on gathering information about these factors to help robots find good grasp positions.

In real-world scenarios, accurately determining object shapes can be hard due to occlusions or limited information. Analytic grasp samplers use force closure analysis [1] or differentiable simulation [2] to find good grasp poses, but they rely on knowing the object’s shape, which isn’t always possible. To address this limitation, new approaches have emerged in diverse manners, all without needing detailed prior knowledge about object shapes. These approaches include fitting shape primitives to an object’s partial point cloud [3], [4], estimating the composition of multiple shape primitives using neural networks (NN) [5], completing object shape using NN [6], or employing encoder-decoder architectures to implicitly learn object shape [7]. These methods




Fig. 1: Experimental setup used to evaluate our feasible dynamic grasping framework (left: static object, right: dynamic object). An Intel RealSense camera mounted on the robot’s end-effector is used to obtain a partial point cloud of the object and its pose (optitrack is used to track the object in dynamic scenarios) with no prior knowledge about the object category.

hold potential but often address only certain aspects of object shape, lacking vital information for effective control. Surface details and distance functions are critical for stable grasping and avoiding collisions. Recent efforts aim to learn the complete signed distance function (SDF) [8], which aids in filtering and validating data. Additionally, distance functions assist in smooth gripper operation and enable strategies for re-grasping if the object shifts during manipulation.

When addressing the robot’s kinematic constraints and collisions, two distinct approaches emerge: the explicit sampling of grasp poses achieved by learning kinematics and collisions implicitly [9], and the transformation of the grasp sampling challenge into a multi-objective optimization through the acquisition of a grasp pose cost or distance function [10], [8]. The first method is computationally fast but limited to specific situations, hindering its applicability across diverse scenarios. In contrast, the latter approach is more adaptable, as it can consider different constraints in the control optimization. It does not suffer from the challenges of grasp selection problem and can be applied to dynamic objects, unlike the linear bandit method [11] for static objects. However, extending this approach to task-specific grasp sampling, accounting for affordances, presents difficulties. Furthermore, SE(3) grasp pose distances don’t necessarily correlate with the nearest pose in joint space.

In scenarios involving human-robot interaction (HRI), it is crucial to address dynamic objects and implement reactive control strategies [12]. One viable approach is to optimize the end-effector [13] or sampled trajectory [14] based on the object’s pose. Recent research has explored leveraging the consistency of learned grasp latent vectors to match grasp poses across time steps [15], [16]. Furthermore, some use RL for object handover control policies [17]. While these methods show promise in learning the dynamics of grasping

<sup>†</sup>Ho Jin Choi and Nadia Figueroa are with the Department of Mechanical Engineering and Applied Mechanics, School of Engineering and Applied Science, University of Pennsylvania, Pennsylvania, PA 19104 USA (email: {cr139139,nadiafig}@seas.upenn.edu)




Fig. 2: **Feasible Dynamic Grasping Framework** pipeline. Given a partial point cloud of an object and its pose, we reconstruct the shape using the Gaussian Process Distance Field (Section II-A) with a refined distance metric as in Section III-A. Next, we sample poses from an SE(3) equivariant grasp sampler (Section III-B). Finally, the discrete set of feasible grasp poses is used to i) construct a continuous Riemannian GMM for dynamic task-space control (Section III-D.1) or ii) convert to joint-space via IK to construct a continuous GP for dynamic control in joint-space (Section III-D.2). The grasp sampling and controller components can operate asynchronously, allowing for real-time adaptation.

and reducing modeling complexity, their direct applicability to diverse scenes or methods remains a challenge.

**Contribution** This work introduces a sequential modular approach for sampling 6D grasp poses, reactive motion planning, and controlling dynamic objects (Figure 2). Our pipeline assumes a segmented point cloud of a rigid object, simplifying initial modeling. It consists of three main steps:

- 1) **Object Shape Reconstruction:** Reconstruct the object's shape using Gaussian Process distance field [18], providing an approximate SDF for resampling the complete point cloud for comprehensive representation.
- 2) **SE(3) Equivariant Grasp Sampling and Filtering:** Perform computationally efficient SE(3) equivariant grasp sampling and post-filtering with Gaussian Mixture Model-based reachability analysis for enhancing grasp pose quality and robustness.
- 3) **Implicit Grasp Pose Selection and Control:** Control the robot towards the implicit grasp pose using gradients from probability or distance functions, employing either task space or joint space methods for smooth motion to do automatic grasp selection.

By integrating these steps, our pipeline offers an effective solution for generating and evaluating grasp poses and facilitating motion planning and control in dynamic environments (Figure 1).

## II. PRELIMINARIES

Next we described the probabilistic and mathematical tools we use to build our feasible dynamic grasping framework.

### A. Gaussian Process Distance Field (GPDF)

Gentil et al. [18] proposed using Gaussian processes (GPs) to estimate an object or entire scene's distance field solely from its surface points. Let  $\mathbf{x} \in \mathbb{R}^m$  be a point and  $\mathbf{x}'$  be points lying on the surface  $\mathcal{S} \in \mathbb{R}^m$ . The occupancy field  $o(\mathbf{x}) \in \mathbb{R}$  is modeled as  $o(\mathbf{x}) \sim \mathcal{GP}(0, k_o(\mathbf{x}, \mathbf{x}'))$ , where  $k_o(d)$  is the covariance kernel. The inferred occupancy from point  $\mathbf{x}$  is calculated as:

$$\hat{o}(\mathbf{x}) = k_o(\mathbf{x}, \mathbf{X})(K_o(\mathbf{X}, \mathbf{X}) + \sigma_o^2 I)^{-1} \mathbf{1} \quad (1)$$

where  $K_o(\mathbf{X}, \mathbf{X})$  is the covariance matrix of the observed surface points  $\mathbf{X}$ ,  $\sigma_o^2$  is the noise variance, and  $\mathbf{1}$  is a vector of ones. Define the inverse function  $r$  as  $r(k_o(\mathbf{x}, \mathbf{x}')) = \|\mathbf{x} - \mathbf{x}'\| = d$ . Then, the distance field between  $\mathbf{x}$  and the surface  $\mathcal{S}$  is given by:

$$\hat{d}(\mathbf{x}) = r(\hat{o}(\mathbf{x})) \quad (2)$$

This can be understood as reversing the effect of the kernel function, which essentially represents the distance between two points when there is only one surface point and a query point involved. The gradient of the distance field, useful for obtaining surface normals, is analytically derived as:

$$\nabla d(\mathbf{x}) = \frac{\partial r}{\partial o} (\nabla k_o(\mathbf{x}, \mathbf{X})(K_o(\mathbf{X}, \mathbf{X}) + \sigma_o^2 I)^{-1} \mathbf{1}) \quad (3)$$

GPs provide an uncertainty measure expressed as:

$$\begin{aligned} \mathbf{var}(o(\mathbf{x})) &= k_o(\mathbf{x}, \mathbf{x}) \\ &- k_o(\mathbf{x}, \mathbf{X})(K_o(\mathbf{X}, \mathbf{X}) + \sigma_o^2 I)^{-1} k_o(\mathbf{X}, \mathbf{x}) \end{aligned} \quad (4)$$

This measure can assess the certainty of surface points, aiding in shape completion, as detailed in Section III-A. This SDF representation and uncertainty give us other measures to further filter out good grasp poses.

### B. SE(3) equivariance

The SE(3) Lie group is commonly used in robotics, especially for describing end-effector and grasp poses. It provides a unified representation consisting of rotation  $R \in SO(3)$  and translation  $t \in \mathbb{R}^3$ . A point in this representation is denoted as  $T = \begin{bmatrix} R & t \\ 0 & 1 \end{bmatrix} \in SE(3)$ . For grasp samplers focusing on object-centric approaches [19], [20], SE(3) equivariance is crucial. This ensures that feasible grasp poses of an object remain consistent relative to the object's pose. Let  $f: \mathbb{R}^{3 \times n} \rightarrow SE(3)$  be a grasp sampler, where  $\mathbf{P} \in \mathbb{R}^{3 \times n}$  represents the object's point cloud and  $G = \begin{bmatrix} R_G & t_G \\ 0 & 1 \end{bmatrix} \in SE(3)$  is a grasp pose. SE(3) equivariance requires that  $f(\mathbf{P}) = G$  satisfies:

$$f(R\mathbf{P} + t) = RG + t \quad (5)$$

for any arbitrary rotation  $R \in SO(3)$  and translation  $t \in \mathbb{R}^3$ . This is often achieved by centering the point cloud  $\mathbf{P}$  around its mean,  $\frac{1}{n} \sum_i \mathbf{P}_i$ , to obtain translation-invariant input. Then, any SO(3) equivariant neural network, such as Vector Neurons (VNN) [21], is utilized. Finally, the output is re-centered by adding the mean of the point cloud to the output. In this work, we adopt a similar approach to sample grasp poses, but we use the grasp pose representation introduced by M. Sundermeyer et al. [9], which is based on contact points. By utilizing this, we can not only train efficiently but also infer efficiently, giving us an advantage in dynamic grasping.

### C. Gaussian Distribution on Riemannian manifold

To generate continuous distributions of SE(3) poses, which we require in both the grasp pose post-filtering stage (Section III-C) and the task-space reactive control formulation (Section III-D.1), we use the Riemannian Expectation-Maximization (EM) algorithm by Calinon [22], suitable for Riemannian manifolds like SE(3) with a left-invariant metric. The Gaussian distribution on Riemannian manifold  $\mathcal{M}$  is:

$$\mathcal{N}_{\mathcal{M}}(x|\mu, \Sigma) = ((2\pi)^d |\Sigma|)^{-\frac{1}{2}} e^{-\frac{1}{2} \text{Log}_{\mu}(x) \Sigma^{-1} \text{Log}_{\mu}(x)} \quad (6)$$

Here,  $\mu \in \mathcal{M}$  is the origin of the tangent space,  $x \in \mathcal{M}$  is a point on the manifold,  $\Sigma \in \mathcal{T}_{\mu}\mathcal{M}$  is a covariance in the tangent space of  $\mu$ , and  $d$  is the dimension of the tangent space (6 for SE(3)). The mean  $\mu$  of  $N$  points on the manifold is iteratively computed as:

$$u = \frac{1}{N} \sum_{i=1}^N \text{Log}_{\mu}(x_i), \quad \mu = \text{Exp}_{\mu}(u) \quad (7)$$

After convergence, the covariance  $\Sigma$  is computed as:

$$\Sigma = \frac{1}{N} \sum_{i=1}^N \text{Log}_{\mu}(x_i) \text{Log}_{\mu}^T(x_i) \quad (8)$$

Note that the logarithmic and exponential operations in the Riemannian manifold are different from those in the Lie group for SE(3):

$$\text{Log}_{\mu}(x) = [\text{Log}_{\mathcal{G}}(\mu^{-1}x)]_{\mathcal{G}}^{\vee} = v \in \mathfrak{se}(3) \quad (9)$$

$$\text{Exp}_{\mu}(v) = \mu \text{Exp}_{\mathcal{G}}([v]_{\mathcal{G}}^{\wedge}) = x \in \text{SE}(3) \quad (10)$$

$\text{Log}_{\mathcal{G}}$  and  $\text{Exp}_{\mathcal{G}}$  are a logarithmic and exponential operation of the Lie group, converting SE(3) matrix into screw axis matrix representation and vice versa.  $[\cdot]_{\mathcal{G}}^{\vee}$  and  $[\cdot]_{\mathcal{G}}^{\wedge}$  are operations that change the screw axis matrix to screw axis compact vector form and vice versa in the Lie algebra.

### III. METHOD

Following we detail each step in our feasible dynamic grasping pipeline outline in Figure 2.

#### A. Shape Completion with Distance Field Refinement

The Matern  $\nu = 1/2$  kernel, given by  $k_o(d) = \sigma^2 \exp(-\frac{d}{l})$  with  $\sigma = 1$ , provides an approximate signed distance function (SDF) for an object, where  $l$  controls interpolation degree. This result is due to GP's interpolation abilities and the unique property of the Matern  $\nu = 1/2$  kernel, which provides an analytic solution for  $d$ . Other kernels lack analytic inverses or yield non-bijective outputs like  $\sqrt{d^2}$ . Despite its capability to estimate the sign of the distance, its accuracy diminishes near and inside the object, as in Figure 3 (left). To enhance accuracy, we employ an iterative refinement inspired by sphere marching. This process relies on two assumptions: 1) The estimated distance field can be represented as a composite function,  $d(\mathbf{x}) = g(\text{SDF}(\mathbf{x}))$ , where  $g: \mathbb{R} \rightarrow \mathbb{R}$  is a monotonically increasing function that goes through the origin and SDF denotes the actual signed distance field; 2) The absolute value of the estimated




Fig. 3: A 1D slice of the distance field of a sphere point cloud. The image on the left shows a comparison between the real SDF and the initial estimation from the GPDF. On the right, the image shows an improved estimation after five iterations of gradient descent, employing the ray marching concept.




Fig. 4: The reconstruction is done using the Gaussian process. The initial two images on the left display the mesh of a mustard bottle. In the subsequent two images, partial point clouds of the object are displayed. Finally, the last image shows the reconstructed point cloud of the object. Areas with low shape uncertainty are colored blue, while regions with high shape uncertainty are highlighted in red.

distance is less than or equal to the actual distance, i.e.,  $|d(\mathbf{x})| \leq |\text{SDF}(\mathbf{x})|$ . The iterative process involves projecting a point  $\mathbf{x}$  onto the object's surface using:

$$\mathbf{x} \leftarrow \mathbf{x} - d(\mathbf{x}) \frac{\nabla d(\mathbf{x})}{\|\nabla d(\mathbf{x})\|_2} \quad (11)$$

Accumulating  $d(\mathbf{x})$  values throughout the iterations yields a refined estimate of the distance from the original  $\mathbf{x}$ . The normalized gradient of the distance field becomes the updated point's surface normal. This iterative process converges without overshooting due to the constant Eikonal function property of the signed distance function, ensuring:

$$\begin{aligned} \frac{\nabla d(\mathbf{x})}{\|\nabla d(\mathbf{x})\|_2} &= \frac{g' \nabla \text{SDF}(\mathbf{x})}{\|g' \nabla \text{SDF}(\mathbf{x})\|_2} \\ &= \frac{g' \nabla \text{SDF}(\mathbf{x})}{g' \|\nabla \text{SDF}(\mathbf{x})\|_2} = \nabla \text{SDF}(\mathbf{x}) \end{aligned} \quad (12)$$

If the second condition is met, the process prevents overshooting and converges, as seen in Figure 3. We determine the best parameter  $l$  for GP using the Acronym dataset [23], generating partial point clouds for 8872 objects. For each  $l$  value, we create a GPDF and calculate the average distance between the complete object's point cloud and the GP-derived surface, akin to Chamfer distance. A small  $l$  results in a large mean distance, indicating poor reconstruction of unseen object parts. Conversely, an excessively large  $l$  causes over-interpolation, losing high-curvature details. Although adaptive  $l$  selection based on point cloud size is possible, we fix it at  $l = 0.3m$  through a grid search ( $l = [0, 0.6]$ ). This process aids in obtaining a rough object shape, as shown in Figure 4. Unlike DeepSDF [24], GPDF lacks semantic object information but offers versatility for both objects and scenes.




Fig. 5: Architecture of the data-driven grasp sampler, numbers (\*) denote the respective number of layers in each component.

### B. Data-driven grasp sampler

We use two VNN PointNet encoders, one VNN decoder for  $SO(3)$  equivariant features and one linear decoder for  $SO(3)$  invariant features, as depicted in Figure 5. The input to our neural architecture, denoted as  $\mathbf{P} \in \mathbb{R}^{n \times 2 \times 3}$ , includes surface points  $\mathbf{x}_s \in \mathbb{R}^{n \times 3}$  and their surface normals  $\nabla d(\mathbf{x}_s) \in \mathbb{R}^{n \times 3}$  extracted from GP. Surface points are generated through the creation of an  $8 \times 8 \times 8$  grid within the expanded bounding box of the partial point cloud, centered at its midpoint. This configuration can be adjusted according to the camera’s orientation. The network’s primary objective is twofold: first, to determine whether a given point in  $\mathbf{x}_s$  is a potential contact point on one side of the gripper; second, if viable, to estimate two columns of the rotation matrix and the width between the contact points. To validate if this grasp pose estimation is feasible, we get points on the gripper, transform points based on the grasp pose, and then input them into GP to verify whether they lie outside the object’s surface. Once we ascertain that all points are indeed situated outside the object, we proceed to retrieve the surface normals of the paired contact points  $\mathbf{x}$  and  $\mathbf{x}' = \mathbf{x} - w\mathbf{b}$  from GP. Subsequently, we examine whether the cosine similarity between the surface normals of these contact pairs closely approximates  $-1$ . This similarity metric serves as an indicative measure of force closure for jaw grippers.

### C. Reachable-Space Riemannian Mixture Model (RMM)

Reachability serves as a crucial metric for a robot manipulator to assess the feasibility of reaching a desired workspace pose  $G \in SE(3)$  with a given joint configuration  $q \in \mathbb{R}^n$ . It has various applications, such as guiding mobile manipulators to adjust their legs or wheels [25], [26] and optimizing grasp poses for fixed manipulators [13], [14]. Kim et al. [13] proposed a method utilizing a Gaussian Mixture Model (GMM) to pre-learn reachability, which expedited the process by avoiding explicit inverse kinematics computations with little loss in accuracy. This was accomplished by utilizing end-effector samples from a dataset of joint configurations, thereby enabling the optimization of end-effector poses in conjunction with grasp pose GMM. They fitted GMM by flattening out the first two columns of the rotation matrix and translation, which gave better results than using Euler angles or quaternions. In this work, we instead learn the reachable space based on the Riemannian Gaussian distribution introduced in Section II-C for better modeling. To build an offline reachable-space Riemannian Gaussian




Fig. 6: Gray grasp poses represent the outputs from the grasp sampler, the four red grasp poses correspond to clusters’ means from the RMM. The green trajectory of the end-effector illustrates a gradient ascent outcome, converging towards a feasible point.

Mixture Model (RMM), we use forward kinematics to derive end-effector poses. From 10 million joint configurations, we downsample to about 3 million, excluding self-collisions and poses below 0.1 m above the table. This will prevent the selection of grasp poses that cause collisions with the table. We determine 64 clusters using the Bayesian Information Criterion (BIC) with K-means++ initialization. Poses with probabilities below the 99% threshold are considered unreachable. If the environment is more cluttered one could fit a GPDF to the surface boundary and use it to filter out poses.

### D. Continuous Grasp Pose Distributions for Control

After filtering grasp poses, multiple viable ones may remain. Due to discrete samples, establishing stable control gradients amid disturbances is challenging. We propose two control strategies based on continuous distribution of grasp poses: task-space grasp representation with RMM and joint-space grasp representation with GPDF.

1) *Task-Space Grasp Distribution*: RMM is applied to sampled grasp poses, prioritizing more plausible ones. This allows computation of the gradient for the robot’s end-effector pose via gradient ascent on the RMM and differential inverse kinematics. In the case of RMM for  $SE(3)$ , the probability of pose  $x$  is given by:

$$\mathcal{P}(x|\mu, \Sigma) = \sum_{i=1}^n \pi_i \mathcal{N}_{\mathcal{M}}(x|\mu_i, \Sigma_i) \quad (13)$$

where  $\pi_i$ ,  $\mu_i$ , and  $\Sigma_i$  are weight, mean, and covariance of cluster  $i$ . The gradient is computed as follows:

$$\frac{\partial \mathcal{P}(x|\theta)}{\partial x} = \sum_{i=1}^n \pi_i \mathcal{N}_{\mathcal{M}}(x|\theta_i) J_r^{-1}(\mu^{-1})(-\Sigma^{-1} \text{Log}_{\mu_i}(x)) \quad (14)$$

with  $\theta_* = \{\mu_*, \Sigma_*\}$  and  $J_r$  is the right Jacobian of  $SE(3)$  [27]. Since obtaining  $J_r$  analytically can be challenging, we rely on numerical methods for gradient computation. Subsequently, we use this gradient to compute joint-space velocity controller:

$$\dot{q} = J(q)^\dagger \frac{\partial \mathcal{P}(x|\theta)}{\partial x} \quad (15)$$

where  $J(q)$  is the manipulator Jacobian and  $^\dagger$  denotes the Pseudo-Inverse. Eq. 15 acts as gradient ascent to the closest grasp pose (local maxima of GMM) from the robot’s current pose (Figure 6). However, there are two problems with using raw GMM gradients: zero gradient in local minima and zero




Fig. 7: Joint-space gradient field derived from GPDF using a 2D robot’s IK solutions. (left) End-effector targets as red points, accompanied by their respective IK solutions. (center-right) joint-space gradient field with varying interpolation value  $l$ .

gradient for poses distant from grasp poses. To escape from the local minima, we sample surrounding points and add a disturbance if they are identified as local minima. To solve the second problem, we use log probability instead,

$$\frac{\partial \log(\mathcal{P}(x|\theta))}{\partial x} = \frac{\sum_{i=1}^n \pi_i \mathcal{N}_{\mathcal{M}}(x|\theta_i) J_r^{-1}(\mu^{-1})(-\Sigma^{-1} \text{Log}_{\mu_i}(x))}{\sum_{i=1}^n \pi_i \mathcal{N}_{\mathcal{M}}(x|\theta_i)} \quad (16)$$

This gradient offers a numerical advantage when the probability is low, as multiplying a constant to all  $\mathcal{N}_{\mathcal{M}}(x|\theta_i)$  yields the same results. For dynamic objects, we adjust means with the object and compute the gradient. With new samples, we iteratively update the RMM. To ensure reachability, we transform means and refit with reachable grasp poses.

2) *Joint-Space Grasp Distribution*: Alternatively, if we have an inverse kinematics (IK) solver for a  $n$ -DoF robot with joint limits that can yield diverse solutions of  $q \in \mathbb{R}^n$  for a set of end-effector poses, we can create a  $n$ -D distance field using GPDF on the IK solutions. As GPDF is a monotonically increasing function it can be considered a potential function with multiple equilibria. We thus propose the following joint-space velocity controller,

$$\dot{q} = -d_q(q) \nabla d_q(q), \quad (17)$$

obtained from the GPDF fitted on the IK solutions,  $d_q$ , as shown in Figure 7 for 2-DoF example. Since sampling multiple IK solutions within a short time frame is challenging, we use a combination of the learning-based method, IKFlow [28], and the batched-damped pseudo-inverse method with batched forward kinematics and Jacobian calculation [29]. IKFlow gives multiple IK solutions that already, considering self-collision and joint limits at nearly 100 Hz using normalizing flow. Although its solutions may lack accuracy compared to the pseudo-inverse method, we initially sample IK solutions with IKFlow and refine them with the pseudo-inverse method for 3 iterations. For dynamic objects, we adjust grasp poses based on the object pose or get new grasp poses and resample IK solutions with the same seed for IKFlow, ensuring smooth transitions from previous solutions. To ensure reachability, we discard unreachable grasp poses.

#### IV. EXPERIMENTS

##### A. Experimental Setup

Our setup includes a KUKA iiwa 7 manipulator with a SAKE gripper and a Realsense D435 depth camera, utilized in both real robot and simulation experiments. Both task and joint space grasp representations are executed on an Intel i7-11700K without a GPU. However, for comparison, the grasp

TABLE I: Grasp Success Rate for Static Objects in Simulation

Method	Accuracy (%)
Grasp diffusion [8] + Energy function	67.23 (29.59)
Grasp diffusion [8] + Task-Space distribution	72.46 (19.84)
Grasp sampler (ours) + Task-Space distribution	77.60 (24.54)
(*) + Post-filtering	88.60 (15.75)
(*) + Reachability analysis	<b>90.00 (11.49)</b>

diffusion model runs on an Nvidia RTX 3080 GPU. Our primary objectives are twofold: firstly, to evaluate the quality of grasp poses generated by the grasp sampler pipeline, and secondly, to assess the feasibility of dynamic object grasping using joint and task space representations.

To evaluate grasp poses, we conduct experiments in both simulation and on a real robot, recognizing the potential dynamic nature of objects. Initially, experiments assume static objects with one-shot grasp sample predictions due to practical constraints, such as on-hand camera limitations where objects might go out of sight. We acknowledge the necessity of addressing dynamic scenarios in future research.

1) *Grasp Sampler Test*: To evaluate grasp poses, we conduct experiments in both simulation and on a real robot, assuming that the object remains static. In both simulation and real-world experiments, we utilize the YCB datasets [30], ensuring that the data used for evaluation were not part of the training dataset for the data-driven grasp sampler.

In simulation, we use Pybullet with fine-tuned friction coefficients, testing with 50 different YCB objects. Objects are initialized at random poses within a 10cm cube located at (x: 40cm, y: 0cm, z: 60cm) relative to the robot’s base coordinate. Our methodology involves capturing segmented RGBD images, sampling and filtering grasp poses, employing task-space representation for stability assessment, executing grasps, applying gravity simulation, and transitioning to a lifting pose. For benchmarking, we compare our approach against the SE(3) grasp diffusion model [8], which takes a partial point cloud as input unlike neural grasp distance [10], and an alternative method of sampling 100 discrete grasp poses from the grasp diffusion model and using our task-space representation. We justify the comparison with the SE(3) grasp diffusion model due to its prediction of the signed distance function (SDF) of the object and utilization of continuous grasp representation.

The success criterion is determined by whether the object remains stable for at least 2 seconds after the robot transitions to the lifting pose. This criterion aims to capture not only the initial grasp success but also the ability to maintain stability during subsequent manipulation steps. And we do this 10 times for each object. For the real robot experiment (Figure 1), we provide a segmented RGBD image to the robot using the method described in [31], sample grasp poses, convert them to joint space grasp representations, and assess the robot’s ability to successfully grasp the objects.

2) *Dynamic Object Test*: To assess the robot’s ability to grasp a moving object in simulation, we create circular, linear, and sinusoidal trajectories with constant object rotation (see accompanying video). We repeat testing five times for

TABLE II: Comparison for Dynamic Object Grasping in Simulation

Trajectory	Method	Self collision (%) ↓	Joint limit violations (%) ↓	Manipulability ↑	Gripper-object distance (cm) ↓
Circular	Joint space (ours)	<b>1.08 (0.57)</b>	<b>57.06 (6.52)</b>	<b>0.0936 (0.0017)</b>	16.65 (0.96)
	Task space (ours)	6.72 (5.55)	89.64 (0.74)	0.0729 (0.0022)	<b>3.31 (0.25)</b>
	Grasp diffusion [8]	38.14 (11.37)	95.47 (0.04)	0.0766 (0.0025)	17.63 (0.79)
Linear	Joint space (ours)	<b>0.00 (0.00)</b>	<b>55.49 (5.34)</b>	0.0887 (0.0052)	6.44 (0.24)
	Task space (ours)	3.84 (2.12)	62.90 (30.58)	0.0785 (0.0076)	<b>2.70 (0.19)</b>
	Grasp diffusion [8]	5.60 (0.77)	86.91 (0.89)	<b>0.0942 (0.0035)</b>	16.12 (0.55)
Sinusoidal	Joint space (ours)	<b>0.00 (0.00)</b>	<b>55.64 (3.82)</b>	<b>0.0884 (0.0021)</b>	12.47 (0.45)
	Task space (ours)	1.03 (1.97)	91.63 (5.63)	0.0847 (0.0033)	<b>7.10 (0.77)</b>
	Grasp diffusion [8]	28.88 (3.99)	89.36 (1.17)	0.0754 (0.0032)	14.85 (1.36)

each trajectory, evaluating joint limit violation, manipulability ( $m = \sqrt{\det(J(q)J(q)^T)}$ ), and self-collision without constraints. In real-world tests (Figure 1), we segment the RGBD image of the object, track its movement relative to the robot’s base frame using an Optitrack tracking system, and control the robot based on sampled grasp poses to verify grasp success. For grasp diffusion model, we use energy function directly to calculate the end-effector gradient.

### B. Experimental Results and Discussion

1) *Grasp quality*: In simulation, we achieved a 90.00% success rate in grasping out of 500 trials involving 50 distinct YCB objects, as detailed in Table I. Without post-filtering, grasp success significantly dropped, emphasizing the need for validating sampled grasp poses. Initially, the grasp diffusion model had poor results using an energy function due to non-geodesic trajectories in task space with gradient descent. Substituting the energy function with our task-space distribution improved performance but slowed inference to over 10 seconds, much longer than the 58 Hz with the energy function. However, even with this adjustment, performance remained inferior to our grasp sampler without filtering.

Our data-driven grasp sampler, running within 10 ms for 1024 points on CPU, may yield inaccurate samples since it is a small model. Full shape information from the distance field helps mitigate these errors. However, this can lead to overly conservative results, offering few grasp poses whereas the diffusion model can create infinite number of samples. Similar to our approach with IKFlow, we can sample IK solutions from data-driven methods and refine them numerically for grasp pose optimization [2] for improvement. Reachability analysis marginally improves sampler accuracy as the objects were already near the robot, where most grasp poses are reachable. Post-filtering and reachability analysis take less than 5 ms. However, grasp pose optimization becomes a bottleneck due to its  $O(n^3)$  complexity, taking 0.25 seconds for 512 points and 1-2 seconds for 1024 points.

In real robot testing with static objects, 14 out of 16 grasps were successful for 8 different objects. Failures occurred with the drill due to its heavy unbalanced weights and the banana due to excessive grasp strength causing slip. Similar failures happened with 19 out of 24 grasps for eight dynamic objects, despite good initial grasp contacts, due to object shape considerations.

TABLE III: Control Frequency of Task vs Joint Space Controllers

Number of GMM means	1	2	3	4	5	6	7	8	9	10
Frequency (Hz)	106.6	99.4	95.7	90.2	88.6	85.2	83.1	81.1	79.4	78.2
Number of GP’s samples	10	20	30	40	50	60	70	80	90	100
Frequency (Hz)	22.4	18.3	16.7	17.2	16	15.1	15.8	14.8	14.1	13.9

### 2) Comparison between task and joint space methods:

The task space method surpasses the joint space method in computational efficiency, as demonstrated in Table III. Specifically, the task space method shows notably faster execution times, rendering it more appealing. Notably, its frequency demonstrates nearly linear decrease with increasing GMM means, typically requiring only 4 to 8 means for adequate object grasp pose representation. In contrast, the frequency of the joint space method doesn’t follow a straightforward linear relationship with the number of IK solutions, largely due to parallel computing advantages.

Interestingly, the joint space method can exhibit speed enhancements even with increasing samples, owing to parallel computations and memory allocations at every power of 2. However, achieving smooth control with the joint space method necessitates a substantial number of samples, typically around 100, to ensure gradual joint pose changes and avoid abrupt transitions. From a robot kinematics standpoint, the joint space method outperforms the task space method, as highlighted in Table II. This superiority arises from the task space method’s susceptibility to singularities, while the joint space method can navigate around singularities, joint limits, and self-collision, resulting in higher manipulability scores. The grasp diffusion model often underperforms because it calculates the end-effector gradient similar to the task space method, resulting in an unnatural robot trajectory due to the gradient field it generates. Consequently, the strength of diffusion models primarily lies in offline optimization.

Although it might appear that the joint space method could outperform the task space method solely by achieving higher speeds, this assumption oversimplifies the comparison. For instance, in an HRI scenario like object handovers, the joint space method may be less intuitive for humans, resulting in perceived inefficiency due to larger distances between the gripper frame and the object frame during motions. In our real robot object handover tests, the task space method was preferred in this regard. The suitability of each method depends on the task’s requirements and constraints.

## V. CONCLUSION

In conclusion, our research focuses on enhancing robots’ object-grasping abilities in dynamic scenarios. We stress the importance of considering factors like object shape, potential collisions, and robotic motion when devising effective grasping strategies. We introduced a step-by-step approach that incorporates Gaussian Process Distance Field for object shape reconstruction, SE(3) equivariant grasp sampling and filtering, and Riemannian Mixture Models for grasp pose selection and control, aimed at improving robotic grasping efficiency. The proposed method provides a viable solution for generating and evaluating grasp poses, while also facilitating motion planning and control in dynamic environments, with potential applications in more complex settings.

## REFERENCES

- [1] V.-D. Nguyen, "Constructing force-closure grasps," *The International Journal of Robotics Research*, vol. 7, no. 3, pp. 3–16, 1988.
- [2] D. Turpin, T. P. Zhong, S. Zhang, G. Zhu, J. Liu, E. Heiden, M. Macklin, S. Tsogkas, S. Dickinson, and A. Garg, "Fast-grasp'd: Dexterous multi-finger grasp generation through differentiable simulation," *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 8082–8089, 2023.
- [3] Y. Wu, W. Liu, Z. Liu, and G. S. Chirikjian, "Learning-free grasping of unknown objects using hidden superquadrics," *ArXiv*, vol. abs/2305.06591, 2023.
- [4] S.-H. Kim, T. Ahn, Y. Lee, J. Kim, M. Y. Wang, and F. C. Park, "Dsnet: A deformable model-based supervised learning algorithm for grasping unknown occluded objects," *IEEE Transactions on Automation Science and Engineering*, vol. 20, pp. 1721–1734, 2023.
- [5] Y. Lin, C. Tang, F.-J. Chu, and P. A. Vela, "Using synthetic data and deep networks to recognize primitive shapes for object grasping," *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 10494–10501, 2019.
- [6] J. Varley, C. DeChant, A. Richardson, J. Ruales, and P. K. Allen, "Shape completion enabled robotic grasping," *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2442–2447, 2016.
- [7] A. Mousavian, C. Eppner, and D. Fox, "6-dof graspnet: Variational grasp generation for object manipulation," *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 2901–2910, 2019.
- [8] J. Uraïn, N. Funk, J. Peters, and G. Chalahatzaki, "Se(3)-diffusionfields: Learning smooth cost functions for joint grasp and motion optimization through diffusion," *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5923–5930, 2022.
- [9] M. Sundermeyer, A. Mousavian, R. Triebel, and D. Fox, "Contact-graspnet: Efficient 6-dof grasp generation in cluttered scenes," *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 13438–13444, 2021.
- [10] T. Weng, D. Held, F. Meier, and M. Mukadam, "Neural grasp distance fields for robot manipulation," *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1814–1821, 2022.
- [11] L. Wang, Y. Xiang, and D. Fox, "Manipulation trajectory optimization with online grasp synthesis and selection," *ArXiv*, vol. abs/1911.10280, 2019.
- [12] A. Billard, S. S. Mirrazavi Salehian, and N. Figueroa, *Learning for Adaptive and Reactive Robot Control: A Dynamical Systems Approach*. Cambridge, USA: MIT Press, 2022.
- [13] S. Kim, A. Shukla, and A. Billard, "Catching objects in flight," *IEEE Transactions on Robotics*, vol. 30, pp. 1049–1065, 2014.
- [14] I. Akinola, J. Xu, S. Song, and P. K. Allen, "Dynamic grasping with reachability and motion awareness," *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 9422–9429, 2021.
- [15] J. Liu, R. Zhang, H. Fang, M. Gou, H. Fang, C. Wang, S. Xu, H. Yan, and C. Lu, "Target-referenced reactive grasping for dynamic objects," *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 8824–8833, 2023.
- [16] H. Fang, C. Wang, H. Fang, M. Gou, J. Liu, H. Yan, W. Liu, Y. Xie, and C. Lu, "Anygrasp: Robust and efficient grasp perception in spatial and temporal domains," *ArXiv*, vol. abs/2212.08333, 2022.
- [17] S. Christen, W. Yang, C. P'erez-D'Arpino, O. Hilliges, D. Fox, and Y.-W. Chao, "Learning human-to-robot handovers from point clouds," *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9654–9664, 2023.
- [18] C. L. Gentil, O.-L. Ouabi, L. Wu, C. Pradalier, and T. Vidal-Calleja, "Accurate gaussian process distance fields with applications to echolocation and mapping," *ArXiv*, vol. abs/2302.13005, 2023.
- [19] A. Simeonov, Y. Du, A. Tagliasacchi, J. B. Tenenbaum, A. Rodriguez, P. Agrawal, and V. Sitzmann, "Neural descriptor fields: Se(3)-equivariant object representations for manipulation," *2022 International Conference on Robotics and Automation (ICRA)*, pp. 6394–6400, 2021.
- [20] B. Sen, A. Agarwal, G. Singh, B. Brojeshwar, S. Sridhar, and M. Krishna, "Scarp: 3d shape completion in arbitrary poses for improved grasping," *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3838–3845, 2023.
- [21] C. Deng, O. Litany, Y. Duan, A. Poulencard, A. Tagliasacchi, and L. J. Guibas, "Vector neurons: A general framework for so(3)-equivariant networks," *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 12180–12189, 2021.
- [22] S. Calinon, "Gaussians on riemannian manifolds: Applications for robot learning and adaptive control," *IEEE Robotics & Automation Magazine*, vol. 27, pp. 33–45, 2019.
- [23] C. Eppner, A. Mousavian, and D. Fox, "Acronym: A large-scale grasp dataset based on simulation," *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6222–6227, 2020.
- [24] J. J. Park, P. R. Florence, J. Straub, R. A. Newcombe, and S. Lovegrove, "DeepSDF: Learning continuous signed distance functions for shape representation," *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 165–174, 2019.
- [25] S. Jauhri, J. Peters, and G. Chalahatzaki, "Robot learning of mobile manipulation with reachability behavior priors," *IEEE Robotics and Automation Letters*, vol. 7, pp. 8399–8406, 2022.
- [26] N. Vahrenkamp, T. Asfour, and R. Dillmann, "Robot placement based on reachability inversion," *2013 IEEE International Conference on Robotics and Automation*, pp. 1970–1975, 2013.
- [27] G. S. Chirikjian, "Stochastic models, information theory, and lie groups, volume 2," 2012.
- [28] B. Ames, J. Morgan, and G. D. Konidaris, "Ikflow: Generating diverse inverse kinematics solutions," *IEEE Robotics and Automation Letters*, vol. 7, pp. 7177–7184, 2021.
- [29] S. Zhong, T. Power, and A. Gupta, "PyTorch Kinematics," 3 2023.
- [30] B. Çalli, A. Walsman, A. Singh, S. S. Srinivasa, P. Abbeel, and A. M. Dollar, "Benchmarking in manipulation research: The ycb object and model set and benchmarking protocols," *ArXiv*, vol. abs/1502.03143, 2015.
- [31] C. Zhang, D. Han, Y. Qiao, J. U. Kim, S.-H. Bae, S. Lee, and C. S. Hong, "Faster segment anything: Towards lightweight sam for mobile applications," *arXiv preprint arXiv:2306.14289*, 2023.