

# $\mathcal{T}(\mathcal{R}, \mathcal{O})$ Grasp: Efficient Graph Diffusion of Robot-Object Spatial Transformation for Cross-Embodiment Dexterous Grasping

Xin Fei<sup>1,2\*</sup>, Zhixuan Xu<sup>1,2\*</sup>, Huaicong Fang<sup>3</sup>, Tianrui Zhang<sup>1</sup>, Lin Shao<sup>1,2†</sup>  
<sup>1</sup>Department of Computer Science, National University of Singapore <sup>2</sup>RoboScience  
<sup>3</sup>College of Control Science and Engineering, Zhejiang University

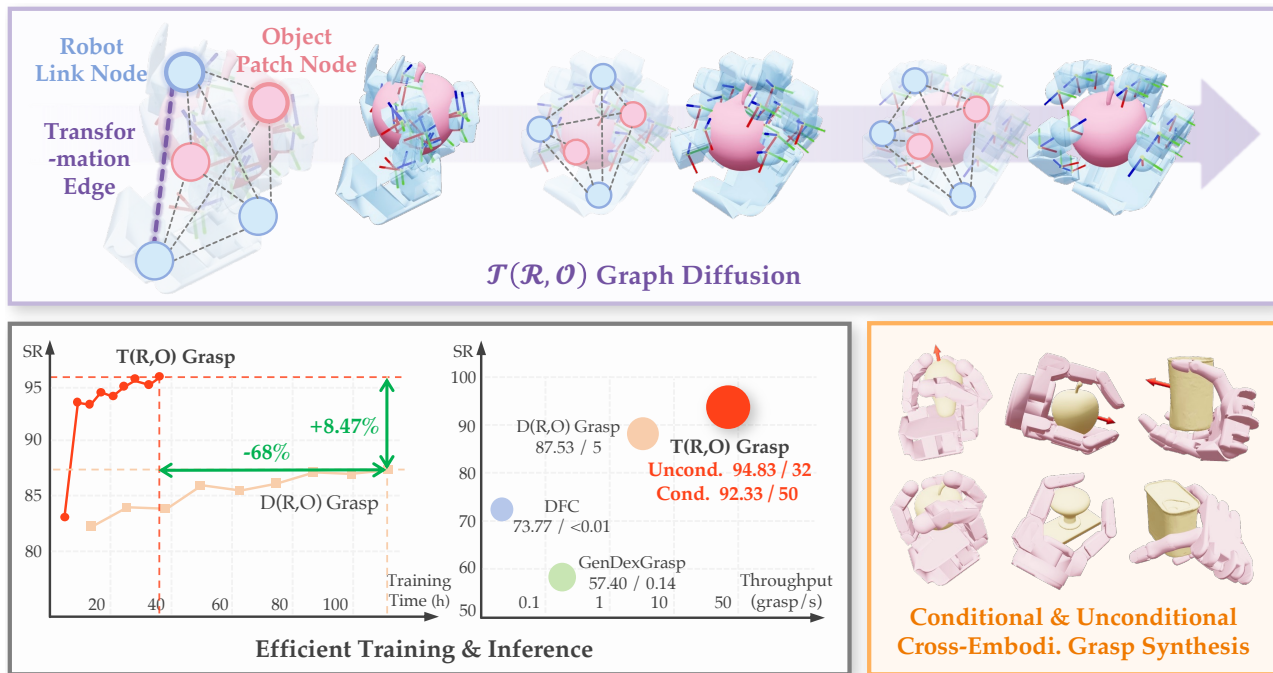


Fig. 1: Given object point cloud and hand URDF,  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp efficiently supports both conditioned and unconditioned grasp synthesis utilizing a graph diffusion model. Compared with  $\mathcal{D}(\mathcal{R}, \mathcal{O})$  Grasp [1], our method achieves superior performance with lower memory usage, significantly higher inference speed and throughput.

**Abstract**—Dexterous grasping remains a central challenge in robotics due to the complexity of its high-dimensional state and action space. We introduce  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp, a diffusion-based framework that efficiently generates accurate and diverse grasps across multiple robotic hands. At its core is the  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph, a unified representation that models spatial transformations between robotic hands and objects while encoding their geometric properties. A graph diffusion model, coupled with an efficient inverse kinematics solver, supports both unconditioned and conditioned grasp synthesis. Extensive experiments on a diverse set of dexterous hands show that  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp achieves average success rate of 94.83%, inference speed of 0.21s, and throughput of 41 grasps per second on an NVIDIA A100 40GB GPU, substantially outperforming existing baselines. In addition, our approach is robust and generalizable across embodiments while significantly reducing memory consumption. More importantly, the high inference speed enables closed-loop dexterous manipulation, underscoring the potential of  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp to scale into a foundation model

for dexterous grasping. The code, appendix, and videos are available at <https://nus-lins-lab.github.io/trograspweb/>.

## I. INTRODUCTION

Grasping with dexterous hands is a fundamental capability for achieving precise, human-level manipulation. Yet, efficiently generating diverse and high-quality grasps remains a longstanding challenge, largely due to the high dimensionality of dexterous hands and the difficulty of ensuring both stability and precision. Existing learning-based research has primarily explored two types of representations: *robot-centric* and *object-centric*. *Robot-centric* approaches [2]–[6], such as representing grasp with wrist poses or joint values, enable rapid inference by mapping observations directly to control signals. However, they often suffer from poor sample efficiency and limited transferability, as the learned representations are implicitly tied to specific hand embodiments and fail to generalize across novel morphologies. In contrast, *object-centric* paradigms focus on describing grasps in terms of object features such as contact points or

\* denotes equal contribution

† denotes corresponding author

This research is supported by the Ministry of Education, Singapore, under the Academic Research Fund Tier 1 (FY2024)

affordances [7]–[13]. This representation naturally promotes transferability across diverse objects and robotic platforms. However, it is often brittle under partial observations and typically requires an additional optimization stage, such as fingertip IK or constrained fitting of contact maps, to convert abstract predictions into feasible joint configurations, making the process computationally expensive and time-consuming.

Recently,  $\mathcal{D}(\mathcal{R}, \mathcal{O})$  Grasp [1] introduced an *interaction-centric* representation for dexterous grasping. Unlike *robot-centric* or *object-centric approaches*, an *interaction-centric representation* characterizes a grasp by directly modeling the spatial relationships between the robot hand and the object, embedding both entities into a shared interaction space. This formulation combines the strengths of robot-centric and object-centric paradigms: it retains embodiment awareness while capturing transferable object-level features. As a result, it facilitates generalization across different hand embodiments and object shapes, preserves the fine-grained geometric structure of potential contacts, and remains more robust under partial observations. In  $\mathcal{D}(\mathcal{R}, \mathcal{O})$ , this is achieved by constructing a distance matrix between points sampled on the hand and the object. However, the point-to-point feature maps in  $\mathcal{D}(\mathcal{R}, \mathcal{O})$  Grasp require substantial GPU memory. Thus, the representation is resource-intensive and difficult to scale up. In addition,  $\mathcal{D}(\mathcal{R}, \mathcal{O})$  relies on a suitable initial status as input, which makes the model performance vulnerable to infeasible initialization.

To overcome the limitations of  $\mathcal{D}(\mathcal{R}, \mathcal{O})$  representation, we propose  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp, an efficient graph diffusion model for cross-embodiment dexterous grasp synthesis based on  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph representation. Formally, for any object-hand pair, we formulate a unified Transformation Graph between the **Robot** and **Object**, namely the  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph. In  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  graph, nodes encode object and link geometries together with their spatial transformations in object global frame for a grasp, while edges consist of relative spatial transformations between nodes. Compared to *robot-centric* representation, such definition is unified and generalizable across all objects and dexterous hand embodiments. Unlike  $\mathcal{D}(\mathcal{R}, \mathcal{O})$  that requires memory-intensive point-to-point representation,  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph establishes relations between object patches and hand links, which significantly reduces memory usage and enables faster training and inference.

Built upon  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  representation, we implement an efficient transformer-based graph diffusion model for  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph generation, thereby enabling cross-embodiment dexterous grasp synthesis either unconditionally or conditioned on a hand status. In extensive experiments,  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp achieves the grasping success rate of 94.83%, with an average inference speed of 0.21s and a throughput of 41 grasps per second on a NVIDIA A100 40GB GPU. More importantly, the fast inference speed allows  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp to support *closed-loop dexterous grasping and manipulation*, a key requirement for real-world deployment. In conclusion, our main contributions are summarized as:

- 1) We introduce the  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph, a unified and generalizable representation for modeling interactions across

all object–hand pairs. This representation is significantly more memory-efficient than the  $\mathcal{D}(\mathcal{R}, \mathcal{O})$  formulation and establishes a foundation for scaling toward large dexterous grasping foundation models.

- 2) We develop an efficient graph-diffusion model built on the  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph that supports both conditioned and unconditioned grasp synthesis. This design not only mitigates sensitivity to infeasible initializations but also enables reliable closed-loop dexterous manipulation.
- 3) We conduct extensive experiments both in simulation and real world, demonstrating that  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp significantly outperforms all baselines on success rate for dexterous grasping while delivering substantially higher efficiency than existing methods.

## II. RELATED WORK

### A. Representation for Dexterous Grasping

Grasping with dexterous robotic hands holds great potential for achieving human-like precision in manipulation. However, it remains highly challenging due to the inherent high dimensions. Achieving a unified model for cross-embodiment dexterous grasping requires an efficient and robust representation that consistently encodes the geometry and interactions of objects and robotic hands. A natural and straightforward approach is to represent a grasp with hand joint values [2]–[4], [14]. While these predictions can be directly executed, the representation of joint values lacks a canonical definition, as a single hand can admit multiple valid joint parameterizations. In addition, its generalization across embodiments is limited, and its efficiency is hindered by the prohibitive computational cost of reinforcement learning in high-dimensional action space.

Another stream of research formulates grasp representation in an object-centric way, typically by encoding object contact points or contact maps [7]–[11], [15]. Though these methods allow generalization across embodiments and objects, they typically rely on contact-based optimization, which hinders real-time grasp synthesis. More recently,  $\mathcal{D}(\mathcal{R}, \mathcal{O})$  Grasp [1] pioneers to introduce an interaction-based representation by predicting a point-to-point distance matrix between objects and robotic hands, which significantly facilitates robust generalization across diverse objects and hand embodiments. However, the  $\mathcal{D}(\mathcal{R}, \mathcal{O})$  representation faces several limitations: (1) the point-to-point feature maps constructed prior to  $\mathcal{D}(\mathcal{R}, \mathcal{O})$  estimation incur substantial memory overhead, which is memory-inefficient and hinders its scalability; (2) the prediction relies heavily on the initial pose (as it conditions on latent outputs from a CVAE [16]), which can degrade performance when the initial pose is suboptimal or infeasible for a valid grasp; (3)  $\mathcal{D}(\mathcal{R}, \mathcal{O})$  representation requires additional optimization to obtain SE(3) transformations of each link. To overcome these limitations, we propose  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph as a unified and efficient representation which centers on spatial transformation between hand links and object patches. In addition,

we implement a lightweight graph diffusion model to support both unconditioned and conditioned grasp synthesis.

### B. Grasp Synthesis with Generative Models

With the success of generative models, numerous researchers train generators on large-scale datasets with sample refinement to perform grasp synthesis. Contact-GraspNet [17] efficiently generates collision-free grasps in cluttered scenes, while GraspME [18] introduces a manifold estimator to model feasible grasp in object manifold. Neural Grasp Distance Fields [19] leverages implicit distance representations for continuous grasp synthesis, and SE(3)-DiffusionFields [20] further employs diffusion-based learning to jointly optimize grasp and motion through smooth cost functions. Despite these advances, all prior methods remain constrained to 6-DoF pose of a single robotic gripper, which is fundamentally inadequate for dexterous hand grasping due to the lack of a suitable representation. In this paper, our proposed  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph provides a unified and lightweight representation for dexterous hands with arbitrary geometry and numbers of links, enabling high-dimensional grasp synthesis with generative models beyond the 6-DoF robotic gripper.

## III. METHOD

Given an object point cloud  $P^O \in \mathbb{R}^{N \times 3}$ , URDF description of the dexterous hand, and the point clouds of  $L$  links  $\{P_i^R\}_{i=1}^L$  in the local frame of each link,  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp generates diverse grasps by first predicting SE(3) transformations of all links while temporarily relaxing joint constraints. Then, the predicted transformations are refined and converted into executable joint values through Pyroki-based inverse kinematics [21], where joint constraints are enforced to ensure feasibility.

**Method Overview.** An overview of our proposed framework is illustrated in Fig. 2. We formulate dexterous grasp synthesis as the denoising process of our proposed  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph (Sec. III-A). Since spatial transformations are defined consistently across different objects and hand embodiments, this graph serves as a unified representation of a diffusion-based generative process (Sec. III-B), where noisy link poses are progressively denoised into valid grasps. Furthermore, we elaborate on the training and inference process of  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp in Sec. III-C, which enables efficient and accurate grasp synthesis both unconditionally (without external constraints) and conditionally (guided by an initial hand status).

### A. $\mathcal{T}(\mathcal{R}, \mathcal{O})$ Graph

Learning cross-embodiment dexterous grasping requires a unified representation to characterize the relation between dexterous hands and objects. Compared to joint values, which vary substantially across embodiments, spatial transformations offer a more robust and consistent representation cross different objects and embodiments. Therefore, we construct  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph based upon spatial SE(3) transformation of object patches and dexterous hand links. In  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph, nodes represent the geometric property

of each object patch or hand link, along with its spatial transformation with respect to the object global frame for a grasp. Edges encode the relative transformations between object patches and hand links.

Since SE(3) is a non-linear manifold, direct application of the standard diffusion process is not feasible. Therefore, we adopt the Lie group representation of SE(3) to model spatial transformations throughout our framework. According to [22], the Lie algebra of SE(3) is isomorphic to the vector space  $\mathbb{R}^6$ , enabling transformation between the SE(3) Lie group and vector space  $\mathbb{R}^6$  via exponential and logarithmic mapping operations as follows:

$$\mathbf{T} = \exp(\boldsymbol{\psi}), \quad \boldsymbol{\psi} = \log(\mathbf{T}), \quad (1)$$

where  $\boldsymbol{\psi} \in \mathbb{R}^6$ ,  $\mathbf{T} \in \text{SE}(3)$  denote rotation vector and homogeneous transformation matrix, respectively. Hence, we represent SE(3) transformations as 6D rotation vectors in our  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph during the diffusion process, and map them back to SE(3) after inference. In the following, we describe the detailed components of the  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph.

**Object Nodes.** Given an object point cloud, we first partition it into  $P$  patches using farthest point sampling. To capture the spatial and geometric characteristics of the object to be grasped, we employ the pretrained VQ-VAE [23] encoder for 3D fracture assembly in [24] to extract geometry tokens  $\{f_i^O\}_{i=1}^P$  from the normalized patches within a unit sphere. Object nodes are then constructed by concatenating patch center locations  $\{c_i^O\}_{i=1}^P$ , global object scale  $s^O$ , and corresponding patch geometry tokens as Eq. 2, where  $d^O$  denotes feature dimension of object nodes.

$$N^O = \{\text{concat}(c_i^O, s^O, f_i^O) | i = 1, \dots, P\} \in \mathbb{R}^{P \times d^O}, \quad (2)$$

**Link Nodes.** Link nodes encode both the geometric properties of each link and its spatial transformation with respect to the object's global frame for each grasp. To represent geometry, we apply the Basis Point Set (BPS) algorithm [25] to the link point clouds  $\{P_i^R\}_{i=1}^L$ . Specifically,  $B$  basis points are uniformly sampled within a unit sphere, and each link point cloud is normalized into the same unit sphere. The BPS feature of a link is then computed as a vector of length  $B$ , where each entry corresponds to the minimum distance between the point cloud and a basis point. This encoding produces fixed-length features invariant to the number of input points, and can be implemented efficiently in PyTorch, making it both scalable and generalizable across diverse dexterous hand embodiments.

Next, an MLP-based encoder embeds the BPS feature together with the link center  $\{c_i^R\}$  and scale  $\{s_i^R\}$  to obtain the geometry embedding for each link. The geometry embedding is further concatenated with the 6D rotation vector  $\{\psi_i^R\}$  of the link global pose with respect to the object frame in each state to formulate the link nodes within the graph:

$$b_i = \text{MLP}(\text{BPS}(P_i^R), c_i^R, s_i^R), \quad i = 1, \dots, L, \quad (3)$$

$$N^R = \{\text{concat}(\psi_i^R, b_i) | i = 1, \dots, L\} \in \mathbb{R}^{L \times d^R}, \quad (4)$$

**Edges.** Since the relative poses of object nodes are fixed for rigid body grasping, no edges are defined between object

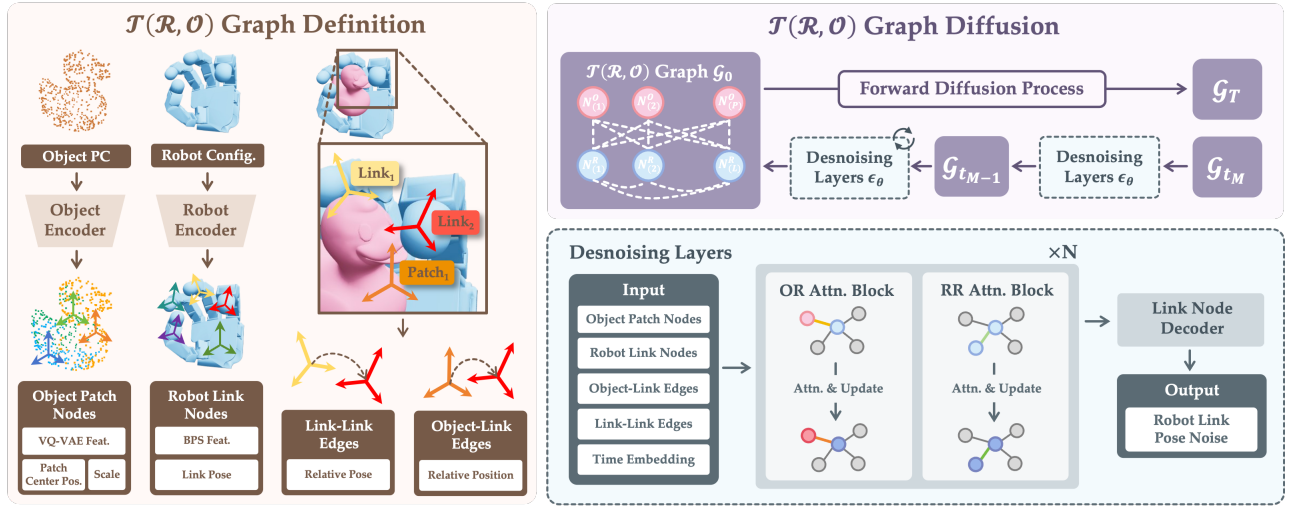


Fig. 2: Overview of  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp: We first define  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph to represent spatial transformations of robotic links and objects with auxiliary geometry information. Next, we introduce a graph diffusion model that enables both unconditioned and conditioned grasp synthesis.

nodes. Instead, we define edges only between object nodes and link nodes, as well as between link nodes. Notably, we establish complete connection for stronger spatial transformation constraints. Link-link edges are defined as the relevant spatial transformation between two link nodes. To ensure consistency across different objects, we assign all object transformation matrices  $\{T_i^O\}_{i=1}^P$  an identity rotation  $R^O = I$  and only use its patch center  $\{c_i^O\}_{i=1}^P$  as translation. The construction of graph edges is then formulated as:

$$\begin{aligned} T_{i,j}^{OR} &= (T_i^O)^{-1} T_j^R, & T_{j,k}^{RR} &= (T_j^R)^{-1} T_k^R, & (5) \\ E^{OR} &= \{\log(T_{i,j}^{OR})\}, & E^{RR} &= \{\log(T_{j,k}^{RR})\}, & (6) \end{aligned}$$

where  $E^{OR} \in \mathbb{R}^{6 \times P \times L}$  and  $E^{RR} \in \mathbb{R}^{6 \times L(L-1)/2}$  represent object-link edges and link-link edges, respectively.

Integrating spatial transformations between object patches and links with rich geometry embeddings, our proposed  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph is defined in Eq. 7, which serves as the foundation structure for subsequent graph diffusion.

$$\mathcal{G} = (N^O, N^R, E^{OR}, E^{RR}). \quad (7)$$

### B. $\mathcal{T}(\mathcal{R}, \mathcal{O})$ Graph Diffusion Model

We employ a graph diffusion model that operates on our proposed  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph, enabling the generation of diverse and precise grasps from random noise that can generalize across different objects and embodiments. Diffusion Probabilistic Models (DDPM) [26] generate data by gradually denoising Gaussian noise through a stochastic process, but require many inference steps. To achieve more efficient inference, we adopt the Denoising Diffusion Implicit Model (DDIM) [27], a variant of DDPM that accelerates sampling, allows controllable generation diversity, and naturally supports conditioned generation.

**Forward Process.** As illustrated in Sec. III-A, object nodes ( $N^O$ ) remain constant during both training and inference, whereas edges ( $E^{OR}, E^{RR}$ ) have to satisfy the spatial constraints between nodes. Consequently, noise injection and

denoising are applied exclusively to link poses  $\Psi^R = \{\psi_i^R\}$ . Given a noise-free graph  $\mathcal{G}_0 = (N_0^O, N_0^R, E_0^{OR}, E_0^{RR})$  derived from a valid grasp, we first gradually add Gaussian noise with variance schedule  $\beta_1 < \dots < \beta_T$  to the link poses as follows:

$$p(\Psi_t^R | \Psi_{t-1}^R) := \mathcal{N}(\sqrt{1 - \beta_t} \Psi_{t-1}^R, \beta_t I), \quad (8)$$

Similar to [26], link poses at arbitrary timestamp  $t$  can be directly sampled from the noisy-free graph (as illustrated in Eq. 8), where  $\alpha_t := 1 - \beta_t$  and  $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$ .

$$p(\Psi_t^R | \Psi_0^R) := \mathcal{N}(\sqrt{\bar{\alpha}_t} \Psi_0^R, (1 - \bar{\alpha}_t) I), \quad (9)$$

The noisy link nodes  $N_t^R$  are obtained by concatenating the noisy pose vectors with the original BPS embeddings. With the noisy link nodes at timestamp  $t$ , the corresponding edges are recomputed as the relative spatial transformation between graph nodes:

$$N_t^R = \{\text{concat}((\psi_{i,t}^R, b_i) | i = 1, \dots, L)\}, \quad (10)$$

$$E_t^{OR} = \text{Edge}(N^O, N_t^R), \quad E_t^{RR} = \text{Edge}(N_t^R), \quad (11)$$

where  $\text{Edge}(\cdot)$  denotes the operation that generates edges from the given graph nodes. Thus, the noisy  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph at step  $t$ ,  $\mathcal{G}_t = (N^O, N_t^R, E_t^{OR}, E_t^{RR})$ , is derived from the noisy-free graph  $\mathcal{G}_0$ .

**Denoising Model.** To predict the link node noise from a noisy  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph  $\mathcal{G}_t$ , we implement a transformer-based graph denoising model  $\epsilon_\theta(\mathcal{G}_t, t)$ . Each graph element is first encoded into token representation  $\{z^O, z^R, z^{OR}, z^{RR}\}$  for object nodes, link nodes, object-link and link-link edges, augmented with timestamp embedding  $\phi(t)$ . These tokens are then processed by a stack of graph denoising layers, where Object–Robot (OR) and Robot–Robot (RR) attention mechanisms iteratively update link nodes and all edges.

In both OR and RR attention, target nodes ( $X_{\text{tgt}}$ ) and connecting edges ( $E$ ) are refined by aggregating information from the corresponding source nodes ( $X_{\text{src}}$ ). For OR attention, we set  $X_{\text{src}} = z^O, X_{\text{tgt}} = z^R$  and  $E = E^{OR}$ ,

incorporating object geometry and spatial information into link updates. While for RR attention, we set  $X_{\text{src}} = X_{\text{tgt}} = z^R$  with  $E = E^{RR}$ , allowing the model to reason link-wise spatial relations. Formally, the attention process can be formulated as Eq. 12, where  $h_Q, h_K, h_V$  are learnable projections, and  $\text{attn}(\cdot)$  denotes attention-based aggregation. After all denoising layers update the graph, the output head aggregates the intermediate graph representations from all layers to predict the noise on link nodes, which is then used to recover clean link poses for valid grasps. A detailed figure of attention structure is provided in the Appendix.

$$\begin{aligned} Q &= h_Q(X_{\text{tgt}}) \in \mathbb{R}^{N_{\text{tgt}} \times d}, \\ K &= h_K(X_{\text{src}}) \in \mathbb{R}^{N_{\text{src}} \times d}, \\ V &= h_V(\text{concat}(X_{\text{tgt}}, X_{\text{src}}, E)) \in \mathbb{R}^{N_{\text{tgt}} \times N_{\text{src}} \times d}, \\ X_{\text{tgt}} &\leftarrow \text{MLP}(X_{\text{tgt}} + \text{attn}(Q, K, V)), \\ E &\leftarrow \text{MLP}(V, E), \end{aligned} \quad (12)$$

**Reverse Process.** Following the DDIM inference strategy,  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph Diffusion Model learns transitions  $p_\theta(G_{t_{i-1}}|G_{t_i})$  from random noisy  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph, where  $\theta$  denotes the learnable parameters, and  $\mathcal{T}^S = \{t_i\}_{i=1}^M$  is an ascending sequence of diffusion timesteps. According to [27], the distribution of link poses ( $\Psi^R$ ) in the reverse process can be formulated as:

$$\Psi_{t'}^R = \sqrt{\bar{\alpha}_{t'}} \hat{\Psi}_0^R + \sqrt{1 - \bar{\alpha}_{t'} - \sigma_t^2} \epsilon_\theta(G_t, t) + \lambda \sigma_t \epsilon_t, \quad (13)$$

$$\hat{\Psi}_0^R = \frac{1}{\sqrt{\bar{\alpha}_t}} (\Psi_t^R - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(G_t, t)), \quad (14)$$

where  $t, t' \in \mathcal{T}^S$  are diffusion timesteps with  $t' < t$ ,  $\epsilon_\theta(G_t, t)$  denotes the predicted noise from the denoiser,  $\sigma_t = \sqrt{(1 - \bar{\alpha}_{t'}) / (1 - \bar{\alpha}_t)} \sqrt{1 - \bar{\alpha}_t / \bar{\alpha}_{t'}}$  is the variance term defined in [27], and  $\lambda$  controls the level of random noise at each inference step. Once  $\Psi_t^R$  is denoised to  $\Psi_{t'}^R$ , link nodes and all edges are generated according to Eq. 11 to maintain node–edge spatial consistency. Through this iterative process, a complete noisy  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph can be denoised into a clean graph, which represents a valid grasp.

### C. Training and Inference

$\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp supports end-to-end training, allowing the model to learn spatial transformations across multiple object-embodiment pairs. During inference,  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  enables grasp synthesis either unconditionally or conditioned on an initial hand status. In contrast to  $\mathcal{D}(\mathcal{R}, \mathcal{O})$  [1] Grasp, it removes the reliance on an initial reference hand status, thereby preventing any suboptimal or infeasible initial poses from degrading grasp quality. Moreover,  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  allows dynamic adjustment of DDIM inference steps and incorporates the GPU-based PyroKi [21] optimization, resulting in a much more light-weighted and efficient inference process.

**Loss Function.** The training objective of diffusion models is to maximize the data log-likelihood via a variational lower bound, which can be simplified to minimize the noise prediction error according to [26]. Therefore, the training loss of the denoising model can be simplified as:

$$L := \gamma_p \|\epsilon^p - \epsilon_\theta^p(G_t, t)\|^2 + \gamma_r \|\epsilon^r - \epsilon_\theta^r(G_t, t)\|^2, \quad (15)$$

where  $\epsilon^p, \epsilon^r$  represent position and rotation noise on link poses, and  $\gamma_p, \gamma_r$  denote corresponding loss weight.

**Unconditioned Inference.** Benefiting from its diffusion-based structure,  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  can generate precise grasps without relying on an initial pose, thereby avoiding the risk of poor initialization degrading grasp quality as in  $\mathcal{D}(\mathcal{R}, \mathcal{O})$ . For efficient and adaptive inference, we adopt the DDIM [27] sampling scheme. We sample  $M$  diffusion timestamps  $\mathcal{T}^S = \{t_i\}_{i=1}^M$  from the full schedule  $\{1, \dots, T\}$  and apply the reverse process as described in Sec. III-B.

**Conditioned Inference.** Grasp synthesis under specific conditions are important in real-world deployment. Based on guidance strategies in diffusion inference [28]–[31],  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  naturally supports grasp synthesis conditioned on diverse grasp configurations, such as initial status and contact points. For example, given an initial hand status, we first inject Gaussian noise  $\epsilon$  to the initial link poses to an intermediate timestamp  $t^*$  corresponding to the noise level  $\bar{\alpha}_{t^*}$  as Eq. 16, where  $\Psi_{\text{init}}^R$  denotes link nodes initialized from given hand status.

$$\Psi_{t^*, \text{init}}^R = \sqrt{\bar{\alpha}_{t^*}} \Psi_{\text{init}}^R + \sqrt{1 - \bar{\alpha}_{t^*}} \epsilon, \quad (16)$$

Starting from the intermediate status  $\Psi_{t^*, \text{init}}^R$ , we employ the gradient of geodesic distance on  $\text{SO}(3)$  to guide the predicted palm orientation toward the palm orientation of the initial status  $R_{\text{init}} \in \text{SO}(3)$ :

$$L_{\text{geo}}(R_1, R_2) = \arccos\left(\frac{\text{tr}(R_1 R_2^T) - 1}{2}\right), \quad (17)$$

$$\hat{\Psi}_0^R \leftarrow \hat{\Psi}_0^R - s(t) \nabla_{\hat{\Psi}_0^R} L_{\text{geo}}(R_{\text{init}}, \hat{\Psi}_{0, \text{palm}}^R), \quad (18)$$

$$\epsilon_\theta(G_t, t) = \frac{1}{\sqrt{1 - \bar{\alpha}_t}} (\Psi_t^R - \sqrt{\bar{\alpha}_t} \hat{\Psi}_0^R), \quad (19)$$

where  $s(t) \in [0, g_s]$  controls the guidance strength for each sampling step and  $\hat{\Psi}_{0, \text{palm}}^R$  denotes the palm rotation matrix extracted from the denoised estimation of link poses at the current timestep. Through gradient-based guidance, the generated grasps are encouraged to follow the initial hand status, while ensuring that the hand–object transformations correspond to a reliable grasp. Notably, such conditioned grasp synthesis further empowers  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp to support closed-loop manipulation in dynamic scenarios.

**Pyroki Joint Optimization.** After obtaining the clean link poses for a potential grasp, we formulate a inverse kinematics problem to compute the joint values  $\mathbf{q}$  that best realize these predicted poses:

$$\min_{\mathbf{q}} \sum_{i=1}^L \|\text{FK}_i(\mathbf{q}) - \psi_i^R\|_2, \quad \text{s.t. } \mathbf{q} \in [\mathbf{q}_{\text{min}}, \mathbf{q}_{\text{max}}]. \quad (20)$$

where  $\text{FK}_i(\cdot)$  and  $\psi_i^R$  denote the forward kinematics and predicted pose of link  $i$ , respectively, and  $[\mathbf{q}_{\text{min}}, \mathbf{q}_{\text{max}}]$  denotes the joint limits. To achieve efficient IK, we integrate Pyroki [21], a highly efficient, modular, and extensive toolkit for solving kinematics optimization problems, into our inference pipeline. Once the scripts are pre-compiled in JAX [32], IK process with Pyroki toolkit can achieve real-time performance at over 30 FPS on our tested dexterous hands (i.e., Allegro, Barrett and Shadowhand).

Method	Success Rate (%) $\uparrow$				Efficiency (sec.) $\downarrow$			Diversity (rad.) $\uparrow$		
	Barrett	Allegro	ShadowHand	Avg.	Barrett	Allegro	ShadowHand	Barrett	Allegro	ShadowHand
DFC [3]	86.30	76.21	58.80	73.77	>1800	>1800	>1800	0.532	<b>0.454</b>	0.435
GenDexGrasp [33]	67.00	51.00	54.20	57.40	14.67	25.10	19.34	0.488	0.389	0.318
ManiFM [8]	-	42.60	-	42.60	-	9.07	-	-	0.288	-
DRO-Grasp [1]	87.30	92.30	83.00	87.53	0.49	0.47	0.98	0.513	0.397	<b>0.441</b>
<b>TRO-Grasp (cond.)</b>	89.80	93.70	93.50	92.33	<b>0.19</b>	<b>0.21</b>	<b>0.22</b>	<b>0.561</b>	0.430	0.326
<b>TRO-Grasp (uncond.)</b>	<b>91.90</b>	<b>94.00</b>	<b>98.60</b>	<b>94.83</b>	0.23	0.23	0.25	0.542	0.370	0.292

TABLE I: Overall performance and inference efficiency comparison with baselines: For ManiFM [8], evaluation is limited to the available Allegro hand checkpoint.

#### IV. EXPERIMENTS

In this section, we perform a series of experiments aimed at addressing the following questions (Q1-Q7):

Q1: How successful and diverse is the  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  unconditioned grasp synthesis? (Sec. IV-B)

Q2: How does the initial status influence the model performance in conditioned grasp synthesis? (Sec. IV-C)

Q3: Is  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  representation generalizable across multiple dexterous hand embodiments? (Sec. IV-B)

Q4: Is  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  representation robust to partial and incomplete object observations? (Sec. IV-B)

Q5: How computationally efficient is  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp during training and inference? (Sec. IV-D)

Q6: Does  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp support closed-loop dexterous manipulation in dynamic environments? (Sec. IV-E)

Q7: How does  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  perform in real-world? (Sec. IV-F)

##### A. Experimental Setting

**Dataset.** Following  $\mathcal{D}(\mathcal{R}, \mathcal{O})$  Grasp, we evaluate on a filtered subset of CMapDataset [33], which includes dexterous hand embodiments with 3–5 fingers (Barrett, Allegro and ShadowHand). The training dataset consists of 14,011 grasp demonstrations on 48 objects, while the testing dataset evaluates 100 grasps per object across 10 unseen objects. For every object, we sample 65,536 points from the mesh file, and randomly sample 512 points in each training iteration.

**Evaluation Metrics.** Similar to  $\mathcal{D}(\mathcal{R}, \mathcal{O})$  Grasp, we report the following three metrics:

- *Success Rate*: a grasp is considered successful in Isaac Gym if the object displacement remains below 2 cm under six directional perturbations.
- *Efficiency*: measured as the computational time per grasp, including both DDIM inference and Pyroki IK process.
- *Diversity*: quantified as the standard deviation of joint values across all successful grasps.

**Implementation Details.** For  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph construction, each object is encoded into  $P = 25$  nodes, and link nodes are zero-padded to  $L = 25$  across all embodiments to enable parallel computing. In the  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph Diffusion process, we adopt a linear scheduler for the noise variance and set the total number of diffusion steps to  $T = 1,000$ . During inference, we sample  $M = 20$  steps and set  $\lambda = 0.2$  to inject random noise for more diverse grasp synthesis. For the loss function, we set  $\gamma_p = \gamma_r = 1.0$  and apply a sine-based guidance strength  $s(t)$  for conditioned grasp synthesis. All

Method	Success Rate (%) $\uparrow$			Diversity (rad.) $\uparrow$		
	Barrett	Allegro	ShadowHand	Barrett	Allegro	ShadowHand
Single	91.60	92.30	96.00	0.528	0.372	0.281
Multi	<b>91.90</b>	<b>94.00</b>	<b>98.60</b>	0.542	0.370	<b>0.292</b>
Partial	86.80	93.40	94.60	<b>0.548</b>	<b>0.391</b>	0.290

TABLE II: Robust and generalizable performance of  $\mathcal{T}(\mathcal{R}, \mathcal{O})$ : Single, Multi, and Partial denote training and evaluation on single-hand, multi-hand, and partial object point cloud datasets, respectively.

$\mathcal{T}(\mathcal{R}, \mathcal{O})$  models are trained for 300 epochs on an NVIDIA A100 40GB GPU with Adam optimizer and step learning rate scheduler. Additional implementation details and parameter settings are provided in the Appendix.

##### B. Unconditioned grasp synthesis

**Overall Performance.** As illustrated in Tab. I, our method substantially outperforms all baselines in grasping success rate, achieving an average increase of 7.3% over  $\mathcal{D}(\mathcal{R}, \mathcal{O})$  and up to 15.6% on ShadowHand. Fig. 3 shows diverse and accurate grasp on unseen objects across Allegro, Barrett and Shadowhand. In Fig. 4, we further visualize the diffusion inference process of  $\mathcal{T}(\mathcal{R}, \mathcal{O})$ , where noisy link poses gradually converge to form a valid grasp. In addition,  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  accelerates inference by  $3\times$  in average compared to  $\mathcal{D}(\mathcal{R}, \mathcal{O})$  Grasp, owing to more efficient  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  representation and DDIM-based diffusion inference, along with the real-time Pyroki IK toolkit.

**Cross-embodiment Generalizability.** To validate the generalizability of the  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  representation, we conduct grasp synthesis across multiple embodiments. As shown in Tab. II,  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp achieves superior performance through multi-embodiment training, demonstrating that the  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph representation, together with the graph diffusion model, generalizes effectively across diverse dexterous hands. These results further suggests the feasibility of scaling  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp into a large-scale foundation model for dexterous grasping by training it on extensive multi-hand grasp demonstrations.

**Partial Observation.** In practical scenarios, object observations are often incomplete due to occlusions or limited viewpoints. To assess the robustness of  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp under such conditions, we conduct both training and testing with one-side 50% point cloud inputs for all objects. As shown in Tab. II,  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp maintains high performance despite partial observations, demonstrating its robustness for real-world deployment with incomplete observations.



Fig. 3: Diverse and accurate unconditioned grasp synthesis of  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  across multiple dexterous hand embodiments.

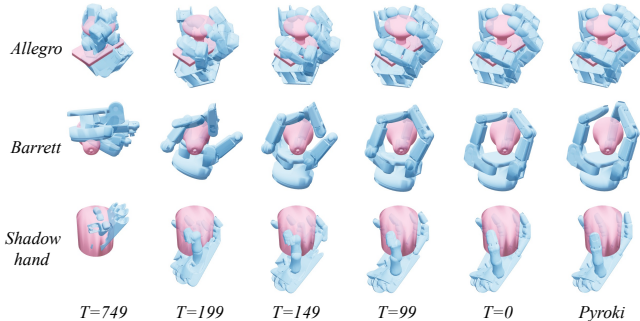


Fig. 4: Visualization of unconditioned DDIM inference across timesteps. “Pyroki” indicates the grasp after IK with the Pyroki [21] toolkit.

### C. Conditioned grasp synthesis

Conditioned grasp synthesis is crucial for real-world deployment scenarios, such as grasping objects from a table or retrieving them from a shelf. In this experiment, we incorporate gradient-based guidance into  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp, enabling grasp synthesis conditioned on the given initial hand status. As indicated in Tab. I, conditioned grasp synthesis trades a slight drop in success rate for higher grasp diversity and faster inference. Visualization in Fig. 5 demonstrates that the generated grasps align well with the palm directions of given hand status. A diverse set of initialization grasping poses increases the risk of generating infeasible grasps, while simultaneously promoting grasp diversity. In addition, starting the diffusion inference process at an intermediate timestep enables faster inference at approximately 5 FPS, indicating the potential of incorporating  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  framework into closed-loop dexterous hand manipulation (Sec. IV-E).

### D. Efficiency Analysis

In this section, we provide a comprehensive efficiency analysis of  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp in both training and inference on an NVIDIA A100 40GB GPU. To evaluate memory usage, we fix the batch size to 4 in training and 25 in inference, and report the average memory usage per sample for fair comparison. To evaluate throughput, we increase the batch size of all models to the maximum that fits into 40GB memory. As indicated in Tab. III,  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  requires only 43% training memory and 21% inference memory compared to  $\mathcal{D}(\mathcal{R}, \mathcal{O})$ , while possessing  $2\times$  more model parameters with stronger representation capability. In addition,  $\mathcal{T}(\mathcal{R}, \mathcal{O})$




Fig. 5: grasp synthesis conditioned on initial hand status: Red arrow denotes the palm direction of pose guidance.

Method	DRO	TRO (cond.)	TRO (uncond.)
Infer Mem. (GB/sample)	0.38	0.08	0.08
Infer throughput (sample/s)	5	50	32
Model Param. (M)	14.00	28.06	28.06
Train Mem. (GB/sample)	4.10	1.77	1.77
Train throughput (sample/s)	7	22	22

TABLE III: Training and Inference Efficiency Comparison.

demonstrates much higher throughput and inference speed owing to its lightweight architecture, DDIM-based inference strategy, and Pyroki-based IK optimization. As a result, it can be deployed on real-world robotic platforms with limited computational resources.

### E. Closed-Loop Grasping Evaluation

By leveraging its fast inference speed and the ability to perform grasp synthesis conditioned on previous grasp poses,  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp achieves conditioned synthesis at an average of 5 FPS (Sec. IV-C), making closed-loop grasping feasible. Specifically, given a discrete time sequence  $\{0, T, \dots, nT\}$ ,  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  uses the pose from the previous timestamp  $\Psi_{(i-1)T}^R$  as guidance to predict all link poses at the current timestamp  $\Psi_{iT}^R$ . We conduct closed-loop grasping evaluation with the Allegro hand on 10 unseen objects. The time interval is set to 0.25s with a total of 30 evaluation steps. At each interval, the object either translates at a constant velocity (5 cm/s) or undergoes random perturbations involve translational noise of up to 1.25 cm and rotational noise up to  $30^\circ$  per step. As shown in Tab. IV, closed-loop  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp achieves average success rates of 89.3% under constant-velocity motion and 91.0% under random perturbations, while open-loop approaches completely fail (0% success rate). These results demonstrate the capability of  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp as a general framework for closed-loop dexterous grasping across diverse dynamic scenarios.

### F. Real-Robot Experiments

We conduct real-world experiments with a uFactory xArm7 robot equipped with a Robotera XHand and LEAP Hand (as illustrated in Fig. 6). Object point cloud observations are obtained using an overhead Intel RealSense D435 camera. We train our model on 7,800 grasp demonstrations and evaluate 100 grasps on 10 previously unseen objects for each dexterous hand. During inference, we leverage FoundationPose [34] to track the object pose.  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp achieves an average success rate of 91% and 90%

Apple		Camera		Cylinder		Door Knob		Rubber Duck	
Const.	Rand.	Const.	Rand.	Const.	Rand.	Const.	Rand.	Const.	Rand.
100.0	100.0	76.7	93.3	93.3	96.7	96.7	90.0	96.7	93.3
Water Bottle		Baseball		Pear		Meat Can		Soup Can	
Const.	Rand.	Const.	Rand.	Const.	Rand.	Const.	Rand.	Const.	Rand.
56.7	70.0	100.0	96.7	93.3	83.3	93.3	90.0	86.7	96.7

TABLE IV: Success rate of closed-loop grasp synthesis of Allegro hand on 10 unseen objects.




Fig. 6: Real-world experiment setting.

on XHand and LEAP Hand respectively, while enabling closed-loop grasping in dynamic environments. Additional experimental details, results, and videos are available in the appendix and on the project website <https://nus-lins-lab.github.io/trograsweb/>.

## V. CONCLUSION

In this paper, we tackled the challenge of efficient and generalizable dexterous grasp synthesis. We highlighted the limitations of robot-centric, object-centric, and interaction-centric representations, and introduced the  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Graph—a unified, memory-efficient representation of object–robot interactions that generalizes across embodiments. Built upon this representation, we proposed a transformer-based graph diffusion model that enables both conditioned and unconditioned grasp generation, as well as closed-loop manipulation. Extensive experiments in simulation and on hardware show that  $\mathcal{T}(\mathcal{R}, \mathcal{O})$  Grasp outperforms all baselines on success rate with high efficiency and closed-loop manipulation. This work not only advances grasping methods but also lays the groundwork for large-scale dexterous grasping foundation models, with future directions in broader manipulation tasks, multimodal feedback, and large-scale pretraining.

## REFERENCES

- [1] Z. Wei, Z. Xu, J. Guo, Y. Hou, C. Gao, Z. Cai, J. Luo, and L. Shao, “D(r,o) grasp: A unified representation of robot and object interaction for cross-embodiment dexterous grasping,” *arXiv preprint arXiv:2410.01702*, 2024.
- [2] W. Wan, H. Geng, Y. Liu, Z. Shan, Y. Yang, L. Yi, and H. Wang, “Unidexgrasp++: Improving dexterous grasping policy learning via geometry-aware curriculum and iterative generalist-specialist learning,” in *ICCV*, 2023, pp. 3891–3902.
- [3] T. Liu, Z. Liu, Z. Jiao, Y. Zhu, and S.-C. Zhu, “Synthesizing diverse and physically stable grasps with arbitrary hand structures using differentiable force closure estimator,” *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 470–477, 2021.
- [4] H. Jiang, S. Liu, J. Wang, and X. Wang, “Hand-object contact consistency reasoning for human grasps generation,” in *ICCV*, 2021.
- [5] L. Huang, H. Zhang, Z. Wu, S. Christen, and J. Song, “Fungrasp: functional grasping for diverse dexterous hands,” *IEEE Robotics and Automation Letters*, 2025.
- [6] Y. Liu, Y. Yang, Y. Wang, X. Wu, J. Wang, Y. Yao, S. Schwertfeger, S. Yang, W. Wang, J. Yu *et al.*, “Realdex: Towards human-like grasping for robotic dexterous hand,” *arXiv preprint arXiv:2402.13853*, 2024.

- [7] L. Shao, F. Ferreira, M. Jorda, V. Nambiar, J. Luo, E. Solowjow, J. A. Ojea, O. Khatib, and J. Bohg, “Unigrasp: Learning a unified model to grasp with multifingered robotic hands,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2286–2293, 2020.
- [8] Z. Xu, C. Gao, Z. Liu, G. Yang, C. Tie, H. Zheng, H. Zhou, W. Peng, D. Wang, T. Chen, Z. Yu, and L. Shao, “Manifold foundation model for general-purpose robotic manipulation of contact synthesis with arbitrary objects and robots,” 2024.
- [9] Y. Xu, W. Wan, J. Zhang, H. Liu, Z. Shan, H. Shen, R. Wang, H. Geng, Y. Weng, J. Chen *et al.*, “Unidexgrasp: Universal robotic dexterous grasping via learning diverse proposal generation and goal-conditioned policy,” in *CVPR*, 2023, pp. 4737–4746.
- [10] J. Varley, J. Weisz, J. Weiss, and P. Allen, “Generating multi-fingered robotic grasps via deep learning,” in *IROS*, 2015, pp. 4415–4420.
- [11] M. Attarian, M. A. Asif, J. Liu, R. Hari, A. Garg, I. Gilitschenski, and J. Tompson, “Geometry matching for multi-embodiment grasping,” in *Conference on Robot Learning*. PMLR, 2023, pp. 1242–1256.
- [12] H.-S. Fang, H. Yan, Z. Tang, H. Fang, C. Wang, and C. Lu, “Anydex-grasp: General dexterous grasping for different hands with human-level learning efficiency,” *arXiv preprint arXiv:2502.16420*, 2025.
- [13] F. Zhao, D. Tsetserukou, and Q. Liu, “Graingrasp: Dexterous grasp generation with fine-grained contact guidance,” in *ICRA*, 2024.
- [14] T. Zhu, R. Wu, X. Lin, and Y. Sun, “Toward human-like grasp: Dexterous grasping via semantic representation of object-hand,” in *ICCV*, 2021, pp. 15 721–15 731.
- [15] S. Liu, Y. Zhou, J. Yang, S. Gupta, and S. Wang, “Contactgen: Generative contact modeling for grasp generation,” in *ICCV*, 2023, pp. 20 609–20 620.
- [16] K. Sohn, H. Lee, and X. Yan, “Learning structured output representation using deep conditional generative models,” *NeurIPS*, 2015.
- [17] M. Sundermeyer, A. Mousavian, R. Triebel, and D. Fox, “Contact-grasnet: Efficient 6-dof grasp generation in cluttered scenes,” 2021.
- [18] J. Hager, R. Bauer, M. Toussaint, and J. Mainprice, “Graspme - grasp manifold estimator,” in *RO-MAN*, 2021, pp. 626–632.
- [19] T. Weng, D. Held, F. Meier, and M. Mukadam, “Neural grasp distance fields for robot manipulation,” in *ICRA*, 2023, pp. 1814–1821.
- [20] J. Urain, N. Funk, J. Peters, and G. Chalahatzaki, “Se(3)-diffusionfields: Learning smooth cost functions for joint grasp and motion optimization through diffusion,” *ICRA*, 2023.
- [21] C. M. Kim\*, B. Yi\*, H. Choi, Y. Ma, K. Goldberg, and A. Kanazawa, “Pyroki: A modular toolkit for robot kinematic optimization,” 2025.
- [22] J. Sola, J. Deray, and D. Atchuthan, “A micro lie theory for state estimation in robotics,” *arXiv preprint arXiv:1812.01537*, 2018.
- [23] A. Van Den Oord, O. Vinyals *et al.*, “Neural discrete representation learning,” *NeurIPS*, vol. 30, 2017.
- [24] Z. Wang, J. Chen, and Y. Furukawa, “Puzzlefusion++: Auto-agglomerative 3d fracture assembly by denoise and verify,” *arXiv preprint arXiv:2406.00259*, 2024.
- [25] S. Prokudin, C. Lassner, and J. Romero, “Efficient learning on point clouds with basis point sets,” in *ICCV*, 2019, pp. 4332–4341.
- [26] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.
- [27] J. Song, C. Meng, and S. Ermon, “Denoising diffusion implicit models,” *arXiv preprint arXiv:2010.02502*, 2020.
- [28] J. Ho and T. Salimans, “Classifier-free diffusion guidance,” *arXiv preprint arXiv:2207.12598*, 2022.
- [29] A. Bansal, H.-M. Chu, A. Schwarzschild, S. Sengupta, M. Goldblum, J. Geiping, and T. Goldstein, “Universal guidance for diffusion models,” in *CVPR*, 2023, pp. 843–852.
- [30] Y. Guo, H. Yuan, Y. Yang, M. Chen, and M. Wang, “Gradient guidance for diffusion models: An optimization perspective,” *NeurIPS*, vol. 37, pp. 90 736–90 770, 2024.
- [31] C. Meng, Y. He, Y. Song, J. Song, J. Wu, J.-Y. Zhu, and S. Ermon, “SDEdit: Guided image synthesis and editing with stochastic differential equations,” in *ICLR*, 2022.
- [32] J. Bradbury, R. Frostig, P. Hawkins *et al.*, “JAX: Composable transformations of Python+NumPy programs,” <https://github.com/google/jax>, 2018, version 0.3.13, accessed 2026-03-05.
- [33] P. Li, T. Liu, Y. Li, Y. Geng, Y. Zhu, Y. Yang, and S. Huang, “Gendexgrasp: Generalizable dexterous grasping,” in *ICRA*, 2023, pp. 8068–8074.
- [34] B. Wen, W. Yang, J. Kautz, and S. Birchfield, “Foundationpose: Unified 6d pose estimation and tracking of novel objects,” in *CVPR*, 2024, pp. 17 868–17 879.