

Dexterous Non-prehensile Manipulation using Environment Guided Diffusion Models

Yao, Shaoxiong; Xing, Chengy; Jha, Devesh K.; Romeres, Diego

TR2026-073 June 04, 2026

Abstract

Humans naturally use non-prehensile actions, such as sliding and poking, to manipulate objects that are not immediately graspable. Enabling robots with a similar capability could significantly broaden the range of environments in which they can operate effectively. In this work, we propose a novel framework for generating non-prehensile manipulation actions conditioned on both object and environmental geometry. Our method builds on a guided diffusion formulation that adapts manipulation behaviors to the environment by encoding object and environment contacts as differentiable loss terms during training-free sampling. To train and evaluate this approach, we develop a simulation-based pipeline for collecting diverse manipulation behaviors across objects with different geometries and contact conditions. Experiments in several challenging environments show that the learned diffusion model adapts effectively to environmental context and achieves around $2\times$ higher success rates compared to baseline approach.

*ICRA Workshop on Dexterity with Multifingered Hands: Hardware, Sensing, and Skills
2026*

© 2026 MERL. This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Dexterous Non-prehensile Manipulation using Environment Guided Diffusion Models

Shaoxiong Yao^{1,2}, Chengyi Xing¹, Devesh K. Jha^{*,3}, Diego Romeres^{*,1}

Abstract—Humans naturally use non-prehensile actions, such as sliding and poking, to manipulate objects that are not immediately graspable. Enabling robots with a similar capability could significantly broaden the range of environments in which they can operate effectively. In this work, we propose a novel framework for generating non-prehensile manipulation actions conditioned on both object and environmental geometry. Our method builds on a guided diffusion formulation that adapts manipulation behaviors to the environment by encoding object and environment contacts as differentiable loss terms during training-free sampling. To train and evaluate this approach, we develop a simulation-based pipeline for collecting diverse manipulation behaviors across objects with different geometries and contact conditions. Experiments in several challenging environments show that the learned diffusion model adapts effectively to environmental context and achieves around $2\times$ higher success rates compared to baseline approach.

I. INTRODUCTION

Humans naturally use diverse manipulation strategies to handle objects in challenging environments. These strategies extend beyond grasping to include non-prehensile behaviors such as sliding, poking that leveraging environmental contact [14]. For example, when picking up a thin object from a tabletop, a person may first slide it toward the table edge to create a configuration easier for grasp as shown in Fig. 1.

Developing comparable capabilities on robotic hardware remains challenging, especially for dexterous hands with many degrees of freedom. While prior work has explored the automatic generation of such behaviors for simpler end-effectors such as parallel grippers [3], extending this ability to dexterous robotic hands remains difficult. Learning-based methods have shown strong promise for automatically generating dexterous robotic grasping actions [22]. Among them, diffusion models [8] are particularly attractive because they can represent complex and diverse action distributions, making them well suited for modeling the multiple valid behaviors that may arise in manipulation. However, existing methods primarily focus on grasp synthesis and are not capable of generating diverse manipulation behaviors under varying extrinsic contact conditions.

¹Mitsubishi Electric Research Laboratories. ²Department of Computer Science, Grainger College of Engineering, UIUC. ³Samsung Research America. Shaoxiong Yao's email: syao16@illinois.edu.

*D. K. Jha and D. Romeres contributed equally to advising.

D. K. Jha was with Mitsubishi Electric Research Laboratories when this work was conducted and is currently with Samsung Research America. D. Romeres is presently with Amazon IRG: dromeres@amazon.com.

This work was done during an internship at Mitsubishi Electric Research Laboratories.




Fig. 1. **Non-prehensile manipulation conditioned on the environment.** Humans naturally use non-prehensile actions, such as sliding, to make initially ungraspable objects graspable (top). We develop a framework that conditions on object and environmental geometry to generate such non-prehensile manipulations (middle). By combining the generated manipulation configurations with grasping actions and motion planning, our method can make an initially ungraspable object into a graspable one (bottom).

In this paper, we propose a novel method for generating non-prehensile manipulation behaviors conditioned on environmental contact. Our approach is based on a guided diffusion formulation that captures non-prehensile manipulation behaviors. We first develop a simulation-based data generation pipeline to produce diverse manipulation behaviors across objects with varying geometries and contact configurations. We then train a diffusion model to jointly represent these behaviors and evaluate it across multiple challenging scenarios. Our results show that the proposed approach substantially outperforms a model trained only on grasping behaviors. We also conduct ablation studies to highlight the contribution of key design choices in our guided diffusion pipeline.

II. RELATED WORK

Early studies have shown that humans use diverse manipulation behaviors when interacting with objects in different environmental conditions [6, 14]. Inspired by this capability, prior work has extensively explored model-based approaches, developing motion planning methods that can generate com-

plex manipulation behaviors [2, 3, 25]. These methods are able to synthesize manipulation sequences that combine different behaviors, such as grasping and non-prehensile interactions. However, they have not yet scaled effectively to high-DoF dexterous hands, where the large action space makes planning much more difficult. For dexterous hands, existing model-based optimization methods remain largely limited to grasping and pick-up tasks [1, 29].

More recently, learning-based methods have achieved substantial progress in dexterous grasp generation for objects with complex geometries. Prior works, including GenDex-Grasp [10], DexGraspNet [20], and $\mathcal{D}(\mathcal{R}, \mathcal{O})$ [21], have shown strong capability in generating dexterous grasps across diverse objects. Beyond grasp generation, simulation-to-real transfer and reinforcement learning have also enabled dexterous in-hand manipulation skills such as in-hand rotation [16] and lid twisting [12]. Among these methods, diffusion models have become particularly promising for dexterous hand generation. As a class of probabilistic generative models, they are appealing because guidance formulations [5, 7] make it possible to adapt generated robot configurations to task-specific objectives and constraints. Recent approaches such as UGG [13] and DexGrasp-Anything [26] show that diffusion models can handle physical constraints, while EvolvingGrasping [28] incorporates human preference alignment into the generation process. Diffusion models have also been integrated with teleoperation systems [24]. Very recent works explored diffusion models for non-prehensile manipulation, but these approaches either rely on inefficient sample-and-select pipelines [11] or fit the model to specific task [9].

Despite this progress, relatively little prior work has studied dexterous manipulation generation conditioned on environmental information, where the goal is to produce behaviors beyond grasping by explicitly leveraging environmental contact. In this paper, we develop a diffusion-based framework with guided sampling to generate non-prehensile manipulation behaviors conditioned on both object geometry and the surrounding environment.

III. METHOD

We represent the robot hand configuration as $q \in \mathbb{R}^{9+d}$, where the first 9 dimensions encode the hand translation and 6D rotation representation [27], and the remaining d dimensions specify the joint angles of the dexterous hand. We assume the hand is defined by a URDF model, with each link geometry represented as a triangle mesh. The inputs also include the object and environment geometry, represented by triangle meshes \mathcal{M}_{obj} and \mathcal{M}_{env} , as well as signed distance fields $\phi_{\text{obj}}(\cdot)$ and $\phi_{\text{env}}(\cdot)$, respectively.

An overview of the proposed approach is shown in Figure 2. For simplicity, we assume that a high-level planner provides a feasible desired object motion. Given this motion, the goal of our system is to generate a hand configuration q that can establish appropriate contact with the object to achieve the desired manipulation objective.

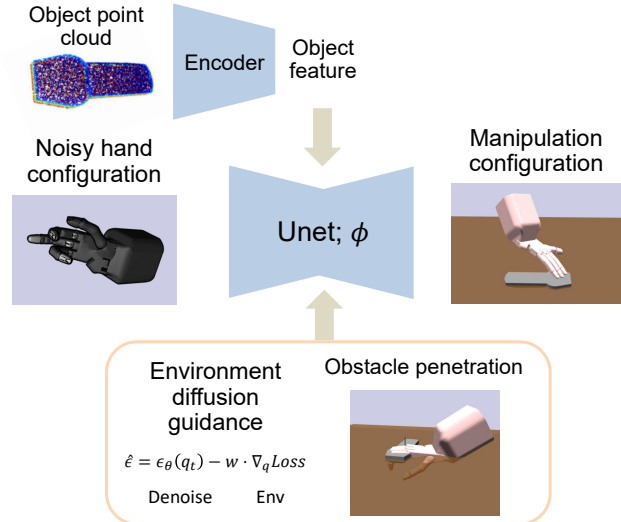


Fig. 2. **Method overview.** We proposed a guided diffusion framework that is capable to generate non-prehensile manipulation behaviors conditioned on environment geometry.




Fig. 3. **Example hand configurations.** This figure shows diverse non-prehensile manipulation configurations from our dataset, highlighting different contact strategies for objects with varying geometries.

A. Dexterous Manipulation Dataset Generation Procedure

The first stage of our framework constructs a simulation dataset of diverse manipulation behaviors across objects with varying geometries. Our data generation pipeline builds on the optimization-based grasp generation procedure of DexGraspNet [20] and extends it to generate non-prehensile manipulation actions.

For non-prehensile actions, we densely sample hand-surface

keypoints $\{x_i(q)\}_{i=1}^N$, whose positions are computed via forward kinematics, and initialize a set of pre-contact hand configurations. For each initialization, we identify nearby object surface points and optimize the hand configuration to encourage contact with the object. Specifically, we minimize

$$E(q) = \sum_{i,j} \|x_i(q) - p_j\|^2 + \sum_i \phi_{\text{obj}}(x_i(q))^2, \quad (1)$$

where $x_i(q)$ is the position of the i -th hand keypoint under configuration q , p_j denotes sampled object surface points, and $\phi_{\text{obj}}(\cdot)$ is the signed distance field of the object. We optimize this objective using gradient descent and stop once the fingertips make contact with the object. This objective encourages the hand to form meaningful contact with the object surface, as shown in Fig. 3.

B. Object-conditioned Diffusion Model Formulation

We model dexterous hand manipulation generation as a conditional diffusion process that predicts a hand configuration from object geometry. We sample a point cloud O from the object mesh \mathcal{M}_{obj} and learn a conditional distribution $P(q | O)$ over non-prehensile manipulation configurations. Given a clean configuration q_0 , we obtain a noisy sample $q_t = q_0 + \epsilon$ through the forward diffusion process and train a network $\epsilon_\theta(q_t, O, t)$ to predict the injected noise conditioned on O and the diffusion step t . The training loss is

$$\mathcal{L} = \mathbb{E}_{q_0, O, \epsilon, t} \left[\|\epsilon - \epsilon_\theta(q_t, O, t)\|^2 \right]. \quad (2)$$

At inference time, we use a deterministic DDIM sampler [19] to iteratively update a noisy sample q_t at diffusion step t to a less noisy sample q_{t-1} :

$$q_{t-1} = \frac{\sqrt{\alpha_{t-1}}}{\sqrt{\alpha_t}} (q_t - \sqrt{1 - \alpha_t} \epsilon_\theta(q_t, O, t)), \quad (3)$$

where α_t is the noise schedule at step t . After the reverse diffusion process, the final sample q_0 is distributed according to the original clean data distribution. At this stage, the formulation conditions only on object geometry. In the following section, we show how environmental information can be incorporated into the sampling process.

C. Guided-diffusion Sampling using Environment Conditions

After training a diffusion model to generate diverse manipulation behaviors, we introduce a guided sampling method that steers the reverse process with differentiable task losses. Following prior works [4, 23], we define the target distribution induced by a differentiable loss $L(q)$ as

$$F(q) = Z \exp(-L(q)), \quad (4)$$

where Z is a normalizing constant, so that configurations with lower loss are more likely under the target distribution.

Following the spherical Gaussian constraint formulation of [4], let $\hat{q}_0(q_t)$ be the estimated clean sample corresponding




Fig. 4. **Projection step.** The left figure shows a diffusion-generated hand configuration that collides with the object, while the right figure shows the configuration after projection with no penetration.

to q_t . We then guide the DDIM update using the normalized gradient of the loss:

$$q_{t-1} = \text{DDIM}(q_t, \epsilon_\theta(q_t, t)) - s \frac{\nabla_{q_t} L(\hat{q}_0(q_t))}{\|\nabla_{q_t} L(\hat{q}_0(q_t))\|}, \quad (5)$$

where s controls the guidance strength. We omit several constant terms from the original derivation for clarity. This guidance mechanism biases sampling toward manipulation behaviors that better satisfy the desired physical objectives.

a) *Non-penetration guidance:* We define a differentiable non-penetration loss during generation, similar to DexGrasp-Anything [26], to discourage collisions with the robot hand. This includes self-collision, penetration into the object, and collision with the environment. We define the penalty over sampled hand surface points $x_i(q)$ under configuration q as

$$\mathcal{L}_{\text{nonpen}}(q) = \sum_i \left[\phi_{\text{env}}^-(x_i(q))^2 + \phi_{\text{obj}}(x_i(q))^2 + \phi_{\text{hand}}^-(x_i(q))^2 \right], \quad (6)$$

Here, $\phi_{\text{env}}(\cdot)$, $\phi_O(\cdot)$, and $\phi_{\text{hand}}(\cdot)$ denote the signed distance fields of the environment, object, and hand. $\phi^-(x) = \min(\phi(x), 0)$ denotes the negative part of the signed distance. This loss encourages collision-free hand configurations.

b) *Project step to satisfy physical constraints:* Since the non-penetration loss provides only soft guidance, we empirically observe that the sampled configurations do not always fully satisfy the physical constraints. We introduce a projection step after diffusion to correct the generated configuration and enforce feasibility. Specifically, we iteratively optimize the hand configuration using gradients of the loss function $\mathcal{L}_{\text{nonpen}}(q)$, while performing collision checks at each step, until no penetration remains. As shown in Fig. 4, we find this projection step to be particularly effective when the diffusion output contains only minor penetration at the end of sampling.

IV. EXPERIMENTS

In this section, we evaluate whether the proposed guided diffusion framework can generate effective non-prehensile manipulation behaviors conditioned on environmental information. We consider three representative environments in which the robot cannot directly grasp the object and must instead rely on non-prehensile manipulation strategies as shown in Fig. 5.




Fig. 5. **Experiment scenarios.** We evaluate our method in several challenging environments where the robot hand cannot directly grasp the object and has to use non-prehensile manipulation actions.

A. Setups

Our method is evaluated on the Shadow Hand Lite platform [18]. We construct our dataset from a subset of object geometries in DexGraspNet [20], using 51 objects for training and 42 for testing. We focus on the `remote` and `cellphone` categories, as these objects are relatively thin and therefore difficult to grasp directly from a tabletop. For training, we generate 41,635 grasp configurations and 62,594 non-prehensile manipulation configurations. For the diffusion model implementation, we use a PointNet [15] encoder for the object point cloud and a 1D U-Net [17] as the conditional denoising network.

B. Evaluations

We evaluate our method in multiple scenarios, as illustrated in Fig. 5. Specifically, we consider settings in which the object is placed on a tabletop, constrained between layers, or located on a shelf, representing different types of environmental constraints that may arise during manipulation.

A successful manipulation configuration must satisfy two criteria: (1) the maximum penetration is no greater than 1 mm, and (2) the robot establishes sufficient contact with the object to enable stable non-prehensile manipulation, such that the object-robot contact can generate enough force to overcome environmental friction. We compare our method against a modified version of DexGrasp-Anything [26], which is trained only on grasp configurations and does not observe non-prehensile manipulation behaviors during training. We further conduct an ablation study to evaluate the effectiveness of the projection step using the same set of environments.

C. Results

We show qualitative generation results of our method in Fig. 6 and Fig. 7. The results demonstrate that our approach can generate non-prehensile manipulation configurations that adapt to different object geometries and environmental constraints. For quantitative evaluation, we compare our method on all 42 test objects across three evaluation scenarios. As shown in Tab. I, training on both grasping and non-prehensile manipulation actions allows our method to generate successful manipulation configurations at a much higher success rate than DGA, which is trained only on grasping behaviors. Notably, in the challenging shelf environment, our approach achieves




Fig. 6. **Comparison of generated hand configurations.** The DexGrasp-Anything (DGA) baseline fails to generate non-prehensile manipulation behaviors, as shown on the left. With guidance alone, the sampled configurations may still violate the manipulation constraints and fail to contact the object, as shown in the middle. In contrast, with constraint projection, our method generates configurations that satisfy non-penetration and touch the object as shown in green ball, as shown on the right.




Fig. 7. **Other non-prehensile manipulation hand configurations.** Results in the *object between layers* and *object on shelf* environments show that our method can adaptively generate non-prehensile manipulation configurations across diverse environments.

over 60% success rate, while the baseline all fails. This result highlights the importance of diverse manipulation data for learning more generalizable dexterous manipulation strategies in complex environments.

In the ablation study, we examine the effect of the projection step after guided diffusion on the success rate, as reported in Tab. II. We find that guidance alone is often insufficient to produce configurations that are both task-feasible and collision-free. Adding the projection step substantially improves success by correcting configurations that are already close to feasible solutions. These results validate the importance of the projection step in manipulation configuration generation.

V. CONCLUSION

This paper presents a novel approach for generating non-prehensile manipulation behaviors conditioned on environmental context. We show that the proposed method outperforms existing baselines in generating feasible manipulation configurations across diverse environments. In future work, we plan to integrate our approach with motion planning to generate complete manipulation sequences and further evaluate the system in real-world settings.

TABLE I
QUANTITATIVE RESULTS COMPARISON IN THREE TESTING SCENARIOS.

	Table	Layers	Shelf
DGA [26]	31.7%	39.0%	0.0%
Ours	73.2%	78.0%	63.4%

TABLE II
ABLATION STUDY RESULTS UNDER THREE TESTING SCENARIOS.

Scenario	Method	Non-penetration	Manipulation feasibility
Table	Guidance only	26.8%	0.0%
	Guidance + projection	100%	73.2%
Layers	Guidance only	0.0%	0.0%
	Guidance + projection	100%	78.0%
Shelf	Guidance only	24.4%	0.0%
	Guidance + projection	73.2%	63.4%

REFERENCES

- [1] Jiayi Chen, Yubin Ke, and He Wang. Bodex: Scalable and efficient robotic dexterous grasp synthesis using bilevel optimization. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 01–08. IEEE, 2025.
- [2] Xianyi Cheng, Eric Huang, Yifan Hou, and Matthew T. Mason. Contact mode guided sampling-based planning for quasistatic dexterous manipulation in 2d. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6520–6526, 2021. doi: 10.1109/ICRA48506.2021.9560766.
- [3] Xianyi Cheng, Eric Huang, Yifan Hou, and Matthew T. Mason. Contact mode guided motion planning for quasi-dynamic dexterous manipulation in 3d. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 2730–2736, 2022. doi: 10.1109/ICRA46639.2022.9811872.
- [4] Hyungjin Chung, Jeongsol Kim, Michael Thompson McCann, Marc Louis Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=OnD9zGAGT0k>.
- [5] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
- [6] Thomas Feix, Javier Romero, Heinz-Bodo Schmiedmayer, Aaron M. Dollar, and Danica Kragic. The grasp taxonomy of human grasp types. *IEEE Transactions on Human-Machine Systems*, 46(1):66–77, 2016. doi: 10.1109/THMS.2015.2470657.
- [7] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021. URL <https://openreview.net/forum?id=qw8AKxfYbI>.
- [8] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [9] Abhinav Kumar, Thomas Power, Fan Yang, Sergio Aguilera Marinovic, Soshi Iba, Rana Soltani Zarrin, and Dmitry Berenson. Diffusion-informed probabilistic contact search for multi-finger manipulation. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 14277–14283, 2025. doi: 10.1109/ICRA55743.2025.11127844.
- [10] Puhao Li, Tengyu Liu, Yuyang Li, Yiran Geng, Yixin Zhu, Yaodong Yang, and Siyuan Huang. Gendexgrasp: Generalizable dexterous grasping. In *2023 IEEE international conference on robotics and automation (ICRA)*, pages 8068–8074. IEEE, 2023.
- [11] Yunshuang Li, Yiyang Ling, Gaurav S. Sukhatme, and Daniel Seita. Learning geometry-aware nonprehensile pushing and pulling with dexterous hands. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2026.
- [12] Toru Lin, Zhao-Heng Yin, Haozhi Qi, Pieter Abbeel, and Jitendra Malik. Twisting lids off with two hands. In *8th Annual Conference on Robot Learning*, 2024. URL <https://openreview.net/forum?id=3wBqoPfoeJ>.
- [13] Jiaxin Lu, Hao Kang, Haoxiang Li, Bo Liu, Yiding Yang, Qixing Huang, and Gang Hua. Ugg: Unified generative grasping. In *European Conference on Computer Vision*, pages 414–433. Springer, 2024.
- [14] Yuzuko C. Nakamura, Daniel M. Troniak, Alberto Rodriguez, Matthew T. Mason, and Nancy S. Pollard. The complexities of grasping in the wild. In *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*, pages 233–240, 2017. doi: 10.1109/HUMANOIDS.2017.8246880.
- [15] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.
- [16] Haozhi Qi, Ashish Kumar, Roberto Calandra, Yi Ma, and Jitendra Malik. In-hand object rotation via rapid motor adaptation. In Karen Liu, Dana Kulic, and Jeff Ichnowski, editors, *Proceedings of The 6th Conference on Robot Learning*, volume 205 of *Proceedings of Machine Learning Research*, pages 1722–1732. PMLR, 14–18 Dec 2023. URL <https://proceedings.mlr.press/v205/qi23a.html>.
- [17] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [18] Shadow Robot Company. *Shadow Dexterous Hand Lite Technical Specification*. Shadow Robot Company, March 2022. URL https://www.shadowrobot.com/wp-content/uploads/2022/03/shadow_dexterous_hand_lite_technical_specification.pdf. Release: February 2021.
- [19] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=St1giarCHLP>.
- [20] Ruicheng Wang, Jialiang Zhang, Jiayi Chen, Yinzhen Xu, Puhao Li, Tengyu Liu, and He Wang. Dexgraspnet: A large-scale robotic dexterous grasp dataset for general

- objects based on simulation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11359–11366, 2023. doi: 10.1109/ICRA48891.2023.10160982.
- [21] Zhenyu Wei, Zhixuan Xu, Jingxiang Guo, Yiwen Hou, Chongkai Gao, Zhehao Cai, Jiayu Luo, and Lin Shao. $\mathcal{D}(\mathcal{R}, \mathcal{O})$ grasp: A unified representation of robot and object interaction for cross-embodiment dexterous grasping. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4982–4988, 2025. doi: 10.1109/ICRA55743.2025.11127754.
- [22] Yinzhen Xu, Weikang Wan, Jialiang Zhang, Haoran Liu, Zikang Shan, Hao Shen, Ruicheng Wang, Haoran Geng, Yijia Weng, Jiayi Chen, et al. Unidexgrasp: Universal robotic dexterous grasping via learning diverse proposal generation and goal-conditioned policy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4737–4746, 2023.
- [23] Lingxiao Yang, Shutong Ding, Yifan Cai, Jingyi Yu, Jingya Wang, and Ye Shi. Guidance with spherical gaussian constraint for conditional diffusion. In *International Conference on Machine Learning*, 2024.
- [24] Zhao-Heng Yin, Changhao Wang, Luis Pineda, Francois Robert Hogan, Chaithanya Krishna Bodduluri, Akash Sharma, Patrick Lancaster, Ishita Prasad, Mrinal Kalakrishnan, Jitendra Malik, Mike Lambeta, Tingfan Wu, Pieter Abbeel, and Mustafa Mukadam. Dexterity-Gen: Foundation Controller for Unprecedented Dexterity. In *Proceedings of Robotics: Science and Systems*, Los Angeles, CA, USA, June 2025. doi: 10.15607/RSS.2025.XXI.103.
- [25] Mengchao Zhang, Devesh K Jha, Arvind U Raghunathan, and Kris Hauser. Simultaneous Trajectory Optimization and Contact Selection for Multi-Modal Manipulation Planning. In *Proceedings of Robotics: Science and Systems*, Daegu, Republic of Korea, July 2023. doi: 10.15607/RSS.2023.XIX.044.
- [26] Yiming Zhong, Qi Jiang, Jingyi Yu, and Yuexin Ma. Dexgrasp anything: Towards universal robotic dexterous grasping with physics awareness. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 22584–22594, 2025. doi: 10.1109/CVPR52734.2025.02103.
- [27] Yi Zhou, Connelly Barnes, Lu Jingwan, Yang Jimei, and Li Hao. On the continuity of rotation representations in neural networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [28] Yufei Zhu, Yiming Zhong, Zemin Yang, Peishan Cong, Jingyi Yu, Xinge Zhu, and Yuexin Ma. Evolvinggrasp: Evolutionary grasp generation via efficient preference alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11665–11674, October 2025.
- [29] René Zurbrugg, Andrei Cramariuc, and Marco Hutter. Graspqp: Differentiable optimization of force closure for diverse and robust dexterous grasping. In *Conference on Robot Learning (CoRL)*, 2025. URL <https://graspqp.github.io/>.