

A Surprisingly Efficient Representation for Multi-Finger Grasping

Hengxu Yan*, Hao-Shu Fang*, Cewu Lu

Abstract—The problem of grasping objects using a multi-finger hand has received significant attention in recent years. However, it remains challenging to handle a large number of unfamiliar objects in real and cluttered environments. In this work, we propose a representation that can be effectively mapped to the multi-finger grasp space. Based on this representation, we develop a simple decision model that generates accurate grasp quality scores for different multi-finger grasp poses using only hundreds to thousands of training samples. We demonstrate that our representation performs well on a real robot and achieves a success rate of 78.64% after training with only 500 real-world grasp attempts and 87% with 4500 grasp attempts. Additionally, we achieve a success rate of 84.51% in a dynamic human-robot handover scenario using a multi-finger hand.

I. INTRODUCTION

Grasping is a critical function that enables robots to interact with the physical world, and numerous scenarios require more than parallel antipodal grasping. Multi-finger grasping can further enable in-hand robotic manipulation and play a crucial role in the overall grasp problem.

So far, most methods generate multi-finger grasp pose candidates directly from raw observations, such as point clouds or RGB images, as demonstrated in [1, 2, 3, 4, 5, 6, 7, 8]. Specifically, [6] proposes a methodology for estimating the pose of an object within a scene and subsequently learning an offset based on the object’s pose to determine the grasp pose. Similarly, [7, 8] focus on predicting the contact area between the object and a multi-finger hand. However, a limitation common to these methods is their capacity to grasp only a single object. Additionally, the multi-finger grasp pose presents significant complexities and variations due to the high degree of freedom of the multi-finger hand, making the labeling of grasp data extremely costly. Furthermore, since there are different types of multi-finger grippers available on the market, establishing a multi-finger grasping dataset for each type of gripper is costly. Consequently, achieving satisfactory grasping performance on multiple grippers is currently impractical.

In our study, we introduce a novel intermediate representation aimed at the generation of multi-finger grasp poses, emphasizing minimal data requirements while preserving robust performance. Our methodology is structured into two principal phases. The initial phase involves the prediction of dense circular antipodal representations distributed throughout the observed scene. These representations facilitate the

Hengxu Yan, Hao-Shu Fang, Cewu Lu are with Shanghai Jiao Tong University, Cewu Lu is the corresponding author. Email: HengxuYan@sjtu.edu.cn, fhaoshu@gmail.com, lucewu@sjtu.edu.cn

*The first two authors contributed equally to this work.




Fig. 1. The circular antipodal representations for multi-finger grasping. The less transparent the representation is, the higher the circular antipodal score.

sampling of multiple candidates for multi-finger grasp poses. For each designated grasp location upon an arbitrary curved surface within the scene, this representation is characterized by two key parameters: **the antipodal scores** and **the minimum antipodal grasp widths**, given different in-plane rotations and grasp depths along the grasp viewpoint. It implicitly characterizes the surface geometry of the object. Subsequently, the representation corresponding to each grasp candidate is input into a decision model, which serves the purpose of identifying the specific multi-finger grasp type. To simplify this process, we discretize the multi-finger grasp space, generating a set of feasible multi-finger grasp poses based on a predefined taxonomy of 16 grasp types. These generated poses serve as the pool from which the decision model selects an appropriate grasp type. This procedural pipeline effectively reduces the complexity of the decision space, all the while maintaining the flexibility of the multi-finger hand. To validate the efficacy of our representation design, we compare our representation with several network features, providing empirical evidence of its internal structural effectiveness. Fig. 1 shows the circular antipodal representation for multi-finger grasping.

We conduct experiments in the real world to verify the effectiveness of the representation, including grasping static objects in cluttered scenes and reactive grasping in a handover setting. For grasping in cluttered scenes, we randomly pour diverse objects onto a plate, and a robot empties the plate with a multi-finger hand. During each grasp attempt, our method generate multiple high-quality multi-finger grasp poses for a partial-view point cloud. The best grasp pose for execution is selected by considering both collision avoidance and grasp quality. We achieve a success rate of 78.64%

after training with only 500 real-world grasp attempts and 87% with 4500 on over 300 unseen objects of varying sizes, shapes, and materials.

For the reactive multi-finger hand-over case, an in-hand camera is attached to the robot’s wrist, enabling the execution of closed-loop control. This is achieved by continuously predicting the grasp pose based on the current observation and subsequently servoing to the target pose. We show that we can achieve the temporal smoothness property for the intermediate representation, which makes it easy to grasp dynamic objects. In a handover setting, we report an 84.51% success rate on 30 objects.

We summarize our primary contributions in the following aspects:

- Identifying and developing an efficient representation tailored for multi-finger grasping.
- Demonstrating that our representation surpasses the performance of other unstructured network features.
- Achieving exceptional performance with our method, necessitating only a few hundred grasp attempts.
- Illustrating the temporal smoothness property of our representation and its successful application in reactive grasping scenarios.

II. RELATED WORK

To date, most of the research in this field has been focusing on the problem of parallel antipodal grasping, with only a small subset of studies investigating multi-finger grasping. As a result, this section will be divided into two parts, each addressing one of these two topics: multi-finger grasping and parallel antipodal grasping.

A. Multi-finger Grasping

During the early stages of multi-finger grasp research, the primary focus is on developing effective force closure strategies, as demonstrated by studies such as [9, 10, 11, 12, 13, 14]. However, a major challenge with this approach is the requirement for an accurate model of each object in the environment [15]. This can limit the model’s applicability and make it difficult to generalize. Moreover, the overall complexity of the model can lead to significant computational costs, resulting in only a small number of feasible grasp poses being generated.

Subsequently, many learning-based approaches emerge, which can overcome these challenges and achieve greater efficiency. Due to the high cost of obtaining data through multi-finger grasping, many studies resort to collecting trial-and-errors in the simulated environment [16], where a domain gap inevitably exists. Some studies borrow demonstrations of human grasping [17, 18] and use model-free methods to imitate human learning, while others adopt pure reinforcement learning approaches. Additionally, some studies focus on generating efficient multi-finger hand grasp poses dataset [19, 20]. For instance, [21] proposes a large-scale dataset of 1.32 million grasp poses for 5355 objects, spanning more than 133 categories of objects. Meanwhile, [22] constructs

a dataset that can train a model to maintain contact with objects as much as possible.

Most of studies previously mentioned predominantly rely on raw perception and necessitate substantial data usage. Often, these studies are narrowly focused on grasping either a single object or a limited array of objects that are methodically arranged, thereby constraining their broader applicability. In contrast, our model, employing an intermediate representation, markedly reduces data requirements and demonstrates the capability to grasp a multitude of objects that are randomly stacked together.

B. Parallel Antipodal Grasping

Parallel antipodal grasping involves generating multiple grasp poses for a given scene. Early research focuses on 2D planes and predicts 4 degrees of freedom, as seen in studies such as [23, 24, 25]. However, this approach may miss many potentially effective grasp poses. DexNet2.0 [26] introduces significant advancements in 3D grasp poses, while GPD [27] becomes important for generating 6 degrees of freedom poses in complete space. Subsequently, [28] proposes an end-to-end network that can generate dense 6 degrees of freedom poses for the entire scene, and they also develop a large parallel antipodal grasp dataset [29]. Many studies explore the potential of related parallel antipodal grasping techniques, such as [30, 31, 32]. As the performance of parallel antipodal grasping approaches that of humans [33], this work aims to build a multi-finger grasping method based on the parallel antipodal grasping approach of [32, 33].

In dynamic scenes, previous works predominantly use a fixed trajectory or select the nearest pose at two adjacent moments as the tracking pose [34, 35, 36, 37, 38]. In contrast, this paper leverages the circular antipodal representation with the highest similarity between the two frames to generate the tracking pose.

III. PROBLEM FORMULATION

For a grasp-based multi-finger hand, we can define its pose \mathcal{G} as

$$\mathcal{G} = [\mathbf{R} \ \mathbf{t} \ \mathbf{B}], \quad (1)$$

where $\mathbf{R} \in \mathbb{R}^{3 \times 3}$ denotes the multi-finger hand rotation, $\mathbf{t} \in \mathbb{R}^{3 \times 1}$ denotes the translation of multi-finger hand, they are the same as the definition of parallel antipodal gripper (for specific correspondence, please refer to Sec. IV-B.1), and \mathbf{B} denotes the freedom of multi-finger hand. Because multi-finger grasping is a decision problem on a continuous joint space, we simplify the degrees of freedom of the multi-finger hand into $\mathbf{C} \in \mathbb{R}^{c \times 1}$ alternative grasp types and $\mathbf{W} \in \mathbb{R}^{w \times 1}$ alternative grasp width in this paper. So we can equate Eqn. (1) as

$$\mathcal{G} = [\mathbf{R} \ \mathbf{t} \ w \ c]. \quad (2)$$

Let \mathcal{E} denote the environment containing all objects and robots in the scene, \mathcal{P} denote the point cloud taken by the camera from a certain angle and \mathcal{M} denote the representation model, the representation $\hat{\mathcal{R}}$ of scene can be denoted as

$$\hat{\mathcal{R}} = \mathcal{M}(\mathcal{P}). \quad (3)$$




Fig. 2. (a) 2D circular antipodal representations. The depth of color denotes the level of antipodal scores. (b) 3D circular antipodal representations. (c) The coordinate system of parallel antipodal and multi-finger grasping

Let $q(\hat{\mathcal{R}}, \mathcal{G}, \mathcal{E})$ denotes a binary variable indicating whether an object is successfully grasped or not given grasp pose \mathcal{G} and representation $\hat{\mathcal{R}}$. Our goal is to find a pose \mathcal{G}_i from a set of candidate poses $\mathbf{G} = \{\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_n\}$ that maximizes the grasp success rate:

$$\mathcal{G}_i = \arg \max_{\mathcal{G} \in \mathbf{G}} \text{Prob}(q = 1 | \hat{\mathcal{R}}, \mathcal{G}, \mathcal{E}). \quad (4)$$

IV. METHOD

In this section, we will introduce our method in two parts. The first part will cover the circular antipodal representation, and the second part will delve into the detailed design of our algorithm.

A. Circular Antipodal Representation

We first consider a 2D object. Given a grasp point \mathcal{X} , our representation is defined as a set of circular antipodal scores and widths:

$$\hat{\mathcal{R}}_{\mathcal{X}} = \{(s_a, w_a) \mid a = 1, 2, \dots, \mathcal{A}\}, \quad (5)$$

where s_a denotes an antipodal score along the a -th in-plane rotation and w_a denotes the minimum antipodal grasp width for the corresponding rotation. The in-plane circular is discretized into \mathcal{A} angles in total. Fig. 2(a) illustrates the antipodal score given a tangent plane \mathcal{H} .

Then, for a real-world 3D object, given a grasp point and grasp approach direction, the circular antipodal representation can be defined similarly except that we add an extra depth dimension along the grasp viewpoint:

$$\hat{\mathcal{R}}_{\mathcal{X}} = \{(s_a^d, w_a^d) \mid a = 1, 2, \dots, \mathcal{A}, d = 1, 2, \dots, \mathcal{D}\}. \quad (6)$$

We discretize the depth dimension along the grasp viewpoint into \mathcal{D} classes, as illustrated in Fig. 2(b).

B. Algorithm Details

1) *Mapping Principles*: To begin with, in order to map the circular antipodal representation to the multi-finger grasp type, we first select the parallel antipodal grasp pose that has the highest antipodal score within the representation. The multi-finger grasp pose is linked to the representation by sharing the same $\mathbf{R}, \mathbf{t}, w$ with the selected parallel antipodal grasp pose. Fig. 2(c) illustrates the coordinate system of both parallel antipodal and multi-finger grasping. For each representation, 16 grasping types are generated, which encompass both precision grasping and power grasping, and cover a wide range of human postures. Fig. 3 illustrates the different grasp types.




Fig. 3. Multi-finger grasp type, every row denotes that we use the same depth but a different finger type to grasp, and every column denotes the same type but a different grasp depth.

2) *Representation Model*: Upon initial examination, the task of obtaining a circular antipodal representation for a partial-view point cloud may seem formidable. However, recent advancements in grasp pose detection algorithms, driven by the application of deep learning techniques, have yielded substantial progress in this regard. These algorithms have demonstrated exceptional proficiency in swiftly and accurately predicting antipodal scores. Noteworthy among these advances are GraspNet-Baseline [28] and GSNet [32], both of which decompose the grasp pose detection challenge into three distinct components: the prediction of grasp point, approach direction, and in-plane antipodal scores and widths along the approach direction. Importantly, the outputs of these algorithms at their final layer inherently provide the circular antipodal representation that is integral to our study. Consequently, for the purposes of this paper, we have opted to employ the readily available GSNet [32] as our designated representation model.

3) *Decision Model*: The architecture of the decision model is illustrated in Fig. 4. The network comprises seven fully connected layers and a skip connection, which is simple and concise. The input to the model is the circular antipodal representation $\hat{\mathcal{R}}_{\mathcal{X}}$ of each grasp point that is selected. The output of the last layer in the model is a list of scores $O \in \mathbb{R}^{\mathcal{A} \times \mathcal{D} \times \mathcal{C}}$ which denotes the grasp quality score of every multi-finger grasp type in \mathcal{C} and maps to different antipodal scores in the representation $\hat{\mathcal{R}}_{\mathcal{X}}$, where \mathcal{A}, \mathcal{D} and \mathcal{C} are predefined as 12, 5 and 16, respectively. The loss function is defined as:

$$L = -\frac{1}{\mathbf{Z}} \sum_{z=1}^{\mathbf{Z}} y_c \log(p_{o,c}). \quad (7)$$

Here, y_c is the binary indicator (0 or 1) if multi-finger grasp type c successfully grasp an object given a and d , $p_{o,c}$ is the predicted probability of c for grasp quality score $o \in O$, and \mathbf{Z} is the batch size. The remaining multi-finger grasp types and different a, d are ignored when calculating the

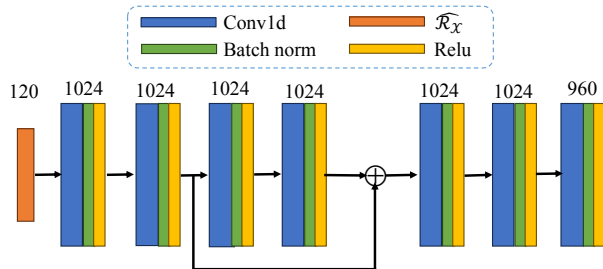


Fig. 4. The architecture of decision model

loss because we do not know if they can grasp successfully (Please refer to the data collection and annotation Sec. V-B). All grasp points share the same decision model parameters.

4) *Detection Post-processing*: After obtaining grasp quality scores of each grasp point from the decision model, we select the top 500 multi-finger grasp poses \mathbf{G} according to the grasp quality scores of all grasp points for collision detection and ultimately determine the grasp pose \mathcal{G}_i . Among them, collision detection is performed by voxelizing the pre-shaped 3D multi-finger hand model with a voxel size of 0.3 cm, and checking whether there are intersections between the hand voxels and scene point cloud, using Open3D library [39]. We aim to improve the grasp success rate by selecting the grasp quality scores of multi-finger grasp poses with high confidence levels, which are typically above 0.9. We augment the point cloud of a scene by randomly translating or rotating it 10 times to generate additional multi-finger grasp proposals.

5) *Dynamic Grasping*: we initialize our approach by selecting a circular antipodal representation $\hat{\mathcal{R}}_{\mathcal{X}}^0$ along with its corresponding multi-finger grasp pose through the utilization of the representation model and decision model for the first frame. Subsequently, we harness the association module within the AnyGrasp framework [33] to construct the correspondence matrix between $\hat{\mathcal{R}}_{\mathcal{X}}^0$ and a set of circular antipodal representations $\{\hat{\mathcal{R}}_{\mathcal{X}_0}, \hat{\mathcal{R}}_{\mathcal{X}_1}, \dots, \hat{\mathcal{R}}_{\mathcal{X}_n}\}$ derived from candidate poses \mathbf{G} generated in the subsequent frame. We then select the multi-finger grasp pose which antipodal representation has the highest similarity in the matrix for the next frame. This approach ensures temporal continuity in the context of reactive grasping.

V. EXPERIMENTS

A. Implementation Details

To evaluate the performance of our multi-finger grasping, we set up several experimental platforms in a real-world environment and conduct our experiments.

For static grasping, we deploy the UR5 robot arm and Intel’s RealSense D415 camera to capture the entire scene’s point cloud. The camera is positioned overhead in the scene. The gripper is a five-finger dexterous hand manufactured by Inspire*, which has 6 motors, one for each of the four fingers, and two for the thumb. The hand consists of 12 underactuated joints, with a weight of approximately 0.5 kg.

*<http://www.inspire-robots.com/product/frwz/>

For dynamic grasping, we employ the Flexiv Rizon robotic arm, which can update the servo target smoothly. The camera we use is the Intel RealSense L515, which is mounted between the multi-finger hand and the end-link of the robotic arm. We opt for this in-hand setting because images captured from an overhead view would be heavily occluded by the robot when grasping. Additionally, we choose the Intel RealSense L515 camera because it can generate point clouds for objects at close range.

For model training, since we adopt the off-the-shelf GSNet [32] as the representation model, we only need to train the decision model. The decision model is trained for only 20 epochs, leveraging the Adam optimizer [40]. The learning rate follows a segmented descent strategy, and the batch size \mathbf{Z} is set to 128 to optimize training efficiency.

The training and inference of all the models in the following experiments are performed on the Ubuntu 20.04 system, using an Intel i9-10900K CPU and an NVIDIA 2080 Ti GPU. All the code is written in Python language.

B. Data Collection and Annotation

As mentioned earlier, most of the data related to multi-finger grasping in previous works are generated within a simulation environment. Nevertheless, significant deviations may arise due to the inherent differences between the simulation and real environments. Consequently, certain poses may be graspable in simulation, but ungraspable in the real world. For this reason, the multi-finger grasp data obtained from simulation still requires validation within a real environment. So we directly collect data in the real world.

Let us provide an overview of the complete data collection pipeline. Initially, we select around 100 objects ranging in size from $3 \times 3 \times 4 \text{ cm}^3$ to $20 \times 8 \times 15 \text{ cm}^3$ and randomly place some objects within a scene. We then use the representation model to generate dense circular antipodal representations $\hat{\mathcal{R}}$ for the scene, along with randomly selecting a multi-finger grasp type c in Fig. 3 to generate multi-finger grasp pose \mathcal{G} for each grasp point. Following this, we perform collision detection and select the multi-finger grasp pose with the highest antipodal score for actual grasping. During this process, we record whether the grasp is successful, the point cloud of the scene, and relevant information associated with the corresponding grasp point. We collect 5000 grasp attempts in total.

As discussed in Sec. IV-B.1, there is a relationship between the parallel antipodal grasping and multi-finger grasping, such that they share the same $\mathbf{R}, \mathbf{t}, w$. As a result, we can achieve a success rate of over 58% during data collection. Our primary objective is to identify and address the remaining unsuccessful multi-finger grasps by leveraging the decision model to learn the inconsistencies between the parallel gripper and the five-finger hand.

We randomly select 500 samples from the entire dataset to create the evaluation dataset, while the remaining samples comprise the training dataset. We construct test scenarios that represent common objects found in daily life, including their shape, material, and texture etc. The data is divided into




Fig. 5. Performance of different training data on the evaluation dataset




Fig. 6. Performance of different training data on real-world grasping

six categories: hardware, snack, ragdoll, household, toy, and adversarial objects, which contain over 300 objects ranging in size from $2.5 \times 2.5 \times 2.5 \text{ cm}^3$ to $36 \times 4 \times 11.5 \text{ cm}^3$. From the collected data distribution, we observe that deeper and more fingers utilized during grasping increase the likelihood of success but also raise the possibility of collision detection failures.

C. Experimental Results

1) *Static Scenes*: To evaluate the learning efficiency of our model, we commence with an analysis of its success rate on the evaluation set, utilizing training datasets of varied sizes. The results for each experiment are obtained from three training runs, each consisting of a random subset of the entire training dataset. The results are presented in Fig. 5. We observe that the performance of our model improves as the size of the training dataset increases. When the amount of trial-and-error data exceeds 500, our model achieves a success rate of over 73%. Furthermore, the average success rate of our model can reach 88% when using all training samples.

We then evaluate our method in real-world grasping. Due to the time-consuming nature of real robot grasping, we only test models trained with three different dataset volumes. The overall experimental results are consistent with those obtained from the evaluation dataset and are presented in Fig. 6. The reason for the decrease in the success rate of hardware and toys is attributed to a minor element of randomness in the grasp process. The grasp video can be found in the Appendix (S1-S2, S4-S7).

Next, we evaluate the success rate of using different representations. Models are tested on real-world grasping

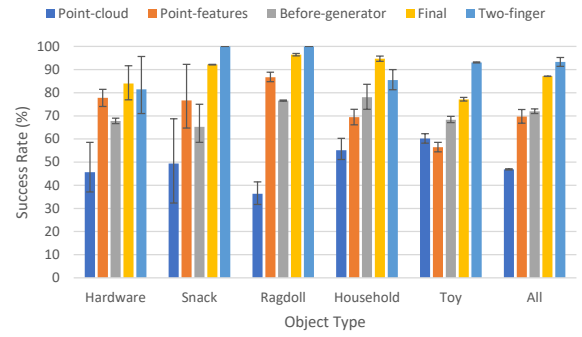


Fig. 7. Performance of different features on real-world grasping




Fig. 8. (a) Performance of different training data on the adversarial object. (b) Performance of different representations on the adversarial object.

while trained on the full data set. The results are shown in Fig. 7. The representation “Point-cloud” only uses the Minkowski Engine [41] to obtain the point cloud features of the entire scene. The representation “Point-features” means the output features of the third last MLP block in the representation model, the “Before-generator” is the second last, and the “Final” refers to our structured representation. We feed the representation of the grasp point into the decision model in Sec. IV-B.3 for learning. Our results show that the performance of using the structured circular antipodal representation outperforms other intermediate features by a large margin, with a final success rate of 87%. The video of the grasping is available in Appendix (S1-S2, S8-S9, S11-S14).

We conduct further evaluation of our model on challenging adversarial objects. The experimental dataset comprises 13 samples from DexNet2.0 [26] and 49 samples from EGAD [42]. The results of our model with different trial-and-error numbers and different representation inputs are shown in Fig. 8. We can observe from the videos in Appendix (S3, S5, S7, S10, S12, S14) that our method achieves the best grasp performance on adversarial objects when using our structured representation. Notably, we also report the results of grasping with the antipodal grasp pose generated by the representation model, which is denoted as “Two-finger” in Fig. 8(b). We can see that our multi-finger grasping can outperform its representation model with the aid of an extra decision model.

Finally, we tested the time taken by our system in different phases. The results are shown in Table. I. The mean picks per hour (MPPH) of the entire system can reach 167.78. We can see that both the representation model and decision model run efficiently. The most time-consuming part of our system




Fig. 9. Dynamic grasping for hardware and ragdoll, the first column is the initial setting.

TABLE I. The time spent in each grasping procedure, RM means the time required by the representation model to generate the representations, DM is the decision model, CD is collision detection, and GAP is the time required by a robot to grasp and place an object.

Grasping Procedure	RM	DM	CD	GAP
Time (s)	0.2	0.01	19	2.24

TABLE II. The success rate of each category for dynamic grasping.

Category	Hardware	Snack	Ragdoll
Success rate (%)	83.33	100.00	90.90
Category	Household	Toy	Adversarial
Success rate (%)	71.43	90.90	83.33

is collision detection, which costs nearly 20 seconds. The current collision detection runs on the CPU, and we aim to address this problem in the future.

2) *Dynamic Scenes*: Temporal continuity is crucial in dynamic scenes. An advantage of our framework is that the temporal continuity for multi-finger grasping can be inherited from that of the representation model. Fortunately, the temporal block introduced in [33] is capable of ensuring this. As a result, we use the output of that block as our circular antipodal representations to conduct the handover experiments. We first generate the circular antipodal representations for the first frame and select the best multi-finger grasp pose grasp quality score by the decision model. Subsequently, the temporal block in [33] is used to track the translation and orientation of the representation. The target multi-finger pose is updated according to the translation and orientation of the associated representation.

We randomly selected five objects from each category in the static scene and combined them to create the real-world test grasp dataset for the handover experiment. In total, 30 objects were collected. We performed the handover evaluation on the entire real-world grasping twice and achieved a final success rate of 84.51%. The results for each category are shown in Table. II, and Fig. 9 illustrates two samples of dynamic grasping. The video can be found in Appendix (S16).

VI. CONCLUSION

In this paper, we propose an efficient multi-finger grasp method that utilizes supervised learning to grasp objects in

complex environments. With a representation model and a decision model, we can generate multi-finger grasp poses. Our training data is collected by trials and errors in the real world that makes it easy to be adapted for different hands. Extensive real robot experiments are conducted to evaluate the effectiveness of our method.

While our proposed multi-finger grasp method has yielded impressive results, we recognize that there is ample room for further improvement. For instance, the multi-finger grasp poses we have defined are not exhaustive, and there are many more poses that human hands can achieve that our method currently cannot replicate. To achieve the same level of dexterity as human hands, we will require more comprehensive datasets, as well as multi-finger hand with greater degrees of freedom. These advancements will enable us to tackle more complex manipulation tasks that build upon the multi-finger hand grasp problem.

APPENDIX All experiments related to real-world scenes in this work are recorded in the form of videos. Since the videos are a bit long, we have accelerated a portion of them. All video links can be viewed below.

- S1: <https://youtu.be/Rgg-HsM2LNA>
- S2: <https://youtu.be/hHFuDTxQh9E>
- S3: https://youtu.be/aWCX-DOfW_o
- S4: <https://youtu.be/frJQw7jovmc>
- S5: <https://youtu.be/eHKvOVoZI6c>
- S6: <https://youtu.be/w4zFK75rKXg>
- S7: <https://youtu.be/kG0k8HOUFzA>
- S8: <https://youtu.be/RZPGgEOBxqM>
- S9: https://youtu.be/QynUVdQdK_0
- S10: <https://youtu.be/cDEc10D07BM>
- S11: <https://youtu.be/4tXBtFnt6Tc>
- S12: <https://youtu.be/dfOCi0yN4sY>
- S13: <https://youtu.be/WNb5JJykQtQ>
- S14: <https://youtu.be/ew6JlynFoJc>
- S15: <https://youtu.be/8FztVFRcvMY>
- S16: <https://youtu.be/LdkniuMXyuk>, <https://youtu.be/M2VwJj18FHE>

ACKNOWLEDGEMENTS This work was supported by the National Key Research and Development Project of China (No.2022ZD0160102, No.2021ZD0110704), Shanghai Artificial Intelligence Laboratory, XPLOER PRIZE grants.

REFERENCES

- [1] Jacob Varley et al. “Generating multi-fingered robotic grasps via deep learning”. In: *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2015, pp. 4415–4420. DOI: 10.1109/IROS.2015.7354004.
- [2] Yinzhen Xu et al. “UniDexGrasp: Universal Robotic Dexterous Grasping via Learning Diverse Proposal Generation and Goal-Conditioned Policy”. In: *arXiv preprint arXiv:2303.00938* (2023).
- [3] Ruicheng Wang et al. “DexGraspNet: A Large-Scale Robotic Dexterous Grasp Dataset for General Objects Based on Simulation”. In: *2023 IEEE International Conference on Robotics and Automation (ICRA)*. 2023, pp. 11359–11366. DOI: 10.1109/ICRA48891.2023.10160982.
- [4] Puhao Li et al. “GenDexGrasp: Generalizable Dexterous Grasping”. In: *2023 IEEE International Conference on Robotics and Automation (ICRA)*. 2023, pp. 8068–8074. DOI: 10.1109/ICRA48891.2023.10160667.
- [5] Weikang Wan et al. “UniDexGrasp++: Improving Dexterous Grasping Policy Learning via Geometry-aware Curriculum and Iterative Generalist-Specialist Learning”. In: *arXiv preprint arXiv:2304.00464* (2023).
- [6] Jens Lundell et al. “Multi-FinGAN: Generative Coarse-To-Fine Sampling of Multi-Finger Grasps”. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)* (2020), pp. 4495–4501.
- [7] Hanwen Jiang et al. “Hand-object contact consistency reasoning for human grasps generation”. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021, pp. 11107–11116.
- [8] Dylan Turpin et al. “Grasp’d: Differentiable contact-rich grasp synthesis for multi-fingered hands”. In: *European Conference on Computer Vision*. Springer, 2022, pp. 201–221.
- [9] Dmitry Berenson and Siddhartha S. Srinivasa. “Grasp synthesis in cluttered environments for dexterous hands”. In: *Humanoids 2008 - 8th IEEE-RAS International Conference on Humanoid Robots*. 2008, pp. 189–196. DOI: 10.1109/ICHR.2008.4755944.
- [10] A. Sahbani, S. El-Khoury, and P. Bidaud. “An overview of 3D object grasp synthesis algorithms”. In: *Robotics and Autonomous Systems* 60.3 (2012), pp. 326–336.
- [11] Alberto Rodriguez, Matthew T Mason, and Steve Ferry. “From caging to grasping”. In: *The International Journal of Robotics Research* 31.7 (2012), pp. 886–900.
- [12] Carlos Rosales et al. “On the synthesis of feasible and prehensile robotic grasps”. In: *2012 IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 550–556.
- [13] Domenico Prattichizzo et al. “On the manipulability ellipsoids of underactuated robotic hands with compliance”. In: *Robotics and Autonomous Systems* 60.3 (2012), pp. 337–346.
- [14] Tengyu Liu et al. “Synthesizing Diverse and Physically Stable Grasps With Arbitrary Hand Structures Using Differentiable Force Closure Estimator”. In: *IEEE Robotics and Automation Letters* 7.1 (2021), pp. 470–477.
- [15] Jeannette Bohg et al. “Data-Driven Grasp Synthesis—A Survey”. In: *IEEE Transactions on Robotics* 30.2 (2014), pp. 289–309. DOI: 10.1109/TRO.2013.2289018.
- [16] Yuzhe Qin, Hao Su, and Xiaolong Wang. “From one hand to multiple hands: Imitation learning for dexterous manipulation from single-camera teleoperation”. In: *IEEE Robotics and Automation Letters* 7.4 (2022), pp. 10873–10881.
- [17] Samarth Brahmabhatt et al. “Contactgrasp: Functional multi-finger grasp synthesis from contact”. In: *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 2386–2393.
- [18] Mia Kovic, Danica Kragic, and Jeannette Bohg. “Learning task-oriented grasping from human activity datasets”. In: *IEEE Robotics and Automation Letters* 5.2 (2020), pp. 3352–3359.
- [19] Yu-Wei Chao et al. “DexYCB: A benchmark for capturing hand grasping of objects”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2021, pp. 9044–9053.
- [20] Omid Taheri et al. “GRAB: A dataset of whole-body human grasping of objects”. In: *European conference on computer vision*. Springer, 2020, pp. 581–600.
- [21] Ruicheng Wang et al. “DexGraspNet: A Large-Scale Robotic Dexterous Grasp Dataset for General Objects Based on Simulation”. In: *arXiv preprint arXiv:2210.02697* (2022).
- [22] Tianqiang Zhu et al. “Toward human-like grasp: Dexterous grasping via semantic representation of object-hand”. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021, pp. 15741–15751.
- [23] Antonio Bicchi and Vijay Kumar. “Robotic grasping and contact: A review”. In: *IEEE International Conference on Robotics and Automation (ICRA)*. Vol. 1. IEEE, 2000, pp. 348–353.
- [24] Hao Dang and Peter K Allen. “Semantic grasping: Planning robotic grasps functionally suitable for an object manipulation task”. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2012, pp. 1311–1317.
- [25] J. Bohg et al. “Data-Driven Grasp Synthesis—A Survey”. In: *IEEE Transactions on Robotics (T-RO)* 30.2 (2014), pp. 289–309.
- [26] Jeffrey Mahler et al. “Dex-Net 2.0: Deep Learning to Plan Robust Grasps with Synthetic Point Clouds

- and Analytic Grasp Metrics”. In: *Proceedings of Robotics: Science and Systems (RSS)*. Cambridge, Massachusetts, July 2017.
- [27] Andreas ten Pas et al. “Grasp pose detection in point clouds”. In: *The International Journal of Robotics Research (IJRR)* 36.13-14 (2017), pp. 1455–1473.
- [28] Hao-Shu Fang et al. “GraspNet-1Billion: A Large-Scale Benchmark for General Object Grasping”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020.
- [29] Hao-Shu Fang et al. “Robust grasping across diverse sensor qualities: The GraspNet-1Billion dataset”. In: *The International Journal of Robotics Research* 42.12 (2023), pp. 1094–1103.
- [30] Minghao Gou et al. “Rgb matters: Learning 7-dof grasp poses on monocular rgb-d images”. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2021, pp. 13459–13466.
- [31] Hongjie Fang et al. “TransCG: A Large-Scale Real-World Dataset for Transparent Object Depth Completion and a Grasping Baseline”. In: *IEEE Robotics and Automation Letters* 7.3 (2022), pp. 7383–7390.
- [32] Chenxi Wang et al. “Grasping discovery in clutters for fast and accurate grasp detection”. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021, pp. 15964–15973.
- [33] Hao-Shu Fang et al. “AnyGrasp: Robust and Efficient Grasp Perception in Spatial and Temporal Domains”. In: *Transaction on Robotics (T-RO)* (2023).
- [34] Seungsu Kim, Ashwini Shukla, and Aude Billard. “Catching objects in flight”. In: *IEEE Transactions on Robotics* 30.5 (2014), pp. 1049–1065.
- [35] Arjun Menon, Benjamin Cohen, and Maxim Likhachev. “Motion planning for smooth pickup of moving objects”. In: *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2014, pp. 453–460.
- [36] Iretoiyo Akinola et al. “Dynamic grasping with reachability and motion awareness”. In: *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2021, pp. 9422–9429.
- [37] Naresh Marturi et al. “Dynamic grasp and trajectory planning for moving objects”. In: *Autonomous Robots* 43.5 (2019), pp. 1241–1256.
- [38] Douglas Morrison, Juxi Leitner, and Peter Corke. “Closing the Loop for Robotic Grasping: A Real-time, Generative Grasp Synthesis Approach”. In: *Robotics: Science and Systems*. 2018.
- [39] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. “Open3D: A Modern Library for 3D Data Processing”. In: *arXiv:1801.09847* (2018).
- [40] Diederik P Kingma and Jimmy Ba. “Adam: A method for stochastic optimization”. In: *arXiv preprint arXiv:1412.6980* (2014).
- [41] Christopher Choy, JunYoung Gwak, and Silvio Savarese. “4D Spatio-Temporal ConvNets: Minkowski Convolutional Neural Networks”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2019, pp. 3075–3084.
- [42] Douglas Morrison, Peter Corke, and Jürgen Leitner. “Egad! an evolved grasping analysis dataset for diversity and reproducibility in robotic manipulation”. In: *IEEE Robotics and Automation Letters* 5.3 (2020), pp. 4368–4375.