

# Grasping a Handful: Sequential Multi-Object Dexterous Grasp Generation

Haofei Lu , Graduate Student Member, IEEE, Yifei Dong , Graduate Student Member, IEEE, Zehang Weng , Florian T. Pokorny , Member, IEEE, Jens Lundell , Member, IEEE, and Danica Kragic 

**Abstract**—We introduce the sequential multi-object robotic grasp sampling algorithm SeqGrasp that can robustly synthesize stable grasps on diverse objects using the robotic hand’s partial Degrees of Freedom (DoF). We use SeqGrasp to construct the large-scale Allegro Hand sequential grasping dataset SeqDataset and use it for training the diffusion-based sequential grasp generator SeqDiffuser. We experimentally evaluate SeqGrasp and SeqDiffuser against the state-of-the-art non-sequential multi-object grasp generation method MultiGrasp in simulation and on a real robot. The experimental results demonstrate that SeqGrasp and SeqDiffuser reach an 8.71%–43.33% higher grasp success rate than MultiGrasp. Furthermore, SeqDiffuser is approximately 1000 times faster at generating grasps than SeqGrasp and MultiGrasp.

**Index Terms**—Grasping, dexterous manipulation, data sets for robot learning.

## I. INTRODUCTION

**D**EXTEROUS grasping has been extensively studied in robotics [1], [2]. The field has evolved through two primary research directions. First, numerous methods have been developed for generating robotic grasps [3], [4], [5], [6], [7], [8], from analytical approaches [7] to data-driven generative models including Generative Adversarial Networks (GANs) [5], diffusion models [3], [4], and Variational Autoencoders (VAEs) [6], [8]. Second, a complementary line of research has focused on understanding human grasping behaviors, including intention recognition [9], real-time tracking of hand-object interactions [10], grasp taxonomy classification [11], and learning task-oriented grasping strategies from human demonstrations [12].

The primary strategy in dexterous grasp generation has been to bring the robotic hand close to the object and then




Fig. 1. Sequential multi-object grasping.

simultaneously envelop it with all fingers. While this strategy often results in efficient and successful grasp generation, it simplifies dexterous grasping to resemble parallel-jaw grasping, thereby underutilizing the many DoF of multi-fingered robotic hands [7]. In contrast, grasping multiple objects with a robotic hand, particularly in a sequential manner that mirrors human-like dexterity, as shown in Fig. 1, is still an unsolved problem.

In this work, we introduce SeqGrasp, a novel hand-agnostic algorithm for generating sequential multi-object grasps. Our approach utilizes an optimization-based method to sequentially determine single-object grasp poses using a subset of the hand’s DoF. As the grasp sequence progresses, the DoF engaged in previous grasps are frozen, leaving only the remaining DoF available for subsequent object grasps. To only engage a subset of the hand’s DoF for each grasp, we propose an Opposition Space (OS) selection strategy that enables stable grasping using only a pair of links. Using SeqGrasp, we construct the large-scale dataset SeqDataset containing 870 K penetration-free Allegro hand grasps across 509 objects, with up to four sequentially grasped objects. Finally, we train the conditional sequential grasping diffusion model SeqDiffuser on SeqDataset to enable sequential grasping on novel objects.

We experimentally evaluate SeqGrasp, SeqDiffuser, and the state-of-the-art simultaneous multi-object grasping method MultiGrasp [13] in simulation and on physical hardware. The simulation results revealed that SeqGrasp and SeqDiffuser perform on par with MultiGrasp for picking one or two objects while outperforming it when picking three to four objects. Moreover, SeqDiffuser demonstrates superior efficiency, generating 256 grasps within one second compared to approximately 1000 seconds for SeqGrasp and MultiGrasp. For the real-world experiments, we replicated the grasp sequences proposed by the

Received 27 March 2025; accepted 8 September 2025. Date of publication 24 September 2025; date of current version 13 October 2025. This article was recommended for publication by Associate Editor M. Liarokapis and Editor J Borrás Sol upon evaluation of the reviewers’ comments. This work was supported in part by Swedish Research Council, in part by Knut and Alice Wallenberg Foundation, and in part by the European Research Council under Grant ERC-BIRD-884807. (Corresponding author: Haofei Lu.)

Haofei Lu, Yifei Dong, Zehang Weng, Florian T. Pokorny, and Danica Kragic are with Robotics, Perception, and Learning (RPL), KTH, 114 28 Stockholm, Sweden (e-mail: haofeil@kth.se; yifeid@kth.se; zehang@kth.se; fpokorny@kth.se; dani@kth.se).

Jens Lundell is with Robotics, Perception, and Learning (RPL), KTH, 114 28 Stockholm, Sweden, and also with Robotics and Autonomous Systems, University of Turku, 20014 Turku, Finland (e-mail: jens.lundell@utu.fi).

Project page: <https://yulihn.github.io/SeqGrasp/>.

This article has supplementary downloadable material available at <https://doi.org/10.1109/LRA.2025.3614051>, provided by the authors.

Digital Object Identifier 10.1109/LRA.2025.3614051

methods on a real Allegro Hand attached to a Franka Panda Robot. These results align with the simulation findings, demonstrating that SeqGrasp reaches a 43.33% higher grasp success rate than MultiGrasp.

Our contributions can be summarized as follows:

- SeqGrasp, a novel algorithm for sequential multi-object grasp generation.
- SeqDataset, a large-scale dataset for sequential multi-object dexterous grasping.
- SeqDiffuser, an efficient conditional sequential grasp diffusion model.
- Extensive simulation and real-world experiments demonstrating the feasibility and effectiveness of SeqGrasp.

## II. RELATED WORK

The work presented here spans dexterous grasping, multi-object grasping, and datasets for dexterous grasping and we review these areas below.

### A. Dexterous Grasping

**Analytical methods:** Early research in dexterous grasping generated stable grasps by optimizing a grasp quality metric such as the force-closure metric [15], [16]. Although these methods are theoretically sound, they are computationally demanding because (i) the many DoF dexterous hands cause high-dimensional search spaces [17], and (ii) the quality metrics are expensive to compute [18]. Consequently, some methods have focused on reducing the search space by imposing constraints on the hand [19], [20] or restricting it to joint configuration subspaces [17]. Another line of work has proposed a computationally cheap differentiable force-closure estimator [7], [18], which has the advantage of being hand-agnostic. This work extends the differentiable force-closure measure from [7] to sequential multi-object grasping.

**Data-driven methods:** Recent advancements in machine learning have significantly improved dexterous grasp generation [3], [4], [6], [21]. Nowadays, deep generative models can generate thousands of dexterous grasps on previously unseen and partially observed objects within seconds [3], [4], [6], [21], something analytical methods cannot. Still, only a few data-driven dexterous grasping methods have been developed for multi-object grasping [13], [22]. In this work, we train a new diffusion-based sequential multi-object grasp sampler SeqDiffuser inspired from [3] on our own optimization-generated sequential multi-object dataset SeqDataset.

### B. Multi-Object Grasping

Multi-object grasping presents unique challenges due to the complex multi-object interactions and the high-dimensional configuration space spanned by the hand and the objects. Some prior parallel-jaw multi-object grasping methods [23], [24] explored multi-object push grasps where scattered objects are first pushed together to facilitate multi-object grasping. However, these methods are limited by their reliance on simple shape primitives and parallel-jaw grippers. In comparison, our work can handle objects of diverse shapes and sizes.

A few works have addressed dexterous multi-object grasping [13], [22], [25] where [13] targets simultaneous multi-object grasping while [22], [25] targets sequential multi-object grasping. Li et al. [13] proposed MultiGrasp a two-stage simultaneous

multi-object dexterous grasping framework where a generative grasp sampler proposed poses to simultaneously pick many objects, followed by a learned policy for executing the pick. Simultaneously grasping multiple objects is an efficient strategy because it avoids repeated hand repositioning, which reduces total execution time, but the main limitation is that objects must be spatially close and of similar size and shape. In comparison, our method can handle scattered objects of different shapes and sizes by sequentially picking one at a time. The other works that do sequential multi-object grasping [22], [25] restrict the grasping to a maximum of two objects [22] or to primitive object shapes such as cylinders or spheres [25]. In comparison, our method can handle up to four objects of complex shapes and sizes. We also use our method to collect the largest sequential multi-object grasping dataset to date.

### C. Datasets for Dexterous Grasping

Large-scale dexterous grasp datasets [5], [7], [13], [14], [21], [26], [27], [28] have significantly advanced the training of data-driven dexterous grasping methods. However, most of these datasets target single-object grasping [5], [7], [21], [26], [27], [28], with only a few for multi-object grasping [13], [14]. As shown in Table I, these existing multi-object datasets are small and predominantly focus on simultaneous rather than sequential grasping. Therefore, we collect the new large-scale sequential multi-object grasping dataset SeqDataset<sup>1</sup> using our method SeqGrasp. SeqDataset is, to date, by far the largest multi-object grasping dataset.

## III. PROBLEM FORMULATION

The problem addressed in this work is sequential multi-object grasping, which we define as follows:

*Definition 1 (Sequential multi-object grasping):* A sequential multi-object grasp is a grasp where one object is grasped at a time using a subset of the dexterous hand’s DoF, while previously grasped objects, if any, remain fixed to the hand.

To contrast, *simultaneous multi-object grasping* addresses how to grasp multiple objects *simultaneously*, typically utilizing all the DoF of the hand [13].

We formulate the sequential multi-object grasping problem as generating a sequence of  $N$  grasps  $\mathcal{G} = \{\mathbf{g}_i\}_{i=1}^N$  for picking a sequence of  $N$  objects  $\mathcal{O} = \{\mathbf{O}_i\}_{i=1}^N$ , where each  $\mathbf{g}_n \in \mathcal{G}$  is restricted to a specific subset  $\mathcal{OS}_n$  of the hand’s total DoF and  $N \geq 2$ . Mathematically, this can be described as

$$\operatorname{argmin}_{\mathbf{g}_n} E(\mathbf{g}_n, \mathbf{O}_n, \mathcal{G}_{n-1}, \mathcal{O}_{n-1}, \mathcal{OS}_n), \forall n = 1, \dots, N, \quad (1)$$

where  $\mathcal{G}_{n-1} = \{\mathbf{g}_i\}_{i=1}^{n-1}$ ,  $\mathcal{O}_{n-1} = \{\mathbf{O}_i\}_{i=1}^{n-1}$ ,  $\mathcal{G}_0 = \emptyset$ , and  $\mathcal{O}_0 = \emptyset$ .  $E$  in (1) is a differentiable function that quantifies how well grasp  $\mathbf{g}_n$  can pick object  $\mathbf{O}_n$  with the DoF  $\mathcal{OS}_n$  given all previously generated grasps  $\mathcal{G}_{n-1}$  and objects  $\mathcal{O}_{n-1}$ . We assume all previously grasped objects  $\mathcal{O}_{n-1}$  remain fixed in the hand when generating a grasp on the next object  $\mathbf{O}_n$ .

In this work, we represent  $\mathcal{OS}_n$  as an opposition space (Section IV-A), each object  $\mathbf{O} \in \mathcal{O}$  as a triangular mesh, and each grasp  $\mathbf{g} \in \mathcal{G}$  as a vector  $\mathbf{g} = [\mathbf{p}, \mathbf{r}, \boldsymbol{\theta}] \in \mathbb{R}^{9+K}$ , where  $\mathbf{p} \in \mathbb{R}^3$  is the hand’s base position,  $\mathbf{r} \in \mathbb{R}^6$  is the hand’s base orientation in a 6D continuous representation [29], and  $\boldsymbol{\theta} \in \mathbb{R}^K$

<sup>1</sup>The dataset is available at <https://yulihn.github.io/SeqGrasp/>.

TABLE I  
MULTI-OBJECT GRASPING DATASETS

Dataset	Hand	Data Collection	Objects	Obj. Set/Seq.	Grasps	Strategy	Max. Obj. Grasped	Open-source
MOG [14]	Human&Barrett	Sim./Real	11	\	28K	Simu./Seq.	\	✗
Grasp'Em [13]	Shadow	Sim.	8	36	90K	Simultaneous	2	✓
SeqDataset(Ours)	Allegro	Sim.	509	2400	870K	Sequential	4	✓

### Algorithm 1: SeqGrasp.

```

Input : Object sequence  $\mathcal{O}$ , OSes  $\mathcal{OS}$ ,  $N_{step}$ , and  $p_{accept}$ .
Output: The optimized grasp sequence  $\mathcal{G}^*_n$ .
1  $n = 1$ ;
2 while  $\mathcal{OS} \neq \emptyset$  and  $n \leq N$  do
3    $\mathcal{OS}_n \sim \mathcal{U}(\mathcal{OS})$ ;
4    $\{\mathbf{x}_j\}_{j=1}^2 \sim \mathcal{U}(\mathcal{S}_n)$ ;
5   for  $s = 1$  to  $N_{step}$  do
6      $\Delta = \partial E(\mathbf{g}_n, \mathbf{O}_n, \mathcal{G}_{n-1}, \mathcal{O}_{n-1}, \{\mathbf{x}_j\}_{j=1}^2) / \partial \mathbf{g}_n$ ;
7      $\mathbf{g}_n \leftarrow \text{MALA}(\mathbf{g}_n, \mathbf{J}_n, \Delta)$ ;
8      $\{\mathbf{x}_j\}_{j=1}^2 \sim f(\mathcal{S}_n, p_{accept})$ ;
9   end
10   $\mathcal{OS} \leftarrow \mathcal{OS} \setminus \mathcal{OS}_n$ ;
11  for  $\mathcal{OS}_j \in \mathcal{OS}$  do
12     $\mathbf{J}_j \leftarrow \mathbf{J}_j \odot (\mathbf{1} - \mathbf{J}_n)$ ;
13    if  $\mathbf{J}_j = \mathbf{0}$  then
14       $\mathcal{OS} \leftarrow \mathcal{OS} \setminus \mathcal{OS}_j$ ;
15    end
16  end
17   $n += 1$ ;
18 end

```



Fig. 2. Multi-object grasping visualizations for different hands. From left to right: Allegro Hand, Shadowhand, MANO.

is the  $K$ -dimensional hand joint angles which are 16 for the Allegro Hand. We assume the shape of all objects in  $\mathcal{O}$  to be fully known. Next, we will introduce SeqGrasp our algorithm for solving (1).

## IV. SEQUENTIAL GRASP GENERATION

Here, we present Algorithm 1 for sequential grasp generation. It includes (i) an opposition space selection strategy (Section IV-A), (ii) an optimization-based grasp synthesis method (Section IV-B), and (iii) an energy-based cost function (Section IV-C). Fig. 2 shown an example of running Algorithm 1 to grasp three different objects with three different dexterous hands.

### A. Opposition Space Selection Strategy

The primary objective in sequential multi-object grasping is to maximize the hand's remaining DoF after each grasp. For this purpose, we propose a grasp planning strategy guided by OSs [11], [25], [30], [31]. An OS is a functional subspace within the hand's kinematic structure formed by pairs of opposing




Fig. 3. (a) Grasps using all seven OSes. From left to right, first row: middle-ring, index-middle, and thumb-index, second row: ring-palm, middle-palm, index-palm, and thumb-palm. (b) Visualization of contact point candidates on Allegro Hand surface. Cyan and pink points denote palm opposition and side opposition contacts, respectively.

surfaces (such as fingertips, lateral surfaces of fingers, or palm surfaces) along with the joints that control these surfaces [25]. It represents regions where opposing forces can be applied to create stable grasps. The number of OSs are hand-dependent and vary based on the kinematic structure. Fig. 3(a) shows the seven different OSs for the Allegro Hand.

Mathematically, each opposition space can be represented as a pair  $\mathcal{OS}_i = \{\mathbf{J}_i, \mathcal{S}_i\}$ , where  $\mathbf{J}_i \in \{0, 1\}^K$  is a binary vector indicating which joints are involved in controlling the opposition space, and  $\mathcal{S}_i \in \mathbb{R}^{3 \times M_i}$  represents the 3D points on the hand where opposing forces can be applied. Fig. 3(b) shows an example of two different  $\mathcal{S}_i$  for the Allegro Hand, where palm and pad oppositions have contact points located on the inner surfaces of fingers and palm and side oppositions have contact points on the fingers' lateral surfaces.

Let  $\mathcal{OS} = \{\mathcal{OS}_i\}_{i=1}^L$  be the set of all OSs. Given this set, Algorithm 1 samples a random OS from it (Line 3) and uses it for subsequent grasp generation (Section IV-B). Once grasp generation is complete, the sampled OS can no longer be used and is thus removed from the available OSs (Line 10).  $\mathbf{J}_i$  of all the remaining OSs are also updated by zeroing out the joints used in  $\mathcal{OS}_n$  (Line 12). Subsequently, all OSs with  $\mathbf{J} = \mathbf{0}$ , meaning that no more controllable joints exist, are removed (Line 14). For instance, in the case of the Allegro Hand, if the thumb-index OS is selected, then both the thumb-palm and index-palm OSs become unavailable due to shared joint constraints.

### B. Optimization-Based Grasp Generation

The next step in the algorithm (Lines 5–9) is to generate a stable, physically plausible, and collision-free grasp that respects the sampled  $\mathcal{OS}_n$ . To achieve this, we formulate  $E$  in (1) as an energy function (Section IV-C) and numerically optimize it using the Metropolis-Adjusted Langevin Algorithm (MALA) [32] (Line 7).

In robotic grasping, MALA has been used to optimize single object grasps  $\mathbf{g}_n$  by iteratively refining  $\mathbf{p}_n$ ,  $\mathbf{r}_n$  and  $\boldsymbol{\theta}_n$  according to Langevin dynamics [7], [18]. However, we must




Fig. 4. (a) **Initialization.** The initial grasp configurations are randomly sampled on the expanded convex hull of the object (bottle) while the previously grasped object (ball) remains in the hand. (b) **Optimization.** During optimization, the grasp is incrementally refined, ultimately securing the target object using the ring-palm OS.

adapt MALA to sequential multi-object grasping. To this end, we propose a new grasp as  $\hat{\mathbf{g}}_n \leftarrow \mathbf{g}_n - \gamma[\mathbf{1}, \mathbf{J}_n] \odot \Delta$ , where  $\gamma$  is the step size,  $\mathbf{1} \in \mathbb{R}^9$  is a padding vector to align the length of  $\mathbf{J}$  with  $\mathbf{g}$ ,  $\Delta = \partial E / \partial \mathbf{g}_n$  is the energy gradient, and  $\odot$  is the element-wise (Hadamard) product.  $\hat{\mathbf{g}}_n$  is accepted if  $\alpha \geq u$ , where  $u \sim \mathcal{U}([0, 1])$  and

$$\alpha = \frac{E(\hat{\mathbf{g}}_n, \mathbf{O}_n, \mathcal{G}_{n-1}, \mathcal{O}_{n-1}, \{\mathbf{x}_j\}_{j=1}^2)}{E(\mathbf{g}_n, \mathbf{O}_n, \mathcal{G}_{n-1}, \mathcal{O}_{n-1}, \{\mathbf{x}_j\}_{j=1}^2)}. \quad (2)$$

The above procedure is repeated for a fixed number of steps where, at each step,  $\{\mathbf{x}_j\}_{j=1}^2$  is re-sampled with probability  $p_{\text{accept}}$  (Line 8). This resampling process helps accelerate convergence and escape from local minimas [7], [18].

We initialize  $\mathbf{g}_n$  at a randomly sampled position on the expanded convex hull of the target object  $\mathbf{O}_n$  as exemplified in Fig. 4(a). If  $n = 1$ , then  $\theta_1$  is set to a natural open-hand and collision-free posture, as demonstrated in Fig. 3(b), all fingers are fully extended to create sufficient space for grasping objects with any OS. While for  $n \geq 2$ ,  $\theta_n = \theta_{n-1}$ . A visual example of the optimization process when grasping a second object is shown in Fig. 4(b).

### C. Energy Function

Numerically optimizing the energy function in (1) should result in stable, collision-free, joint-respecting, and OS-respecting grasps. We design the following energy function to capture all of these behaviors

$$E = \mathbf{w}^T [E_{\text{fc}} \quad E_{\text{dis}} \quad E_{\text{hop}} \quad E_{\text{hsp}} \quad E_{\text{joint}} \quad E_{\text{oop}}]^T, \quad (3)$$

where  $\mathbf{w} \in \mathbb{R}^6$  is a weight vector controlling the relative importance of the force-closure  $E_{\text{fc}}$ , contact distance  $E_{\text{dis}}$ , hand-object penetration  $E_{\text{hop}}$ , hand self-penetration  $E_{\text{hsp}}$ , joint limits  $E_{\text{joint}}$ , and object-object penetration  $E_{\text{oop}}$  energy terms.

The force-closure term ( $E_{\text{fc}}$ ) encourages the grasp to be in force-closure equilibrium [33]. The definition for ( $E_{\text{fc}}$ ) originates from [7], where zero friction and uniform contact force magnitudes are assumed because it significantly speeds up the computation. We follow the same definition and write it as

$$E_{\text{fc}}(\{\mathbf{x}_j\}_{j=1}^2) = \|\mathbf{G}\mathbf{c}\|^2, \quad (4)$$

where  $\mathbf{c} = [\mathbf{c}_1^T, \mathbf{c}_2^T]^T \in \mathbb{R}^{6 \times 1}$  represents the concatenated contact normals at each contact point  $\{\mathbf{x}_j\}_{j=1}^2$ .  $\mathbf{G}$  is defined as:

$$\mathbf{G} = \begin{bmatrix} \mathbf{I} & \mathbf{I} \\ [\mathbf{x}_1]_{\times} & [\mathbf{x}_2]_{\times} \end{bmatrix}, \quad (5)$$

where  $\mathbf{I}$  represents the identity matrix, and  $[\mathbf{x}_j]_{\times}$  ( $1 \leq j \leq 2$ ) denotes the skew-symmetric matrix formed from the contact point  $\mathbf{x}_j$ .

The contact distance and penetration terms ( $E_{\text{dis}}$  &  $E_{\text{hop}}$ ) encourage the hand-object contacts to occur close to the object surface without penetrating it. The contact distance is mathematically defined as

$$E_{\text{dis}}(\{\mathbf{x}_j\}_{j=1}^2, \mathbf{O}_n) = \sum_{j=1}^2 d(\mathbf{x}_j, \mathbf{O}_n), \quad (6)$$

where  $d(\mathbf{x}_j, \mathbf{O}_n) = \min_{\mathbf{v} \in \mathbf{O}_n} \|\mathbf{x}_j - \mathbf{v}\|_2$  is the shortest point-mesh distance. Similarly, the hand-object penetration term is defined as:

$$E_{\text{hop}}(\mathbf{g}_n, \mathbf{O}_n) = \sum_{\mathbf{v} \in \mathcal{V}_{\text{hop}}(\mathbf{H}_{\mathbf{g}}, \mathbf{O}_n)} d(\mathbf{v}, \mathbf{O}_n), \quad (7)$$

where  $d(\mathbf{v}, \mathbf{O}_n) = \min_{\mathbf{v}_1 \in \mathbf{O}_n} \|\mathbf{v} - \mathbf{v}_1\|_2$  and  $\mathcal{V}_{\text{hop}}(\mathbf{H}_{\mathbf{g}}, \mathbf{O}_n)$  is the set of points on the hand surface pointcloud  $\mathbf{H}_{\mathbf{g}} \in \mathbb{R}^{3 \times M_h}$  that penetrate the object  $\mathbf{O}_n$ .

The self-collision and joint limit terms ( $E_{\text{hsp}}$  &  $E_{\text{joint}}$ ) encourage physical feasibility. We define these as

$$E_{\text{hsp}}(\mathbf{g}_n) = \sum_{\mathbf{v}_1, \mathbf{v}_2 \in \mathcal{V}_{\text{hsp}}(\mathbf{H}_{\mathbf{g}}), \mathbf{v}_1 \neq \mathbf{v}_2} \max(\|\mathbf{v}_1 - \mathbf{v}_2\|_2, 0), \quad (8)$$

$$E_{\text{joint}}(\mathbf{g}_n) = \|(\theta - \theta^{\text{upper}})^+\|_1 + \|(\theta^{\text{lower}} - \theta)^+\|_1, \quad (9)$$

where  $\mathcal{V}_{\text{hsp}}(\mathbf{H}_{\mathbf{g}})$  denotes all surface points of the hand that are self-penetrating,  $(\cdot)^+$  denotes the element-wise operation  $\max(\cdot, 0)$ , and  $\theta^{\text{upper}}$  and  $\theta^{\text{lower}}$  denote the upper and lower limits of all joints.

Finally, the term ( $E_{\text{oop}}$ ) minimizes object-object penetration. It is defined as

$$E_{\text{oop}}(\{\mathbf{O}_i\}_{i=1}^n) = \sum_{i=1}^{n-1} \sum_{\mathbf{v} \in \mathcal{V}_{\text{oop}}(\mathbf{O}_i, \mathbf{O}_n)} d(\mathbf{v}, \mathbf{O}_n), \quad (10)$$

where  $d(\mathbf{v}, \mathbf{O}_n) = \min_{\mathbf{v}_1 \in \mathbf{O}_n} \|\mathbf{v} - \mathbf{v}_1\|_2$ , and  $\mathcal{V}_{\text{oop}}(\mathbf{O}_i, \mathbf{O}_n)$  are the inter-penetrating surface points between the previously grasped object  $\mathbf{O}_i$  and the current object  $\mathbf{O}_n$ .

## V. DATASET GENERATION

We use SeqGrasp to generate our large-scale dataset SeqDataset<sup>2</sup> containing 4.9 million sequential Allegro Hand grasps on over 509 objects from DexGraspNet [7]. Each object is resized to fit within the Allegro Hand by scaling its axis-aligned bounding box to be between 0.06 and 0.10 meters. Then, from the 509 resized objects, we generate 600 unique object sets containing four objects each. The objects in these sets are randomly permuted four times, resulting in 2,400 unique object sequences.

We run Algorithm 1 on the 2,400 unique object sequences with  $\mathbf{w} = [50, 50, w_{\text{hop}}, 5, 1, 5]^T$ , where  $w_{\text{hop}}$ , that penalizes hand-object penetration, grows linearly from 5 to  $5e^2$ . The optimization runs for 6,000 iterations per grasp. We validate

<sup>2</sup>The dataset is available at <https://yulihn.github.io/SeqGrasp/>.

TABLE II  
STATISTICS OF SEQDATASET. ALL NUMBERS OF GRASPS ARE SHOWN IN THOUSAND

OSes	Success	Total	Success Rate (%)
Middle-Ring	100.83	704.97	14.30
Index-Middle	97.43	700.90	13.90
Thumb-Index	97.29	1691.41	5.75
Ring-Palm	147.33	323.31	45.57
Middle-Palm	152.77	404.84	37.74
Index-Palm	194.13	546.60	35.52
Thumb-Palm	85.86	485.80	17.67

Num. Obj. Grasped	Consumed	Total	Consumed Rate (%)
One	0.00	444.72	0.00
Two	9.75	251.52	3.88
Three	69.20	139.31	49.68
Four	40.09	40.09	100.00

the generated grasp sequences in the physics simulator Isaac Gym following the setup in [7]. This setup initializes all object densities to 500 kg/m<sup>3</sup>, the friction coefficient to 2.0, the hand to the generated grasp configuration, and objects as free-floating. Then, the hand closes until contact is established with the object, which happens when the distance between the hand and the object is less than 2mm. Finally, a grasp stability test is evaluated by linearly applying an acceleration of 9.8 m/s<sup>2</sup> in all six orthogonal directions for 100 consecutive simulation steps. A grasp sequence is successful if all grasped objects remain in contact with the hand after the grasp stability test and the maximum penetration depth is less than 1cm. Otherwise, the grasp sequence is a failure. This procedure ultimately produced 4.9 million grasps, of which 870K (17.82%) were successful. Only successful grasps were retained in SeqDataset. The entire data generation process requires approximately 2,500 GPU hours on an NVIDIA A100.

The statistics of SeqDataset are presented in Table II. The results demonstrate that grasps using palm opposition achieved significantly higher success rates, suggesting that these OSs are important in sequential multi-object grasping. Notably, grasps using thumb-index and thumb-palm OSs display lower success rates than other OSs, which deviates from previous research findings that underscore the thumb’s dominant role in human hand manipulation tasks [34]. We hypothesize that this discrepancy may stem from the biomechanical differences between the Allegro and the human hand. Finally, SeqDataset’s grasp consumption, which indicates no more available OSs ( $\mathcal{OS} = \emptyset$ ), aligns with its objective of supporting multi-object grasping, with a significant portion of cases (96.12%) showing potential for sequential grasping of three to four objects, indicating efficient utilization of the hand’s DoF.

## VI. CONDITIONAL SEQUENTIAL GRASP DIFFUSER

Finally, we introduce SeqDiffuser, a diffusion-based sequential grasp generation method trained on SeqDataset. The architecture of SeqDiffuser is similar to the sampler from [3], with the distinction that SeqDiffuser is also conditioned on the OS to enable sequential multi-object grasping.

For generating the grasp  $\mathbf{g}_n$  on the  $n$ -th object  $\mathbf{O}_n$ , we first sample  $\mathcal{OS}_n \sim \mathcal{OS}$ . The selected OS is encoded as a one-hot feature vector  $\mathbf{f}_n \in \{\mathbf{f}^i\}_{i=1}^L$ . This feature vector is concatenated with the grasp  $\mathbf{g}_n$  forming the augmented grasp representation  $\tilde{\mathbf{g}}_n = [\mathbf{f}_n, \mathbf{p}_n, \mathbf{r}_n, \boldsymbol{\theta}_n]$ . Of all the elements in  $\tilde{\mathbf{g}}_n$ , we only want to diffuse  $\mathbf{p}_n, \mathbf{r}_n$ , and the subset of  $\boldsymbol{\theta}$  that corresponds to  $\mathcal{OS}_n$ , i.e.,  $\mathbf{J}_n$ . Therefore, we create the binary diffusion selection vector

$\mathbf{k}_n = [\mathbf{0}, \mathbf{1}, \mathbf{J}_n]$ , where  $\mathbf{0} \in \mathbb{R}^L$  and  $\mathbf{1} \in \mathbb{R}^9$ . Note that in the following, the subscript  $t$  of  $\mathbf{g}_t$  denotes the timestep in the diffusion process, while the sequential grasp step  $n$  is omitted from  $\tilde{\mathbf{g}}_n$  and  $\mathbf{k}_n$  for clarity.

The forward process for adding noise to a successful grasp  $\tilde{\mathbf{g}}_0$  over  $T$  timesteps is

$$q(\tilde{\mathbf{g}}_{1:T}|\tilde{\mathbf{g}}_0) = \prod_{t=1}^T q(\tilde{\mathbf{g}}_t|\tilde{\mathbf{g}}_{t-1}), \quad (11)$$

$$q(\tilde{\mathbf{g}}_t|\tilde{\mathbf{g}}_{t-1}) = \mathcal{N}(\tilde{\mathbf{g}}_t; \mathbf{k} \odot \sqrt{1 - \beta_t} \tilde{\mathbf{g}}_{t-1}, \beta_t \mathbf{I}), \quad (12)$$

where  $\beta_t$  is the scheduled noise variance at time step  $t$ . To reconstruct the original  $\tilde{\mathbf{g}}_0$  from  $\tilde{\mathbf{g}}_T$ , SeqDiffuser learns the reverse diffusion process by estimating the Gaussian noise at each step  $t$  and progressively removing it:

$$p_\psi(\tilde{\mathbf{g}}_{0:T}|\mathbf{h}_\mathbf{O}) = p(\tilde{\mathbf{g}}_T) \prod_{t=1}^T p_\psi(\tilde{\mathbf{g}}_{t-1}|\tilde{\mathbf{g}}_t, \mathbf{h}_\mathbf{O}), \quad (13)$$

$$p_\psi(\tilde{\mathbf{g}}_{t-1}|\tilde{\mathbf{g}}_t, \mathbf{h}_\mathbf{O}) = \mathcal{N}(\tilde{\mathbf{g}}_{t-1}; \mathbf{k} \odot \hat{\mu}_\psi(\tilde{\mathbf{g}}_t, \mathbf{h}_\mathbf{O}, t), \hat{\Sigma}_\psi(\tilde{\mathbf{g}}_t, \mathbf{h}_\mathbf{O}, t)), \quad (14)$$

where  $\hat{\mu}_\psi(\tilde{\mathbf{g}}_t, \mathbf{h}_\mathbf{O}, t)$  and  $\hat{\Sigma}_\psi(\tilde{\mathbf{g}}_t, \mathbf{h}_\mathbf{O}, t)$  are the learnable mean and variance of a Gaussian distribution and  $\mathbf{h}_\mathbf{O}$  is the Basis Point Set (BPS) [35] encoded feature of the target object  $\mathbf{O}$ . In the forward and reverse processes, noise is only added or removed at positions where  $\mathbf{k}$  is nonzero. The loss function is thereby formulated as:

$$\mathcal{L}_\epsilon = \|\mathbf{k} \odot \hat{\epsilon}_t - \mathbf{k} \odot \epsilon_t\|^2, \quad (15)$$

where  $\hat{\epsilon}_t = \epsilon_\psi(\tilde{\mathbf{g}}_t, \mathbf{h}_\mathbf{O}, t)$  and  $\epsilon_t$  are the estimated noise and ground-truth noise, respectively.

## VII. EXPERIMENTS

We experimentally evaluate SeqGrasp and SeqDiffuser in both simulation and the real world. The specific questions we aim to address with the experiments are:

- 1) How well can SeqGrasp and SeqDiffuser generate successful and diverse grasps?
- 2) What is the difference between simultaneously and sequentially grasping multiple objects?
- 3) Are the generated grasps executable on real hardware?

We compare our method to the optimization-based sampler MultiGrasp from [13], which generates simultaneous grasps on clustered objects. As such, for MultiGrasp, we must first sample clustered object configurations and then generate multi-object grasps directly on the object cluster. In contrast, SeqGrasp and SeqDiffuser do not require objects to be spatially close as they generate grasps sequentially based on previously successful ones. While this comparison is not entirely fair, we believe that comparing these two strategies still offers valuable insights.

For training SeqDiffuser, we split SeqDataset into an 80% training set and a 20% test set, ensuring that no training objects were used in the experimental evaluation. The object point clouds are obtained by sampling 2048 points on the object mesh surfaces using Farthest Point Sampling (FPS).

TABLE III

**SIMULATION RESULTS.** MULG, SEQG, SEQD, G'EM, AND SDATA ARE SHORT FOR THE MULTIGRASP, SEQGRASP, SEQDIFFUSER, GRASP'EM, AND SEQDATASET, RESPECTIVELY. THE  $-I$  FOLLOWING THE METHOD NAME DENOTES THE NUMBER OF OBJECTS USED FOR GRASP GENERATION.  $\uparrow$  ( $\downarrow$ ), THE HIGHER (LOWER), THE BETTER

Method	SR $\uparrow$		Pene. $\downarrow$		Eff. $\downarrow$	Div. $\uparrow$
	SData	G'Em	SData	G'Em	Avg.	Avg.
MulG-1	<b>66.84</b>	<b>65.39</b>	<b>1.14</b>	<b>1.27</b>	600	0.284
SeqD-1	46.95	45.23	5.55	5.59	<b>0.8</b>	0.321
SeqG-1	50.04	40.78	1.73	2.16	900	<b>0.332</b>
MulG-2	<b>22.46</b>	16.48	2.30	2.83	750	0.347
SeqD-2	18.83	23.00	5.84	6.28	<b>0.8</b>	0.359
SeqG-2	21.21	<b>32.03</b>	<b>2.14</b>	<b>1.78</b>	900	<b>0.367</b>
MulG-3	10.78	3.55	3.39	4.04	900	0.340
SeqD-3	11.05	9.22	5.93	6.47	<b>0.8</b>	0.334
SeqG-3	<b>19.49</b>	<b>21.05</b>	<b>2.23</b>	<b>2.23</b>	900	<b>0.349</b>
MulG-4	0.90	0.47	5.17	6.27	1000	<b>0.329</b>
SeqD-4	<b>3.01</b>	1.68	6.18	6.69	<b>0.8</b>	0.293
SeqG-4	2.93	<b>5.04</b>	<b>2.70</b>	<b>2.62</b>	900	0.312



Fig. 5. **Qualitative results.** For SeqGrasp (first row) and SeqDiffuser (second row), the grasp sequences are visually indicated by a color gradient, transitioning from lighter to darker shades. In contrast, for MultiGrasp (third row), the color gradient is only used to differentiate the objects. For SeqGrasp and SeqDiffuser, we only show consumed grasps, that is, when  $\mathcal{O}_S = \emptyset$ .

### A. Simulation Experiments

We evaluated all generated grasps in the simulation experiments in Isaac Gym [36]. We used two object sets: (1) all eight objects from Grasp'Em [13] and (2) a random selection of eight validation objects from SeqDataset. We randomly generated ten four-object sequences for each object set, and, per object, we generated 256 grasps, resulting in 10,240 grasps per method.

We used the following metrics to assess the quality of the generated grasps:

- 1) **Success rate (SR) in percent:** The same success criteria as described in Section V.
- 2) **Maximum penetration depth (Pene.) in mm:** The maximum interpenetration distance between the hand and all grasped objects.
- 3) **Diversity (Div.) in radians:** Grasp diversity is determined by calculating the standard deviation of  $g$  across all successful grasps.
- 4) **Efficiency (Eff.) in seconds:** The computational time required to generate a batch of 256 grasps on an NVIDIA A100.

The quantitative results are presented in Table III, while Fig. 5 qualitatively illustrates a few grasps. The results demonstrate that SeqGrasp achieves the highest success rate and the lowest penetration depth when grasping two or more objects.

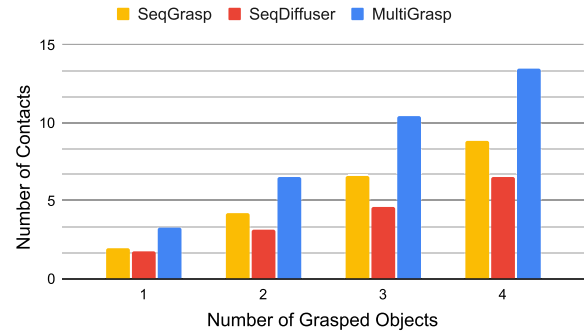


Fig. 6. **Number of hand-object contacts** for different methods.




Fig. 7. Left: Two object sets for the grasp stability test. Right: The dynamic stability test procedure.

However, the lower performance of SeqGrasp on single-object grasps is attributed to the fewer contact points compared to MultiGrasp, as shown in Fig. 6. For single-object grasps, SeqGrasp and SeqDiffuser establish approximately two contacts on average, whereas MultiGrasp achieves three or more due to it engaging all the available DoF. However, as the number of objects increases, the overall hand-object penetration increases for MultiGrasp but not for SeqGrasp.

We observe a notable decline in performance across all methods when transitioning from three-object to four-object grasps. We hypothesize that this decline occurs because the three previously grasped objects occupy substantial space within the Allegro Hand, pushing the fourth object grasp to the limits of the hand's kinematic redundancy. Additionally, as the number of grasped objects increases, object-object interactions grow exponentially, making the task considerably more challenging, a finding also reported in [13]. Nevertheless, SeqGrasp demonstrates superior performance in scenarios involving three or more objects.

Another notable observation is that SeqDiffuser generates grasps with high penetration, which aligns with previous work on single-object diffusion-based grasp sampling [3]. Still, SeqDiffuser is valuable because it generates grasps 750–1250 times faster than SeqGrasp and MultiGrasp.

### B. Real-World Experiments

To evaluate grasp stability in the real world, we tested the generated grasps on a real Allegro Hand mounted to a Franka Emika Panda robotic arm. We evaluated the grasps on the two object sets visualized in Fig. 7: one comprising eight 3D printed SeqDataset test objects, and another comprising eight everyday




Fig. 8. **Grasp replication on real hardware.** First row: three-object grasp, second row: four-object grasp.

TABLE IV

**REAL-WORLD EXPERIMENTAL RESULTS.** EACH ENTRY REPORTS THE NUMBER OF SUCCESSFUL TRIALS OUT OF 10 AND THE AVERAGE NUMBER OF OBJECTS SUCCESSFULLY GRASPED IN MEAN  $\pm$  STD AFTER THE TEST

Real exp.		SeqG-3	SeqG-4	SeqD-3	SeqD-4	MulG-3	MulG-4
Printed Obj.	Succ. trials	5/10	5/10	4/10	0/10	1/10	1/7
	Succ. objects	2.3 $\pm$ 0.9/3	3.4 $\pm$ 0.8/4	2.1 $\pm$ 0.9/3	1.5 $\pm$ 0.8/4	1.2 $\pm$ 1.1/3	1.6 $\pm$ 1.5/4
Real Obj.	Succ. trials	5/10	/	0/10	/	0/6	/
	Succ. objects	2.1 $\pm$ 1.2/3	/	1.8 $\pm$ 0.4/3	/	1.0 $\pm$ 0.6/3	/

objects. To enhance surface friction, we coated the printed object with silicone gel and attached an anti-slip grip to the palm.

We followed the procedure outlined in [25] to replicate the grasp on the real hardware. This procedure involves positioning each object as closely as possible with its generated pose and then closing the relevant joints of the hand to retain the objects, as exemplified in Fig. 8. For grasps with severe penetration issues, we placed the object in the closest feasible real-world position. For MultiGrasp and SeqGrasp, we replicated the successful grasps with the lowest energy, while for SeqDiffuser, we randomly chose one of the generated grasps. Due to the larger real-world objects, four-object grasping exceeded the workspace limits of the Allegro Hand. Thus, we only evaluated three-object grasps on that object set. In contrast, both three- and four-object grasping were tested on the 3D-printed object set.

We evaluated the grasps with the grasp stability test shown in Fig. 7. In this test, the arm followed a randomly generated dynamic trajectory executed with the highest achievable acceleration. The arm was then moved to a holding position with the palm facing downward. From this position, the last joint rotated by  $\pm 90^\circ$ . This test comprehensively assessed the grasp stability by evaluating both translational and rotational disturbances.

The experimental results are presented in Table IV. Similar to the simulation results, SeqGrasp outperformed SeqDiffuser and MultiGrasp in both three-object and four-object grasping tasks, achieving an average success rate of 50%. The low success rate of SeqDiffuser (13%) is primarily due to substantial object inter-penetration in the generated grasps, which meant that the closest physically feasible grasp replicable on the real hand differed substantially from the intended configuration. MultiGrasp performed worst, with an average success rate of 8.7%. The reason is its reliance on object-object contacts for grasp stability, which are sensitive to minor object displacements. This resulted in frequent failures, particularly in four-object grasping trials where MultiGrasp achieved no successful trials and failed to generate valid grasps in three of ten attempts (last column in Table IV). Accounting for these additional failed trials lowers the average success rate for MultiGrasp from 8.7% to 6.7%.

The real-world experiments highlight that our sequential grasping strategy enables robust and effective dexterous multi-object manipulation. By generating grasps iteratively and without requiring spatial proximity between objects, SeqGrasp and SeqDiffuser mitigate the challenges posed by object interactions, making it a more reliable and practical approach for real-world applications.

## VIII. LIMITATIONS AND FUTURE WORK

Our work has two main limitations. First of all, neither SeqGrasp nor SeqDiffuser accounts for object sequences as OSes and object orders are uniformly sampled. The random sampling results in a diverse dataset that is useful for training and benchmarking, rather than the most optimal grasp sequence. Learning the optimal grasp sequences from human preferences is an interesting future work direction.

The second limitation is that the grasp generation process does not account for motion planning to the grasp pose. Consequently, it proved extremely challenging to find a collision-free path to the intended grasp pose, which is the primary reason no such experiments were conducted. However, we believe this limitation can be resolved by learning a policy conditioned on the final grasp pose using, e.g., reinforcement learning.

## IX. CONCLUSION

We proposed SeqGrasp, an algorithm for sequentially grasping multiple objects with a dexterous hand. SeqGrasp combines OSs and differentiable-force closure to generate stable grasps that maximize the hand’s remaining DoF after each grasp. Using SeqGrasp, we constructed SeqDataset, currently the largest sequential grasping dataset, comprising 870 K validated grasps across 509 diverse objects. This dataset enabled the training of SeqDiffuser, our diffusion-based sequential multi-object grasp sampler. The experimental evaluations demonstrated that SeqGrasp and SeqDiffuser outperformed the simultaneous multi-object grasping baseline MultiGrasp, achieving an 8.71%–43.33% higher average success rate. Moreover, SeqDiffuser proved to be 750–1250 times faster at generating grasps than SeqGrasp and MultiGrasp. In conclusion, this work demonstrates a fast and stable solution for generating sequential multi-object grasps, which we hope will pave the way for further research in multi-object grasping.

## ACKNOWLEDGMENT

The authors would like to express their sincere gratitude to Li Chen for valuable assistance with the experimental hardware.

## REFERENCES

- [1] J. Bohg, A. Morales, T. Asfour, and D. Kragic, “Data-driven grasp synthesis—A survey,” *IEEE Trans. Robot.*, vol. 30, no. 2, pp. 289–309, Apr. 2014.
- [2] R. Newbury et al., “Deep learning approaches to grasp synthesis: A review,” *IEEE Trans. Robot.*, vol. 39, no. 5, pp. 3994–4015, Oct. 2023.
- [3] Z. Weng, H. Lu, D. Kragic, and J. Lundell, “DexDiffuser: Generating dexterous grasps with diffusion models,” *IEEE Robot. Automat. Lett.*, vol. 9, no. 12, pp. 11834–11840, Dec. 2024.
- [4] J. Lu et al., “UGG: Unified generative grasping,” in *Proc. Eur. Conf. Comput. Vis.*, 2024, pp. 414–433.
- [5] J. Lundell, F. Verdoja, and V. Kyrki, “DDGC: Generative deep dexterous grasping in clutter,” *IEEE Robot. Automat. Lett.*, vol. 6, no. 4, pp. 6899–6906, Oct. 2021.

- [6] Z. Wei et al., “ $\mathcal{D}(\mathcal{R}, \mathcal{O})$  grasp: A unified representation of robot and object interaction for cross-embodiment dexterous grasping,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2025, pp. 4982–4988.
- [7] R. Wang et al., “DexGraspNet: A large-scale robotic dexterous grasp dataset for general objects based on simulation,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2023, pp. 11359–11366.
- [8] P. Li et al., “GenDexGrasp: Generalizable dexterous grasping,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2023, pp. 8068–8074.
- [9] D. Song et al., “Predicting human intention in visual observations of hand/object interactions,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2013, pp. 1608–1615.
- [10] J. Romero, H. Kjellström, and D. Kragic, “Hands in action: Real-time 3D reconstruction of hands in interaction with objects,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2010, pp. 458–463.
- [11] T. Feix, J. Romero, H. -B. Schmiemayer, A. M. Dollar, and D. Kragic, “The GRASP taxonomy of human grasp types,” *IEEE Trans. Hum.-Mach. Syst.*, vol. 46, no. 1, pp. 66–77, Feb. 2016.
- [12] M. Kokic, D. Kragic, and J. Bohg, “Learning task-oriented grasping from human activity datasets,” *IEEE Robot. Automat. Lett.*, vol. 5, no. 2, pp. 3352–3359, Apr. 2020.
- [13] Y. Li et al., “Grasp multiple objects with one hand,” *IEEE Robot. Automat. Lett.*, vol. 9, no. 5, pp. 4027–4034, May 2024.
- [14] Y. Sun, E. Amatova, and T. Chen, “Multi-object grasping-types and taxonomy,” in *Proc. Int. Conf. Robot. Automat.*, 2022, pp. 777–783.
- [15] A. T. Miller and P. K. Allen, “Graspit! A versatile simulator for robotic grasping,” *IEEE Robot. Autom. Mag.*, vol. 11, no. 4, pp. 110–122, Dec. 2004.
- [16] C. Rosales, R. Suárez, M. Gabiccini, and A. Bicchi, “On the synthesis of feasible and prehensile robotic grasps,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2012, pp. 550–556.
- [17] M. T. Ciocarlie and P. K. Allen, “Hand posture subspaces for dexterous robotic grasping,” *Int. J. Robot. Res.*, vol. 28, no. 7, pp. 851–867, 2009.
- [18] 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 Robot. Automat. Lett.*, vol. 7, no. 1, pp. 470–477, Jan. 2022.
- [19] J. Ponce, S. Sullivan, J. -D. Boissonnat, and J. -P. Merlet, “On characterizing and computing three- and four-finger force-closure grasps of polyhedral objects,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 1993, pp. 821–827.
- [20] J. -W. Li, H. Liu, and H. -G. Cai, “On computing three-finger force-closure grasps of 2-D and 3-D objects,” *IEEE Trans. Robot. Autom.*, vol. 19, no. 1, pp. 155–161, Feb. 2003.
- [21] V. Mayer, Q. Feng, J. Deng, Y. Shi, Z. Chen, and A. Knoll, “FFHNet: Generating multi-fingered robotic grasps for unknown objects in real-time,” in *Proc. Int. Conf. Robot. Automat.*, 2022, pp. 762–769.
- [22] S. He et al., “Sequential multi-object grasping with one dexterous hand,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2025.
- [23] W. C. Agboh, J. Ichnowski, K. Goldberg, and M. R. Dogar, “Multi-object grasping in the plane,” in *Proc. Int. Symp. Robot. Res.*, 2022, pp. 222–238.
- [24] T. Yonemaru, W. Wan, T. Nishimura, and K. Harada, “Learning to group and grasp multiple objects,” 2025, *arXiv:2502.08452*.
- [25] K. Yao and A. Billard, “Exploiting kinematic redundancy for robotic grasping of multiple objects,” *IEEE Trans. Robot.*, vol. 39, no. 3, pp. 1982–2002, Jun. 2023.
- [26] J. Chen, Y. Ke, and H. Wang, “BODex: Scalable and efficient robotic dexterous grasp synthesis using bilevel optimization,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2025, pp. 1–8.
- [27] H. Zhang, S. Christen, Z. Fan, O. Hilliges, and J. Song, “GraspXL: Generating grasping motions for diverse objects at scale,” in *Proc. Eur. Conf. Comput. Vis.*, 2024, pp. 386–403.
- [28] Y. Liu et al., “RealDex: Towards human-like grasping for robotic dexterous hand,” in *Proc. 33rd Int. Joint Conf. Artif. Intell.*, 2024, pp. 6859–6867.
- [29] Y. Zhou, C. Barnes, J. Lu, J. Yang, and H. Li, “On the continuity of rotation representations in neural networks,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 5745–5753.
- [30] T. Iberall, “Opposition space as a structuring concept for the analysis of skilled hand movements,” *Exp. Brain Res. Ser.*, vol. 15, pp. 158–173, 1986.
- [31] J. B. Smeets, K. van der Kooij, and E. Brenner, “A review of grasping as the movements of digits in space,” *J. Neurophysiol.*, vol. 122, no. 4, pp. 1578–1597, 2019.
- [32] G. O. Roberts and R. L. Tweedie, “Exponential convergence of langevin distributions and their discrete approximations,” *Bernoulli*, vol. 2, pp. 341–363, 1996.
- [33] E. Rimon and J. Burdick, “On force and form closure for multiple finger grasps,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 1996, pp. 1795–1800.
- [34] G. Cotugno, K. Althoefer, and T. Nanayakkara, “The role of the thumb: Study of finger motion in grasping and reachability space in human and robotic hands,” *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 47, no. 7, pp. 1061–1070, Jul. 2017.
- [35] S. Prokudin, C. Lassner, and J. Romero, “Efficient learning on point clouds with basis point sets,” in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 4332–4341.
- [36] V. Makovychuk et al., “Isaac Gym: High performance GPU-based physics simulation for robot learning,” *NeurIPS Datasets Benchmarks*, 2021.