

# GarmentPile++: Affordance-Driven Cluttered Garments Retrieval with Vision-Language Reasoning

Mingleyang Li<sup>1\*</sup>, Yuran Wang<sup>1\*</sup>, Yue Chen<sup>1</sup>, Tianxing Chen<sup>2</sup>  
 Jiaqi Liang<sup>1</sup>, Zishun Shen<sup>1</sup>, Haoran Lu<sup>3</sup>, Ruihai Wu<sup>1†</sup>, Hao Dong<sup>1†</sup>  
 Peking University<sup>1</sup>, The University of Hong Kong<sup>2</sup>, Northwestern University<sup>3</sup>  
 \*Equal contribution †Corresponding author




Fig. 1: **Overview.** GarmentPile++ mainly includes three stages: 1) **Which to Retrieve.** Leveraging SAM2 (with optionally triggered mask fine tuning), masks of garments are obtained to provide visual cues for aiding VLM-based reasoning on the optimal garment to retrieve under task constraints. 2) **Where to Retrieve.** The Retrieval Affordance Model infers optimal grasp points for the target garment, maximizing single-arm retrieval feasibility while ensuring the garment’s safety and cleanliness. 3) **How to Retrieve.** For large/long garments, single-arm retrieval may be unfeasible. The single-arm grasping condition is therefore fed to VLM to determine whether dual-arm cooperation is required, ensuring smooth garment retrieval.

**Abstract**—Garment manipulation has attracted increasing attention due to its critical role in home-assistant robotics. However, the majority of existing garment manipulation works assume an initial state consisting of only one garment, while piled garments are far more common in real-world settings. To bridge this gap, we propose a novel garment retrieval pipeline that can not only follow language instruction to execute safe and clean retrieval but also guarantee exactly one garment is retrieved per attempt, establishing a robust foundation for the execution of downstream tasks (e.g., folding, hanging, wearing). Our pipeline seamlessly integrates vision-language reasoning with visual affordance perception, fully leveraging the high-level reasoning and planning capabilities of VLMs alongside the generalization power of visual affordance for low-level actions. To enhance the VLM’s comprehensive awareness of each garment’s state within a garment pile, we employ visual segmentation model (SAM2) to execute object segmentation on the garment pile for aiding VLM-based reasoning with sufficient visual cues. A mask fine-tuning mechanism is further integrated to address scenarios where the initial segmentation results are suboptimal. In addition, a dual-arm cooperation framework is deployed to address cases involving large or long garments, as well as excessive garment sagging caused by incorrect grasping point determination, both of which are strenuous for a single arm to handle. The effectiveness of our pipeline are consistently demonstrated across diverse tasks and varying scenarios in both real-world and simulation environments. Project page: <https://garmentpile2.github.io/>.

## I. INTRODUCTION

As a key aspect of general home-assistant robot capabilities, garment manipulation is highly challenging because of

its near-infinite space of self-deformable states and complex kinematic and dynamic properties. These properties become even more pronounced in cluttered scenarios, where garments are prone to entanglement with more complicated states, making successful retrieval challenging. Although significant progress has been made in garment manipulation including folding [38], [4], flinging [13], [29], hanging [31], [8], or dressing humans [26], [16], most of the existing work assumes the manipulation of a single garment. However, garment piles are far more common in everyday environments, significantly limiting the real-world deployability of these approaches. Therefore, the ability to retrieve individual garment from a pile is of critical importance.

A recent work, GarmentPile [35], presents a promising approach for cluttered garment manipulation using a per-point affordance model to predict retrieval actionability. Furthermore, For scenarios where the overly complex cluttered state of garments leaves no suitable manipulation region, an additional adaptation model is trained to perform simple adjustments to the initial garment pile, boosting the performance of method. However, due to the specificity of their task scenarios (e.g., placing garments into washing machine), their approach doesn’t consider a dedicated mechanism for single garment retrieval, often resulting in the retrieval of multiple garments simultaneously, which in turn limits the applicability of some downstream tasks. Moreover, due to its reliance on visual affordance alone, the method lacks

integration with language, significantly constraining its flexibility. Furthermore, The garment retrieval process is based solely on a single arm, which struggles to handle large/long garments, therefore compromises the robustness and overall performance of the method in various types of garments.

To overcome these limitations, we propose *GarmentPile++*, a novel pipeline which not only leverages the high-level reasoning capabilities of VLMs, but also exploits the low-level manipulation accuracy and generalization of visual affordance models. As shown in Fig. 1, we structure the process into three stages: **(1) Which to Retrieve:** By integrating visual segmentation model (SAM2 [23]) with VLM-based reasoning, we identify the garment which is suitable for retrieval under task constraints. Furthermore, to address the potential inaccurate segmentation caused by heavy occlusion and color similarity in garment pile, we introduce a mask fine-tuning procedure, which utilizes SAM2 video tracking and point-prompt segmentation to get better segmentation results. **(2) Where to Retrieve:** To maximize the feasibility of single-arm retrieval while ensuring garment safety and cleanliness, the Retrieval Affordance Model is trained to execute the garment grasping, in which the affordance representation is enhanced by incorporating segmentation mask feature into the garment pile’s point cloud, enabling seamless integration into a language-guided manipulation pipeline. **(3) How to Retrieve:** The dual-arm system is deployed for handling different garments and boosting the robustness of pipeline. After master-arm grasp and lift, the current scene situation will be reasoned by VLM for deciding whether to trigger dual-arm cooperation or not. If triggered, a tracking selection procedure will be used for selecting cooperation point for slave-arm grasp. What’s more, the grasp of more than one garments will be detected by VLM to terminate current retrieval and continue next attempt.

We design two typical task scenarios encountered in daily life: open-boundary and closed-boundary. Within these scenarios, we define two retrieval tasks: 1) retrieve garments sequentially, 2) retrieve specific garment. These settings collectively cover nearly all practical cases of garment pile retrieval in everyday environments. Extensive experiments demonstrate the effectiveness of our pipeline.

In summary, our main contributions are as follows:

- we propose a novel pipeline that follows language instruction to execute safe garment retrieval while ensuring exactly one garment is retrieved per attempt, establishing a robust foundation for downstream tasks.
- Our pipeline seamlessly combines vision-language reasoning with visual affordance, fully exploiting the high-level reasoning and planning capabilities of VLMs alongside the generalization power of affordance for low-level grasping.
- Through extensive experiments in both real-world and simulated environments, we demonstrate the effectiveness of our method by successfully tackling two distinct tasks across the two designed scenarios.

## II. RELATED WORK

### A. Deformable Object and Garment Manipulation

Deformable object manipulation poses significant challenges due to complex dynamics and high-dimensional state spaces [6]. Compared with manipulating 1D cables and ropes [25], [36], [40] or square-shaped fabrics [11], [17], [33], [30], garments with diverse categories, shapes, and deformations are significantly more challenging to handle. Current research primarily focuses on manipulating single garments, proposing methods for folding [1], [38], unfolding [3], [13], [18], hanging [7], [32], and dressing [16], [26], [28]. Recent advancements in simulation platforms like *GarmentLab* [19] and *DexGarmentLab* [29] have enabled systematic evaluation of garment manipulation techniques. However, most studies focus on isolated garments in structured environments. In this paper, we address the more challenging problem of cluttered garment manipulation.

### B. Cluttered Environment Manipulation

Robotic manipulation in cluttered environments is a fundamental capability for tasks such as grasping [27], [39], retrieving [37], [14], and rearranging [12], [9], yet it remains highly challenging due to severe occlusions, tight object arrangements, and unstable configurations. Previous works have addressed these difficulties through methods such as Visual-Language-Action Model (VLA) [10], multimodal reasoning [22], and synthetic simulation data [5]. However, These works primarily focus on cluttered rigid object scenarios without considering cluttered deformable object scenarios. *GarmentPile* [35] represents the first attempt at garment manipulation in cluttered environments, but suffers from low compatibility with downstream tasks, single-arm constraints, and lack of language guidance. Our *GarmentPile++* addresses these limitations by integrating VLM-based understanding with dual-arm coordination for robust garment manipulation in cluttered scenarios.

## III. PROBLEM FORMULATION

For a cluttered garment pile, at certain time step  $t$ , given an RGB-D observation  $\mathcal{I}_t \in \mathbb{R}^{H \times W \times 3}$ ,  $\mathcal{D}_t \in \mathbb{R}^{H \times W}$  and an open-ended language query  $\mathcal{T}$ , the goal of *GarmentPile++* is to provide a policy  $\pi(a_t^1, a_t^2 \mid \mathcal{I}_t, \mathcal{D}_t, \mathcal{T})$ , which converts  $\mathcal{T}$  into the target grasping object  $\mathcal{X}_t$ , and generates actions via  $(a_t^1, a_t^2) = G(\mathcal{X}_t, \mathcal{I}_t, \mathcal{D}_t)$ . Herein,  $a_t^1$  denotes the action of master arm and  $a_t^2$  denotes the action of slave arm. Specially,  $a_t^2$  can be set to None in scenarios where cooperative operation is not requisite. The policy  $\pi$  will retrieve garments one by one in a safe and neat manner until  $\mathcal{T}$  is satisfied.

With reference to previous deformable manipulation studies [34], [35], we adopt pick-and-place as action primitive. Since place positions are typically predefined (e.g., placing garments on the table), we use the grasp point  $p_{retrieve} \in \mathbb{R}^3$  with heuristic retrieval orientation as retrieval action  $a_t$ .

Our tasks consist of two stages:

- 1) identify the target garment  $\mathcal{X}_t$  according to  $\mathcal{T}$  and generate master-arm grasping action  $a_t^1$ .

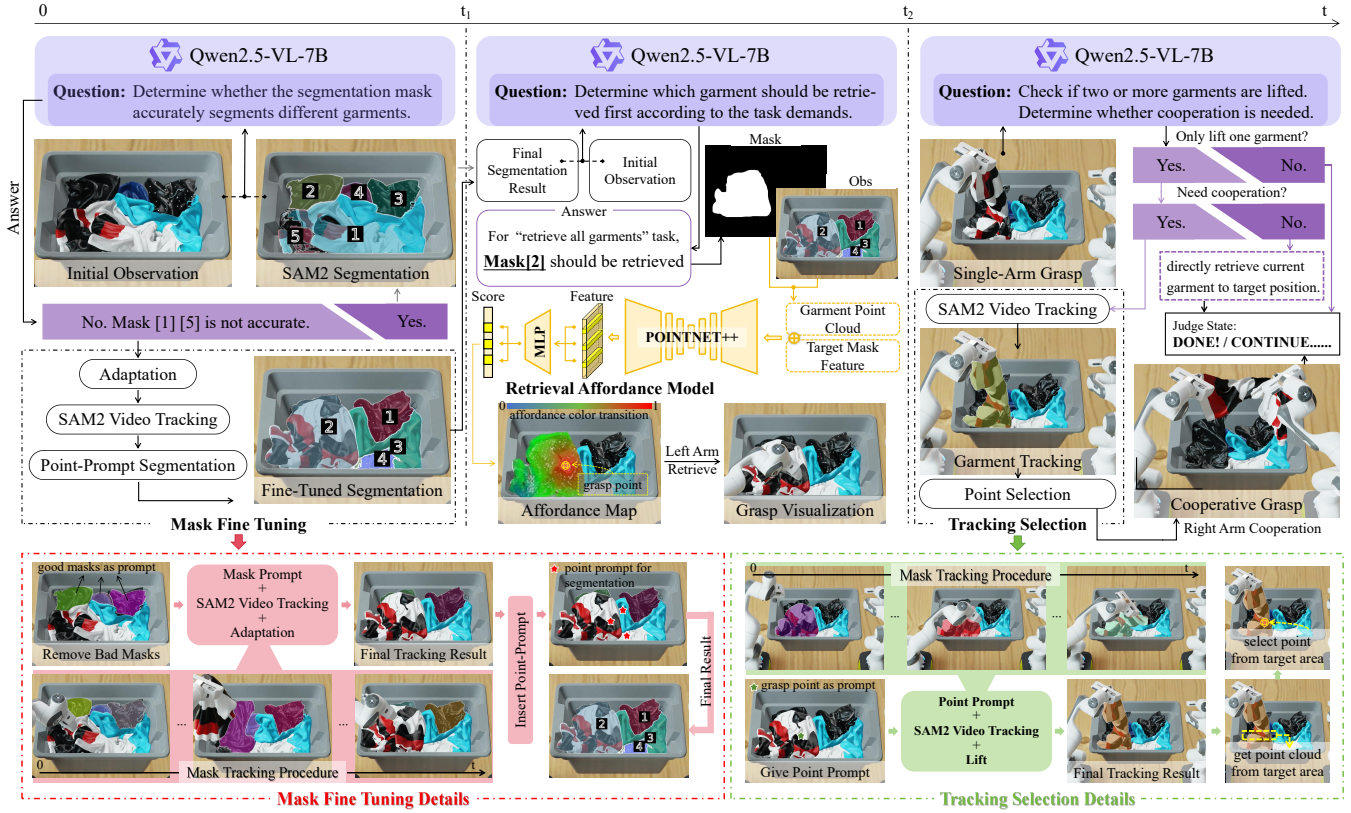


Fig. 2: **GarmentPile++ Pipeline.** Given the initial observation, SAM2 segments the garment pile. Due to pile complexity, masks may be inaccurate, thus optionally triggering mask fine tuning procedure (Bottom-Left, red) based on VLM reasoning. The finalized mask and observation are then passed to the VLM to select the target garment under task constraints. Treating the target mask as garment pile point-cloud feature, the Retrieval Affordance Model predicts single-arm grasp points, which is used for single-arm retrieval. Then the post-grasp observation is re-queried by the VLM to decide on dual-arm cooperation; if required, a tracking-selection module (Bottom-Right, green) produces cooperative grasp points to ensure successful retrieval. What’s more, if VLM detects the single-arm grasp of more than one garment, the current retrieval will be terminated.

- determine whether dual-arm cooperation is required for  $\mathcal{X}_t$  and generate slave-arm cooperation action  $a_t^2$ .

In stage I, we define affordance map  $\mathcal{A}_{retrieve}$ , each digit normalized to  $[0, 1]$ , indicating per-point actionability of  $\mathcal{X}_t$  for retrieval. The point with the highest score is selected as  $p_{retrieve}$  for  $a_t^1$ . In stage II, if cooperation is required, the tracking selection procedure will generate  $p_{retrieve}$  for  $a_t^2$ . Otherwise,  $a_t^2$  will be set to None.

## IV. METHOD

### A. Overview

**GarmentPile++** fuses the vision-language reasoning capacity of VLMs with pre-trained instance segmentation model and affordance model. The overall pipeline is shown in Fig. 2. We first describe the details of garment segmentation leveraging SAM2 and selection via VLM (Section IV-B). Then we introduce the architectural design of Retrieval Affordance Model, alongside its inference workflow and training procedure (Section IV-C). Finally, we delve into the details about master-slave arms cooperation mechanism in GarmentPile++ (Section IV-D).

### B. Vision-Language-Guided Segmentation and Selection

**Single-Image Garment Segmentation.** We employ a SAM2-based model  $\mathcal{S}_{image} : \mathbb{R}^{H \times W \times 3} \rightarrow \{0, 1\}^{H \times W}$  to

perform scene segmentation. It segments all objects in the scene and produces masks. Using the depth map  $\mathcal{D}_t$ , we filter out implausible masks (e.g., overly fragmented or nearly planar) to obtain the final set  $M_{1:N'}$ . For each  $M_i$ , we choose a location on the mask to place a numeric visual marker  $i$  and render annotated images  $\mathcal{I}_t^{border, label}, \mathcal{I}_t^{mask, label}$  (Border refers to the border of the mask, drawn with a white line.). The annotated images, together with the VLM’s OCR capability, enable the VLM to refer to each region by its ID.

**Mask-Adjustment Decision.** To mitigate SAM2’s limitations with garment pile segmentation caused by heavy occlusion and color similarity (e.g., several garments fused into one mask, ragged boundaries), we ask the VLM to determine whether any adjustment is required and, if so, which masks are in need of adjustment. Let the VLM decision function be  $\mathcal{F}^{adjust}$ , such that  $\{I \mid I \subseteq [1:N']\} \ni \mathcal{A}^{adjust} = \mathcal{F}^{adjust}(\mathcal{I}_t^{border, label}, \mathcal{I}_t^{mask, label})$  where  $i \in \mathcal{A}^{adjust}$  means mask  $M_i$  is in need of adjustment.

**Mask Fine Tuning.** If  $\mathcal{A}^{adjust} \neq \emptyset$ , We command the robot to pinch the garment at a random point within  $\bigcup_{i \in \mathcal{A}^{adjust}} M_i$ , lift, shake, and release the garment while recording a video. Using SAM2’s *VideoPredictor*, we track  $M_{\{1:N'\} \setminus \mathcal{A}^{adjust}}$  over the sequence, then regenerate masks on the last frame as follows: we start from the image

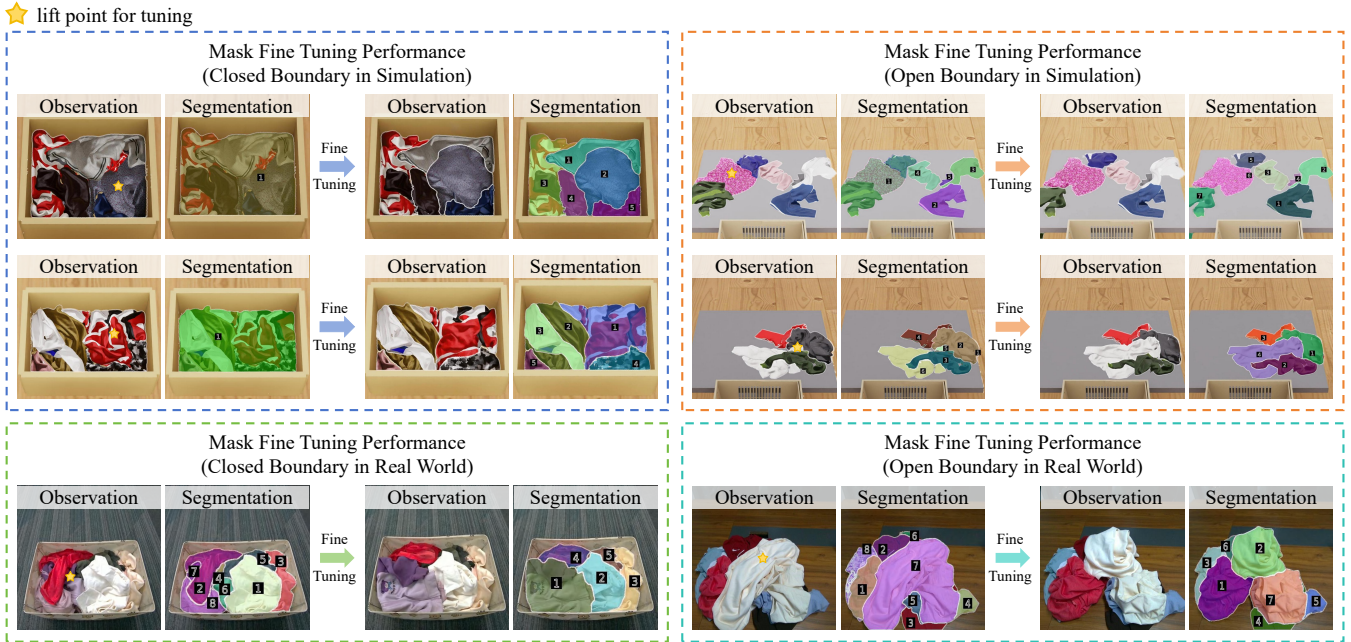


Fig. 3: **Segmentation Performance before and after Mask Fine Tuning Procedure.** Initial segmentation result (left-side images of each part) may be flawed : e.g., a mask covering multiple garments (mask 1 in blue part row 1, 2; mask 1 in orange part row 1), fragmented masks (mask 5 in orange part row 2; mask 4,6,7 in green part), or color confusion (mask 7 in cyan part covering two white garments). ☆ denotes the lift point for tuning. After fine tuning, the adjusted masks (right-side images of each part) are clearly more reasonable, with each mask generally corresponding to a single complete garment.



Fig. 4: **Reasoning Procedure.** For sequential garment retrieval (rows 1 and 3), the VLM infers the most suitable garment to retrieve (e.g., topmost or with simpler entanglement) to maximize overall success. For specific garment retrieval (row 2), it intends to remove obstructing garments so that the target garment is quickly exposed and ready for retrieving.

center and treat any pixel not covered by a mask as a point prompt for SAM2 to generate a new mask. After filtering out implausible masks and applying Non-Maximum Suppression (NMS), we obtain the final set  $M_{1:N}$  and corresponding numeric markers  $i$ . Fig. 3 shows the performance of mask fine tuning procedure.

**Vision-Language-Guided Selection.** Let the VLM selection function be  $\mathcal{F}^{selection}$ , such that  $\{1:N\} \ni n^{selection} = \mathcal{F}^{selection}(\mathcal{I}_t^{border,label}, \mathcal{T})$ . Hence,  $M_{n^{selection}}$  is the mask

of the target garment  $\mathcal{X}_t$ . Fig. 4 illustrates the reasonability of VLM-based reasoning.

### C. Affordance-Guided Garment Retrieval

The  $\mathcal{M}_{aff} : (\mathcal{O}, M_{1:N}, n^{selection}) \rightarrow [0, 1]^{\mathcal{O}} = \mathcal{A}^{retrieve}$  is our Retrieval Affordance Model, which takes the point-cloud observation  $\mathcal{O}$  of the garment clutter and the partial results in Section IV-B as input and produces the affordance map  $\mathcal{A}^{retrieve}$  indicating the per-point actionability of  $\mathcal{X}_t$  for retrieval with a single arm, as shown in Fig. 5.

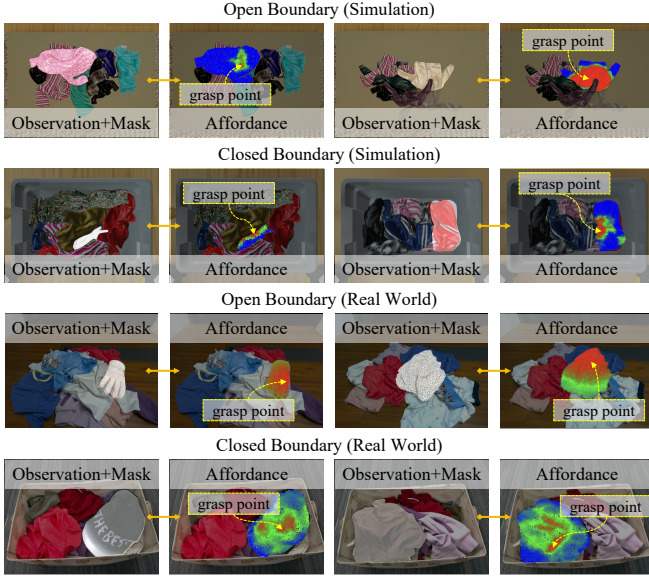


Fig. 5: **Affordance Visualization.** The color transition (blue  $\rightarrow$  green  $\rightarrow$  red) indicates an increasing affordance level. Beyond generalization across different garment types, the affordance further captures: **garment geometry** (row 1, center regions of garments are typically preferable.), **subtle structures** (row 4, garment wrinkle regions easier to grasp exhibit higher affordance level.), **spatial relations** (row 3 left, grasping left / right region of glove is equivalent, but the right side is preferable due to less garments, as well as tops side of garment in row 3 right; row 2, central regions far from closed boundaries are preferable for grasp safety.)

Since  $\mathcal{O}$  is projected from  $\mathcal{I}_t$ , we denote by  $\mathcal{O}^{M_i}$  the point-cloud observation on  $\mathcal{I}_t$  covered by mask  $M_i$ .  $\mathcal{M}_{\text{aff}}$  first employs the PointNet++ [21] backbone feature extractor  $\mathbf{F}_{\text{retrieve}}$ . It takes as input the 3D coordinates in  $\mathcal{O}^{M_{1:N}}$  with the 1-dimensional feature [1 if  $p \in \mathcal{O}^{M_{n\text{selection}}}$  else 0] and extracts the feature  $f_p^{\text{retrieve}} \in \mathbb{R}^{128}$  for each point  $p$  in  $\mathcal{O}^{M_{1:N}}$ . It then feeds these features into multi-layer perceptrons (MLPs) and, after normalization with sigmoid activation, yields the 1-dimensional affordance prediction  $\hat{g}_p^{\text{retrieve}} \in [0, 1]$  for  $p$ . the point  $p$  with highest score on  $\mathcal{O}^{M_{n\text{selection}}}$  is selected as retrieval point  $p_{\text{retrieve}}$ .

As for training procedure, We denote by  $g_p^{\text{retrieve}} \in \{0, 1\}$  the ground-truth retrieval affordance score on  $p$ . We directly perform single-arm grasping at  $p$ , and assign  $g_p^{\text{retrieve}}$  to 1 if the grasp-and-place succeeds, and to 0 if it fails. The loss is computed using Binary Cross Entropy(BCE) Loss [24], i.e.,

$$L_{\text{retrieve}} = - \left( g_p^{\text{retrieve}} \cdot \log(\hat{g}_p^{\text{retrieve}}) + (1 - g_p^{\text{retrieve}}) \cdot \log(1 - \hat{g}_p^{\text{retrieve}}) \right). \quad (1)$$

The points  $p$  participating in training satisfy  $p \in \mathcal{O}^{M_{n\text{selection}}}$ , since ultimately only the highest-scoring point in  $\mathcal{O}^{M_{n\text{selection}}}$  is selected for the retrieval action. By retaining the remaining points which belong to  $\mathcal{O}$  but not to  $\mathcal{O}^{M_{n\text{selection}}}$ , the point-level features extracted by PointNet++ can not only capture the local structure of  $\mathcal{O}^{M_{n\text{selection}}}$  but also aggregate information about the global

structure and inter-garment relations. Such information is crucial for predicting  $\hat{g}_p^{\text{retrieve}}$ .

#### D. State-Based Arm Cooperation

As mentioned in IV-C,  $p_{\text{retrieve}}$ , the highest-scoring point on  $\mathcal{O}^{M_{n\text{selection}}}$ , is selected for retrieval action. According to the position of  $p_{\text{retrieve}}$  in the scene, we choose the arm on the closer side as the master arm and use it to lift the garment at  $p_{\text{retrieve}}$ .

We record the lifting process on video and employ SAM2’s Video Predictor to track  $M^{\text{selection}}$  from start to end until  $M^{\text{picked}}$ . This avoids redundant re-segmentation of the scene and prevents other garments with similar colors from being mistakenly merged into the same mask in the post-lifting 2D image.

Next, we assign the numeric visual marker 1 to mask  $M^{\text{picked}}$ . The VLM’s cooperation function  $\mathcal{F}^{\text{cooperation}}$  then determines whether two or more garments are lifted and whether master-slave arms cooperation is required (e.g., considering whether the garment is too long), i.e.,  $(\{0, 1\} \ni x^{\text{error}}, \{0, 1\} \ni x^{\text{dual}}) = \mathcal{F}^{\text{cooperation}}(\mathcal{I}_t^{\text{border, label}}, \mathcal{I}_{t'}^{\text{border, label}})$ , where  $t$  is the time step defined in Section IV-B, and  $t'$  denotes the post-lifting time step. In this setting, we prompt the VLM to align two garments in the scene to better identify the lifted garment.

$x^{\text{error}} = 1$  means two or more garments are lifted, which results in termination of current retrieval and begin of next attempt. When  $x^{\text{error}} = 0$ , only one garment is lifted, allowing the cooperation process to proceed. If  $x^{\text{dual}} = 0$ , no master-slave arms cooperation is required; the master-arm directly delivers the garment to the target position. Otherwise, we use the function  $\mathcal{P}^{\text{dual}} : (\mathcal{D}_{t'}, M^{\text{picked}}) \rightarrow p^{\text{dual}}$  to select the grasping point of the slave-arm. Specifically, all points in  $M^{\text{picked}}$  are sorted along the  $z$ -axis, and a pick point is sampled from the bottom. Finally, the slave-arm grasps the selected point to lift, and then both arms cooperatively deliver the garment horizontally to the target location.

## V. EXPERIMENT

### A. Simulation Experiments

**Environment Setup.** We use DexGarmentLab [29] built on Isaac Sim 4.5.0 [20] as our backbone simulation environment, which supports both multi-garments interactions and garment-robot interactions. We construct two different types of representative scenario for garment retrieval:

- **Open-Boundary Scenario:** In this scenario, garments are randomly scattered across relatively open planar surfaces (e.g., tables, floors, beds, etc.). This implies that the stacking state of garment piles is relatively insignificant, as spatial constraints are absent.
- **Closed-Boundary Scenario:** In this scenario, garments are randomly placed into containers with enclosed boundaries (e.g., baskets, boxes, etc.). This implies that, due to spatial constraints, the garment piles will exhibit a relatively severe stacking state.

We load 9 categories (dress, trousers, tops, skirt, socks, glove, hat, scarf and underpants) of 153 different garments

TABLE I: Comparison of baselines on garment retrieval tasks.

Method	Retrieve Garments Sequentially				Retrieve a Piece of Specific Garment							
	Open Boundary		Closed Boundary		Hat (Open Boundary)		Glove (Closed Boundary)		Green Tops (Open Boundary)		Black Trousers (Closed Boundary)	
	ASR <sub>A</sub> ↑	AMS	ASR <sub>A</sub> ↑	AMS	ASR <sub>B</sub> ↑	AMS	ASR <sub>B</sub> ↑	AMS	ASR <sub>B</sub> ↑	AMS	ASR <sub>B</sub> ↑	AMS
ThinkGrasp	0.641	↘	0.413	↘	0.517	2.017	0.593	1.439	0.542	1.067	0.585	1.171
GarmentPile	0.819	↘	0.792	↘	0.623	2.798	0.659	3.152	0.729	1.957	0.545	1.394
Qwen (only)	0.829	↘	0.720	↘	0.708	1.858	0.739	1.797	0.787	1.160	0.686	1.081
Our Method	<b>0.904</b>	↘	<b>0.874</b>	↘	<b>0.892</b>	2.196	<b>0.925</b>	2.142	<b>0.877</b>	1.106	<b>0.862</b>	1.147

from ClothesNet [41]. In the Open-Boundary Scenario, 6–16 garments will be loaded to form garment piles; in the Closed-Boundary Scenario, the quantity of garments employed will be 3–8. What’s more, there are two types of typical tasks conducted in each scenario:

- *Task A: Retrieve Garments Sequentially.* The robot is required to continue retrieving until no garments remain in the observation field.
- *Task B: Retrieve a Piece of Specific Garment.* The robot needs to retrieve until the target garment (e.g., hat, glove, green tops, black trousers, etc.) is retrieved.

**Baselines.** Owing to the lack of relevant studies in cluttered garments retrieval with VLM reasoning, we select representative works in the fields of VLM-guided grasping and cluttered garment retrieval, and adapt them to align with the requirements of our tasks. The baselines and their introductions are shown as below: (1) **ThinkGrasp** [22]. According to the current RGB observation, the modified ThinkGrasp first give a brief introduction about the target garment, which is then used by LangSAM to segment and output the mask. The VLM select a grasp point inside the mask and guide the single arm to grasp and retrieve the garment. (2) **GarmentPile** [35]. The robot continues retrieving garment only according to the point cloud of garment piles without VLM reasoning, until the task is finished. (3) **Qwen (only)** [2]. Qwen2.5-VL-7B is used as the same with our method in this baseline. It first infers an appropriate grasping point based on the given RGB observation, and then guides the single arm to grasp and lift the garment. Subsequently, the current RGB observation is captured and sent to Qwen2.5-VL-7B to determine whether to use the dual arms and which point to select for cooperation grasp.

We evaluate performance using Average Success Rate (ASR) for both tasks, with Average Motion Steps (AMS) as an additional efficiency metric for Task B.

*Average Success Rate (ASR):* For Task A, ASR measures the probability of successfully grasping any single garment:

$$ASR_A = \frac{\text{Number of successfully grasped garments}}{\text{Total number of loaded garments}} \quad (2)$$

For Task B, ASR represents the probability of successfully retrieving the target garment while maintaining proper handling of all previously grasped items:

$$ASR_B = \frac{\text{Number of completed tasks}}{\text{Total number of tasks}} \quad (3)$$

*Average Motion Steps (AMS):* This metric quantifies the efficiency of our approach by measuring the average number

of motion steps required to complete each task:

$$AMS = \frac{\text{Total motion steps to complete all tasks}}{\text{Total number of tasks}} \quad (4)$$

**Ablations.** To demonstrate the necessity of the proposed adaptation module, we compare with the following ablated versions: (1) **w/o Mask Fine Tuning** that removes the mask fine-tuning part and directly uses masks generated by SAM2 scene segmentation. (2) **w/o Affordance** that removes the affordance part and randomly selects grasp point from target garment mask. (3) **w/o Tracking Selection** that removes the tracking selection part and directly queries Qwen2.5-VL-7B using single-arm lift image to select cooperative point for dual arm. (4) **w/o Dual Arm** that removes dual-arm cooperation part and only uses single for garment retrieval. (5) **w/o Affordance & Dual Arm** that removes both affordance and dual-arm cooperation part and randomly selects grasp point for single arm to retrieve garment.

For all the ablations, the ASR metric is used for evaluation; while for w/o Affordance baseline and w/o Mask Fine Tuning baseline, the Probability of Dual Arm (PDR) metric is additionally used for evaluating the probability of triggering dual-arm collaborative operations, which represents the efficiency of pipeline.

**Results and Analysis.** Tab. I and Tab. II present the baseline comparison results and ablation study results for our method, correspondingly. Fig. 3, Fig. 4 and Fig. 5 respectively show the visualization performance of the mask fine-tuning part, reasoning procedure part and affordance part in our proposed method.

Tab. I demonstrates that our method outperforms all baselines. For **ThinkGrasp**, due to the complex states and diverse colors of garments, it is challenging to guide LangSAM to obtain appropriate and accurate masks via linguistic descriptions. Additionally, ThinkGrasp only adopts single-arm grasping, resulting in relatively low success rates across all tasks. For **GarmentPile**, it leverages affordance to identify the most accessible single-arm grasping position on the garment pile. In terms of success rate, it generally outperforms ThinkGrasp. However, the lack of linguistic guidance leads to no clear direction when determining garment grasping positions, resulting in a significantly higher AMS than other methods and thus lower efficiency. As for **Qwen (only)**, it incorporates both VLM reasoning and dual-arm cooperation modules, contributing to a better success rate than ThinkGrasp and GarmentPile. Nevertheless, it is highly affected by the garment stacking state: in the Open Boundary scenario, where stacking is mild and garment structures are

TABLE II: Ablation of our methods on two typical tasks

Method	Retrieve a Piece of Green Tops (Open Boundary)		Retrieve Garments Sequentially (Closed Boundary)	
	ASR <sub>B</sub> ↑		ASR <sub>A</sub> ↑	
w/o Tracking Selection	0.669		0.739	
w/o Dual Arm	0.613		0.634	
w/o Affordance & Dual Arm	0.543		0.564	
	ASR <sub>B</sub> ↑	PDR ↓	ASR <sub>A</sub> ↑	PDR ↓
w/o Mask Fine Tuning	0.841	0.794	0.857	0.624
w/o Affordance	0.823	0.843	0.810	0.595
Our Method	<b>0.877</b>	<b>0.766</b>	<b>0.874</b>	<b>0.551</b>

relatively distinct, Qwen achieves a high success rate; in the Closed Boundary scenario, severe garment stacking limits its reasoning capability, leading to a marked drop in success rate. **Our method** integrates multiple components, including mask fine-tuning, affordance, VLM reasoning, and dual-arm cooperation. It achieves state-of-the-art (SOTA) performance for cluttered garment manipulation across all task scenarios and its AMS shows slight differences from other baselines, ensuring both high success rate and efficiency.

Table II presents a quantitative analysis of the ablation study, demonstrating the effectiveness of key components in our method, including *mask fine tuning*, *tracking selection*, *dual-arm cooperation*, and *affordance*. Notably: 1) The comparison between the "w/o Dual Arm" and "w/o Affordance & Dual Arm" reveals that integrating the affordance component enhances the success rate of single-arm grasping. This demonstrates that the affordance guides the robotic arm to prioritize manipulating the more suitable region of garments, preventing excessive sagging. 2) Based on this observation, the further comparison between "w/o Affordance" and "Our Method" (in terms of ASR and PDR) shows that the affordance component not only improves the overall success rate but also reduces the probability of triggering dual-arm cooperation, which boosts efficiency. 3) Besides, the comparison between "w/o Mask Fine Tuning" and "Our Method" (in terms of ASR and PDR) establishes that our mask fine tuning procedure enhances the accuracy of segmented masks, which enables the affordance component to locate better points for retrieval according to the precise garment mask, ultimately leading to higher ASR score and lower PDR score. 4) Additionally, other results also indicate that dual-arm cooperation and tracking selection significantly elevate the overall task success rate, which further underscores the necessity of adopting dual-arm cooperation for cluttered garments manipulation, and exposes the deficiency in the understanding and reasoning capabilities of current VLMs for garments without auxiliary information cues.

Fig. 3, Fig. 4 and Fig. 5 further illustrate the effectiveness of *mask fine tuning*, *VLM reasoning* and *affordance*. In Fig. 3, the adjusted mask enables more accurate segmentation of individual garments in the scene, thereby supporting the subsequent reasoning and judgment of the VLM. In previous works [22], [15], since the task scenarios only involved rigid objects where segmentation is relatively less challenging, no additional mask adjustment was required. However, in scenarios with cluttered garments, accurate segmentation

TABLE III: Real-World Performance Comparison

Real-World Task	Metrics	ThinkGrasp	Qwen(only)	Our Method
Retrieve Garments Sequentially (Open Boundary)	ASR <sub>A</sub> ↑	9/20 (45%)	16/20 (80%)	<b>19/20 (95%)</b>
Retrieve Garments Sequentially (Closed Boundary)	ASR <sub>A</sub> ↑	7/20 (35%)	11/20 (55%)	<b>17/20 (85%)</b>
Retrieve a Piece of Hat (Open Boundary)	ASR <sub>B</sub> ↑	5/10 (50%)	7/10 (70%)	<b>9/10 (90%)</b>
Retrieve a Piece of Purple Tops (Closed Boundary)	ASR <sub>B</sub> ↑	4/10 (40%)	5/10 (50%)	<b>8/10 (80%)</b>

of garments cannot be guaranteed, which necessitates the adoption of the mask fine-tuning method. In Fig. 4, it is evident that the VLM can reasonably infer the mask of the garment to be grasped based on task requirements, which further validates the rationality of the overall system. In Fig. 5, the visualized affordance clearly demonstrates its generalization ability and awareness of garment structure and relations. This not only improves the manipulation success rate but also reduces the probability of triggering dual-arm cooperation, thereby enhancing efficiency.

### B. Real-World Experiments




Fig. 6: Real-World Experiments Setting.

We use ARX x7s as real-world robot and RealSense D405 as head camera for images capture. The real-world experiments setting are shown in Fig. 6. We have already provided some real-world performances on *mask fine tuning*, *VLM reasoning* and *affordance* in Fig. 3, Fig. 4 and Fig. 5 to illustrate the effectiveness of our method. Tab. III reports ASR on baselines and our method, which further demonstrates the excellent performance of GarmentPile++. Supplementary video and project website show more demonstrations.

## VI. CONCLUSIONS

we propose a novel cluttered garment retrieval pipeline that fully utilizes VLM's powerful capability of reasoning and strong generalization ability of affordance to retrieve garment one by one based on language instructions, which establish robust foundation for the execution of downstream tasks (e.g., folding, hanging, wearing). As our method relies on visual segmentation model, it may suffer from performance degradation under low-light conditions or for garments with complex patterns/colors due to segmentation

errors. Nevertheless, it still contributes significantly to building garment manipulation pipeline in home environments.

#### ACKNOWLEDGMENT

This project was supported by National Natural Science Foundation of China (62376006) and National Youth Talent Support Program (8200800081).

#### REFERENCES

- [1] Yahav Avigal, Lars Berscheid, Tamim Asfour, Torsten Kröger, and Ken Goldberg. Speedfolding: Learning efficient bimanual folding of garments. In *IROS*, 2022.
- [2] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, and etc. Qwen2.5-vl technical report, 2025.
- [3] Alper Canberk, Cheng Chi, Huy Ha, Benjamin Burchfiel, Eric Cousineau, Siyuan Feng, and Shuran Song. Cloth funnels: Canonicalized-alignment for multi-purpose garment manipulation. *ICRA*, 2023.
- [4] Haonan Chen et al. Metafold: Language-guided multi-category garment folding framework via trajectory generation and foundation model. *arXiv preprint arXiv:2503.08372*, 2025.
- [5] Tianxing Chen, Zanxin Chen, Baijun Chen, Zijian Cai, Yibin Liu, Qiwei Liang, Zixuan Li, Xianliang Lin, Yiheng Ge, Zhenyu Gu, et al. Robotwin 2.0: A scalable data generator and benchmark with strong domain randomization for robust bimanual robotic manipulation. *arXiv preprint arXiv:2506.18088*, 2025.
- [6] Tianxing Chen, Kaixuan Wang, Zhaohui Yang, Yuhao Zhang, Zanxin Chen, Baijun Chen, Wanxi Dong, Ziyuan Liu, Dong Chen, Tianshuo Yang, et al. Benchmarking generalizable bimanual manipulation: Robotwin dual-arm collaboration challenge at cvpr 2025 meis workshop. *arXiv preprint arXiv:2506.23351*, 2025.
- [7] Wei Chen, Dongmyoung Lee, Digby Chappell, and Nicolas Rojas. Learning to grasp clothing structural regions for garment manipulation tasks. In *IROS*, 2023.
- [8] Yuxing Chen, Songlin Wei, Bowen Xiao, Jiangran Lyu, Jiayi Chen, Feng Zhu, and He Wang. Robohanger: Learning generalizable robotic hanger insertion for diverse garments. *IEEE Robotics and Automation Letters*, 10(9):8922–8929, 2025.
- [9] Sang Hun Cheong, Brian Y Cho, Jinhwi Lee, ChangHwan Kim, and Changjoo Nam. Where to relocate?: Object rearrangement inside cluttered and confined environments for robotic manipulation. In *ICRA*, 2020.
- [10] Shengliang Deng, Mi Yan, Songlin Wei, Haixin Ma, Yuxin Yang, Jiayi Chen, Zhiqi Zhang, Taoyu Yang, Xuheng Zhang, Heming Cui, et al. Graspvla: a grasping foundation model pre-trained on billion-scale synthetic action data. *arXiv preprint arXiv:2505.03233*, 2025.
- [11] Aditya Ganapathi et al. Learning dense visual correspondences in simulation to smooth and fold real fabrics. In *ICRA*, 2021.
- [12] Ankit Goyal, Arsalan Mousavian, Chris Paxton, Yu-Wei Chao, Brian Okorn, Jia Deng, and Dieter Fox. Ifor: Iterative flow minimization for robotic object rearrangement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14787–14797, 2022.
- [13] Huy Ha and Shuran Song. Flingbot: The unreasonable effectiveness of dynamic manipulation for cloth unfolding. In *Conference on Robot Learning*. PMLR, 2022.
- [14] Baichuan Huang, Shuai D Han, Jingjin Yu, and Abdeslam Boularias. Visual foresight trees for object retrieval from clutter with nonprehensile rearrangement. *IEEE Robotics and Automation Letters*, 7(1):231–238, 2021.
- [15] Runyu Jiao, Alice Fasoli, Francesco Giuliani, Matteo Bortolon, Sergio Povoli, Guofeng Mei, Yiming Wang, and Fabio Poiesi. Free-form language-based robotic reasoning and grasping. *arXiv preprint arXiv:2503.13082*, 2025.
- [16] Stelios Kotsovolis and Yiannis Demiris. Model predictive control with graph dynamics for garment opening insertion during robot-assisted dressing. In *ICRA*, 2024.
- [17] Robert Lee et al. Learning arbitrary-goal fabric folding with one hour of real robot experience. *arXiv:2010.03209*, 2020.
- [18] Yinxiao Li, Danfei Xu, Yonghao Yue, Yan Wang, Shih-Fu Chang, Eitan Grinspun, and Peter K Allen. Regrasping and unfolding of garments using predictive thin shell modeling. In *ICRA*, 2015.
- [19] Haoran Lu, Ruihai Wu, Yitong Li, Sijie Li, Ziyu Zhu, Chuanruo Ning, Yan Shen, Longzan Luo, Yuanpei Chen, and Hao Dong. Garmentlab: A unified simulation and benchmark for garment manipulation. In *Advances in Neural Information Processing Systems*, 2024.
- [20] NVIDIA. Isaac Sim.
- [21] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *NeurIPS*, 2017.
- [22] Yaoyao Qian, Xupeng Zhu, Ondrej Biza, Shuo Jiang, Linfeng Zhao, Haojie Huang, Yu Qi, and Robert Platt. Thinkgrasp: A vision-language system for strategic part grasping in clutter, 2024.
- [23] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024.
- [24] Usha Ruby, Vamsidhar Yendapalli, et al. Binary cross entropy with deep learning technique for image classification. *Int. J. Adv. Trends Comput. Sci. Eng.*, 9(10), 2020.
- [25] Daniel Seita, Pete Florence, Jonathan Tompson, Erwin Coumans, Vikas Sindhwani, Ken Goldberg, and Andy Zeng. Learning to rearrange deformable cables, fabrics, and bags with goal-conditioned transporter networks. In *ICRA*, 2021.
- [26] Zhanyi Sun, Yufei Wang, David Held, and Zackory Erickson. Force-constrained visual policy: Safe robot-assisted dressing via multi-modal sensing. *RA-L*, 2024.
- [27] Chenxi Wang, Hao-Shu Fang, Minghao Gou, Hongjie Fang, Jin Gao, and Cewu Lu. Graspness discovery in clutters for fast and accurate grasp detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15964–15973, 2021.
- [28] Yufei Wang et al. One policy to dress them all: Learning to dress people with diverse poses and garments. In *RSS*, 2023.
- [29] Yuran Wang, Ruihai Wu, Yue Chen, Jiarui Wang, Jiaqi Liang, Ziyu Zhu, Haoran Geng, Jitendra Malik, Pieter Abbeel, and Hao Dong. Dexgarmentlab: Dexterous garment manipulation environment with generalizable policy, 2025.
- [30] Thomas Weng, Sujay Man Bajracharya, Yufei Wang, Khush Agrawal, and David Held. Fabricflownet: Bimanual cloth manipulation with a flow-based policy. In *CoRL*, 2022.
- [31] Ruihai Wu, Haoran Lu, Yiyan Wang, Yubo Wang, and Hao Dong. Unigarmentmanip: A unified framework for category-level garment manipulation via dense visual correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16340–16350, 2024.
- [32] Ruihai Wu, Haoran Lu, Yiyan Wang, Yubo Wang, and Hao Dong. Unigarmentmanip: A unified framework for category-level garment manipulation via dense visual correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2024.
- [33] Ruihai Wu, Chuanruo Ning, and Hao Dong. Learning foresightful dense visual affordance for deformable object manipulation. In *IEEE International Conference on Computer Vision (ICCV)*, 2023.
- [34] Ruihai Wu, Chuanruo Ning, and Hao Dong. Learning foresightful dense visual affordance for deformable object manipulation, 2023.
- [35] Ruihai Wu, Ziyu Zhu, Yuran Wang, Yue Chen, Jiarui Wang, and Hao Dong. Garmentpile: Point-level visual affordance guided retrieval and adaptation for cluttered garments manipulation, 2025.
- [36] Yilin Wu, Wilson Yan, Thanard Kurutach, Llerel Pinto, and Pieter Abbeel. Learning to manipulate deformable objects without demonstrations. *RSS*, 2020.
- [37] Kechun Xu, Shuqi Zhao, Zhongxiang Zhou, Zizhang Li, Huaijin Pi, Yifeng Zhu, Yue Wang, and Rong Xiong. A joint modeling of vision-language-action for target-oriented grasping in clutter. *arXiv preprint arXiv:2302.12610*, 2023.
- [38] Han Xue, Yutong Li, Wenqiang Xu, Huanyu Li, Dongzhe Zheng, and Cewu Lu. Unifolding: Towards sample-efficient, scalable, and generalizable robotic garment folding. In *CoRL*, 2023.
- [39] Andy Zeng et al. Robotic pick-and-place of novel objects in clutter with multi-affordance grasping and cross-domain image matching. *IJRR*, 2022.
- [40] Kaifeng Zhang, Baoyu Li, Kris Hauser, and Yunzhu Li. Adaptigraph: Material-adaptive graph-based neural dynamics for robotic manipulation. In *RSS*, 2024.
- [41] Bingyang Zhou, Haoyu Zhou, Tianhai Liang, Qiaojun Yu, and etc. Clothesnet: An information-rich 3d garment model repository with simulated clothes environment. *ICCV*, 2023.