

# OCRA: Object-Centric Learning with 3D and Tactile Priors for Human-to-Robot Action Transfer

Kuanning Wang<sup>1\*</sup>, Ke Fan<sup>1\*</sup>, Yuqian Fu<sup>1†</sup>, Siyu Lin<sup>1</sup>, Hu Luo<sup>1</sup>  
 Daniel Seita<sup>2</sup>, Yanwei Fu<sup>1</sup>, Yu-Gang Jiang<sup>1</sup>, Xiangyang Xue<sup>1†</sup>

**Abstract**—We present OCRA, an Object-Centric framework for video-based human-to-Robot Action transfer that learns directly from human demonstration videos to enable robust manipulation. Object-centric learning emphasizes task-relevant objects and their interactions while filtering out irrelevant background, providing a natural and scalable way to teach robots. OCRA leverages multi-view RGB videos, the state-of-the-art 3D foundation model VGGT, and advanced detection and segmentation models to reconstruct object-centric 3D point clouds, capturing rich interactions between objects. To handle properties not easily perceived by vision alone, we incorporate tactile priors via a large-scale dataset of over one million tactile images. These 3D and tactile priors are fused through a multimodal module (ResFiLM) and fed into a Diffusion Policy to generate robust manipulation actions. Extensive experiments on both vision-only and visuo-tactile tasks show that OCRA significantly outperforms existing baselines and ablations, demonstrating its effectiveness for learning from human demonstration videos.

## I. INTRODUCTION

“Man is the most imitative of living creatures, and through imitation learns his earliest lessons”, observed Aristotle. Humans naturally acquire skills by observing others, motivating the growing paradigm of human-to-robot imitation learning, particularly from videos, which offer a scalable and information-rich source of supervision. Inspired by this, existing work can be broadly categorized into interface-based [1], [2], egocentric vision based [3], [4], and object-centric based [5]. Among these, object-centric methods have shown generalization by focusing on the primary objects while reducing sensitivity to less relevant or distracting backgrounds. However, approaches such as [5] rely on object-specific priors, such as 6-DOF poses, which limit the range of manipulable objects, and the absence of richer 3D geometry weakens the ability to capture object interactions. Other object-centric methods (e.g., [6]) use 3D point clouds, but these are collected through costly, time-consuming teleoperation rather than scalable human videos.

In this paper, we propose OCRA, an Object-Centric framework for video-based human-to-Robot Action transfer that combines multi-view RGB inputs, 3D scene reconstruction, and tactile priors to enable robust perception and manipulation. Our key contribution is the extraction of rich object-centric representations by integrating 3D information with tactile priors, which are then used to generate effective

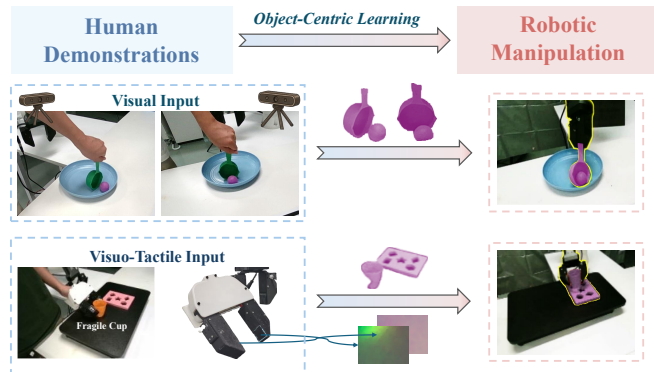


Fig. 1: **Overview of OCRA.** OCRA leverages object-centric learning from multi-view human demonstration videos and tactile sensing to enable robust execution of diverse manipulation tasks transferred from humans to robots.

robot actions directly from human demonstration videos, achieving strong generalization in challenging manipulation tasks. Specifically, we begin by capturing third-person human manipulation videos. Since single-view RGB often misses critical spatial geometry and interaction details, we record multi-view RGB sequences to provide richer scene coverage. We then leverage VGGT [7], a state-of-the-art 3D foundation model, to reconstruct object-centric 3D scenes and estimate depth, producing more complete and consistent representations than raw depth maps, which are prone to noise, missing regions, and limited coverage. To further isolate task-relevant objects, we apply advanced detection and segmentation models [8], [9] to track objects of interest, and combine these object-centric masks with metric depth to extract SE(3) transformations of manipulated objects. By integrating these object-centric visual representations, OCRA effectively removes background distractions and enables human-to-robot action transfer.

While RGB video provides rich visual information, many manipulation tasks also require tactile perception to capture object properties that vision alone cannot reliably discern. To address this, OCRA integrates tactile priors collected using a portable tactile gripper, building a dataset of over one million tactile images from diverse object manipulations to pretrain a tactile encoder. This multimodal integration enables the robot to accurately perceive object properties such as texture and weight, allowing robust execution of tasks like sorting objects by their properties and safely adjusting grasp force when handling delicate items.

To fully leverage both visual and tactile information, we present a novel multimodal fusion module, ResFiLM, which integrates 3D and tactile priors to produce rich object-centric

<sup>1</sup> Fudan University, China. <sup>2</sup> Thomas Lord Department of Computer Science, University of Southern California, USA.

\* indicates equal contributions.

† indicates the corresponding author.

representations. These representations are then integrated by our diffusion policy [10] model to generate robust and precise manipulation actions, enabling effective human-to-robot action transfer. Extensive experiments show that OCRA consistently outperforms existing baselines and ablations on both vision-only and visuo-tactile tasks, demonstrating strong performance on challenging manipulation scenarios, thereby validating the effectiveness of our multimodal, object-centric approach.

We summarize our main contributions as,

- 1) *Object-Centric Learning from Multi-View Videos*: We propose OCRA, an object-centric framework that focuses on task-relevant objects in multi-view human demonstration videos, using 3D reconstruction and object-centric masks to filter background distractions and capture object interactions for robust human-to-robot action transfer.
- 2) *Integration of 3D and Tactile Priors*: We combine 3D object-centric representation produced by 3D foundation model VGGT with tactile priors pretrained from a large-scale tactile dataset via a fusion module (ResFiLM) and a diffusion policy, enabling the robot to perceive object properties and generate precise manipulation actions.
- 3) *Extensive Empirical Validation*: We validate OCRA across vision-only and visuo-tactile tasks. Our method outperforms baselines and ablations, and the experimental results and analyses highlight its effectiveness and robustness.

## II. RELATED WORK

### A. Imitation Learning from Human Demonstrations Videos

Traditional imitation learning methods, such as ACT [11], Diffusion Policy [10], and vision-language-action models [12], [13], primarily focus on behavior cloning [14] from robotic demonstrations. Due to the high cost of collecting such data, recent research increasingly turns to learning from human demonstration videos. Pretraining approaches like R3M [3] and VIP [4] address this by leveraging large-scale egocentric human demonstration datasets (e.g., Ego4D [15]) to train visual encoders, which are then fixed for downstream robotic policy learning. To further narrow the gap between human and robot domains, UMI [1] and Fast-UMI [2] introduce a universal interface that standardizes the eye-on-hand viewpoint across both human demonstrations and robot executions, enabling more direct policy transfer. However, this restricts demonstrations to the hand-mounted first-person perspective. Alternatively, pose-based methods [5] utilize diffusion models and FoundationPose [16] to learn manipulation trajectories from human demonstrations, but they require object-specific priors for accurate pose estimation and additional depth images, while the limited 3D input for policies further hinders their ability to capture complex object interactions.

### B. Object-Centric Learning in Robotics

Object-centric learning [17] seeks to represent individual objects rather than entire images by decomposing scenes into

semantically coherent regions. Such object-level representations enable robots to perceive and interact with their environments in a manner analogous to human object-centered reasoning. A diverse body of work on object discovery has proposed methods for extracting discrete, meaningful object representations [18], [19], [20], [21], [22], [23], which have shown effectiveness in reinforcement learning and robotic manipulation [24], [25]. More recently, approaches such as FOCUS [26] and SOLD [27] have embedded these object-centric representations within world models to boost robotic learning. Other lines of research employ explicit object descriptions, such as bounding boxes [28], [29], 6-DoF poses [30], [31], or additional annotations, to support downstream robotic tasks. Notably, GROOT [6] and VIOLA [32] construct object representations from segmented or detected instances, thereby enhancing robots’ ability to generalize to novel situations. Progress in this direction has been further accelerated by recent open-world visual foundation models [8], [33], [9], [34], [35]. Similarly, our work builds on these advances by leveraging visual foundation models to obtain object masks and integrating them with multi-view 3D reconstruction and tactile sensing, yielding richer object-centric representations tailored for human-to-robot action transfer.

### C. Tactile Sensing Hardware and Learning in Robotics

Resistive, capacitive, and piezoelectric tactile sensors detect mechanical stimuli via pressure-induced changes in electrical properties. Instead, optical tactile system like GelSight [36] and GelSlim [37] uses a camera to detect soft material deformations, producing optical flow for tactile sensing. We use a custom optical tactile sensor inspired by GelSight that generates dense 3D signals via optical flow for normal and tangential forces. Built on open-source code, it is low-cost, user-friendly, and mounted on each finger of a parallel gripper for rich contact sensing. 3D-ViTac [38] combines visual and tactile data to fuse 3D point clouds and uses a diffusion model to learn complex manipulation tasks. It outperforms vision-only methods, even when using low-cost robotic arms. However, it needs teleoperation for the precise states of the robot. FreeTacMan [39] is a wearable, robot-free visuo-tactile demonstration system that records tactile signals along with end-effector poses tracked by a high-precision motion capture system, which limits its applicability in scenarios without such costly tracking infrastructure.

## III. METHOD

**Overview.** First, we capture multi-view videos of object manipulation, performed either with the human hand alone or with a gripper controlled by the human. Next, we extract the object-centric 3D prior through 3D reconstruction and object segmentation. The tactile prior is obtained by feature extraction using a large-scale pretrained encoder. Finally, we fuse the visual and tactile signals and use them as conditional inputs to a diffusion model, enabling generative modeling of the object trajectory.

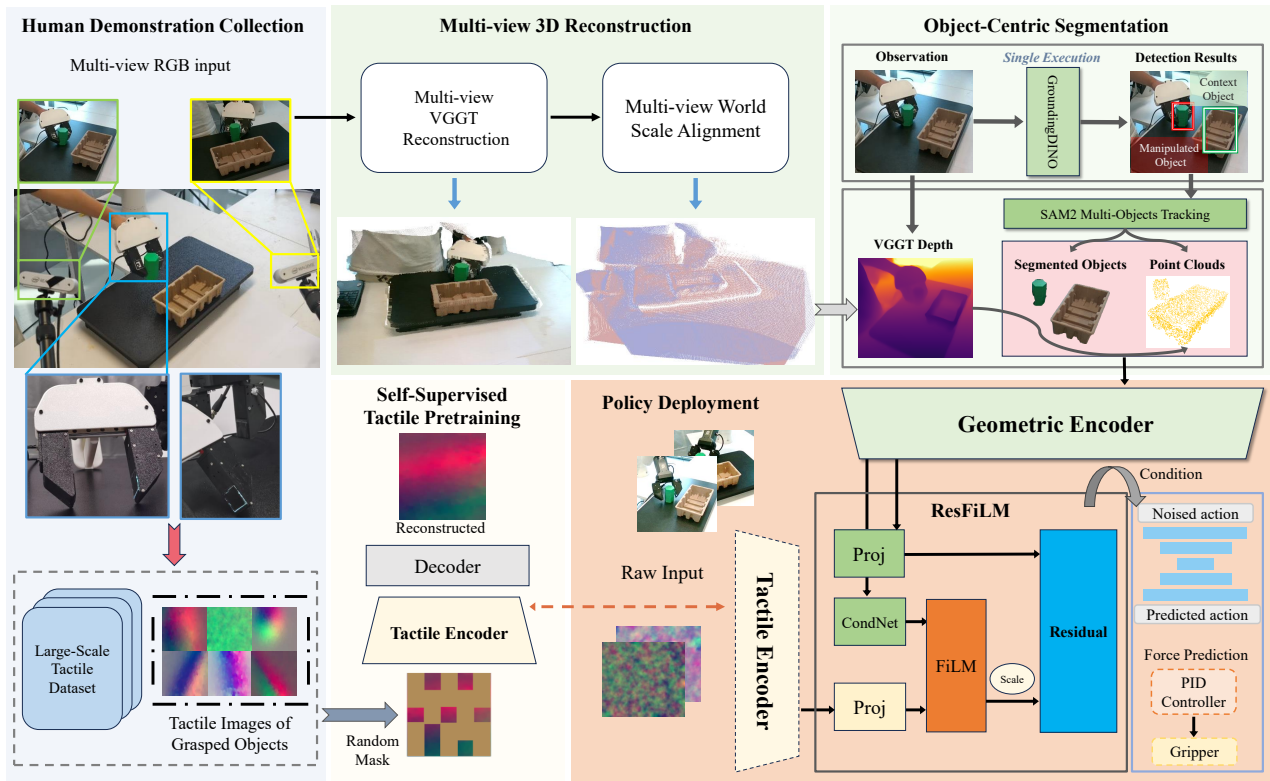


Fig. 2: **Framework.** *The left column* illustrates our human demonstration collection system. Two RGB cameras capture demonstration videos, while the blue box highlights a portable tactile gripper for tactile data collection, which also records fingertip tactile images used to build our large-scale tactile dataset (shown at the bottom). *The first row* depicts how OCRA processes multi-view RGB inputs to obtain object-centric 3D priors. We first reconstruct the 3D scene using VGGT, followed by bi-view metric depth prediction for world-scale alignment. GroundingDINO and SAM2 then provide object segmentation masks, divided into a Manipulable Object Mask (for target objects) and a Context Object Mask (for surrounding objects). These are used to extract visual object-centric representations across modalities (segmentation, point cloud). *The middle of the second row* shows tactile-prior extraction via Tactile Encoder pretraining under a Masked Autoencoder paradigm. *The right of the second row* presents policy deployment. Multi-view RGB and tactile images are encoded into geometric and tactile features, which are fused by ResFiLM and passed to a Diffusion Policy. The policy predicts actions through iterative denoising of noisy action samples.

### A. Human Demonstration Collection System

**Hardware Design.** We designed an easy-to-use system to collect human demonstrations, as shown in the left column of Fig. 2. For the visual information, we used 2 external cameras to capture RGB information from a third-person perspective during human demonstrations of object manipulation. The task-relevant objects were simultaneously captured by the two cameras, whose relative positions were calibrated.

**Vision-Based Tactile Sensor.** To capture tactile information, we employ an elastomeric sensing layer embedded with tracer particles, providing a compact, low-cost design with high-resolution sensing capability, following [40]. Under external loading, the particles exhibit motion patterns governed by the Poisson effect, which vary with the applied force. By analyzing the resulting displacement field within the material, we extract tactile signals. To ensure stability and reduce errors from ambient light, we add a black light-shielding mask and controlled LED lighting.

### B. Object-Centric 3D Prior

**Multi-view 3D Reconstruction.** We used multiple RGB perspectives; however, simply using RGB perspectives overlooks the relationship between images. To make more effective use of this 3D prior, we used a 3D reconstruction model

to extract the scene point cloud. We adopt the state-of-the-art multi-view reconstruction algorithm VGGT [7], which estimates scene geometry from multi-view RGB images and performs real-time processing of RGB frames to produce depth maps, point maps, and pixel-wise tracking. This approach not only reduces the negative impact of sensor noise but also makes the method more practical, since RGB images are much easier to obtain and more widely available than RGB-D data.

Considering the scale ambiguity in the point maps by VGGT, we measure the transformation matrix between the two camera views and feed it into the depth prediction branch. By comparing the relative camera poses predicted by VGGT with the measured camera poses, we calibrate the depth prediction and achieve metric depth estimation.

**Object-Centric Point Cloud Segmentation.** To obtain object point clouds, we leverage advanced vision foundation models. Specifically, a text description is first provided to GroundingDINO [8] to generate target bounding boxes in the initial frame. These bounding boxes then serve as prompts for SAM2 [9], enabling real-time segmentation and tracking of the targets throughout the two third-person videos  $V^{view,1}$  and  $V^{view,2}$ . We segment all task-relevant objects in the video, dividing the resulting masks into the manipulated object mask  $M^{manip}$  and the context object mask  $M^{ctx}$ . Since two

cameras are used, we perform segmentation separately on each view, crop the corresponding depth regions, and back-project them into 3D space according to their respective camera poses. The resulting point clouds from both views are then fused, yielding the manipulated object point cloud  $P^{manip}$  and the context object point cloud  $P^{ctx}$ . The whole process is illustrated in Fig. 2.

**Object Point Cloud Encoder.** We extract features by concatenating the object-centric point clouds  $P^{manip}/P^{ctx}$  with  $M^{manip}/M^{ctx}$ , which indicate manipulated and context objects, as input to the Geometric Encoder, following DP3 [41].

**Object Transformation Estimation.** To obtain *pose transformation* of the the manipulate object between adjacent frames, we apply the Iterative Closest Point (ICP) algorithm to compute the rotation and translation  $T_t \in SE(3)$  at each time step  $t$ , aligning the manipulated object point clouds  $P_t^{manip}$  and  $P_{t+1}^{manip}$  computed from  $V^{view}$ . Benefiting from multi-view reconstruction, the object point clouds provide rich geometry, enabling reliable and effective ICP-based transformation estimation across frames.

### C. Tactile Prior

**Flow-based Tactile.** The *displacement field* of the sensor is computed by comparing the optical flow between the first frame  $V_0^{tac}$  and the current frame  $V_t^{tac}$ . The first frame is initialized under unloaded conditions, while the current frame may be either loaded or unloaded. We use the Dense Inverse Search (DIS) algorithm to estimate the optical flow. The optical flow is further decoupled into 3-D force vectors  $F_t^{tac} \in \mathbb{R}^{240 \times 320 \times 3}$  with a U-Net, following [40].

**Large-Scale Tactile Pretraining.** We use our portable gripper to collect a large-scale dataset of one million tactile images, covering diverse objects, gripping forces, gripping postures, and manipulation tasks. Among them, a large subset consists of high-quality samples with coarse annotations of the manipulated objects and their corresponding task labels. We adapt a ViT-based Tactile Encoder to extract features from tactile images. To mitigate the domain gap, we do not use pretrained weights from natural images; instead, we pretrain the encoder on our self-collected tactile dataset. The encoder is trained in a self-supervised manner using the Masked Autoencode [42] paradigm.

### D. Generative Trajectory Diffusion Modeling

**Diffusion Framework.** We model the distribution of future object transformations using a generative model over  $SE(3)$  transformations. This approach captures complex, multi-modal dynamics conditioned on past observations, and outperforms direct per-step regression methods. Recent Diffusion Policies sample joint-space action chunks via Langevin dynamics; to decouple from robot configuration, we instead predict the manipulated object’s rigid-body motion conditioned on past observations.

$$T_t^{k-1} = \alpha_k(T_t^k - \gamma_k \varepsilon_\theta(O_t, T_t^k, k) + \sigma_k \varepsilon^k), \quad (1)$$

where  $T_t^k \in \mathbb{R}^{T \times 4 \times 4}$ ,  $\varepsilon^k$  is a Gaussian noise, and  $\alpha_k, \gamma_k, \sigma_k$  are parameters of reverse process. The corresponding loss is

defined as:

$$\mathcal{L} = \mathbb{E}_{k, \varepsilon^k} \|\varepsilon^k - \varepsilon_\theta(T_t^0 + \varepsilon^k, k, O_t)\|_2^2. \quad (2)$$

**Conditioning Network.** Learning from human demonstration videos is hindered by the embodiment gap, the discrepancy between human demonstrators during training and robotic agents during deployment. To overcome this challenge, we utilize a fused object-centric feature representation  $\mathbf{f}_{oc}$ , which isolates and emphasizes the object of interest to ensure consistency across both phases.

To integrate tactile features into the object-centric point cloud feature, we design a lightweight **ResFiLM (Residual FiLM)** fusion module. Given a point-cloud feature  $\mathbf{f}_{pc} \in \mathbb{R}^D$  and aggregated tactile features  $\mathbf{f}_t \in \mathbb{R}^D$ , the module first projects them into a common latent space. Conditioned on the visual feature, we generate FiLM parameters  $(\gamma, \beta)$  that modulate the tactile feature, yielding a tactile embedding  $\mathbf{f}_t' = \gamma \cdot \mathbf{f}_t + \beta$ . The final fusion is performed via a gated residual connection:

$$\mathbf{f}_{oc} = \mathbf{f}_{pc} + \alpha(\gamma \odot \mathbf{f}_t' + \beta), \quad (3)$$

where  $\alpha$  is a learnable scaling coefficient. This design ensures stable training and allows tactile signals to act as adaptive residual corrections to the visual representation.

### E. Robot Control

We use the external camera’s extrinsic parameter to transfer the motion of the object in the camera’s frame to the robot’s frame. We transfer the  $SE(3)$  transformations of the object to the ones of the robot gripper. At inference step  $t$ , given the sampled transformation  $T_t^0$  (we omit the diffusion step 0 hereafter), the object is moved by the cumulative transformation  $\prod_{i=1}^t T_i = T_t \cdots T_1$ .

In addition to the predicted transformation, we augment the action representation with tactile feedback by appending the average contact force from the left and right finger pads. This enables the policy to explicitly utilize force information during control, rather than relying solely on geometric motion cues. The gripper is controlled with a PID-based scheme that takes the contact force predicted from demonstrations (rescaled for calibration) as the reference. Tactile feedback is incorporated during grasping to maintain a stable hold, enabling reliable force-controlled manipulation.

## IV. EXPERIMENTS

### A. Experimental Setup

**Hardware.** A workstation with an NVIDIA RTX 6000 Ada GPU is used for policy training and evaluation. A low-cost FAIRINO Robot 5 and Jodell Gripper are used for real-world manipulation. We have an NVIDIA A800 GPU for tactile encoder pretraining. We use 2 RealSense-D455 as the visual cameras. The human-held portable gripper is 3D-printed.

**Implementation.** Depending on task complexity, we collected 20 to 40 human demonstrations for each task, ensuring sufficient diversity in both demonstrations and evaluations (Time cost in Tab. IV). Our policy is trained by learning rate  $1e-4$ , with the observation horizon 2, and predicted steps

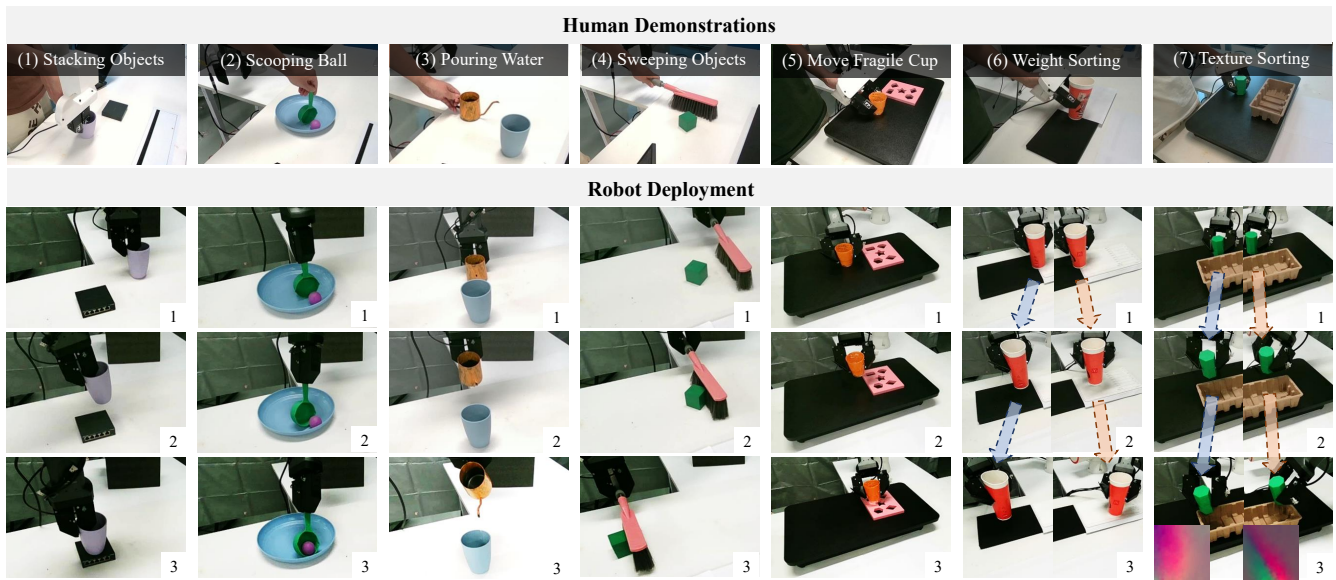


Fig. 3: **Experiments Visualization.** The first row shows a human demonstrator using either a hand or a portable gripper to collect demonstrations. The left four tasks are vision-only, while the right three are visuo-tactile. The following rows show the robot’s execution of our policy. The last two columns illustrate visuo-tactile tasks under different input conditions: in the Robot Deployment subfigures, the left and right insets correspond to different execution attempts, with arrows of different colors indicating motion directions. In the Weight Sorting task, cups are guided to distinct target locations based on mass; in the Texture Sorting task, objects are sorted by surface texture. For the Texture Sorting column (column 7, inset 3), we additionally show tactile images from grasps on different textures. These results demonstrate that vision alone is insufficient for reliable discrimination, highlighting the critical role of tactile perception in accurate decision-making.



Fig. 4: **Experimental setup for our system.** For the right part, the top half shows the robot and portable gripper, which uses the same tactile device. The bottom half shows our experimental objects. We display the camera setting in Fig. 2. For the left part, the image shows a subset of the objects we used to collect our large-scale tactile image dataset.

4 or 8, decided by the task. We begin each trial by moving the robot and gripper to a predefined pre-grasp pose.

**Evaluation Protocols.** Evaluation on 7 tasks—4 vision-only and 3 requiring tactile input. Each task runs 10 trials; we report human-judged mean success, discarding trials with external disturbances or hardware failures, so rates reflect intrinsic policy skill in precise, contact-rich manipulation:

**Vision-Only Tasks:** (1) *Stacking Objects.* This task requires the robot to grasp a plastic cup and place it on a box. A trial succeeds if the cup stays on the box’s top surface. (2) *Scooping Ball.* This task requires the robot to use a spoon to scoop a ball from a plate. Success requires lifting the ball and keeping it in the spoon. (3) *Pouring Water.* This task requires the robot to manipulate a teapot and attempt to pour water into a nearby cup. The task is considered successful if water enters the cup. (4) *Sweeping Objects.* This task requires the robot to use a broom to sweep a block across the table surface. Moving the block in the intended direction with a steady broom motion is regarded as a success.

**Visuo-Tactile Tasks:** (5) *Move Fragile Cup.* This task requires the robot to grasp and move a lightweight, fragile cup to a target place (a pink box), testing its ability to manipulate delicate objects. A trial is successful if the cup is placed without falling or noticeable deformation (e.g., crushing or bending). (6) *Weight Sorting.* This task requires the robot to recognize a given cup’s weight (empty or weighted). Placing the empty cups in the white area or the weighted cups in the black area is regarded as a success. (7) *Texture Sorting.* The robot distinguishes and sorts two visually similar objects with different geometric textures (hexagonal vs. octagonal prisms), placing each on its designated side of the tray. A trial is successful if both objects are correctly positioned.

### B. Baselines

For those vision-only tasks, we evaluate our method without tactile information, where all tactile-related modules are removed and the policy is conditioned solely on 3D point clouds. We consider two baseline variants: (a) **RGB-based Diffusion Policy**, abbreviated as “DP (RGB)”: The policy is trained conditioned only on RGB images without 3D point clouds. The object trajectories are still used as supervision during training to ensure consistent guidance across methods. (b) **3D-based Diffusion Policy without Object-Centric Modeling**, abbreviated as “DP (w/o OC)”: We extract 3D scene point clouds from VGGT, and feed them directly into the policy through our point cloud encoder, without applying object-centric modeling.

For those visuo-tactile tasks, we compare our complete OCRA against “**Ours without Tactile**”, where all tactile-related modules are discarded, and the policy relies solely on point cloud observations.

Task	Method	Success	Proc. Fail.	Out. Fail.
Stack	Ours	100%	0%	0%
	DP (RGB)	0%	100%	0%
	DP (w/o OC)	10%	50%	40%
Scoop	Ours	70%	30%	0%
	DP (RGB)	0%	70%	30%
	DP (w/o OC)	0%	100%	0%
Pour	Ours	70%	0%	30%
	DP (RGB)	0%	70%	30%
	DP (w/o OC)	30%	20%	50%
Sweep	Ours	100%	0%	0%
	DP (RGB)	90%	0%	10%
	DP (w/o OC)	50%	50%	0%

TABLE I: **Real-world manipulation task evaluation.** Outcomes are categorized into Success, Process Failure (Proc. Fail.), and Outcome Failure (Out. Fail.).

## V. RESULTS

### A. Vision-Only Results

We show our Vision-Only results in Tab. I. Our OCRA method achieves an 85.0% success rate across four challenging manipulation tasks. These tasks span a diverse range of difficulties: fluid dynamics in *Pouring Water*, dynamic object handling and tool-object interactions in *Scooping Ball* and *Sweeping Objects*, and precise object manipulation in *Stacking Cup*. The results highlight that our approach can robustly address tasks requiring reasoning over both rigid and non-rigid object behaviors. Our qualitative results could be seen in Fig. 3. Notably, in the highly challenging *Scooping Ball* scenario, we further evaluated the case where the ball is placed at the center of the bowl, a configuration that induces unstable dynamics. Remarkably, we observed that OCRA could robustly follow the ball’s motion through repeated collisions and still accomplish a successful scoop.

We further compare against baseline methods. The RGB-based Diffusion Policy achieves partial success; for instance, it can sweep objects, but often fails in other tasks. Typical errors include cups being lifted but moved randomly in stacking, and the spoon scooping prematurely and pushing the ball away. In pouring tasks, misalignment between the spout and the container is common. Overall, the policy struggles to generalize from limited demonstrations, with performance further degraded by variations in background, lighting, and manipulator differences (human versus robot). This indicates the benefit of incorporating 3D information.

Similar to the RGB-based variant, the Diffusion Policy without object-centric modeling achieves a substantially lower success rate than OCRA. While 3D spatial information offers modest gains in complex tasks such as *Pouring Water* and *Stacking Objects*, generalization remains limited due to point cloud noise and train-test discrepancies. In *Sweeping Objects*, performance is even worse than the RGB baseline, often stalling mid-execution. This supports our motivation for adopting an object-centric learning approach.

**More Analysis of the Performance Gap.** To further investigate the relatively low performance of the two baselines, we examined their failure cases and categorized them into two types: *outcome failures* and *process failures*. Outcome failures occur when a full trajectory is executed but ends

Method	Fragile Cup	Weight Sort	Texture Sort	Average
Ours	90%	90%	90%	90%
Ours(W/o Tactile)	60%	30%	30%	40%

TABLE II: Success rates of different methods on three evaluated real-world manipulation tasks

Method	Success	Proc. Fail.	Out. Fail.
<b>Ours</b>	90%	0%	10%
Ours (w/o Pretrain)	30%	20%	50%
Ours (w/o Fusion)	30%	40%	30%

TABLE III: **Ablations on the Texture Sort task.** Outcomes categorized as Success, Process Failure (Proc. Fail.), and Outcome Failure (Out. Fail.).

inaccurately. Process failures arise when the robot cannot complete a coherent trajectory. The two baselines predominantly suffer from process failures, indicating an inability to generate or sustain coherent trajectories, while our method eliminates such failures, demonstrating robustness in producing executable trajectories.

In summary, these comparison studies demonstrate that incorporating an object-centric 3D representation is critical for enabling robust generalization across diverse manipulation tasks. Our OCRA significantly outperforms baselines by leveraging spatial information and object centric understanding, thereby achieving state-of-the-art success rates in both rigid and non-rigid object interactions.

### B. Visuo-Tactile Results

Main visuo-tactile results are shown in Tab. II, where our OCRA with tactile information achieves a 90% success rate. See the qualitative results in Fig. 3. RGB-based and 3D-based DP fail on *Fragile Cup* and *Texture Sort*, with less than 30% success rate on *Weight Sort*, so we omit them in the table. Specifically, in the *Move Fragile Cup* task, our OCRA enables the gripper to securely grasp the cup with slight, non-damaging deformation, allowing stable lifting and placement without breakage in all trials; in the *Weight Classification* task, cups are consistently grasped without damage. Heavy cups are placed on the right and empty cups on the left, supported by tactile shear force signals that provide discriminative weight information [43]; in the *Texture Sort* task, the two prismatic objects are visually similar, making vision alone insufficient. Tactile sensing enabled reliable category discrimination by guiding placement decisions, while vision planned the trajectory.

Besides, for OCRA without tactile feedback, the average success rate dropped to 40%. Vision alone could not distinguish subtle shape differences or detect cup weight, causing frequent errors in the sorting tasks. In the fragile cup task, inconsistent grasping forces led to cup deformation, unnatural trajectories, and placement failures.

### C. Ablations

After establishing the vision-only baseline, we further investigate how tactile information contributes to policy learning in the visuo-tactile setting. To this end, we conduct two ablation studies: (a) **Without tactile pretraining:** The tactile encoder is randomly initialized, while the normalization statistics remain derived from the large-scale dataset;




Fig. 5: **Visualization of camera views.** We use the two cameras marked by green boxes during both the demonstration collection and testing. To evaluate the view generalization ability of OCRA, the left green-boxed camera remains fixed, while the camera in both green and yellow boxes is replaced according to the arrow direction.

Method	Stack	Sweep	Scoop	Pour	Tactile Tasks
Ours	8	8	8	8	5
Teleoperation	20	20	N/A	N/A	N/A

TABLE IV: **Demonstration collection time (seconds per task).** “N/A” denotes tasks that are not applicable or cannot be performed by the method. Our approach achieves substantially lower demonstration time and greater adaptability compared to teleoperation.

(b) **Without Residual FiLM Fusion:** The fusion module is removed, and point cloud features are directly concatenated with tactile features. As shown in Tab. III.

Our analysis indicates that our proposed masked tactile pre-training enhances the feature discriminability of tactile inputs, enabling the model to extract more informative and noise-resilient features, thereby improving the stability of both trajectory generation and final stage execution. Furthermore, the design of the fusion module is critical for generating reliable trajectories. When point cloud and tactile features are simply concatenated, the model is easily misled by noisy tactile signals, resulting in erroneous and unstable trajectories, proving the necessity of our Residual FiLM. Similar phenomenon is observed in Weight Sort Task.

#### D. View Generalization

A common challenge in visuomotor learning is sensitivity to camera viewpoints, where even small perspective changes can affect perception and action execution. To evaluate whether our method overcomes this challenge, we replaced one of the original cameras with a novel third-view camera and tested it on the *Stacking Objects* task, which critically depends on recognizing object spatial relationships. In 10 out of 10 trials, our method successfully completed the task with smooth trajectories. Remarkably, performance remained stable even when the cup is partially outside the new camera’s field of view, indicating that our approach reuses learned perceptual representations in a viewpoint-invariant manner rather than memorizing specific camera configurations. Qualitative results in Fig. 5 further illustrate robust generalization under unseen viewpoints.

#### E. Further Analysis

**Demonstration Collection Efficiency.** Tab. IV compares demonstration times across tasks. Our method consistently




Fig. 6: **Visualization of depth reconstruction quality.** We compare the depth map generated by VGGT with the camera ground-truth (GT) depths from the RGBD cameras on the Texture Sort task. The GT depth contains numerous holes and often produces misaligned point clouds due to measurement instability (see red points in the figure), whereas VGGT yields smoother depth maps and more consistent, regular point clouds.

Task	Depth Setting	Success Rate
Stack	Ours	100%
	Single-View GT Depth	60%
Texture Sort	Ours	90%
	Bi-view GT Depth	50%

TABLE V: **Depth analysis across tasks.** Success rates for our reconstruction compared with the camera’s ground-truth (GT) depth baselines.

requires only 8 seconds per demonstration (5 seconds for tactile tasks), achieving over a  $2\times$  speedup compared to teleoperation (20 seconds). Moreover, teleoperation is infeasible for *Scoop* and *Pour* due to complex object–object dynamics, and for tactile tasks due to the lack of direct force feedback. In contrast, our approach handles all tasks efficiently, highlighting both its speed and adaptability.

**Flexibility to Ground-Truth Depth During Inference.** Our framework is designed as a depth-free approach, leveraging multi-view RGB reconstruction for both data processing and deployment. Nevertheless, when ground-truth depth is available (e.g., from depth sensors), it can be directly used during testing with mild performance drop (less than 20%) on Stacking, Sweeping and Texture/Weight Sort.

**Comparing OCRA with Ground-Truth Depth.** Tab. V presents a comparison between our OCRA with 3D reconstruction and OCRA models trained and tested using ground-truth depth baselines. On the *Stack* task, which relies solely on vision, single-view depth fails to capture object–object interactions, leading to collisions. On the *Texture Sort* task, which combines vision and touch, bi-view depth yields only 50% success, as small objects suffer from missing regions in the depth maps and merging views introduces misalignment from calibration errors. As shown in Fig. 6, ground-truth depth often suffers from holes and noisy point clouds, while VGGT yields smoother and more consistent geometry. This richer reconstruction enables more reliable execution, explaining the performance gap in Tab. V.

## VI. CONCLUSIONS

We presented OCRA, an object-centric framework for human-to-robot imitation learning directly from RGB demonstration videos. By integrating multi-view reconstruction with VGGT, object-centric point cloud segmentation and feature extraction, large-scale tactile data collection and pre-training, and generative trajectory diffusion modeling with a lightweight multimodal (3D+Tactile) fusion module, OCRA provides rich object-centric representations that enable robust

action transfer from human to robot. Experiments on both vision-only and visuo-tactile tasks validate the effectiveness of this design, consistently outperforming baselines.

## ACKNOWLEDGMENT

Prof. Yanwei Fu is the core project leader. The paper is supported by Science and Technology Commission of Shanghai Municipality (No. 24511103100), and National Natural Science Foundation of China (62521004).

## REFERENCES

- [1] C. Chi, Z. Xu, C. Pan, E. Cousineau, B. Burchfiel, S. Feng, R. Tedrake, and S. Song, "Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots," *arXiv preprint arXiv:2402.10329*, 2024.
- [2] Z. Zhaxizhuoma, K. Liu, C. Guan, Z. Jia, Z. Wu, X. Liu, T. Wang, S. Liang, P. Chen, P. Zhang *et al.*, "Fastumi: A scalable and hardware-independent universal manipulation interface with dataset," in *Conference on Robot Learning*. PMLR, 2025, pp. 3069–3093.
- [3] S. Nair, A. Rajeswaran, V. Kumar, C. Finn, and A. Gupta, "R3m: A universal visual representation for robot manipulation," *arXiv preprint arXiv:2203.12601*, 2022.
- [4] Y. J. Ma, S. Sodhani, D. Jayaraman, O. Bastani, V. Kumar, and A. Zhang, "Vip: Towards universal visual reward and representation via value-implicit pre-training," *arXiv preprint arXiv:2210.00030*, 2022.
- [5] C.-C. Hsu, B. Wen, J. Xu, Y. Narang, X. Wang, Y. Zhu, J. Biswas, and S. Birchfield, "Spot: Se (3) pose trajectory diffusion for object-centric manipulation," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 4853–4860.
- [6] Y. Zhu, Z. Jiang, P. Stone, and Y. Zhu, "Learning generalizable manipulation policies with object-centric 3d representations," *arXiv preprint arXiv:2310.14386*, 2023.
- [7] J. Wang, M. Chen, N. Karaev, A. Vedaldi, C. Rupprecht, and D. Novotny, "Vggg: Visual geometry grounded transformer," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 5294–5306.
- [8] S. Liu, Z. Zeng, T. Ren, F. Li, H. Zhang, J. Yang, Q. Jiang, C. Li, J. Yang, H. Su *et al.*, "Grounding dino: Marrying dino with grounded pre-training for open-set object detection," in *ECCV*, 2024.
- [9] N. Ravi, V. Gabeur, Y.-T. Hu, R. Hu, C. Ryali, T. Ma, H. Khedr, R. Rädle, C. Rolland, L. Gustafson *et al.*, "Sam 2: Segment anything in images and videos," *arXiv preprint arXiv:2408.00714*, 2024.
- [10] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," *The International Journal of Robotics Research*, vol. 44, pp. 1684 – 1704, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:257378658>
- [11] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning fine-grained bimanual manipulation with low-cost hardware," *arXiv preprint arXiv:2304.13705*, 2023.
- [12] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter *et al.*, " $\pi_0$ : A vision-language-action flow model for general robot control," *arXiv preprint arXiv:2410.24164*, 2024.
- [13] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. Foster, G. Lam, P. Sanketi *et al.*, "Openvla: An open-source vision-language-action model," *arXiv preprint arXiv:2406.09246*, 2024.
- [14] D. A. Pomerleau, "Alvin: An autonomous land vehicle in a neural network," *NeurIPS*, 1988.
- [15] K. Grauman, A. Westbury, E. Byrne, Z. Chavis, A. Furnari, R. Girdhar, J. Hamburger, H. Jiang, M. Liu, X. Liu *et al.*, "Ego4d: Around the world in 3,000 hours of egocentric video," in *CVPR*, 2022.
- [16] B. Wen, W. Yang, J. Kautz, and S. T. Birchfield, "Foundationpose: Unified 6d pose estimation and tracking of novel objects," *CVPR*, 2024.
- [17] J. Yuan, T. Chen, B. Li, and X. Xue, "Compositional scene representation learning via reconstruction: A survey," *IEEE TPAMI*, 2023.
- [18] J. Jiang, S. Janghorbani, G. de Melo, and S. Ahn, "Scalar: Generative world models with scalable object representations," in *ICLR*, 2020.
- [19] Z. Lin, Y.-F. Wu, S. V. Peri, W. Sun, G. Singh, F. Deng, J. Jiang, and S. Ahn, "Space: Unsupervised object-oriented scene representation via spatial attention and decomposition," in *ICLR*, 2020.
- [20] F. Locatello, D. Weissenborn, T. Unterthiner, A. Mahendran, G. Heigold, J. Uszkoreit, A. Dosovitskiy, and T. Kipf, "Object-centric learning with slot attention," *NeurIPS*, 2020.
- [21] M. Seitzer, M. Horn, A. Zadaianchuk, D. Zietlow, T. Xiao, C.-J. Simon-Gabriel, T. He, Z. Zhang, B. Schölkopf, T. Brox *et al.*, "Bridging the gap to real-world object-centric learning," *ICLR*, 2023.
- [22] T. Kipf, G. F. Elsayed, A. Mahendran, A. Stone, S. Sabour, G. Heigold, R. Jonschkowski, A. Dosovitskiy, and K. Greff, "Conditional object-centric learning from video," *arXiv preprint arXiv:2111.12594*, 2021.
- [23] G. Singh, Y.-F. Wu, and S. Ahn, "Simple unsupervised object-centric learning for complex and naturalistic videos," *NeurIPS*, 2022.
- [24] A. Zadaianchuk, M. Seitzer, and G. Martius, "Self-supervised visual reinforcement learning with object-centric representations," *arXiv preprint arXiv:2011.14381*, 2020.
- [25] J. Yoon, Y.-F. Wu, H. Bae, and S. Ahn, "An investigation into pre-training object-centric representations for reinforcement learning," *arXiv preprint arXiv:2302.04419*, 2023.
- [26] S. Ferraro, P. Mazzaglia, T. Verbelen, and B. Dhoedt, "Focus: object-centric world models for robotic manipulation," *Frontiers in Neuro-robotics*, vol. 19, p. 1585386, 2025.
- [27] M. Mosbach, J. N. Ewertz, A. Villar-Corrales, and S. Behnke, "Sold: Reinforcement learning with slot object-centric latent dynamics," *arXiv preprint arXiv:2410.08822*, 2024.
- [28] C. Devin, P. Abbeel, T. Darrell, and S. Levine, "Deep object-centric representations for generalizable robot learning," in *ICRA*, 2018.
- [29] D. Wang, C. Devin, Q.-Z. Cai, F. Yu, and T. Darrell, "Deep object-centric policies for autonomous driving," in *ICRA*, 2019.
- [30] J. Tremblay, T. To, B. Sundaralingam, Y. Xiang, D. Fox, and S. Birchfield, "Deep object pose estimation for semantic robotic grasping of household objects," *arXiv preprint arXiv:1809.10790*, 2018.
- [31] X. Deng, Y. Xiang, A. Mousavian, C. Eppner, T. Bretl, and D. Fox, "Self-supervised 6d object pose estimation for robot manipulation," in *ICRA*, 2020.
- [32] Y. Zhu, A. Joshi, P. Stone, and Y. Zhu, "Viola: Imitation learning for vision-based manipulation with object proposal priors," in *CoRL*, 2023.
- [33] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo *et al.*, "Segment anything," in *ICCV*, 2023.
- [34] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin, "Emerging properties in self-supervised vision transformers," in *ICCV*, 2021.
- [35] M. Oquab, T. Darcet, T. Moutakanni, H. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby *et al.*, "Dinov2: Learning robust visual features without supervision," *arXiv preprint arXiv:2304.07193*, 2023.
- [36] W. Yuan, S. Dong, and E. H. Adelson, "Gelsight: High-resolution robot tactile sensors for estimating geometry and force," *Sensors*, 2017.
- [37] E. Donlon, S. Dong, M. Liu, J. Li, E. Adelson, and A. Rodriguez, "Gelslim: A high-resolution, compact, robust, and calibrated tactile-sensing finger," in *IROS*, 2018.
- [38] B. Huang, Y. Wang, X. Yang, Y. Luo, and Y. Li, "3d-vitac: Learning fine-grained manipulation with visuo-tactile sensing," in *CoRL*, 2024. [Online]. Available: <https://api.semanticscholar.org/CorpusID:273707578>
- [39] L. Wu, C. Yu, J. Ren, L. Chen, R. Huang, G. Gu, and H. Li, "Freetac-man: Robot-free visuo-tactile data collection system for contact-rich manipulation," *arXiv preprint arXiv:2506.01941*, 2025.
- [40] H. Luo, X. Zhang, Y. Xu, J. Yu, X. Ma, and W.-M. Chen, "A novel flexible multidimensional force platform for assessing regional plantar pressure and shear stress under the foot during gait," *IEEE Sensors Journal*, 2024.
- [41] Y. Ze, G. Zhang, K. Zhang, C. Hu, M. Wang, and H. Xu, "3d diffusion policy: Generalizable visuomotor policy learning via simple 3d representations," *arXiv preprint arXiv:2403.03954*, 2024.
- [42] K. He, X. Chen, S. Xie, Y. Li, P. Doll'ar, and R. B. Girshick, "Masked autoencoders are scalable vision learners," *CVPR*, 2022.
- [43] J. H. Kang, S. Joshi, R. Huang, and S. K. Gupta, "Robotic compliant object prying using diffusion policy guided by vision and force observations," *IEEE RA-L*, 2025.