

# ROPA: Synthetic Robot Pose Generation for RGB-D Bimanual Data Augmentation

Jason Chen, I-Chun Arthur Liu, Gaurav S. Sukhatme\*, Daniel Seita  
 Viterbi School of Engineering, University of Southern California




Fig. 1: ROPA performs offline data augmentation for bimanual imitation learning. White arrows indicate pose differences between the original and augmented images. Red regions represent ROPA generated images and states every  $k$  timesteps at  $t+k$  and  $t+2k$ , while blue regions show the original dataset. RGB and depth image pairs are captured at the same timesteps, with the top row displaying depth colormap and the bottom row showing standard RGB images.

**Abstract**—Training robust bimanual manipulation policies via imitation learning requires demonstration data with broad coverage over robot poses, contacts, and scene contexts. However, collecting diverse and precise real-world demonstrations is costly and time-consuming, which hinders scalability. Prior works have addressed this with data augmentation, typically for either eye-in-hand (wrist camera) setups with RGB inputs or for generating novel images without paired actions, leaving augmentation for eye-to-hand (third-person) RGB-D training with new action labels less explored. In this paper, we propose **Synthetic Robot Pose Generation for RGB-D Bimanual Data Augmentation (ROPA)**, an offline imitation learning data augmentation method that fine-tunes Stable Diffusion to synthesize third-person RGB and RGB-D observations of novel robot poses. Our approach simultaneously generates corresponding joint-space action labels while employing constrained optimization to enforce physical consistency through appropriate gripper-to-object contact constraints in bimanual scenarios. We evaluate our method on 5 simulated and 3 real-world tasks. Our results across 2625 simulation trials and 300 real-world trials demonstrate that ROPA outperforms baselines and ablations, showing its potential for scalable RGB and RGB-D data augmentation in eye-to-hand bimanual manipulation. Our project website is available at: <https://ropaug.github.io/>.

## I. INTRODUCTION

Bimanual manipulation is critical for a wide range of daily tasks [1] such as lifting large objects [2, 3], handling deformable objects [4]–[7], and opening containers [8, 9]. These activities require coordinated motion of both arms and awareness of nearby objects and surfaces. A third-person

camera can capture both arms and the surrounding scene from one viewpoint, which is appealing for vision-based imitation learning [10]. Recent imitation learning approaches show that large demonstration datasets yield increasingly general bimanual policies [11]–[13], but collecting sufficiently large and diverse action-labeled data is costly, which limits scalability.

Data augmentation has emerged as a promising strategy to address this bottleneck, particularly in single-arm manipulation, where novel views and corresponding action labels can be synthesized offline [14]. However, extending these techniques to bimanual manipulation, depth images, and beyond wrist-mounted (eye-in-hand) viewpoints [15, 16], introduces new difficulties. These include enforcing scene-wide visual consistency across two robot arms, preserving action label correctness, and handling the higher degrees-of-freedom (DOF).

In this paper, we propose **Synthetic Robot Pose Generation for RGB-D Bimanual Data Augmentation (ROPA)**, a data augmentation method for vision-based imitation learning of bimanual manipulation. Our key insight in ROPA is to adapt Pose-Guided Person Image Synthesis techniques [17] (originally developed for human pose generation) by conditioning a diffusion model on robot joint configurations. This enables the diffusion model to yield realistic and consistent third-person views for varied robot arm poses. We maintain a tight correspondence between augmented visual observations and valid actions, so that policies trained on the augmented dataset remain executable on real hardware.

We evaluate ROPA across 5 simulated and 3 real-world bimanual manipulation tasks, based on the PerAct2 [18] benchmark in simulation and a physical bimanual UR5 setup. See Figure 1 for an example rollout. Compared to an Action Chunking with Transformers (ACT) [19] baseline trained only on raw demonstrations, data augmentation using ROPA enables ACT to achieve higher success rates across bimanual manipulation tasks that require coordinated motion and fine-grained precision. By enabling third-person vision-based data augmentation for bimanual systems, ROPA scales training data and broadens coverage without extra manual data collection.

The contributions of this paper include:

- ROPA, a novel pose-guided image synthesis method for offline data augmentation in bimanual manipulation, which supports learning from both RGB and RGB-D data.
- Depth image synthesis that generates depth maps consistent with the augmented joint positions and RGB images.
- Action-consistent augmentation that outputs images and joint-space labels, with constraints to ensure feasibility.
- Simulation and real-world experiments showing that bimanual policies with ROPA achieve significantly improved performance over baseline methods and ablations.

## II. RELATED WORK

### A. Bimanual Manipulation Methods

Compared to single-arm setups, bimanual systems introduce significant complexity due to larger degrees of freedom and coordination requirements. Classical approaches rely on geometric methods [6] or constraint-based representations [20], but assume strong task-specific priors that limit generalization. This has motivated learning-based methods using reinforcement learning and imitation learning. While reinforcement learning shows progress [21]–[24], scaling to robust real-world bimanual policies remains difficult due to high-dimensional control and heavy environment interaction requirements [25].

Imitation learning has therefore emerged as a dominant strategy for “robotic foundation model” manipulation systems [11, 12], since demonstrations directly capture desired patterns without exploration [10]. Advances in algorithms [19, 26], hardware [27]–[29], and datasets [30] have accelerated this trend. Despite this progress, data scalability remains a bottleneck. Our contribution is orthogonal to advances in teleoperation and algorithm design; we study imitation learning from third-person views where data augmentation improves robustness without requiring new demonstrations.

### B. Data Augmentation for Imitation Learning

Data augmentation synthetically expands training data at lower cost than collecting new data. In vision classification, simple transformations (e.g., crops, jitters) with preserved labels improve generalization [31]. However, imitation learning suffers from error accumulation, driving policies into unseen states [32, 33] which motivates augmentation methods specifically for robotics.

Visual augmentation approaches include altering visual context through structured scene augmentations [34]–[37],

changing backgrounds or objects while keeping robot actions fixed. In contrast, we augment robot poses. Other approaches leverage 3D Gaussian Splatting [38] but require high-fidelity scene reconstruction, or focus on state-based inputs [39, 40] and wrist-camera observations [15, 16, 41]. We study third-person image-based learning and generate full robot poses for richer environmental context.

Unlike approaches that augment new robot states and actions through simulation rollouts [42, 43], our method operates offline without requiring interactive data collection. Recent diffusion-based methods like VISTA [14] and RoVi-Aug [44] enable novel-view synthesis and cross-embodiment transfer. However, these approaches do not generate new action labels and have not been evaluated on bimanual systems. While RoVi-Aug generates multiple views of the same robot pose, ROPA generates multiple robot poses from up to four camera viewpoints. Finally, there is relatively limited work studying depth augmentation for robotics. A recent effort [45] uses standard image augmentation techniques to RGB-D images but does not generate new robot poses or action labels.

### C. Pose-Guided Image Synthesis

Pose-guided image synthesis generates realistic images of people in target poses while preserving identity and background context [17]. Recent advances leverage diffusion models for higher fidelity and controllability [46, 47]. Diffusion-based pose generation has also begun appearing in robotics. GENIMA [48] fine-tunes Stable Diffusion with ControlNet [49] to visualize future joint-action targets as images, then translates them into executable actions. Other approaches employ ControlNet for temporal consistency during action sequence generation [50]. However, these methods focus on visuomotor control rather than generating image-action pairs for imitation learning data augmentation. Existing works explore differentiable robot rendering [51] using 3D Gaussian Splatting, but such approaches [52] focus on robot pose reconstruction from images and videos rather than pose generation. We build on ControlNet’s flexibility for guiding diffusion models with structural cues and apply it to bimanual manipulation with skeleton-based conditioning to produce realistic and consistent images.

## III. PROBLEM STATEMENT

Our focus is on vision-based, eye-to-hand imitation learning for bimanual manipulation, where a policy  $\pi_\theta$  with parameters  $\theta$  is trained from expert demonstration data using third-person RGB or RGB-D camera images.

We denote the source camera image at time  $t$  by  $I_t^s$  (where  $s$  denotes source) and the target camera image by  $I_t^d$  (where  $d$  denotes target), with the camera intrinsic parameters and extrinsic parameters (camera pose relative to world frame). We also denote the left and right robot arms using superscripts  $l$  and  $r$ , and indicate their base positions relative to world frame as  $b^l$  and  $b^r$ , respectively. The source camera image serves as the policy input which produces actions  $\mathbf{a}_t = \pi_\theta(I_t^s)$ . The action  $\mathbf{a}_t = (\mathbf{a}_t^l, \mathbf{a}_t^r)$  specifies target joint positions for the left

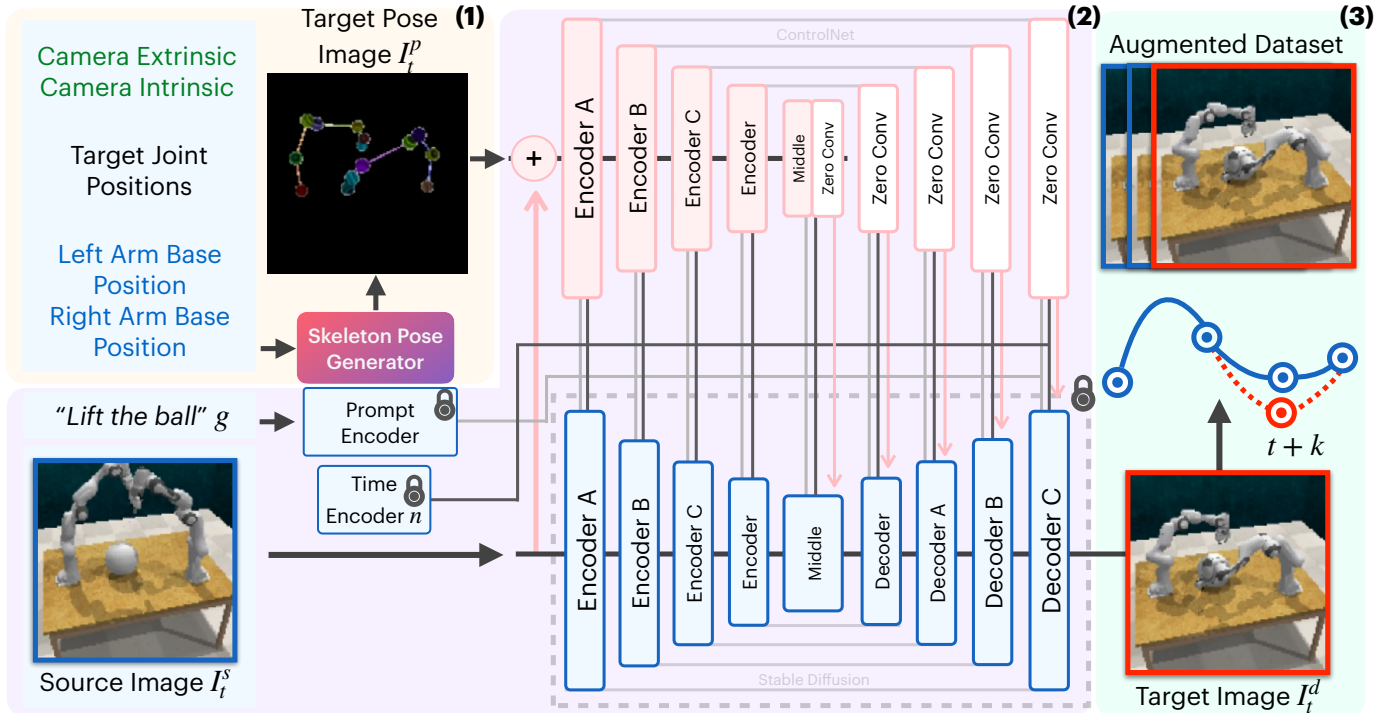


Fig. 2: **ROPA Overview.** (1) The Skeleton Pose Generator takes camera extrinsics and intrinsics, target joint positions, and left and right robot base positions to generate a skeleton pose image  $I_t^p$  representing the target joint configuration. (2) The source image  $I_t^s$  and language goal  $g$  are fed into Stable Diffusion (the bottom U-Net model), while the generated skeleton pose serves as control input to ControlNet (the top U-Net model), producing the target image  $I_t^d$ . The locked icons represent frozen parameters. (3) The original dataset is duplicated and **generated target states** replace the original states every  $k$  timesteps (see Section IV-C for more details), with updated corresponding action labels. This augmented dataset is combined with the **original dataset** to train a bimanual manipulation policy.

and right arms from the policy. Training relies on a dataset of  $M$  expert demonstrations  $\mathcal{D} = \{\tau_1, \dots, \tau_M\}$ , where each  $\tau_i$  is a sequence of third person camera image observations and corresponding actions:  $\tau_i = (I_1^s, \mathbf{a}_1^l, \mathbf{a}_1^r, \dots, I_T^s, \mathbf{a}_T^l, \mathbf{a}_T^r)$  for a demonstration with  $T$  time steps.

#### IV. METHOD: ROPA

ROPA leverages a conditional diffusion model to generate novel robot poses while preserving scene context. The diffusion model takes in a source camera image  $I_t^s$ , target skeleton pose image  $I_t^p$ , and language goal  $g$  as input, and outputs the target camera image  $I_t^d$ . The key objective is to translate a source robot image into a desired target pose. This is accomplished through a three-stage process: fine-tuning Stable Diffusion on expert demonstrations (Section IV-A), generating the skeleton pose for the new robot poses (Section IV-B), and generating action labels to synthesize new robot poses and reconstructing the dataset (Section IV-C). Figure 2 provides an overview.

##### A. Pose Guided Robot Image Synthesis (PGRIS)

Inspired by Pose Guided Person Image Synthesis (PG-PIS) [17], we formulate the problem of **Pose Guided Robot Image Synthesis (PGRIS)**. PGRIS aims to synthesize novel robot configurations while maintaining environmental consistency. Formally, given a source camera image  $I_t^s$ , a target

skeleton pose image  $I_t^p$  (described in Section IV-B), and a language goal  $g$ , PGRIS seeks to learn a mapping:

$$I_t^d = f_\psi(I_t^s, I_t^p, g), \quad (1)$$

where  $f_\psi$  is a conditional generative model parameterized by  $\psi$ . In our implementation,  $f_\psi$  corresponds to a **Stable Diffusion 2.1** [54] model fine-tuned with ControlNet [49]. ControlNet employs a two-stream architecture: a frozen Stable Diffusion U-Net that processes noisy latents and text embeddings, and a copy of the Stable Diffusion U-Net that serves as the trainable encoder stream processing conditioning images. The source image  $I_t^s$  and target pose image  $I_t^p$  are encoded into latent representations  $z_s = E(I_t^s)$  and  $z_{pose} = E(I_t^p)$  using the pretrained Stable Diffusion encoder  $E$ . The language goal  $g$  is embedded via a text encoder [55]:  $c_{text} = \xi(g)$ .

The trainable ControlNet module  $\mathcal{C}_\phi$  processes the conditioning inputs  $(z_s, z_{pose})$  and injects pose-aware feature maps into the frozen U-Net denoiser  $\varepsilon_\phi$  through zero-initialized convolution layers [49]. These zero-convolutions encode spatial constraints that guide the denoising process to the desired robot configuration while preventing harmful noise from affecting the network during early training phases. At training time, Gaussian noise  $\varepsilon$  is added to the target latent  $z_d = E(I_t^d)$  at diffusion step  $n$  following the DDPM noise schedule [56], producing  $z_{d,n} = \sqrt{\alpha_n}z_d + \sqrt{1 - \alpha_n}\varepsilon$ . The model is trained to




Fig. 3: **Skeleton Pose Ablations and Visualization.** Comparison of different skeleton pose formats: (1) ROPA’s skeleton pose, (2) OpenPose [53] inspired skeleton pose, and (3) an all white Skeleton Pose (less visual contrast). (4) demonstrates precise alignment between ROPA’s skeleton pose and the source image. (5) shows a source image input for multi-view generation, while (6) displays the skeleton pose for both robots overlaid on that same source image.

predict the noise term by minimizing the following loss:

$$\mathcal{L} = \mathbb{E}_{z_d, z_s, z_{pose}, g, n, \epsilon} \left\| \epsilon - \epsilon_\phi(z_d, n; c_{text}, n; \mathcal{C}_\phi(z_s, z_{pose})) \right\|_2^2 \quad (2)$$

During inference, the model iteratively denoises a random Gaussian latent with ControlNet guidance to produce the final target image  $I_t^d = D(z_d^*)$  using DDIM sampling [57]. We train the diffusion model on expert demonstrations  $\mathcal{D}$  for a fixed number of iterations, then use the trained model with sampled camera perturbations to synthesize novel views from  $\mathcal{D}$ .

### B. Skeleton Pose Generator

In order to train our diffusion model from Section IV-A, we need a scalable method to render large numbers of images showing diverse robot skeleton poses that maintain geometric consistency. Inspired by conventions in human pose estimation frameworks such as OpenPose [53], we create a skeletal representation of each arm’s kinematic chain using cylinders for bone segments and spheres for joint connectors. This provides explicit geometric structure that enables the diffusion model to understand spatial relationships and maintain kinematic consistency during image generation. We apply striped textures from GENIMA [48] to the spheres (see panel 1 in Figure 3) to provide visual cues that help Stable Diffusion learn joint orientations and rotational states for image generation. Our approach does not require explicit 3D information because depth-varying spheres and textural cues help encode spatial relationships for the diffusion model.

To generate the target skeleton pose  $I_t^p$ , we apply an end-effector transformation  $\Delta\rho$  (Section IV-C shows how we generate the transformation) to the source end-effector poses and solve inverse kinematics to obtain the target joint positions. We then configure the virtual arms to these joint positions and render the scene using the configured virtual camera, capturing the skeleton pose image with identical dimensions  $H \times W$  as the source camera image  $I_t^s$ .

To achieve this, we leverage PyRender [58], a Python-based 3D rendering library, to create skeletal visualizations of the robot arms. We configure PyRender’s virtual camera to match

the intrinsic and extrinsic parameters from the source camera that captured  $I_t^s$ , ensuring geometric consistency between the rendered skeleton and the real camera view (see panel 4 in Figure 3). We load URDF models of both robot arms in the virtual environment and position them at their respective base locations  $b^l$  and  $b^r$  to match the source camera image setup.

### C. Action Labeling and Dataset Construction

The previous approach enables image synthesis but lacks constraint-enforced action sampling to ensure valid perturbations. To avoid arm collisions or uncoordinated perturbations in coordinated tasks, we leverage camera pose sampling for bimanual manipulation by decomposing tasks into contactless and contact-rich states. For contact detection, we adopt the force-based approach of Zhang et al. [59], modeling external torque residuals as  $\mu_{ext} = \mu_{motor} - \hat{\mu}_{model}$  using an autoregressive model. Contact is detected when actual torque exceeds dynamic thresholds for consecutive timesteps. This approach offers advantages over vision-based detection: (1) it works with third-person cameras where contact points may be obscured, and (2) it avoids the computational overhead of external modules for detection and segmentation (e.g., using a vision-language model for detecting visual contacts).

We apply pose perturbations  $\Delta\rho$  to source joint positions, where  $\Delta\rho = {}_{p_s}T_{p_d}$  represents the end-effector transformation from source pose  $p_s$  to target pose  $p_d$ . The joint-space perturbations are computed by solving inverse kinematics for the transformed end-effector poses. The perturbation strategy employs uniform sampling for contactless states and constrained optimization for contact phases to ensure kinematically feasible poses while preserving contact relationships. The perturbations  $\Delta\rho$  are used to compute perturbed actions  $\tilde{\mathbf{a}}_t = (\tilde{\mathbf{a}}_t^l, \tilde{\mathbf{a}}_t^r)$  from the original  $\mathbf{a}_t$ , where the perturbed actions correspond to the joint positions required to achieve the transformed end-effector poses. We duplicate the original dataset, replacing original states with augmented out-of-distribution states every  $k$  timesteps ( $t+k$ ,  $t+2k$ , etc.) to mitigate compounding errors in behavior cloning policies, where we set  $k=8$  based on trials




Fig. 4: **Depth Image Generation.** Condensed variation of the pipeline in Fig. 2 for depth image synthesis. (1) Source depth colormap image input to Stable Diffusion. (2) RGB target image and skeleton pose provide conditioning inputs to ControlNet. (3) Generated target depth image.

with different  $k$  values. This process generates an augmented dataset  $\mathcal{D}$  of novel images paired with corresponding action labels, which we combine with the original dataset.

#### D. Multi-Modal and Multi-View Extensions

We extend our approach to incorporate RGB-D data, as depth information is crucial for robot manipulation by providing geometric context and enabling 3D spatial understanding.

For depth image augmentation, we adapt the PGRIS architecture to generate depth-consistent synthetic data. Instead of using an RGB source image  $I_t^s$ , we input a depth colormap as the source image (see Figure 4, panel 1). To ensure geometric consistency between synthesized RGB and depth modalities, we condition ControlNet on both the corresponding RGB target image and target pose image  $I_t^p$  (see Figure 4, panel 2). We modify ControlNet to accept this multimodal input by concatenating the two RGB images (6 channels total) as input. This dual conditioning ensures generated depth images maintain geometric correspondence with RGB images (see Figure 4, panel 3). The final augmented dataset contains paired RGB-depth images that are geometrically consistent.

We also extend ROPA beyond single-camera third-person viewpoints to support multi-camera setups with up to four cameras. Since multiple camera viewpoints are crucial for avoiding occlusions and improving spatial robustness, we address this challenge by tiling four camera observations into a single composite image (see Figure 3, panels 5 and 6). This tiling approach enables multi-view consistent image generation, aligning with GENIMA’s [48] observations that tiling multiple viewpoints improves consistency.

## V. SIMULATION EXPERIMENTS

### A. Simulation Experiment Setup

We conduct simulation experiments using the open-source code from PerAct2 [18], which builds upon RLbench [60] to provide a benchmark for bimanual manipulation. We simplify certain tasks by reducing axes of variation (e.g., constraining the workspace) to enhance the performance of the Action Chunking with Transformers (ACT) baseline [19]. We use the following simulation tasks:

- **Coordinated Lift Ball (CLB):** Lift a randomly spawned ball.
- **Coordinated Lift Tray (CLT):** Lift a tray with an item above.
- **Coordinated Push Box (CPB):** Push a large box to a target area.
- **Bimanual Straighten Rope Easy (BSR):** Place both rope endpoints in their target areas.
- **Coordinated Put Item in Drawer (CPID):** Place an item inside one of three specified drawers.

We compare ROPA against these strong baselines:

- **ACT [19] (w/o Augmentation):** A state-of-the-art imitation learning method for bimanual manipulation trained on 100 original demonstrations.
- **ACT [19] (more data):** Trained on the original dataset along with 100 additional demonstrations without data augmentation, providing a performance reference for ACT with expanded expert data.
- **VISTA [14]:** A diffusion-based novel view synthesis method that uses ZeroNVS [61] to augment third-person viewpoints from a single third-person view. We extend this to a bimanual manipulation setting.

We ensure fair comparisons by using identical training, validation, and test data with consistent environment seeds across all methods. Demonstrations are generated via RLbench’s waypoint-based motion planner. To maintain consistency with our RGB-only evaluation policy, we adapt ACT to handle both RGB and depth modalities. We modify the input architecture by expanding from 3 to 6 channels. For VISTA, we train ZeroNVS with VISTA’s default fine-tuning parameters.

To our knowledge, no prior work has addressed depth-based data augmentation for robotic manipulation. Existing depth estimation models such as DepthAnything [62] are unsuitable as baselines since they require a set of reference images for fine-tuning, which would be obtained with ROPA-generated RGB images, but this would create a circular dependency. Consequently, we do not provide baselines for the depth augmentation component of our method.

### B. Simulation Results

Table I shows the success rates of different methods in simulation, with example synthesized images shown in Figure 5. ROPA outperforms the baselines on all 5 tasks, achieving substantial improvements across both RGB and RGB-D modalities. Baseline methods, particularly VISTA,




Fig. 5: **Synthesized images in simulation.** We present synthesized images from the Coordinated Lift Ball (CLB) task across two timesteps. The blue bordered images show the original RGB and RGB-D images, while the red bordered images represent the generated target image RGB and RGB-D images conditioned on the corresponding skeleton pose shown below.

| Method                   | CLB         | CLT         | CPB         | BSR         | CPID        |
|--------------------------|-------------|-------------|-------------|-------------|-------------|
| RGB ACT (w/o augment.)   | 41.3        | 10.7        | 43.3        | 21.3        | 17.3        |
| RGB Fine-tuned VISTA     | 52.7        | 1.6         | 54.3        | 10.3        | 15.7        |
| RGB ACT (more data)      | 48.0        | 12.0        | 42.7        | 17.3        | 10.3        |
| <b>RGB ROPA (ours)</b>   | <b>68.0</b> | <b>30.7</b> | <b>62.7</b> | <b>24.3</b> | <b>30.7</b> |
| RGB-D ACT (w/o augment.) | 56.7        | 14.3        | 49.7        | 26.7        | 17.6        |
| RGB-D ACT (more data)    | 33.3        | 9.7         | 47.0        | 20.3        | 16.3        |
| <b>RGB-D ROPA (ours)</b> | <b>72.0</b> | <b>35.0</b> | <b>61.3</b> | <b>29.0</b> | <b>35.0</b> |

TABLE I: **Simulation Results.** Success rates of ROPA compared with four baselines, averaged over three random seeds. All methods are evaluated using the ACT policy. The top half of the table shows training with RGB images, while the bottom half shows training with RGB images and depth maps. Task names and baseline details are given in Section V-A.

struggle with precision-demanding tasks that require low error tolerance, such as grasping trays or manipulating ropes. ROPA’s ability to generate out-of-distribution states enables more robust performance in these failure-prone scenarios, with better error recovery capabilities. Another important finding is that simply increasing training data does not guarantee performance improvements. This highlights the importance of our data augmentation approach, which shows consistent improvements over naive data scaling strategies. We also find that incorporating depth into ACT leads to small performance boost compared to RGB only. By augmenting both the depth and RGB modalities, we see slight but measurable performance improvements across all tasks. Table II also shows that ROPA maintains strong performance in multi-camera environments.

### C. Generalization Experiments

We evaluate the diffusion model’s ability to generalize to unseen tasks and objects. The model is trained on five PerAct2 tasks: Pick Up Notebook, Sweep Dust Pan, Coordinated Push Box, Pick Up Plate, and Coordinated Lift Tray, then tested on our CLB task using three approaches: (1) **Zero-Shot** direct image synthesis

| Method                         | CLB         |
|--------------------------------|-------------|
| One Camera ACT (w/o augment.)  | 41.3        |
| <b>One Camera ROPA (ours)</b>  | <b>68.0</b> |
| Four Camera ACT (w/o augment.) | 72.0        |
| <b>Four Camera ROPA (ours)</b> | <b>80.0</b> |

TABLE II: **Multi-Camera Results.** Comparison of single vs. four-camera setups on the Coordinated Lift Ball (CLB) task. Four-camera setup uses tiled observations from four viewpoints. Success rates averaged over three seeds.

| Method                 | CLB  | Method                        | CLB         |
|------------------------|------|-------------------------------|-------------|
| Zero-Shot              | 50.7 | ROPA with Joint Position Only | 46.7        |
| Few-Shot (10 demos)    | 58.3 | ROPA with White Skeleton      | 32.0        |
| Fine-tuned (100 demos) | 68.0 | ROPA with OpenPose Skeleton   | 62.6        |
|                        |      | <b>ROPA (ours)</b>            | <b>68.0</b> |

TABLE III: **Generalization Experiments.** Simulation results averaged over 3 seeds.

TABLE IV: **Ablation Experiments.** Simulation results averaged over 3 seeds.

without fine-tuning, (2) **Few-Shot** fine-tuned with 10 CLB demonstrations, and (3) **Fine-tuned** trained on 100 CLB demonstrations directly. Table III shows that direct fine-tuning on the target dataset yields the best performance. However, zero-shot and few-shot approaches achieve reasonable results when target data is limited, though image artifacts increase with fewer training demonstrations.

### D. Ablation Studies

We conduct ablation studies in simulation to address two key questions: (1) How does using Skeleton Pose as input compare to raw joint positions? and (2) How does varying the Skeleton Pose influence image generation quality?

**Skeleton Pose vs. Raw Joint Positions:** We modify ControlNet to replace Skeleton Pose input with direct 14-dimensional joint position vectors, processed via MLP layers and convolutions to map positions to spatial features. Table IV shows ROPA performs best with Skeleton Pose compared to *ROPA with Joint Position Only*, suggesting skeleton representation provides richer spatial context and kinematic constraints that guide diffusion more effectively than raw positional data.

**Pose Variation Impact:** We evaluate different skeleton pose representations in Figure 3. *ROPA with White Skeleton* uses an all-white skeleton, causing significant performance degradation by eliminating visual contrast and anatomical structure that guide diffusion, providing minimal conditioning signal. *ROPA with OpenPose Skeleton* achieves strong performance through clear joint connectivity and structure. However, combining this with GENIMA yields the strongest results because GENIMA’s sphere textures provide better visual cues for learning joint orientations and rotational states.

## VI. REAL-WORLD EXPERIMENTS

### A. Real-World Experiment Setup

We use two UR5 robots with Robotiq 2F-85 grippers and an Intel RealSense D415 RGB-D camera. An experienced roboticist collected 30 demonstrations per task using GELLO [63] (see Figure 6). We use the following tasks:

- **Lift Ball:** Lift a ball above a region.




Fig. 6: **Real-world setup.** The system features dual UR5 robotic arms in a bimanual configuration, each equipped with a Robotiq 2F-85 gripper. An Intel RealSense D415 RGB-D camera provides visual perception.

- **Push Block:** Push a block past the front of a workspace.
- **Lift Drawer:** Lift a drawer above a region.

| Method                   | Lift Ball      | Lift Drawer    | Push Block     |
|--------------------------|----------------|----------------|----------------|
| RGB Fine-tuned VISTA     | 11 / 20        | 1 / 20         | 18 / 20        |
| RGB ACT (w/o augment.)   | 13 / 20        | 4 / 20         | 14 / 20        |
| RGB ROPA (ours)          | <b>17 / 20</b> | <b>10 / 20</b> | <b>20 / 20</b> |
| RGB-D ACT (w/o augment.) | 15 / 20        | 5 / 20         | 17 / 20        |
| RGB-D ROPA (ours)        | <b>19 / 20</b> | <b>13 / 20</b> | <b>20 / 20</b> |

TABLE V: **Real-world Experiments.** Results comparing ROPA with baselines, with 20 trials per method and task combination.

### B. Real-World Results

Real-world results are presented in Table V, with synthesized images in Figure 7. ROPA demonstrates substantially improved robustness across all tasks compared to baselines.

**Lift Ball:** Baseline methods frequently freeze upon ball contact, with excessive gripping force sometimes triggering safety limits. Even when successful, baselines often achieve imperfect arm alignment, resulting in loose grasps prone to dropping. ROPA achieves more reliable and stable grasping.

**Lift Drawer:** Fine-tuned VISTA struggles with precise manipulation, failing to grasp drawers properly or close grippers entirely. This highlights difficulties that this baseline has with fine motor control requirements.

**Push Block:** ROPA demonstrates strong performance when blocks are positioned at greater distances or challenging orientations. We observe that the baseline methods frequently get stuck or fail to initiate contact. This shows how ROPA improves robustness in manipulation tasks requiring precise control and recovery from suboptimal initial conditions.

### VII. LIMITATIONS

Like all methods, ROPA has limitations. The reliance on robot poses for image generation necessitates some observability of the robotic system, potentially reducing effectiveness in scenarios with significant occlusions. Additionally, ROPA requires camera calibration for skeleton rendering. Like other image generation frameworks, the diffusion model may produce artifacts or poorly synthesized objects, though this

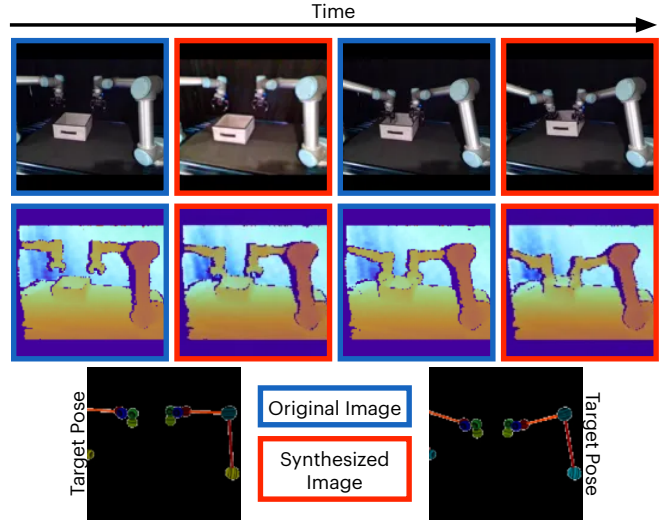


Fig. 7: **Synthesized images in the real world.** We present synthesized images from the `Lift Drawer` task across two timesteps. The blue bordered images show the original RGB and RGB-D images, while the red bordered images represent the generated target image RGB and RGB-D images conditioned on the corresponding skeleton pose shown below.

could be addressed by incorporating additional conditioning modalities such as detected edges or segmentation masks.

### VIII. CONCLUSION

We present ROPA, a novel data augmentation approach for bimanual robotic manipulation in imitation learning. Through evaluation in both simulation and real-world settings, we demonstrate consistent improvements over baselines. Our work highlights the potential of pre-trained diffusion models in robotics, mirroring their revolutionary impact on image generation. We hope this research inspires further development of robust data augmentation techniques for robotic manipulation.

### IX. ACKNOWLEDGMENTS

Huge thanks to SLURM Lab for discussions and helpful writing feedback. The authors used ChatGPT (OpenAI) for grammar refinement and language polishing.

### REFERENCES

- [1] F. Krebs and T. Asfour, "A Bimanual Manipulation Taxonomy," in *IEEE Robotics and Automation Letters (RA-L)*, 2022.
- [2] S. Stepputtis, M. Bandari, S. Schaal, and H. Ben Amor, "A System for Imitation Learning of Contact-Rich Bimanual Manipulation Policies," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022.
- [3] F. Xie, A. Chowdhury, M. C. De Paolis Kaluza, L. Zhao, L. L. Wong, and R. Yu, "Deep Imitation Learning for Bimanual Robotic Manipulation," in *Neural Information Processing Systems (NeurIPS)*, 2020.
- [4] Y. Avigal, L. Berscheid, T. Asfour, T. Kröger, and K. Goldberg, "SpeedFolding: Learning Efficient Bimanual Folding of Garments," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022.
- [5] J. Grannen, Y. Wu, B. Vu, and D. Sadigh, "Stabilize to Act: Learning to Coordinate for Bimanual Manipulation," in *Conference on Robot Learning (CoRL)*, 2023.
- [6] J. Maitin-Shepard, M. Cusumano-Towner, J. Lei, and P. Abbeel, "Cloth Grasp Point Detection Based on Multiple-View Geometric Cues with Application to Robotic Towel Folding," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2010.
- [7] T. Weng, S. Bajracharya, Y. Wang, K. Agrawal, and D. Held, "FabricFlowNet: Bimanual Cloth Manipulation with a Flow-based Policy," in *Conference on Robot Learning (CoRL)*, 2021.
- [8] I.-C. A. Liu, S. He, D. Seita, and G. Sukhatme, "VoxAct-B: Voxel-Based Acting and Stabilizing Policy for Bimanual Manipulation," in *Conference on Robot Learning (CoRL)*, 2024.

- [9] T. Lin, Z.-H. Yin, H. Qi, P. Abbeel, and J. Malik, "Twisting Lids Off with Two Hands," in *Conference on Robot Learning (CoRL)*, 2024.
- [10] T. Osa, J. Pajarinen, G. Neumann, J. A. Bagnell, P. Abbeel, and J. Peters, "An Algorithmic Perspective on Imitation Learning," *F&T in Robotics*, 2018.
- [11] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, and et al., " $\pi_0$ : A Vision-Language-Action Flow Model for General Robot Control," in *Robotics: Science and Systems (RSS)*, 2025.
- [12] P. Intelligence, K. Black, N. Brown, J. Darphinian, K. Dhabalia, D. Driess, A. Esmail, M. Equi, C. Finn, and et al., " $\pi_{0.5}$ : a Vision-Language-Action Model with Open-World Generalization," in *Conference on Robot Learning (CoRL)*, 2025.
- [13] S. Liu, L. Wu, B. Li, H. Tan, H. Chen, Z. Wang, K. Xu, H. Su, and J. Zhu, "RDT-1B: a Diffusion Foundation Model for Bimanual Manipulation," in *International Conference on Learning Representations (ICLR)*, 2025.
- [14] S. Tian, B. Wulfe, K. Sargent, K. Liu, S. Zakharov, V. Guizilini, and J. Wu, "View-Invariant Policy Learning via Zero-Shot Novel View Synthesis," in *Conference on Robot Learning (CoRL)*, 2024.
- [15] I.-C. A. Liu, J. Chen, G. Sukhatme, and D. Seita, "D-CODA: Diffusion for Coordinated Dual-Arm Data Augmentation," in *Conference on Robot Learning (CoRL)*, 2025.
- [16] X. Zhang, M. Chang, P. Kumar, and S. Gupta, "Diffusion Meets DAgger: Supercharging Eye-in-hand Imitation Learning," in *Robotics: Science and Systems (RSS)*, 2024.
- [17] L. Ma, X. Jia, Q. Sun, B. Schiele, T. Tuytelaars, and L. V. Gool, "Pose Guided Person Image Generation," in *Neural Information Processing Systems (NeurIPS)*, 2017.
- [18] M. Grotz, M. Shridhar, T. Asfour, and D. Fox, "PerAct2: Benchmarking and Learning for Robotic Bimanual Manipulation Tasks," *arXiv:2407.00278*, 2024.
- [19] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware," in *Robotics: Science and Systems (RSS)*, 2023.
- [20] L. Ureche and A. Billard, "Constraints extraction from asymmetrical bimanual tasks and their use in coordinated behavior," *Robotics and Autonomous Systems*, vol. 103, pp. 222–235, 2018.
- [21] Y. Chen, Y. Yang, T. Wu, S. Wang, X. Feng, J. Jiang, S. M. McAleer, H. Dong, Z. Lu, and S.-C. Zhu, "Towards Human-Level Bimanual Dexterous Manipulation with Reinforcement Learning," in *Neural Information Processing Systems (NeurIPS)*, 2022.
- [22] K. Zakka, P. Wu, L. Smith, N. Gileadi, T. Howell, X. B. Peng, S. Singh, Y. Tassa, P. Florence, A. Zeng, and P. Abbeel, "Robopianist: Dexterous piano playing with deep reinforcement learning," in *Conference on Robot Learning (CoRL)*, 2023.
- [23] R. Chitnis, S. Tulsiani, S. Gupta, and A. Gupta, "Efficient bimanual manipulation using learned task schemas," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [24] Y. Li, C. Pan, H. Xu, X. Wang, and Y. Wu, "Efficient Bimanual Handover and Rearrangement via Symmetry-Aware Actor-Critic Learning," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2023.
- [25] N. Chernyadev, N. Backshall, X. Ma, Y. Lu, Y. Seo, and S. James, "BiGym: A Demo-Driven Mobile Bi-Manual Manipulation Benchmark," in *Conference on Robot Learning (CoRL)*, 2024.
- [26] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, "Diffusion Policy: Visuomotor Policy Learning via Action Diffusion," in *Robotics: Science and Systems (RSS)*, 2023.
- [27] Z. Fu, T. Z. Zhao, and C. Finn, "Mobile ALOHA: Learning Bimanual Mobile Manipulation with Low-Cost Whole-Body Teleoperation," in *Conference on Robot Learning (CoRL)*, 2024.
- [28] T. Z. Zhao, J. Tompson, D. Driess, P. Florence, K. Ghasemipour, C. Finn, and A. Wahid, "ALOHA Unleashed: A Simple Recipe for Robot Dexterity," in *Conference on Robot Learning (CoRL)*, 2024.
- [29] R. Ding, Y. Qin, J. Zhu, C. Jia, S. Yang, R. Yang, X. Qi, and X. Wang, "Bunny-visionpro: Real-time bimanual dexterous teleoperation for imitation learning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025.
- [30] O. X.-E. Collaboration, "Open X-Embodiment: Robotic Learning Datasets and RT-X Models," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- [31] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Neural Information Processing Systems (NeurIPS)*, 2012.
- [32] S. Ross, G. J. Gordon, and J. A. Bagnell, "A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning," in *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011.
- [33] M. Laskey, J. Lee, R. Fox, A. D. Dragan, and K. Goldberg, "DART: Noise Injection for Robust Imitation Learning," in *Conference on Robot Learning (CoRL)*, 2017.
- [34] Z. Chen, Z. Mandi, H. Bharadhwaj, M. Sharma, S. Song, A. Gupta, and V. Kumar, "Semantically Controllable Augmentations for Generalizable Robot Learning," in *International Journal of Robotics Research (IJRR)*, 2024.
- [35] C. Yuan, S. Joshi, S. Zhu, H. Su, H. Zhao, and Y. Gao, "RoboEngine: Plug-and-Play Robot Data Augmentation with Semantic Robot Segmentation and Background Generation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025.
- [36] H. Bharadhwaj, J. Vakil, M. Sharma, A. Gupta, S. Tulsiani, and V. Kumar, "RoboAgent: Generalization and Efficiency in Robot Manipulation via Semantic Augmentations and Action Chunking," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- [37] T. Yu, T. Xiao, A. Stone, J. Tompson, A. Brohan, S. Wang, J. Singh, C. Tan, D. M. J. Peralta, B. Ichter, K. Hausman, and F. Xia, "Scaling Robot Learning with Semantically Imagined Experience," in *Robotics: Science and Systems (RSS)*, 2023.
- [38] S. Yang, W. Yu, J. Zeng, J. Lv, K. Ren, C. Lu, D. Lin, and J. Pang, "Novel demonstration generation with gaussian splatting enables robust one-shot manipulation," in *Robotics: Science and Systems (RSS)*, 2025.
- [39] P. Mitrano and D. Berenson, "Data Augmentation for Manipulation," in *Robotics: Science and Systems (RSS)*, 2022.
- [40] L. Ke, Y. Zhang, A. Deshpande, S. Srinivasa, and A. Gupta, "CCIL: Continuity-based Data Augmentation for Corrective Imitation Learning," in *International Conference on Learning Representations (ICLR)*, 2024.
- [41] A. Zhou, M. J. Kim, L. Wang, P. Florence, and C. Finn, "NeRF in the Palm of Your Hand: Corrective Augmentation for Robotics via Novel-View Synthesis," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [42] A. Mandlekar, S. Nasiriany, B. Wen, I. Akinola, Y. Narang, L. Fan, Y. Zhu, and D. Fox, "MimicGen: A Data Generation System for Scalable Robot Learning using Human Demonstrations," in *Conference on Robot Learning (CoRL)*, 2023.
- [43] Z. Jiang, Y. Xie, K. Lin, Z. Xu, W. Wan, A. Mandlekar, L. Fan, and Y. Zhu, "DexMimicGen: Automated Data Generation for Bimanual Dexterous Manipulation via Imitation Learning," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2025.
- [44] L. Y. Chen, C. Xu, K. Dharmarajan, M. Z. Irshad, R. Cheng, K. Keutzer, M. Tomizuka, Q. Vuong, and K. Goldberg, "RoVi-Aug: Robot and Viewpoint Augmentation for Cross-Embodiment Robot Learning," in *Conference on Robot Learning (CoRL)*, 2024.
- [45] J. Li, W. Wang, Y. Peng, C. Shen, Y. Zhu, and Z. Xu, "Visual robotic manipulation with depth-aware pretraining," in *IEEE International Conference on Robotics and Biomimetics (ROBIO)*, 2024.
- [46] A. K. Bhunia, S. Khan, H. Cholakkal, R. M. Anwer, J. Laaksonen, M. Shah, and F. S. Khan, "Person Image Synthesis via Denoising Diffusion Model," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [47] Y. Lu, M. Zhang, A. J. Ma, X. Xie, and J.-H. Lai, "Coarse-to-Fine Latent Diffusion for Pose-Guided Person Image Synthesis," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [48] M. Shridhar, Y. L. Lo, and S. James, "Generative Image as Action Models," in *Conference on Robot Learning (CoRL)*, 2024.
- [49] L. Zhang, A. Rao, and M. Agrawala, "Adding Conditional Control to Text-to-Image Diffusion Models," in *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023.
- [50] X. Liu, Y. Zhou, F. Weigend, S. Sonawani, S. Ikemoto, and H. B. Amor, "Diff-control: A stateful diffusion-based policy for imitation learning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024.
- [51] R. Liu, A. Canberk, S. Song, and C. Vondrick, "Differentiable robot rendering," in *Conference on Robot Learning (CoRL)*, 2024.
- [52] Y. Labbé, J. Carpentier, M. Aubry, and J. Sivic, "Single-view robot pose and joint angle estimation via render & compare," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [53] Z. Cao, G. Hidalgo, T. Simon, S.-E. Wei, and Y. Sheikh, "Openpose: Realtime multi-person 2d pose estimation using part affinity fields," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [54] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [55] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, "Learning Transferable Visual Models From Natural Language Supervision," in *International Conference on Machine Learning (ICML)*, 2021.
- [56] J. Ho, A. Jain, and P. Abbeel, "Denoising Diffusion Probabilistic Models," in *Neural Information Processing Systems (NeurIPS)*, 2020.
- [57] J. Song, C. Meng, and S. Ermon, "Denoising Diffusion Implicit Models," in *International Conference on Learning Representations (ICLR)*, 2021.
- [58] M. Matl, "PyRender," 2018. [Online]. Available: <https://github.com/mmatl/pyrender/>
- [59] T. Zhang, P. Ge, Y. Zou, and Y. He, "Robot collision detection without external sensors based on time-series analysis," *Journal of Dynamic Systems, Measurement, and Control*, vol. 143, no. 4, 11 2020.
- [60] S. James, Z. Ma, D. Rovick Arroj, and A. J. Davison, "Rlbench: The robot learning benchmark & learning environment," in *IEEE Robotics and Automation Letters (RA-L)*, 2020.
- [61] K. Sargent, Z. Li, T. Shah, C. Herrmann, H.-X. Yu, Y. Zhang, E. R. Chan, D. Lagun, L. Fei-Fei, D. Sun, and J. Wu, "ZeroNVs: Zero-Shot 360-Degree View Synthesis from a Single Image," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [62] L. Yang, B. Kang, Z. Huang, X. Xu, J. Feng, and H. Zhao, "Depth anything: Unleashing the power of large-scale unlabeled data," in *CVPR*, 2024.
- [63] P. Wu, Y. Shentu, Z. Yi, X. Lin, and P. Abbeel, "GELLO: A General, Low-Cost, and Intuitive Teleoperation Framework for Robot Manipulators," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024.