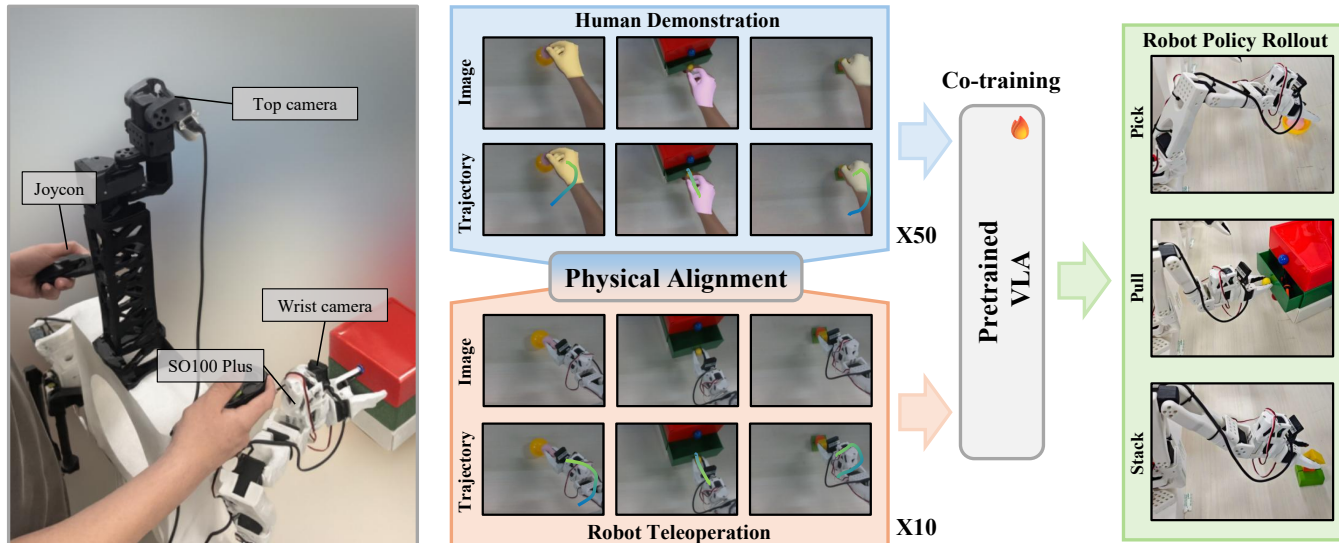


EasyMimic: A Low-Cost Framework for Robot Imitation Learning from Human Videos

Tao Zhang^{1*}, Song Xia^{1*}, Ye Wang^{1*†}, Qin Jin^{1†}



(a) Low-cost hardware: the entire system can be built for less than \$300. (b) Low-cost pipeline: achieves a high success rate in real scenarios with less than 5 minutes of real robot teleoperation data combined with human demonstration data.

Fig. 1: Overview of the EasyMimic framework. The framework learns robotic manipulation from human videos captured with low-cost hardware. To bridge the embodiment gap, it aligns the action and visual spaces via physical alignment. A VLA model is then fine-tuned on the combined data for rapid adaptation to new tasks.

Abstract—Robot imitation learning is often hindered by the high cost of collecting large-scale, real-world data. This challenge is especially significant for low-cost robots designed for home use, as they must be both user-friendly and affordable. To address this, we propose the EasyMimic framework, a low-cost and replicable solution that enables robots to quickly learn manipulation policies from human video demonstrations captured with standard RGB cameras. Our method first extracts 3D hand trajectories from the videos. An action alignment module then maps these trajectories to the gripper control space of a low-cost robot. To bridge the human-to-robot domain gap, we introduce a simple and user-friendly hand visual augmentation strategy. We then use a co-training method, fine-tuning a model on both the processed human data and a small amount of robot data, enabling rapid adaptation to new tasks. Experiments on the low-cost LeRobot platform demonstrate that EasyMimic achieves high performance across various manipulation tasks. It significantly reduces the reliance on expensive robot data collection, offering a practical path for bringing intelligent robots into homes. Project website: <https://zt375356.github.io/EasyMimic-Project/>.

I. INTRODUCTION

In recent years, bringing robots into homes to assist with daily life has been a long-standing vision in embodied intelligence. Equipping robots to master diverse household

tasks requires learning from vast amounts of demonstration data. Traditional data collection methods, such as expert teleoperation [1]–[3], provide high-quality data but are hindered by expensive equipment and complex operation, limiting the widespread adoption of low-cost robots in home environments. This data bottleneck remains a core challenge in the field. While recent Vision-Language-Action (VLA) models [4]–[8] exhibit powerful zero-shot capabilities, unlocking their full potential for precise and reliable manipulation in diverse home settings often requires in-domain fine-tuning. This highlights the urgent need for scalable, low-cost data acquisition that can complement these large models.

A promising solution is to leverage ordinary human videos for imitation learning. People can easily record manipulation demonstrations using devices like mobile phones, providing a nearly zero-cost and massive-scale data source for robot learning [9]–[11]. Our core motivation is therefore to explore a more convenient and lower-barrier method that allows non-expert users to provide effective demonstration using consumer devices, simplifying and accelerating robot learning. However, putting this idea into practice requires addressing two fundamental gaps between humans and robots: 1) **Visual Appearance Gap**: Human hands differ completely from robot grippers in texture and shape. Models trained directly on human videos tend to overfit these human-specific fea-

*Equal contribution, ‡Project lead, †Corresponding author.

¹AIM3 Lab, Renmin University of China.

tures, failing to recognize and act correctly when faced with a real robot gripper. 2) **Action Space Gap**: The kinematic structures, joint limits, and workspaces of human arms and robot manipulators are vastly different. Directly mimicking human motion trajectories can result in unnatural or even unsafe robot movements. While existing work attempts to solve these problems, they often rely on computationally expensive image generation techniques [12], [13] or require costly hand-tracking hardware [9], [14], [15], undermining the low-cost and convenience goals for home robots.

To address the challenges of high data costs and the human-robot embodiment gap, we introduce the **EasyMimic** framework, a simple and efficient imitation learning framework for low-cost robots using only consumer-grade devices. The framework systematically bridges the human-robot gap across two dimensions: action and vision. On the action level, we utilize 3D hand pose estimation to extract key kinematic information and design a stable retargeting algorithm to map it into robot actions. On the visual level, we discard complex generative models and instead adopt a lightweight visual augmentation strategy. By rendering hand meshes with randomized colors, we compel the model to learn cross-embodiment general patterns. We then co-train the model on combination of this processed human data and a small amount of robot teleoperation data. Experiments on a low-cost LeRobot platform demonstrate significant improvements in task success rates.

The main contributions of this work are as follows: (i) We propose a complete, low-cost pipeline that enables robot policy training from human demonstrations captured with standard RGB cameras. (ii) We design an action alignment module that effectively retargets 3D human hand trajectories into executable robot actions. (iii) We employ a lightweight visual augmentation strategy based on hand color randomization to mitigate the visual gap. (iv) Systematic evaluations on a low-cost robot platform demonstrate that co-training on human data and limited robot data bridges the domain gap and significantly improves task performance.

II. RELATED WORK

A. Robot Data Collection

In robotics, teleoperation is a dominant paradigm that offers high safety and low latency, but its data collection efficiency and task coverage are limited [16]–[18]. To improve efficiency and scalability, some methods introduce more advanced data collection hardware or systems [2], [19]–[22]. However, these approaches typically rely on specialized hardware and skilled operators, incurring significant deployment and maintenance costs.

Hand-centric data collection presents another alternative. For example, instrumented gloves can directly record hand poses but suffer from complex calibration and sensor drift [23]. Head-mounted devices can capture egocentric views and hand information simultaneously, but their high cost and reliance on SLAM for spatiotemporal alignment limit their broader application [9], [14], [24], [25]. To address these limitations, our method uses only a single RGB camera

to efficiently acquire hand data, striking a better balance between cost and efficiency while achieving high scalability through the rapid iteration of commodity hardware [26].

B. Learning from Human Videos

To address the challenges of cross-embodiment learning, prior research can be broadly categorized into three directions. The first line of work reduces the appearance gap by synthesizing or editing data to make demonstrations more robot-like, though such pipelines often involve complex rendering and high computational costs [13], [27], [28]. The second line focuses on distilling embodiment-agnostic, high-level information from human videos, such as reward functions or sub-task plans, which requires robust object perception and temporal abstraction capabilities [29]–[32]. The third line builds unified representations or policies to align human and robot demonstrations at the feature and control levels, thereby mitigating discrepancies in appearance and action jointly [9], [33]–[35]. Our work follows this third direction. We employ lightweight rendering to standardize hand appearance and narrow the appearance gap, while simultaneously aligning action spaces with inverse kinematics (IK) to map hand trajectories into executable robot actions, enabling low-cost and effective cross-embodiment transfer.

C. Vision Language Action Model

Driven by the rapid progress of Multimodal Large Language Models (MLLMs), Vision-Language-Action (VLA) models have emerged as a research hotspot in embodied intelligence. By co-training on internet-scale vision-language data and large-scale robot trajectories, these models integrate powerful language understanding and visual perception with robot action generation, achieving unprecedented generalization. Numerous studies have shown that these models exhibit strong potential in open-vocabulary understanding, zero-shot task execution, and long-horizon planning [4]–[6], [8], [36]–[41]. Nevertheless, adapting these general-purpose models to new robot embodiments for specific tasks at low cost remains a critical challenge. The core bottleneck lies in the prohibitive expense of collecting demonstration data for each new embodiment. Our work directly addresses this challenge by introducing a low-cost, scalable pipeline that leverages human video data as an alternative to expensive robot demonstrations.

III. METHOD

We present **EasyMimic**, a framework that enables low-cost robots to efficiently leverage human video demonstrations. The overall framework is illustrated in Figure 1. It consists of three core components: 1) low-cost collection of both human and robot demonstrations using consumer-grade hardware; 2) physical alignment across the action and visual domains to bridge the embodiment gap; and 3) a co-training strategy that effectively fuses human and robot data to train a unified policy model.

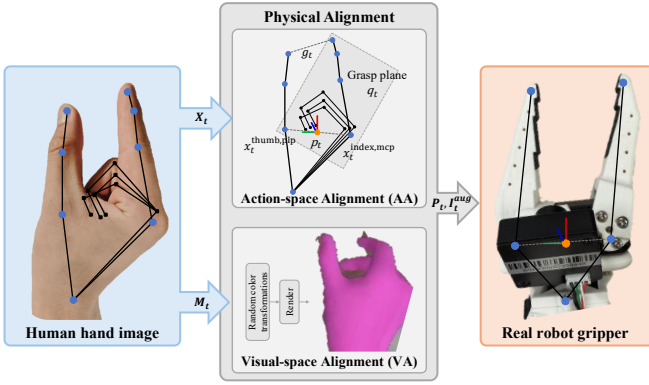


Fig. 2: Physical alignment process. Human hand keypoints and meshes are extracted from videos. Hand motion is retargeted to robot actions via the action space alignment module, while the hand mesh is augmented through the visual space alignment module to bridge the physical gap between humans and robots.

A. Data Collection Systems and Hardware Design

In designing the system, we adhere to the principles of low cost and accessibility. For the robot platform, we use the LeRobot SO100 manipulator [42], [43]. To address the limitations of its original five-degree-of-freedom (5-DoF) design in end-effector orientation control, such as singularities in certain workspace areas that prevent arbitrary poses, we added an extra joint, upgrading it to a six-degree-of-freedom (6-DoF) configuration. This extension significantly expands its effective workspace. We name the upgraded manipulator SO100-Plus. It has a payload capacity of up to 1 kg and a reach of 400 mm, making it well-suited for various tabletop manipulation tasks.

For data collection, we use a Nintendo Joy-Con for robot teleoperation and deploy two consumer-grade RGB cameras: one fixed in a first-person view (shared by human and robot data collection), and the other mounted on the robot’s wrist to capture close-up manipulation details. The entire system can be assembled for under \$300, far more cost-effective than traditional research platforms. To process human videos, we leverage the advanced HaMeR model [44] to extract 3D hand information from monocular inputs. We utilize this model to precisely reconstruct the hand morphology for each video frame, providing 21 hand keypoint coordinates X_t^C and 778 hand mesh vertex coordinates V_t^C in the camera coordinate frame \mathcal{C} . This information serves as the foundation for our subsequent alignment process, where keypoints are primarily used for action space mapping and the complete mesh is used for visual space alignment.

B. Physical Alignment

Human-to-robot transfer requires resolving discrepancies in both kinematics and appearance. As illustrated in Figure 2, our alignment module systematically addresses these challenges across action space and visual space.

1) *Action Space Alignment*: The primary challenge in cross-embodiment imitation is to accurately translate human

hand motion trajectories into executable action sequences for the robot’s end-effector. We need to construct a mapping function from human hand kinematics to robot kinematics. Given a sequence of human hand keypoints $\mathcal{X}_h = \{X_t\}_{t=1}^T$ extracted from a video, our goal is to generate a corresponding sequence of robot end-effector poses $\mathcal{P}_h = \{P_t\}_{t=1}^T = \{(p_t, q_t, g_t)\}_{t=1}^T$, where p_t , q_t , and g_t represent its 3D position, orientation (typically as a quaternion or Euler angles), and gripper state, respectively. This pose sequence constitutes the robot’s state representation.

Position Alignment: Prior works often use the wrist [10] or the midpoint of the fingertips [27] as an anchor point, but these can deviate from the true center of interaction during complex manipulations. For example, the wrist is too far from the object, while fingertips move relative to each other during fine manipulation, leading to an unstable anchor. Inspired by unified representation approaches [45], [46], we select the center of the thenar eminence—defined as the midpoint between the thumb’s proximal interphalangeal (PIP) joint and the index finger’s metacarpophalangeal (MCP) joint—as our retargeting anchor. This point remains relatively stable during grasping, approximating the palm’s center of mass regardless of finger movement. It therefore better represents the core region of hand-object interaction.

Orientation Alignment: We fit a plane through five keypoints (the four joints of the index finger and the thumb’s PIP joint). These five points collectively define the primary orientation of the palm. The plane’s normal vector defines the Z-axis, while the vector from the index finger’s MCP to its PIP joint defines the X-axis. A full 3D coordinate frame is constructed via the cross product, yielding a rotation matrix R_t , which is finally converted into the orientation representation q_t .

Gripper State Alignment: To achieve fine-grained grasping control, we calculate the Euclidean distance d_t between the thumb and index fingertips and normalize it to the range $[0, 1]$ to serve as the gripper state g_t :

$$g_t = \text{clip} \left(\frac{d_t - d_{\min}}{d_{\max} - d_{\min}}, 0, 1 \right) \quad (1)$$

where d_{\min} and d_{\max} are the distance thresholds corresponding to a fully closed and fully open gripper, respectively. This continuous gripper state representation, compared to binary open-close control, enables the robot to perform tasks requiring gentle or partial grasps.

After alignment, we use a pre-calibrated hand-eye transformation matrix T_C^R to convert the states from the camera frame \mathcal{P}_h^C to the robot’s base frame \mathcal{P}_h^R . Finally, inspired by chunk-based prediction [47], we define the action a_t at timestep t as the state at the next timestep, $P_{h,t+1}$, and form an action chunk $A_t = (a_t, \dots, a_{t+h-1})$ from h future actions.

2) *Visual Space Alignment*: The visual discrepancy is another major obstacle in human-to-robot imitation. Instead of employing computationally expensive generative models to translate image styles, we propose a lightweight visual

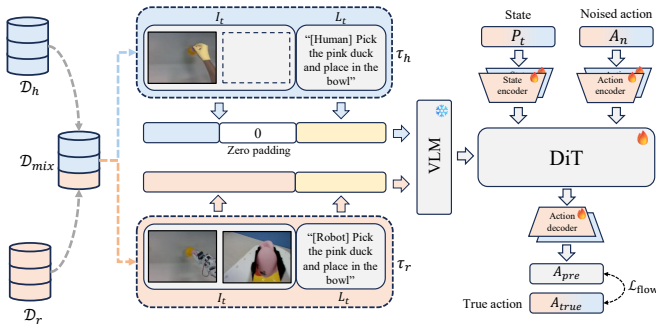


Fig. 3: Co-training strategy. Human demonstration data and robot teleoperation data are mixed during training. A shared DiT module learns a unified policy representation, while separate action encoders and decoders for each embodiment handle their specific data properties.

augmentation strategy that aligns with the philosophy of our low-cost framework. This strategy is based on the idea of domain randomization. Its core objective is to compel the model to ignore task-irrelevant superficial features, such as color and skin texture, and instead learn more fundamental geometric information, such as hand pose and shape.

Specifically, we reuse the extracted 3D hand mesh M_t and render it onto the original image I_t . During the rendering process, we apply random color transformations to the entire hand, generating an appearance-standardized augmented view I_t^{aug} . By using this augmented data during training, the model is passively exposed to diverse embodiment morphologies. This method, without adding any extra complex modules, guides the model at the data level to learn a cross-embodiment visual representation, thereby effectively enhancing its generalization to the robot’s morphology.

C. Training Strategy

To effectively leverage human and robot data, we adopt a co-training strategy. These two data sources have complementary advantages: human data is abundant and easy to acquire, containing rich high-level task semantics, while robot data, though scarce, provides precise low-level control signals consistent with its own kinematic properties. We combine our processed human demonstration dataset \mathcal{D}_h with a small real-world robot dataset \mathcal{D}_r to form a mixed dataset \mathcal{D}_{mix} . As shown in Figure 3, all data first passes through a shared frozen vision-language encoder, followed by a shared Diffusion Transformer (DiT) module for cross-embodiment policy learning.

In our implementation, to handle data imbalance, we use balanced sampling to ensure that each minibatch contains data from both human and robot sources. As the human videos lack the wrist-mounted camera view, we apply zero-padding to the corresponding image data to maintain consistent input dimensions.

Furthermore, while the core DiT module is shared, we design independent encoders and decoders for the state (P_t) and the noised action (A_n) for each embodiment (human and robot). This design allows the model to flexibly handle

differences in dimensionality, numerical ranges, and physical meanings between the two data types, leading to a more stable shared policy learning process.

IV. EXPERIMENTS

In this section, we first introduce the manipulation tasks and experimental setups. We then present a comparison with baseline methods and provide a detailed analysis to validate the effectiveness of our approach.

A. Experimental Setup

Hardware and Tasks. Our experimental platform uses a 6-DoF so100-plus robotic arm equipped with a two-finger gripper. The vision system includes two monocular RGB cameras: one fixed above the robot’s base for a top-down global view, and another mounted on the wrist for an end-effector-centric first-person view. We evaluate our method on four tabletop manipulation tasks.

- **Pick and Place (Pick):** The robot picks up a toy duck and places it into a bowl. Success is scored in two stages: grasping (0.5) and placing (0.5).
- **Pull and Push (Pull):** The robot must first pull open a drawer and then push it closed. The task is evaluated in two stages: pulling (0.5) and pushing (0.5).
- **Stacking (Stack):** The robot stacks a cube on top of a rectangular block, then stacks a triangular pyramid on top of the cube. Each successful stack is awarded 0.5 points.
- **Language Conditioned (LC):** The robot executes natural language instructions that specify both the target object (e.g., by color) and the placement goal (e.g., target container or location). For example: “pick up the pink duck and place it into the yellow bowl”. Scoring follows the same scheme as the Pick task.

Model and Training. We use the pre-trained Gr00T N1.5-3B [8] as the foundation VLA model. The policy network is configured to output absolute actions, which consist of 6-DoF end-effector poses and gripper states. All models are trained for 5,000 gradient steps on a single NVIDIA RTX 4090 GPU using the AdamW optimizer with a learning rate of 1×10^{-4} and a batch size of 32.

Data Collection. For each task, we collect 100 human video demonstrations and 20 robot teleoperation trajectories. Table I summarizes the data collection statistics. The human data collection rate is substantially higher than that of robot teleoperation, highlighting the efficiency and utility of leveraging human data.

Baselines. We validate the effectiveness of EasyMimic, which combines physical alignment with a co-training strategy, by comparing it against the following baselines:

- **Robot-Only (10 traj):** The standard imitation learning baseline, trained solely on 10 robot teleoperation trajectories.
- **Robot-Only (20 traj):** Trained using 20 robot teleoperation trajectories.

Task	H			R		
	#	min	#/min	#	min	#/min
Pick	100	8	12.5	20	10	2.0
Pull	100	10	10.0	20	15	1.3
Stack	100	15	6.7	20	40	0.5
LC	100	12	8.3	20	12	1.7

TABLE I: Data collection overview. The table reports the number of demonstrations (#), total collection time (min), and average collection rate (#/min) for both human (H) and robot (R) data across different tasks.

- **Pretrain-Finetune:** The model is first pre-trained on processed human video data and then fine-tuned on robot data.

Evaluation Metrics. We evaluate each task using a stage-based scoring scheme with a maximum score of 1.0. Each task is executed 10 times, and the average score is reported. Performance is assessed with respect to key execution stages, including object localization, grasping, transportation or operation, and final release or reset.

B. Main Results

1) *Comparison of Training Strategies:* Table II compares the performance of different training strategies. The results show that training with only a small amount of robot data (10 or 20 trajectories) yields limited performance, achieving average success rates of only 0.26 and 0.51, respectively. This indicates that even with a pre-trained VLA, training on a small amount of robot data is insufficient to achieve a satisfactory level of task completion.

In contrast, incorporating human data, either through pre-training or co-training, leads to a substantial performance improvement. Our physical alignment approach combined with the Pretrain-Finetune method achieves an average score of 0.75, significantly outperforming the Robot-Only baselines and demonstrating the value of leveraging human demonstrations. However, this two-stage training process is approximately 1.6 times slower than the co-training approach due to its sequential pre-training and fine-tuning phases. Our **EasyMimic** framework achieves the best performance across all tasks, with an average score of 0.88, surpassing the Pretrain-Finetune method by 0.13 points and the Robot-only (10 trajectories) baseline by 0.62 points. These results highlight that co-training effectively integrates the strengths of both human and robot data, achieving robust task capabilities with minimal robot data while maintaining efficiency.

2) *Language Condition:* The experimental setup of this task includes four target objects: a pink duck, a green duck, a yellow bowl, and a wooden block. During data collection, we systematically acquire demonstrations covering four distinct combinations: placing the pink duck into the yellow bowl, placing the pink duck onto the block, placing the green duck into the yellow bowl, and placing the green duck onto the block. At inference time, the model executes the specified task following natural language instructions. As shown in Table II, our EasyMimic framework significantly improves

Strategy	Pick	Pull	Stack	LC	Avg.
Robot-Only (10 traj.)	0.30	0.30	0.15	0.30	0.26
Robot-Only (20 traj.)	0.60	0.70	0.35	0.40	0.51
Pretrain-Finetune	0.80	0.90	0.50	0.80	0.75
EasyMimic	1.00	0.90	0.70	0.90	0.88

TABLE II: Performance evaluation of different training strategies. Results show the average scores across different tasks.

performance on language-conditioned tasks. Training on 20 robot trajectories alone results in a low score of 0.40. In contrast, EasyMimic, by incorporating human data, boosts the score to 0.90, demonstrating the substantial value of leveraging human demonstrations for complex, language-conditioned manipulation.

C. Further Analysis

1) *Effect of Data Scale:* Figure 4 analyzes the scaling effects of both human and robot data on model performance.

As shown in Figure 4a increasing the amount of human demonstrations consistently improves performance across all tasks. However, we observe diminishing returns beyond 50 demonstrations, suggesting that while human data provides rich and diverse task priors, there exists an optimal scale beyond which additional data yields limited gains.

Figure 4b shows the effect of scaling robot data. Performance increases notably when the number of trajectories grows from 5 to 10 but quickly saturates thereafter, indicating that only a small amount of robot data is needed for effective domain adaptation when complemented with abundant human demonstrations.

The experiments confirm that with sufficient human data (e.g., 50 demonstrations), high performance can be achieved using only a small amount of robot data (e.g., 10-20 trajectories), significantly reducing reliance on expensive robot data collection while maintaining robust task execution.

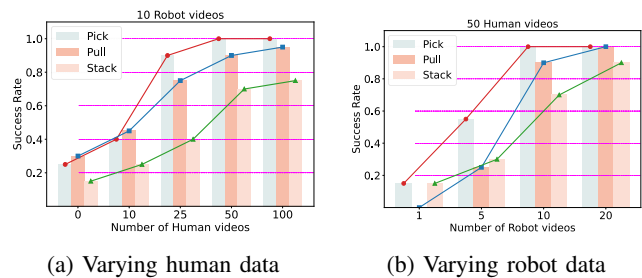


Fig. 4: Effect of Dataset Size. (a) Varying human data with fixed robot data (10 trajectories). (b) Varying robot data with fixed human data (50 videos).

2) *Effect of Alignment:* Action Alignment (AA) is a core component for translating human motions into executable robot actions. To quantify its impact, we compare our full EasyMimic pipeline against a baseline where human data is used without proper AA. As shown in Table III, removing

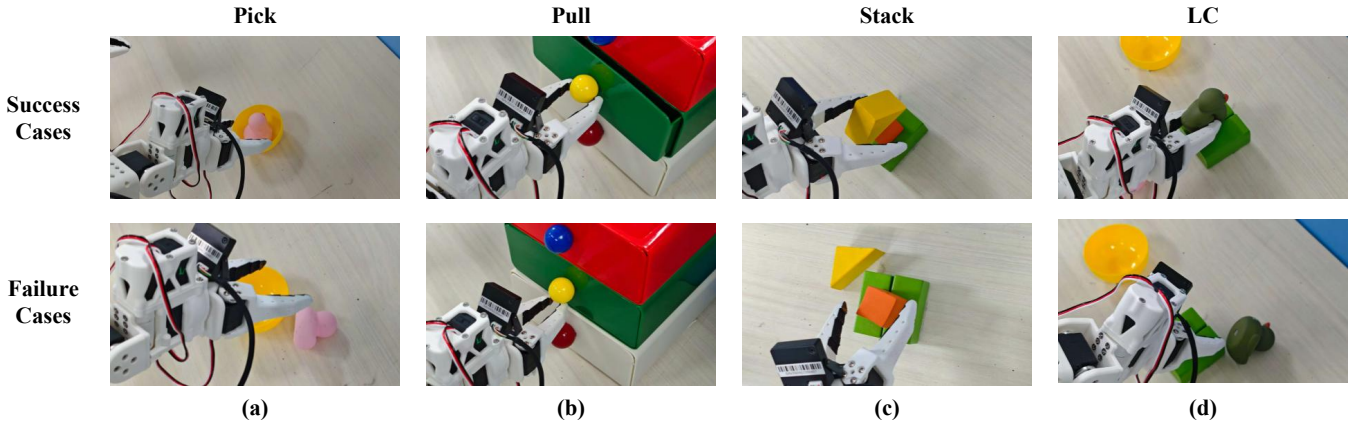


Fig. 5: Case analysis of failure modes across different tasks. (a) Premature gripper release during pick and place. (b) Imprecise handle grasping in drawer manipulation. (c) Collision-induced object falling during stacking. (d) Unstable placement leading to the object falling.




Fig. 6: Visualization of the VA module. (a) Original human hand image. (b) Partial masking of the thumb and index finger (VA-Partial). (c) Full masking of the hand region (VA-Full).

AA leads to a significant performance drop, with an average decrease of 0.27 points across tasks. The performance degradation is particularly pronounced in tasks such as pick-and-stack, which require precise hand rotation. This demonstrates that AA effectively reduces discrepancies between human and robot data, enabling the model to learn more generalizable task representations and strategies.

Visual Alignment (VA) addresses the visual domain gap between human and robot demonstrations. When VA is removed, performance drops sharply, with the average score decreasing by 0.47 points. This highlights that without visual augmentation, cross-embodiment knowledge transfer is severely hindered. Our lightweight VA module enhances the model’s understanding of shared task context, improving execution capabilities.

We further evaluate the design choice of full versus partial hand masking in the VA module. Figure 6(b) illustrates the partial masking strategy (VA-Partial), which only covers the thumb and index finger. As shown in Table IV, VA-Partial performs poorly (average score 0.27), indicating that full-hand visual augmentation is crucial for bridging the visual gap and achieving robust cross-embodiment transfer.

3) *Effect of Independent Action Heads*: Our EasyMimic framework employs independent action heads for human and robot data to account for their distinct characteristics. To validate this design, we compare it against a variant that employs a single shared action head. As shown in Table V,

Strategy	Pick	Pull	Stack	Avg.
EasyMimic	1.00	0.90	0.70	0.87
EasyMimic-AA	0.50	0.80	0.50	0.60
EasyMimic-VA	0.40	0.40	0.40	0.40

TABLE III: Ablation study on the effectiveness of Action Alignment and Visual Alignment strategies.

Strategy	Pick	Pull	Stack	Avg.
EasyMimic	1.00	0.90	0.70	0.87
EasyMimic (VA-Partial)	0.30	0.20	0.30	0.27

TABLE IV: Ablation study on the effectiveness of Visual Alignment strategies.

using independent heads improves the average performance by 0.40 points over the shared-head variant. This improvement suggests that a shared action head can confuse the model, as it cannot effectively differentiate between the two types of action data. In contrast, explicitly separating the heads prevents interference between human and robot data, allowing the model to leverage the unique properties of each source. This design enhances task execution capabilities by ensuring that both data types contribute effectively to learning a unified policy.

Strategy	Pick	Pull	Stack	Avg.
EasyMimic	1.00	0.90	0.70	0.87
EasyMimic (Shared)	0.60	0.40	0.40	0.47

TABLE V: Ablation study on the effect of independent versus shared action heads.

4) *Comparison with and without Pretraining*: Both pre-training and incorporating human demonstration data contribute to performance improvement, and their combination proves particularly effective. We evaluate the effect of completely randomly initializing the action expert of GR00T-N1.5-3B. As shown in Table VI, training on robot data alone

without pretraining yields an average score of only 0.15. Using a pre-trained VLA improves this to 0.25, demonstrating the benefit of large-scale pretraining. Even without pretraining, applying the EasyMimic framework raises the score from 0.15 to 0.53, indicating that EasyMimic is effective even for randomly initialized models. Combining EasyMimic with the pre-trained action expert achieves the highest score of 0.87, highlighting the complementary strengths of human demonstration data and pretraining. Notably, relying solely on pretraining without leveraging human videos remains substantially less effective, underscoring the critical role of human data for downstream task adaptation.

Initialization	Pick	Pull	Stack	Avg.
w/o Pretraining (Robot-Only)	0.20	0.10	0.10	0.15
w/o Pretraining (EasyMimic)	0.70	0.50	0.40	0.53
w/ Pretraining (Robot-Only)	0.30	0.30	0.15	0.25
w/ Pretraining (EasyMimic)	1.00	0.90	0.70	0.87

TABLE VI: Effect of pre-trained VLA initialization.

5) *Generalization to Unseen Objects*: We evaluate zero-shot generalization on the Pick and Place task by training the model exclusively on a pink duck and testing it on two unseen objects: a green duck (unseen color) and a pink cube (unseen geometry and affordance). As shown in Table VII, EasyMimic outperforms the Robot-Only baseline on both unseen objects, demonstrating effective zero-shot transfer across variations in color and shape. This indicates that leveraging human demonstration data enables the model to learn more generalizable task strategies beyond the objects seen during training.

Strategy	Green duck	Pink cube	Avg.
Robot-Only	0.50	0.20	0.35
EasyMimic	0.80	0.50	0.65

TABLE VII: Generalization performance on unseen objects for the Pick and Place task.

6) *Case Analysis*: To gain deeper insight into the model’s behavior, we analyze both successful executions and common failure cases across different tasks. Figure 5 illustrates representative examples of common failure modes: In pick and place tasks, the robot occasionally releases the gripper before reaching the target location, often due to an insufficient understanding of spatial relationships. For drawer manipulation, the model sometimes fails to accurately grasp the handle, particularly when the handle’s orientation differs from training examples. During stacking operations, the robot may inadvertently knock over existing objects while attempting to place new ones, indicating limitations in spatial awareness. The model occasionally places objects at unstable positions near edges, leading to subsequent falls and task failures. After training with the EasyMimic framework, the model can effectively learn these action logics from human demonstrations, leading to improved task execution and higher success rates.

V. CONCLUSIONS

In this paper, we introduce the EasyMimic framework, a low-cost, efficient, and replicable paradigm for robot learning in non-standardized settings. EasyMimic enables robots to acquire manipulation skills by leveraging readily available human videos. By incorporating carefully designed action and visual alignment modules, the framework effectively bridges the morphological and kinematic gaps between human demonstrations and robot execution.

Our experimental results demonstrate that co-training on easily accessible human video data, combined with a minimal amount of robot teleoperation data, allows EasyMimic to achieve superior performance across multiple tabletop manipulation tasks. It significantly outperforms baseline methods that rely solely on limited robot data, validating the effectiveness of our approach. This framework provides a practical solution to the data bottleneck problem in robot learning and has potential to lower the barrier for ordinary users to teach robots new skills.

REFERENCES

- [1] 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.
- [2] Y. Qin, W. Yang, B. Huang, K. Van Wyk, H. Su, X. Wang, Y.-W. Chao, and D. Fox, “Anyteleop: A general vision-based dexterous robot arm-hand teleoperation system,” *arXiv preprint arXiv:2307.04577*, 2023.
- [3] P. Wu, Y. Shentu, Z. Yi, X. Lin, and P. Abbeel, “Gello: A general, low-cost, and intuitive teleoperation framework for robot manipulators,” in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 12156–12163.
- [4] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, T. Jackson, S. Jesmonth, N. J. Joshi, R. Julian, D. Kalashnikov, Y. Kuang, I. Leal, K.-H. Lee, S. Levine, Y. Lu, U. Malla, D. Manjunath, I. Mordatch, O. Nachum, C. Parada, J. Peralta, E. Perez, K. Pertsch, J. Quiambao, K. Rao, M. Ryoo, G. Salazar, P. Sanketi, K. Sayed, J. Singh, S. Sontakke, A. Stone, C. Tan, H. Tran, V. Vanhoucke, S. Vega, Q. Vuong, F. Xia, T. Xiao, P. Xu, S. Xu, T. Yu, and B. Zitkovich, “Rt-1: Robotics transformer for real-world control at scale,” 2023. [Online]. Available: <https://arxiv.org/abs/2212.06817>
- [5] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter, S. Jakubczak, T. Jones, L. Ke, S. Levine, A. Li-Bell, M. Mothukuri, S. Nair, K. Pertsch, L. X. Shi, J. Tanner, Q. Vuong, A. Walling, H. Wang, and U. Zhilinsky, “ π_0 : A vision-language-action flow model for general robot control,” 2024. [Online]. Available: <https://arxiv.org/abs/2410.24164>
- [6] 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.
- [7] A. O’Neill, A. Rehman, A. Maddukuri, A. Gupta, A. Padalkar, A. Lee, A. Pooley, A. Gupta, A. Mandelkar, A. Jain, *et al.*, “Open x-embodiment: Robotic learning datasets and rt-x models: Open x-embodiment collaboration 0,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 6892–6903.
- [8] NVIDIA, J. Bjorck, F. Castaneda, N. Cherniadev, X. Da, R. Ding, L. J. Fan, Y. Fang, D. Fox, F. Hu, S. Huang, J. Jang, Z. Jiang, J. Kautz, K. Kundalia, L. Lao, Z. Li, Z. Lin, K. Lin, G. Liu, E. Llontop, L. Magne, A. Mandelkar, A. Narayan, S. Nasiriany, S. Reed, Y. L. Tan, G. Wang, Z. Wang, J. Wang, Q. Wang, J. Xiang, Y. Xie, Y. Xu, Z. Xu, S. Ye, Z. Yu, A. Zhang, H. Zhang, Y. Zhao, R. Zheng, and Y. Zhu, “Gr00t n1: An open foundation model for generalist humanoid robots,” *arXiv preprint arXiv:2503.14734*, 2025. [Online]. Available: <https://arxiv.org/abs/2503.14734>

- [9] S. Kareer, D. Patel, R. Punamiya, P. Mathur, S. Cheng, C. Wang, J. Hoffman, and D. Xu, "Egomimic: Scaling imitation learning via egocentric video," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 13 226–13 233.
- [10] Y. Liu, W. C. Shin, Y. Han, Z. Chen, H. Ravichandar, and D. Xu, "Immimic: Cross-domain imitation from human videos via mapping and interpolation," *arXiv preprint arXiv:2509.10952*, 2025.
- [11] Y. Qin, Y.-H. Wu, S. Liu, H. Jiang, R. Yang, Y. Fu, and X. Wang, "Dexmv: Imitation learning for dexterous manipulation from human videos," in *European Conference on Computer Vision*. Springer, 2022, pp. 570–587.
- [12] M. Lepert, J. Fang, and J. Bohg, "Phantom: Training robots without robots using only human videos," *arXiv preprint arXiv:2503.00779*, 2025.
- [13] G. Li, Y. Lyu, Z. Liu, C. Hou, J. Zhang, and S. Zhang, "H2r: A human-to-robot data augmentation for robot pre-training from videos," *arXiv preprint arXiv:2505.11920*, 2025.
- [14] R.-Z. Qiu, S. Yang, X. Cheng, C. Chawla, J. Li, T. He, G. Yan, D. J. Yoon, R. Hoque, L. Paulsen, *et al.*, "Humanoid policy~ human policy," *arXiv preprint arXiv:2503.13441*, 2025.
- [15] R. Yang, Q. Yu, Y. Wu, R. Yan, B. Li, A.-C. Cheng, X. Zou, Y. Fang, X. Cheng, R.-Z. Qiu, *et al.*, "Egovla: Learning vision-language-action models from egocentric human videos," *arXiv preprint arXiv:2507.12440*, 2025.
- [16] M. Hersch, F. Guenter, S. Calinon, and A. Billard, "Dynamical system modulation for robot learning via kinesthetic demonstrations," *IEEE Transactions on Robotics*, vol. 24, no. 6, pp. 1463–1467, 2008.
- [17] P. Kormushev, S. Calinon, and D. G. Caldwell, "Imitation learning of positional and force skills demonstrated via kinesthetic teaching and haptic input," *Advanced Robotics*, vol. 25, no. 5, pp. 581–603, 2011.
- [18] H. Li, Y. Cui, and D. Sadigh, "How to train your robots? the impact of demonstration modality on imitation learning," *arXiv preprint arXiv:2503.07017*, 2025.
- [19] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," *The International Journal of Robotics Research*, p. 02783649241273668, 2023.
- [20] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning fine-grained bimanual manipulation with low-cost hardware," 2023. [Online]. Available: <https://arxiv.org/abs/2304.13705>
- [21] 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," *arXiv preprint arXiv:2407.03162*, 2024.
- [22] 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," 2024. [Online]. Available: <https://arxiv.org/abs/2402.10329>
- [23] A. Carfi, M. Alameh, V. Belcamino, and F. Mastrogiovanni, "A modular architecture for imu-based data gloves," in *European Robotics Forum*. Springer, 2024, pp. 53–57.
- [24] Y. Niu, Y. Zhang, M. Yu, C. Lin, C. Li, Y. Wang, Y. Yang, W. Yu, T. Zhang, Z. Li, *et al.*, "Human2locoman: Learning versatile quadrupedal manipulation with human pretraining," *arXiv preprint arXiv:2506.16475*, 2025.
- [25] V. Liu, A. Adeniji, H. Zhan, S. Haldar, R. Bhirangi, P. Abbeel, and L. Pinto, "Egozero: Robot learning from smart glasses," *arXiv preprint arXiv:2505.20290*, 2025.
- [26] K. Grauman, A. Westbury, L. Torresani, K. Kitani, J. Malik, T. Afouras, K. Ashutosh, V. Baiyya, S. Bansal, B. Boote, *et al.*, "Ego-exo4d: Understanding skilled human activity from first-and third-person perspectives," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 19 383–19 400.
- [27] M. Lepert, J. Fang, and J. Bohg, "Phantom: Training robots without robots using only human videos," 2025.
- [28] S. Christen, L. Feng, W. Yang, Y.-W. Chao, O. Hilliges, and J. Song, "Synh2r: Synthesizing hand-object motions for learning human-to-robot handovers," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 3168–3175.
- [29] K. Zakka, A. Zeng, P. Florence, J. Tompson, J. Bohg, and D. Dwibedi, "Xirl: Cross-embodiment inverse reinforcement learning," 2021. [Online]. Available: <https://arxiv.org/abs/2106.03911>
- [30] C. Wang, L. Fan, J. Sun, R. Zhang, L. Fei-Fei, D. Xu, Y. Zhu, and A. Anandkumar, "Mimicplay: Long-horizon imitation learning by watching human play," 2023. [Online]. Available: <https://arxiv.org/abs/2302.12422>
- [31] Y. Zhu, A. Lim, P. Stone, and Y. Zhu, "Vision-based manipulation from single human video with open-world object graphs," 2024. [Online]. Available: <https://arxiv.org/abs/2405.20321>
- [32] 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 *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 4788–4795.
- [33] V. Jain, M. Attarian, N. J. Joshi, A. Wahid, D. Driess, Q. Vuong, P. R. Sanketi, P. Sermanet, S. Welker, C. Chan, I. Gilitschenski, Y. Bisk, and D. Dwibedi, "Vid2robot: End-to-end video-conditioned policy learning with cross-attention transformers," 2024. [Online]. Available: <https://arxiv.org/abs/2403.12943>
- [34] H. Zhou, R. Wang, Y. Tai, Y. Deng, G. Liu, and K. Jia, "You only teach once: Learn one-shot bimanual robotic manipulation from video demonstrations," *arXiv preprint arXiv:2501.14208*, 2025.
- [35] S. Haldar and L. Pinto, "Point policy: Unifying observations and actions with key points for robot manipulation," *arXiv preprint arXiv:2502.20391*, 2025.
- [36] B. Zitkovich, T. Yu, S. Xu, P. Xu, T. Xiao, F. Xia, J. Wu, P. Wohlhart, S. Welker, A. Wahid, *et al.*, "Rt-2: Vision-language-action models transfer web knowledge to robotic control," in *Conference on Robot Learning*. PMLR, 2023, pp. 2165–2183.
- [37] O. M. Team, D. Ghosh, H. Walke, K. Pertsch, K. Black, O. Mees, S. Dasari, J. Hejna, T. Kreiman, C. Xu, J. Luo, Y. L. Tan, L. Y. Chen, P. Sanketi, Q. Vuong, T. Xiao, D. Sadigh, C. Finn, and S. Levine, "Octo: An open-source generalist robot policy," 2024. [Online]. Available: <https://arxiv.org/abs/2405.12213>
- [38] P. Intelligence, K. Black, N. Brown, J. Darphinian, K. Dhabalia, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, M. Y. Galliker, D. Ghosh, L. Groom, K. Hausman, B. Ichter, S. Jakubczak, T. Jones, L. Ke, D. LeBlanc, S. Levine, A. Li-Bell, M. Mothukuri, S. Nair, K. Pertsch, A. Z. Ren, L. X. Shi, L. Smith, J. T. Springenberg, K. Stachowicz, J. Tanner, Q. Vuong, H. Walke, A. Walling, H. Wang, L. Yu, and U. Zhilinsky, " $\pi_{0.5}$: a vision-language-action model with open-world generalization," 2025. [Online]. Available: <https://arxiv.org/abs/2504.16054>
- [39] W. Peebles and S. Xie, "Scalable diffusion models with transformers," 2023. [Online]. Available: <https://arxiv.org/abs/2212.09748>
- [40] 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.
- [41] H. Luo, Y. Feng, W. Zhang, S. Zheng, Y. Wang, H. Yuan, J. Liu, C. Xu, Q. Jin, and Z. Lu, "Being-h0: Vision-language-action pretraining from large-scale human videos," 2025.
- [42] M. Shukor, D. Aubakirova, F. Capuano, P. Kooijmans, S. Palma, A. Zouitine, M. Aractingi, C. Pascal, M. Russi, A. Marafioti, S. Alibert, M. Cord, T. Wolf, and R. Cadene, "Smolvla: A vision-language-action model for affordable and efficient robotics," 2025.
- [43] G. Wang and Z. Lu, "Xlerobot: A practical low-cost household dual-arm mobile robot design for general manipulation," <https://github.com/Vector-Wangel/XLeRobot>, 2025.
- [44] G. Pavlakos, D. Shan, I. Radosavovic, A. Kanazawa, D. Fouhey, and J. Malik, "Reconstructing hands in 3d with transformers," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 9826–9836.
- [45] S. Haldar and L. Pinto, "Point policy: Unifying observations and actions with key points for robot manipulation," 2025.
- [46] J. Ren, P. Sundaresan, D. Sadigh, S. Choudhury, and J. Bohg, "Motion tracks: A unified representation for human-robot transfer in few-shot imitation learning," 2025.
- [47] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning fine-grained bimanual manipulation with low-cost hardware," 2023.