

# Human2Robot: Learning Robot Actions from Paired Human-Robot Videos

Sicheng Xie<sup>1,2\*</sup>, Haidong Cao<sup>1\*</sup>, Zejia Weng<sup>1</sup>, Zhen Xing<sup>1</sup>, Haoran Chen<sup>1</sup>,  
Shiwei Shen<sup>1</sup>, Jiaqi Leng<sup>1</sup>, Zuxuan Wu<sup>1,2†</sup>, Yu-Gang Jiang<sup>1</sup>

<sup>1</sup>Shanghai Key Lab of Intell. Info. Processing, School of CS, Fudan University

<sup>2</sup>Shanghai Innovation Institute

## Abstract

Distilling knowledge from human demonstrations is a promising way for robots to learn and act. Existing methods, which often rely on coarsely-aligned video pairs, are typically constrained to learning global or task-level features. As a result, they tend to neglect the fine-grained frame-level dynamics required for complex manipulation and generalization to novel tasks. We posit that this limitation stems from a vicious circle of inadequate datasets and the methods they inspire. To break this cycle, we propose a paradigm shift that treats fine-grained human-robot alignment as a conditional video generation problem. To this end, we first introduce H&R, a novel third-person dataset containing 2,600 episodes of precisely synchronized human and robot motions, collected using a VR teleoperation system. We then present HUMAN2ROBOT, a framework designed to leverage this data. HUMAN2ROBOT employs a Video Prediction Model to learn a rich and implicit representation of robot dynamics by generating robot videos from human input, which in turn guides a decoupled action decoder. Our real-world experiments demonstrate that this approach not only achieves high performance on seen tasks but also exhibits one-shot generalization to novel positions, objects, instances, and even new task categories.

## 1 Introduction

The ability to observe others, whether humans or animals, and acquire skills to solve new tasks is a fundamental reason why humans excel at tackling a wide range of complex challenges. To enable robots to assist with diverse real-world problems, it is essential that they develop a similar capability, particularly the ability to learn directly from human demonstrations. This has motivated extensive research (Bahl et al. 2023; Srirama et al. 2024; Nair et al. 2022; Wang et al. 2023; Bahl, Gupta, and Pathak 2022; Smith et al. 2019) on how to learn from human demonstrations effectively.

Despite recent progress, current approaches face a fundamental generalization gap. While they perform well on tasks encountered during training, they often fail entirely when presented with human demonstrations of unseen tasks. The prevailing paradigm relies on applying self-supervised (Xu

\*Equal Contribution.

†Corresponding Author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

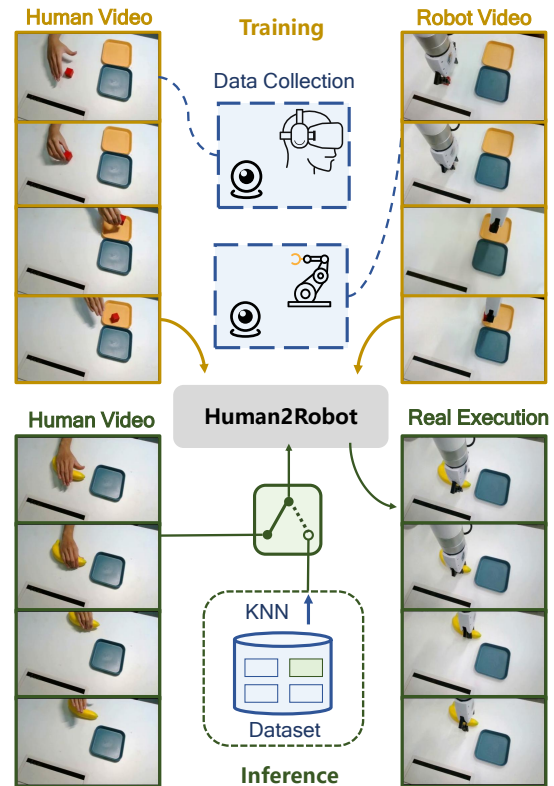


Figure 1: HUMAN2ROBOT: An human-video-conditioned policy, capable of completing seen tasks and one-shot performing unseen tasks with a single human video.

et al. 2023b) or contrastive learning (Jain et al. 2024) methods to coarsely aligned human-robot video pairs. However, the lack of fine-grained supervision in existing datasets fundamentally constrains what models can learn. Rather than uncovering detailed action-level correspondences, models are often limited to capturing global features or high-level task summaries. Consequently, many approaches extract holistic video representations using models such as the Perceiver Resampler (Alayrac et al. 2022), which compresses entire clips into fixed-length embeddings. This severely limits their ability to model the nuanced frame-level temporal

dynamics that are essential for generalization.

Furthermore, even when densely aligned data are available, models that rely on global feature matching remain fundamentally limited in their ability to capture fine-grained spatio-temporal structure. They often discard the frame-by-frame dynamics that are essential for imitating complex tasks. This limitation has created a vicious circle: the lack of fine-grained datasets gives rise to methods that are incapable of leveraging detailed supervision, while the dominance of such methods in turn discourages the development of fine-grained datasets needed to overcome this barrier. As such, we argue that to achieve true generalization where a robot can perform tasks demonstrated by a human but unseen in its own training data, we must break this cycle. This necessitates a paradigm shift in both learning methodology and data curation.

In this paper, we propose that the key to unlocking fine-grained human-robot alignment lies in video generation (Ho et al. 2022; Esser et al. 2023; Blattmann et al. 2023a; Wu et al. 2023; Khachatryan et al. 2023; Xing et al. 2024; Tu et al. 2024a). Instead of merely learning whether two videos are broadly similar or mapping task goals from the human domain to the robot domain, we aim to generate a corresponding robot video directly from a human demonstration. We posit that by training a model to predict the precise frame-by-frame evolution of a robot’s movements, it learns a far richer and more temporally coherent alignment. This generative process forces the model to internalize the intricate dynamics of manipulation, enabling it to understand how a task is done, not just what is done.

To train such a model would require a dataset with densely aligned and fine-grained human-robot video pairs. Manually annotating such data is prohibitively labor-intensive. To overcome this, we leverage virtual reality teleoperation to create a novel dataset, which we call H&R. By enhancing existing teleoperation systems with improved coordinate system matching, we achieve a seamless mapping between the operator’s hand movements and the robot arm’s motion. This allows us to efficiently collect a large-scale third-person dataset of 2,600 episodes, featuring perfectly synchronized videos of human hands and robot arms across 4 types of basic tasks and 6 long-horizon tasks.

Building on the proposed dataset, we introduce HUMAN2ROBOT, a framework designed to leverage the proposed H&R dataset. It employs a two-stage training process. First, we train a Video Prediction Model (VPM) built upon a pretrained Stable Diffusion model, which learns to translate a human video into a robot video. This model includes a spatial UNet for feature extraction, a behavior extractors for motion and position encoding, a spatial-temporal UNet that uses a spatial-temporal UNet architecture to explicitly model motion and temporal dynamics, generating a rich latent representation that captures core dynamics of the task. Second, we train an action decoder that conditions on the predictive representations generated by the VPM to output robot actions. This two-stage design effectively learn human-robot alignment and leverages the implicitly learned robot-dynamics features to guide the final policy learning.

With this carefully designed framework, HU-

MAN2ROBOT not only excels on seen tasks but also demonstrates remarkable one-shot generalization to novel object positions, appearances, and even entirely new task types and backgrounds. Furthermore, we introduce a KNN-based inference method that allows HUMAN2ROBOT to perform previously seen tasks with high precision, even without a human demonstration at test time.

In summary, our main contributions include:

- We present H&R, the first dataset featuring perfectly aligned videos of human hands and robotic arms across a variety of tasks, enabling high-fidelity learning from human demonstrations.
- We introduce HUMAN2ROBOT, an end-to-end generative framework utilizes a two-stage training process, which excels in carefully selected tasks, even with variations in positions, appearances, instances, backgrounds and different task types.
- We propose a KNN+HUMAN2ROBOT method, which integrates KNN for task prediction, enabling to perform tasks even without human videos as input. This further enhances the scalability and flexibility of the system.

## 2 Related Work

**Teleoperation.** Recently, VR-based methodologies (Iyer et al. 2024; Ding et al. 2024) have attracted considerable attention due to their cost-effectiveness, efficiency and versatility. However, these methods focus on controlling the robot, whereas our goal of using VR is to capture perfectly aligned videos of humans and robots, which are essential for robotic imitation.

**Learning from Human Videos.** Researchers have recently been attempting to leverage existing human-centric video datasets to enhance robot policy learning (Liu et al. 2018; Smith et al. 2020; Chen, Nair, and Finn 2021; Shaw, Bahl, and Pathak 2022; Zeng et al. 2024; Zhang et al. 2025b). Researchers propose to learn representations from human videos to assist in task execution (Xiao et al. 2022; Wang et al. 2023; Majumdar et al. 2023). However, these approaches need strong prior knowledge and struggle to transfer to robots (Nair et al. 2022; Bahl et al. 2023; Bahl, Gupta, and Pathak 2022). Meanwhile, some human-video-conditioned methods (Bharadhwaj et al. 2024; Jain et al. 2024; Xu et al. 2023b) have focused on aligning representations across human and robot videos, they are neither efficient nor capable of generalization. In contrast, our approach uses paired data and diffusion models to achieve strong generalization capabilities.

**Diffusion Models for Video Generation.** Current research has achieved remarkable performance of video generation (Ho et al. 2022; Esser et al. 2023; Blattmann et al. 2023a; Wu et al. 2023; Khachatryan et al. 2023; Xing et al. 2024, 2023a, 2025, 2023b; Tu et al. 2024a,b,c). Moreover, diffusion models for video stylization (Ye et al. 2025; Liu et al. 2023) and human image animation (Hu 2024; Tu et al. 2024a) have also shown its promising potential for visual and motion transfer. Moreover, diffusion model not only performs impressively in generative domains but also shows promising applications in visual understanding task. Some

notable approaches involve using diffusion models for object detection (Chen et al. 2023; Zhang et al. 2025a), image segmentation (Xu et al. 2023a), and visual representation learning (Yu et al. 2024; Weng et al. 2024). Thus, we aim to utilize the strong capabilities of diffusion model to learn human-robot alignment and leverages the implicitly learned robot-dynamics features to guide the final policy learning.

### 3 H&R Dataset

#### 3.1 Coordinate Alignment for Teleoperation

Although teleoperation has become a mainstream method for data collection, no one has yet attempted to use it for collecting paired data of humans and robots. We aim to provide a feasible solution for this.

We identify two core issues. The first issue is that the coordinate system of the robotic arm is not aligned with that of the human hand, which results in a mismatch between the movements of the human hand and the robotic arm in the video. For example, the hand may move across the entire screen, while the robotic arm only moves halfway. This visual difference could lead to difficulties in policy learning. The second issue is the embodiment gap between the human hand and the robotic gripper, which makes teleoperation unsuitable for certain tasks, such as screwing.

To address the issue of coordinate system alignment, we record corresponding three points in the real world of the robotic arm and human hand. We leverage these points to establish a shared coordinate system with consistent scale for both the robotic arm and the human hand. This enables the movement range of the robotic arm to match that of the human hand, as illustrated in Figure 1. Details on how to establish a shared coordinate system can be found in the Supplementary Materials.

However, the embodiment gap remains a challenging issue for current teleoperation system, and we leave it as future work. Therefore, at the data level, we focus on collecting pick-and-place data with human and robotic arm alignment. In our experiments, we evaluate how well the model generalizes after being trained on relatively simple data.

#### 3.2 Statistics of H&R Dataset

Based on our teleoperation method, we propose our H&R dataset, the first dataset featuring paired human and robot videos. Human events vary from 4 types of basic tasks and 6 long-horizon tasks, as shown in Figure 2. The whole dataset includes 2,600 episodes, and each episode contains frames ranging from 300 to 600. To the best of our knowledge, the H&R dataset is the first video dataset that ensures perfectly aligned video between human and robot.

## 4 HUMAN2ROBOT

In this section, we describe the two-stage training process of our HUMAN2ROBOT. First, we explore training a VPM to generate robot videos based on human videos, implicitly learning the corresponding robotic arm actions from human movements. Next, we design an action decoder that aggregates the predictive action representations within the VPM and outputs the corresponding robot actions.




Figure 2: Dataset Overview. (L) The ratio of four basic task types and long tasks. (R) Platform environment and the object instances used.

#### 4.1 Stage1 : Video Prediction Model (VPM)

We first explored the possibility of generating robot videos from human videos. Inspired by (Hu 2024), our VPM explores a Spatial UNet (S-UNet in short) and a Spatial-Temporal UNet (ST-UNet in short), working collaboratively to learn from humans. In particular, the S-UNet extracts features from the robot arm, which is further fed into the ST-UNet for temporal modeling. VPM also contains a behavior extractors that estimate position and motion clues from human videos. Figure 3 gives an overview of our pipeline. Below, we introduce the architecture of VPM in detail.

**Behavior Extractor.** Recall that we aim to learn from human demonstrations, and thus it is important to extract motion and position information from human videos. This is achieved by Behavior Extractor, which contains four convolution layers ( $4 \times$  kernels, with a stride of 2). Behavior Extractor takes the human image  $o_i^h$  ( $i \in [0, T]$ ) or the entire video of human  $o_{0:T}^h$  as inputs depending on the training stage, we will talk about it later in training section.

**Spatial UNet.** The S-UNet, consisting of 4 upsampling blocks and 4 downsampling blocks, extracts useful clues from the robotic arm. In particular, each block, known as the Spatial Layer (S-Layer in short), leverages self-attention to learn features tailored for robot arms, as well as cross-attention with pre-extracted CLIP embedding to incorporate semantic clues. The weights of the S-UNet are initialized from those of the Stable Diffusion (Rombach et al. 2022) to ease the learning process. The features derived by S-UNet provide a condition for predicting future frames.

While S-UNet introduces a similar number of parameters as the denoising UNet, it only extracts features once during the entire process, unlike diffusion-based video generation where each frame undergoes denoising multiple times. Therefore, it does not significantly increase computational overhead during inference.

**Spatial-Temporal UNet.** The ST-UNet aims to predict future frames, requiring the model to learn temporal dynamics. Thus, each block of the ST-UNet has a S-Layer which is the same as the one in S-UNet, followed by an additional Temporal Layer (T-Layer in short). More specifically, the ST-UNet takes the features from the behavior extractors and noise latent as inputs. In addition, each S-Layer in ST-UNet also concatenates the features of the corresponding layer of the S-UNet, exploring the robot arm as references for prediction. The T-Layer focuses on temporal modeling so as to produce high-fidelity future frames.

**VAE Encoder and Decoder.** HUMAN2ROBOT builds upon

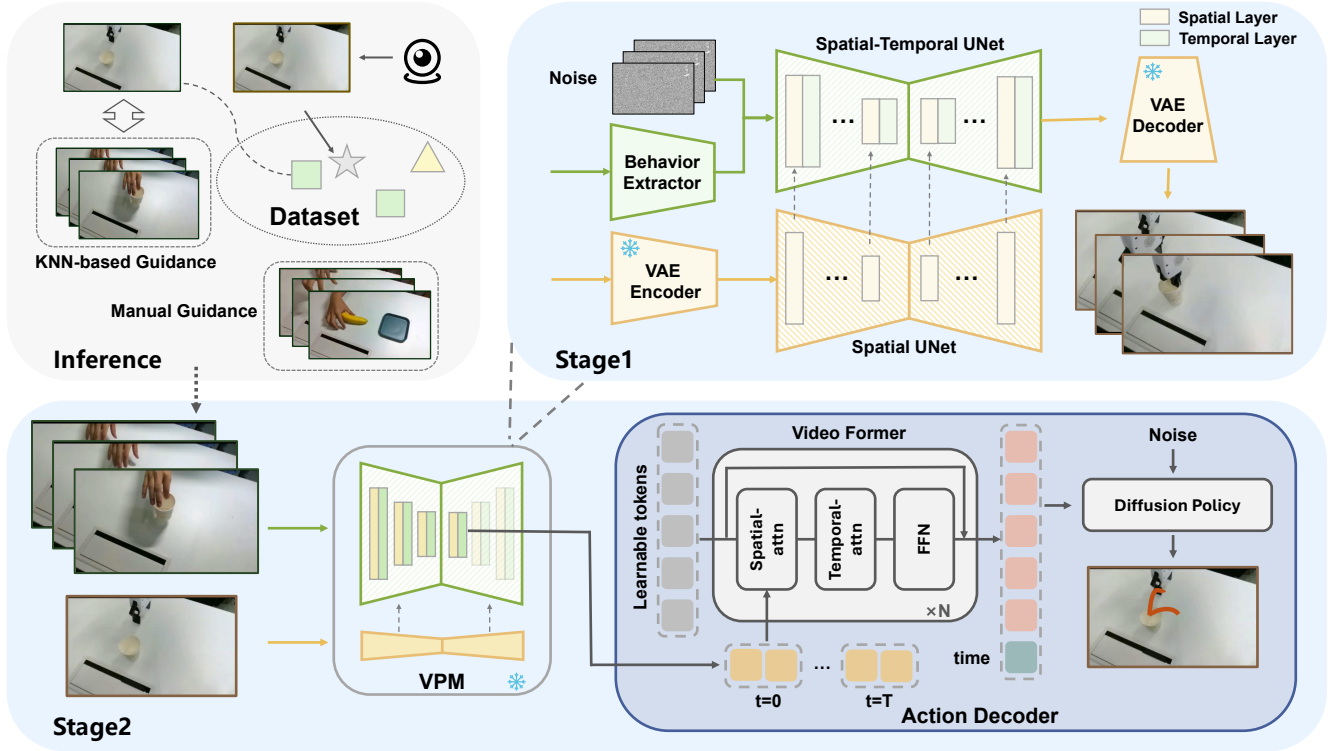


Figure 3: Architecture overview of HUMAN2ROBOT. Our approach consists of two training stages. In the first stage, we train a Video Prediction Model (VPM) to generate robotic arm videos conditioned on human videos. In the second stage, we freeze the VPM and train an action decoder to predict robot actions based on the motion features generated by the VPM.

Stable Diffusion (Rombach et al. 2022), which consists of an encoder and a decoder. The encoder turns images into latent embeddings for fast denoising, and the decoder maps the latents back to images. During training, both the encoder and decoder are kept frozen.

**Training strategy.** To generate higher-quality videos, the VPM training process consists of two stages. In the first stage, we focus on image generation to acquire basic generative capabilities. Next, we train it to generate videos.

In the first stage, by taking the first frame of robot  $o_0^r$ , the first frame of human  $o_0^h$  and future frame of human  $o_i^h$  ( $i \in [0, T]$ ) as inputs to predict the future frame of robot  $o_i^r$ , where  $T$  refer to the maximum length of video generation. In this stage, we train the S-UNet, the Behavior Extractors, and the ST-UNet without the temporal layers. The S-Layer in both the ST-UNet and S-UNet are instantiated using the powerful open-source Stable Diffusion model (SD) (Rombach et al. 2022), while the Behavior Extractors are initialized with Gaussian weights, except for the final projection layer, which utilizes zero initialized convolution. In the second stage, we concentrate the training effort on the temporal layer to generate video. In this stage, we take a 30-frame segment of human video  $o_{0:T}^h$  and the first frame of robot  $o_0^r$  as inputs to predict the future robot observation  $o_{0:T}^r$ . The entire human video segment is first passed into Behavior Extractor, which is subsequently added to noise latent and fed to our ST-UNet. The weights of the VAE En-

coder and Decoder, as well as the CLIP image encoder are frozen all the time.

The optimization objectives for this stage are as follows:

$$L_G = \mathbb{E}_{z_{ts}, c, \epsilon, ts} (\|\epsilon - \epsilon_\theta(z_{ts}, c, ts)\|_2^2), \quad (1)$$

the parameters of the formula represent the latent state  $z_{ts}$  at time step  $ts$ , which is obtained from ST-UNet. The conditioning variable  $c$  includes the first frame  $o_0^r$  of the robot and a human video  $o_{0:T}^h$ .  $\epsilon$  denotes the added noise, and  $\epsilon_\theta(z_{ts}, c, ts)$  represents the model’s predicted noise.

## 4.2 Stage2 : Action Decoding.

After the first stage of training, our VPM model is able to visually predict the corresponding robotic arm actions based on human movements. However, fully denoising an entire video is time-consuming, with most of the time spent on reconstructing image-level details, which are unrelated to manipulation. Increasingly, recent research (Hu et al. 2024; Wen et al. 2024; Zhu et al. 2025) has shown that the features of a generative model after a single denoising step already contain sufficient motion features to guide the action head’s planning. Inspired by these works, we treat the pretrained VPM as a vision encoder to extract the predicted robotic arm action information. This means adding intensity-invariant noise ( $t = K$ ) to the human video and using the first denoised result as the feature. As shown in the Figure 5, we can see

that the one-step denoising already contains a lot of action and position information.

Additionally, previous work (Weng et al. 2024; Xiang et al. 2023) has emphasized that the upsampling layers of diffusion models often contain more effective information. Therefore, we use the outputs of the upsampling layers of VPM as features; more precisely, we use those from the first upsampling layer.

In summary, during this stage, we freeze the parameters of the VPM and treat it as a visual encoder to train the subsequent action decoder. We use the first upsampling layer output  $F_{VPM}$ , after the initial denoising step, as the action prior features for the action decoder.

**Action Decoder.** Inspired by recent video conditioned works (Hu et al. 2024; Luo and Lu 2025), our action decoder consists of two parts: the Video Former (Blattmann et al. 2023b) and the Diffusion Policy (Chi et al. 2023).

The Video Former uses a learnable token  $Q$  to aggregate the video features  $F_{VPM}$  into a feature  $F_{VF} \in \mathbb{R}^{n \times L}$  with a fixed-length  $n$  and specified channel dimension  $L$ . Formally, this branch can be expressed as follows:

$$F_{VF} = \text{FFN}(\text{Temp-Attn}(\text{Spat-Attn}(Q, F_{VPM}))). \quad (2)$$

We incorporate the aggregated representation  $F_{VF}$  into the diffusion transformer blocks through cross-attention layers. The diffusion policy seeks to reconstruct the original actions  $a_0$  from the noisy actions  $a_k = \sqrt{\beta_k}a_0 + \sqrt{1 - \beta_k}\epsilon$ , where  $\epsilon$  represents white noise, and  $\beta_k$  is the noisy coefficient at step  $k$ . The optimization objectives for diffusion policy are as follows:

$$L_A = \mathbb{E}_{a_k, F_{VF}, \epsilon, k} (\|\epsilon - \epsilon_\phi(a_k, F_{VF}, k)\|_2^2). \quad (3)$$

### 4.3 KNN + HUMAN2ROBOT.

To avoid the need to explicitly provide human demonstration videos for seen task, we use a k-nearest neighbors (KNN) approach to identify the most probable task for the current scene. We retrieve the human demonstration video corresponding to the closest matching features as the conditioning input to guide the task execution. Specifically, we use DINOv2 (Oquab et al. 2023) and CLIP (Radford et al. 2021) as feature extractors to capture features from the first frame of each robotic arm video in the training set. During prediction, we select the  $n$  closest features based on the current environment, and the episode with the most frequent and closest match is chosen as the conditioning input, which is depicted in the Inference section of Figure 3.

## 5 Experiments

### 5.1 Experimental Setups

**Task Definition.** As mentioned in section 3.1, even with our adjustments, current teleoperation systems still cannot collect data for difficult tasks, such as screwing, when paired human and robot arm videos are required. Therefore, we mainly focus training and testing on pick-and-place tasks. While individual tasks are relatively simple, we aim to demonstrate the generalization ability of our model after

training on such tasks. This goal aligns with the spirit of learning from humans, that is, to learn from simple tasks and then generalize to more complex or even new tasks.

In addition to testing on tasks from the training set, we also test on many unseen tasks, including variations in appearance, position, instance, background, task combinations, and entirely new tasks. The task scenarios and their corresponding descriptions are shown in the Figure 4.

**Baselines.** Since HUMAN2ROBOT is characterized by learning from humans and video generation pretraining, we compare it with the following baselines:

- **Diffusion Policy (Chi et al. 2023):** A action learning policy using action diffusers with CLIP language conditions. For simplicity, we refer to it as DP.
- **XSkill (Xu et al. 2023b):** A human-video-conditioned policy through self-supervised learning.
- **Video Prediction Policy (Hu et al. 2024):** A language-conditioned policy that uses video prediction for pretraining. For simplicity, we refer to it as VPP.
- **Action Decoder w. Human:** Condition the action decoder of HUMAN2ROBOT on human videos instead of the features extracted by VPM. The human videos are processed by ResNet18 (He et al. 2016) and used as input to action decoder.
- **HUMAN2ROBOT w/o. Pretrain:** HUMAN2ROBOT without the video generation pretraining in Section 4.1.
- **HUMAN2ROBOT w. KNN:** HUMAN2ROBOT with our KNN method proposed in Section 4.3, which enables to perform tasks without explicit demonstrations.

**HUMAN2ROBOT Training Details.** As mentioned in Section 4, we used a two-stage training method. In the first stage, we trained a video prediction model, focusing on generating robotic arm videos based on human videos. During this stage, we pre-trained on 2,600 task videos, including some long tasks, such as picking up two blocks. Training for the first stage took 3 days using 4 NVIDIA A100 GPUs. In the second stage, we trained HUMAN2ROBOT on simple tasks, as the seen tasks shown in the Figure 4. This took about 6 hours using 8 NVIDIA A100 GPUs. However, for the writing task mentioned in Section 5.4, we train HUMAN2ROBOT solely on play data of writing for about 6 hours using 8 NVIDIA A100 GPUs. It is worth noting that the play data of writing here consists of random movements, without actually writing any characters.

### 5.2 Main results

**Quantitative Results.** The comparison on the basic tasks are show in Table 1. However, the DP baseline appears to converge primarily on push and pull tasks and performs poorly on the other tasks. Although XSkill conditions on human videos, it merely treats them as task labels. As a result, it can complete seen tasks but does not fully exploit the information in the human videos, leading to unstable performance and success rates that are 30–50% lower than HUMAN2ROBOT. VPP also employs video-generation pretraining and attains success rates close to those of HUMAN2ROBOT. Nevertheless, because




Figure 4: Task overview. We train the models on seen tasks and test them on different generalization ability level.

	Push & Pull	Pick & Place	Rotation	Average	Generalization	XSkill	VPP	H2R(ours)
<i>DP (Chi et al. 2023)</i>	50	20	15	28	<i>Appearance</i>	0	80	<b>100</b>
<i>XSkill (Xu et al. 2023b)</i>	70	40	50	53	<i>Position</i>	20	50	<b>80</b>
<i>VPP (Hu et al. 2024)</i>	95	70	75	80	<i>Instance</i>	0	0	<b>70</b>
<i>Action Decoder w. Human</i>	50	10	10	23	<i>Background</i>	0	0	<b>80</b>
<i>HUMAN2ROBOT w/o Pretrain</i>	20	10	0	10	<i>Combination</i>	0	0	<b>50</b>
<i>HUMAN2ROBOT w. KNN</i>	90	75	80	82	<i>Brand-New</i>	0	0	<b>70</b>
<b>HUMAN2ROBOT (ours)</b>	<b>100</b>	<b>90</b>	<b>95</b>	<b>95</b>				

Table 1: Left: multi-task success rates on basic tasks. Right: generalization success rates. 20 trails for each task.

VPP is language-conditioned—whereas HUMAN2ROBOT is video-conditioned, providing richer, fine-grained motion cues—HUMAN2ROBOT still leads by 10–20 percentage points in success rate. Overall, our proposed HUMAN2ROBOT achieves the highest success rate across all tasks, highlighting the benefits of conditioning on human videos coupled with video-generation pretraining.

**KNN Results.** As shown in Table 1, HUMAN2ROBOT with KNN outperforms all other baselines across all tasks, demonstrating that even without direct demonstrations, HUMAN2ROBOT can still achieve strong performance. However, compared to HUMAN2ROBOT, the KNN method shows a 10–20% decrease in success rate, which we consider to be within an acceptable range. Overall, these competitive results demonstrate that HUMAN2ROBOT is both efficient and accurate on seen tasks.

**Visualizations of Predictive Representations.** Since we repurpose the video prediction model as a visual encoder and extract predictive representations with a single forward pass, we examine the quality of these representations. In Figure 5, we visualize the ground-truth future alongside 1-step predictions and 30-step denoised outputs. The visualizations show that a 1-step denoised result already contains sufficient motion information for downstream tasks, validating the soundness of our approach. In addition, the 30-step (fully denoised) result is very close to the GT Robot video, demonstrating the effective design of our VPM.

### 5.3 Ablation Study

**Effectiveness of VPM.** Since our H&R dataset contains paired videos of humans and robotic arms, it naturally includes paired human observations and the corresponding robot actions. Therefore, we design *Action Decoder w. Human* to predict robot actions directly from human videos. Specifically, we use the human video  $o_{0:T}^h$  as a conditioning input to the action decoder to predict the corresponding robotic actions  $a_{0:T}$ .

However, in our experiments, this approach produced highly jittery executions, and the model was insensitive to grasping motions in the human videos, frequently failing to complete grasps. As a result, the average success rate was only 23%. We attribute this to the complexity of human motions: without any prior for learning the correspondence, it is difficult to infer the correct mapping from human videos. By contrast, HUMAN2ROBOT leverages video generation to learn the action correspondence, providing the downstream action decoder with a more reliable motion prior.

**Effectiveness of Video Pretraining.** We designed HUMAN2ROBOT *w/o. Pretraining* to assess the contribution of video-generation pretraining. In this variant, we initialize the VPM with the initial parameters described in Section 4.1, freeze the VPM, and train only the downstream action decoder. As shown in the table, this approach is almost incapable of completing tasks: it achieves 20% success on the simplest push-and-pull tasks and only 10% on pick-and-place.

We attribute this poor performance to Stable Diffusion




Figure 5: Visualization of VPM results. We can observe that that a 1-step denoised result already contains sufficient motion information for downstream tasks. In addition, the 30-step (fully denoised) result is very close to the GT Robot video, demonstrating the effective design of our Video Prediction Model (VPM).

initialization fails to extract task-relevant features and, in effect, injects additional noise into the observations—thereby degrading the policy. By contrast, the pretrained HUMAN2ROBOT delivers strong performance, demonstrating the effectiveness of video pretraining.

## 5.4 Generalization

To evaluate generalization ability, we test not only basic generalization types—position, appearance, instance, and background—but also stronger changes: combined and brand-new tasks. As shown in Table 1, compared with the baselines, HUMAN2ROBOT not only achieves substantially higher metrics on position and appearance, but also maintains solid performance in the other four settings where XSkill and VPP fail. We attribute this generalization to two factors: (1) unlike the self-training approach of XSkill, our method leverages paired human–robot data from H&R, enabling direct learning of the correspondence between human hands and robot arms; and (2) unlike VPP, our policy is human-video-conditioned rather than language-conditioned, so human videos supply richer, fine-grained motion cues, allowing strong generalization even when training on simple tasks and environments.

**Appearance Generalization.** We tested the generalization ability on objects with different colors, materials, and textures. HUMAN2ROBOT and VPP maintained their core capabilities without being affected, achieving success rates of 100% and 80%, respectively. However, XSkill could barely complete the tasks, showing limited generalization.

**Position Generalization.** Due to small variations in object placement within the dataset, all three models exhibit some degree of positional generalization. However, when the positional shift is large, XSkill achieves only a 20% success rate, VPP reaches only 50%, whereas HUMAN2ROBOT succeeds on 80% of the tested positions.

**Instance Generalization.** We tested the generalization ability on different instances. During training, all the models only encountered data for picking up blocks and pens. We tested whether the models could generalize to picking up objects such as ping pong balls and bananas. Due to the different shapes of the instances, their placement positions also varied slightly, adding complexity to the task. HUMAN2ROBOT still achieved a 70% success rate on this task, while the other baselines were unable to complete it.

**Background Generalization.** We also added many irrelevant objects and introduced unseen backgrounds to test the

generalization ability. We found that both baselines were unable to make correct predictions under these conditions, while HUMAN2ROBOT was still able to predict correctly, achieving a success rate of 80%.

**Task Combination.** We believe that a model with strong generalization ability should not only be able to complete tasks it has seen before, but also be capable of handling unseen tasks. At the beginning, we think such a model should have the ability to learn from each short task and then be able to complete long tasks composed of these short tasks. The challenge in this task lies in the fact that tasks from the previous stage may affect the execution of the subsequent stage. For example, we designed a task that involves pulling the plate and placing the cube onto it. Since the plate is moved, the placement position of the block changes, introducing significant difficulty. Therefore, only HUMAN2ROBOT, which has access to the corresponding human video, is capable of completing these tasks.

**Brand-New Task.** Teaching a robot to write is a challenging problem for current learning-based approaches. The core issue lies in the vast number of characters, with writing each character being a brand-new task. We cannot teach the model to write each character individually, so it must possess strong generalization capabilities. To demonstrate the performance of HUMAN2ROBOT in writing, we created a special dataset where each trajectory involves random movements on a desktop without writing any specific characters. We used this data to train the models to learn the correspondence between human and robots, and then during inference, we instructed the models to write brand-new characters, such as "H" or "R." Experimental results show that only HUMAN2ROBOT can learn the correspondence from meaningless data and complete the task of writing character.

## 6 Conclusion

In this paper, we use VR teleoperation systems to collect perfectly paired human-robot data, and create the H&R dataset. We then present HUMAN2ROBOT, which leverages a video prediction model (VPM) for human-robot alignment, and the robot-dynamics features captured turn out to be effective in guiding action decoding. Our evaluations demonstrate that HUMAN2ROBOT excels in seen tasks and unseen tasks.

## Acknowledgments

This work was supported by the Science and Technology Commission of Shanghai Municipality (No. 24511103100).

## References

- Alayrac, J.-B.; Donahue, J.; Luc, P.; Miech, A.; Barr, I.; Hasson, Y.; Lenc, K.; Mensch, A.; Millican, K.; Reynolds, M.; et al. 2022. Flamingo: a visual language model for few-shot learning. *NeurIPS*.
- Bahl, S.; Gupta, A.; and Pathak, D. 2022. Human-to-robot imitation in the wild. *arXiv preprint arXiv:2207.09450*.
- Bahl, S.; Mendonca, R.; Chen, L.; Jain, U.; and Pathak, D. 2023. Affordances from human videos as a versatile representation for robotics. In *CVPR*.
- Bharadhwaj, H.; Dwibedi, D.; Gupta, A.; Tulsiani, S.; Doersch, C.; Xiao, T.; Shah, D.; Xia, F.; Sadigh, D.; and Kirmani, S. 2024. Gen2act: Human video generation in novel scenarios enables generalizable robot manipulation. *arXiv preprint arXiv:2409.16283*.
- Blattmann, A.; Rombach, R.; Ling, H.; Dockhorn, T.; Kim, S. W.; Fidler, S.; and Kreis, K. 2023a. Align your latents: High-resolution video synthesis with latent diffusion models. In *CVPR*.
- Blattmann, A.; Rombach, R.; Ling, H.; Dockhorn, T.; Kim, S. W.; Fidler, S.; and Kreis, K. 2023b. Align your latents: High-resolution video synthesis with latent diffusion models. In *ICCV*.
- Chen, A. S.; Nair, S.; and Finn, C. 2021. Learning generalizable robotic reward functions from "in-the-wild" human videos. *RSS*.
- Chen, S.; Sun, P.; Song, Y.; and Luo, P. 2023. Diffusiondet: Diffusion model for object detection. In *ICCV*.
- Chi, C.; Xu, Z.; Feng, S.; Cousineau, E.; Du, Y.; Burchfiel, B.; Tedrake, R.; and Song, S. 2023. Diffusion policy: Visuomotor policy learning via action diffusion. *RSS*.
- Ding, R.; Qin, Y.; Zhu, J.; Jia, C.; Yang, S.; Yang, R.; Qi, X.; and Wang, X. 2024. Bunny-visionpro: Real-time bimanual dexterous teleoperation for imitation learning. *arXiv preprint arXiv:2407.03162*.
- Esser, P.; Chiu, J.; Atighehchian, P.; Granskog, J.; and Germanidis, A. 2023. Structure and content-guided video synthesis with diffusion models. In *ICCV*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- Ho, J.; Salimans, T.; Gritsenko, A.; Chan, W.; Norouzi, M.; and Fleet, D. J. 2022. Video diffusion models. In *NeurIPS*.
- Hu, L. 2024. Animate anyone: Consistent and controllable image-to-video synthesis for character animation. In *CVPR*.
- Hu, Y.; Guo, Y.; Wang, P.; Chen, X.; Wang, Y.-J.; Zhang, J.; Sreenath, K.; Lu, C.; and Chen, J. 2024. Video prediction policy: A generalist robot policy with predictive visual representations. *arXiv preprint arXiv:2412.14803*.
- Iyer, A.; Peng, Z.; Dai, Y.; Guzey, I.; Haldar, S.; Chintala, S.; and Pinto, L. 2024. Open teach: A versatile teleoperation system for robotic manipulation. *arXiv preprint arXiv:2403.07870*.
- Jain, V.; Attarian, M.; Joshi, N. J.; Wahid, A.; Driess, D.; Vuong, Q.; Sanketi, P. R.; Sermanet, P.; Welker, S.; Chan, C.; et al. 2024. Vid2robot: End-to-end video-conditioned policy learning with cross-attention transformers. *arXiv preprint arXiv:2403.12943*.
- Khachatryan, L.; Movsisyan, A.; Tadevosyan, V.; Henschel, R.; Wang, Z.; Navasardyan, S.; and Shi, H. 2023. Text2Video-Zero: Text-to-Image Diffusion Models are Zero-Shot Video Generators. In *ICCV*.
- Liu, G.; Xia, M.; Zhang, Y.; Chen, H.; Xing, J.; Wang, Y.; Wang, X.; Yang, Y.; and Shan, Y. 2023. Stylecrafter: Enhancing stylized text-to-video generation with style adapter. *arXiv preprint arXiv:2312.00330*.
- Liu, Y.; Gupta, A.; Abbeel, P.; and Levine, S. 2018. Imitation from observation: Learning to imitate behaviors from raw video via context translation. In *ICRA*.
- Luo, H.; and Lu, Z. 2025. Learning Video-Conditioned Policy on Unlabelled Data with Joint Embedding Predictive Transformer. In *ICLR*.
- Majumdar, A.; Yadav, K.; Arnaud, S.; Ma, J.; Chen, C.; Silwal, S.; Jain, A.; Berges, V.-P.; Wu, T.; Vakili, J.; et al. 2023. Where are we in the search for an artificial visual cortex for embodied intelligence? In *NeurIPS*.
- Nair, S.; Rajeswaran, A.; Kumar, V.; Finn, C.; and Gupta, A. 2022. R3M: A Universal Visual Representation for Robot Manipulation. In *CoRL*.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *CVPR*.
- Shaw, K.; Bahl, S.; and Pathak, D. 2022. VideoDex: Learning Dexterity from Internet Videos. *CoRL*.
- Smith, L.; Dhawan, N.; Zhang, M.; Abbeel, P.; and Levine, S. 2019. Avid: Learning multi-stage tasks via pixel-level translation of human videos. *arXiv preprint arXiv:1912.04443*.
- Smith, L.; Dhawan, N.; Zhang, M.; Abbeel, P.; and Levine, S. 2020. Avid: Learning multi-stage tasks via pixel-level translation of human videos. *arXiv preprint arXiv:1912.04443*.
- Srirama, M. K.; Dasari, S.; Bahl, S.; and Gupta, A. 2024. Hrp: Human affordances for robotic pre-training. *arXiv preprint arXiv:2407.18911*.
- Tu, S.; Dai, Q.; Cheng, Z.-Q.; Hu, H.; Han, X.; Wu, Z.; and Jiang, Y.-G. 2024a. Motioneditor: Editing video motion via content-aware diffusion. In *CVPR*.
- Tu, S.; Dai, Q.; Zhang, Z.; Xie, S.; Cheng, Z.-Q.; Luo, C.; Han, X.; Wu, Z.; and Jiang, Y.-G. 2024b. Motionfollower: Editing video motion via lightweight score-guided diffusion. *arXiv preprint arXiv:2405.20325*.

- Tu, S.; Xing, Z.; Han, X.; Cheng, Z.-Q.; Dai, Q.; Luo, C.; and Wu, Z. 2024c. StableAnimator: High-Quality Identity-Preserving Human Image Animation. *arXiv preprint arXiv:2411.17697*.
- Wang, C.; Fan, L.; Sun, J.; Zhang, R.; Fei-Fei, L.; Xu, D.; Zhu, Y.; and Anandkumar, A. 2023. Mimicplay: Long-horizon imitation learning by watching human play. *arXiv preprint arXiv:2302.12422*.
- Wen, Y.; Lin, J.; Zhu, Y.; Han, J.; Xu, H.; Zhao, S.; and Liang, X. 2024. Vidman: Exploiting implicit dynamics from video diffusion model for effective robot manipulation. *NeuIPS*.
- Weng, Z.; Yang, X.; Xing, Z.; Wu, Z.; and Jiang, Y.-G. 2024. Genrec: Unifying video generation and recognition with diffusion models. *arXiv preprint arXiv:2408.15241*.
- Wu, J. Z.; Ge, Y.; Wang, X.; Lei, S. W.; Gu, Y.; Shi, Y.; Hsu, W.; Shan, Y.; Qie, X.; and Shou, M. Z. 2023. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In *ICCV*.
- Xiang, W.; Yang, H.; Huang, D.; and Wang, Y. 2023. Denoising diffusion autoencoders are unified self-supervised learners. In *ICCV*.
- Xiao, T.; Radosavovic, I.; Darrell, T.; and Malik, J. 2022. Masked visual pre-training for motor control. *arXiv preprint arXiv:2203.06173*.
- Xing, Z.; Dai, Q.; Hu, H.; Chen, J.; Wu, Z.; and Jiang, Y.-G. 2023a. Svformer: Semi-supervised video transformer for action recognition. In *CVPR*.
- Xing, Z.; Dai, Q.; Hu, H.; Wu, Z.; and Jiang, Y.-G. 2024. Simda: Simple diffusion adapter for efficient video generation. In *CVPR*.
- Xing, Z.; Dai, Q.; Weng, Z.; Wu, Z.; and Jiang, Y.-G. 2025. Aid: Adapting image2video diffusion models for instruction-guided video prediction. In *ICCV*.
- Xing, Z.; Dai, Q.; Zhang, Z.; Zhang, H.; Hu, H.; Wu, Z.; and Jiang, Y.-G. 2023b. Vidiff: Translating videos via multi-modal instructions with diffusion models. *arXiv preprint arXiv:2311.18837*.
- Xu, J.; Liu, S.; Vahdat, A.; Byeon, W.; Wang, X.; and De Mello, S. 2023a. Open-vocabulary panoptic segmentation with text-to-image diffusion models. In *CVPR*.
- Xu, M.; Xu, Z.; Chi, C.; Veloso, M.; and Song, S. 2023b. Xskill: Cross embodiment skill discovery. In *CoRL*.
- Ye, Z.; Huang, H.; Wang, X.; Wan, P.; Zhang, D.; and Luo, W. 2025. Stylemaster: Stylize your video with artistic generation and translation. In *CVPR*.
- Yu, S.; Kwak, S.; Jang, H.; Jeong, J.; Huang, J.; Shin, J.; and Xie, S. 2024. Representation Alignment for Generation: Training Diffusion Transformers Is Easier Than You Think. *arXiv preprint arXiv:2410.06940*.
- Zeng, J.; Bu, Q.; Wang, B.; Xia, W.; Chen, L.; Dong, H.; Song, H.; Wang, D.; Hu, D.; Luo, P.; et al. 2024. Learning Manipulation by Predicting Interaction. *arXiv preprint arXiv:2406.00439*.
- Zhang, H.; Wang, Z.; Zeng, D.; Wu, Z.; and Jiang, Y.-G. 2025a. DiffusionAD: Norm-guided one-step denoising diffusion for anomaly detection. *TPAMI*.
- Zhang, S.; Xu, Z.; Liu, P.; Yu, X.; Li, Y.; Gao, Q.; Fei, Z.; Yin, Z.; Wu, Z.; Jiang, Y.-G.; et al. 2025b. Vlabench: A large-scale benchmark for language-conditioned robotics manipulation with long-horizon reasoning tasks. In *ICCV*.
- Zhu, C.; Yu, R.; Feng, S.; Burchfiel, B.; Shah, P.; and Gupta, A. 2025. Unified world models: Coupling video and action diffusion for pretraining on large robotic datasets. *arXiv preprint arXiv:2504.02792*.