

Hierarchical Diffusion Policy for Kinematics-Aware Multi-Task Robotic Manipulation

Xiao Ma, Sumit Patidar, Iain Haughton, Stephen James
Dyson Robot Learning Lab

{xiao.ma, sumit.patidar, iain.haughton, stephen.james}@dyson.com

Abstract

This paper introduces *Hierarchical Diffusion Policy (HDP)*, a hierarchical agent for multi-task robotic manipulation. HDP factorises a manipulation policy into a hierarchical structure: a high-level task-planning agent which predicts a distant next-best end-effector pose (NBP), and a low-level goal-conditioned diffusion policy which generates optimal motion trajectories. The factorised policy representation allows HDP to tackle both long-horizon task planning while generating fine-grained low-level actions. To generate context-aware motion trajectories while satisfying robot kinematics constraints, we present a novel kinematics-aware goal-conditioned control agent, *Robot Kinematics Diffuser (RK-Diffuser)*. Specifically, RK-Diffuser learns to generate both the end-effector pose and joint position trajectories, and distill the accurate but kinematics-unaware end-effector pose diffuser to the kinematics-aware but less accurate joint position diffuser via differentiable kinematics. Empirically, we show that HDP achieves a significantly higher success rate than the state-of-the-art methods in both simulation and real-world.¹

1. Introduction

Learning efficient visual manipulation strategies in robotics is challenging due to diverse environments, objects, and robot trajectories. The choice of policy representation strongly influences agent performance.

One way of parameterising the policy is to directly map visual observations to robot commands, e.g., joint position or velocity actions [18, 22, 27, 39]. These approaches make the least assumptions of the task and environment and retain the flexible control of the over-actuated, but they often suffer from low sample efficiency and poor generalisation ability, especially for long-horizon tasks [20, 34].

Recent advances in learning next-best-pose (NBP) agents [6, 7, 15–17, 20, 34, 43] have significantly improved




Figure 1. We introduce HDP, a hierarchical agent for robotic manipulation. At the high-level, HDP learns to predict the next-best end-effector pose. Conditioned on the current and the predicted pose (red), a diffusion model generates an action trajectory for the robot to follow (blue). In contrast, the trajectories generated by classic planners (yellow) cannot be executed due to violating environment constraints, e.g., the hinge of the box.

the sample efficiency and performance for robotic manipulation. Instead of learning continuous actions, NBP agents directly predict a distant “keyframe” [17], a next-best end-effector pose, and use a predefined motion planner to compute a trajectory for the agent to follow. However, as the motion planner is unaware of the task context, it will fail to perform tasks that require understanding the environment context, e.g., dynamics. For example, in Fig. 1 to open the box, the agent has to understand the unknown physics properties of the hinge, e.g., the resistance force, and only a specific curved trajectory can be successfully executed.

In this work, we introduce Hierarchical Diffusion Policy (HDP), a hierarchical multi-task agent that combines the best of both worlds. HDP factorises a manipulation policy by chaining a high-level NBP agent with a low-level learned controller. At the high level, HDP takes the 3D visual observations and language instructions as the inputs, and predicts a 6-DoF next-best end-effector pose. At the high level, HDP entails the capability of understanding the visual environment and language instructions and performing long-horizon task-level decision-making. At the low level, given the high-level 6-DoF end-effector pose action as a goal, HDP casts the control task as a context-aware 6-DoF pose-reaching task. We introduce a novel kinematics-aware low-level agent, *Robot Kinematics Diffuser (RK-Diffuser)*, a diffusion-based policy [5] that directly generates the mo-

¹Code and videos are available in our [project page](#).




Figure 2. We focus on learning multi-task language-guided agent for robotic manipulation. Unlike a standard motion planner that only samples an arbitrary trajectory to the end pose.

tion trajectory via conditional sampling and trajectory inpainting. Specifically, instead of generating the end-effector pose trajectories as in Chi et al. [5], Xian et al. [40] and solving the robot inverse kinematics, RK-Diffuser learns both end-effector pose and robot joint position diffusion, and distill the accurate but kinematics-unaware end-effector pose trajectory into the joint position trajectory via differentiable robot kinematics. RK-Diffuser achieves accurate trajectory generation and maximum control flexibility, while avoiding violating the robot kinematic constraints, which is a common issue of inverse kinematics solvers.

In our experiments, we empirically analyse HDP on a wide range of challenging manipulation tasks in RL-Bench [19]. We show that (1) RK-Diffuser generally achieves a higher success rate on goal-conditioned motion generation. (2) The proposed hierarchical agent, HDP, outperforms the flat baseline agents and other hierarchical variants. (3) HDP can be directly trained on a real robot with only 20 demonstrations on a challenging oven opening task with a high success rate.

2. Related Works

2.1. End-to-End Visual Manipulation Agents

End-to-end manipulation approaches [18, 22, 27, 39] make the least assumption about objects and tasks, and learn a direct mapping from RGB images to a robot action, but tend to be sample-inefficient. Two recent approaches have emerged to combat this sample inefficiency: (1) the next-best pose (NBP) action mode that learns to directly predict a distant “keyframe” [17]; (2) 3D action-value maps [20] that aligns the 3D task space and the action space, by learning 3D voxel-based action-value maps as policies, and extracting actions by taking the coordinates of the voxel with the highest value. Such a structured action space significantly reduces the amount of data needed and the generalisation of the learned policy. In particular, built with Transformers backbones, Shridhar et al. [34] and Gervet et al. [6] are able

to take in language tokens as input, and develop language-conditioned policies. In this work, without loss of generality, we choose PerAct [34] as our high-level language-conditioned agent for various tasks. Taking the predicted 6-DoF NBP as input, RK-Diffuser naturally works as a low-level policy of PerAct. Similar to our work, James and Abbeel [15] combines a high-level C2F-ARM [20] with a low-level agent that learns to rank a set of sampled trajectories by human heuristics. This approach has been shown to work on a series of challenging manipulation tasks, but it is computationally heavy and not scalable, conditioned on predefined motion generators. We show that HDP achieves strong multi-task manipulation capabilities with both kinematics awareness and high accuracy.

2.2. Diffusion Models

The diffusion model is a powerful class of generative models that learn to approximate the data distribution by iterative denoising processes. They have shown impressive results on both the conditional and unconditional image, video, and 3D object generation [10, 11, 26, 32, 35, 38]. In the field of decision making, diffusion models have recently been adopted as a powerful policy class [1, 5, 21, 23, 37]. Specifically, Diffusion Policies [5] learn to generate diverse multi-modal trajectories for robot manipulation by conditional generation with imitation learning. Concurrent with our work, Xian et al. [40] proposes ChainedDiffuser as a hierarchical agent. As we show in our experiments, the gripper-pose diffusion policy in ChainedDiffuser relies on inverse kinematics solvers to generate robot joint actions, which, however, is susceptible to prediction errors and might violate the kinematic constraints of the robot. On the contrary, the proposed RK-Diffuser learns both the end-effector pose and joint position trajectories and refines the joint position trajectories by distilling the end-effector poses.

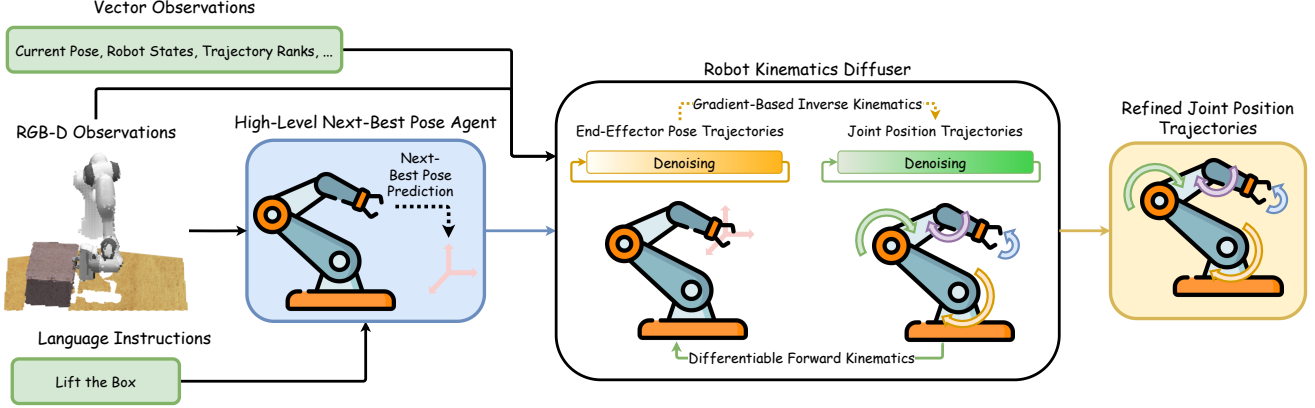


Figure 3. Overview of Hierarchical Diffusion Policy (HDP). HDP is a multi-task hierarchical agent for kinematics-aware robotic manipulation. HDP consists of two levels: a high-level language-guided agent and a low-level goal-conditioned diffusion policy. From left to right, the high-level agent takes in 3D environment observations and language instructions, then predicts the next-best end-effector pose. This pose guides the low-level RK-Diffuser. The RK-Diffuser subsequently generates a continuous joint-position trajectory by conditional sampling and trajectory inpainting given the next-best pose and environment observations. To generate kinematics-aware trajectories, RK-Diffuser distills the accurate but less flexible end-effector pose trajectories into joint position space via differentiable robot kinematics.

2.3. Differentiable Physics for Decision Making

Differentiable physics simulation constructs each simulation step as a differentiable computational graph, such that the environment steps are fully differentiable with respect to network parameters [4, 12, 13, 42]. Learning decision-making policies via differentiable physics has shown to be more efficient and generalisable compared with standard non-differentiable environments [3, 41] with the physics priors as an inductive bias to the gradients. Similar to differentiable physics, we make use of the differentiable robot kinematics models [45] to distill the accurate but less reliable end-effector pose trajectory to the joint position space.

3. Preliminaries

3.1. Diffusion Models

Diffusion models are a powerful family of generative models that consist of forward and backward Markov-chain diffusion processes. Consider a real data distribution $q(x)$ and a sample $x^0 \sim q(x)$ drawn from it. The forward diffusion processes adds Gaussian noise to x^0 in K steps, which gives a sequence of noisy samples $\{x^i\}_{i=1}^K$. In DDPM [10], the noise is controlled by a variance scheduler

$$q(x^k|x^{k-1}) = \mathcal{N}(x^k; \sqrt{1 - \beta^k}x^{k-1}, \beta^k \mathbf{I}), \quad (1)$$

$$q(x^{1:K}|x^0) = \prod_{k=1}^K q(x^k|x^{k-1}). \quad (2)$$

where β^1, \dots, β^K are scheduler parameters. Theoretically, x^∞ will distribute as an isotropic Gaussian. To reconstruct distribution $q(x)$, diffusion models learn a conditional dis-

tribution $p_\theta(x^{t-1}|x^t)$ and generate new samples by

$$p_\theta(x^{0:K}) = p(x^K) \prod_{k=1}^K p_\theta(x^{k-1}|x^k), \quad (3)$$

$$p_\theta(x^{k-1}|x^k) = \mathcal{N}(x^{k-1}; \mu_\theta(x^k, k), \Sigma_\theta(x^k, k)), \quad (4)$$

where $p(x^K) = \mathcal{N}(\mathbf{0}, \mathbf{I})$ under the condition that $\prod_{k=1}^K (1 - \beta^k) \approx 0$. The model can be trained by maximising the evidence lower bound (ELBO)

$$\mathbb{E}_{x^0}[\log p_\theta(x^0)] \geq \mathbb{E}_q \left[\log \frac{p_\theta(x^{0:K})}{q(x^{1:K}|x^0)} \right] \quad (5)$$

In the context of decision making, diffusion policies consider a trajectory of actions $a_{1:T} = \{a(t)\}_{t=1}^T$, and learn a conditional distribution $p_\theta(a_{1:T}^{k-1}|a_{1:T}^k, \{c_i\}_{i=1}^N)$, where $\{c_i\}_{i=1}^N$ are N additional conditions for policy learning, e.g., RGB observations, point cloud, robot states, etc. For simplicity, we abuse the notation and denote $a_{1:T}^k$ as \mathbf{a}^k .

3.2. Differentiable Kinematics

Differentiable simulation aims to encode the physics simulation steps as a differentiable computational graph. Take a simple point-mass system as an example.

$$y_t = y_{t-1} + \Delta t * v_t, \quad v_{t+1} = v_t + \Delta t * \frac{F}{m} \quad (6)$$

where the force F is the input to the system, m is the mass, v is the speed, and y is the position of the point. Importantly, such a system is differentiable and we can optimise the input force f by the gradients from positions y . Similarly in the context of robotics, conditioned on a predefined URDF

model of a robot, the end-effector pose s_p of a robot can be obtained by a differentiable forward kinematics function f_K as $s_p = f_K(s_j)$, where s_j is the joint angles. Thus, given a loss function $L(s_p)$ over the gripper poses, the joint positions can be directly updated by gradients $\frac{\partial L(s_p)}{\partial s_j}$.

4. Hierarchical Diffusion Policy

The overall pipeline of HDP is shown in Fig. 3.

Problem Definition. We aim to learn a HDP policy $\pi(\mathbf{a} \mid o, l)$, which processes the RGB-D observation o and language instruction l , specifying the task, to predict a hybrid action \mathbf{a} . Here, \mathbf{a} consists of a trajectory $\mathbf{a}_{\text{joint}} = \{a(0), a(1), \dots, a(T)\}$ and gripper opening / closing action a_{grip} , where T is the trajectory length and $a(i) \in \mathbb{R}^N$, with N denoting the number of the robot joints. For brevity, we symbolise actions \mathbf{a} without temporal index in the episode.

Factorised Hierarchical Policy. To tackle long-horizon context-aware manipulation tasks, we factorise the policy $\pi(\mathbf{a} \mid o, l)$ into a hierarchical structure. Specifically, $\pi(\mathbf{a} \mid o, l) = \pi_{\text{high}}(a_{\text{high}} \mid o, l) \circ \pi_{\text{low}}(\mathbf{a} \mid o, a_{\text{high}})$. Here, the high-level action $a_{\text{high}} = (a_{\text{pose}}, a_{\text{grip}}) \sim \pi_{\text{high}}$, consists of (1) the end-effector pose action $a_{\text{pose}} = (a_{\text{trans}}, a_{\text{rot}})$, with translation action $a_{\text{trans}} \in \mathbb{R}^3$ and quaternion rotation action $a_{\text{rot}} \in \mathbb{R}^4$; and (2) a binary gripper action $a_{\text{grip}} \in \mathbb{R}$. Conditioned on the high-level action a_{high} , we parameterise the low-level policy $\pi_{\text{low}}(\mathbf{a} \mid a_{\text{high}}, o)$ with RK-Diffuser, and learn to generate accurate joint position trajectories $\mathbf{a}_{\text{low}} = \mathbf{a}_{\text{joint}}$. Such a factorisation offloads the complex and expensive task-level understanding from language instructions to the high-level agent, leaving only control to be learned by a simple, goal-conditioned low-level agent. During inference, HDP works in a sequential manner and we take $\mathbf{a} = \{\mathbf{a}_{\text{joint}}, a_{\text{grip}}\}$ as the output.

4.1. Dataset Preparation

We assume access to a multi-task dataset $\mathcal{D} = \{\xi_i\}_{i=1}^{N_D}$, containing a total of N_D expert demonstrations paired with $\mathcal{D}_l = \{l_i\}_{i=1}^{N_D}$ language descriptions. Note that a single task might have multiple variations, each with different description, e.g., “open the middle drawer” or “open the bottom drawer”. Each demonstration $\xi = \{\mathbf{a}_{\text{demo}}, \mathbf{o}_{\text{demo}}\}$, consists of an expert trajectory \mathbf{a}_{demo} and resulting observation \mathbf{o}_{demo} . To enable the training of both the high-level policy π_{high} and the low-level RK-Diffuser π_{low} , the action \mathbf{a}_{demo} includes: (1) end-effector poses \mathbf{a}_{pose} ; (2) gripper opening / closing action a_{grip} ; and (3) joint positions $\mathbf{a}_{\text{joint}}$. The observation \mathbf{o}_{demo} includes multi-view calibrated RGB-D camera observations and robot states.

Keyframe Discovery. Referring to prior works [17, 20], training the high-level agent on all trajectory points is inefficient and instead we apply a Keyframe discovery method introduced in James and Davison [17]. Scanning through

each trajectory ξ , we extract a set of K_ξ keyframe indices $\{k_i\}_{i=1}^{K_\xi}$ that capture the principal bottleneck end-effector poses. Specifically, a frame is considered as a keyframe if (1) the joint velocity is close to 0; and (2) the gripper open / close state remains unchanged. Unlike prior works, which only keep keyframes for training, we maintain the keyframe indices and extract different segments of data to train both high-level and low-level agents. The details will be discussed in the following sections.

4.2. High-Level Next-Best Pose Agent

For the high-level policy $a_{\text{high}} = (a_{\text{pose}}, a_{\text{grip}}) \sim \pi_{\text{high}}(a \mid o, l)$, we utilise a next-best pose agent [17] with structured action representations. In this work, to parameterise π_{high} and fulfil this objective, we employ Perceiver-Actor (PerAct) [34]. PerAct is a language conditioned Behaviour Cloning (BC) agent with Transformer [36] backbones. PerAct achieves its high sample-efficiency, generalisability and accuracy through the use of a high-resolution voxel scene representations to predict 3D voxel-based action-value maps. To tackle the large number of visual and language tokens, PerAct adopts PerceiverIO [14], which encodes the inputs with a small set of latent vectors and reduces the computational complexity.

Action Spaces. PerAct uses discrete action spaces for all action heads, including (1) a discrete policy head over the voxels for a_{trans} and (2) a pair of discrete policies for a_{rot} and a_{grip} . Continuous actions are reconstructed by converting the discrete indices according to the action space ranges.

Model Training. For the high-level agent, we use only the keyframes for training. In addition, following Shridhar et al. [34], we use demo augmentation and translation augmentation to generate more samples. The network is optimised by behaviour cloning losses, i.e., cross-entropy losses in the discrete action space

$$\mathcal{L}_{\text{high}} = -\mathbb{E}_{k \sim \xi, \xi \sim \mathcal{D}} [\log \pi_{\text{high}}(a_{\text{demo}}(k) \mid o, l)] \quad (7)$$

where $a_{\text{demo}}(k)$ is the expert action of the keyframe k .

4.3. Low-Level RK-Diffuser

Given the predicted high-level action a_{high} , we perform conditional trajectory generation with RK-Diffuser through denoising diffusion processes. Standard diffusion policy for robotic manipulation considers end-effector pose diffusion

$$\begin{aligned} p_\theta(\mathbf{a}_{\text{pose}}^{k-1} \mid \mathbf{a}_{\text{pose}}^k, C_{\text{pose}}) \\ = \mathcal{N}(\mathbf{a}_{\text{pose}}^{k-1}; \boldsymbol{\mu}_\theta(\mathbf{a}_{\text{pose}}^k, C_{\text{pose}}, k), \boldsymbol{\Sigma}_\theta(\mathbf{a}_{\text{pose}}^k, C_{\text{pose}}, k)) \end{aligned} \quad (8)$$

where C_{pose} consists of the conditional variables, including the known start pose $a(0)_{\text{pose}}^0$, the predicted next-best

pose $\hat{\mathbf{a}}_{\text{pose}}^0(T)$ by the high-level agent, the low-dimensional state s of the robot, the end-effector pose, the gripper open amount, and the point cloud of the environment \mathbf{v} .

Besides using the start and next-best pose as conditional variables to the networks, we inpaint the trajectory with the start pose and the predicted next-best pose at each denoising step. This end-effector pose diffusion allows the inpainting operation to act as a hard constraint for the diffusion process, which guarantees the last step in the trajectory will always be aligned with the output of the high-level agent.

Prior to execution, the end-effector pose trajectory must undergo processing by an inverse kinematics (IK) solver to determine corresponding joint positions. However, the predicted end-effector pose trajectory lacks kinematics awareness and there is a high likelihood of it violating the kinematic constraints. Consider, for example, that each step of the predicted trajectory has a probability p to violate the IK constraints. For a trajectory of length T , the probability of the trajectory might violate the constraint is $p_{\text{error}} = 1 - (1 - p)^T$, and $\lim_{T \rightarrow \infty} p_{\text{error}} = 1$. As we show in our experiments, IK error contributes to most of the failure cases in end-effector pose trajectory diffusion.

Kinematics-Aware Diffusion. As an alternative to using IK solvers, the robot could be operated through joint position control. This approach provides direct and complete control of the robot. However, learning a trajectory diffusion model in the joint position space is challenging. In the case of end-effector pose diffusion models, we can impose accurate and strong constraints with the predicted next-best pose $\hat{\mathbf{a}}_{\text{pose}}^0(T)$. However, for an over-actuated 7-DoF robot arm, a 6-DoF end-effector pose $\hat{\mathbf{a}}_{\text{pose}}^0(T)$ might have an infinite number of corresponding joint positions $\hat{\mathbf{a}}_{\text{joint}}^0(T)$, which makes it difficult to perform inpainting for joint position diffusion. As we show in our experiments, the naive joint position diffusion model tends to be less accurate for goal-conditioned control, especially for the end poses.

To tackle this issue, we introduce Robot Kinematics Diffuser (RK-Diffuser). Similar to Xian et al. [40], RK-Diffuser learns an end-effector pose diffusion model $p_{\theta}(\mathbf{a}_{\text{pose}}^{k-1} | \mathbf{a}_{\text{pose}}^k, C_{\text{pose}})$ which generates accurate but less reliable end-effector pose trajectories. RK-Diffuser further learns an additional joint position diffusion model

$$\begin{aligned} & p_{\phi}(\mathbf{a}_{\text{joint}}^{k-1} | \mathbf{a}_{\text{joint}}^k, C_{\text{pose}}) \\ &= \mathcal{N}(\mathbf{a}_{\text{joint}}^{k-1}; \boldsymbol{\mu}_{\theta}(\mathbf{a}_{\text{joint}}^k, C_{\text{pose}}, k), \boldsymbol{\Sigma}_{\theta}(\mathbf{a}_{\text{joint}}^k, C_{\text{pose}}, k)) \end{aligned} \quad (9)$$

where we use the same set of conditional variables C_{pose} for conditional generation, but for inpainting, we only fix the initial joint action $\mathbf{a}_{\text{joint}}^0(0)$.

For action trajectories sampled from each learned policy, $\mathbf{a}_{\text{pose}}^0 \sim p_{\theta}(\mathbf{a}_{\text{pose}}^0 | \mathbf{a}_{\text{pose}}^1, C_{\text{pose}})$ and $\mathbf{a}_{\text{joint}}^0 \sim p_{\phi}(\mathbf{a}_{\text{joint}}^0 | \mathbf{a}_{\text{joint}}^1, C_{\text{pose}})$, we can build such a mapping by treating

the differentiable robot kinematics model f_K as a function $\hat{\mathbf{a}}_{\text{pose}}^0 = f_K(\mathbf{a}_{\text{joint}}^0)$. During inference, initialised with a near-optimal solution $\mathbf{a}_{\text{joint}}^0$, we can optimise the joint positions $\mathbf{a}_{\text{joint}}^0$ to predict end-effector poses $\hat{\mathbf{a}}_{\text{pose}}^0$ that are close to $\mathbf{a}_{\text{pose}}^0$ using gradients

$$\mathbf{a}_{\text{joint}}^0 \leftarrow \mathbf{a}_{\text{joint}}^0 - \alpha \frac{\partial \|\mathbf{a}_{\text{pose}}^0 - \hat{\mathbf{a}}_{\text{pose}}^0\|}{\partial \mathbf{a}_{\text{joint}}^0}, \quad (10)$$

where α is the learning rate. This gives a trajectory $\mathbf{a}_{\text{joint}}^{0*}$ that does not violate the kinematics constraint of the robot while achieving a high accuracy for manipulation tasks.

Networks. The low-level RK-Diffuser takes as input the start pose, the end pose, the RGB-D image of the first step observation, a vector of the robot low-dimensional states, and the trajectory rank. For the RGB-D image, we first convert it to a point cloud in the world frame and extract the features with PointNet++ [29]; for the other vector features, we use 4 layers of MLP. For the temporal encoding network, we found a temporal Conv1D UNet used in Janner et al. [21] performs well and has no clear performance gap between the commonly adopted Transformer backbones.

Model Training. When training diffusion models, we aim to maximise the ELBO of the dataset (Eqn. 5). However, taking the predicted next-best poses from the high-level policy π_{high} is inefficient as the prediction might be sub-optimal and slow. To alleviate this issue, for each demonstration ξ , we construct sub-trajectories $\{\xi(i)\}_{i=1}^K$ by chunking the trajectory ξ with the detected keyframe indices $\{k_i\}_{i=1}^{K_{\xi}}$. Next, we relabel each keyframe as a sub-goal of the training trajectory. This aligns with the training of the high-level agent π_{high} , and in practice, π_{high} and π_{low} can be optimised at the same time. The relabeling idea also resembles the Hindsight Experience Replay [2], which has been shown to be effective in learning hierarchical policy learning [24, 28]. Specifically, we have

$$\begin{aligned} \mathcal{L}_{\text{low}} &= -\beta_1 \mathcal{L}_{\text{pose}} - \beta_2 \mathcal{L}_{\text{joint}} - \beta_3 \mathcal{L}_{\text{joint} \rightarrow \text{pose}} \\ \mathcal{L}_{\text{pose}} &= \mathbb{E}_{q, \xi(i) \sim \mathcal{D}} \left[\log \frac{p_{\theta}(\mathbf{a}_{\text{pose}}^{0:K} | \xi(i))}{q(\mathbf{a}_{\text{pose}}^{1:K} | \mathbf{a}_{\text{pose}}^0, \xi(i))} \right] \\ \mathcal{L}_{\text{joint}} &= \mathbb{E}_{q, \xi(i) \sim \mathcal{D}} \left[\log \frac{p_{\phi}(\mathbf{a}_{\text{joint}}^{0:K} | \xi(i))}{q(\mathbf{a}_{\text{joint}}^{1:K} | \mathbf{a}_{\text{joint}}^0, \xi(i))} \right] \\ \mathcal{L}_{\text{joint} \rightarrow \text{pose}} &= \mathbb{E}_{q, \xi(i) \sim \mathcal{D}} \left[\log \frac{p_{\phi}(\mathbf{a}_{\text{pose}}^{0:K} | \xi(i))}{q(\mathbf{a}_{\text{pose}}^{1:K} | \mathbf{a}_{\text{pose}}^0, \xi(i))} \right], \end{aligned} \quad (11)$$

where β_1 , β_2 , and β_3 are weighting parameters and $\xi(i)$ is a sub-trajectory sampled from the dataset with start and end relabeled to two nearby keyframes. In particular, $\mathcal{L}_{\text{joint} \rightarrow \text{pose}}$ is made possible by predicting the end-

effector poses from joint positions via differentiable kinematics $\hat{\mathbf{a}}_{\text{pose}}^{0:K} = f_K(\mathbf{a}_{\text{joint}}^{0:K})$. This allows us to train a joint-position trajectory which better regularizes the joint positions with the kinematics as an inductive bias.

Trajectory Ranking. During training, most of the manipulation algorithms use sampling-based motion planners whose trajectories might be sub-optimal. In RK-Diffuser, we propose to add an additional conditional variable for each sub-trajectory, a trajectory rank $r_\xi = \frac{d_{\text{Euclidean}}}{d_{\text{travel}}}$, where $d_{\text{Euclidean}}$ is the Euclidean distance between the start and end pose and d_{travel} is the travelled distance between the start and end pose. Intuitively, an optimal trajectory, ignoring the kinematics constraint of the robot, should have $r_\xi = 1$. To encourage RK-Diffuser to generate near-optimal trajectories, we set $r_\xi = 1$ during inference. An analysis of the influence of trajectory ranking is in the appendix.

4.4. Practical Implementation Choices

For the high-level agent π_{high} , different from the past work [15, 34], we ignore the $a_{\text{collision}}$, which is a binary variable used to indicate whether the motion planner should perform collision avoidance because the low-level RK-Diffuser is trained to generate collision-aware optimal trajectories. For the low-level agent, different from most of the diffusion models that learn to predict a noise prediction model and learn to reconstruct the noise during the denoising steps, we follow Ramesh et al. [31] and observe that empirically directly predicting the original actions $\mathbf{a}_{\text{pose}}^0$ and $\mathbf{a}_{\text{joint}}^0$ is giving better performance. Besides, when truncated by the keyframe indices, the sub-trajectories might have different lengths. To tackle this issue, we resample each trajectory into a length of 64 for batched training. More implementations and discussions are in the appendix.

5. Experiments

In our experiments we show the following: **(1)** HDP outperforms the state-of-the-art methods across all RL Bench tasks; **(2)** in general, hierarchical agents outperform simple low-level continuous control policies; and **(3)** task-aware planning is important for many manipulation tasks, in particular those involving articulated objects.

In addition to this, we perform a series of ablation studies and show: **(1)** IK errors contribute to the majority of the failure cases of end-effector pose diffusion policy; **(2)** joint position diffusion is less accurate without the access to last joint position inpainting; and **(3)** 3D information and the corresponding feature extraction module are critical to the performance of RK-Diffuser.

Finally, we show HDP is capable of solving challenging real-world tasks efficiently and effectively on an open oven task with only 20 demonstrations.

For all simulation experiments, we use 100 demonstrations from RL Bench [19] for each task and train for 100K




Figure 4. Trajectory visualisations of the open box task.

iterations. On a real robot, we show HDP can learn efficiently and effectively with only 20 demonstrations.

5.1. Trajectory Visualisations

Firstly, we aim to understand why learning a low-level controller is necessary. In Fig. 4, we visualise the trajectory of an open box task in RL Bench. RRT learns a trajectory that correctly reaches the goal pose. Nevertheless, without understanding the task context, the trajectory generated by RRT will cause the lid of the box to fall from the gripper. To visualise the joint position trajectories of both the vanilla joint position diffusion policy and RK-Diffuser, we further predict the end-effector poses from the joint positions. Although the joint position diffusion policy understands the task context, without direct inpainting with the next-best joint position, the trajectory will be less accurate. RK-Diffuser distills the accurate end-effector poses to the joint positions via differentiable kinematics, which achieves both high prediction accuracy and kinematic awareness.

5.2. Simulation Experiments

We aim to compare HDP against (1) the state-of-the-art low-level control behaviour cloning agents, including ACT [44] and the vanilla Diffusion Policy [5]; (2) the high-level next-best-pose agent with a fixed local planner, PerAct. In addition, we aim to demonstrate the benefit of the proposed RK-Diffuser against alternatives, including: (1) **Planner**: a hybrid planner of fixed linear paths and standard RRT, which is the default setup used in RL Bench; (2) **Planner + Bezier**: in which an additional head is added to the PerAct backbone with a discrete output trained to choose the most appropriate trajectory generation method at each episode step, akin to Learned Path Ranking (LPR) [15] in the behaviour cloning setting; (3) **Diffuser**: the vanilla Diffuser [21] framed as a goal-conditioned joint-position diffusion model. More details of the baseline algorithms are available in the appendix. We choose 11 RL Bench tasks ranging from simple context-unaware grasping tasks to challenging tasks that require interacting with articulated objects. We present the results in Tab. 1 and make the following observations.

HDP outperforms the state-of-the-art methods across RL Bench tasks. As shown in Tab. 1, HDP achieves an overall 80.2% success rate across 11 RL Bench tasks. In partic-

Table 1. Success Rates (%) on RL Bench Tasks. For **red** tasks, we expect no improvement of HDP over baselines; with **blue** tasks, we expect HDP to outperform many of the baselines.

| | reach target | take lid off saucepan | pick up cup | toilet seat up | open box | open door | open drawer | open grill | open microwave | open oven | knife on board | overall |
|---------------------------|--------------|-----------------------|-------------|----------------|-----------|-----------|-------------|------------|----------------|-----------|----------------|--------------|
| ACT | 50 | 45 | 46 | 6 | 12 | 5 | 26 | 1 | 11 | 0 | 0 | 18.36 |
| Diffusion Policy | 43 | 25 | 24 | 5 | 4 | 22 | 28 | 9 | 7 | 0 | 0 | 15.18 |
| PerAct + Planner | 100 | 100 | 86 | 0 | 0 | 64 | 68 | 54 | 32 | 0 | 76 | 57.72 |
| PerAct + Planner + Bezier | 96 | 100 | 72 | 80 | 8 | 48 | 84 | 76 | 20 | 4 | 36 | 56.73 |
| PerAct + Diffuser | 100 | 94 | 84 | 80 | 82 | 88 | 84 | 82 | 20 | 18 | 52 | 71.27 |
| HDP | 100 | 96 | 82 | 86 | 90 | 94 | 90 | 88 | 26 | 58 | 72 | 80.18 |

ular, we observe on simple tasks (**red**), that require no accurate trajectory control, most of the baselines have achieved a competent performance. However, when it comes to the more challenging tasks (**blue**), HDP maintains its performance while the baselines mostly fail, due to either a lack of understanding in the task context or inaccurate motion trajectory generation.

Hierarchical agents outperform simple low-level continuous control policies. Comparing ACT and the vanilla Diffusion Policy with hierarchical agents, we observe that hierarchical agents consistently outperform the former. Empirically, both ACT and the Diffusion Policy fail to accurately detect intermediate keyframes, such as the handle of a drawer or an oven. This error is amplified due to distribution shift, which is a common issue of behaviour cloning agents in long-horizon tasks. In contrast, the hierarchical agent, with PerAct at the high level, achieves better generalisation and simplifies the optimisation task of the low-level agent. When trained on a multi-task setting, both ACT and Diffusion Policy fail to manage different skills and generalise to unseen test examples. However, all algorithms achieve a low performance on the open microwave task. We observe that this task has a highly diverse final end-effector pose distribution, which causes the high-level policy to have a high variance and generate inaccurate next-best poses. This error is then propagated to the low-level agents. Further exploration of this issue is left for future study.

Learned low-level agents achieve better performance than motion planners. In particular, we note that even with accurate predictions of the next-best pose, a lack of task understanding by the planner often leads to trajectories deviating from the desired optimal trajectory. For instance, while PerAct + Planner achieves 0% success rate on the open box task it regularly succeeds in grasping the box lid. The predicted trajectory consistently exceeds the turning radius of the lid hinge, leading to the failure. This issue is exacerbated by strict kinematic limitations. For example, in the same task, PerAct + Planner + Bezier performs poorly because, unlike in the lift toilet seat task, the smooth opening curves, prompted by the additional head of PerAct, are kinematically infeasible. On the contrary, the learned trajectory

capture the task context as demonstrated by the data and result in superior performance on a greater number of tasks.

5.3. Ablation Studies

We perform ablation studies on the selected RL Bench tasks to further understand the proposed low-level agent, RK-Diffuser. Since the high-level agent has been well-studied in prior works [34], we swap it with an expert and only focus on the performance of the low-level agents. We present the results in Tab. 2.

Sampling-based motion planners might fail without understanding the task context. As a sampling-based planner, RRT achieves a strong performance on simple tasks that only require goal information. However, for tasks that require a fine-grained trajectory, e.g., toilet seat up, RRT fails completely. As shown in Sect. 5.1, we see that trajectories generated by RRT might easily violate the task constraints.² One could handcraft task-specific constraints, but it is not generalisable across tasks.

IK errors contribute to most of the failure cases of end-effector pose diffusion policy. The Pose Diffusion denotes learning a diffusion policy directly over the end-effector pose trajectories and generate robot controls by solving the inverse kinematics. We observe that although Pose Diffusion achieves strong performance on several tasks, e.g., open microwave, it suffers from an overall 24.55% IK error rate. Specifically, most of the IK errors are caused by invalid quaternions and contribute to 75% of its failure cases. In particular, the IK error rate increases as the control difficulty increases. This explains the importance of learning joint position trajectories, instead of end-effector poses.

Joint position diffusion is less accurate without the access to last joint position inpainting. As in Sect. 4.3, an end-effector pose will have multiple corresponding joint positions, and hence, it is infeasible for a joint position diffusion model to perform the last step inpainting. In our ablations, we show that it achieves a worse performance than RK-Diffuser, especially on challenging tasks, e.g., open oven.

²RLBench uses a hybrid motion planner of RRT and predefined linear paths by default. To reproduce the RRT result, we disable the linear path trajectories manually.

Table 2. Ablation Study: Success Rates (%) / IK Error Rates (%) of low-level agents with the ground-truth next-best poses. For **red** tasks, we expect no improvement of HDP over baselines; with **blue** tasks, we expect HDP to outperform many of the baselines.

| | reach target | take lid off saucepan | pick up cup | toilet seat up | open box | open door | open drawer | open grill | open microwave | open oven | knife on board | overall |
|-----------------|----------------|-----------------------|---------------|----------------|----------------|---------------|----------------|---------------|----------------|---------------|----------------|------------------|
| RRT | 100 / 0 | 100 / 0 | 95 / 0 | 0 / 0 | 0 / 0 | 0 / 0 | 0 / 0 | 0 / 0 | 0 / 0 | 0 / 0 | 0 / 0 | 26.82 / 0 |
| Pose Diffusion | 100 / 0 | 85 / 6 | 93 / 0 | 93 / 4 | 88 / 8 | 24 / 68 | 3 / 88 | 64 / 22 | 98 / 0 | 9 / 62 | 82 / 12 | 67.18 / 24.55 |
| Joint Diffusion | 100 / 0 | 100 / 0 | 91 / 0 | 95 / 0 | 100 / 0 | 74 / 0 | 15 / 0 | 62 / 0 | 75 / 0 | 13 / 0 | 85 / 0 | 73.64 / 0 |
| RKD-RGB | 100 / 0 | 96 / 0 | 78 / 0 | 40 / 0 | 98 / 0 | 94 / 0 | 78 / 0 | 36 / 0 | 80 / 0 | 0 / 0 | 94 / 0 | 72.18 / 0 |
| RKD-ResNet | 100 / 0 | 100 / 0 | 95 / 0 | 92 / 0 | 100 / 0 | 93 / 0 | 100 / 0 | 86 / 0 | 21 / 0 | 43 / 0 | 88 / 0 | 83.45 / 0 |
| RK-Diffuser | 100 / 0 | 100 / 0 | 98 / 0 | 100 / 0 | 100 / 0 | 95 / 0 | 100 / 0 | 90 / 0 | 88 / 0 | 75 / 0 | 94 / 0 | 94.55 / 0 |




Figure 5. Real-robot execution sequences. For both tasks, the robot needs to accurately predict the trajectories that understand the task context conditioned on languages. As appliances have high resistance force, a slight deviation from the expected trajectory would cause the robot to fail because of exceeding the joint torque limit.

3D information and the corresponding feature extraction module are critical to the performance of RK-Diffuser. As mentioned earlier in Sect. 4.3, RK-Diffuser uses a PointNet++ for point cloud feature extraction. For RKD-RGB, we discard the depth information and use a pretrained ResNet50 to extract the image features; for RKD-ResNet, we ablate using a ResNet to extract features from the RGB-D image. We observe that both achieve worse performance when compared to the original RK-Diffuser, which indicates that understanding the 3D environment is necessary for generalisable and accurate control. We believe there are alternative representations and leave it for future study.

5.4. Real Robot Experiment

We also conducted a real-world experiment on an opening oven task and a sorting objects into drawer task with a Franka Panda 7 DoF arm. We use 2 RealSense D415 cameras that captures the scene. For each sub-task we collect 10 demonstrations. Both tasks require the robot to accurately locate the target and control all its joints, especially the orientation of the wrist at every time step, otherwise, given the high resistance force of the oven, the arm will halt due to exceeding the joint torque limit. As a summary, HDP achieves 100% success rate for the opening oven task and 94% success rate for the sorting object into drawer task. Due to the nature of demo collection, we observe high variance in the demonstrated trajectories for the task. Intuitively, this leads to sub-optimal and highly diverse next-best pose predictions from the high-level agent, PerAct, some of which

are out of distribution for RK-Diffuser. Interestingly, however, there appears to be minimal impact on RK-Diffuser, and the method is still capable of generalising to these unseen poses and generating accurate trajectories. Detailed results are in the appendix and are best viewed via the supplementary video.

6. Conclusion

We present HDP, a hierarchical agent for kinematics-aware robotic manipulation. HDP factorises a policy: at the high-level, a task-planning agent predicts the next-best end-effector pose, and at the low-level, RK-Diffuser performs goal-conditioned prediction of a joint position trajectory that connects to the predicted next-best pose. To achieve both kinematics-awareness and high prediction accuracy, RK-Diffuser distills the accurate but less reliable end-effector pose trajectory to the joint position trajectory via differentiable kinematics. We show that HDP achieves state-of-the-art performance on a set of challenging RL-Bench manipulation tasks. On a real robot, HDP learns to solve both opening oven and sorting objects into drawer task. Although we have demonstrated some robustness of RK-Diffuser to out-of-distribution poses, the nature of behaviour cloning for longer-horizon tasks suggests that error accumulation could lead to significant distribution shifts and ultimate failure. Future works could explore improving the framework by designing more unified structures that minimises the compounding error.

References

- [1] Anurag Ajay, Yilun Du, Abhi Gupta, Joshua B. Tenenbaum, Tommi S. Jaakkola, and Pulkit Agrawal. Is conditional generative modeling all you need for decision making? In *The Eleventh International Conference on Learning Representations*, 2023.
- [2] Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. Hindsight experience replay. *Advances in neural information processing systems*, 30, 2017.
- [3] Siwei Chen, Xiao Ma, and Zhongwen Xu. Imitation learning as state matching via differentiable physics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7846–7855, 2023.
- [4] Siwei Chen, Yiqing Xu, Cunjun Yu, Linfeng Li, Xiao Ma, Zhongwen Xu, and David Hsu. Daxbench: Benchmarking deformable object manipulation with differentiable physics. In *The Eleventh International Conference on Learning Representations*, 2023.
- [5] Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin Burchfiel, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. In *Proceedings of Robotics: Science and Systems (RSS)*, 2023.
- [6] Theophile Gervet, Zhou Xian, Nikolaos Gkanatsios, and Katerina Fragkiadaki. Act3d: Infinite resolution action detection transformer for robotic manipulation. *arXiv preprint arXiv:2306.17817*, 2023.
- [7] Ankit Goyal, Jie Xu, Yijie Guo, Valts Blukis, Yu-Wei Chao, and Dieter Fox. Rvt: Robotic view transformer for 3d object manipulation. *arXiv preprint arXiv:2306.14896*, 2023.
- [8] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016.
- [9] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.
- [10] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [11] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022.
- [12] Yuanming Hu, Luke Anderson, Tzu-Mao Li, Qi Sun, Nathan Carr, Jonathan Ragan-Kelley, and Frédo Durand. DiffTaichi: Differentiable programming for physical simulation. *arXiv preprint arXiv:1910.00935*, 2019.
- [13] Zhiao Huang, Yuanming Hu, Tao Du, Siyuan Zhou, Hao Su, Joshua B Tenenbaum, and Chuang Gan. Plasticinelab: A soft-body manipulation benchmark with differentiable physics. *arXiv preprint arXiv:2104.03311*, 2021.
- [14] Andrew Jaegle, Sebastian Borgeaud, Jean-Baptiste Alayrac, Carl Doersch, Catalin Ionescu, David Ding, Skanda Koppala, Daniel Zoran, Andrew Brock, Evan Shelhamer, et al. Perceiver io: A general architecture for structured inputs & outputs. *arXiv preprint arXiv:2107.14795*, 2021.
- [15] Stephen James and Pieter Abbeel. Coarse-to-fine q-attention with learned path ranking. *arXiv preprint arXiv:2204.01571*, 2022.
- [16] Stephen James and Pieter Abbeel. Coarse-to-fine q-attention with tree expansion. *arXiv preprint arXiv:2204.12471*, 2022.
- [17] Stephen James and Andrew J Davison. Q-attention: Enabling efficient learning for vision-based robotic manipulation. *IEEE Robotics and Automation Letters*, 7(2):1612–1619, 2022.
- [18] Stephen James, Andrew J Davison, and Edward Johns. Transferring end-to-end visuomotor control from simulation to real world for a multi-stage task. In *Conference on Robot Learning*, pages 334–343. PMLR, 2017.
- [19] Stephen James, Zicong Ma, David Rovick Arrojo, and Andrew J Davison. Rlbench: The robot learning benchmark & learning environment. *IEEE Robotics and Automation Letters*, 5(2):3019–3026, 2020.
- [20] Stephen James, Kentaro Wada, Tristan Laidlow, and Andrew J Davison. Coarse-to-fine q-attention: Efficient learning for visual robotic manipulation via discretisation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13739–13748, 2022.
- [21] Michael Janner, Yilun Du, Joshua B Tenenbaum, and Sergey Levine. Planning with diffusion for flexible behavior synthesis. *arXiv preprint arXiv:2205.09991*, 2022.
- [22] Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, Ethan Holly, Mrinal Kalakrishnan, Vincent Vanhoucke, et al. Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation. *arXiv preprint arXiv:1806.10293*, 2018.
- [23] Bingyi Kang, Xiao Ma, Chao Du, Tianyu Pang, and Shuicheng Yan. Efficient diffusion policies for offline reinforcement learning. *arXiv preprint arXiv:2305.20081*, 2023.
- [24] Andrew Levy, George Konidaris, Robert Platt, and Kate Saenko. Learning multi-level hierarchies with hindsight. *arXiv preprint arXiv:1712.00948*, 2017.
- [25] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [26] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 35:5775–5787, 2022.
- [27] Jan Matas, Stephen James, and Andrew J Davison. Sim-to-real reinforcement learning for deformable object manipulation. In *Conference on Robot Learning*, pages 734–743. PMLR, 2018.
- [28] Hai Nguyen, Zhihan Yang, Andrea Baisero, Xiao Ma, Robert Platt, and Christopher Amato. Hierarchical reinforcement learning under mixed observability. In *International Workshop on the Algorithmic Foundations of Robotics*, pages 188–204. Springer, 2022.
- [29] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017.

- [30] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [31] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022.
- [32] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- [33] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- [34] Mohit Shridhar, Lucas Manuelli, and Dieter Fox. Perceiver-actor: A multi-task transformer for robotic manipulation. In *Conference on Robot Learning*, pages 785–799. PMLR, 2023.
- [35] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*, 2022.
- [36] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [37] Zhendong Wang, Jonathan J Hunt, and Mingyuan Zhou. Diffusion policies as an expressive policy class for offline reinforcement learning. In *The Eleventh International Conference on Learning Representations*, 2023.
- [38] Daniel Watson, William Chan, Ricardo Martin-Brualla, Jonathan Ho, Andrea Tagliasacchi, and Mohammad Norouzi. Novel view synthesis with diffusion models. *arXiv preprint arXiv:2210.04628*, 2022.
- [39] Yilin Wu, Wilson Yan, Thanard Kurutach, Lerrel Pinto, and Pieter Abbeel. Learning to manipulate deformable objects without demonstrations. *arXiv preprint arXiv:1910.13439*, 2019.
- [40] Zhou Xian, Nikolaos Gkanatsios, Theophile Gervet, and Katerina Fragkiadaki. Unifying diffusion models with action detection transformers for multi-task robotic manipulation. In *7th Annual Conference on Robot Learning*, 2023.
- [41] Jie Xu, Miles Macklin, Viktor Makoviychuk, Yashraj Narang, Animesh Garg, Fabio Ramos, and Wojciech Matusik. Accelerated policy learning with parallel differentiable simulation. In *International Conference on Learning Representations*, 2022.
- [42] Jie Xu, Sangwoon Kim, Tao Chen, Alberto Rodriguez Garcia, Pulkit Agrawal, Wojciech Matusik, and Shinjiro Sueda. Efficient tactile simulation with differentiability for robotic manipulation. In *Conference on Robot Learning*, pages 1488–1498. PMLR, 2023.
- [43] Mandi Zhao, Pieter Abbeel, and Stephen James. On the effectiveness of fine-tuning versus meta-reinforcement learning. *Advances in Neural Information Processing Systems*, 35:26519–26531, 2022.
- [44] Tony Z Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023.
- [45] Sheng Zhong, Thomas Power, and Ashwin Gupta. PyTorch Kinematics. 2023.