

# DISCO: Language-Guided Manipulation with Diffusion Policies and Constrained Inpainting

Ce Hao<sup>\*1</sup>, Kelvin Lin<sup>1</sup>, Zhiwei Xue<sup>1</sup>, Siyuan Luo<sup>1</sup>, Harold Soh<sup>\*1,2</sup>

**Abstract**—Diffusion policies have demonstrated strong performance in generative modeling, making them promising for robotic manipulation guided by natural language instructions. However, generalizing language-conditioned diffusion policies to open-vocabulary instructions in everyday scenarios remains challenging due to the scarcity and cost of robot demonstration datasets. To address this, we propose DISCO, a framework that leverages off-the-shelf vision-language models (VLMs) to bridge natural language understanding with high-performance diffusion policies. DISCO translates linguistic task descriptions into actionable 3D keyframes using VLMs, which then guide the diffusion process through constrained inpainting. However, enforcing strict adherence to these keyframes can degrade performance when the VLM-generated keyframes are inaccurate. To mitigate this, we introduce an inpainting optimization strategy that balances keyframe adherence with learned motion priors from training data. Experimental results in both simulated and real-world settings demonstrate that DISCO outperforms conventional fine-tuned language-conditioned policies, achieving superior generalization in zero-shot, open-vocabulary manipulation tasks. Videos see website: [sites.google.com/view/disco2025](https://sites.google.com/view/disco2025).

## I. INTRODUCTION

Robotic manipulation remains a fundamental challenge due to the complexity of handling diverse objects across varying contexts. Recent advancements in diffusion-based policies have significantly improved robot capabilities, enabling the successful execution of multi-modal, long-horizon, and multi-object tasks [1]–[3]. A key next step is language-conditioned control of diffusion policies, allowing robots to follow natural language instructions for more flexible and adaptable behavior [4], [5]. By integrating language as a control mechanism, an otherwise unconditional diffusion policy can be adapted to specific task contexts, paving the way for more generalizable and interactive robot manipulation.

Existing approaches directly input language instructions into diffusion policies to guide action generation [6]–[8]. However, fine-tuned language-conditioned policies often struggle with unseen tasks and open-vocabulary instructions [6], [9], [10]. This limitation becomes evident in tasks like object grasping, where even policies trained on large multi-modal datasets frequently fail to execute correct grasps when faced with novel natural language instructions [11], [12].

In this paper, we propose a framework to integrate general-purpose “off-the-shelf” vision-language models (VLMs) and diffusion policy for open-vocabulary robot manipulation,

This research is supported by A\*STAR under its National Robotics Programme (NRP) (Award M23NBK0053).

<sup>1</sup> School of Computing, National University of Singapore.

<sup>2</sup> Smart Systems Institute, NUS.

\*correspondence to [cehao@u.nus.edu](mailto:cehao@u.nus.edu) and [harold@nus.edu.sg](mailto:harold@nus.edu.sg)

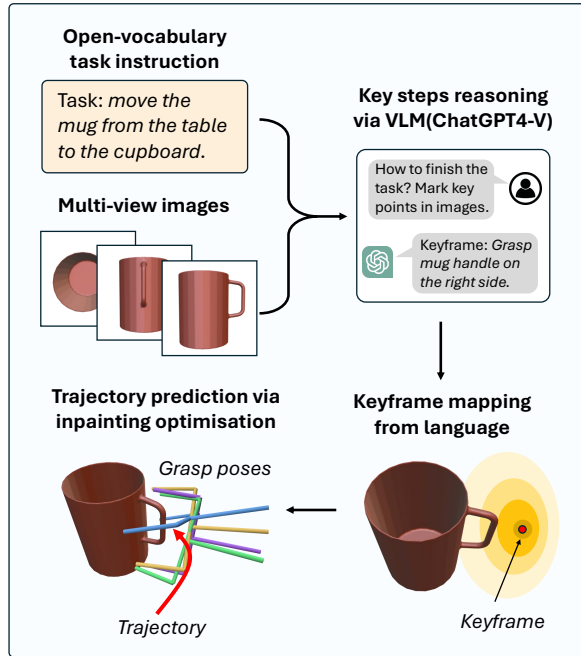


Fig. 1: An overview of the DISCO framework, which interprets task descriptions and multi-view observations to generate keyframes in 3D space, which are used to guide the diffusion-based generation of actions (e.g., grasp poses) through constrained inpainting optimization.

which we call **DISCO** (**D**iffusion **I**npainting with **S**emantic **K**eyframes and **C**onstrained **O**ptimization) As in Figure. 1, our key idea is that by integrating VLM-generated *keyframes* as actionable guides in a diffusion process, we can bridge the gap between high-level linguistic instructions and low-level robot actions. Different from prior work [11], the keyframes are *not* precise subgoals but serve as “coarse” guidance. Here, we use a general VLM model (e.g., ChatGPT4) to convert task descriptions into keyframes in 3D space. Then, we employ an inpainting method that uses these keyframes to guide the policy’s action sequence generation. Since the keyframes serve only as rough guides, standard inpainting [13], [14] — which enforces adherence to the keyframes — can fail to generate successful trajectories. This occurs since keyframes can be inaccurate or far outside the support of the training distribution [15]. To further enhance robustness, we propose an inpainting optimization method that constrains the generated points to lie within regions of sufficient support during diffusion. This approach balances keyframe adherence and remaining within the support of the diffusion model, which results in better performance for novel

or ambiguous task descriptions.

We validate our open-vocabulary diffusion policy, DISCO, using simulated benchmarks [16], [17] and a real-robot grasping testbed [18]. The simulation experiments include adapted language-guided block pushing, object grasping, and long-horizon tasks such as the Franka Kitchen and multi-task Calvin benchmark. These tasks require generating continuous actions, including end-effector positions, velocities, and grasp poses. Compared to fine-tuned diffusion policy baselines, DISCO achieves comparable performance on seen tasks while significantly outperforming them on tasks with novel language instructions. Additionally, we conduct an ablation study comparing the use of curves generated by vision-language models (VLMs) as intermediate representations with the keyframes used in DISCO [19]. Finally, we evaluate zero-shot transfer to real-world language-guided grasping, where DISCO successfully executes grasping in real-time, manipulating various objects under human instructions without fine-tuning.

In summary, the contribution of this paper is threefold:

- 1) We proposed the DISCO framework that integrates off-the-shelf VLMs to diffusion policy via keyframes for language-guided robot manipulation.
- 2) We designed the inpainting optimization method to circumvent inaccurate and out-of-distribution keyframes generated by open-vocabulary instructions.
- 3) We conducted extensive experiments in both simulation and real-world robots, showcasing DISCO’s performance on continuous action and grasping pose generation over baselines.

## II. RELATED WORKS

**Language-conditioned robot manipulation**, where natural language is used to guide robot behavior, has been receiving increased attention, partially due to the advances made in large-language models. One line of prior work has applied vision-language models (VLMs) to convert linguistic descriptions into executable trajectories [4], [6], [20], [21]; however, direct trajectory generation with general-purpose VLMs often results in unreliable or infeasible motions, especially in open-vocabulary or unstructured settings. Other methods use VLMs to transform language descriptions into code, affordances, or actionable trajectory tokens [10], [11], [22], [23], but affordance approaches typically require task-specific fine-tuning and do not naturally generalize to multi-view 3D reasoning, while code-generation methods do not provide interpretable intermediate signals for action guidance. To address these challenges, our method introduces keyframes as a more reliable and generalizable intermediate representation to bridge high-level language commands and low-level trajectories.

**Diffusion-based policies** are a rich policy class that can represent multi-modal and high-dimensional action distributions. They have been shown to be excellent for imitation learning, where they can be trained to exhibit intricate and complex motion [1]–[3], [18]. Recent advancements have expanded diffusion policy applications to include goal-conditioned

and language-conditioned tasks [7], [8], [24], providing a structured approach for integrating explicit objectives and linguistic instructions into action generation [4], [25]. However, these methods typically require extensive fine-tuning for each specific task, limiting their ability to generalize to unseen, open-vocabulary tasks [26], [27]. To tackle this, we propose a zero-shot transfer strategy using inpainting, complemented by constrained optimization to handle out-of-distribution keyframes resulting from novel task descriptions.

**Inpainting** was originally developed for image completion [28], but has been adapted for use with diffusion policies to generate actions rather than images [13], [29]. These methods have also been applied to goal-conditioned planning, using goals as known conditions to guide action generation [3], [14], [30]. In our approach, we apply inpainting within the diffusion policy framework to generate trajectories and grasp poses, employing keyframes as guidance for action sequences.

## III. PRELIMINARIES

### A. Problem Formulation

We define the language-conditioned manipulation task as controlling the robot to manipulate objects to achieve a target state under a language specification (or task description)  $l$ . The language specifications are “open-vocabulary” in that they contain words and phrases in natural language, beyond a small predefined set of labels. The behavior of the robot should comply with the specification. In this work, we have access to image observations  $O$  of the robot and objects. The set of observations is assumed to be sufficient to complete the task; for example, multiple views may be available to provide sufficient state information. We seek a policy  $\pi$  that, at each time step  $t$ , generates an action (sequence)  $\mathbf{a}$  conditioned on the current observation  $O_t$  and language description  $l$ , i.e.,  $\mathbf{a}_t \sim \pi(\mathbf{a} | O_t, l)$ . The task is deemed successful when the target is reached and the robot trajectories satisfy the language specifications. In the following, we will drop the explicit dependence on the time step to avoid clutter.

### B. Diffusion Policy

Diffusion Policy [1], [8], [31] applies diffusion-based generative models towards robot action generation. In brief, the model assumes that  $p_\theta(\mathbf{x}_0) := \int p_\theta(\mathbf{x}_{0:N}) d\mathbf{x}_{1:N}$ , where  $\mathbf{x}_1, \dots, \mathbf{x}_N$  are latent variables of the same dimensionality as the data  $\mathbf{x}_0 \sim p(\mathbf{x}_0)$ . A forward diffusion chain gradually adds noise to the data  $\mathbf{x}_0 \sim q(\mathbf{x}_0)$  in  $N$  steps with a predefined variance schedule  $\beta_i$ , as,

$$q(\mathbf{x}_{1:N} | \mathbf{x}_0) := \prod_{t=1}^N q(\mathbf{x}_t | \mathbf{x}_{t-1}), \quad (1)$$

$$q(\mathbf{x}_i | \mathbf{x}_{i-1}) := \mathcal{N}(\mathbf{x}_i; \sqrt{1 - \beta_i} \mathbf{x}_{i-1}, \beta_i \mathbf{I}). \quad (2)$$

In DDPM [29], the reverse process is modeled as a Gaussian distribution, and the reverse diffusion chain is constructed as,

$$p_\theta(\mathbf{x}_{0:N}) := \mathcal{N}(\mathbf{x}_N; \mathbf{0}, \mathbf{I}) \prod_{t=1}^N p_\theta(\mathbf{x}_{i-1} | \mathbf{x}_i) \quad (3)$$

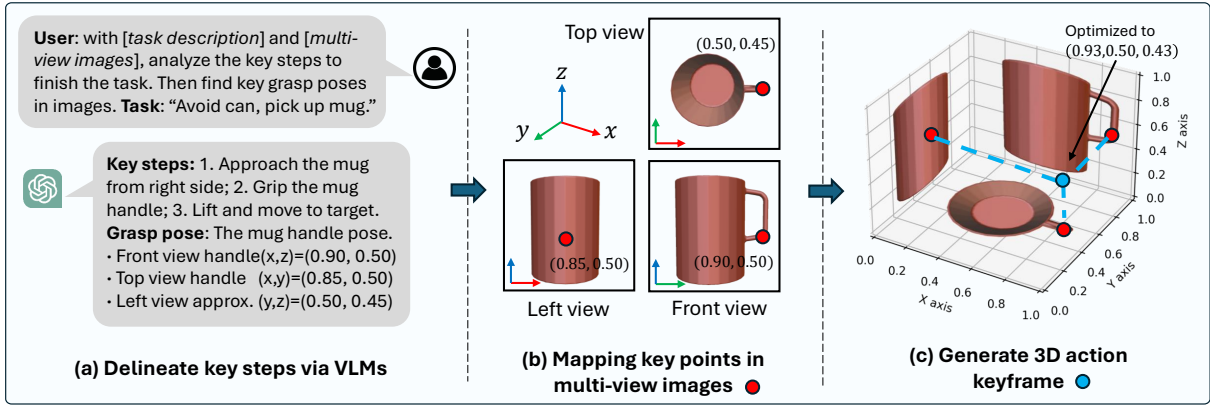


Fig. 2: DISCO predicts keyframes by (a) delineating the key steps to finish the task; (b) mapping key point positions in images; and (c) generating the 3D action keyframe based on multi-view key points.

$$p_{\theta}(\mathbf{x}_{i-1} | \mathbf{x}_i) = \mathcal{N}(\mathbf{x}_{i-1}; \mu_{\theta}(\mathbf{x}_i, i), \Sigma_{\theta}(\mathbf{x}_i, i)), \quad (4)$$

which is then optimized by maximizing the evidence lower bound defined as  $\mathbb{E}_q \left[ \ln \frac{p_{\theta}(\mathbf{x}_{0:N})}{q(\mathbf{x}_{1:N} | \mathbf{x}_0)} \right]$ . After training, sampling from the diffusion model consists of sampling  $\mathbf{x}_N \sim p(\mathbf{x}_N)$  and running the reverse diffusion chain to go from  $i = N$  to 0. The diffusion policy is formulated as the reverse process of a conditional diffusion model to generate actions,

$$\begin{aligned} \pi_{\theta}(\mathbf{a} | O) &= p_{\theta}(\mathbf{a}_{0:N} | O) \\ &= \mathcal{N}(\mathbf{a}_N; 0, I) \prod_{i=1}^N p_{\theta}(\mathbf{a}_{i-1} | \mathbf{a}_i, O), \end{aligned} \quad (5)$$

where the end sample of the reverse chain,  $\mathbf{a}_0$ , is the action used for execution.

### C. Inpainting

Inpainting [3], [28], [32] was originally formulated for image generation, where part of an image is known  $\mathbf{x}^{\text{known}}$  and the other part is “masked” (or unknown)  $\mathbf{x}^{\text{unknown}}$ . Standard inpainting integrates the  $\mathbf{x}^{\text{known}}$  and  $\mathbf{x}^{\text{unknown}}$  during the reverse process to reconstruct the missing regions [13],

$$\mathbf{x}_{i-1}^{\text{known}} \sim p_{i-1}^{\text{known}} = \mathcal{N}(\sqrt{\bar{\alpha}_i} \mathbf{x}_0^{\text{known}}, (1 - \bar{\alpha}_i) I), \quad (6)$$

$$\mathbf{x}_{i-1}^{\text{unknown}} \sim q_{i-1}^{\text{unknown}} = \mathcal{N}(\mu_{\theta}(\mathbf{x}_i, i), \Sigma_{\theta}(\mathbf{x}_i, i)), \quad (7)$$

$$\mathbf{x}_{i-1} = \mathbf{m} \odot \mathbf{x}_{i-1}^{\text{known}} + (1 - \mathbf{m}) \odot \mathbf{x}_{i-1}^{\text{unknown}}. \quad (8)$$

In Eqn. (6) and (7),  $\mathbf{x}_{i-1}^{\text{known}}$  and  $\mathbf{x}_{i-1}^{\text{unknown}}$  denote the known part of the image and the unknown inpainting parts, sampled from the probabilities  $p_{i-1}^{\text{known}}$  and  $q_{i-1}^{\text{unknown}}$  at  $(i - 1)$ th step.  $\bar{\alpha}_i = \prod_{s=1}^i (1 - \beta_s)$ , where  $\beta_s$  controls the variance of the Gaussian noise added at each forward step.  $\mu_{\theta}$  and  $\Sigma_{\theta}$  represent the scheduling parameters and parameters of the learned diffusion distribution [13], [29].  $\mathbf{m}$  represents the inpainting mask, which takes values of 1 for known regions and 0 for unknown regions;  $\odot$  denotes the element-wise production. During the reverse diffusion,  $\mathbf{x}_{i-1}$  is constructed by merging  $\mathbf{x}_{i-1}^{\text{known}}$  and  $\mathbf{x}_{i-1}^{\text{unknown}}$ . This integration ensures the content generated for the unknown regions is conditioned on the known regions, leading to a coherent output.

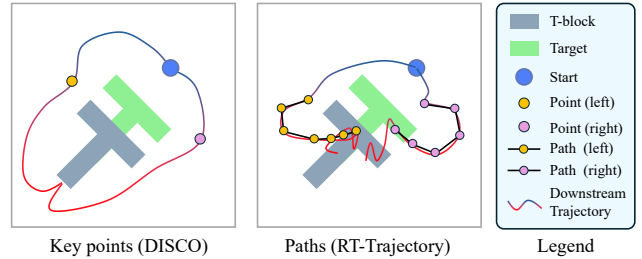


Fig. 3: Comparison of VLM-generated points versus paths on the Push-T task. The agent is instructed to push the block from the left or right side. The figures depict the generated points/paths by ChatGPT4o and the resulting trajectories using DISCO (for points) and RT-Trajectory [19] (for paths). We found ChatGPT4o generated key points more reliably than paths.

## IV. METHOD: DIFFUSION INPAINTING WITH SEMANTIC KEYFRAMES AND CONSTRAINED OPTIMIZATION

This section describes DISCO, our contribution towards zero-shot language-conditioned diffusion. We introduce the keyframe generation method (Sec. IV-A) followed by the constrained inpainting optimization (Sec. IV-B) to address issues that arise due to poor or invalid keyframes generated by open-vocabulary descriptions. DISCO is summarized in **Algorithm 1**.

### A. Keyframe Generation from Language Descriptions

As illustrated in Figure 2, we employ general multi-modal vision-language models (VLMs), such as ChatGPT4, to analyze the task description  $l$  and observation  $O$ , and generate action keyframes using three main steps:

- 1) **Delineating the key steps** necessary to complete the task [20], [21]. For example, given the task “Avoid the can and pick up the mug”, the VLM generates two key steps: “First move to the mug from the right side of the can” then “Grasp the mug handle and lift it.”
- 2) **Mapping key points**  $\mathbf{p}^{\text{key}}$  on the observations  $O$ . We ask the VLMs to map corresponding key points  $\mathbf{p}^{\text{key}}$  based on the above steps by marking the relevant points on the images  $O$ . For example, grasping the mug handle would involve marking the handle region. We use points rather

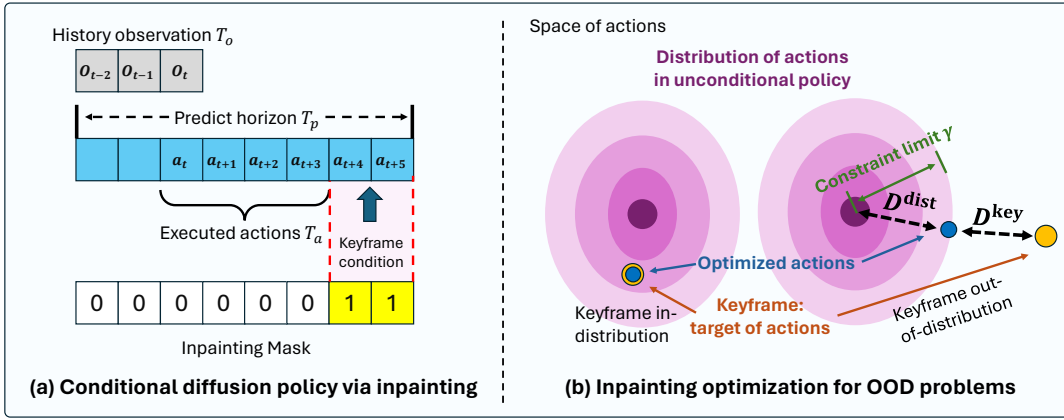


Fig. 4: (a): Inpainting controls the trajectory generation of diffusion policy via keyframes. (b): Our inpainting optimization method constrains the generated actions to lie within regions of sufficient support (high likelihood).  $D^{key}$  denotes the distance of masked action to keyframe in Eqn. (12);  $D^{dist}$  represents the negative log-likelihood constraint in Eqn. (13).

than paths because preliminary tests indicated that points were generated more robustly by current off-the-shelf VLMs (see illustrative example in Fig. 3). DISCO can potentially be used on paths, which would be interesting future work. Moreover, note that the points marked by VLMs are not exact grasp points (unlike [11], [33], [34]) and are not used to directly plan manipulation trajectories. Instead, they serve to provide “coarse” information that guides trajectory generation.

- 3) **Generating action keyframes  $\mathbf{a}^{key}$**  in the environment. Given the key points of the end-effector, we convert them to the keyframes  $\mathbf{a}^{key}$  based on the policy and environment. Depending on the exact application,  $\mathbf{a}^{key}$  could represent the 3D position, or be converted to end-effector velocity or joint configuration. In practice, we first calibrate a mapping from image keypoints to 3D positions, then use the robot’s current state and inverse kinematics to obtain the actionable keyframe.

### B. Conditional Diffusion Policy via Inpainting

After generating the keyframes  $\mathbf{a}^{key}$ , we employ inpainting to perform keyframe-conditioned diffusion-based action generation  $\mathbf{a} \sim \pi(\mathbf{a}|O, \mathbf{a}^{key}) = \mathcal{N}(\mathbf{a}_N; 0, I) \prod_{i=1}^N p_\theta(\mathbf{a}_{i-1} | \mathbf{a}_i, O, \mathbf{a}^{key})$  (from Eqn. (5)), where inpainting of the conditional diffusion model is formulated as,

$$\mathbf{a}_{i-1} = \mathbf{m} \odot \mathbf{a}_{i-1}^{known} + (1 - \mathbf{m}) \odot \mathbf{a}_{i-1}^{unknown} \quad (9)$$

$$\mathbf{a}_{i-1}^{known} \sim p_{i-1}^{known} = \mathcal{N}(\sqrt{\bar{\alpha}_i} \mathbf{a}_0^{known}, (1 - \bar{\alpha}) I) \quad (10)$$

$$\mathbf{a}_{i-1}^{unknown} \sim q_{i-1}^{unknown} = \mathcal{N}(\mu_\theta(\mathbf{a}_i, O, i), \Sigma_\theta(\mathbf{a}_i, O, i)) \quad (11)$$

The diffusion policy is executed in a receding-horizon control fashion [1]. As illustrated in Figure 4(a), at time  $t$ , the diffusion policy utilizes historical observations  $O_t$  with length  $T_o$  to generate a predicted multi-step action sequence  $\mathbf{a}_t$  with length  $T_p$ , and executes the first  $T_a$  actions. At time  $t + T_a$ , the policy enters the next cycle, repeating the observation and trajectory generation until the task is completed. During the generation of the action sequence, inpainting is applied by setting the known part as the keyframes, i.e.,  $\mathbf{a}_0^{known} = \mathbf{a}^{key} \cdot \mathbf{m}$ .

### Algorithm 1: DISCO for Manipulation

- 1 **Inputs:** Environment, language description  $l$ , observation  $O_0$ , VLM with prompts, diffusion policy model ( $\mu_\theta$  and  $\Sigma_\theta$ );
- 2 Reasoning key points for task:  $\mathbf{p}^{key} \leftarrow \text{VLM}(O_0, l)$ ;
- 3 Generate action keyframe:  $\mathbf{a}^{key} \leftarrow \text{Env.}(\mathbf{p}^{key})$ ;
- 4 Formulate initial history observation by padding  $O = [0, 0, \dots, O_0]$ , and set  $t = 0$ ;
- 5 **while** task is **NOT Done** **do**
- 6      $\mathbf{a}_N \sim \mathcal{N}(0, I)$ ;
- 7     **for**  $i$  in  $\{N, \dots, 1\}$  **do**
- 8         Perform reversed diffusion to calculate  $\mu_\theta(\mathbf{a}_i, O, i), \Sigma_\theta(\mathbf{a}_i, O, i)$ ;
- 9         Perform inpainting optimization  $\mathbf{a}_{i-1}$  in Eqn. (14) and (10), given  $\mathbf{a}_0^{known} = \mathbf{a}^{key}$ ;
- 10     **end for**
- 11     Obtain  $\mathbf{a}_0$ , execute  $\mathbf{a}_0[t : t + T_a - 1]$  in Env.;
- 12     Update time step,  $t \leftarrow t + T_a$ ;
- 13     Collect new observation  $O = [O_{t-T_o+1}, \dots, O_t]$ ;
- 14 **end while**

In this setup, the mask  $\mathbf{m}$  for the first  $T_o + T_a - 1$  steps is set to 0, and subsequent steps are set to 1. Thus,  $\mathbf{a}^{key}$  serves to condition the generation of  $\mathbf{a}^{unknown}$ . When one task contains multiple keyframes, DISCO selects the first keyframe and sequentially switches to the following ones when the current keyframe is achieved (two consecutive actions remain almost unchanged). For grasp pose generation, where only a single action is executed, inpainting optimization is applied by fixing the 3D position from the keyframe and sampling the rotation component with the diffusion model, ensuring the final 6D grasp pose is both conditioned and feasible.

### C. Inpainting Optimization for Out-of-Distribution Keyframes

Vanilla inpainting can produce unsuccessful robot behavior due to “poor” keyframes — it forces the generation to comply with the keyframe which leads to suboptimal  $\mathbf{a}^{unknown}$ . This results in the robot state moving towards regions that are not

well supported by the training data and consequently, and the produced behavior fails to accomplish the task.

To address this issue, we propose a constrained inpainting optimization approach. Intuitively, we adjust the generation process to ensure that when faced with poor keyframes, the generated action sequences remain within the data distribution. This approach helps sustain task performance even under the challenging conditions introduced by novel task descriptions. More formally, during the reserve step of the diffusion model, we replace Eqn. (9) with,

$$\mathbf{a}_{i-1} = \arg \min_{\mathbf{a}_{i-1}} D^{\text{key}}(\mathbf{m} \odot \mathbf{a}_{i-1}^{\text{known}}, \mathbf{m} \odot \mathbf{a}_{i-1}), \quad (12)$$

$$\text{s.t. } -\log(q_{i-1}^{\text{unknown}}(\mathbf{a}_{i-1})) \leq \gamma_{i-1}, \quad (13)$$

where,  $\mathbf{a}_{i-1}^{\text{known}}$  denotes the known part in Eqn. (10) and  $q_{i-1}^{\text{unknown}}$  is density in Eqn. (11).  $D^{\text{key}}$  measures the distance between the generated  $\mathbf{a}_{i-1}$  and the  $\mathbf{a}_{i-1}^{\text{known}}$  (keyframe  $\mathbf{a}^{\text{key}}$ ). We design the distance function  $D^{\text{key}}$  according to the policy and task environment. We constrain the negative log-likelihood  $-\log(q_{i-1}^{\text{unknown}}(\mathbf{a}_{i-1}))$  to be bounded by hyperparameter  $\gamma_i$ .

Figure 4(b) illustrates inpainting optimization: the left diagram shows that the keyframe is within a high-likelihood region of  $p(\mathbf{a}_{i-1}^{\text{known}})$ , i.e.,  $-\log(q_{i-1}^{\text{unknown}}(\mathbf{a}_{i-1})) \leq \gamma_{i-1}$ . In this scenario, inequality Eqn. (13) is always satisfied, and thus the formula reduces to vanilla inpainting (Eqn. (9)). In the right diagram, since  $-\log(q_{i-1}^{\text{unknown}}(\mathbf{a}_{i-1})) \geq \gamma_{i-1}$ , constrained optimization via Eqn (13) results in a generated action sequence that lies within the training data distribution but also minimizes the distance to the keyframe. We find this enables the policy to complete the task while maximally satisfying the language specification.

The distance function  $D^{\text{key}}$  and  $\gamma_i$  are application-dependent. In our experiments, we set  $D^{\text{key}}$  as the L-2 norm distance function and  $q_{i-1}^{\text{unknown}}$  is Gaussian distribution. As such, we can simplify general constrained inpainting to the convex optimization problem,

$$\mathbf{a}_{i-1} = \operatorname{argmin}_{\mathbf{a}_{i-1}} \sum_{k, \mathbf{m}^{[k]=1} } \left( \mathbf{a}_{i-1}^{[k]} - \mathbf{a}_{i-1}^{\text{known}, [k]} \right)^2, \quad (14)$$

$$\text{s.t. } \frac{1}{2} [\mathbf{a}_{i-1} - \mu_{\theta}(\mathbf{a}_i, \hat{i})]^T \cdot \Sigma_{\theta}^{-1}(\mathbf{a}_i, O, \hat{i}) \cdot [\mathbf{a}_{i-1} - \mu_{\theta}(\mathbf{a}_i, O, \hat{i})] \leq \gamma'_{i-1}, \quad (15)$$

$$\gamma'_i = \gamma_i - \frac{1}{2} \log[(2\pi)^d \Sigma_{\theta}(\mathbf{a}_i, O, \hat{i})], \quad (16)$$

where superscript  $[k]$  denotes the  $k_{\text{th}}$  element of the vector.  $d$  denotes the dimension of the actions.  $\gamma'_i$  is the normalized constraint based on  $\gamma_i$  and the covariance matrix (see the Appendix for detailed ablation study of  $\gamma_i$ ).

## V. EXPERIMENT

In this section, we discuss experiments designed to evaluate two key hypotheses related to DISCO:

- H1:** Using VLM-generated keyframes enables effective language-guided completion of tasks.
- H2:** Constrained inpainting improves task completion compared to standard inpainting, especially on novel task descriptions.

We conducted three sets of experiments. In the following, we first describe diffusion policy-related environments for trajectory and grasp pose generation (Sec. V-A). Then, we tested DISCO on a large-scale language-conditioned manipulation benchmark, CALVIN (Sec. V-B). Finally, we evaluated zero-shot transfer from simulation to a real robot setup (Sec. V-D). Experiment details are on website.


### A. Diffusion Policy Environments

**Experiment and Environment Setup.** As shown in Figure 5, we adapted four simulation environments for language-conditioned manipulation tasks. We generated corresponding language descriptions and subgoals for the action sequences in the dataset, and categorized tasks into “seen tasks” and “unseen tasks”. Seen tasks are those with demonstrations and corresponding language descriptions in the dataset. Unseen tasks had *new* language descriptions, e.g., tasks specified in a different way or using synonyms, or manipulating objects in a different order than in the dataset. In the 2D environments (push-T and block-push), DISCO generates key points in the 2D plane, which are converted to the positional keyframes. In the 3D grasp pose and Kitchen environments, DISCO observes the scene from three views to generate key points and 3D positional keyframes as in Section IV-A.

To evaluate H1, we compared DISCO against an *unconditional diffusion policy* [1], [17], [37] and alternative methods for language-conditioning, i.e., classifier-based diffusion models with a fine-tuned *language-conditioned* diffusion policy (LangDiff) [8] or a *goal-conditioned* diffusion policy (GoalDiff) [7]. We also conducted comparative tests to keyframe adherence using standard or *vanilla inpainting* (Vanilla inpt.) [13] to evaluate H2.

**Results.** Overall, the results support both H1 and H2. As shown in Table 5(b), DISCO achieved comparable performance to both language and goal conditioned diffusion policies on *seen tasks* — the diffusion baselines performed better in tasks like push-T, push-block, and Franka kitchen (single task), but on the Franka kitchen composite/multi tasks and grasp pose generation, vanilla inpainting and DISCO achieved higher success rates. On *unseen tasks*, the success rates of the DISCO are significantly higher than the other baseline methods. One potential reason is that the inpainting methods leverage the keyframes generated by VLMs, which have been trained on large amounts of text data and were able to produce reasonable keyframes in many circumstances. Performance is further enhanced when using constrained inpainting across all the tasks, supporting H2 and indicating the importance of generating action sequences within high-likelihood regions of the diffusion policy. DISCO’s ability to handle unseen objects depends on both the VLM’s keyframe reasoning and the coverage of the downstream policy. When the policy is trained only on specific objects, generalization to new categories is limited. However, with unconditional diffusion models trained on diverse objects [37], DISCO can generate effective grasp poses even for unseen categories. DISCO’s two-stage setup is flexible in that it enables us to swap in diffusion models dependent on downstream context.

(a) Simulation experimental environments.




(b) Results of Simulation Experiments (success rate  $\uparrow$ )

		Push-T	Push-block	Kitchen		Grasp Pose Generation			
				Single	Multi	Mug	Bottle	Ham.	Fork
Seen tasks <sup>1</sup>	No condition	0.54	0.50	0.12	0.03	0.47	0.31	0.40	0.48
	LangDiff [8]	<b>0.98</b>	<b>1.00</b>	<b>1.00</b>	0.31	0.91	0.84	0.78	0.96
	GoalDiff [7]	0.96	<b>1.00</b>	<b>1.00</b>	0.27	0.88	<b>0.95</b>	0.70	0.93
	Vanilla inpt.	0.82	<b>1.00</b>	0.71	0.35	0.90	0.88	0.85	0.96
	DISCO(ours)	0.87	<b>1.00</b>	0.71	<b>0.42</b>	<b>0.92</b>	0.92	<b>0.97</b>	<b>0.97</b>
Unseen tasks <sup>2</sup>	LangDiff	0.63	0.63	-	0.06	0.02	0.04	0.02	0.24
	GoalDiff	0.83	0.70	-	0.06	0.00	0.08	0.00	0.23
	Vanilla inpt.	0.89	0.70	-	0.06	0.02	0.30	0.58	0.55
	DISCO (ours)	<b>0.94</b>	<b>0.75</b>	-	<b>0.31</b>	<b>0.65</b>	<b>0.51</b>	<b>0.80</b>	<b>0.87</b>

Fig. 5: Experiments in diffusion policy-related environments. (a): Environments include Push-T [35], Block-push [36], Franka kitchen [16], and Grasp Pose generation [17] (details in website). (b): The policies were only trained in the seen tasks and zero-shot transferred to unseen tasks. Success rates are calculated over 50 trials. Overall, we observed DISCO, which combines VLM-generated keyframes with constrained inpainting, achieved robust performance across seen and unseen tasks.

(a) Calvin Benchmark. ABC  $\rightarrow$  D zero-shot transfer tasks.



(b) Results in Calvin benchmark (success rate  $\times 100 \uparrow$ )

	Task completed in a row					Avg. Len
	1	2	3	4	5	
RT-1 [10]	53.3	22.2	9.4	3.8	1.3	0.90
RoboFlamingo [38]	82.4	61.9	46.6	33.1	23.5	2.48
SuSIE [39]	87.0	69.0	49.0	38.0	26.0	2.69
3D Diffuser Actor [31]	92.2	78.7	63.9	51.2	41.2	3.27
Vanilla inpt.	92.9	78.6	64.9	51.4	41.9	3.30
DISCO (ours)	<b>94.7</b>	<b>82.9</b>	<b>71.0</b>	<b>58.8</b>	<b>49.4</b>	<b>3.57</b>

Fig. 6: Experiments in Calvin Benchmark [40]. Results of DISCO were calculated in  $3 \times 1000$  trials.

### B. CALVIN Benchmark

**Experiment and Environment Setup.** The CALVIN benchmark [40] (Fig. 6(a)) was built for language-conditioned policy learning in long-horizon robot manipulation tasks. Each instruction chain comprises five language instructions that need to be executed sequentially. Our experiments involved the *zero-shot generalization setup*, where models are trained in environments ABC and tested in D. We implemented DISCO with ChatGPT4 as keyframe generator and 3D Diffuser Actor [31] as the low-level diffusion policy. DISCO observes and marks key points in the top and front views and synthesis target end-effector position as the keyframe to guide the policy. We compared DISCO with SOTA language-conditioned policies RT-1 [10], RoboFlamingo [38], and SuSIE [39] and DISCO with vanilla inpainting method. Following prior work [40], we report the success rate and the average length of completed sequential tasks.

**Results.** As shown in Table 6(b), DISCO achieved the best

performance on the zero-shot transfer (ABC  $\rightarrow$  D) tasks, with an average length of 3.57, improving upon the SOTA baseline policy. Vanilla inpainting did not improve performance — this was primarily due to the inaccurate keyframes that caused the suboptimal action generation. We observed that DISCO mainly improved success rates in two types of scenarios: those requiring (i) additional observational information and (ii) common-sense understanding of the language instructions. For example, the task *Place the grasped block into the slider* requires the robot to first grasp and lift one block and then place it into the slider. However, CALVIN does not provide precise observations of the slider (only whether it “left” or “right”) and the policy has to infer its position and whether it is open/closed from the training dataset. DISCO exploits the image processing capabilities of VLMs to infer the state/position of the slider and generate sufficiently useful keyframes. This led to an increase of task success from 70% to 90%. For the task *Push the red/blue/pink block to the left/right side*, the language instruction was not in the dataset and requires the policy to understand that it is required to push the block in a particular direction. The baseline policy struggled with this task, achieving only a 50% success rate, but by using the common-sense understanding of VLMs, DISCO improved the success rate to 85%.

### C. Additional Analysis

**Keyframe quality analysis.** We investigate how the clarity of task instructions influences the quality of keyframe generation—specifically, whether the resulting keyframes are reasonable and satisfy the intended task conditions. For seen tasks (e.g., “pick up the mug by the handle”), instructions are explicit and conditions are well defined. In contrast, unseen tasks present moderate ambiguity (e.g., “Give me the bottle”) or high ambiguity (e.g., “Hand over the hammer to me”), where object references or conditions are indirect or missing. We assess keyframe quality using two metrics: (1) the improvement in DISCO’s success rate over the vanilla inpainting method, and (2) the normalized deviation of generated actions from the provided keyframes ( $D^{\text{key}}$  in

TABLE I: Keyframe Quality Analysis

Instruction clearness	Task portion	Success rate improvement ↓	Keyframe deviation ↓
Clear	0.44	0.04	0.07
Moderate	0.40	0.12	0.25
Ambiguous	0.16	0.22	0.41

Eqn. (12)). As summarized in Table I, most tasks involve clear instructions that yield high-quality keyframes, with minimal improvement or deviation required. As ambiguity increases, DISCO demonstrates larger gains over vanilla inpainting, and the deviation from the diffusion policy distribution increases accordingly. In highly ambiguous cases, keyframes become nearly out-of-distribution, making the inpainting optimization essential for correcting and refining the guidance. This analysis highlights how instruction ambiguity directly affects keyframe quality and underscores the importance of constrained inpainting in handling suboptimal guidance.

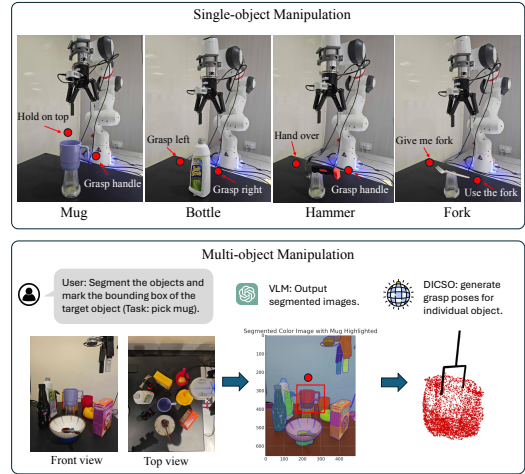
**Comparison with Visual-Language-Action (VLA) Models.** VLA models, such as OpenVLA [41],  $\pi_0$  [42], and DeeR-VLA [43], have recently gained popularity as an alternative approach for language-guided robotics, relying on fine-tuning large vision-language models to generate robot actions end-to-end via auto-regressive outputs. While these models provide a powerful solution, DISCO adopts a different strategy by decoupling high-level reasoning and low-level control: VLMs are used to generate interpretable keyframes, which then guide a diffusion policy through inpainting optimization. We compare DISCO with the recent DeeR-VLA model in the CALVIN zero-shot setup (ABC→D), where DISCO achieves a higher average task completion length (3.57) than DeeR-VLA (2.90). This result highlights the advantage of VLM-guided diffusion policies for compositional and open-vocabulary manipulation tasks.

#### D. Transfer to Real Robot Grasping

**Experiment and Environment Setup.** We transferred a diffusion-based grasp pose generation model trained in simulation to a real-world Franka Emika robot (Figure 7). We applied DISCO to generate 3D keyframes given real images and point clouds, then generated conditional grasp poses via constrained inpainting. We used a *single* grasp diffusion model [37] trained in simulation to generalize across all object categories without any parameter fine-tuning. Example language instructions for seen tasks are “give me the mug from top”, and “grasp the hammer handle”, while unseen tasks include “lift the bottle” and “hand over the fork”.

**Results.** Overall, the results in Table. 7(b) support both H1 and H2. In the *single-object* manipulation experiments, The goal-conditioned methods achieved only a 10% success rate in seen tasks and 0% in unseen tasks. On seen tasks, vanilla inpainting and DISCO (inpainting optimization) methods achieved success rates of 55% to 85%, respectively. In unseen tasks where vanilla inpainting faltered (15%), DISCO maintained a 65% success rate. In the *multi-object* manipulation experiment, DISCO still achieved fairly good grasping in seen tasks with success rate 65%. The tasks with “grasping from the side” instruction only achieved 40% successes, but

(a) Real-robot experiments setup



(b) Results in Real-robot Experiments (success rate ↑)

	Single-object task		Multi-object task	
	Seen	Unseen	Seen	Unseen
GoalDiff [7]	0.10	0.0	0.05	0.0
Vanilla inpt.	0.55	0.15	0.30	0.05
DISCO (ours)	<b>0.85</b>	<b>0.65</b>	<b>0.65</b>	<b>0.25</b>

Fig. 7: Real-robot experiments. We transferred DISCO from simulation to real-robot scenarios with multiple objects. In the single-object tasks, DISCO generated language-conditioned grasp poses for one object, while in multi-object tasks, we first segmented the target object from the cluster, and then applied DISCO. Each object was tested with 20 trials.

when asked to pick objects “from the top” success rates rose to 80% as this avoided collisions with other objects. We admit this limitation that the keyframes generated by DISCO did not consider the potential collisions when the end-effector moves to it. Finally, DISCO could also understand the unseen instruction in cluttered scenarios and achieve certain successful grasping (25 %).

#### VI. CONCLUSION, DISCUSSION, AND FUTURE WORK

In this paper, we propose DISCO, a zero-shot, open-vocabulary framework for robot manipulation using language descriptions. By leveraging VLMs, our method guides diffusion policies through actionable keyframes and addresses difficult out-of-distribution issues with an inpainting optimization technique. This approach is able to match (and sometimes surpass) traditional fine-tuned methods in both simulated and real-world settings. We believe DISCO represents a key step towards effective language-conditioned robot behavior. While it performs well in our experiments, there are areas for improvement. At present, DISCO does not account for occlusions during reasoning, which could be addressed by incorporating object pose information, applying joint constraints, or using failure data to replan and gather more information. DISCO is also heavily reliant on the underlying diffusion policy, which could be further improved by novel flow-matching models. Another limitation is that when DISCO executes a sequence of sub-tasks, keyframe switching is based on an empirical thresholding rule, which can occasionally lead

to backtracking or suboptimal transitions. We are actively working to build upon DISCO to address these limitations.

#### REFERENCES

- [1] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," *arXiv preprint arXiv:2303.04137*, 2023.
- [2] Y. Ze, G. Zhang, K. Zhang, C. Hu, M. Wang, and H. Xu, "3d diffusion policy," *arXiv preprint arXiv:2403.03954*, 2024.
- [3] M. Janner, Y. Du, J. B. Tenenbaum, and S. Levine, "Planning with diffusion for flexible behavior synthesis," *arXiv preprint arXiv:2205.09991*, 2022.
- [4] H. Zhou, X. Yao, Y. Meng, S. Sun, Z. Bing, K. Huang, and A. Knoll, "Language-conditioned learning for robotic manipulation: A survey," *arXiv preprint arXiv:2312.10807*, 2023.
- [5] C. Hao, A. Xiao, Z. Xue, and H. Soh, "Chd: Coupled hierarchical diffusion for long-horizon tasks," *arXiv preprint arXiv:2505.07261*, 2025.
- [6] S. Stepputtis, J. Campbell, M. Phielipp, S. Lee, C. Baral, and H. Ben Amor, "Language-conditioned imitation learning for robot manipulation tasks," *Advances in Neural Information Processing Systems*, vol. 33, pp. 13 139–13 150, 2020.
- [7] M. Reuss, M. Li, X. Jia, and R. Lioutikov, "Goal-conditioned imitation learning using score-based diffusion policies," *arXiv preprint arXiv:2304.02532*, 2023.
- [8] H. Ha, P. Florence, and S. Song, "Scaling up and distilling down: Language-guided robot skill acquisition," in *Conference on Robot Learning*. PMLR, 2023, pp. 3766–3777.
- [9] F. Liu, K. Fang, P. Abbeel, and S. Levine, "Moka: Open-vocabulary robotic manipulation through mark-based visual prompting," *arXiv preprint arXiv:2403.03174*, 2024.
- [10] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu *et al.*, "Rt-1: Robotics transformer for real-world control at scale," *arXiv preprint arXiv:2212.06817*, 2022.
- [11] M. Shridhar, L. Manuelli, and D. Fox, "Perceiver-actor: A multi-task transformer for robotic manipulation," in *Conference on Robot Learning*. PMLR, 2023, pp. 785–799.
- [12] K. Lin, C. Agia, T. Migimatsu, M. Pavone, and J. Bohg, "Text2motion: From natural language instructions to feasible plans," *Autonomous Robots*, vol. 47, no. 8, pp. 1345–1365, 2023.
- [13] A. Lugmayr, M. Danelljan, A. Romero, F. Yu, R. Timofte, and L. Van Gool, "Repaint: Inpainting using denoising diffusion probabilistic models," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 11 461–11 471.
- [14] L. Rout, A. Parulekar, C. Caramanis, and S. Shakkottai, "A theoretical justification for image inpainting using denoising diffusion probabilistic models," *arXiv preprint arXiv:2302.01217*, 2023.
- [15] Z. Yang, J. Mao, Y. Du, J. Wu, J. B. Tenenbaum, T. Lozano-Pérez, and L. P. Kaelbling, "Compositional diffusion-based continuous constraint solvers," *arXiv preprint arXiv:2309.00966*, 2023.
- [16] P. Sharma, L. Mohan, L. Pinto, and A. Gupta, "Multiple interactions made easy (mime): Large scale demonstrations data for imitation," in *Conference on robot learning*. PMLR, 2018, pp. 906–915.
- [17] J. Urain, N. Funk, J. Peters, and G. Chalvatzaki, "Se (3)-diffusionfields: Learning smooth cost functions for joint grasp and motion optimization through diffusion," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 5923–5930.
- [18] K. Chen, E. Lim, K. Lin, Y. Chen, and H. Soh, "Behavioral refinement via interpolant-based policy diffusion," *arXiv preprint arXiv:2402.16075*, 2024.
- [19] J. Gu, S. Kirmani, P. Wohlhart, Y. Lu, M. G. Arenas, K. Rao, W. Yu, C. Fu, K. Gopalakrishnan, Z. Xu *et al.*, "Rt-trajectory: Robotic task generalization via hindsight trajectory sketches," *arXiv preprint arXiv:2311.01977*, 2023.
- [20] D. Driess, F. Xia, M. S. Sajjadi, C. Lynch, A. Chowdhery, B. Ichter, A. Wahid, J. Tompson, Q. Vuong, T. Yu *et al.*, "Palm-e: An embodied multimodal language model," *arXiv preprint arXiv:2303.03378*, 2023.
- [21] W. Huang, C. Wang, R. Zhang, Y. Li, J. Wu, and L. Fei-Fei, "Voxposer: Composable 3d value maps for robotic manipulation with language models," *arXiv preprint arXiv:2307.05973*, 2023.
- [22] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, X. Chen, K. Choremanski, T. Ding, D. Driess, A. Dubey, C. Finn *et al.*, "Rt-2: Vision-language-action models transfer web knowledge to robotic control," *arXiv preprint arXiv:2307.15818*, 2023.
- [23] J. Liang, W. Huang, F. Xia, P. Xu, K. Hausman, B. Ichter, P. Florence, and A. Zeng, "Code as policies: Language model programs for embodied control," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 9493–9500.
- [24] J. Ho and T. Salimans, "Classifier-free diffusion guidance," *arXiv preprint arXiv:2207.12598*, 2022.
- [25] E. Zhang, Y. Lu, W. Wang, and A. Zhang, "Language control diffusion: Efficiently scaling through space, time, and tasks," *arXiv preprint arXiv:2210.15629*, 2022.
- [26] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman *et al.*, "Do as i can, not as i say: Grounding language in robotic affordances," *arXiv preprint arXiv:2204.01691*, 2022.
- [27] P. Zhi, Z. Zhang, M. Han, Z. Zhang, Z. Li, Z. Jiao, B. Jia, and S. Huang, "Closed-loop open-vocabulary mobile manipulation with gpt-4v," *arXiv preprint arXiv:2404.10220*, 2024.
- [28] O. Elharrouss, N. Almaadeed, S. Al-Maadeed, and Y. Akbari, "Image inpainting: A review," *Neural Processing Letters*, vol. 51, pp. 2007–2028, 2020.
- [29] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.
- [30] C.-F. Yang, H. Xu, T.-L. Wu, X. Gao, K.-W. Chang, and F. Gao, "Planning as in-painting: A diffusion-based embodied task planning framework for environments under uncertainty," *arXiv preprint arXiv:2312.01097*, 2023.
- [31] T.-W. Ke, N. Gkanatsios, and K. Fragkiadaki, "3d diffuser actor: Policy diffusion with 3d scene representations," *arXiv preprint arXiv:2402.10885*, 2024.
- [32] A. Liu, M. Niepert, and G. V. d. Broeck, "Image inpainting via tractable steering of diffusion models," *arXiv preprint arXiv:2401.03349*, 2023.
- [33] M. Shridhar, L. Manuelli, and D. Fox, "Cliport: What and where pathways for robotic manipulation," in *Conference on robot learning*. PMLR, 2022, pp. 894–906.
- [34] X. Li, M. Zhang, Y. Geng, H. Geng, Y. Long, Y. Shen, R. Zhang, J. Liu, and H. Dong, "Manipllm: Embodied multimodal large language model for object-centric robotic manipulation," *arXiv preprint arXiv:2312.16217*, 2023.
- [35] P. Florence, C. Lynch, A. Zeng, O. A. Ramirez, A. Wahid, L. Downs, A. Wong, J. Lee, I. Mordatch, and J. Tompson, "Implicit behavioral cloning," in *Conference on Robot Learning*. PMLR, 2022, pp. 158–168.
- [36] N. M. Shafiqullah, Z. Cui, A. A. Altanzaya, and L. Pinto, "Behavior transformers: Cloning  $k$  modes with one stone," *Advances in neural information processing systems*, vol. 35, pp. 22 955–22 968, 2022.
- [37] J. Urain, N. Funk, J. Peters, and G. Chalvatzaki, "Se(3)-diffusionfields: Learning smooth cost functions for joint grasp and motion optimization through diffusion," *IEEE International Conference on Robotics and Automation (ICRA)*, 2023.
- [38] X. Li, M. Liu, H. Zhang, C. Yu, J. Xu, H. Wu, C. Cheang, Y. Jing, W. Zhang, H. Liu *et al.*, "Vision-language foundation models as effective robot imitators," *arXiv preprint arXiv:2311.01378*, 2023.
- [39] K. Black, M. Nakamoto, P. Atreya, H. Walke, C. Finn, A. Kumar, and S. Levine, "Zero-shot robotic manipulation with pretrained image-editing diffusion models," *arXiv preprint arXiv:2310.10639*, 2023.
- [40] O. Mees, L. Hermann, E. Rosete-Beas, and W. Burgard, "Calvin: A benchmark for language-conditioned policy learning for long-horizon robot manipulation tasks," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 7327–7334, 2022.
- [41] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. Foster, G. Lam, P. Sanketi *et al.*, "Openvla: An open-source vision-language-action model," *arXiv preprint arXiv:2406.09246*, 2024.
- [42] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter *et al.*, "pi\_0: A vision-language-action flow model for general robot control," *arXiv preprint arXiv:2410.24164*, 2024.
- [43] Y. Yue, Y. Wang, B. Kang, Y. Han, S. Wang, S. Song, J. Feng, and G. Huang, "Deer-vla: Dynamic inference of multimodal large language models for efficient robot execution," *Advances in Neural Information Processing Systems*, vol. 37, pp. 56 619–56 643, 2024.
- [44] P. Dhariwal and A. Nichol, "Diffusion models beat gans on image synthesis," *Advances in neural information processing systems*, vol. 34, pp. 8780–8794, 2021.