

APEX: Ambidextrous Dual-Arm Robotic Manipulation Using Collision-Free Generative Diffusion Models

Apan Dastider*, Hao Fang* and Mingjie Lin*

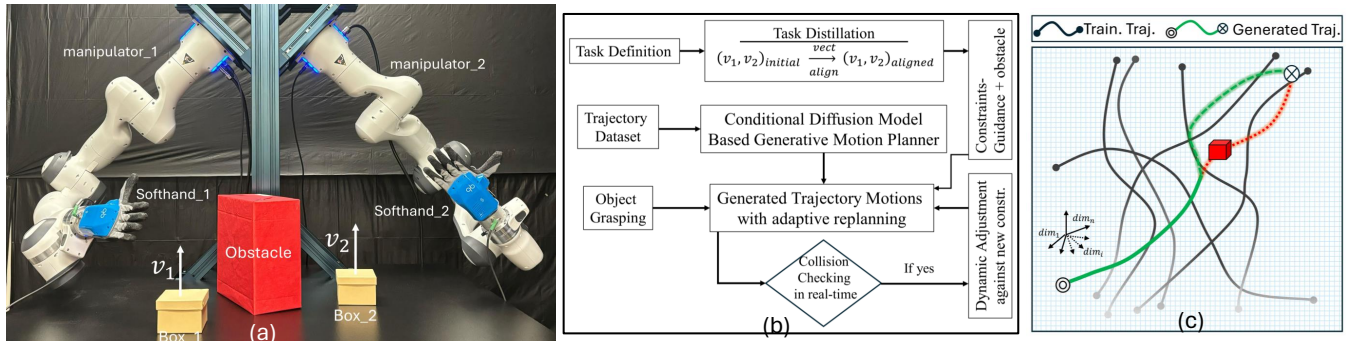


Fig. 1: a) Dual-arm robotic manipulation platform (b) Overall diagram of our algorithm. (c) Our diffusion models learn the distribution of trajectory using training data (black line) and generate a new planning trajectory (green dashed line). Further, it adapts the generated trajectory (red dashed line) to avoid the potential obstacles (red box).

Abstract— Dexterous manipulation, particularly adept coordinating and grasping, constitutes a fundamental and indispensable capability for robots, facilitating the emulation of human-like behaviors. Integrating this capability into robots empowers them to supplement and even supplant humans in undertaking increasingly intricate tasks in both daily life and industrial settings. Unfortunately, contemporary methodologies encounter serious challenges in devising manipulation trajectories owing to the intricacies of tasks, the expansive robotic manipulation space, and dynamic obstacles. We propose a novel approach, APEX, to address all these difficulties by introducing a collision-free latent diffusion model for both robotic motion planning and manipulation. Firstly, we simplify the complexity of real-life ambidextrous dual-arm robotic manipulation tasks by abstracting them as aligning two vectors. Secondly, we devise latent diffusion models to produce a variety of robotic manipulation trajectories. Furthermore, we integrate obstacle information utilizing a classifier-guidance technique, thereby guaranteeing both the feasibility and safety of the generated manipulation trajectories. Lastly, we validate our proposed algorithm through extensive experiments conducted on the hardware platform of ambidextrous dual-arm robots. Our algorithm consistently generates successful and seamless trajectories across diverse tasks, surpassing conventional robotic motion planning algorithms. These results carry significant implications for the future design of diffusion robots, enhancing their capability to tackle more intricate robotic manipulation tasks with increased efficiency and safety. Complete video demonstrations of our experiments can be found in <https://sites.google.com/view/apex-dual-arm/home>.

Index Terms— Ambidextrous dual-arm robotic manipulation, Latent diffusion model, Collision-free, Classifier guidance

I. INTRODUCTION

Ambidextrous dual-arm robotic manipulation requires sophisticated coordination and control to synchronize the movements of both arms and hands seamlessly. Unlike well-studied single-arm robots, ambidextrous robots must coordinate the actions of two separate arms and two hands while navigating complex environments and interacting with dynamic objects. This heightened level of complexity ne-

cessitates advanced planning and control strategies to ensure successful task execution and obstacle avoidance. However, realizing such ambidextrous dual-arm robotic manipulation in real world is complicated [1]. First, different from the prior 7-DoF robotic arm model, the ambidextrous dual-arm robotic system allows more joint freedom, which requires more parameters and faces complex constraints to comprehensively describe the joint configuration space [2], [3]. Second, the aim of the ambidextrous dual-arm robotic system is to perform more complicated human-like tasks, such as packaging, stacking, steering liquid in a cup, and etc; which impose more dexterous manipulation requirements such as smoothness and diversity. Third, the real-time moving obstacles should be avoided during manipulation, which requires the algorithm to work under a close-loop fashion [4]. Fourth, it is crucial that motion planning algorithms exhibit computational efficiency to support real-time adaptive control schemes [5]–[7]. All these challenges collectively impede the advancement of motion planning algorithms in dual-arm robotic manipulation tasks.

To date, diffusion models have emerged as a powerful tool in robotics, owing to their flexibility and multi-modality. However, their application in the context of dual-arm robotic manipulation remains relatively unexplored, particularly concerning real-time performance and the diversity of manipulation tasks. In this paper, we propose a collision-free diffusion model to address the aforementioned challenges. The key contributions of our work are outlined below:

- Abstracting real-life ambidextrous dual-arm robotic manipulation tasks as two vector alignment problems, a conceptualization that significantly simplifies the tasks and underscores the model’s applicability across various real-world scenarios.
- Development of latent diffusion models for generating diverse robotic manipulation trajectories. Additionally, we integrate obstacle information using the classifier-guidance technique, ensuring the practicality and safety of the resulting manipulation trajectories.
- Validation of the proposed algorithm on a hardware

*Department of Electrical and Computer Engineering, University of Central Florida, Orlando, FL, 32816, USA (E-mail: apan.dastider@ucf.edu)

Sample Tasks which can be solved by APEX

Task Distillation Applied on Simple Vector Alignment (TD-SVA) Task Abstraction			Task Distillation Applied on Complicated Vector Alignment Task (TD-DVA) Abstraction		
TASK: Ball in a Cup (peg and hole task)	TASK : Stacking Package	TASK : Open Lid of a Box	TASK : Opening Cap of Bottle(CCW Motion)	TASK : Steering Liquid in A Cup (CW Rotation)	TASK : Pouring Liquid from one cup to another cup

Fig. 2: Solving the real-world dexterous dual-arm robotic manipulation tasks. The first row indicates various robotic manipulation tasks. The second row is the initial position of the manipulation tasks. The third row distills the manipulation tasks by abstracting them as vector alignment problems, which can be solved using our proposed algorithm. The last row shows the successful completeness of the manipulation tasks. We focus here on first two tasks.

platform featuring ambidextrous dual-arm robots. Our algorithm consistently produces successful and seamless trajectories across multiple tasks, outperforming conventional robotic motion planning algorithms. These findings carry implications for the future design of diffusion robots, particularly in tackling more complex robotic manipulation tasks effectively.

II. RELATED WORK

A. Dual-arm robotic manipulation

Dual-arm robotic manipulation has been traditionally formulated as a classical optimal control problem, which primarily involves minimizing a desired cost function using a simplified model of the dynamical system [5], [8]–[10]. Second, researchers also primarily relied on the classic sampling framework of Rapidly-exploring Random Trees (RRT) algorithms [11], [12], which mainly develops graph traversal methods in the robotic configuration space, whose trajectory planning ability becomes limited and useless if the dimension of the robotic configuration space becomes higher. Recently, reinforcement learning (RL) algorithms [13], [14], have been introduced by maximizing a reward function for accomplishing such dual-arm manipulation tasks. Imitation Learning (IL) algorithms initially gather a substantial number of task samples, such as human demonstrations, and then employs deep learning architectures like convolutional/recurrent neural networks [15], [16] or transformers [17] to train the model through teacher-forcing training. However, the above learning-based methods may be severely suffered by the heterogeneous data distribution shift, i.e., the training and test data are sampled from different manipulation tasks or the obstacles distribution is biased, and limited trajectory generation diversity, which ultimately decreases their potential advantages. Although these approaches have proven highly effective in some cases, they do not address highly complex real-world working spaces with potentially dynamic obstacles or the full coordination of robotic arms and hands.

B. Diffusion models

Diffusion models have been illustrated with great potential as generative models of images [18] and videos [19]

and outperform other generative models such as adversarial networks (GANs) [20] and variational autoencoders (VAEs) [21]. Ho et al [22] proposes denoising diffusion probabilistic models (DDPMs) by treating the training procedure as minimizing the mean-square error of noise at each time step. To facilitate the sampling time, latent diffusion models (LDMs) are proposed to first integrate autoencoders (AEs) for latent representation and then to train the diffusion model in the corresponding latent space [23]. After sampling, the generated latent representation will be sent through a decoder to recover the generated images. For most of the diffusion models, the underlying backbone architecture is selected as convolutional neural networks (CNNs), such as U-Nets [24]. To date, there are few works that incorporate the diffusion models for robotic motion planning [25]–[28]. For instance, [25] utilized a zero-shot learning mechanism with pre-trained generative models to complete object arrangement tasks. On the contrary, [27] implemented a reward mechanism along with diffusion model-based trajectory generation and the authors proposed a reward-based guidance for optimization with maximum utility for task completion.

III. METHODS

A. Distill ambidextrous dual-arm manipulation tasks as vector alignment problems

Many day-to-day tasks such as stacking packages or pouring liquid from one cup to another cup can be summarized and simplified. In our work, we distill the dual-arm manipulation tasks by vector alignment problems. Here, we demonstrated several dual-arm manipulation tasks as well as their distillations in Fig.2. We first used the well-established YOLO algorithm (a typical type of convolutional neural network in computer vision) for object detection from the captured frame using the depth sensor [29]. The YOLO algorithm localized the object of interest at each time frame and detected the position and orientations of those objects. Then we drew simple 2D vectors v_1 and v_2 along the orientations of the objects as depicted in the upper row of Fig.2. Once we captured the orientations of two vectors, we can convert the complex manipulation task to a simple vector alignment problem. Therefore, this distillation reduces computationally

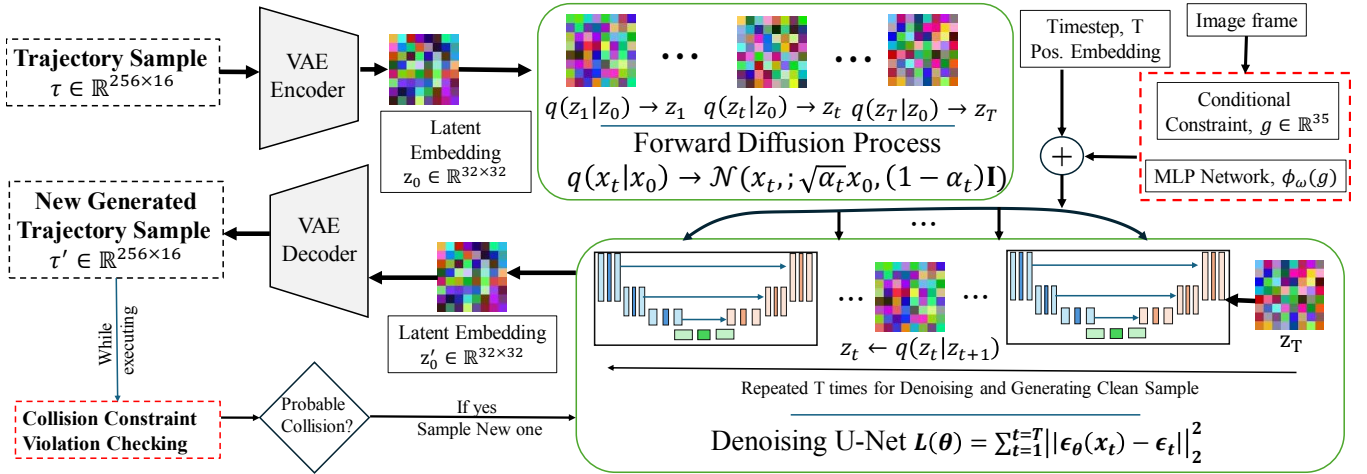


Fig. 3: The diagram of our ambidextrous dual-arm robotic manipulation algorithm using latent diffusion models to generate the collision-free trajectory. The input of our trajectory data x first goes into a pre-trained variational autoencoder to have a low-dimensional latent embedding z_0 , which will be used to train a diffusion model. In the sampling procedure, we randomly generate Gaussian noise trajectory z_T and pass it through reverse denoising steps to generate clean trajectory latent embedding z_0' , which can be decoded using a pre-trained VAE decoder. Further, to ensure the generated trajectory x' is collision-free, we add two mechanisms on top of the latent diffusion models. First, we use the obstacle information g as the conditional constraint in both training and sampling process. Second, our trajectory generation process works in a close-loop manner, where we check the possibility of collision violation and resample a new trajectory if necessary.

heavy image-based motion planning problems to a simple vector-conditioned trajectory planning problem. Take the first task (Ball in a cup) for an illustration, the two vectors are randomly placed initially, where we are required to develop a motion planning algorithm such that initial two vectors are aligned in the desired positions and opposite orientations. Once our proposed algorithm finishes the task, we can measure the robot configuration space using the measured initial state and final state for the performance evaluation of our algorithm. Further, we also consider dynamic obstacles during our tasks for the mimic of more complicated real-world robotic motion planning tasks.

Let $\mathcal{S} \subseteq \mathbb{R}^n$ and $\mathcal{U} \subseteq \mathbb{R}^m$ define the state space and control input for our dual-arm robotic system. The robotic system evolves through the following discrete time-dynamics,

$$\mathcal{S}_{t+1} = f(\mathcal{S}_t, u_t) \quad (1)$$

where $f(\cdot)$ is a function controlling the discrete-time system progression for action u_t from state \mathcal{S}_t to \mathcal{S}_{t+1} . We assume the $\mathcal{S}_{\text{safe}}$ and $\mathcal{S}_{\text{unsafe}}$ are the collision-free and collision space with $\mathcal{S}_{\text{safe}} \cap \mathcal{S}_{\text{unsafe}} = \emptyset$, respectively. During the task distillation, we get the initial state vector, $\mathcal{S}_{\text{ini}} \in \mathcal{S}_{\text{safe}}$ and target vector, $\mathcal{S}_{\text{goal}} \in \mathcal{S}_{\text{safe}}$ and current point obstacle information. Together, all the above information constitutes a conditional guidance vector g_t . The goal of our algorithm is to generate a trajectory, $\tau = \{\tau_0, \dots, \tau_{T-1}\}$ conditioned on g_t such that it is safe against dynamically moving obstacles ($\tau \in \mathcal{S}_{\text{safe}}$) and drive the two-arm robotic system reaching the goal state $\mathcal{S}_{\text{goal}}$.

B. Diffusion-based deep generative learning

In this section, we describe the main component of our diffusion-based deep generative model for the synthesis of alignment trajectories. We consider the following Gaussian diffusion models, which include a forward process and a reverse process. The forward noise process describes conditional probability distribution $q(x_t|x_0)$ given the initial data distribution $q(x_0)$, which can be the collected trajectory data in our case.

$$q(x_t|x_0) = \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t}x_0, (1 - \bar{\alpha}_t)\mathbf{I}), \quad (2)$$

where $\bar{\alpha}_t$ are predefined constants for introducing noise at each forward process [22]. Using the property of Gaussian distribution and parameterization trick, one can directly sample x_t using (2) by $x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{(1 - \bar{\alpha}_t)}\epsilon_t$, where ϵ_t follows standard Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$. The reverse diffusion process is to learn a set of neural network parameters θ that reverse the forward noise process,

$$p(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t), \Sigma_\theta(x_t)). \quad (3)$$

Therefore, the goal of the diffusion models is to ensure the prediction $p_\theta(x_0)$ using learned neural network parameters θ that can recover the original data distribution $q(x_0)$ as close as possible, which can be quantified using evidence lower bound (ELBO) [22]. Thus, one can maximize the lower bound of the log-likelihood of x_0 with the following loss function

$$L(\theta) = -\mathbb{E}_{x_0:T \sim q} [\log p_\theta(x_0:T) - \log q(x_1:T|x_0)]. \quad (4)$$

However, the above loss function (4) has been recognized as hard for fast training since it focuses on image reconstructions. Essentially, the diffusion model needs to care more about the reverse denoising process by focusing on cleaning the forward noise corruption, which inspires the design of noise prediction work [22],

$$L(\theta) = \sum_{t=1}^T \|\epsilon_\theta(x_t) - \epsilon_t\|_2^2. \quad (5)$$

Therefore, we utilize this simple mean-square error loss through our training process. In the sampling case, we randomly sample x_T from standard Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$ and pass through the reverse denoising process to gradually clean the noisy trajectories step by step $p(x_{t-1}|x_t)$ (see Figure 3). The above equations (2), (3), (5) describe how we construct the diffusion model for the generation of new trajectories. Also, the backbone of the diffusion model follows the U-Net architecture, a typical type of symmetric conventional convolutional neural networks (CNNs).

However, the above naive diffusion model meets two difficulties for the robotic trajectory generations. First, the robotic planning trajectory usually has high temporal dimensions. Thus, directly training and sampling on the high-dimensional

trajectory space is hard and time-consuming. Second, the manipulation of robotics usually involves the interaction of dynamic environments, meaning that the robot needs to react to the environmental obstacles and replan the trajectories to avoid any potential collisions. Therefore, we leverage the latent diffusion models by first training a variational autoencoder to embed the high-dimensional trajectory as $z = \text{Encoder}(\tau)$. Then, we train the proposed diffusion model in the embedded space. Last, we use a trained decoder $\tau = \text{Decoder}(z)$ to map the sampled low-dimensional trajectory embedding back to the original trajectories for later manipulation of ambidextrous dual-arm robots. To deal with the second difficulty, we integrate the classifier guidance technique as the real-time feedback condition to guide the generation of safe and feasible trajectories (see next section).

C. Obstacle guidance as feedback to close the loop

The most advanced diffusion robot works mainly applied the diffusion framework for synthesizing planning trajectories without considering dynamic environmental obstacles [27], [28], [30], [31]. Our second innovative contribution is to use the classifier guidance technique [20], [32] as the real-time feedback condition to guide the generation of safe and feasible trajectories. Specifically, we assume at certain steps, our camera sensor can take the image of the current interaction of the robotic arms and environment (denoted as g_t), abstractly representing the current alignment procedure and potential environmental obstacles. The conditional diffusion models using the obstacle guidance technique can be developed under the Bayes theorem $p(x|g) \propto p(x)p(g|x)$,

$$\nabla_x \log p_\theta(x|g) = \nabla_x \log p_\theta(x) + \gamma \nabla_x \log p_\theta(g|x), \quad (6)$$

where γ controls the guidance scale; g is the interaction information that we described above. Put equation (6) into the diffusion model, it turns out that we only need to modify noise prediction at the sampling stage via

$$\hat{\epsilon}_\theta(x_t) = \epsilon_\theta(x_t) - \gamma \sqrt{(1 - \alpha_t)} \nabla_{x_t} \log p_\theta(g|x). \quad (7)$$

In the extreme case $\gamma = 0$, this conditional diffusion model degrades to the previous diffusion model, which can generate diverse but unsafe manipulation trajectories. On the contrary, we impose the conditional generation of collision-free trajectory as we increase the guidance scale γ . Therefore, by integrating the classifier-guidance technique as obstacle-guidance, we essentially build a closed-loop trajectory generation procedure, where we adapt the prior generated trajectory at certain time steps using the real-time obstacle feedback.

To this end, we complete the development of our proposed algorithm. A summary of the algorithm can be found in the algorithm 1 and figure 3.

D. Platforms and implementation details

We utilized two high-fidelity 7-DoF Franka Emika Panda robot arms and assembled them on two separate table-top for smoother and collision-free operation. We used two anthropomorphic five-fingered hands named QB SoftHand2 Research in left-right hand combinations. To locate objects and track dynamic obstacles with depth information, we used an Intel RealSense Depth Camera D435i. On each frame captured by depth camera, we run YOLO for object detection and obstacle localization. Later, we detected the orientation and position of vectors associated with each

Algorithm 1: APEX

TRAINING

Input: Dataset, $\mathcal{D} : \{\text{Trajectory } \tau_i, \text{Guidance } g_i\}_{i=1}^N$,
Conditional Diffusion Model, $\epsilon_\theta; \bar{\alpha}_t, T$

Output: Trained model, $\tilde{\epsilon}_\theta$

Function $\text{train}(\mathcal{D}, \bar{\alpha}_t, \epsilon_\theta, T)$:

```

while not converged do
  # sample batch of traj. and respective
  # conditional vector
   $(\tau_0, g_0) \sim \mathcal{D}, \epsilon \sim \mathcal{N}(0, I), t \sim \mathcal{U}(1, T)$ 
  # concat time-embedding and  $g_i$ 
   $t\_cond = \text{concat}(\text{pos\_emb}(t), \phi_\omega(g_i))$ 
  # create noisy trajectory
   $\tau_\epsilon(t) \leftarrow \sqrt{\bar{\alpha}_t(t)}\tau_0 + \sqrt{1 - \bar{\alpha}_t(t)}\epsilon$ 
  # MSE Error calc.
   $L_\theta = \|\epsilon - \epsilon_\theta(\tau_\epsilon(t), t\_cond)\|_2^2$ 
  # Gradient Update
   $\theta \leftarrow \theta + \nabla_\theta L_\theta$ 

```

SAMPLING

Input: Frame f_0 , Trained Model, $\tilde{\epsilon}_\theta, \lambda_{obs}$

Output: Feasible Trajectory, τ^*

Function $\text{execute}(f_0, \tilde{\epsilon}_\theta, \lambda_{obs})$:

```

# Task Distil. for Vector Alignment and obs info
 $g_0 \leftarrow \text{vect\_align}(f_0)$ 
# Sample Noisy Trajectory
 $\tau_T \leftarrow \mathcal{N}(0, I)$ 
# Noise Prediction
 $\epsilon_T \leftarrow \epsilon_\theta(\tau_T, \text{concat}(\text{pos\_emb}(T), \phi_\omega(g_0)))$ 
# Initial Safe Trajectory Calc.
 $\tau_0 \leftarrow \text{traj\_calc}(\alpha_t, \epsilon_T, \tau_T)$ 
while not done do
  CALCULATE:
  Distance,  $d_{obs} \leftarrow GJK(\text{robot}, \text{obstacle})$ 
  # Collision Checking
  if  $d_{obs} < \lambda_{obs}$  then
    UPDATE:
    # Sample new trajectory for collision
    # avoidance
     $\tau_{new} \leftarrow \text{generate\_traj}(f_{curr}, \tilde{\epsilon}_\theta)$ 
  else
    EXECUTE:
    # Action Execution following generated
    # traj.
    Dual Arm-hand Control,
     $\tau_i \rightarrow \{[\theta_0^k, \dots, \theta_6^k]_{k=1,2}, \theta_h^1, \theta_h^2\}_i$ 

```

object. To differentiate the obstacle from objects, we used color-based segmentation on detected frames. By applying the conventional hand-to-eye calibration, we transported the $[x, y, z]_o$ information from the camera reference frame to the robot coordinate system. Our algorithm utilized this feedback from depth sensors to avoid obstacles in real-world scenarios by computing the distance between the 3D geometry of the obstacle and the collision meshes of robot links using the Gilbert–Johnson–Keerthi distance (GJK) [33] algorithm. This distance feedback triggers the controller to search for a new trajectory if chances of collision occur while completing the vector alignment task. To validate our approach, we replicated the exact model of the Franka Emika Panda Arms with QB SoftHand2 in the Robot Operating System (ROS). We utilized the libfranka and Franka ROS, to establish low-latency and low-noise communication protocols for data processing and parallel execution among simulation environments and real hardware. QB SoftHand2 research hands were operated through serial communication in parallel for object grasping and placement into the assigned alignment by associated vectors.

We collected high-dimensional trajectory τ_i samples for dual-arm manipulation setup from ROS, where each data sample τ_i^j contains total 16 float values – 14 joint values from two robotic arms $[\theta_0^k, \dots, \theta_6^k]_{k=1,2}$ and 2 synergy joint values




Fig. 4: Snapshots of hardware demonstration from vertical (a)-(e) and horizontal alignment task (f)-(j). (a)-(b) Robotic arms grasp the boxes from a table. (c)-(d) Right robotic arm avoids possible collision (red box). (e) Successful task completed. (f)-(g) Two arms take the cup and the ball to prepare a task in hand. (h) A dynamic obstacle appears in the workspace and the left robotic arm faces probable collision, (i) The Left arm moves over the obstacle to avoid collision. Now, the right arm gets into a probable collision with the moving red box. (j) Both arms dodge the obstacles and complete the task.




Fig. 5: Denoising diffusion process for the generation of manipulation trajectories. From left to right, we gradually show how the denoising diffusion process generates a smooth and feasible trajectory to reach the predefined goal position.

from two robotic hands $[\theta_h^1, \theta_h^2]$. We only use robotic hands for grasping tasks. Thus, we only exploit the first synergy of robotic hands. As a result, each robotic hand can either close or open, where the synergy value indirectly controls the 19 self-healing finger joints. The conditional constraint vector has 35 float values – 16 values for initial or current system state which contains $\{[\theta_0^k, \dots, \theta_6^k]_{k=1,2}, \theta_h^1, \theta_h^2\}_{ini}$, 16 values for final or target state $\{[\theta_0^k, \dots, \theta_6^k]_{k=1,2}, \theta_h^1, \theta_h^2\}_{tar}$, 3 values for obstacle point $[x, y, z]_o$. Then, the conditional vector g goes through a MLP network $\phi_\omega(g)$ parameterized by ω and the output of $\phi_\omega(\cdot)$ applies residual connections to each convolution layer in the U-Net(see Fig. 3), and later also uses positional embedding of each timestep during the model training. After obstacle detection, the GJK algorithm localizes the nearest point on obstacle body by calculating the distance from full robot meshes. We collected 500k collision-free trajectory samples with different starting and ending alignment configurations with only static obstacles from our simulation platform. Each trajectory samples consist 256 discrete trajectory steps which are from the starting state to the final ending state. Also, the dense trajectory samples allow us to perform adaptive planning from any initial state to the target goal state and handle dynamic obstacles simultaneously. We ensure that most of the workspace are covered while gathering trajectory samples to earn the generalizability required to successfully complete the assigned robotic tasks. Our algorithm was implemented using the Pytorch library and ran on a Lambda QUAD GPU workstation equipped with an Intel Core-i9-9820X processor and 4 Nvidia 3080 – ti GPU machines. It took 10 hours to

train our collision-free generative diffusion models.

E. Experiments setup and evaluation metrics

We design both comprehensive Monte Carlo computer simulations and hardware experiments to evaluate our proposed APEX algorithm for achieving dual-arm alignment manipulations. Our experiments mainly focus on two aspects: A) the vertical alignment, such as stacking, which can simplify the package sorting and stacking tasks in real life; B) the horizontal alignment, such as peg and hole task, which can be used to represent any insertion task in real life. We believe that by illustrating the performance on the above vertical and horizontal alignment tasks, we essentially evaluate our algorithms across a wide range of dual-arm robotic manipulation tasks. Further, we also implement typical existing robotic motion planning algorithms in our experiments, where we include the RRT [34] and GPMP [35], a transformer-based imitation learning algorithm for motion planning [17] and our prior VAE-based graph traversal planning algorithm [4]. For object grasping, we utilize advanced functionalities from ROS MoveIt Grasp API.

We define the following evaluation metrics to effectively quantify the performance. First, we define the failure ratio,

$$R_f = \frac{F}{N} \times 100\%, \quad (8)$$

where F counts the failed tasks and N denotes the total number of tasks that we performed. Therefore, smaller R_f means the generated trajectories successfully avoid obstacles and realize better alignment performances. Then, we quantify

the smoothness of the executing trajectory of dual-arm robots by summing up the joint space configuration change in sampled trajectory as

$$S_r = \sum_{i=0}^{N-1} \|\Theta_i - \Theta_{i+1}\|_2 \quad (9)$$

where Θ_i is joint space configuration at τ_i . Last, we use the Frechet inception distance (FID) to measure the generation diversity of our obstacle-guidance diffusion models,

$$\text{FID} = \|\mu_R - \mu_G\|^2 + \text{Tr}(\Sigma_R + \Sigma_G - 2(\Sigma_R \Sigma_G)^{\frac{1}{2}}), \quad (10)$$

where μ_R and μ_G denote the feature-wise mean of the oracle and generated trajectories; Σ_R and Σ_G are the covariance matrix for the real and generated feature trajectories; $\text{Tr}(\cdot)$ refers to the trace operation in linear algebra. Notice that researchers widely use the above FID score to quantify the realism and diversity in evaluating the performance of generative models. Here, we focus on the trajectory data and use pre-trained VAE model on our collected dataset. Therefore, we use the mean μ_G and variance Σ_G calculated from the VAE model for the measurement of the FID score.

IV. RESULTS

We use comprehensive experiments to evaluate our proposed APEX algorithm. Specifically, our experiments seek to investigate the following questions:

- 1) Can the ambidextrous dual-arm robots successfully complete the manipulation tasks using APEX?
- 2) What are the qualitative advantages and the essential trade-off of APEX compared to SOTA algorithms?
- 3) Why the obstacle guidance and closed-loop fashion are important for safe planning?
- 4) Are the trajectories generated by our algorithm diverse?

A. Successful and smooth ambidextrous dual-arm robotic manipulation trajectory

We first presented horizontal and vertical task demos to visualize the trajectories generated by our proposed algorithm. Fig. 4 presented the successful ambidextrous dual-arm robotic manipulation performance in both vertical and horizontal tasks. Here, we took the horizontal tasks for example. We observed that the two arm-hand robots were far away and randomly placed in the configuration space at the beginning (see Fig. 4(f)). After our diffusion model generated a trajectory, two robots gradually began to align resulting in a smooth trajectory (see Fig. 4(g)). During the middle point of the alignment, an obstacle was imposed manually, which may result in the potential collision (see Fig. 4(h)-(i)). Therefore, our robots stopped for avoiding any collision and generated a new trajectory, which lifted the two arms up to avoid the imposed hard obstacles (see Fig. 4(h)-(i)). Finally, our two arm-hand robots successfully realized horizontal alignment, which completed the manipulation task ((see Fig. 4(j)). Similar successful and smooth ambidextrous dual-arm robotic manipulation trajectories are observed in our vertical tasks (see Fig. 4(a)-(e)).

Meanwhile, we also looked into the denoising diffusion process in our APEX algorithm. The beginning manipulation trajectory ($t = 0$) was randomly sampled as a Gaussian image, which led to a rigid and aimless trajectory. Then, the denoising diffusion process gradually cleaned the noise trajectory (snapshots at $t = 64$ and $t = 128$). Finally, the cleaned trajectory enabled a successful completion of the task




Fig. 6: (a) Collision While doing horizontal alignment task, (b) Two arms are stuck in the middle and can not replan, (c) Incorrect orientation position for vertical alignment task




Fig. 7: (a) Comparison of failure ratio with SOTA algorithms, (b) Comparison of the smoothness with SOTA algorithms.

and the planar movement of the arm-hand robots was smooth ($t = 256$). The above two illustrations of manipulation tasks provide interesting collision-free trajectory generation results and show promising advantages of our proposed generative diffusion model.

B. Comparison results with SOTA algorithms

Next, we comprehensively evaluated our proposed algorithms for achieving various vertical and horizontal tasks from different initial positions and orientations. Further, to illustrate the potential advantages of our proposed algorithm, we compare our algorithm with other SOTA motion planning algorithms to evaluate the potential advantages.

We started by looking some failure cases of both vertical and horizontal alignments using other algorithms. Fig. 6 provides various failure examples. For example, Fig. 6(b) showed that the traditional RRT seemed simply generated half of the trajectories, where two robotic arms stopped during the alignment; Fig. 6(a) showed that the imitation learning algorithm can not deal with the obstacles, leading to a collision; Fig. 6(c) showed that how DAMON failed to complete the alignment task and ended in a faulty configuration for box placement task. In total, fig. 7(a) gave the comprehensive comparison of different algorithms in executing vertical and horizontal tasks. We observed that our prior work DAMON and imitation learning (IL) algorithms achieved better performance compared to the RRT and GPMP (Failure ratio: DAMON 32.65 v.s. IL 25.79 v.s. RRT 48.12 v.s. GPMP 39.26). However, our proposed diffusion-based generative models achieved extraordinary performance with the lowest failure rate across both vertical and horizontal tasks (Failure ratio (avg): vertical 5.57 and horizontal 8.95). We argue that the possible reason why the vertical failure ratio is even lower than the horizontal task is due to the more complicated orientation alignment in horizontal tasks. For example, the vertical alignment task mimics the stacking manipulation, which does not involve orientation manipulations, whereas horizontal alignment tasks such as ball-insertion-in-a-cup require more complicated orientational manipulations. Above all, our APEX algorithm achieved the best alignment performance across all comparison algorithms and tasks.

Second, smoothness is also very important for avoid-




Fig. 8: Snapshots of hardware demonstration. (a)-(c) Without guidance, the robots can only operate under small obstacles and partially complete the task. (d)-(e) The right robotic arm collides with a larger obstacle can not complete the task.

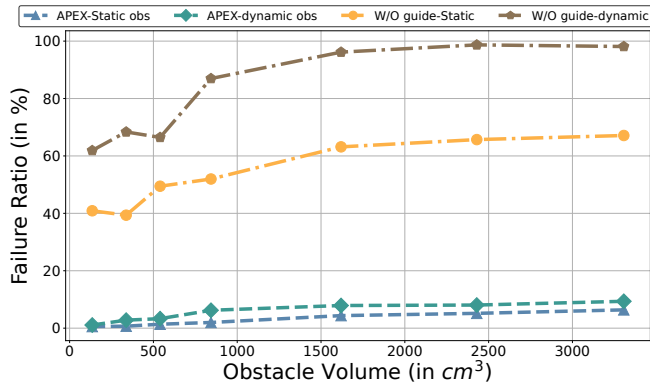


Fig. 9: Comparison of Failure Ratio with APEX and guidance-free stable diffusion algorithm for static and dynamic obstacle including awkward robotic manipulation behaviors. Prior demos illustrate that our diffusion models can generate smooth trajectories. One possible mechanism is the reverse diffusion process, where we use the trained noise prediction network to gradually denoise from a pure Gaussian noise trajectory. Advanced imitation learning (IL) algorithms also have mechanisms to ensure the smoothness of trajectory since the goal of IL is to mimic the training trajectories with mean square error loss. Fig. 7(b) gave the overall smoothness metric in our experiments, where we observed that our diffusion models achieved the smallest value indicating the most smooth manipulation trajectories (S_r ours: 2.36 v.s. DAMON 6.25 v.s. IL 5.23 v.s. RRT 8.72 v.s. GPMP 6.87). Last, we also observed a potential trade-off using our proposed diffusion models. The training and sampling time (also as trajectory generation time) is slightly longer than RRT and GPMP methods, and comparable with DAMON and IL methods.

In total, the above experimental results showed that our proposed diffusion models have many promising advantages such as lower failure rate and smooth trajectory, and outperformed other algorithms.

C. Importance of obstacle guidance as real-time feedback to avoid obstacles and achieve alignment

In Methods section III-C, we emphasized that integrating the obstacle guidance technique is important as it functions as the real-time closed-loop feedback condition to guide the generation of safe and feasible trajectories. Therefore, we use an ablation study to comprehensively evaluate the functionality of the obstacle guidance component.

Similar as the prior section, we first provide the manipulation demos for dynamical obstacles. Here we focused on two algorithms, our proposed algorithm and our algorithm without the obstacle guidance component. From the demos (see Fig. 8(a)-(c)), we observed that if the obstacles were small, it was still possible for the no-guidance diffusion model to generate a safe and feasible manipulation trajectory. It can be understood if the sampler was lucky enough since the obstacles only occupied a small area in the configuration

space. However, when we increased the size and number of obstacles or even enabled dynamical obstacles, the chance level using the diffusion model without guidance dropped drastically, leading to the collision trajectories (see Figure 8(e)). On the contrary, our obstacle-guidance diffusion models can incorporate real-time video sensor images as the conditional guidance for the generation of future potential collision-free trajectories, which consistently resulted in lower failure ratios even in dynamical obstacle cases. We summarized the experimental results in Fig. 9, where our methods achieved a significantly lower failure ratio compared to the no-guidance diffusion model (Failure ratio in static for the largest obstacle: ours 6.56 v.s. no-guidance diffusion model 67.19; Failure ratio in dynamic and large obstacle: ours 9.36 v.s. no-guidance diffusion model 98.45). We further showed in Fig. 9 that the failure ratio for the no-guidance diffusion model increased as the obstacle volume increased. On the contrary, APEX showed consistently low failure ratio across static and dynamic obstacles. Thus, our ablation study confirmed the essential contribution of the obstacle guidance technique in generating safe and feasible ambidextrous dual-arm robotic manipulation trajectories.

D. Diversity of the trajectory generation

Last, we show two interesting findings. It is believed that by using different initial Gaussian noise, the diffusion models can generate various types of images or videos. Therefore, we also test our obstacle-guidance diffusion models by giving different initial Gaussian noise in the sampling process. We observed that our models indeed can generate multiple feasible and collision-free manipulation trajectories (see Fig. 10(a)), meaning that our model essentially learned complicated manipulation-task dependent distributions and realized effective multi-modal sampling from the learned distribution. Compared with other algorithms, our methods still have the most diverse trajectory generation performance. Further, we compute the FID for measurement of the generation diversity. Our FID results also confirmed that our proposed obstacle-guidance diffusion achieved the highest score across all the comparison algorithms (see Fig 10(b)). Second, regarding to the FID score comparison between our obstacle-guidance model and no-guidance diffusion model. We observed that our obstacle-guidance model has the lower FID score but faster trajectory generation speed (the FID score: ours 21.51 v.s. no-guidance diffusion model 48.63; the computational time ours 8.36s v.s. no-guidance diffusion model 19.63s), which revealed another trade-off that we sacrifice part of the generation diversity but both improve the sampling speed and ensure the feasibility and safety of the generated trajectories.

V. DISCUSSION AND CONCLUSION

Effectively coordinating dual-arm movements while controlling 5-finger robotic hands to perform manual tasks poses




Fig. 10: (a) Several collision-free trajectories are generated from APEX. The \oplus denotes the starting position and ∇ denotes the goal position. The upper and lower panels are in X and Y coordinations. (b) FID comparison. serious challenges for real-time execution in a cluttered environment. APEX successfully leverages deep generative models for robotic motion planning to synthesize diverse feasible trajectories precisely and accurately. This achievement leads to low failure rates and flexible planning trajectories, fundamentally differing from traditional robotic planning algorithms and prior imitation learning-based deep neural networks.

APEX introduces three major innovations. Firstly, our task distillation dramatically simplifies training costs and enhances algorithm efficiency, thus facilitating adaptive obstacle avoidance, which proves crucial in addressing dynamic obstacles. Secondly, the task demonstrations of APEX using latent diffusion models [23], mainly consisting of a Variational Autoencoder (VAE) and conventional U-Net structure, were trained and evaluated separately as the initial trial of our algorithms. Lastly, to further ensure that the generated trajectory is feasible and collision-free, we incorporate the obstacle-guidance technique by utilizing real-time obstacle and target information to construct conditional diffusion models.

APEX exhibits the capability to generate diverse feasible trajectories with lower failure rates for completing both ambidextrous vertical and horizontal dual-arm robotic manipulation tasks compared to state-of-the-art approaches. Given the generalizability of our proposed algorithm, we also foresee future explorations in more intricate tasks and robotic systems such as robotic dogs [36].

REFERENCES

- [1] M. Abbas *et al.*, "A systematic review on cooperative dual-arm manipulators: modeling, planning, control, and vision strategies," *International Journal of Intelligent Robotics and Applications*, vol. 7, no. 4, p. 683–707, Jul. 2023.
- [2] X. Zhao *et al.*, "A dual-arm robot cooperation framework based on a nonlinear model predictive cooperative control," *IEEE/ASME Transactions on Mechatronics*, pp. 1–13, 2023.
- [3] A. Dastider, H. Fang, and M. Lin, "Retro: Reactive trajectory optimization for real-time robot motion planning in dynamic environments," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 8764–8770.
- [4] A. Dastider and M. Lin, "Damon: Dynamic amorphous obstacle navigation using topological manifold learning and variational autoencoding," in *IROS*, 2023, pp. 251–258.
- [5] Y. He, M. Wu, and S. Liu, "A distributed optimal control framework for multi-robot cooperative manipulation in dynamic environments," *Journal of Intelligent and Robotic Systems*, vol. 105, no. 1, Apr. 2022. [Online]. Available: <http://dx.doi.org/10.1007/s10846-022-01621-4>
- [6] H. Fang and Y. Yang, "Designing and validating a robust adaptive neuromodulation algorithm for closed-loop control of brain states," *Journal of Neural Engineering*, vol. 19, no. 3, p. 036018, 2022.
- [7] H. Fang *et al.*, "Robust adaptive deep brain stimulation control of in-silico non-stationary parkinsonian neural oscillatory dynamics," *Journal of Neural Engineering*, 2024.
- [8] C. Zhou *et al.*, "Topp-mpc-based dual-arm dynamic collaborative manipulation for multi-object nonprehensile transportation," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 999–1005.
- [9] H. Fang and Y. Yang, "Predictive neuromodulation of cingulo-frontal neural dynamics in major depressive disorder using a brain-computer

interface system: A simulation study," *Frontiers in Computational Neuroscience*, vol. 17, p. 1119685, 2023.

- [10] Y. Huang *et al.*, "A differential dynamic programming-based approach for balancing energy and time optimality in motion planning," in *2023 59th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2023, pp. 1–7.
- [11] W. Shi *et al.*, "Obstacle avoidance path planning for the dual-arm robot based on an improved rrt algorithm," *Applied Sciences*, vol. 12, no. 8, 2022.
- [12] N. Vahrenkamp *et al.*, "Humanoid motion planning for dual-arm manipulation and re-grasping tasks," in *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009, pp. 2464–2470.
- [13] L. Liu *et al.*, "A collaborative control method of dual-arm robots based on deep reinforcement learning," *Applied Sciences*, vol. 11, no. 4, 2021.
- [14] M. Alles and E. Aljalbout, "Learning to centralize dual-arm assembly," *Frontiers in Robotics and AI*, vol. 9, p. 830007, 2022.
- [15] F. Xie *et al.*, "Deep imitation learning for bimanual robotic manipulation," in *Advances in Neural Information Processing Systems*, H. Larochelle *et al.*, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 2327–2337.
- [16] H. Kim, Y. Ohmura, and Y. Kuniyoshi, "Goal-conditioned dual-action imitation learning for dexterous dual-arm robot manipulation," *T-RO*, pp. 1–20, 2024.
- [17] —, "Transformer-based deep imitation learning for dual-arm robot manipulation," in *IROS*, 2021, pp. 8965–8972.
- [18] F.-A. Croitoru *et al.*, "Diffusion models in vision: A survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 9, pp. 10 850–10 869, 2023.
- [19] J. Ho *et al.*, "Video diffusion models," in *Advances in Neural Information Processing Systems*, vol. 35. Curran Associates, Inc., 2022, pp. 8633–8646.
- [20] P. Dhariwal and A. Nichol, "Diffusion models beat gans on image synthesis," in *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., vol. 34. Curran Associates, Inc., 2021, pp. 8780–8794.
- [21] A. Dastider, H. Fang, and M. Lin, "Unified control framework for real-time interception and obstacle avoidance of fast-moving objects with diffusion variational autoencoder," *arXiv preprint arXiv:2209.13628*, 2022.
- [22] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.
- [23] R. Rombach *et al.*, "High-resolution image synthesis with latent diffusion models," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 10 684–10 695.
- [24] O. Ronneberger, P. Fischer, and T. 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*. Springer, 2015, pp. 234–241.
- [25] I. Kapelyukh, V. Vosylius, and E. Johns, "Dall-e-bot: Introducing web-scale diffusion models to robotics," *IEEE Robotics and Automation Letters*, vol. 8, no. 7, pp. 3956–3963, 2023.
- [26] T. Power *et al.*, "Sampling constrained trajectories using composable diffusion models," in *IROS 2023 Workshop on Differentiable Probabilistic Robotics: Emerging Perspectives on Robot Learning*, 2023.
- [27] M. Janner *et al.*, "Planning with diffusion for flexible behavior synthesis," in *International Conference on Machine Learning*, 2022.
- [28] J. Carvalho *et al.*, "Motion planning diffusion: Learning and planning of robot motions with diffusion models," in *IROS*, 2023, pp. 1916–1923.
- [29] Z. Sun *et al.*, "Bifa-yolo: A novel yolo-based method for arbitrary-oriented ship detection in high-resolution sar images," *Remote Sensing*, vol. 13, no. 21, 2021.
- [30] C. Chi *et al.*, "Diffusion policy: Visuomotor policy learning via action diffusion," *ArXiv*, vol. abs/2303.04137, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:257378658>
- [31] M. Reuss *et al.*, "Goal conditioned imitation learning using score-based diffusion policies," in *Robotics: Science and Systems*, 2023.
- [32] J. Ho and T. Salimans, "Classifier-free diffusion guidance," *arXiv preprint arXiv:2207.12598*, 2022.
- [33] A. H. Khan *et al.*, "Obstacle Avoidance and Tracking Control of Redundant Robotic Manipulator: An RNN-Based Metaheuristic Approach," *IEEE Trans. Industr. Inform.*, vol. 16, no. 7, pp. 4670–4680, 2020.
- [34] O. Adiyatov and H. A. Varol, "A novel rrt*-based algorithm for motion planning in dynamic environments," in *2017 IEEE International Conference on Mechatronics and Automation*, 2017, pp. 1416–1421.
- [35] M. Mukadam, X. Yan, and B. Boots, "Gaussian process motion planning," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 9–15.
- [36] Y. Zhao *et al.*, "Intelligent control of multilegged robot smooth motion: A review," *IEEE Access*, vol. 11, pp. 86 645–86 685, 2023.