

# Learning Diffusion Policies from Demonstrations For Compliant Contact-rich Manipulation

Malek Aburub<sup>1,2\*</sup>, Cristian C. Beltran-Hernandez<sup>1\*</sup>, Tatsuya Kamijo<sup>1</sup>, Masashi Hamaya<sup>1</sup>

**Abstract**—Robots hold great promise for performing repetitive or hazardous tasks, but achieving human-like dexterity, especially in contact-rich and dynamic environments, remains challenging. Rigid robots, which rely on position or velocity control, often struggle with maintaining stable contact and applying consistent force in force-intensive tasks. Learning from Demonstration has emerged as a solution, but tasks requiring intricate maneuvers, such as powder grinding, present unique difficulties. This paper introduces Diffusion Policies For Compliant Manipulation (DIPCOM), a novel diffusion-based framework designed for compliant control tasks. By leveraging generative diffusion models, we develop a policy that predicts Cartesian end-effector poses and adjusts arm stiffness to maintain the necessary force. Our approach enhances force control through multimodal distribution modeling, improves the integration of diffusion policies in compliance control, and extends our previous work by demonstrating its effectiveness in real-world tasks. We present a detailed comparison between our framework and existing methods, highlighting the advantages and best practices for deploying diffusion-based compliance control.

## I. INTRODUCTION

Robots hold significant potential to enhance daily life by performing repetitive or hazardous tasks that require dexterous manipulation. A particularly challenging class of tasks known as contact-rich manipulation refers to those involving sustained physical contact between the robot’s end-effector and objects, where precise and dynamic force modulation is essential to achieve task success. Examples of contact-rich manipulation include powder grinding, pencil erasing, and precise peg insertion—tasks characterized by continuous interaction forces that must be carefully regulated over extended periods. Achieving human-like dexterity in such conditions remains difficult, especially for rigid robots, which typically operate through position or velocity commands alone.

Learning from Demonstration (LfD) techniques [1] have emerged as a promising approach, enabling robots to acquire complex tasks by observing human experts and replicating their precise actions. However, tasks that require intricate maneuvers and high force, such as powder grinding in Fig. 1, pose unique challenges. In such scenarios, rigid robots often fail to sustain controlled contact and force over extended periods, limiting their effectiveness in demanding applications.

\* Equal contribution.

<sup>1</sup>OMRON SINIC X Corporation, Tokyo, Japan

<sup>2</sup>Department of Engineering Science, Osaka University, Japan

Corresponding authors:

cristian.beltran [at] sinicx.com

This work was partly supported by the JST-Mirai Program (Grant Number JPMJMI21G2).

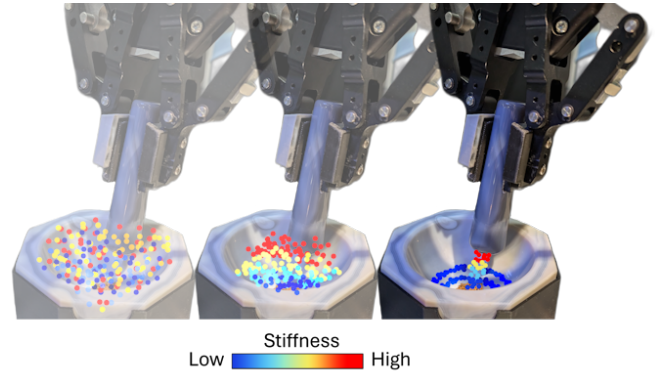


Fig. 1: **DIPCOM Denoising Process**: illustration of how DIPCOM predicts end-effector positions and adjusts arm stiffness during multiple timesteps, enabling consistent force during contact and smooth transitions between movements.

In such cases, rigid robots necessitate mechanical compliance mechanisms [2] to make the physical interaction safer at the expense of making it harder to precisely control the position of the tool, i.e., the pestle.

To address these challenges, we adopt compliance control schemes that enable rigid robots to handle tasks involving direct contact by modulating forces through external sensors [3]. In our prior work [4], a method called Comp-ACT was proposed that combines force information and compliance control with a policy based on Variational AutoEncoders and the Action Chunking with Transformers strategy (VAE-ACT). The Comp-ACT policies perform well in various tasks but struggle to solve long-horizon tasks with repetitive behaviors, such as grinding powder. Therefore, we extend compliant LfD policies by incorporating diffusion models, enabling robots to adaptively regulate force during task execution while maintaining precision and stability, especially during long-horizon tasks. Our approach builds on the robot demonstration system for contact-rich manipulation introduced in [4], improving its performance in handling force-intensive tasks.

In this work, we propose Diffusion Policies For Compliant Manipulation (DIPCOM), a novel diffusion-based framework that leverages generative models to handle compliant control tasks. Our choice of diffusion-based policies is motivated by their potential to capture multi-modal action distributions and produce a more diverse behavior than VAE-ACT-based policies [5]. Diffusion models progressively add noise to data in a forward stochastic process. By training a neural network

known as a denoising model, they approximate the original data and iteratively generate new actions. We implement diffusion models into compliance control to develop a policy capable of predicting both the Cartesian end-effector (EE) pose and simultaneously adjusting the stiffness of the robot arm to apply the necessary force for the task.

To summarize our contributions:

- A novel diffusion-based framework for rigid robots to learn complex contact-rich manipulations from demonstrations via compliance control. We show that the expressiveness of diffusion models, along with their ability to capture multimodal distributions, enhances force control.
- A comprehensive comparison of DIPCOM and Comp-ACT [4] on challenging real-world tasks, highlighting the specific tasks where diffusion-based policies excel and outlining best practices for their application.

## II. RELATED WORKS

### A. Learning from Demonstrations for Contact-rich Manipulation

Learning from demonstration has emerged as a promising approach for teaching robots complex contact-rich manipulation skills [6]. Researchers have explored using force/torque (F/T) sensing and haptic feedback to capture expert demonstrations to learn tasks such as grasping [7], ironing [8], pouring [9], and peg-in-hole insertion [10]. Recent work has focused on sample-efficient methods that can learn from a few demonstrations through transformer-based models and innovative teleoperation interfaces such as ALOHA [11] and Universal Manipulation Interface (UMI) [12].

While most prior approaches focus on position control and mechanical compliance to allow robots a degree of safety while performing contact-rich tasks, this work focuses on active compliance control to perform contact-rich manipulation tasks.

In our prior work [4], we proposed an intuitive teleoperation interface for collecting demonstrations and the Compliance Control via Action Chunking with Transformers (Comp-ACT) method to learn compliance control policies from a few demonstrations. The current study utilizes the same teleoperation interface to collect demonstrations for training a DIPCOM policy.

Similar to the work of Drolet et al. [13] and Zhao et al. [14], where several imitation learning methods, including Action Chunking with Transformers (ACT) [11] and Diffusion policies [15], were compared on bimanual manipulation tasks, this study compares the performance of Comp-ACT [4] with the proposed DIPCOM with a focus on rigid robots tackling challenging real-world contact-rich tasks.

### B. Diffusion Policies

Diffusion models, introduced by Ho et al. [16], are probabilistic generative models that transform random noise into meaningful samples from a target distribution. For a more comprehensive survey on diffusion models, see [17].

In robotics, diffusion models have effectively captured multi-modal actions and have been applied to various domains, including motion planning [18], navigation [19], human-robot interaction [20], and grasping tasks [21]. In manipulation, Chi et al. [15] demonstrated strong results in visuomotor policy learning from demonstrations using diffusion models.

Further advances include Liu et al. [22], who leveraged language-annotated play data to enable skill acquisition, while Reuss et al. [23] conditioned diffusion models on goal states to enhance skill learning from play data for task completion.

Recent developments in diffusion-based and compliance control robotic methods have explored diverse approaches tailored to specific task domains and control abstractions. Methods such as TacDiffusion [24] focus on direct force (wrench) prediction for precision insertion tasks, while Reactive Diffusion Policy [25] integrates hierarchical slow-fast control with tactile feedback to enable reactive manipulation. Other approaches, including ForceMimic [26] and Adaptive Compliance Policy [27], emphasize force-centric learning or continuous spatial-temporal compliance adaptation, often using different data collection modalities or targeting deformable object manipulation. Our work, DIPCOM, uniquely combines diffusion-based generative modeling with discrete stiffness mode compliance control to address long-horizon repetitive contact-rich tasks.

## III. METHODOLOGY

We introduce Diffusion Policies For Compliant Manipulation (DIPCOM), a novel method for learning variable compliance control from demonstrations using diffusion models. Our approach predicts target EE poses and robot stiffness parameters conditioned on current observations, including the contact force data. A compliance controller uses these predictions to compute the final joint position commands that allow robots to move compliantly at the predicted stiffness. Fig. 2 illustrates the architecture of DIPCOM.

### A. Problem Formulation

Learning from demonstration (LfD) aims to enable robots to acquire new skills by autonomously observing and imitating human-provided demonstrations. In our approach, the policy learns to predict a sequence of absolute Cartesian EE pose and stiffness parameters given current observation  $\mathcal{O}$  that include RGB images  $\mathcal{I} \in \mathbb{R}^{H \times W \times 3}$ , the latest F/T sensor reading  $\mathcal{F} \in \mathbb{R}^6$ , and proprioception data  $\mathcal{S} \in \mathbb{R}^9$ . The predicted action  $\mathcal{A} = \{\mathbf{p}, \mathbf{g}, \mathbf{k}\}$  comprises three components: the absolute EE pose, the gripper action, and the stiffness parameter. The absolute EE pose, denoted as  $\mathbf{p} = \{\mathbf{r}, \mathbf{o}\} \in \mathbb{R}^9$ , consists of a position vector  $\mathbf{r} \in \mathbb{R}^3$  and an orientation vector  $\mathbf{o} \in \mathbb{R}^6$ . For more details on this 6D orientation representation, refer to Sect. III-C. The gripper action  $\mathbf{g} \in \mathbb{R}^1$  represents the desired non-binary opening width of the gripper. Finally, the stiffness parameter  $\mathbf{k} \in \mathbb{R}^6$  corresponds to the diagonal elements of the stiffness matrix. Each robot action  $\mathcal{A}$  is then fed into a compliance controller.

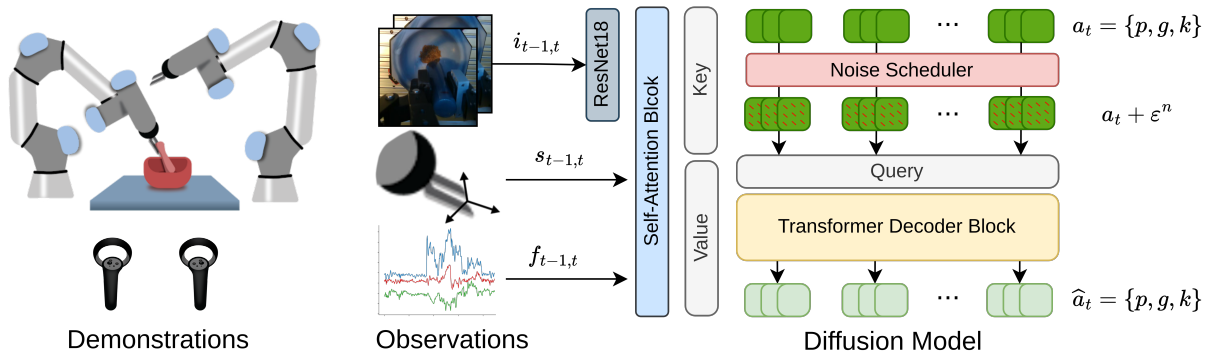


Fig. 2: **Policy Framework:** Left: Dataset collection framework. Middle: Observations  $O$  include images  $i_{t-1,t}$ , robot Cartesian pose  $s_{t-1,t}$ , and measured force/torque  $f_{t-1,t}$ , all encoded using a self-attention transformer. Right: During training, actions  $a_t$ —comprising the end-effector pose  $p$ , gripper pose  $g$ , and stiffness  $K$ —are processed through a noise scheduler that adds Gaussian noise  $\epsilon$  over time steps  $n$ . These noisy actions are then input into the transformer decoder block. During inference, Gaussian noise replaces the training noise, and the transformer decoder block predicts the actions  $\hat{a}_t$

### B. Data Collection

The teleoperation system presented in [4] was used for data collection. The system uses Virtual Reality (VR) controllers to provide the reference EE pose and stiffness that are input to the compliance controller. During the demonstration, the operator can switch between two pre-selected stiffness modes using the grip button on the side of the controller. The action  $\mathcal{A}$  from the VR controllers and the observation  $\mathcal{O}$  from the robots and cameras are collected and stored.

### C. Orientation Representation for Cartesian Action Space

As in our previous work [4], we represent the robot’s state observations and actions in absolute Cartesian pose. In [4], orientations were encoded using axis-angle (rotation vector) representations, but this choice has two drawbacks: (1) multiple rotation vectors can represent the same orientation, requiring the model to handle all equivalent cases, and (2) the representation is inherently discontinuous. As noted by Zhou et al. [28], neural networks perform better with continuous representations, and they introduced 5D and 6D continuous alternatives. To address the issues of ambiguity and discontinuity, we adopt the 6D orientation representation from [28], which uses the first two columns of the rotation matrix and provides a continuous, unique encoding for each orientation. For further details, please refer to [28].

### D. Diffusion Policies For Compliant Manipulation

Our dataset is inherently multi-modal, containing a variety of observations and the corresponding actions that must be predicted. To address this, we aim to learn a policy distribution  $\pi(\mathcal{A}|I, F, S)$  from a task-specific demonstration dataset. We introduce Diffusion Policies For Compliant Manipulation, a classifier-free conditional diffusion model designed to generate actions  $\mathcal{A}$  based on observations  $\mathcal{O}$ .

The diffusion process involves two stages: a forward process that incrementally adds noise to the data and a learnable inverse diffusion process that recovers the original data from noise conditioned on the input observations. The

forward process progressively corrupts the data, while the inverse process, learned by the model, removes the noise step by step.

We implement the Denoising Diffusion Implicit Model (DDIM) formulation from Song et al. [29] for noise scheduling and denoising. This approach makes the denoising process deterministic, facilitating more efficient inference and allowing for flexible adjustment of inference steps to optimize performance.

The architecture of DIPCOM is shown in Fig.2. It follows the design of [15], using an encoder-decoder transformer [30] with a ResNet18 vision backbone (without pre-trained weights). Images are processed through the vision backbone, concatenated with force and robot state data, and input into the transformer’s encoder. Cross-attention is applied to the noisy actions in the decoder.

The model is trained to predict actions  $\hat{a}_0$ , using the mean squared error loss function:

$$L_{\text{sample}} = \|a_t^0 - \hat{a}_t^0\|^2$$

During inference, the model iteratively predicts the original sample using the formula:

$$a_t^{n-1} = \sqrt{\beta_{n-1}} \hat{a}_t^0 + \sqrt{1 - \beta_{n-1}} \cdot \frac{a_t^n - \sqrt{\beta} \hat{a}_t^n}{\sqrt{1 - \beta}}$$

where  $a_t^n$  represents the noisy data at time step  $n$ ,  $\beta$  is the cumulative product of noise scales up to time  $n$ , and  $\hat{a}_t^0$  is the estimated original data.

### E. Action Sequence Generation

Applying diffusion models to contact-rich manipulation presents unique challenges due to the robot’s prolonged interaction with the object’s surface. In such tasks, variable forces are exerted, and any inconsistencies in action prediction can negatively impact performance. Previous works [15], [12] typically predict a fixed number of actions, using part of the horizon while discarding the rest. This can lead to instability, such as jerky or abrupt movements,




Fig. 3: Contact-rich manipulation tasks used for evaluation. A - Powder grinding. B - Pencil eraser. C - Bimanual round peg insertion. D - Bimanual cuboid peg insertion.

especially during transitions between inference steps. While this might be acceptable for simpler tasks like pick-and-place or rearranging, it becomes problematic for contact-rich manipulation.

To address this, we operate at a higher action-prediction frequency than previous works, which typically cap predictions at around 16 actions per horizon and run at a maximum of 20 Hz. In contrast, we predict substantially longer horizons, averaging 72 steps while applying 32 of them, and tune the execution frequency per task, typically around 30 Hz. This larger prediction horizon, combined with a higher control frequency, enables finer temporal resolution and more precise control in contact-rich scenarios.

Another challenge in using diffusion over longer horizons is the accumulation of error, especially when only a portion of the predicted actions is applied, which can lead to instability. To mitigate this, we utilize the Temporal Ensemble introduced in [11] for ACT and extend it to diffusion: the unused portion of each predicted horizon is carried forward and blended with the next inference, smoothing transitions and minimizing erratic movements for stable performance in contact-rich tasks.

#### IV. EXPERIMENTS

The proposed method DIPCOM was evaluated on challenging contact-rich tasks that require applying force carefully and consistently to solve them. The tasks are described in Sect. IV-C and illustrated in Fig. 3. Additionally, we compare the performance of the proposed method against a baseline described below.

##### A. Baseline Method

The prior work, Comp-ACT [4], was used as the baseline method. The baseline policy is trained as the decoder of a conditional variational autoencoder (CVAE). It consists of the transformer encoder, which synthesizes all the observational data, and the transformer decoder, which generates a sequence of actions, namely Cartesian EE pose and stiffness

TABLE I: Task conditions

Task	Task Conditions						
	# of demos.	# of views	# of arms	Stiffness Modes			
				Position		Rotation	
low	high	low	high				
A	40	1	1	300	800	100	150
B	60	2	1	800	1200	150	300
C / D	60	3	2	R1: 800	1200	150	300
				R2: 200	800	100	150

parameters. Similarly to our approach, the baseline uses F/T observations, Cartesian EE poses, and RGB images to predict actions that are passed to the compliance controller. Comp-ACT policies were run at a control frequency of 50 Hz.

##### B. Experimental Setup

The robotic system used for experimentation consisted of two UR5e robot arms (Universal Robots A/S, 2024) with built-in F/T sensors on their wrists. Two cameras, RealSense SR305 (Intel Corporation, 2024), are attached to the wrist of each arm. Another static third-person view camera is placed in front of the robots to capture a wider field of view. Both the baseline and the proposed method were trained and evaluated on a PC with AMD Ryzen™ 9 7900 CPU and Nvidia RTX 4080 GPU.

The demonstration data from each task were collected from three co-authors to add variety to the datasets, and all demonstrations were recorded at 50 Hz.

##### C. Tasks description and results

This section describes the tasks used for experimentation alongside the results obtained after training one policy per task for each model. The general conditions considered for each task are reported in Table I, such as the number of demonstrations and stiffness modes.

**A - Powder grinding:** The task starts with the robot already holding the ceramic pestle and the granular powder placed at the center of the ceramic mortar. A single camera was




Fig. 4: Powder Grinding performance by the DIPCOP policy. The policy imitates the demonstrated behavior of pausing every few seconds to look at the powder’s state before continuing the grinding process. Position Z indicates the height of the tip of the pestle relative to the mortar.




Fig. 5: A - Powder Grinding: Force profile comparison between the demonstrations, the proposed method DIPCOP, and the baseline Comp-ACT. A bold line and color-shaded area represent the average and standard deviation normal contact force, respectively.

TABLE II: A - Powder grinding results

	Percentage of fine powder produced	
	Average	Standard Deviation
Human demonstrations	76.67%	8.4%
DIPCOP	<b>55.88%</b>	13.54%
CompACT	9.96%	1.39%

used for this task, the one attached to the robot’s wrist, as the robot would obstruct any other external view of the inside of the mortar. For this reason, the demonstrations consisted of pressing the pestle against the mortar and performing circular motions for a few seconds, then moving the robot out of the mortar to get a clear view of the state of the powder, and then repeating this process for about 80 seconds. We measured the fine powder produced after each demonstration and compared it against the policy performance for this task. For practical purposes, instant coffee powder was used for experimentation as it comes in a granular form that can be ground to a finer powder. Similarly, a tea strainer was used to sift the powder.

**Results:** As shown in Table II, our proposed method achieved the highest percentage of fine powder produced

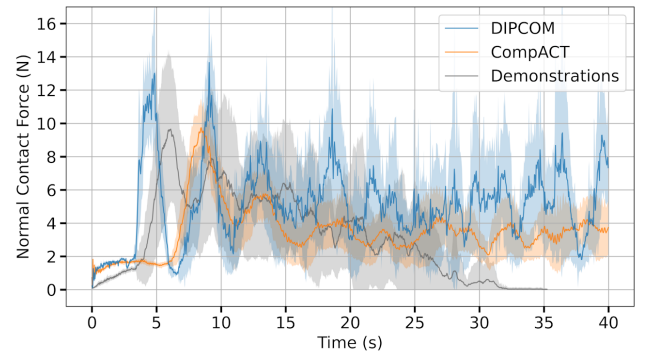


Fig. 6: B - Pencil Eraser. Force profile comparison between the demonstrations, the proposed method DIPCOP, and the baseline Comp-ACT.

TABLE III: B - Pencil eraser task results

Method	Percentage Erased Average (SD)	Success Rate
Comp-ACT	26.0% (16.6%)	0.0%
DIPCOP	<b>77.32% (19.48%)</b>	<b>52.3%</b>

at 56 % compared to the 10 % achieved by Comp-ACT. Both methods applied a similar magnitude of force against the powder and mortar, as illustrated in Fig. 5. Nevertheless, our DIPCOP policy obtained better results by reproducing the demonstrated behavior of moving the pestle in circular motions and periodically checking the state of the powder before repeating the grinding action, as shown in Fig. 4. On the contrary, the Comp-ACT policy struggled to reproduce the circular motions with the pestle, mostly staying in the same place during the entire duration of each test.

**B - Pencil eraser:** The task’s goal is to use a rubber eraser to remove pencil markings from a notepad. The task begins with the word "OSX" written on the notepad fixed to the table while the robot already holds the rubber eraser. The robot has to gently but firmly press the eraser against the paper to remove the pencil marks without damaging the paper. During data collection, the demonstrator rubbed

TABLE IV: Success rate for bimanual insertion tasks

Task Name	Comp-ACT	DIPCOM
C - Bimanual Round Insertion	100%	100%
D - Bimanual Cuboid Insertion	95%	95%

the eraser from right to left in a straight line, then lifted the eraser from the paper, moved back to the right side of the marks, and repeated several times until the mark was removed completely. The task was evaluated with two metrics. First, the success rate where success was defined by whether the pencil marks were completely erased or not. Second, we compute the percentage of erased pencil marks by analyzing the similarity of the paper’s initial and final state after each rollout. To measure how much of the pencil marks were erased, we take photos of the paper before and after the robot erases it. The photos are converted to black and white to highlight the pencil marks. The dark pixels represent the pencil marks. These results are reported in Table III.

**Results:** As shown in Fig. 6, the Comp-ACT policy tends to apply less overall force to the paper but struggles to completely erase the marks on the paper. On the contrary, our DIPCOM policy applied more force to the paper while simultaneously aiming better toward the pencil marks. As a result, our proposed method achieved a 52.3% success rate, while Comp-ACT could not complete the task on any of the 20 rollouts. Additionally, on average, our proposed method can erase 51% more marks than Comp-ACT.

**C / D - Bimanual insertion tasks:** The task begins with one arm already holding the peg while the other arm similarly holds the mating part. The goal is to align, insert, and release the peg. The peg and hole have a tolerance of 2 mm. The peg insertion task is already challenging, but coordinating both arms to solve the task significantly increases the difficulty. Success is defined as the correct insertion and release of the peg for these tasks.

**Results:** As shown in Table IV, both methods achieve a similar performance, where the round peg insertion was successful every single time, but both models struggle to solve the cuboid peg insertion 5% of the times. Nevertheless, the behavior displayed by each method is very different. The Comp-ACT policy consistently approaches the task in the same manner. Conversely, the DIPCOM policy seemingly attempts different trajectories, similar to the diversity of demonstrations provided. Interestingly, while the demonstrators typically followed a consistent strategy of fixing the pose of the mating part with one arm before inserting the peg with the other, DIPCOM introduced motor skills not seen in the demonstrations. The DIPCOM policy continuously adjusted the supporting arm during insertion, assisting the peg-holding arm to ensure smooth and safe task execution. In contrast, Comp-ACT stuck rigidly to the demonstrated strategy.

## V. DISCUSSION

As mentioned above, DIPCOM and Comp-ACT achieved similar overall performance in the bimanual insertion tasks, but their execution behaviors differed noticeably. These differences were more evident in longer-horizon tasks—such as powder grinding and pencil erasure—that require repetitive, non-linear motions driven by recurring observations. Comp-ACT initially executed these tasks correctly but often failed to sustain smooth, repetitive movements, where it frequently froze during actions like the up-and-down motion for erasing or circular motions for grinding. In contrast, DIPCOM maintained fluid and adaptive behavior throughout, despite showing slightly higher variance. This contrast aligns with findings by Jia et al. [5], who reported that diffusion models preserve multimodal action distributions through iterative denoising, whereas VAE-based policies compress behaviors through a latent bottleneck that can suppress valid strategy variations. In long-horizon tasks where demonstrations contain behavioral variation (e.g., different grinding patterns or inspection timing), diffusion better retains this diversity, enabling more flexible execution.

Another observation was about force application during these tasks. Both policies applied force more conservatively than the human demonstrators. However, despite having a larger standard deviation, DIPCOM’s diffusion-based policy better mimicked the demonstrators’ force patterns. Comp-ACT, on the other hand, applied force more consistently but tended to fall into cyclic force patterns, especially in long-horizon tasks.

**Limitations:** During training and fine-tuning, we observed that DIPCOM is considerably more sensitive to hyperparameter choices than Comp-ACT. Diffusion-based policies are computationally demanding, which makes them more susceptible to fluctuations in the robot’s control frequency and can degrade performance if not tuned carefully. Increasing the prediction horizon improves the effective operating frequency by reducing how often inference must be performed, but it also makes the system more sensitive to the choice of horizon. Likewise, the temporal ensemble requires careful balancing: large ensemble windows introduce latency and slow responsiveness, while small ones provide limited smoothing and can result in jittery behavior. This heightened sensitivity makes DIPCOM more challenging to stabilize than Comp-ACT, motivating future work on more systematic hyperparameter tuning and on exploring alternative action representations, such as the relative trajectory formulation proposed in [12].

It is important to highlight that, in this study, the number of demonstrations used to train the policies was relatively small compared to other works exploring similar methods [15][14]. An interesting future research avenue is to explore the effects of large-scale datasets for compliant manipulation tasks, aiming to improve policy generalization across extended task variations. For instance, instead of training a policy for each specific powder grinding task, we aim to develop a single policy that can perform across different grinding tasks,

broadening its applicability and robustness.

## VI. CONCLUSIONS

This study introduced Diffusion Policies For Compliant Manipulation (DIPCOM), a diffusion-based framework designed for compliant control tasks, particularly for rigid robots. Our approach demonstrates how diffusion policies effectively capture the multimodality of the data, predicting Cartesian end-effector poses while adjusting the arm's stiffness to apply the required contact forces. We provide guidelines for implementing diffusion policies for compliant control and outline best practices for optimizing performance. Through extensive experimental evaluations, we showcased DIPCOM's strengths in various contact-rich tasks, highlighting its advantages over previous methods. Moving forward, we aim to explore the policy's ability to generalize across different task variations by expanding the dataset. We plan to enhance force processing and refine the policy architecture to improve force-aware inference.

## REFERENCES

- [1] H. Ravichandar, A. S. Polydoros, S. Chernova, and A. Billard, "Recent advances in robot learning from demonstration," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, pp. 297–330, 2020. [Online]. Available: <https://doi.org/10.1146/annurev-control-100819-063206>
- [2] Y. Nakajima, M. Hamaya, Y. Suzuki, T. Hawaii, F. von Drigalski, K. Tanaka, Y. Ushiku, and K. Ono, "Robotic powder grinding with a soft jig for laboratory automation in material science," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 2320–2326.
- [3] A. Calanca, R. Muradore, and P. Fiorini, "A review of algorithms for compliant control of stiff and fixed-compliance robots," *IEEE/ASME Transactions on Mechatronics*, vol. 21, no. 2, pp. 613–624, 2015.
- [4] T. Kamijo, C. C. Beltran-Hernandez, and M. Hamaya, "Learning variable compliance control from a few demonstrations for bimanual robot with haptic feedback teleoperation system," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024.
- [5] X. Jia, D. Blessing, X. Jiang, M. Reuss, A. Donat, R. Lioutikov, and G. Neumann, "Towards diverse behaviors: A benchmark for imitation learning with human demonstrations," in *International Conference on Learning Representations (ICLR)*, 2024. [Online]. Available: <https://openreview.net/forum?id=6pPYRXKPPw>
- [6] A. Billard, S. Calinon, R. Dillmann, and S. Schaal, "Robot programming by demonstration," in *Springer Handbook of Robotics*, 2008, pp. 1371–1394. [Online]. Available: [https://doi.org/10.1007/978-3-540-30301-5\\_60](https://doi.org/10.1007/978-3-540-30301-5_60)
- [7] A. M. Schmidts, D. Lee, and A. Peer, "Imitation learning of human grasping skills from motion and force data," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2011, pp. 1002–1007.
- [8] P. Kormushev, S. Calinon, and D. G. Caldwell, "Imitation learning of positional and force skills demonstrated via kinesthetic teaching and haptic input," *Advanced Robotics*, vol. 25, no. 5, pp. 581–603, 2011.
- [9] L. Rozo, P. Jiménez, and C. Torras, "A robot learning from demonstration framework to perform force-based manipulation tasks," *Intelligent Service Robotics*, vol. 6, no. 1, pp. 33–51, 2013.
- [10] Y. Wang, C. C. Beltran-Hernandez, W. Wan, and K. Harada, "Robotic imitation of human assembly skills using hybrid trajectory and force learning," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 11 278–11 284.
- [11] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning fine-grained bimanual manipulation with low-cost hardware," in *Robotics: Science and Systems (RSS)*, 2023. [Online]. Available: <https://doi.org/10.15607/RSS.2023.XIX.016>
- [12] C. Chi, Z. Xu, C. Pan, E. Cousineau, B. Burchfiel, S. Feng, R. Tedrake, and S. Song, "Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots," *arXiv preprint arXiv:2402.10329*, 2024.
- [13] M. Drolet, S. Stepputtis, S. Kailas, A. Jain, J. Peters, S. Schaal, and H. Ben Amor, "A comparison of imitation learning algorithms for bimanual manipulation," *IEEE Robotics and Automation Letters (RA-L)*, vol. 9, no. 10, pp. 8579–8586, 2024.
- [14] T. Z. Zhao, J. Tompson, D. Driess, P. Florence, S. K. S. Ghasemipour, C. Finn, and A. Wahid, "ALOHA Unleashed: A simple recipe for robot dexterity," in *Annual Conference on Robot Learning (CoRL)*, 2024.
- [15] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," in *Robotics: Science and Systems (RSS)*, 2023. [Online]. Available: <https://doi.org/10.15607/RSS.2023.XIX.026>
- [16] J. Ho, A. Jain, and P. Abbeel, "Denosing diffusion probabilistic models," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, pp. 6840–6851, 2020.
- [17] L. Yang, Z. Zhang, Y. Song, S. Hong, R. Xu, Y. Zhao, W. Zhang, B. Cui, and M.-H. Yang, "Diffusion models: A comprehensive survey of methods and applications," *ACM Computing Surveys*, vol. 56, no. 4, pp. 1–39, 2023.
- [18] J. Carvalho, A. T. Le, M. Baierl, D. Koert, and J. Peters, "Motion planning diffusion: Learning and planning of robot motions with diffusion models," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 1916–1923.
- [19] H. Ryu, J. Kim, H. An, J. Chang, J. Seo, T. Kim, Y. Kim, C. Hwang, J. Choi, and R. Horowitz, "Diffusion-edfs: Bi-equivariant denoising generative modeling on se (3) for visual robotic manipulation," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024, pp. 18 007–18 018.
- [20] E. Ng, Z. Liu, and M. Kennedy, "Diffusion co-policy for synergistic human-robot collaborative tasks," *IEEE Robotics and Automation Letters (RA-L)*, vol. 9, no. 1, pp. 215–222, 2024.
- [21] 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 *IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 5923–5930.
- [22] W. Liu, Y. Du, T. Hermans, S. Chernova, and C. Paxton, "Structdiffusion: Language-guided creation of physically-valid structures using unseen objects," in *Robotics: Science and Systems (RSS)*, 2023. [Online]. Available: <https://doi.org/10.15607/RSS.2023.XIX.031>
- [23] M. Reuss, M. Li, X. Jia, and R. Lioutikov, "Goal-conditioned imitation learning using score-based diffusion policies," in *Robotics: Science and Systems (RSS)*, 2023. [Online]. Available: <https://doi.org/10.15607/RSS.2023.XIX.028>
- [24] Y. Wu, Z. Chen, F. Wu, L. Chen, L. Zhang, Z. Bing, A. Swikir, S. Haddadin, and A. Knoll, "Tacdiffusion: Force-domain diffusion policy for precise tactile manipulation," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 11 831–11 837.
- [25] H. Xue, J. Ren, W. Chen, G. Zhang, Y. Fang, G. Gu, H. Xu, and C. Lu, "Reactive diffusion policy: Slow-fast visual-tactile policy learning for contact-rich manipulation," *arXiv preprint arXiv:2503.02881*, 2025.
- [26] W. Liu, J. Wang, Y. Wang, W. Wang, and C. Lu, "Forcemimic: Force-centric imitation learning with force-motion capture system for contact-rich manipulation," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 1105–1112.
- [27] Y. Hou, Z. Liu, C. Chi, E. Cousineau, N. Kuppusswamy, S. Feng, B. Burchfiel, and S. Song, "Adaptive compliance policy: Learning approximate compliance for diffusion guided control," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 4829–4836.
- [28] Y. Zhou, C. Barnes, J. Lu, J. Yang, and H. Li, "On the continuity of rotation representations in neural networks," in *IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, 2019, pp. 5745–5753.
- [29] J. Song, C. Meng, and S. Ermon, "Denosing diffusion implicit models," in *International Conference on Learning Representations (ICLR)*, 2021.
- [30] A. Vaswani, "Attention is all you need," *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.