

Adaptive Compliance Policy: Learning Approximate Compliance for Diffusion Guided Control

Yifan Hou¹ Zeyi Liu¹ Cheng Chi¹ Eric Cousineau² Naveen Kuppaswamy²

Siyuan Feng² Benjamin Burchfiel² Shuran Song¹

¹ Stanford University ² Toyota Research Institute

adaptive-compliance.github.io




Fig. 1: **Compliance Requirements.** [Left] Flipping an item requires the robot to follow an arc trajectory (blue) while maintaining contact force. This demands low stiffness in pushing directions (K_2) and high stiffness elsewhere (K_1). [Right] Wiping a vase necessitates 3D compliance adjustments in both end-effectors to 1) hold the vase, 2) trace the marking, and 3) apply appropriate force without damage. Our algorithm aims to model these spatial-, temporal-, and task-dependent compliance requirements from human demonstration data.

Abstract—Compliance plays a crucial role in manipulation, as it balances between the concurrent control of position and force under uncertainties. Yet compliance is often overlooked by today’s visuomotor policies that solely focus on position control. This paper introduces Adaptive Compliance Policy (ACP), a novel framework that learns to dynamically adjust system compliance both spatially and temporally for given manipulation tasks from human demonstrations, improving upon previous approaches that rely on pre-selected compliance parameters or assume uniform constant stiffness. However, computing full compliance parameters from human demonstrations is an ill-defined problem. Instead, we estimate an approximate compliance profile with two useful properties: avoiding large contact forces and encouraging accurate tracking. Our approach enables robots to handle complex contact-rich manipulation tasks and achieves over 50% performance improvement compared to state-of-the-art visuomotor policy methods. Project website with result videos: adaptive-compliance.github.io.

I. INTRODUCTION

Manipulation often requires the concurrent control of both position and force to achieve the desired outcome. This joint objective can be captured by the concept of mechanical compliance [32, 37, 14], where a low compliance prioritizes position accuracy regardless of external forces, while a high compliance allows large position deviation in response to external forces, making the system “soft” during interaction.

The desired compliance for a robotics system is not a static property; rather, it varies drastically depending on the task objectives and the system’s state. For instance, consider the flipping task in Fig. 1, the desired compliance:

- **Varies temporally.** For example, The system needs to be less compliant before contact to prioritize precise position

tracking while becoming compliant upon contact.

- **Varies spatially.** For example, during the pivoting stage, the system should be compliant only in the pushing directions (i.e., K_2 direction) while maintaining high stiffness in other directions to follow the arc motion (e.g., low compliance in K_1 direction).
- **Varies from task to task.** If we change to a different task, such as wiping a vase in Fig. 1 [Right], both temporal and spatial properties of the compliance will change in order to satisfy the unique 3D motion and force requirements.

While the desired compliance can be obtained from optimization given physical measurements of the manipulation problem [18, 16], it remains a challenge to obtain compliance parameters directly from human demonstrations. Compliance describes how force and motion variations in all directions are related, a typical demonstration trajectory does not contain all the information. Prior work often requires detailed known dynamics parameter [13] or multiple demonstration trajectories on the exact same task to statistically estimate one robot’s compliance [4, 25, 11, 12]. These approaches cannot handle new scene configurations or unexpected perturbations. As a result, compliant policies either rely on pre-selected compliance parameters for the target tasks [20] or assume uniform constant stiffness across all directions [26].

In this work, we introduce **Adaptive Compliance Policy (ACP)**, a sensorimotor policy that learns to dynamically adjust the system compliance both *spatially* and *temporally* for a *given manipulation task* from human demonstrations.

Specifically, the algorithm represents the compliance profile by an additional stiffness value and a virtual target




Fig. 2: **Method Comparisons.** [LEFT] shows the comparison between a) a typical visuomotor policy [6], b) a typical force-based compliant policy [26], and c) Adaptive Compliance Policy. [Right] Visualization of virtual target (orange squares) and reference poses (yellow circles) inferred by Adaptive Compliance Policy. The directional difference (orange arrows) between the virtual and reference poses encodes compliance direction.

pose in addition to the original reference pose (e.g., robot end-effector pose) predicted by the policy. The directional difference between the virtual and reference targets encodes the spatial distribution of stiffness. Finally, the predicted reference pose and stiffness can be executed by a standard low-level high-rate compliance controller to achieve robust and adaptive compliance behaviors.

Instead of estimating the exact human compliance, we derive a simple rule to obtain a useful compliance profile that guarantees the avoidance of large internal forces while encourages precise tracking, under mild assumptions about the tasks. This simple rule allows us to approximate varying stiffness for every demonstration episode with different object variations and scene configurations. Then, by learning from a collection of demonstrations, the policy can summarize the typical compliance profile for a specific task and quickly adjust compliance based on visual and force feedback.

We systematically evaluate the performance of our algorithm on two real world contact-rich manipulation tasks: object flipping and vase wiping. Our method achieves over 50% increase in performance compared to state-of-the-art visuomotor policy methods. In summary, the main contribution of the paper includes:

- Adaptive Compliance Policy formulation that is able to dynamically adjust compliance to maintain desired contact modes despite uncertainties and disturbances.
- A kinesthetic teaching system that allows demonstrations with varying compliance profiles.
- A method to compute spatial-, temporal-varying compliance labels from human demonstrations, making ACP training practical and scalable.

II. RELATED WORK

Contact-rich manipulation with active compliance. There is a long history of work on utilizing robot compliance for robust contact-rich manipulation [30, 39, 38, 15]. Given modeling information such as contact geometry and friction, the stiffness that provides the maximum robustness can be computed efficiently [18, 17, 16]. Although proved in many occasions, these methods typically require careful modeling

work to setup. Our work utilizes similar mechanical models but derives a model-free method that is easier to scale.

Learning compliance from Reinforcement Learning (RL). RL can effectively learn compliance controllers by exploring force-motion variations [36, 19, 3, 2, 34, 5]. However, these policies need to be retrained for any scene variations. Moreover, most of existing work uses fixed-parameter low-level compliance controllers like impedance [26] or admittance [23] controllers, which lack robustness against disturbances.

Learning compliance from human Human stiffness during a manipulation task can be estimated from sufficient repetitions of the same motion [8, 11, 12, 41, 42, 43]. The stiffness is considered to be either proportional to the force covariance [11, 12] or inversely proportional to the position covariance [4, 25]. These methods typically do not work for visuomotor policy learning since every demonstration is different. [13] proposed a stiffness estimator that works with one single demonstration. However, it requires perfect knowledge of human mass and damping, which is hard to obtain. Prior work also attempted to give the demonstrator the ability to specify stiffness [24, 20] without spatial variation.

Learning visuomotor policies with force feedback. To effectively encode the temporal relations in the force/torque data, Prior work has explored many different encoding methods such as causal convolution [26], TCN [2], LSTM [9], VAE [1], and self-attention [22, 23]. However, despite taking force as input, most of the existing visuomotor policies are still solely focused on predicting the position of the robot [29, 27] with uniform constant compliance [26] that fails to capture the spatial- and temporal- variations of the compliance requirements for delicate manipulation task.

III. METHOD

In this section, we introduce our pipeline, from collecting human demonstrations to designing the Adaptive Compliance Policy. We first introduce related physical concepts in §III-A, then describe a kinesthetic teaching system with compliance control §III-B. Next, §III-C introduces how we annotate compliance labels for the human demonstration data. Finally, §III-D explain the Adaptive Compliance Policy design.

A. Preliminaries: Modeling

Robot Compliance Modeling. Compliance expands the action space of a robot [21, 32]. Consider a N dimensional system described by position $x \in \mathbb{R}^N$ and force $f \in \mathbb{R}^N$. Compliance is the elastic behavior between force and motion, which is typically modeled by a spring-mass-damper system:

$$f = M\ddot{x} + K(x - x_{ref}) + K_D\dot{x}. \quad (1)$$

The three terms on the right-hand side represent inertia force, spring force, and damping force, respectively. x_{ref} is a reference position, at which the spring force is zero. The compliant behavior is described by the inertia matrix $M \in \mathbb{R}^{N \times N}$, stiffness matrix $K \in \mathbb{R}^{N \times N}$ and damping matrix $K_D \in \mathbb{R}^{N \times N}$, which can be user-specified if the compliance $\mathbf{1}$ is implemented by control. In other words, they can be added to the action space of a high level policy [3, 2, 8], while a low level compliance controller implements the “virtual” stiffness/damping/inertia. This is beneficial when the high level policy cannot be fast enough to exhibit compliance.

We use admittance control [31] for our compliance controller, which takes in force feedback and outputs position targets. This works for our high accuracy position-controlled robot (UR5e). Robots with force control interface may also use impedance control [14], or some hybrid force-motion control schemes [32, 37].

Manipulation Modeling Next we model how a manipulation system interacts with a robot, treating the robot as a black box. We make the following assumption:

Assumption I: *Contact force dominates. Other types of force, such as inertia force, friction and gravity, are negligible comparing with contact force.*

This is not to be confused with the compliance control of the robot itself, which is fast and dynamic. We ensure Assumption I by avoiding fast robot motion and using lightweight objects. Consider a robot with N degrees of freedom, making n contacts with the environment. Denote $\lambda \in \mathbb{R}^n$ as the vector of contact normal forces, the *Newton’s Second Law* can be written as following with Assumption I:

$$J^T \lambda = f, \quad (2)$$

where J is the Contact Jacobian matrix that maps contact force into the robot generalized force space. While used in derivation, our method does not need to compute it. Denote $v \in \mathbb{R}^N$ as the generalized velocity vector, the Jacobian J can also describe the velocity constraint imposed by contacts:

$$Jv \geq 0. \quad (3)$$

B. Demonstration Collection

We choose kinesthetic teaching instead of teleoperation to collect human demonstrations in order for the operator to easily demonstrate variable compliance behavior under direct haptic feedback (see Fig. 3). The setup for one arm includes one robot manipulator to provide accurate position feedback, one RGB camera to record visual information, and one force torque sensor mounted near the robot hand.

During demonstration, we specify low stiffness, low damping and low mass for the robot compliance controller so the




Fig. 3: **Data Collection with Haptic Feedback.** We designed a kinesthetic teaching system with low-stiffness compliance that allows the operator to demonstrate variable compliance behavior with direct haptic feedback.

operator can move the robot freely. Low damping and mass are achievable during demonstration because the human hand provides a natural external stabilization. During testing, we increase robot damping and virtual mass to maintain stability of the admittance controller. We also found it necessary to set the tool center point (TCP) near the handle, so the robot can rotate under external force in a way intuitive to the operator.

C. Estimating Compliance from Demonstrations

Human uses varying stiffness during manipulation. High stiffness provides position accuracy under force disturbances. Low stiffness is also necessary, since the high stiffness controls impose velocity constraints that might conflict with the contact constraints defined in Eq. 3, during which a huge internal force may be generated [17].

However, it is often an ill-conditioned problem to estimate compliance parameters from a single human demonstration due to the lack of variations [4]. Previous work on stiffness estimation [13] assumes perfect knowledge of damping and virtual mass and sufficient variations in force-motion signals. These requirements are not met in kinesthetic teaching, where the human hand changed the effective damping and mass of the robot hand, and the demonstration could have constant force or position for a period of time.

To simplify the problem, we also use pre-selected values for mass and damping. Then instead of estimating the true human stiffness, we propose to find a stiffness matrix with the following properties:

- It avoids huge internal forces in manipulation.
- It encourages accurate tracking of the desired motion.

C.1 Stiffness direction: We propose the following simple strategy to choose stiffness direction in the generalized space: *Use a low stiffness k_{low} in the direction of the force feedback, and a high stiffness k_{high} in all other directions.*

Next we explain why this simple rule works. From rigid body mechanics [33], we know that the rows of Contact Jacobian J represents the directions of contact normal forces, which forms a polyhedral convex cone in the generalized force space. We make two more assumptions:

Assumption II: *Nonzero contact force: all made contacts should have nonzero contact forces.*

Assumption III: No pinching contacts: the cone formed by rows of the Contact Jacobian J is contained in its dual cone.

Assumption II can be satisfied by making contacts clearly in demonstrations. Assumption III means the contacts on the robot are not too restrictive. Fig. 4 shows some examples:




Fig. 4: **Pinching Examples.** Grey shape represents a robot tool, blue shape represents a frictionless environment. First three examples are not pinching contact, the last one is.

We propose the following theorem under the assumptions:

Theorem 1. For a robot under external contact described by Eq. 2, there exists a solution v that satisfies the contact constraint 3 as long as it does not control its velocity in the direction of feedback force f in the generalized space.

Proof. Not controlling velocity in the force direction means the velocity has a free component:

$$v = v_0 + kf = v_0 + kJ^T \lambda, \quad (4)$$

where k is an arbitrary scaling factor, v_0 denotes the components of generalized velocity in other directions. Due to Assumption II, the contact force λ must have all positive components, $J^T \lambda$ represents a ray strictly inside the cone formed by rows of J . Assumption III says this cone is contained in its dual cone $\{x \in \mathbb{R}^N | Jx \geq 0\}$, so

$$JJ^T \lambda > 0. \quad (5)$$

Then $JV = JV_0 + kJJ^T \lambda > 0$ for large enough k .

Theorem 1 shows that one-dimensional low stiffness control suffices to avoid constraint violation, so we can use high stiffness in other directions to improve position tracking. Let $K_0 \in \mathbb{R}^N$ be a diagonal matrix with $[k_{low}, k_{high}, \dots, k_{high}]$ on its diagonal, and $S \in \mathbb{R}^N$ be a matrix whose columns form an orthonormal basis of \mathbb{R}^N with its first column as $f/|f|$. The stiffness matrix can be written as:

$$K = SK_0S^{-1} \quad (6)$$

We use k_{high} in all directions when $|f|$ is small.

C.2 Stiffness Magnitude The high stiffness k_{high} should support accurate position tracking in those directions, which can be set empirically. The low stiffness value should be zero, however, since the low stiffness direction is estimated from noisy force signal, we found it helpful to let the stiffness decrease continuously with the force magnitude:

$$k_{low} = \begin{cases} k_{max}, & |f| < f_{min} \\ k_{max} - (k_{max} - k_{min}) \frac{|f| - f_{min}}{f_{max} - f_{min}}, & f_{min} \leq |f| \leq f_{max} \\ k_{min}, & |f| > f_{max} \end{cases} \quad (7)$$

where $k_{max}, k_{min}, f_{max}$ and f_{min} are parameters determined by the hardware system.

D. Adaptive Compliance Policy

We formulate the policy as a diffusion process for both reference action and target stiffness. The policy runs in a receding-horizon manner [6], where an action trajectory is predicted using recent observations: 1) fisheye RGB images, 2) robot end-effector poses, and 3) force/torque data.

1) *Inputs and Encoding:* We implement two encoding strategies for the force/torque data: 1) temporal encoding via causal convolution [35], which helps capture causal relations from sequential data like force. We use force readings from the past 32 timesteps and pass them through a 5-layer causal convolution network, outputting a 768-dimensional vector; 2) FFT encoding, where we convert each dimension of the 6D force/torque readings into a 2D spectrogram. These spectrograms ($6 \times 30 \times 17$) are passed to a ResNet-18 model with a modified input channel of 6. A coordinate convolution layer [28] is added at the beginning to handle translational invariances. The images from the past two timesteps are resized to 224×224 with random cropping then encoded using a CLIP-pretrained ViT-B/16 model [10].

Both image and force encodings are passed to a transformer encoder layer, using self-attention to learn adaptive visual-force representations. This representation is concatenated with robot end-effector poses from the past 3 timesteps and fed to the downstream diffusion policy head as a condition following [6].

2) *Outputs and decoding:* Our policy output encodes the position target, the stiffness matrix, and the reference force in a 19-dimensional vector per robot arm:

- Reference pose: 9D pose vector following convention in [6], where the last six elements are the top two rows of a rotation matrix;
- Virtual target pose: 9D pose representing the actual target for the low level compliance controller to track;
- A scalar value representing the stiffness magnitude in the low stiffness direction.

The virtual target pose is computed such that the robot will exert the reference force if it reaches the reference pose while tracking the virtual target with the given stiffness. It essentially changes a force target into a position target. The benefit is to have a uniform target representation across different robots: an impedance controlled robot without FT sensor can also execute the virtual target.

During training, we first pass the whole episode of wrench data through a moving average filter with one second window size, then compute the stiffness from it using Eq. 6, and finally compute the virtual target following a 3D mechanical spring. The heavy wrench filtering has two benefits: 1) it makes the virtual target smooth; 2) it gives the action labels hindsight information about contacts about to be made, which is crucial for smooth contact engaging motions.

At inference time, the full stiffness matrix is reconstructed following Eq. 6 by replacing the force direction with the direction from the reference pose to the virtual target. Finally both the stiffness matrix and the virtual target are sent to the low level compliance controller for execution.

IV. EXPERIMENTS

We evaluate our method in two contact-rich manipulation tasks whose success depends on the maintenance of suitable contact modes. We use the UR5e robot with one GoPro RGB camera and one ATI mini-45 force torque sensor. The GoPro camera streams images at 60Hz, the robot receives Cartesian pose commands and send pose feedback at 500Hz, while the ATI sensor streams force-torque data at up to 7000Hz.

We evaluate the following four policies, all trained on the same dataset with the same number of epochs:

- ACP: Adaptive Compliance Policy, our approach;
- ACP w.o. FFT: same as ACP but with force encoded using temporal convolution [35, 26] instead of FFT.
- Stiff policy: Diffusion policy [6, 7] with additional force input. Outputs target positions.
- Compliant policy: Same as the [Stiff policy] except that the low level controller has a uniform stiffness $k = 500\text{N/m}$. Relying on low level robot compliance is common in visuomotor policies [26, 40, 44, 23].

Both ACP variations uses compliance in 3D translational space, i.e. $N = 3$, though our method formulations works for 6D compliance too if needed. For all polices, we use two frames of RGB image and three frames of end-effector poses. The policy is not sensitive to these numbers when they are small. [ACP w.o. FFT] uses 32 frames of wrench data sampled at 250Hz, while all other three policies uses one second of data at 7000Hz.

A. Task I: Item Flipping

The task is to flip up an item with a point finger by pivoting it against a corner of a fixture (i.e., a wall), as exemplified in Fig. 2. The task has two challenges: 1) The finger needs to consistently maintain contact force during the flipping motion regardless of the item’s shape, weight, and fixture locations; 2) Larger contact force will cause the fixture to slide on the floor, making it harder to maintain good contact. We trained each method on 230 episodes of demonstrations collected on 15 different items for 300 epochs.

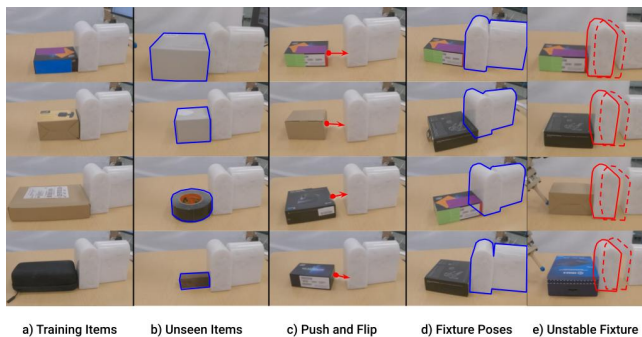


Fig. 5: **Flipping Scenarios.** We test the policy under a variety of settings that require the policy to adapt to different and unseen object geometries (a,b), configuration (c,d) and react to unexpected perturbations caused by fixture movements (e).

Test Scenarios. (Fig. 5) We ran 20 tests in each of the five scenarios below, making 100 tests per algorithm:

- Training Items: Items appeared in training data.

- Unseen Items: Items not seen in training.
- Push&Flip: Items start 5cm away from the fixture. This scenario requires the policy to first push the free-standing item towards the fixture then start the pivoting action.
- Varied Fixture Pose: Two different fixture poses.
- Unstable Fixture: Lighter fixture that causes unstable movements during the pivoting process that require the policy to quickly adapt.

Results. The success rate is shown in Tab. I. Success is defined as the item being rotated greater than 70 degrees. Typical failure cases include items falling off the finger, motion getting stuck, or triggering robot force violation.

TABLE I
FLIPPING-UP SUCCESS RATES (%)

	Train Items	Unseen Items	Push & Flip	Fixture Pose	Unstable Fixture	All
ACP	90	95	95	100	100	96
ACP w.o. FFT	90	100	100	95	90	95
Compliant Policy	80	15	15	5	0	23
Stiff Policy	20	0	5	35	10	14

Findings. The two baselines [Compliant Policy] and [Stiff Policy] have a few successes when they can exploit the passive compliance in the system. They effectively applies a force when the predicted trajectory is in collision with a deformable item. When the item is rigid, or when the position uncertainty is not in a convenient direction, the baseline policies break the contacts and fail the task. On the contrary, both variations of ACP tolerates a large range of position uncertainties while maintain the needed contacts.

In this task, [ACP w.o FFT] has similar performance (95/100) as ACP (96/100), indicating that the force spectrum encoding is similarly useful to convolution in this task. This makes sense as the force signal does not contain much high frequency component in the flipping motion.

B. Task II: Vase Wiping

The vase is randomly placed in front of the bimanual robot. The right upper side of the vase is marked by a random drawing using random colored dry markers. For this task, we equip each robot arm with a 3D-printed tool with two piece of kitchen sponges as wipers. The demonstrated motion uses the left arm to hold the vase while the right arm performs the wiping. We collected 200 demonstrations with various vase poses, marking shapes, and colors. Each demonstration includes one to five wipes to fully clean the markings.

Test Scenarios. We compared with the same set of algorithm alternatives as in the Item Flipping task. For each algorithm, we ran a total of 16 tests in the following four scenarios:

- Small Mark×5: easier cases that need only one wipe.
- Large Mark×5: require multiple wipes.
- Perturbation before contact (PbC)×4: move the vase right before the tool comes in contact with the vase to disturb the contact-engaging motion.
- Perturbation after contact (PaC)×2: move the vase after the tool is engaged to disturb the wiping motion.




Fig. 6: **Flipping Results.** [Top] Predicted stiffness in the world frame. The robot first decreases X-axis stiffness while approaching from the same direction, then gradually shifts the compliance to Z-axis, aligning with the contact normal. [Bottom] Predicted reference pose (yellow dot), virtual target poses (orange dot) and compliance direction (yellow line).

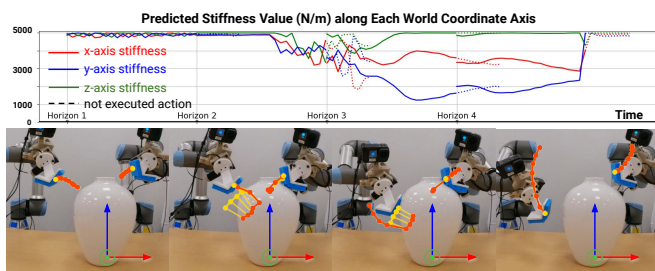


Fig. 7: **Wiping Results.** [Top] Predicted stiffness of the wiping arm in world frame. [Bottom] Predicted reference (yellow dot) and virtual poses (orange dot). The Wiping arm’s compliant direction roughly aligns with the contact normal. Despite some errors, the wiping succeeded as the virtual target pose still pulls arm towards the vase.

Results. The quantitative results are summarized in Tab. II. All policies demonstrated wiping behaviors. However, the effects of the wipes are different. We define success as the mark being cleaned within three wipes, where we consider a vase “clean” if the remaining marks are within $1\text{cm} \times 1\text{cm}$.

TABLE II
WIPING SUCCESS RATES (%)

	Small	Large	PbC	PaC	All
ACP	100	80	100	100	93.75
ACP w.o. FFT	100	60	75	100	81.25
Compliant Policy	60	20	25	100	43.75

Findings. Compared with the [Stiff Policy], [ACP] safely engages and maintains contacts during the wipes for its compliance. [Stiff Policy]’s contact force magnitude varies greatly with the accuracy of the position action. While the first few tests were successful, the robot broke its tool during the fourth test, as shown in Fig. 8, middle row.

Compared with the [Compliant Policy], [ACP] maintains accurate tracking of the desired motion for being stiff in the motion directions, as shown in Fig. 8, top row. The wiping motion of the [Compliant Policy] deviates from the position target because it is sensitive to the friction from the vase. As a result, the sliding motion is insufficient to make a high quality wipe. The comparison of wiping results is shown in Tab. II, which excludes the [Stiff Policy] since it broke the robot tool frequently.

Compared with [ACP w.o. FFT], our policy with the




Fig. 8: **Wiping Comparisons.** [Top] APC: maintains contact and follows desired trajectory. [Middle] Stiff Policy: Position noise causes excessive force that breaks the tool. [Bottom] Compliant Policy: Safe contact, but friction hinders wiping position accuracy and eventually loses contact.

FFT encoding performs better and finishes the task with fewer wipes. We observed that the FFT encoding, when used together with RGB encoding via cross-attention, makes better decision on the next best wiping location.

V. CONCLUSIONS

In this work, we show that Adaptive Compliance Policy is an effective visuomotor policy to do compliant manipulation. Extensive result in real-world shows that our approach is able to extract useful compliance from human demonstration in a variety of manipulation tasks, and thereby significantly improve the success rate of two contact-rich manipulation tasks.

ACKNOWLEDGMENT

This work was supported in part by the Toyota Research Institute, NSF Award #2143601, #2037101, and #2132519. We would like to thank Google and TRI for the UR5 robot hardware. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the sponsors.

REFERENCES

- [1] Marina Y Aoyama et al. “Few-shot learning of force-based motions from demonstration through pre-training of haptic representation”. In: *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2024, pp. 12839–12845.
- [2] Cristian C Beltran-Hernandez et al. “Variable compliance control for robotic peg-in-hole assembly: A deep-reinforcement-learning approach”. In: *Applied Sciences* 10.19 (2020), p. 6923.
- [3] Cristian Camilo Beltran-Hernandez et al. “Learning force control for contact-rich manipulation tasks with rigid position-controlled robots”. In: *IEEE Robotics and Automation Letters* 5.4 (2020), pp. 5709–5716.
- [4] Sylvain Calinon, Irene Sardellitti, and Darwin G Caldwell. “Learning-based control strategy for safe human-robot interaction exploiting task and robot redundancies”. In: *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE. 2010, pp. 249–254.
- [5] Chunyang Chang et al. “Impedance adaptation by reinforcement learning with contact dynamic movement primitives”. In: *2022 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM)*. IEEE. 2022, pp. 1185–1191.
- [6] Cheng Chi et al. “Diffusion Policy: Visuomotor Policy Learning via Action Diffusion”. In: *Proceedings of Robotics: Science and Systems (RSS)*. 2023.
- [7] Cheng Chi et al. “Diffusion Policy: Visuomotor Policy Learning via Action Diffusion”. In: *The International Journal of Robotics Research* (2024).
- [8] Zhen Deng et al. “Learning human compliant behavior from demonstration for force-based robot manipulation”. In: *2016 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE. 2016, pp. 319–324.
- [9] Junfeng Ding, Chen Wang, and Cewu Lu. “Transferable force-torque dynamics model for peg-in-hole task”. In: *arXiv preprint arXiv:1912.00260* (2019).
- [10] Alexey Dosovitskiy et al. “An image is worth 16x16 words: Transformers for image recognition at scale”. In: *arXiv preprint arXiv:2010.11929* (2020).
- [11] Jianghua Duan et al. “Learning compliant manipulation tasks from force demonstrations”. In: *2018 IEEE International Conference on Cyborg and Bionic Systems (CBS)*. IEEE. 2018, pp. 449–454.
- [12] Jianghua Duan et al. “Sequential learning unification controller from human demonstrations for robotic compliant manipulation”. In: *Neurocomputing* 366 (2019), pp. 35–45.
- [13] Giorgio Grioli and Antonio Bicchi. “A non-invasive, real-time method for measuring variable stiffness”. In: (2011).
- [14] Neville Hogan. “Impedance control: An approach to manipulation”. In: *1984 American control conference*. IEEE. 1984, pp. 304–313.
- [15] Yifan Hou, Zhenzhong Jia, and Matthew T Mason. “Fast planning for 3d any-pose-reorienting using pivoting”. In: *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2018, pp. 1631–1638.
- [16] Yifan Hou, Zhenzhong Jia, and Matthew T Mason. “Manipulation with shared grasping”. In: *arXiv preprint arXiv:2006.02996* (2020).
- [17] Yifan Hou and Matthew T Mason. “An efficient closed-form method for optimal hybrid force-velocity control”. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2021, pp. 11125–11131.
- [18] Yifan Hou and Matthew T Mason. “Robust execution of contact-rich motion plans by hybrid force-velocity control”. In: *2019 International Conference on Robotics and Automation (ICRA)*. IEEE. 2019, pp. 1933–1939.
- [19] Mrinal Kalakrishnan et al. “Learning force control policies for compliant manipulation”. In: *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE. 2011, pp. 4639–4644.
- [20] Tatsuya Kamijo, Cristian C Beltran-Hernandez, and Masashi Hamaya. “Learning Variable Compliance Control From a Few Demonstrations for Bimanual Robot with Haptic Feedback Teleoperation System”. In: *arXiv preprint arXiv:2406.14990* (2024).
- [21] Oussama Khatib. “A unified approach for motion and force control of robot manipulators: The operational space formulation”. In: *IEEE Journal on Robotics and Automation* 3.1 (1987), pp. 43–53.
- [22] Heecheol Kim et al. “Training robots without robots: deep imitation learning for master-to-robot policy transfer”. In: *IEEE Robotics and Automation Letters* 8.5 (2023), pp. 2906–2913.
- [23] Colin Kohler et al. “Symmetric models for visual force policy learning”. In: *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2024, pp. 3101–3107.
- [24] Klas Kronander and Aude Billard. “Learning compliant manipulation through kinesthetic and tactile human-robot interaction”. In: *IEEE transactions on haptics* 7.3 (2013), pp. 367–380.
- [25] Klas Kronander and Aude Billard. “Online learning of varying stiffness through physical human-robot interaction”. In: *2012 IEEE International Conference on Robotics and Automation*. IEEE. 2012, pp. 1842–1849.
- [26] Michelle A Lee et al. “Making sense of vision and touch: Self-supervised learning of multimodal representations for contact-rich tasks”. In: *2019 International conference on robotics and automation (ICRA)*. IEEE. 2019, pp. 8943–8950.
- [27] Hao Li et al. “See, hear, and feel: Smart sensory fusion for robotic manipulation”. In: *arXiv preprint arXiv:2212.03858* (2022).
- [28] Rosanne Liu et al. “An intriguing failing of convolutional neural networks and the coordconv solution”.

- In: *Advances in neural information processing systems* 31 (2018).
- [29] Zeyi Liu et al. “ManiWAV: Learning Robot Manipulation from In-the-Wild Audio-Visual Data”. In: *arXiv preprint arXiv:2406.19464* (2024).
- [30] Tomas Lozano-Perez, Matthew T Mason, and Russell H Taylor. “Automatic synthesis of fine-motion strategies for robots”. In: *The International Journal of Robotics Research* 3.1 (1984), pp. 3–24.
- [31] J Maples and Joseph Becker. “Experiments in force control of robotic manipulators”. In: *Proceedings. 1986 IEEE International Conference on Robotics and Automation*. Vol. 3. IEEE. 1986, pp. 695–702.
- [32] Matthew T Mason. “Compliance and force control for computer controlled manipulators”. In: *IEEE Transactions on Systems, Man, and Cybernetics* 11.6 (1981), pp. 418–432.
- [33] Richard M Murray, Zexiang Li, and S Shankar Sastry. *A mathematical introduction to robotic manipulation*. CRC press, 2017.
- [34] Michael Noseworthy et al. “FORGE: Force-Guided Exploration for Robust Contact-Rich Manipulation under Uncertainty”. In: *arXiv preprint arXiv:2408.04587* (2024).
- [35] Aaron van den Oord. “WaveNet: A Generative Model for Raw Audio”. In: *arXiv preprint arXiv:1609.03499* (2016).
- [36] Tifanny Portela et al. “Learning force control for legged manipulation”. In: *IEEE International Conference on Robotics and Automation (ICRA)* (2024).
- [37] Marc H Raibert and John J Craig. “Hybrid position/force control of manipulators”. In: (1981).
- [38] Naoyuki Sawasaki and Hirochika INOUE. “Tumbling objects using a multi-fingered robot”. In: *Journal of the Robotics Society of Japan* 9.5 (1991), pp. 560–571.
- [39] Masaru Uchiyama and Pierre Dauchez. “A symmetric hybrid position/force control scheme for the coordination of two robots”. In: *Proceedings. 1988 IEEE international conference on robotics and automation*. IEEE. 1988, pp. 350–356.
- [40] Chen Wang et al. “Mimicplay: Long-horizon imitation learning by watching human play”. In: *arXiv preprint arXiv:2302.12422* (2023).
- [41] Ning Wang, Chuize Chen, and Alessandro Di Nuovo. “A framework of hybrid force/motion skills learning for robots”. In: *IEEE Transactions on Cognitive and Developmental Systems* 13.1 (2020), pp. 162–170.
- [42] Koki Yamane et al. “Soft and rigid object grasping with cross-structure hand using bilateral control-based imitation learning”. In: *IEEE Robotics and Automation Letters* (2023).
- [43] Chao Zeng et al. “Generalization of robot force-relevant skills through adapting compliant profiles”. In: *IEEE Robotics and Automation Letters* 7.2 (2021), pp. 1055–1062.
- [44] Yifeng Zhu et al. “Viola: Imitation learning for vision-based manipulation with object proposal priors”. In: *Conference on Robot Learning*. PMLR. 2023, pp. 1199–1210.