

# MOMA-Force: Visual-Force Imitation for Real-World Mobile Manipulation

Taozheng Yang\*, Ya Jing\*, Hongtao Wu\*, Jiafeng Xu\*, Kuankuan Sima†,  
Guangzeng Chen, Qie Sima†, Tao Kong‡  
ByteDance Research

<https://visual-force-imitation.github.io>

**Abstract**—In this paper, we present a novel method for mobile manipulators to perform multiple contact-rich manipulation tasks. While learning-based methods have the potential to generate actions in an end-to-end manner, they often suffer from insufficient action accuracy and robustness against noise. On the other hand, classical control-based methods can enhance system robustness, but at the cost of extensive parameter tuning. To address these challenges, we present MOMA-Force, a visual-force imitation method that seamlessly combines representation learning for perception, imitation learning for complex motion generation, and admittance whole-body control for system robustness and controllability. MOMA-Force enables a mobile manipulator to learn multiple complex contact-rich tasks with high success rates and small contact forces. In a real household setting, our method outperforms baseline methods in terms of task success rates. Moreover, our method achieves smaller contact forces and smaller force variances compared to baseline methods without force imitation. Overall, we offer a promising approach for efficient and robust mobile manipulation in the real world. Videos and more details can be found on <https://visual-force-imitation.github.io>.

## I. INTRODUCTION

Mobile manipulation combines two fundamental robot capabilities: mobility and manipulation. These two capabilities substantially extend robot applications in the real world compared to static manipulation [1]. For example, mobile manipulation enables robots to complete tasks involving manipulations with large workspaces (*e.g.*, opening closet doors). However, mobile manipulation poses significant challenges when it comes to real-world tasks. The challenges are mainly twofold. First, uncertainties caused by localization and control can lead to potential safety issues, especially in contact-rich tasks. Second, the high-dimensional configuration space makes motion generation and control complex.

Early work assumed prior knowledge of the scene and explored path planning [2] and task and motion planning (TAMP) [3] for mobile manipulators. Whole-body control algorithms have been studied to solve high-dimensional optimization problems [4], [5], [6]. However, the environment is also simplified to ease the optimization complexity. Data-driven methods such as imitation learning [7], [8] and reinforcement learning [9] are becoming popular as they enable robots to learn directly from visual observation. Recent works extend the large-scale pre-training paradigm from

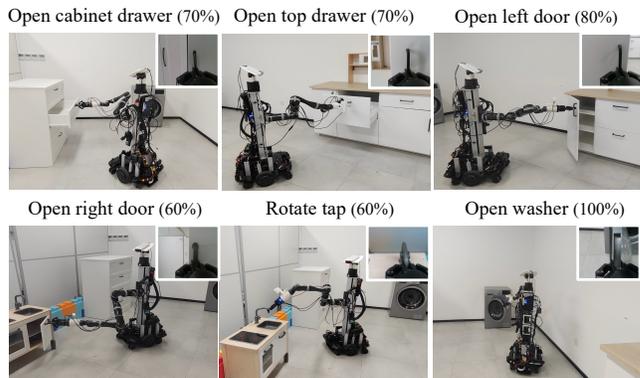


Fig. 1: **Overview.** We perform six real-world contact-rich mobile manipulation tasks. The success rates of our method are shown at the top of each image.

computer vision and natural language processing to static manipulation [7], [8], [10] and navigation [11]. They show promising results in various real robot tasks. Compared to exclusively learning from expensive robot data, this paradigm makes use of existing large-scale dataset [12], [13] to pre-train a visual representation model which is able to learn useful visual properties of the world for efficient policy learning. However, these methods are rarely designed for contact-rich mobile manipulation tasks. And visual imitation learning generally suffers from the “last-centimeter” error [14]. This problem becomes non-negligible in many real-world contact-rich tasks as it may lead to severe mechanical damage to the environment and robots.

In this paper, we argue that wrench  $\mathcal{F} = [m, f] \in \mathbb{R}^6$  ( $m$  is the torque and  $f$  is the force) is critical in terms of task completion and safety. To this end, we propose a novel visual-force imitation method for solving real-world contact-rich mobile manipulation tasks. In contrast to previous work which only imitates kinematic actions, we extend a state-of-the-art visual imitation method [15] to support both kinematic action and force imitation. This enables us to tackle the uncertainties caused by localization and control in mobile manipulation, and thus ensures safer robot execution. Moreover, we leverage whole-body control (WBC) [6], [16] to solve the high-dimensional control problem of mobile manipulators. This allows the robot to track the trajectory generated by the predicted action while regulating the wrench to imitate the expert wrench. We perform multi-task learning on

\*Equal contribution. †Work done during an internship.

‡Corresponding author: Tao Kong (kongtao@bytedance.com).

six household contact-rich tasks (Fig. 1). Results show that our method outperforms several baseline methods without force imitation on task success rate. In particular, our method achieves an average success rate of 73.3% on the six tasks, while the best baseline method achieves 45.0%. Additionally, with force imitation, the average absolute contact wrench and the average wrench variance decrease, indicating a safer and more stable contact between the robot and the object.

The key contributions are summarized as follows:

- We address the challenging problems caused by uncertainty and high-dimensional kinematics in real-world mobile manipulation tasks in an imitation learning context, and propose an effective visual-force imitation method to solve real-world mobile manipulation tasks.
- We present systematic real robot experiments on six contact-rich mobile manipulation tasks and showcase the performance of the proposed system on these challenging tasks.

## II. RELATED WORK

### A. Mobile Manipulation

Previous work explored using task and motion planning (TAMP) to plan and control mobile manipulators [3]. Another line of methods leverages a move-and-act pipeline [17]. However, it is often challenging to find viable kinematic solutions due to constraints from robots and/or environments. Alternatively, whole-body control [18], [19] are able to deal with complex constraints arise from mechanical limitations, manipulability or dynamic obstacle avoidance requirements. Learning-based methods are adopted in mobile manipulation to enable robots with visual feedback. In particular, imitation learning methods try to mimic actions from human demonstrations [20], [21], [22]. Additionally, force information must be carefully considered to improve task completion and safety [17].

### B. Robot Learning from Demonstrations

Learning from demonstrations is a sample-efficient and practical approach for training robots to perform complex manipulation skills. A popular line of work leverages behavior cloning (BC) to efficiently learn from human demonstrations [23], [8], [24], [14], [25], [26]. A recent work trains a transformer to learn multiple tasks with a dataset containing  $\sim 130k$  episodes [22]. Additionally, there has been growing interest in learning general representations for robot manipulation based on pre-trained models [7], [8], [27], [28]. Inspired by these methods, we use a pre-trained model to extract visual representations in this paper. The most related work is a non-parametric method that retrieves expert actions by comparing the representation of the observation image captured in the rollout with those of the observations in the expert trajectories [15]. Our method differs in that we imitate both kinematic actions and wrenches while this work only imitates kinematic actions.

### C. Force Compliant Manipulation

In classical control community, several early works on force compliance control have been developed. Specifically, [29] and [30] provide implementations of impedance and/or admittance control for robot manipulation based on wrench feedback. Manipulation tasks with articulated constraints are very common in household environments, *e.g.*, opening doors and drawers. Force compliance control is generally used to provide the required compliance for tackling these tasks [31]. Force compliance is also essential for dexterous in-hand manipulation and teleoperation [32]. Inspired by the previous work, we believe force imitation is able to enhance the performance of mobile manipulators through improved action accuracy and safer environment interaction.

## III. METHOD

We consider contact-rich mobile manipulation tasks where a robot interacts with the environment or an object within the environment. Besides manipulation, such tasks also necessitate the integration of mobility due to large workspace. Our method mainly consists of two parts: 1) an action-wrench prediction module and 2) an admittance whole-body control module (Fig. 2). At each time step  $t$ , the action-wrench prediction module takes as input the RGB image  $\mathbf{o}_t$  captured from an arm-mounted camera and generates a kinematic action  $\mathbf{a}_t \in SE(3)$ , a gripper action  $\mathbf{g}_t \in \{-1, 0, 1\}$ , a target wrench for the next time step  $\hat{\mathcal{F}}_{t+1} \in \mathbb{R}^6$  and a terminate flag  $\mathcal{T}_t \in \{0, 1\}$  which indicates rollout termination. If  $\mathbf{g}_t = 0$  and  $\mathcal{T}_t = 0$ , the control module takes as input  $\mathbf{a}_t$  and  $\hat{\mathcal{F}}_{t+1}$  and controls the robot to track a target pose generated by  $\mathbf{a}_t$  while regulating the contact wrench to imitate  $\hat{\mathcal{F}}_{t+1}$ . If  $\mathbf{g}_t = 1$  (or  $\mathbf{g}_t = -1$ ), the robot opens (or closes) the gripper. If  $\mathcal{T}_t = 1$ , the rollout terminates the rollout. We assume access to an expert dataset  $D$  containing expert trajectories of *multiple* tasks. Each trajectory consists of multiple frames  $d$  captured at different time steps. Each frame is a tuple  $d = \{\mathbf{o}^e, \mathbf{a}^e, \mathbf{g}^e, \mathcal{F}^e, \mathcal{T}^e\}$  where  $\mathbf{o}^e$ ,  $\mathbf{a}^e$ ,  $\mathbf{g}^e$ ,  $\mathcal{F}^e$ , and  $\mathcal{T}^e$  are the observation image, kinematic action, gripper action, wrench, and terminate flag, respectively. We use superscript  $e$  on variables to indicate that they are from the expert dataset  $D$ . Alg. 1 shows the algorithm of our method.

### A. Action-Wrench Prediction

Our prediction module is built on a state-of-the-art imitation learning method [15]. It consists of two phases: offline observation encoding (Line 1-5 of Alg. 1) and online rollout (Line 6-17). In the offline phase, we leverage a pre-trained vision encoder to project the observation image of every frame in  $D$  to a deep embedding  $\mathbf{z}^e \in \mathcal{Z}^e$  (Line 3-5). The embedding holds a compact representation for the visual observation. We use a recently proposed self-supervised visual representation model [33] as the pre-trained visual encoder and compare the performance of different representation models in Sec. IV.

During the online rollout phase, at each time step  $t$ , the captured observation image  $\mathbf{o}_t$  is also encoded as  $\mathbf{z}_t$  using the same visual encoder (Line 10).  $\mathbf{z}_t$  is used to predict the action

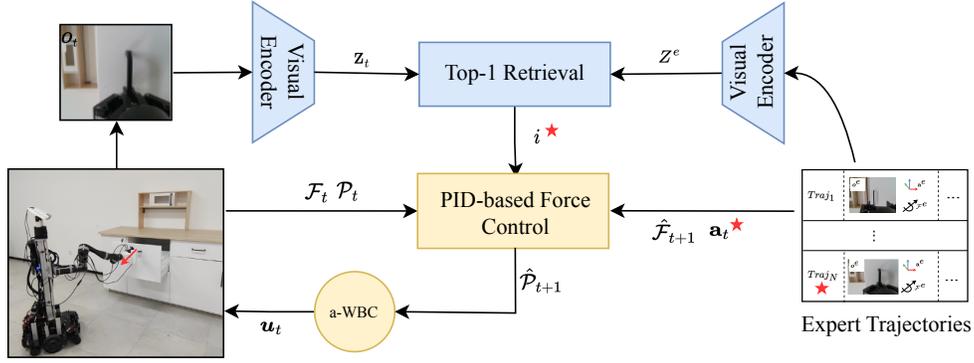


Fig. 2: **MOMA-Force** The observation images of the expert data are converted to representation vectors by a visual encoder. In the rollout, the observation image is converted to a representation vector with the same visual encoder. The action and target wrench are predicted by retrieving the action and wrench of the expert data with top-1 similarity of the representation vectors. We use admittance whole-body control (a-WBC) to control the robot.

---

**Algorithm 1:** VisualForceImitation

---

**Input:**

$D$  - Expert dataset  
 $i_{\max}$  - Maximum rollout iteration

```

1  $Z^e \leftarrow \emptyset$ 
2 for  $d \leftarrow D$  do
3    $(\mathbf{o}^e, \mathbf{a}^e, \mathcal{F}^e) \leftarrow d$ 
4    $\mathbf{z}^e \leftarrow \text{VisualEncoder}(\mathbf{o}^e)$ 
5    $Z^e \leftarrow Z^e \cup \{\mathbf{z}^e\}$ 
6  $i \leftarrow 0$ 
7 while  $i < i_{\max}$  do
8    $i \leftarrow i + 1$ 
9    $(\mathbf{o}_t, \mathcal{F}_t, \mathcal{P}_t) \leftarrow \text{getCurrentState}()$ 
10   $\mathbf{z}_t \leftarrow \text{VisualEncoder}(\mathbf{o}_t)$ 
11   $(\mathbf{a}_t, \hat{\mathcal{F}}_{t+1}, \mathcal{T}_t) \leftarrow \text{Retrieve}(\mathbf{z}_t, Z^e)$ 
12  if  $\mathcal{T}_t = 1$  then
13    return
14   $\hat{\mathcal{P}}_{t+1} \leftarrow \mathcal{P}_t \circ \mathbf{a}_t$ 
15   $\Delta \mathcal{P}_{t+1} \leftarrow \text{ForceControl}(\mathcal{P}_t, \hat{\mathcal{P}}_{t+1}, \mathcal{F}_t, \hat{\mathcal{F}}_{t+1})$ 
16   $\hat{\mathcal{P}}_{t+1} \leftarrow \hat{\mathcal{P}}_{t+1} \circ \Delta \mathcal{P}_{t+1}$ 
17   $\mathbf{u}_t \leftarrow \text{a-WBC}(\hat{\mathcal{P}}_{t+1})$ 
18 return

```

---

and target wrench by comparing with all the embeddings in  $Z^e$  (Line 11). Specifically, we use the cosine distance to compute the similarity between  $\mathbf{z}_t$  and each embedding in  $Z^e$ :

$$\text{sim}(\mathbf{z}_t, \mathbf{z}_i^e) = \frac{\mathbf{z}_t^T \mathbf{z}_i^e}{\|\mathbf{z}_i^e\| \|\mathbf{z}_t\|} \quad (1)$$

where  $\mathbf{z}_i^e$  is the embedding of the  $i$ -th frame in  $D$ . The top-1 frame index in  $D$  is retrieved via maximum similarity:

$$i^* = \arg \max_i \text{sim}(\mathbf{z}_t, \mathbf{z}_i^e) \quad (2)$$

The kinematic action, gripper action, and terminate flag of the top-1 frame are used as the predicted kinematic action, gripper action and terminate flag:  $\mathbf{a}_t = \mathbf{a}_{i^*}^e$ ,  $\mathbf{g}_t = \mathbf{g}_{i^*}^e$ ,  $\mathcal{T}_t =$

$\mathcal{T}_{i^*}^e$ . The wrench of the *next* frame of the top-1 frame is used as the target wrench for the next time step  $\hat{\mathcal{F}}_{t+1} = \mathcal{F}_{i^*+1}^e$ .

### B. Admittance Whole-Body Control

The current pose of the robot end-effector is denoted as  $\mathcal{P}_t = [R_t, \mathbf{p}_t] \in SE(3)$  where  $R_t \in SO(3)$  and  $\mathbf{p}_t \in \mathbb{R}^3$  denote the rotation and translation, respectively. The target end-effector pose for the next time step  $\hat{\mathcal{P}}_{t+1} \in SE(3)$  is computed by  $\hat{\mathcal{P}}_{t+1} = \mathcal{P}_t \circ \mathbf{a}_t$  where  $\circ$  is the group action of  $SE(3)$  (Line 14). However, this pose may not be accurate due to the uncertainties caused by localization and the insufficient accuracy caused by action prediction (Sec. I). And small pose errors in contact-rich tasks may lead to large contact wrenches and even mechanical damage. Therefore, we leverage an admittance control scheme and control the robot to track a pose that augments the original  $\hat{\mathcal{P}}_{t+1}$  with an admittance term  $\Delta \mathcal{P}_{t+1} = (\Delta R_{t+1}, \Delta \mathbf{p}_{t+1}) \in SE(3)$  (Line 15-16). This term compensates for the difference between the wrench at the current time step measured by the force-torque sensor  $\mathcal{F}_t$  and the target wrench  $\hat{\mathcal{F}}_{t+1}$ . It is computed based on the wrench tracking error:

$$\begin{bmatrix} \Delta \mathbf{w}_{t+1} \\ \Delta \mathbf{p}_{t+1} \end{bmatrix} = \mathbf{K}_p \Delta \mathcal{F} + \mathbf{K}_i \int \Delta \mathcal{F} dt + \mathbf{K}_d \frac{\partial \Delta \mathcal{F}}{\partial t} \quad (3)$$

where  $\Delta R_{t+1} = \exp(\widehat{\Delta \mathbf{w}_{t+1}})$ ;  $\Delta \mathcal{F} = \hat{\mathcal{F}}_{t+1} - \mathcal{F}_t$ ;  $\mathbf{K}_p$ ,  $\mathbf{K}_i$  and  $\mathbf{K}_d$  are the proportional, integral, and differential terms, respectively. To obtain the final  $\Delta \mathcal{P}_{t+1}$  applied to the robot, we discard the component parallel to  $\mathbf{a}_t$  and only retain the perpendicular component. This is equivalent to force-position hybrid control in which the position weight and the force weight are set to: 1 and 0 in the desired motion directions; 0 and 1 in the remaining directions.

With the admittance compensation, the augmented target pose  $\hat{\mathcal{P}}_{t+1}$  is sent to a whole-body controller to generate control commands for the mobile manipulator. The whole-

body controller is expressed as:

$$\begin{aligned} \min_{\mathbf{u}} \quad & f = \frac{1}{2} \mathbf{u}^T \mathbf{Q} \mathbf{u} + \mathbf{c}^T \mathbf{u} \\ \text{s.t.} \quad & \mathbf{J} \mathbf{u} = \mathbf{v}_e \\ & \mathbf{A} \mathbf{u} \leq \mathbf{B} \\ & \mathbf{u} \in [\mathbf{u}_{\min}, \mathbf{u}_{\max}] \end{aligned} \quad (4)$$

$\mathbf{u}$  is the decision variable vector which includes velocity control for the base and arm.  $\mathbf{Q}$  incorporates joint velocity costs.  $\mathbf{c} = (\mathbf{0}_b, \mathbf{J}_m^a)$  tries to maximize the manipulability of the arm, where  $\mathbf{0}_b$  is a zero vector and  $\mathbf{J}_m^a$  is the manipulability Jacobian of the arm [18].  $\mathbf{J}$  is the generalized Jacobian of the base and arm.  $\mathbf{v}_e \in \mathbb{R}^6$  is the desired spatial velocity calculated by  $\hat{\mathcal{P}}_{t+1}$  and the current end-effector pose.  $\mathbf{A}$  and  $\mathbf{B}$  implement the joint position constraints. The optimization problem is a QP problem with variables  $\mathbf{u}$ . We use qpOASES [34] to solve the optimization problem in this paper.

#### IV. EXPERIMENTS

We perform extensive real robot experiments on six contact-rich tasks (Fig. 1). Through the experiments, we aim to answer three questions. 1) Is force imitation effective on improving task success rate? 2) Does force imitation effectively reduce undesired contact wrenches and unstable oscillations between the robot and the environment? 3) How do different visual encoders perform in terms of action prediction on real robot data?

##### A. Experiment Setup

1) *System*: The robot platform is shown in Fig. 3(a). It consists of a differential-drive mobile base and a 7-DOF Kinova Gen2 robot arm. Wheel encoder, IMU, and lidar are used for the localization of the mobile base. A Robotiq 2F-85 parallel-jaw gripper and a Robotiq FT 300-S force-torque sensor are mounted on the end-effector of the arm.

2) *Tasks*: We perform experiments on six contact-rich tasks in which the robot manipulates different articulated objects (Fig. 1). The articulation includes prismatic joints (*e.g.*, drawers) and revolute joints (*e.g.*, doors and the tap). *Almost all tasks are impossible to finish by static manipulations*. All tasks require the robot to grasp the object and stay in contact with it during manipulation.

3) *Expert Data Collection*: We collected 30 expert demonstrations for each task. The human expert carried the end effector to collect expert trajectories (Fig. 3(b)). The observation RGB image, the pose of the end effector, the gripper action, and the wrench were recorded for each time step. The pose of the end effector is tracked via an HTC VIVE tracking system. The kinematic action for each frame is computed as the relative transformation between the pose of the next time step and the pose of the current time step. The frames in which experts perform gripper actions have no kinematic actions. That is, if this frame has been retrieved in the rollout, the robot executes the gripper action and the pose of the robot remains unchanged. The terminate flag is 1 for the last frame and 0 otherwise.

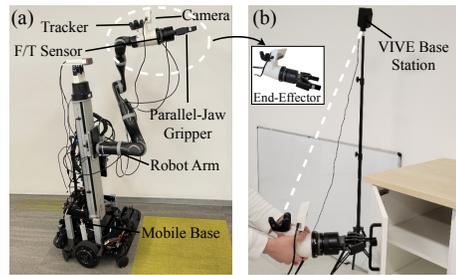


Fig. 3: **Experiment.** (a) The mobile manipulator used in the experiment. (b) Expert trajectory collection.

4) *Pipeline*: All tasks can be divided into three phases: approaching, grasping, and in-contact manipulation. The robot starts at a pre-defined initial pose which is contact-free. We randomize the end-effector initial pose before a rollout starts. The robot first performs predicted actions to approach the object. When the robot is commanded to close the gripper, the robot enters the grasping phase. Finally, the robot enters the in-contact manipulation phase. A rollout terminates if either one of the three conditions is met: 1) the termination flag  $\mathcal{T}_t = 1$ ; 2) the rollout time step exceeds the maximum time step; 3) the force is larger than 40N or the average force in the past 1s is larger than 30N. The last condition ensures safety and avoids mechanical damage to the object and the robot during rollouts. We consider a trial successful if at least 80% of the task has been completed. Each task is carried out 10 times for each method.

##### B. Experiment Results

We denote our visual-force imitation method as MOMA-Force. We compare with three baseline methods. The first baseline method is behavior cloning, denoted as BC. We use the same visual encoder [33] to encode observation images to representation vectors. We follow [7] and train a multi-layer perceptron that takes as input the representation vector and outputs the kinematic action, gripper action and terminate flag. In particular, we train BC in a single-task learning setting. We performed multi-task learning but found the results not good. We also compare with a variant of our method without force imitation, denoted as MOMA-Force w/o FC, to validate the effectiveness of force imitation. In addition, we test a variant of our method without force imitation and rotation, *i.e.*, the action only contains translation  $\mathbf{a}_t \in \mathbb{R}^3$ . We denote this method as MOMA-Force w/o FC & Rot. Qualitative and quantitative results are shown in Fig. 4 and Tab. I, respectively.

1) *Is force imitation effective in improving task success rate?* Our method MOMA-Force achieves the best average success rate among all the comparing baseline methods. Compared to BC, our method obtains a performance gain of 53.33% even though it performs multi-task learning while BC performs single-task learning. BC struggles with the compounding-error problem when approaching the object in a few tasks. The grasping success rates in these tasks are low. In other tasks where grasping is mostly successful, it cannot

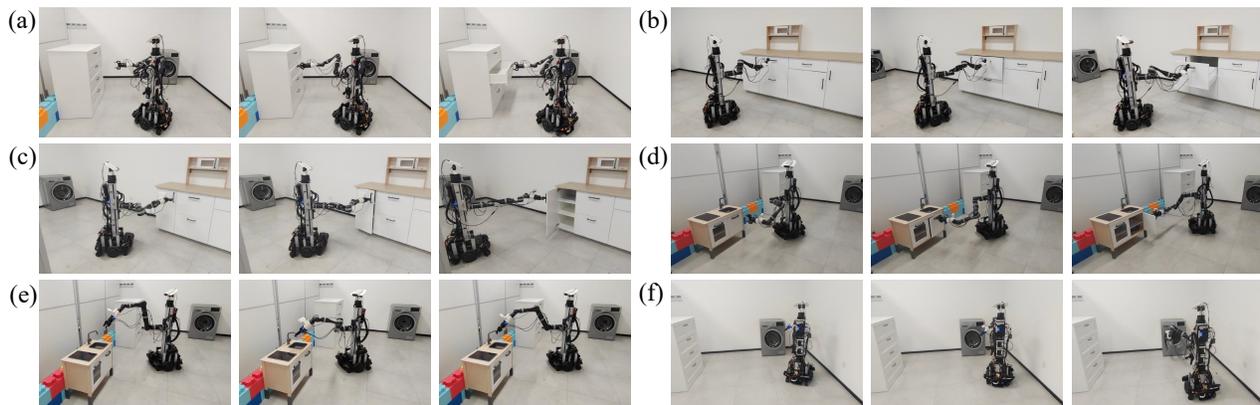


Fig. 4: **Real Robot Experiments.** (a) Open cabinet drawer. (b) Open drawer. (c) Open left door. (d) Open right door. (e) Rotate tap. (f) Open washing machine.

TABLE I: Success Rates of Real Robot Experiments.

| Method                  | Open cabinet drawer | Open drawer | Open left door | Open right door | Rotate tap | Open washing machine | Average      |
|-------------------------|---------------------|-------------|----------------|-----------------|------------|----------------------|--------------|
| BC (single-task) [7]    | 0%                  | 20%         | 40%            | 0%              | 50%        | 10%                  | 20.0%        |
| MOMA-Force w/o FC & Rot | 20%                 | 50%         | 10%            | 0%              | 20%        | 0%                   | 16.7%        |
| MOMA-Force w/o FC       | 70%                 | 70%         | 50%            | 60%             | 10%        | 10%                  | 45.0%        |
| MOMA-Force (Ours)       | 70%                 | 70%         | 80%            | 60%             | 60%        | 100%                 | <b>73.3%</b> |

maintain a small force, resulting in early termination without task completion. Without force imitation, the success rate of MOMA-Force w/o FC decreases by 28.3% on average. This method mainly struggles with tasks involving translation and rotation, *e.g.*, opening the washing machine and rotating the tap. In these tasks, similar to BC which also does not have force imitation, it suffers from early termination due to large forces.

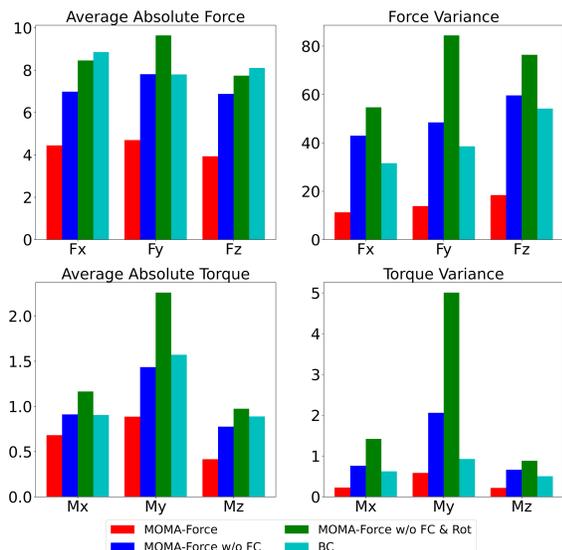


Fig. 5: **Results of Contact Wrenches.** The force imitation results in a reduction on contact wrench and variance.

The **failure cases** of MOMA-Force mainly attribute to inaccurate action prediction. A typical failure mode is unsuccessful grasping. In the expert trajectories, the observation images between frames with small pose changes are very

TABLE II: Mean Squared Error of predicted actions and wrenches.

| Visual Encoder | Translation | Rotation    | Wrench      |
|----------------|-------------|-------------|-------------|
| MVP [8]        | 0.19        | 0.17        | 0.21        |
| CLIP [10]      | 0.16        | 0.17        | 0.15        |
| iBOT [33]      | <b>0.10</b> | <b>0.11</b> | <b>0.10</b> |

similar. Therefore, in some cases, the visual encoder is not able to distinguish and thus the robot cannot retrieve the best expert action.

2) *Is force imitation effective on decreasing the contact wrench between the robot and the environment?* We compare the contact wrench of MOMA-Force and the baseline methods. For all methods, we first compute the average absolute contact wrenches and average wrench variances of all the trials in which the grasping is successfully in the rollout for each of the six tasks. We then average across tasks and report the results in Fig. 5. Lower variances indicate more stable contact. With force imitation, the average absolute contact force and torque of MOMA-Force in x, y, and z-axes are all smaller compared to those of the baseline methods without force imitation. In addition, MOMA-Force has a smaller force variance, indicating less oscillation and more stable contact during the rollout.

3) *How do different pre-trained visual encoders perform on real-robot data?* To compare the effectiveness of different visual pre-trained models, we show the Mean Squared Error (MSE) of using various state-of-the-art pre-trained models as the visual encoder on the test set with 5-fold cross-validation. Tab. II shows the result. MVP [8] is pre-trained via a masked autoencoder from the Internet and egocentric videos. CLIP [10] aims to align the image representation with the paired text through contrastive learning. iBOT [33] gives a good trade-off between masked autoencoder and

contrastive learning with an online tokenizer. Tab. II shows that iBOT achieves the best performance, demonstrating its advantages on our tasks. We note that a concurrent work [35] gives a more sufficient study on visual representations for robot tasks.

## V. CONCLUSION

In this work, we present a novel visual-force imitation learning method for real-world, contact-rich mobile manipulation tasks. We extend a state-of-the-art visual imitation learning method to support both kinematic action and target wrench prediction. We leverage admittance whole-body control to enable robots to track the trajectory generated by the action while regulating the wrench to follow the expert wrench. We implement our method on a high DoF mobile manipulator and perform multi-task learning on six real-world contact-rich tasks. Our method achieves an average success rate of 73.3%, outperforming several baseline methods without force imitation. In addition, the average absolute contact wrench and wrench variance of our method are smaller compared to all the baseline methods.

## REFERENCES

- [1] J. Wong, A. Tung, A. Kurenkov, A. Mandlekar, L. Fei-Fei, S. Savarese, and R. Martín-Martín, “Error-aware imitation learning from teleoperation data for mobile manipulation,” in *Conference on Robot Learning*. PMLR, 2022, pp. 1367–1378.
- [2] W. F. Carriker, P. K. Khosla, and B. H. Krogh, “Path planning for mobile manipulators for multiple task execution,” *IEEE Transactions on Robotics and Automation*, vol. 7, no. 3, pp. 403–408, 1991.
- [3] J. Wolfe, B. Marthi, and S. Russell, “Combined task and motion planning for mobile manipulation,” in *Proceedings of the International Conference on Automated Planning and Scheduling*, vol. 20, 2010, pp. 254–257.
- [4] M. Mittal, D. Hoeller, F. Farshidian, M. Hutter, and A. Garg, “Articulated object interaction in unknown scenes with whole-body mobile manipulation,” *arXiv preprint arXiv:2103.10534*, 2021.
- [5] M. V. Minniti, F. Farshidian, R. Grandia, and M. Hutter, “Whole-body mpc for a dynamically stable mobile manipulator,” *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3687–3694, 2019.
- [6] J. Pankert and M. Hutter, “Perceptive model predictive control for continuous mobile manipulation,” *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6177–6184, 2020.
- [7] S. Nair, A. Rajeswaran, V. Kumar, C. Finn, and A. Gupta, “R3m: A universal visual representation for robot manipulation,” in *6th Annual Conference on Robot Learning*, 2022.
- [8] I. Radosavovic, T. Xiao, S. James, P. Abbeel, J. Malik, and T. Darrell, “Real-world robot learning with masked visual pre-training,” in *6th Annual Conference on Robot Learning*, 2022.
- [9] Z. Fu, X. Cheng, and D. Pathak, “Deep whole-body control: Learning a unified policy for manipulation and locomotion,” in *6th Annual Conference on Robot Learning*.
- [10] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, “Learning transferable visual models from natural language supervision,” in *International conference on machine learning*, 2021.
- [11] R. Ramrakhya, E. Undersander, D. Batra, and A. Das, “Habitat-web: Learning embodied object-search strategies from human demonstrations at scale,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 5173–5183.
- [12] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *2009 IEEE conference on computer vision and pattern recognition*. IEEE, 2009, pp. 248–255.
- [13] K. Grauman, A. Westbury, E. Byrne, Z. Chavis, A. Furnari, R. Girdhar, J. Hamburger, H. Jiang, M. Liu, X. Liu *et al.*, “Ego4d: Around the world in 3,000 hours of egocentric video,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 18 995–19 012.
- [14] E. Jang, A. Irpan, M. Khansari, D. Kappler, F. Ebert, C. Lynch, S. Levine, and C. Finn, “Bc-z: Zero-shot task generalization with robotic imitation learning,” in *Conference on Robot Learning*, 2022.
- [15] J. Pari, N. M. Shafiqullah, S. P. Arunachalam, and L. Pinto, “The surprising effectiveness of representation learning for visual imitation,” *arXiv preprint arXiv:2112.01511*, 2021.
- [16] J.-P. Sleiman, F. Farshidian, and M. Hutter, “Constraint handling in continuous-time ddp-based model predictive control,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 8209–8215.
- [17] M. A. Roa, M. R. Dogar, J. Pages, C. Vivas, A. Morales, N. Correll, M. Gornier, J. Rosell, S. Foix, R. Memmesheimer *et al.*, “Mobile manipulation hackathon: Moving into real world applications,” *IEEE Robotics & Automation Magazine*, vol. 28, no. 2, pp. 112–124, 2021.
- [18] J. Haviland and P. Corke, “Neo: A novel expeditious optimisation algorithm for reactive motion control of manipulators,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1043–1050, 2021.
- [19] J. Haviland, N. Sünderhauf, and P. Corke, “A holistic approach to reactive mobile manipulation,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3122–3129, 2022.
- [20] T. Welschehold, C. Dornhege, and W. Burgard, “Learning mobile manipulation actions from human demonstrations,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 3196–3201.
- [21] G. Kazhoyan, A. Hawkin, S. Koralewski, A. Haidu, and M. Beetz, “Learning motion parameterizations of mobile pick and place actions from observing humans in virtual environments,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 9736–9743.
- [22] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu *et al.*, “Rt-1: Robotics transformer for real-world control at scale,” *arXiv preprint arXiv:2212.06817*, 2022.
- [23] F. Torabi, G. Warnell, and P. Stone, “Behavioral cloning from observation,” in *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, 2018.
- [24] S. Young, D. Gandhi, S. Tulsiani, A. Gupta, P. Abbeel, and L. Pinto, “Visual imitation made easy,” in *Conference on Robot Learning*, 2021.
- [25] P. Florence, C. Lynch, A. Zeng, O. A. Ramirez, A. Wahid, L. Downs, A. Wong, J. Lee, I. Mordatch, and J. Tompson, “Implicit behavioral cloning,” in *Conference on Robot Learning*. PMLR, 2022, pp. 158–168.
- [26] M. Shridhar, L. Manuelli, and D. Fox, “Perceiver-actor: A multi-task transformer for robotic manipulation,” in *Conference on Robot Learning*. PMLR, 2023, pp. 785–799.
- [27] A. Gupta, S. Tian, Y. Zhang, J. Wu, R. Martín-Martín, and L. Fei-Fei, “Maskvit: Masked visual pre-training for video prediction,” *arXiv preprint arXiv:2206.11894*, 2022.
- [28] T. Xiao, I. Radosavovic, T. Darrell, and J. Malik, “Masked visual pre-training for motor control,” *arXiv preprint arXiv:2203.06173*, 2022.
- [29] N. Hogan, “Stable execution of contact tasks using impedance control,” *Proceedings. 1987 IEEE International Conference on Robotics and Automation*, vol. 4, pp. 1047–1054, 1987.
- [30] C. Ott, R. Mukherjee, and Y. Nakamura, “Unified impedance and admittance control,” in *2010 IEEE International Conference on Robotics and Automation*, 2010, pp. 554–561.
- [31] W. Meeussen, M. Wise, S. Glaser, S. Chitta, C. McGann, P. Mihelich, E. Marder-Eppstein, M. Muja, V. Eruhimov, T. Foote *et al.*, “Autonomous door opening and plugging in with a personal robot,” in *2010 IEEE International Conference on Robotics and Automation*. IEEE, 2010, pp. 729–736.
- [32] C. Zeng, S. Li, Y. Jiang, Q. Li, Z. Chen, C. Yang, and J. Zhang, “Learning compliant grasping and manipulation by teleoperation with adaptive force control,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 717–724.
- [33] J. Zhou, C. Wei, H. Wang, W. Shen, C. Xie, A. Yuille, and T. Kong, “ibot: Image bert pre-training with online tokenizer,” *arXiv preprint arXiv:2111.07832*, 2021.
- [34] H. Ferreau, C. Kirches, A. Potschka, H. Bock, and M. Diehl, “qpOASES: A parametric active-set algorithm for quadratic programming,” *Mathematical Programming Computation*, vol. 6, no. 4, pp. 327–363, 2014.
- [35] K. Siddharth, N. Suraj, C. Annie, S. K. Thomas, F. Chelsea, S. Dorsa, and L. Percy, “Language-driven representation learning for robotics,” *arXiv preprint arXiv:2302.12766*, 2023.

### A. Expert Data Details

For each task, we collect 30 expert demonstrations as described in Sec. IV-A. For each expert demonstration, we first filtered out frames that are static. Following [14], the expert action of a frame is computed as the state difference between the frame and its next  $N$ -th frame. We set  $N = 1$  first and then increment  $N$  such that either one of the conditions is met: 1) the translation difference is larger than 1cm; 2) the rotation difference is larger than 0.1rad; 3) the gripper status changes.

### B. Experiment Details

For MOMA-Force and its variants, the maximum rollout iteration is 200; for BC, the maximum iteration is 250. For each trial, the robot is first controlled to move its end-effector to a pre-defined initial pose. We then randomize the pose with a uniform noise  $(\delta x, \delta y, \delta z, \delta \alpha, \delta \beta, \delta \gamma)$  in which  $\delta x$ ,  $\delta y$ , and  $\delta z$  are sampled from  $U(-2, 2)$ cm;  $\delta \alpha$ ,  $\delta \beta$ , and  $\delta \gamma$  are sampled from  $U(-0.06, 0.06)$ rad (Fig. 6).  $\alpha$ ,  $\beta$ , and  $\gamma$  are the Euler angles. For all the tasks, the robot motion can be divided into three phases: approaching, grasping, and in-contact manipulation. In the approaching phase, as the robot is close to the object, the mobile base is locked and the robot only performs arm motion. After the robot grasps the object, force control is activated and the base is unlocked, *i.e.*, the robot enters whole-body control mode. We use PID control to regulate the wrench to the target wrench. The PID parameters are shown in Tab. III. The derivative term  $\partial \Delta \mathcal{F} / \partial t$  in Eq. 3 is determined by the difference between the wrench errors  $\Delta \mathcal{F}$  in two consecutive time steps.

TABLE III: PID Parameters for Force Control.

| Terms | $F_x$ | $F_y$ | $F_z$ | $M_x$ | $M_y$ | $M_z$ |
|-------|-------|-------|-------|-------|-------|-------|
| $K_p$ | 4e-4  | 4e-4  | 4e-4  | 4e-4  | 4e-4  | 4e-4  |
| $K_i$ | 5e-5  | 5e-5  | 5e-5  | 5e-5  | 5e-5  | 5e-5  |
| $K_d$ | 5e-4  | 5e-4  | 5e-4  | 5e-4  | 5e-4  | 5e-4  |

### C. Future Work

For future work, we plan to explore more on extending our method to more complex scenarios including perturbed backgrounds and novel objects. In addition, we find in the experiment that grasping is important in terms of overall task success rate and force control. While we use a parallel-jaw gripper in this work, future work can also investigate using more sophisticated grippers (*e.g.*, robot hands) which can be actively controlled during the course of rollouts.

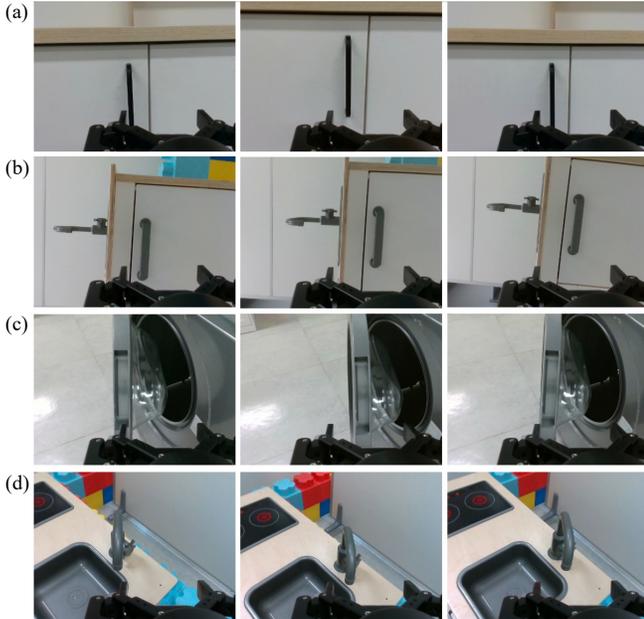


Fig. 6: **Initialization.** We show the initialize observation images from three trials of (a) open left door, (b) open right door, (c) open washing machine, and (d) rotate tap.