

HMC: Learning Heterogeneous Meta-Control for Contact-Rich Loco-Manipulation

Lai Wei*, Xuanbin Peng*, Ri-Zhao Qiu, Tianshu Huang, Xuxin Cheng, Xiaolong Wang
*equal contribution

UC San Diego

<https://loco-hmc.github.io>



Fig. 1: **Rolling out HMC for contact-rich tasks on a humanoid robot.** Compared to naïve position-only policies [5, 26], HMC designs a blending interface (HMC-Controller) that dynamically balances the torques from position tracking and force-aware controllers for contact-rich tasks. Using a heterogeneous architecture, HMC-Policy trains on both large-scale position-only data and fine-grained force-aware demonstrations. (Shown tasks: opening a wheeled cabinet; wiping a whiteboard; pulling open a door; opening a microwave; tightening a nut and lifting a trash can.)

Abstract—Learning from real-world robot demonstrations holds promise for interacting with complex real-world environments. However, the complexity and variability of interaction dynamics often cause purely positional controllers to struggle with contacts or varying payloads. To address this, we propose a Heterogeneous Meta-Control (HMC) framework for Loco-Manipulation that adaptively stitches multiple control modalities: position, impedance, and hybrid force-position. We first introduce an interface, HMC-Controller, for blending actions from different control profiles continuously in the torque space. HMC-Controller facilitates both teleoperation and policy deployment. Then, to learn a robust force-aware policy, we propose HMC-Policy to unify different controllers into a heterogeneous architecture. We adopt a mixture-of-experts style routing to learn from large-scale position-only data and fine-grained force-aware demonstrations. Experiments on a real humanoid robot show over 50% relative improvement vs. baselines on challenging tasks such as compliant table wiping and drawer opening, demonstrating the efficacy of HMC.

I. INTRODUCTION

Robots that walk and manipulate in a seamless whole — known as *loco-manipulation* systems — extend robot autonomy beyond structured environments to enable versatile operations in homes, warehouses, and disaster sites. Recent advancements in large-scale learning methods, such as imitation learning [5, 26, 37], sim-to-real transfer [30, 32], and zero-shot planning [19, 22, 42, 43], have driven significant progress in robot manipulation. Internet-scale pre-trained visual models [33, 40] and sophisticated teleoperation interfaces [12, 14, 20] have notably enhanced robotic capabilities.

However, these state-of-the-art methods primarily rely on *position-only controllers* [5, 26, 37] that track only joint angles or end effector poses, which excel in tracking accuracy yet fundamentally neglect complex interaction dynamics. On

the other hand, real-world tasks inherently involve **contact-rich interactions**, such as wiping surfaces, lifting heavy or delicate objects, or opening drawers secured by magnets. These scenarios demand precise yet adaptable regulation of both *motion* and *force*. Consequently, these controllers frequently generate hazardous oscillations and excessive forces in contact-rich tasks.

Traditional compliance methodologies, including impedance and admittance control [17, 35, 50], as well as hybrid force-position control [41, 44], address some of these limitations by either dynamically adapting stiffness [17, 35] or precisely managing forces in normal and shear directions [41, 44]. However, traditional approaches typically focus on a single, manually-tuned setting, lacking generalizability across different scenes. To address such an issue, recent methods [18, 31, 47] incorporate learning-based strategies to learn force-aware actions in a data-driven fashion. FACTR [31] and RDP [47] use force feedback to handle contact-rich interactions, while ACP [18] use imitation learning to dynamically adjust stiffness parameters. However, existing learning-based methods often rely on a specific type of compliant controller and a small set of domain-specific expert data collected by the exact controller, in addition to requiring expensive force sensors.

This paper identifies three types of challenges to develop a robust force-aware robot policy:

- 1) **Modality Mismatch:** Position, impedance, and hybrid controllers each excel in different task phases, but no single mode is sufficient for the entire task.
- 2) **Data Imbalance:** Large-scale teleoperation data is predominantly positional, whereas demonstrations leveraging compliance or force control are scarce.
- 3) **Abrupt Switching:** Hard, discrete transitions between controllers cause torque discontinuities, leading to unstable and unsafe interactions.

To overcome these challenges, we introduce *Heterogeneous Meta-Control (HMC)*. Our meta controller, HMC-Controller, is a meta-control [46] interface that takes inputs from multiple low-level control modalities (position, impedance, hybrid force-position) and computes unified torque actions. HMC-Controller operates on-the-fly in the torque space, which enables continuous blending of different control profiles based on the evolving task progress and environment. We then build HMC-Policy, which learns all control profiles simultaneously in a heterogeneous fashion. HMC-Policy predicts actions from different control profiles, which are sent to the low-level HMC-Controllers. With a soft Mixture-of-Expert (MoE) routing design, HMC-Policy autonomously weights different controllers across the task execution horizon. Unlike traditional discrete switching strategies, our soft router ensures smooth and interpretable transitions between controllers, prevents expert collapse from imbalanced data distributions, and produces stable behaviors vital for real-world deployment. Furthermore, the heterogeneous design intuitively enables a pre-train-fine-tune paradigm to train on open-sourced position-only demonstrations.

Empirical evaluations on challenging tasks show that

HMC significantly outperforms conventional fixed or discrete-switching baselines. In sum, our contributions are:

- **Unified low-level control interface.** HMC-Controller blends different control profiles continuously in the torque space for teleoperation and policy rollouts.
- **Heterogeneous high-level policy.** HMC-Policy uses heterogeneous learning and MoE-style soft routing, which trains on abundant positional data and fine-grained multi-expert demonstrations.
- **Real-world evaluation.** Evaluations on various contact-rich tasks to show notable enhancements in stability, compliance, and adaptability in real-world scenarios.

II. RELATED WORK

Manipulation Policies. Classical manipulation policies rely on model-based control or motion planning [16, 52], which often fail to generalize in unstructured real-world environments. In contrast, learning-based manipulation policies have seen rapid progress in recent years. Behavior Cloning [6] is a common approach to directly map from observations to actions. Recent developments include transformer-based models [7, 8] and diffusion-based methods [13, 21, 49], which aim to mitigate multi-modality and improve trajectory generation. Despite their success, these methods often rely on low-level positional controllers, making them ill-suited for contact-rich tasks that require reasoning about forceful interactions or multi-stage behaviors.

Whole-body Loco-Manipulation. Loco-manipulation has emerged as an increasingly important direction in robotic control, especially in settings that demand both mobility and dexterous interaction. Early efforts such as [15, 32] primarily addressed static or quasi-static tasks, focusing on balance-constrained manipulation. More recent works including [4, 11, 23, 51] tackle dynamic, contact-rich tasks that require whole-body coordination and physical stability during interaction. These systems highlight the growing demand for controllers that account for both task-specific objectives and the robot’s whole-body dynamics.

Contact-Rich Compliance Control. Contact-rich manipulation often requires compliant behavior to maintain stability and prevent damage during interaction. Traditional model-based methods such as [17, 25, 50] explicitly regulate force and motion based on physical models, but are difficult to tune and deploy in unstructured or dynamic environments. Learning-based approaches have gained traction, typically through reinforcement learning (RL) or learning from demonstration (LfD). Some RL-based methods [3] predict compliance parameters in real time, which are then passed to a low-level controller. Others [38] integrate the RL agent directly into the low-level control loop. LfD methods [1, 18, 24, 34, 45] often aim to infer both target motion and a set of compliance parameters from human demonstrations. Despite progress, these methods face a common challenge: collecting data with reliable force or impedance supervision remains difficult.

Meta Controllers. Meta-controllers aim to automate the selection and coordination of robot skills, enabling agents to

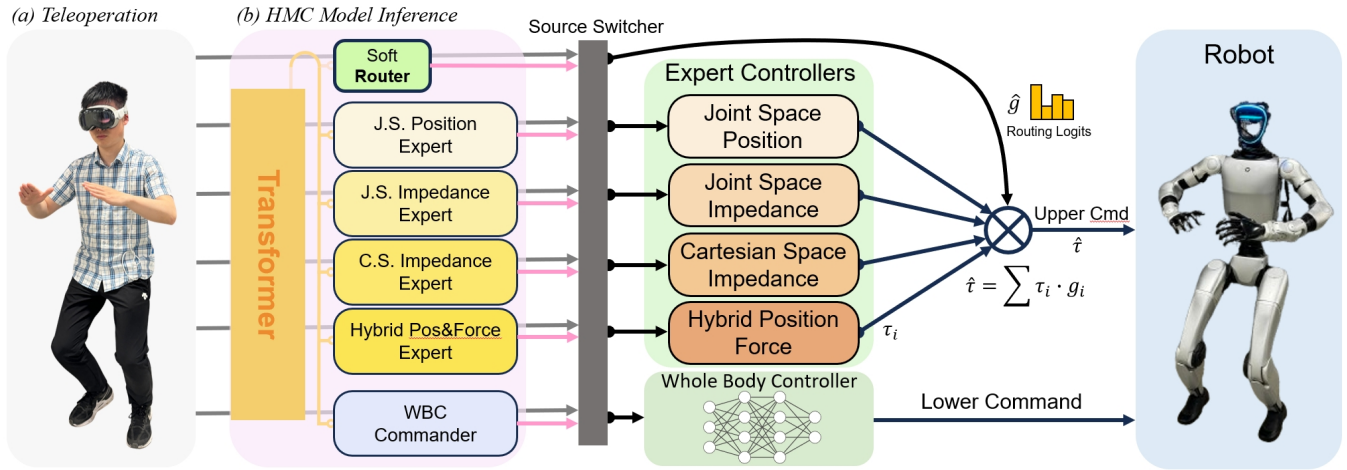


Fig. 2: **System overview.** HMC-Controller accepts inputs from either a VR-based teleoperation system or HMC-Policy inference. In the model inference path, multiple expert heads output corresponding control strategies. In the teleoperation path, shared hand poses are distributed to the expert controllers with different controller-specific parameters. All expert controllers output joint-space torque commands, which are blended via soft routing using predicted soft weights. Finally, the fused torque commands and lower-body joint targets are executed by the robot in real time. (“J.S”: Joint Space. “C.S”: Cartesian Space.)

perform diverse tasks by dynamically invoking the most suitable experts. Prior approaches include skill library methods such as [29, 39], which retrieve or blend existing policies for new tasks. More recent efforts leverage large language models (LLMs) to guide skill selection, parameter inference, or reward specification, as seen in [28, 36]. Building on this line of work, [46] introduces a hierarchical control framework that synthesizes customized models and controllers per task via LLMs, mimicking human expert reasoning. However, these LLM-based methods often face latency constraints and struggle to integrate real-time onboard observations—such as contact feedback or phase transitions—into the control loop. This makes them ill-suited for multi-phase, contact-rich tasks that require fast and reactive expert switching. Our approach similarly adopts a structured hierarchy, but focuses on explicit phase switching and expert policy selection, enabling real-time, feedback-driven adaptation rather than offline synthesis.

III. HMC-CONTROLLER INTERFACE

Our loco-manipulation meta-control system is illustrated in Fig. 2, which provides a unified interface for both teleoperation and deployment. We apply HMC to the upper body of the robot for manipulation, with an off-the-shelf whole-body humanoid controller [27] to enable whole-body dexterity and enlarged workspace.

HMC-Controller. To get the best out of position-based controllers and force-aware controllers, we implement a suite of primitive controllers: pure position, impedance, and hybrid controllers. Each controller follows standard implementation and accepts desired Cartesian and/or joint setpoints along with a controller-specific profile. The outputs are torque commands for motors.

Given the robot joint positions \mathbf{q} , **pure position controller** is a standard PD controller [2] that directly tracks the desired joint positions \mathbf{q}_d with gain coefficients K_p and D_p .

$$\boldsymbol{\tau} = K_p(\mathbf{q}_d - \mathbf{q}) + D_p(\dot{\mathbf{q}}_d - \dot{\mathbf{q}}) \quad (1)$$

Joint-space impedance controller modulates compliance at the joint level with desired joint space stiffness K_q and damping D_q .

$$\boldsymbol{\tau} = K_q(\mathbf{q}_d - \mathbf{q}) + D_q(\dot{\mathbf{q}}_d - \dot{\mathbf{q}}) + \boldsymbol{\tau}_{gravity} \quad (2)$$

Cartesian-space impedance controller regulates end-effector compliance in the Cartesian space using the Jacobian $J(\mathbf{q})$ and cartesian space stiffness K_x and damping D_x :

$$\boldsymbol{\tau} = J^\top [K_x(\mathbf{x}_d - \mathbf{x}) + D_x(\dot{\mathbf{x}}_d - \dot{\mathbf{x}})] + \boldsymbol{\tau}_{gravity} \quad (3)$$

Finally, **hybrid position-force controller** tracks position along free axes while maintaining a specified force \mathbf{f}_d along constrained directions, using selection matrices S_p and S_f :

$$\boldsymbol{\tau} = J^\top \left[S_p \left(K_x(\mathbf{x}_d - \mathbf{x}) + D_x(\dot{\mathbf{x}}_d - \dot{\mathbf{x}}) \right) + S_f(\mathbf{f}_d - \mathbf{f}) \right] + \boldsymbol{\tau}_{gravity} \quad (4)$$

Given actions corresponding to different controllers at the same timestamp, HMC-Controller computes torque for each controller, and performs a soft weighted average to obtain the final torques for execution. A low-pass filter is applied to ensure continuity and stability.

Teleoperation Architecture. We use OpenTV [12] to obtain head (\mathbf{x}_{head}) and hand (\mathbf{x}_{hand}) pose tracking at 50 Hz, which are solved to joint angles \mathbf{q}_d [9, 10]. In addition, our meta-control teleoperation system incorporates a user-friendly dashboard that enables the second operator to switch control modes and adjust parameters (e.g., impedance stiffness) on-the-fly.

As a low-cost teleoperation solution, our system dispenses with dedicated physical control sticks (e.g., [31, 48]) and end-effector force sensors. Instead, to convey cues during object interaction, we perform an online, coarse estimation of contact forces based on joint motor torques and positional errors. Together we visualize these forces on the virtual robot arm within the 3D scene to allow an immersive perception of contact dynamics than a conventional 2D display.

IV. LEARNING ALGORITHM: HETEROGENEOUS META-CONTROL POLICY

Problem Formulation and Motivation. We frame the task of meta-control for loco-manipulation as a *heterogeneous behavioral cloning* problem, where the goal is to replicate expert-level, multi-modal control behaviors demonstrated via teleoperation. Formally, we consider a demonstration dataset consisting of trajectories:

$$\tau = \{(o_t, a_t, g_t)\}_{t=1}^T, \quad a_t = \{\mathbf{EE}, \text{Imp. Stiff}, \dots, \mathbf{F}\} \quad (5)$$

where at each timestep t , the robot receives multi-sensory observations o_t , including visual o_t^{vis} and proprioceptive o_t^{prop} data. The action $a_t \in \mathbb{R}^{N_a}$ is unified across different controllers, including all potential values such as end-effector (EE) pose and attributes (e.g., impedance stiffness, force vectors). For an action in the demonstration, zero-padding is applied for irrelevant controller attributes to maintain a consistent unified vector space. A soft control belief state $g_t \in [0, 1]$ is used to indicate the relative contribution of each expert at time t , where $\sum_{i=1}^N g_{t,i} = 1$.

Critically, teleoperation logs exhibit a highly imbalanced distribution—positional data dominate, while impedance and force modality data are scarce yet crucial for contact-rich interactions. A naive single-head network tends to collapse toward purely positional behavior, neglecting critical compliant behaviors. Conversely, separate specialist networks fail to handle smooth and continuous modality transitions required in real-world applications. To tackle these challenges, we propose a two-stage training strategy and a Soft Mixture-of-Experts (Soft MoE) architecture capable of dynamically blending control modalities smoothly.

A. HMC-Policy: Architecture and Two-Stage Training

Our proposed Heterogeneous Meta-Control Policy (HMC-Policy) approach involves two main stages: (a) a positional pretraining stage leveraging abundant, easily acquired positional data, and (b) a fine-tuning stage utilizing fine-grained multi-expert demonstrations to enable smooth transitions among control modes. The architecture includes a shared Transformer-based encoder that extracts task-relevant latent embeddings, modality-specific expert heads, a soft gating mechanism (router), and a dedicated Whole-Body Control (WBC) commander head, which continuously outputs high-level whole-body directives (e.g., base velocity V_{base} , base height h_{base} , waist posture rpy , etc.) to the low-level HMC whole-body controller, ensuring stable and coordinated locomotion-manipulation behavior.

Positional data are cheap while contact dynamics vary widely; hence we adopt a shared Transformer trunk for

general spatial reasoning and modality-specific experts that specialise in their respective dynamic regimes.

Shared Transformer Trunk. Given the observation $o_t = (o_t^{\text{vis}}, o_t^{\text{prop}})$, we first tokenize each modality using separate tokenizers: a visual tokenizer $f_{\theta_{\text{vis}}}$ and a proprioceptive tokenizer $f_{\theta_{\text{prop}}}$. These tokenizers output fixed-length token embeddings:

$$v_t = f_{\theta_{\text{vis}}}(o_t^{\text{vis}}) \in \mathbb{R}^{N_v \times d}, \quad p_t = f_{\theta_{\text{prop}}}(o_t^{\text{prop}}) \in \mathbb{R}^{N_p \times d}, \quad (6)$$

where v_t and p_t represent token sequences from visual and proprioceptive modalities respectively. The combined tokens are then passed into a Transformer trunk encoder $f_{\theta_{\text{trunk}}}$, which outputs a shared latent embedding $z_t = f_{\theta_{\text{trunk}}}(v_t, p_t) \in \mathbb{R}^D$.

Modality-Specific Expert Heads. We construct N modality-specific experts, each implemented as separate MLPs:

$$a_{t,i} = g_{\theta_i}(z_t), \quad i \in \{\text{pos, imp, hybrid force}\}, \quad (7)$$

where each expert outputs an action profile including the EE pose and modality-specific parameters (joint stiffness for joint-space impedance, Cartesian stiffness for Cartesian-space impedance, force vectors for hybrid pos&force control, etc.).

Soft Router. The router head, parameterized by ϕ , predicts a gating weight (soft belief state) distribution:

$$g_t = \text{softmax}(r_{\phi}(z_t)), \quad g_t \in [0, 1]^N, \quad \sum_{i=1}^N g_{t,i} = 1. \quad (8)$$

Soft routing prevents expert collapse under imbalanced data by encouraging exploring, yields smooth transitions between control modalities, avoids abrupt torque discontinuities and enhances hardware safety during deployment. The final unified action prediction at each timestep is computed by blending the expert outputs weighted by the router outputs $\hat{a}_t = \sum_{i=1}^N g_{t,i} \cdot a_{t,i}$.

Whole-Body Control (WBC) Commander. A WBC commander head, $h_{\psi}(z_t)$, continuously outputs whole-body high-level commands, which are concatenated into the final unified action vector alongside the blended expert predictions.

$$a_t^{\text{WBC}} = h_{\psi}(z_t). \quad (9)$$

Pretraining Stage (Positional Prior Learning). Initially, we freeze all expert heads except the positional expert and WBC commander, and pretrain the Transformer trunk on large positional-only datasets. Positional data is abundant and easier to acquire, naturally covering diverse variations such as initial poses, recovery behaviors, and positional distributions. This embeds a strong positional prior into the shared representation:

$$\mathcal{L}_{\text{pretrain}} = \|a_t^{\text{pos}} - a_t^{\text{pos, gt}}\|_1 + \|a_t^{\text{WBC}} - a_t^{\text{WBC, gt}}\|_1. \quad (10)$$

Fine-tuning Stage (Multi-Expert Integration). In the fine-tuning stage, all heads are unfrozen and co-trained, while the Transformer trunk receives a reduced learning rate.

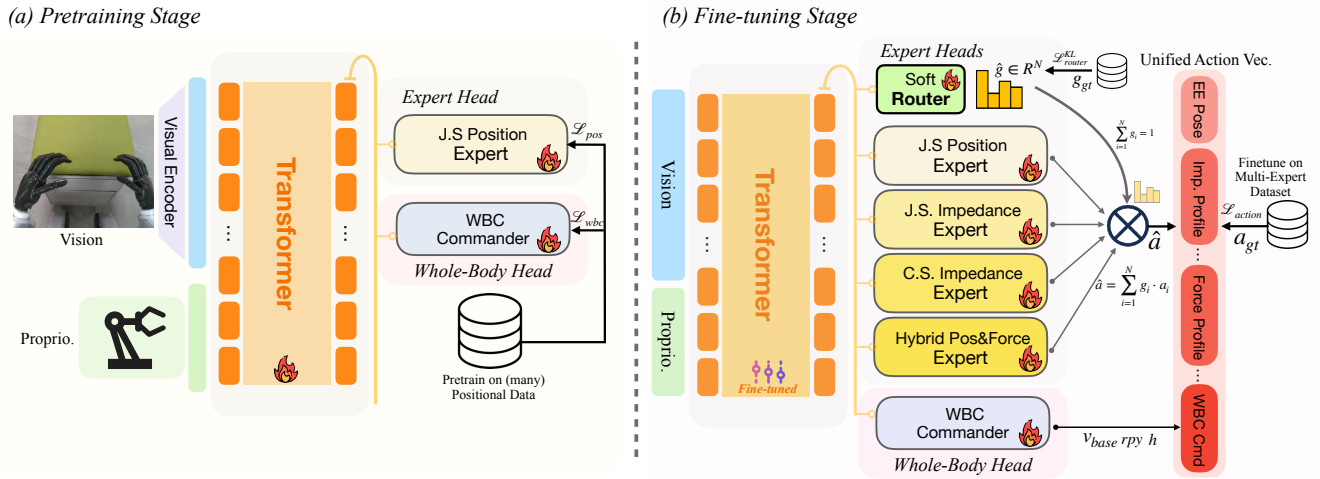


Fig. 3: **Overview of Two-stage HMC.** (a) **Pretraining:** We harness abundant positional demonstrations to train the shared transformer trunk and position expert head, thereby embedding a strong positional prior that boosts generalization. (b) **Fine-tuning Stage:** All parameters are unfrozen and fine-tuned on a smaller, fine-grained multi-expert dataset. A soft routing network learns to blend outputs from multiple experts, producing smooth and adaptive control policies. (“J.S”: Joint Space. “C.S”: Cartesian Space.)

A balanced multi-expert dataset and the soft router keep gradients flowing to every expert, avoiding modality collapse and enabling smooth blending. The final fine-tuning objective is given by

$$\mathcal{L}_{\text{finetune}} = \lambda_{\text{action}} \cdot \underbrace{\|\hat{a}_t - a_t^{gt}\|_1}_{\mathcal{L}_{\text{action}}} + \lambda_{\text{router}} \cdot \underbrace{D_{\text{KL}}(g_t^{gt} \| g_t)}_{\mathcal{L}_{\text{router}}}, \quad (11)$$

where g_t^{gt} in $\mathcal{L}_{\text{router}}$ is a smoothed ground truth signal obtained via low-pass filtering of tele-operated modality IDs, and λ are loss-balancing hyper-parameters.

V. EXPERIMENTS

A. Experiment Setup




Fig. 4: Tasks visualization.

Hardware. We conduct experiments using the Unitree G1 humanoid robot, equipped with two 7-DoF arms for loco-manipulation tasks. The visual information is acquired from an Intel RealSense D435i camera mounted on the robot’s head, and all proprioceptive data are obtained via the Unitree SDK. The robot is controlled through our HMC-Controller interface (shown in Fig. 2), enabling seamless demonstration collection and precise multi-modal control execution.

Task Setup. We evaluate our approach across three contact-rich loco-manipulation tasks (Fig. 4). **Wipe Table:**

the robot wipes the table surface to remove marker traces. This task requires **appropriate force regulation** capability. **Lift Bottle with Both Hands:** The robot is tasked with lifting a bottle using only the bare end-effectors without grippers or hand, only relying on arm friction to secure the object. This task requires both bimanual coordination to maintain closed-chain dynamic contact and lift the bottle stably. **Open Drawer:** The humanoid crouches and inserts its hand into a narrow handle slot (Hand Insertion), then pulls to open a drawer secured by a magnetic latch, requiring over 20N of force to overcome closure resistance (Pulling). This task requires: (1) **stage-aware control** to switch between compliant insertion and forceful pulling stages to achieve high success rate; (2) **dynamic stiffness modulation** to adjust impedance properties during interaction with the drawer; and (3) **whole-body stability** to maintain lower-body balance while the upper body applies large interaction forces.

Baselines. we implement the baselines as follows:

- **ACT (vanilla):** Action Chunking Transformer (ACT) policy, trained solely on positional data.
- **ACT (meta):** ACT policy trained with positional data and controller profiles. Hard routing switch.
- **Stiff Policy:** Our model architecture trained exclusively on stiff positional expert data.
- **Compliant Policy:** Our model architecture trained exclusively on impedance experts data.
- **HMC w/o soft routing:** Our HMC architecture with a hard argmax router instead of soft routing.
- **HMC (from scratch):** Our HMC architecture trained from scratch instead of two stages training.
- **HMC (ours):** The complete implementation of our proposed HMC-Policy.

Evaluation Protocols. For the Open Drawer and Lift Bot-

Method	Lift Bottle(%)		Open Drawer(%)			
	Seen	Unseen	Seen Hand Insert	Pull	Unseen Hand Insert	Pull
Stiff Policy	67	33	80	80	33	20
Compliant Policy	80	40	80	53	67	13
HMC (w/o soft routing)	93	53	80	80	53	47
HMC (from scratch)	87	60	60	60	47	33
HMC (ours)	93	80	87	87	67	67

TABLE I: Ablation experiment results across seen and unseen settings for the Lift Bottle and Open Drawer tasks. The proposed soft routing and position-only pretraining increases success rate considerably. Hand insertion and pulling correspond to two separate stages in opening a drawer.




Fig. 5: **Interpretability of routing logits** in visualization of an Open Drawer Episode. The upper figure is the predicted routing logits, while the lower figure shows the predicted right-hand stiffness (N/m). Five images from the head camera are demonstrated.

the tasks, we evaluate generalization under seen and unseen settings, where the seen settings correspond to configurations encountered during training, and unseen ones involve novel object instances or initial positions for Lift Bottle and Open Drawer tasks respectively. Each policy variant is evaluated over 15 trials, with performance assessed based on the task success rate.

B. Evaluation Results

The general experiment results across different policies are shown in Table II, and the ablation studies result are demonstrated in Table I. We try to answer important questions in our evaluation:

- **Q1:** What advantages does our HMC-Policy have compared to baselines?
- **Q2:** Why do we need HMC-Controller for contact-rich and multi-stage tasks?
- **Q3:** What are benefits of Soft Routing?
- **Q4:** Does pretrain and finetune approach helps training and generalization performance?
- **Q5:** Why do our HMC-Policy excels at contact-rich unseen tasks?

A1: Our HMC consistently outperforms baseline approaches across all three tasks. The stiff policies, though widely used, struggle in contact-rich scenarios—for example, in Wipe Table their rigidity produced excessive torques,

	Wipe (%)	Lift (%)	Drawer (%)
Stiff Policy	33	67	80
ACT (vanilla)	33	60	40
ACT (meta)	47	67	47
HMC (w/o soft routing)	87	93	80
HMC (ours)	93	93	87

TABLE II: Policy success rates on three contact-rich tasks, including wiping table, lifting bottle and opening drawer.

sometimes triggering motor overheating. ACT (meta) improves upon its vanilla version by jointly predicting trajectories and meta-controller profiles, but its single-head design makes joint learning of low-level commands and high-level modes difficult, often causing unstable behaviors in complex tasks. By contrast, HMC separates meta-level decision-making from trajectory generation and dynamically routes to expert controllers, reducing learning complexity and ensuring robust execution in multi-stage manipulation.

A2: HMC-Controller is essential as it enables adaptive control switching across task phases. Contact-rich tasks require distinct strategies at different stages. In the Open Drawer example, insertion benefits from compliance to tolerate alignment errors, while pulling demands stiffness to overcome resistance. As shown in Table I, fixed controllers (stiff-only or compliant-only) fail in unseen settings

since neither adapts to both needs. By contrast, HMC-Policy flexibly routes between expert controllers (Fig. 5), using compliance for insertion and increasing stiffness for pulling, thereby achieving higher success through phase-aware control routing.

A3: Soft routing enables smooth and stable transitions between expert controllers. This is critical in contact-rich tasks, where abrupt switching can destabilize the interaction. In HMC-Controller framework, soft routing allows for continuous blending of control outputs, ensuring robustness during phase transitions. In the Lift Bottle task under unseen settings, we observed that the HMC-Policy without soft routing often triggered an abrupt expert controller switch when the arm made contact with the bottle. Subtle discrepancies between expert controllers caused sudden changes in gripping force, occasionally leading to slippage of the cylindrical bottle. By contrast, soft routing allowed the robot to interpolate control outputs across experts, maintaining stability during transitions.

A4: Pretraining followed by finetuning improves both training stability and generalization. Compared to the training-from-scratch variant in Table I, the pretrain-and-finetune strategy significantly enhances performance, particularly in unseen scenarios. During pretraining on position-only data, the shared stem network learns fundamental task representations and spatial alignment. Subsequent finetuning on heterogeneous controller data focuses learning on expert controller outputs, leveraging a well-initialized stem. This two-stage process mitigates conflicts among expert losses and stabilizes gradients during training. Moreover, the pretrained stem provides a strong prior, improving generalization and reducing the risk of overfitting to training environments.

A5: Our HMC-Policy excels in contact-rich unseen tasks by combining adaptive control routing with strong generalization. Contact-rich tasks often involve unpredictable variations in contact dynamics, requiring adaptive and robust control strategies. Our HMC-Policy framework addresses this by combining meta-level scheduling of diverse expert controllers with smooth transitions enabled by soft routing. Additionally, the pretrain-and-finetune paradigm provides a strong initialization, improving the generalizability with existing position-only data.

VI. CONCLUSION

Contact-rich loco-manipulation exposes the limits of single-mode, position-only control. In this work, we tackled this gap with Heterogeneous Meta-Control (HMC), a unified framework in which HMC-Controller blends heterogeneous commands directly in torque space, while HMC-Policy learns via a two-stage pretrain-finetune paradigm, and Mixture-of-Experts to route those modalities smoothly in response to task phase and contact dynamics. This heterogeneous design compensates for the heavy positional bias in teleoperation data, keeps all experts trainable, and removes the torque discontinuities that plague hard switching.

Real-robot experiments on the Unitree G1 (wiping, bottle lifting, drawer opening) show that HMC-Policy surpasses

various baselines, yielding higher success, smoother forces, and stronger generalization. These results highlight adaptive soft routing as a practical approach to robust loco-manipulation, with future work aimed at scaling to more experts and integrating task-level planners for long-horizon autonomy.

VII. ACKNOWLEDGEMENT

This project was supported, in part, by NSF CAREER Award IIS-2240014, NSF CCF-2112665 (TILOS), gifts from Amazon, Meta, and Qualcomm.

REFERENCES

- [1] M. Aburub et al., “Learning diffusion policies from demonstrations for compliant contact-rich manipulation,” *arXiv preprint arXiv:2410.19235*, 2024.
- [2] K. H. Ang et al., “Pid control system analysis, design, and technology,” *IEEE transactions on control systems technology*, vol. 13, no. 4, pp. 559–576, 2005.
- [3] C. C. Beltran-Hernandez et al., “Learning force control for contact-rich manipulation tasks with rigid position-controlled robots,” *IEEE Robotics and Automation Letters*, vol. 5, no. 4, 5709–5716, Oct. 2020.
- [4] Q. Ben et al., “Homie: Humanoid loco-manipulation with isomorphic exoskeleton cockpit,” *arXiv preprint arXiv:2502.13013*, 2025.
- [5] K. Black et al., “ π_0 : A vision-language-action flow model for general robot control,” *arXiv preprint arXiv:2410.24164*, 2024.
- [6] M. Bojarski et al., *End to end learning for self-driving cars*, 2016.
- [7] A. Brohan et al., *Rt-1: Robotics transformer for real-world control at scale*, 2023.
- [8] A. Brohan et al., *Rt-2: Vision-language-action models transfer web knowledge to robotic control*, 2023.
- [9] S. Caron, *Pink: Python inverse kinematics*, <https://github.com/stephane-caron/pink>, Accessed: 2025-05-23, 2021.
- [10] J. Carpentier et al., “Pinocchio: Fast forward and inverse dynamics for poly-articulated systems,” in *IEEE International Symposium on System Integrations (SII)*, IEEE, 2019, pp. 614–619.
- [11] Z. Chen et al., “Gmt: General motion tracking for humanoid whole-body control,” *arXiv preprint arXiv:2506.14770*, 2025.
- [12] X. Cheng et al., “Open-television: Teleoperation with immersive active visual feedback,” *arXiv preprint arXiv:2407.01512*, 2024.
- [13] C. Chi et al., “Diffusion policy: Visuomotor policy learning via action diffusion,” *The International Journal of Robotics Research*, p. 02 783 649 241 273 668, 2023.
- [14] C. Chi et al., “Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots,” *arXiv preprint arXiv:2402.10329*, 2024.
- [15] Z. Fu et al., *Deep whole-body control: Learning a unified policy for manipulation and locomotion*, 2022.

- [16] H. Guo et al., “Recent trends in task and motion planning for robotics: A survey,” *ACM Comput. Surv.*, vol. 55, no. 13s, Jul. 2023.
- [17] N. Hogan, “Impedance control: An approach to manipulation: Part i—theory,” *Journal of Dynamic Systems, Measurement, and Control*, vol. 107, no. 1, pp. 1–7, Mar. 1985.
- [18] Y. Hou et al., “Adaptive compliance policy: Learning approximate compliance for diffusion guided control,” *arXiv preprint arXiv:2410.09309*, 2024.
- [19] W. Huang et al., “Rekep: Spatio-temporal reasoning of relational keypoint constraints for robotic manipulation,” in *CoRL*, 2024.
- [20] A. Iyer et al., “Open teach: A versatile teleoperation system for robotic manipulation,” *arXiv preprint arXiv:2403.07870*, 2024.
- [21] M. Janner et al., *Planning with diffusion for flexible behavior synthesis*, 2022.
- [22] M. Ji et al., “Graspsplats: Efficient manipulation with 3d feature splatting,” in *CoRL*, 2024.
- [23] M. Ji et al., “Exbody2: Advanced expressive humanoid whole-body control,” *arXiv preprint arXiv:2412.13196*, 2025.
- [24] T. Kamijo et al., “Learning variable compliance control from a few demonstrations for bimanual robot with haptic feedback teleoperation system,” *arXiv preprint arXiv:2406.14990*, 2024.
- [25] A. Q. Keemink et al., “Admittance control for physical human–robot interaction,” *The International Journal of Robotics Research*, vol. 37, no. 11, pp. 1421–1444, 2018.
- [26] M. J. Kim et al., “Openvla: An open-source vision-language-action model,” in *Conference on Robot Learning (CoRL)*, 2024.
- [27] J. Li et al., “Amo: Adaptive motion optimization for hyper-dexterous humanoid whole-body control,” *Robotics: Science and Systems 2025*, 2025.
- [28] J. Liang et al., “Code as policies: Language model programs for embodied control,” in *International Conference on Learning Representations (ICLR)*, 2023.
- [29] Z. Liang et al., *Skilldiffuser: Interpretable hierarchical planning via skill abstractions in diffusion-based task execution*, 2024.
- [30] T. Lin et al., “Sim-to-real reinforcement learning for vision-based dexterous manipulation on humanoids,” *arXiv preprint arXiv:2502.20396*, 2025.
- [31] J. J. Liu et al., *Factr: Force-attending curriculum training for contact-rich policy learning*, 2025.
- [32] M. Liu et al., “Visual whole-body control for legged loco-manipulation,” *arXiv preprint arXiv:2403.16967*, 2024.
- [33] S. Liu et al., “Grounding dino: Marrying dino with grounded pre-training for open-set object detection,” in *ECCV*, 2024.
- [34] W. Liu et al., “Forcemimic: Force-centric imitation learning with force-motion capture system for contact-rich manipulation,” *arXiv preprint arXiv:2410.07554*, 2024.
- [35] T. Lozano-Pérez et al., “Automatic synthesis of fine-motion strategies for robots,” *The International Journal of Robotics Research*, vol. 3, no. 1, pp. 3–24, 1984.
- [36] Y. Ma et al., “Eureka: Human-level reward design via coding large language models,” *arXiv preprint arXiv:2310.01894*, 2023.
- [37] A. O’Neill et al., “Open x-embodiment: Robotic learning datasets and rt-x models,” *arXiv preprint arXiv:2310.08864*, 2023.
- [38] T. Portela et al., “Learning force control for legged manipulation,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2024, pp. 15 366–15 372.
- [39] R.-Z. Qiu et al., *Wildlma: Long horizon loco-manipulation in the wild*, 2025.
- [40] A. Radford et al., “Learning transferable visual models from natural language supervision,” in *ICML*, 2021.
- [41] M. H. Raibert et al., “Hybrid position/force control of manipulators,” *Journal of Dynamic Systems, Measurement, and Control*, vol. 103, no. 2, pp. 126–133, Jun. 1981.
- [42] A. Rashid et al., “Language embedded radiance fields for zero-shot task-oriented grasping,” in *CoRL*, 2023.
- [43] W. Shen et al., “Distilled feature fields enable few-shot language-guided manipulation,” in *CoRL*, 2023.
- [44] M. Uchiyama et al., “A symmetric hybrid position/force control scheme for the coordination of two robots,” in *Proceedings. 1988 IEEE International Conference on Robotics and Automation*, 1988, 350–356 vol.1.
- [45] L. Wei et al., *Ensuring force safety in vision-guided robotic manipulation via implicit tactile calibration*, 2024.
- [46] T. Wei et al., *Meta-control: Automatic model-based control synthesis for heterogeneous robot skills*, 2024.
- [47] H. Xue et al., “Reactive diffusion policy: Slow-fast visual-tactile policy learning for contact-rich manipulation,” *arXiv preprint arXiv:2503.02881*, 2025.
- [48] S. Yang et al., *Ace: A cross-platform visual-exoskeletons system for low-cost dexterous teleoperation*, 2024.
- [49] Y. Ze et al., *3d diffusion policy: Generalizable visuomotor policy learning via simple 3d representations*, 2024.
- [50] T. Zhang et al., “Whole-body compliance control for quadruped manipulator with actuation saturation of joint torque and ground friction,” in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024, pp. 11 124–11 131.
- [51] Y. Zhang et al., “Falcon: Learning force-adaptive humanoid loco-manipulation,” *arXiv preprint arXiv:2505.06776*, 2025.
- [52] C. Zhou et al., *A review of motion planning algorithms for intelligent robotics*, 2021.