Robotics: Science and Systems 2026
Sydney, Australia, July 13–17, 2026

Force Policy: Learning Hybrid Force-Position Control Policy under Interaction Frame for Contact-Rich Manipulation

Hongjie Fang[*,1,3], Shirun Tang[*,1], Mingyu Mei[*,1,4], Haoxiang Qin[3], Zihao He[3], Jingjing Chen[3], Ying Feng[3,5], Chenxi Wang[1], Wanxi Liu[1,2], Zaixing He[4], Cewu Lu[1,2,3,5], Shiquan Wang[1,2,†]
[1]Noematrix [2]Flexiv [3]Shanghai Jiao Tong University [4]Zhejiang University [5]Shanghai Innovation Institute
[*]Equal Contribution [†]Corresponding Author

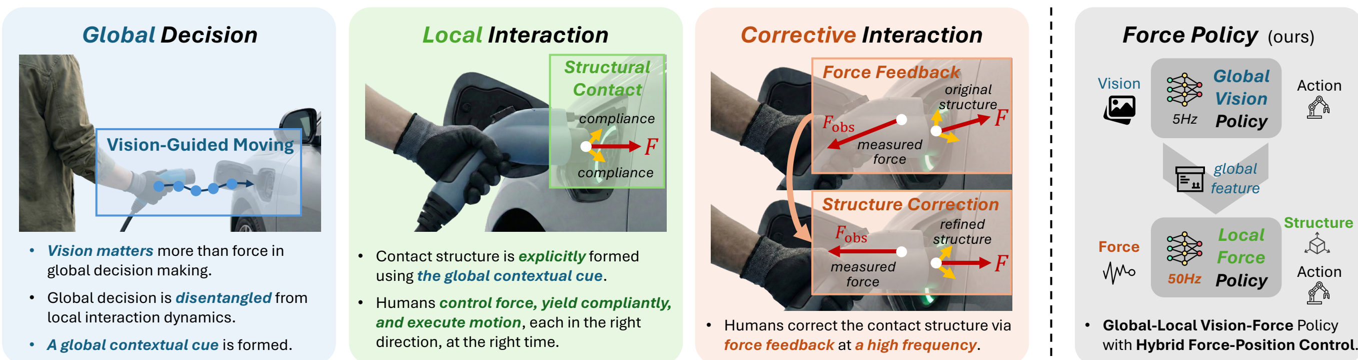

Fig. 1: A Global-Local Vision-Force Policy Inspired by Human Interaction. *(Left)* During electric vehicle (EV) charging, humans **use vision to guide global movement** and coarse alignment, and rely on **force feedback** to continuously adjust the **local contact structure**. *(Right)* This global-local organization motivates a vision-force policy with an explicit contact structure for stable hybrid force-position control.

***Abstract*—Contact-rich manipulation demands human-like integration of perception and force feedback: vision should guide task progress, while high-frequency interaction control must stabilize contact under uncertainty. Existing learning-based policies often entangle these roles in a monolithic network, trading off global generalization against stable local refinement, while control-centric approaches typically assume a known task structure or learn only controller parameters rather than the structure itself. In this paper, we formalize a physically grounded interaction frame, an instantaneous local basis that decouples force regulation from motion execution, and propose a method to recover it from demonstrations. Based on this, we address both issues by proposing *Force Policy*, a global-local vision-force policy in which a global policy guides free-space actions using vision, and upon contact, a high-frequency local policy with force feedback estimates the interaction frame and executes hybrid force-position control for stable interaction. Real-world experiments across diverse contact-rich tasks show consistent gains over strong baselines, with more robust contact establishment, more accurate force regulation, and reliable generalization to novel objects with varied geometries and physical properties, ultimately improving both contact stability and execution quality. Project website: force-policy.github.io.**

I. Introduction

Human dexterity is defined not merely by motion, but by the seamless integration of force modulation and feedback during physical interaction. Replicating this level of competence in robotic systems is the central challenge of contact-rich manipulation. Operations such as tightening a screw [34], peeling a vegetable [14, 29], or polishing a surface [31, 79] impose strict constraints where the robot must simultaneously control motion and interaction forces. Mastering contact-rich manipulation is essential for deploying robots in real-world settings, from factory assembly to everyday household tasks [66].

A defining aspect of how humans succeed in contact-rich tasks is **a global-local organization of perception and control** [13, 72], as illustrated in Fig. 1 (left). At the global level, humans primarily use vision to decide *what to do next*, *where to go*, and *how to align* [26]. At the local level, once contact happens, humans rely on high-frequency force feedback to *refine execution in a structured way*, applying or maintaining interaction where needed, yielding compliantly to uncertainty, and executing motion to proceed [27]. Functionally, this implies a clear division of roles: the global component should generalize across task variations, while the local component should guarantee stable contact. This suggests two core questions for robotics:

(1) Global-local organization. How can we realize this global-local organization, so that global guidance generalizes well while local interaction remains stable?

(2) Interaction structure for control. How can we represent the task interaction structure explicitly, so that it transfers across diverse contact skills?

For **(1)**, many learning-based policies still adopt a monolithic design. Force signals are often appended to the observation [29, 80] or used as auxiliary supervision [82], but perception, planning, and contact refinement are learned jointly in a single end-to-end network. This creates an inherent tradeoff: entangling local refinement with global perception makes it difficult to maintain a dedicated high-frequency interaction loop while improving global generalization by scaling up models. RDP [79] is one of the few works that explicitly introduces a slow-fast policy design, with fast reactive corrections handled by a force-aware tokenizer. However, its slow policy depends on the fast tokenizer rather than treating reactivity as a plug-and-play component, which limits reuse with arbitrary slow policies and complicates upgrading the slow policy due to the dependency of the fast reactive part.

For **(2)**, a substantial body of prior work approaches contact-rich manipulation through the control interface. Classical compliant control [49] provides principled tools such as hybrid position/force control [59] and impedance control [30]. Conventional approaches apply these tools through task-specific environment/contact modeling and manual parameter tuning to match a pre-defined interaction structure. Learning-based variants aim to reduce manual tuning by predicting controller parameters from data, such as admittance gains [62, 77], or contact forces [20, 45]. However, the interaction structure is usually assumed or implicit, and learning is applied mainly to tuning parameters rather than to representing the structure itself. ACP [31] is a partial exception by predicting a stiffness matrix, but its directional formulation is insufficient for multi-axis constrained interactions such as peg insertion.

In this paper, we answer both questions by making the missing structure explicit and building a global-local vision-force policy around it. For **(2)**, we formalize the interaction frame, an instantaneous local basis that captures task-relevant interaction structure, from a physical perspective, and we present a principled approach to recover it directly from non-ideal real-world demonstrations. For **(1)**, we propose ***Force Policy***, a global-local vision-force policy with phase-wise authority, as shown in Fig. 1 (right). A global vision policy drives task progress, while a high-frequency force module estimates the interaction frame and executes hybrid force-position control during contact. This decoupling preserves task-level planning while ensuring stable contact and reliable execution under uncertainty. Extensive real-world experiments show improved robustness and force regulation across diverse contact-rich tasks, while retaining strong generalization to novel objects with varied geometries and physical properties. Overall, ***Force Policy*** combines the generalization of visuomotor policies with the precision of force control, enabling reliable contact-rich manipulation.

II. Related Works

A. Contact-Rich Manipulation

Compared to free-space motion, contact-rich manipulation requires precise force regulation and rapid adaptation to contact constraints [66, 76]. Uncertainties in contact geometry, friction, and material compliance make interaction dynamics difficult to model reliably [21, 49]. Classical robotics addresses these challenges with compliant control frameworks, including impedance control [12, 30], admittance control [37, 40, 63], and hybrid force-position control [17, 38, 59]. Nonetheless, these approaches often rely on hand-tuned parameters and reasonably accurate environment assumptions, which can limit robustness in unstructured settings [73].

Learning-based methods have recently improved contact-rich manipulation by learning interaction strategies directly from data. Reinforcement learning [35, 41, 53, 81] can acquire complex contact behaviors through exploration, but frequently struggles with sim-to-real transfer due to mismatches in perception and contact dynamics [19]. Hence, imitation learning has increasingly leveraged richer sensing to better ground interaction, incorporating modalities like audio [25, 41, 42, 47, 75], tactile [23, 25, 32, 42, 60], and force/torque signals [11, 31, 36, 45, 77, 84], to handle contact-rich manipulations.

B. Force-Aware Manipulation Policies

Recent advances have explored different strategies to incorporate force sensing into learning-based manipulation. A common strategy treats force as an auxiliary observation, fusing it with visual features to condition motion prediction. FoAR [29] dynamically weights visual and force via contact-phase prediction. RDP [79] couples a visual latent diffusion policy with a high-frequency tactile-aware tokenizer for reactive execution. ForceVLA [80] integrates force through a mixture-of-experts representation, while TA-VLA [82] adds torque prediction as an auxiliary objective to promote physically grounded features.

A distinct body of work learns imitation policies within compliant control frameworks, spanning variable impedance and admittance formulations [1, 31, 36, 62, 77, 84]. Even when driven by predicted wrenches, these controllers often tie force realization to motion through a single compliance model, so what the policy should imitate can drift with contact geometry and local stiffness. In contrast, hybrid force-position control provides a more interpretable interface by explicitly separating constrained and free directions: it lets the policy express "*what should be enforced*" (force objectives in constraint directions) and "*what should be achieved*" (position objectives in free directions), yielding clearer credit assignment and more stable learning across contact transitions. While [45] adopts a hybrid scaffold, it applies force control only along a single direction, capturing only a limited form of this factorization.

C. Control Structure Discovery from Demonstrations

The problem of inferring control structures from demonstration data has progressed from analytical geometric modeling [9, 10, 49] to data-driven learning approaches [39, 64, 71]. Across this spectrum, the goal is to recover the underlying control structure of a task, typically a decomposition that enables compliant manipulation. Because this structure is induced by task constraints, inferring the control structure is essentially equivalent to inferring the constraints from demonstration data.

Task frame formalization (TFF) [9, 49] provides the canonical force-motion subspace decomposition, but typically assumes explicit geometry and rigid contact. Subsequent works infer geometric constraints from demonstrations via twist-wrench analysis [65, 67, 68, 69] or by fitting a set of pre-defined constraint models [64], with refinements using interactive perception [44, 58] or augmentation [43]. Some works [20, 45] exploit contact force direction as the force control axis in hybrid force-position control [59]. Recent attempts improve robustness by aggregating frame estimates from multiple model-based fits using statistical fusion [50], or by estimating task-aligned frames via kinematics [54] or power-based objectives [55].

In this work, we generalize the definition of task frame into a physically-grounded interaction frame and propose compact approximations based on energy dissipation, enabling reliable recovery from data for different contact-rich tasks.

III. Theoretical Formulation

A. Interaction Frame: A Physical Perspective

In classical TFF [9], the interaction frame is prescribed based on known geometric models. However, in unstructured environments where geometry is unknown or uncertain, such a prescription is infeasible. To bridge this gap, we reformulate **interaction frame (IF)**, described in [55], *purely from the physical response*, using stiffness as a proxy to estimate the local geometry. We focus on tasks maintaining *topologically invariant contact*, excluding irreversible changes like fracture and plastic deformations. For these *locally conservative* contacts, the environmental stiffness is *geometry-induced*, as resistance arises directly from physical boundaries. This inherent physical causality implies that the principal axes of stiffness structurally align with the local surface geometry, allowing us to recover the geometric frame from physical observations.

Formally, we model the local environment response at the interaction point p_I by *symmetric* stiffness $\mathbf{K}_{\text{env}}(p_I)$. We spectrally decompose $\mathbf{K}_{\text{env}}(p_I)$ into principal stiffnesses λ_i and axes $\mathbf{q}_i$, partitioning the interaction space into the **constraint subspace** $\mathcal{U}(p_I)$ with high stiffness $\lambda_i \gg 0$ and **admissible-motion subspace** $\mathcal{T}(p_I)$ with negligible stiffness $\lambda_i \approx 0$.

Crucially, physical interaction is *goal-directed*. We introduce the task intent $\{\boldsymbol{\xi}^*(p_I), \boldsymbol{\mathcal{W}}^*(p_I)\}$ to represent the *driving factors* of the interaction — specifically, the intended motion twist $\boldsymbol{\xi}^*(p_I)$ and interaction wrench $\boldsymbol{\mathcal{W}}^*(p_I)$. This formulation distinguishes our approach from TFF, where ideal variables in [9] describe the resultant kinematic constraints without differentiating between causal actuation and environmental reaction. In our framework, **the intent represents the active cause compatible with the physical structure**: the motion intent $\boldsymbol{\xi}^*(p_I)$ acts as the driver within the admissible-motion subspace $\mathcal{T}(p_I)$, while the wrench intent $\boldsymbol{\mathcal{W}}^*(p_I)$ acts as the driver against the constraint subspace $\mathcal{U}(p_I)$. We thus define the IF $\Sigma(p_I)$ by anchoring the spectral axes derived from $\mathbf{K}_{\text{env}}(p_I)$ to these driving intents:

$$\Sigma(p_I) \triangleq \Psi\left(\mathbf{K}_{\text{env}}(p_I), \boldsymbol{\xi}^*(p_I), \boldsymbol{\mathcal{W}}^*(p_I)\right). \tag{1}$$

Specifically, we set the z-axis of IF to the direction of the dominant wrench component in $\boldsymbol{\mathcal{W}}^*$. We then define the x-axis by projecting the motion intent $\boldsymbol{\xi}^*$ onto the plane orthogonal to z. In degenerate cases where this projection is zero, *e.g.*, static holding with negligible motion or screw driving with collinear motion, the interaction is effectively transversely *isotropic* in

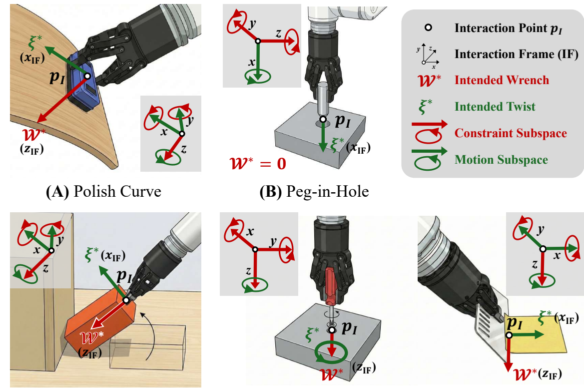

(A) Polish Curve **(B)** Peg-in-Hole **(C)** Flip Box **(D)** Tighten Screw **(E)** Scrape off Sticker

Fig. 2: Interaction Frame for Example Contact-Rich Tasks.

the constraint plane, so the rotation about z is physically irrelevant. We therefore default Ψ to a canonical reference and complete the frame by the right-hand rule.

To illustrate this formulation in practice, Fig. 2 visualizes the derived IF for a representative set of contact-rich tasks. This physics-aware formulation retains the orthogonal decomposition power of TFF but replaces geometric priors with environmental compliance, making it intrinsically suitable for unstructured tasks. Having defined the frame theoretically, we next address the practical challenge of recovering the IF directly from interaction data.

B. Recovering Interaction Frame from Interaction Signals

We aim to recover IF directly from the observed interaction signals $\boldsymbol{\xi}$ and $\boldsymbol{\mathcal{W}}$. Assume that the gravity and inertial forces are compensated from the observed wrench. In the following, we omit the dependence of p_I in the notation for brevity.

Under ideal *rigid and frictionless* contact, the observed twist and wrench perfectly align with the task intent, *i.e.*, $\boldsymbol{\xi} = \boldsymbol{\xi}^*$ and $\boldsymbol{\mathcal{W}} = \boldsymbol{\mathcal{W}}^*$, yielding zero power: $\boldsymbol{P} = \boldsymbol{\mathcal{W}}^\top \boldsymbol{\xi} = \mathbf{0}$. In the non-ideal situations, the intended twist $\boldsymbol{\xi}^*$ may induce a parasitic wrench $\boldsymbol{\mathcal{W}}_c$ along its direction. Since humans tend to align with the principal axes of the environmental stiffness $\mathbf{K}_{\text{env}}$ in expert demonstrations, the intended wrench $\boldsymbol{\mathcal{W}}^*$ may induce a parasitic twist $\boldsymbol{\xi}_c$ along its direction. Hence,

$$\boldsymbol{P} = (\boldsymbol{\mathcal{W}}^* + \boldsymbol{\mathcal{W}}_c)^\top (\boldsymbol{\xi}^* + \boldsymbol{\xi}_c) = \boldsymbol{\mathcal{W}}_c^\top \boldsymbol{\xi}^* + \boldsymbol{\mathcal{W}}^{*\top} \boldsymbol{\xi}_c. \tag{2}$$

The two terms correspond to distinct power sources: $\boldsymbol{\mathcal{W}}_c^\top \boldsymbol{\xi}^*$ acts in $\mathcal{T}$ as a **dissipative residual**, arising from frictional or viscous losses along the intended motion direction, while $\boldsymbol{\mathcal{W}}^{*\top} \boldsymbol{\xi}_c$ acts in $\mathcal{U}$ as a **structural residual**, resulting from environmental stiffness or contact constraints such as compliant deflections along the intended wrench direction. Consequently:

(1) If the structural residual dominates, we obtain $\boldsymbol{\mathcal{W}}^* \approx \boldsymbol{\mathcal{W}}$, and the twist can be orthogonalized against the wrench:

$$\boldsymbol{\xi}^* = \boldsymbol{\xi} - \text{Proj}_{\boldsymbol{\mathcal{W}}^*}(\boldsymbol{\xi}) \approx \boldsymbol{\xi} - \text{Proj}_{\boldsymbol{\mathcal{W}}}(\boldsymbol{\xi}). \tag{3}$$

(2) If the dissipative residual dominates, we obtain $\boldsymbol{\xi}^* \approx \boldsymbol{\xi}$, and the wrench can be orthogonalized against the twist:

$$\boldsymbol{\mathcal{W}}^* = \boldsymbol{\mathcal{W}} - \text{Proj}_{\boldsymbol{\xi}^*}(\boldsymbol{\mathcal{W}}) \approx \boldsymbol{\mathcal{W}} - \text{Proj}_{\boldsymbol{\xi}}(\boldsymbol{\mathcal{W}}). \tag{4}$$

This framework contextualizes the limitations of prior art: [20, 45] exclusively address structural dominance, effectively ignoring the dissipative dominance regime critical for friction-heavy tasks. [50] relies on statistical averaging, which lacks physical interpretability and obscures the intrinsic environmental properties. [55] directly performs power minimization, which may converge to spurious frames that satisfy orthogonality but fail to capture the true task structure.

IV. Method

In this section, we first present how to recover the control structure of contact-rich tasks from demonstrations (§IV-A). We then introduce ***Force Policy***, a global-local vision-force policy inspired by the human organization of perception and control (§IV-B). Finally, we present a dual-policy asynchronous scheduler that manages dual-frequency policy execution during deployment to ensure smooth trajectories (§IV-C).

Let $\tau = (I_{0:T}, s_{0:T}, e_{0:T}, \boldsymbol{\xi}_{0:T}, \boldsymbol{\mathcal{W}}_{0:T})$ denote a demonstration of task $\mathcal{T}$ with horizon T, where I_t is the visual observation, s_t the end-effector pose, and e_t the end-effector state at timestep t. The robot twist $\boldsymbol{\xi}_t \in \mathfrak{se}(3)$ is composed of its linear velocity $\boldsymbol{v}_t$ and angular velocity $\boldsymbol{\omega}_t$, *i.e.*, $\boldsymbol{\xi}_t = [\boldsymbol{v}_t^\top, \boldsymbol{\omega}_t^\top]^\top$, while the wrench $\boldsymbol{\mathcal{W}}_t \in \mathfrak{se}^*(3)$ is composed of its force $\boldsymbol{f}_t$ and moment $\boldsymbol{m}_t$, *i.e.*, $\boldsymbol{\mathcal{W}}_t = [\boldsymbol{f}_t^\top, \boldsymbol{m}_t^\top]^\top$. We assume $\boldsymbol{\mathcal{W}}_t$ is gravity-compensated, and inertial effects are negligible given the smooth motion profile of human demonstrations.

A. Control Structure of Contact-Rich Tasks

Interaction Frame Identification. Consider a Δt patch of a demonstration starting from timestep t. In this work, we do not address the estimation of the exact contact point p_I; instead, we assume the IF origin is anchored at the end-effector position. We aim to recover the IF orientation from visual observations $I_{t:t+\Delta t}$, measured signals $\boldsymbol{\xi}_{t:t+\Delta t}$ and $\boldsymbol{\mathcal{W}}_{t:t+\Delta t}$, together with the task description $\mathcal{T}$. Inferring the ideal interaction frame directly from these signals is ill-posed, as identical local power exchange patterns may arise from different physical mechanisms, such as structural stiffness or surface friction. We therefore adopt an adaptive approximation strategy: we prompt Gemini 3 Pro [22] with the initial visual context I_t and task description $\mathcal{T}$ to classify the dominant power source (dissipative residual or structural residual) based on high-level semantics, and apply the corresponding reconstruction formulation detailed in §III-B.

Task Classification and Control Signal Generation. Given IF, where the x-axis denotes the intended motion twist direction and the z-axis denotes the intended interaction wrench direction, we classify interaction patches into four canonical task modes based on the *IF-aligned signal characteristics* of twist $\boldsymbol{\xi}$ and wrench $\boldsymbol{\mathcal{W}}$:

- **Free**: No sustained contact. The interaction wrench is negligible, *i.e.*, $\|\boldsymbol{\mathcal{W}}\| \approx 0$.
- **Surface**: Contact with a surface imposing a single dominant normal constraint, *i.e.*, $\|\boldsymbol{f}_z\| \gg \|\boldsymbol{f}_{xy}\|$.
- **Insertion**: Highly-constrained insertions with strong force anisotropy from frictions, *i.e.*, $\|\boldsymbol{f}_x\| \gg \|\boldsymbol{f}_{yz}\|$.
- **Rotation**: Rotation-dominated interaction like fastening screws, *i.e.*, $\|\boldsymbol{\omega}_z\| \gg 0$, and $\|\boldsymbol{f}_z\| \gg 0$.

This categorization maps directly to the hybrid control structure via a selection mask $S \in \{0,1\}^6$ [49], where $S_i = 1$ denotes force control and $S_i = 0$ denotes position control. By configuring S to align with the spectral axes of each mode, we enable the requisite motion while regulating contact forces. The specific control parameterization is detailed in Table I.

TABLE I: Control Structure Selection based on Task Modes. The IF x-axis aligns with the primary motion/insertion direction, while the z-axis aligns with the dominant normal.

Interaction	Task Mode	Selection Mask S	Reference Wrench
Free	Free	$[0,0,0,0,0,0]$	$[-,-,-,-,-,-]$
Contact	Surface	$[0,0,1,0,0,0]$	$[-,-,\boldsymbol{f}_z,-,-,-]$
	Insertion	$[0,1,1,0,1,1]$*	$[-,0,0,-,0,0]$
	Rotation	$[0,0,1,0,0,0]$*	$[-,-,\boldsymbol{f}_z,-,-,-]$

* In practice, for insertion tasks, both axial translation and rotation about the insertion x-axis are theoretically inside $\mathcal{T}$, whereas for screw (pure rotation) tasks, only the axial rotation around z-axis belongs to the $\mathcal{T}$. Nevertheless, force and torque control are often applied to these degrees of freedom to cope with high friction and prevent jamming.

B. *Force Policy*: A Global-Local Vision-Force Policy

Global-Local Vision-Force Decomposition. Contact-rich manipulation alternates between two regimes: vision-guided motion in free space and force-governed interaction at contact. Hence, we factor the policy accordingly. *Global motion* is handled by a vision policy that reasons over geometry and long-horizon sequencing; outside contact, force readings are largely uninformative and often dominated by sensor noise or incidental touches [29]. *Local interaction* is handled by a high-frequency force policy that reacts to end-effector proprioception and force/torque signals to regulate contact dynamics. *This decomposition assigns each modality to the scale where it is most informative: vision drives task progress, while force enables responsive, stable interaction control.*

Modeling Local Force Policy. Human interaction with the physical world offers a useful intuition for designing the local interaction policy. Consider the simple act of sliding a finger along the edge of a table. Once contact is established, precise visual perception of the geometry is unnecessary; vision mainly provides *global contextual cues* about where interaction occurs. The interaction itself is regulated through high-frequency proprioception and force feedback: motion sensed by the hand informs how movement evolves, while force feedback reveals contact, resistance, and slip. Therefore, we formulate the fast, local, force policy Π_{local} as:

$$\mathbf{a}_t^{\text{local}} = \Pi_{\text{local}}\Big(\Delta s_{t-T_o+1:t},\ \boldsymbol{\mathcal{W}}_{t-T_o+1:t},\ \phi(I_{t'})\Big), \tag{5}$$

where $\mathbf{a}_t^{\text{local}}$ is the predicted local action, $\Delta s_{t-T_o+1:t}$ and $\boldsymbol{\mathcal{W}}_{t-T_o+1:t}$ denotes the end-effector motion and the force feedback histories of length T_o, respectively. $\phi(I_{t'})$ is the global scene context updated by the global vision policy at $t' \le t$, decoupling the local and global policy inference frequencies.

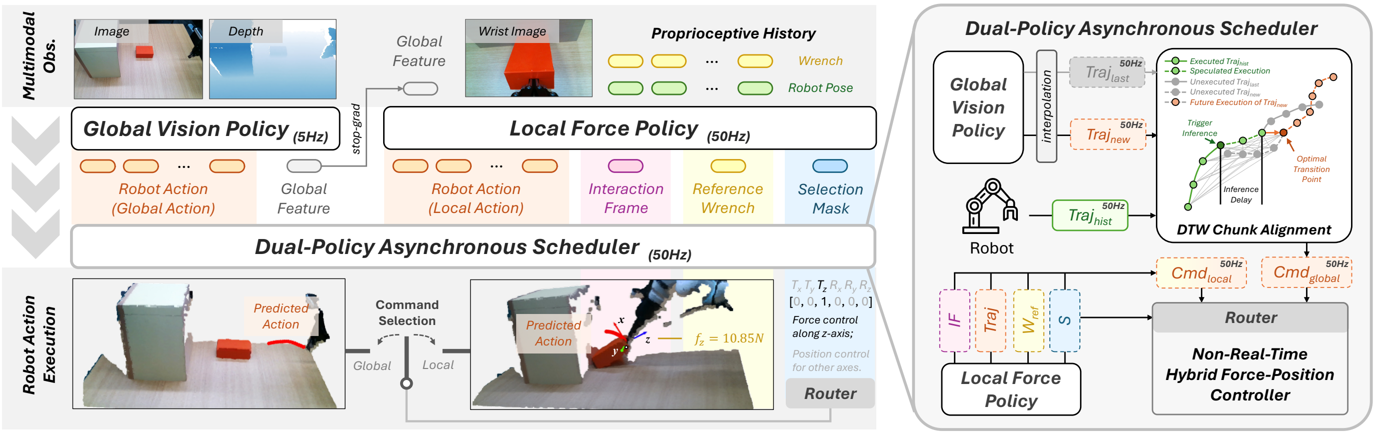

Fig. 3: ***Force Policy* and Dual-Policy Asynchronous Scheduler.** *(Left)* ***Force Policy*** consists of a global vision policy and a local force policy. The global policy provides task-level visual context and global actions, while the local policy predicts interaction structure and local actions to realize hybrid force-position control during contact. *(Right)* The dual-policy asynchronous scheduler switches between the two policies and reduces latency and jerk via model-agnostic chunk alignment using dynamic time warping (DTW) [51].

Each predicted local action explicitly parameterizes force-based interaction: it specifies how motion and force should be regulated in IF, rather than directly commanding joint torques or poses. Formally, we define

$$\mathbf{a}_t^{\text{local}} \triangleq \left(\Sigma_{t+1},\ S_{t+1},\ \mathcal{W}_{t+1}^{\text{ref}},\ \Delta s_{t+1:t+T_a}\right), \tag{6}$$

where S_{t+1} selects the controlled motion and force subspaces in IF Σ_{t+1}, $\Delta s_{t+1:t+T_a}$ specifies the desired motion chunk [83] of horizon T_a, and $\mathcal{W}_{t+1}^{\text{ref}}$ provides the reference wrench that directly governs force regulation during contact.

Modeling Global Vision Policy. The global vision policy Π_{global} governs task-level progression. In free space, it acts as a standard vision-based policy and directly outputs actions. During interaction, control is delegated to the local force policy Π_{local}, and the global vision policy no longer issues commands. Instead, it provides a global visual feature $\phi(I_{t'})$ to condition local interaction control. This modular design allows Π_{global} to be instantiated from existing visuomotor policies [15, 24, 74, 83] or vision-language-action (VLA) models [4, 6, 7] without modifications.

Router. The router determines which policy holds control authority between the global vision policy and the local force policy. Rather than introducing an explicit routing module, we *reuse the predicted selection mask S_{t+1} from the local force policy as an implicit routing signal*. The local force policy dynamically activates force-control axes as needed by assigning non-zero entries in S_{t+1}, thereby switching control from the global policy to engage hybrid force-position regulation. In contrast, unintended contacts do not trigger the selection mask, and control remains with the global vision policy.

***Force Policy*.** As shown in Fig. 3 (left), the overall policy is decomposed into a global vision policy Π_{global} and a local force policy Π_{local}. Π_{global} is responsible for global perception and high-level action generation, while Π_{local} handles local interaction through force-aware feedback at a high frequency. In our implementation, wrist-mounted camera observations are fed into Π_{local}, as they lie within the local interaction field. We instantiate Π_{global} with RISE-2 [24], a generalizable visuomotor policy that relies solely on global 3D visual perception. In its pipeline, we treat the action feature used as the conditions in the diffusion head as the global visual feature $\phi(I_{t'})$. For Π_{local}, the wrist image and proprioceptive history are encoded via a lightweight ResNet [28] and a GRU [18], respectively. Both are conditioned on $\phi(I_t)$ via FiLM [57]. The resulting features are adaptively fused through a gated mechanism and used to denoise the action sequence with an MIP head [56], as well as to predict the interaction frame, reference wrench, and selection mask via an MLP head.

C. Dual-Policy Asynchronous Scheduler

We introduce a dual-policy asynchronous scheduler during deployment to coordinate a global vision policy with a high-frequency local force policy, while preventing global inference latency from degrading closed-loop control. The scheduler (1) runs both policies asynchronously and uses a router with smooth state transitions to avoid discontinuities during policy switching; and (2) compensates for delayed global outputs by time-aligning and merging them into the execution stream in a latency-aware manner, enforcing trajectory smoothness so executed motions remain stable and physically consistent. This smoothing also reduces jerk-induced inertial force transients, improving force-sensing fidelity during manipulation. An overview is shown in Fig. 3 (right).

Asynchronous Inference and Execution. To mask inference latency, the scheduler adopts a multi-threaded asynchronous design that overlaps model inference with trajectory execution, preemptively launching inference at fixed intervals. Action selection follows the router in §IV-B, with hysteresis applied to suppress switching jitter. In implementation, outputs from both the global vision policy and the local force policy are first interpolated to a unified frequency of 50Hz. At this common rate, waypoints are selectively dropped to mitigate chunk mismatch induced by inference latency, as detailed in the following paragraph, followed by acceleration-continuous trajectory planning. This planning step ensures smooth execution and high-fidelity force control, as higher-order continuity

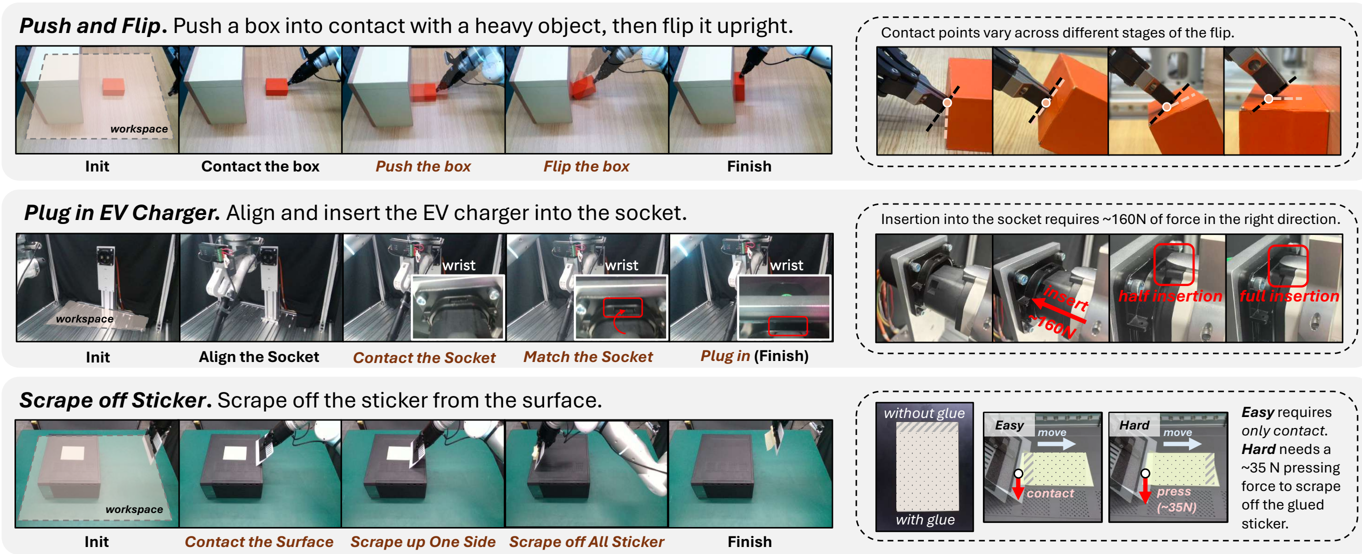

Fig. 4: Tasks. We design three tasks spanning two categories (polishing and insertion) to evaluate different policies for contact-rich manipulation. The descriptions on the right highlight the key challenges of each task compared to similar tasks in prior literature. All tasks require highly accurate force regulation to be successfully completed. We randomize object placement within the workspace area during both data collection and evaluation for each task.

suppresses jerk and reduces inertial disturbances. The resulting trajectories are then dispatched to a non-real-time hybrid force-position controller [70] for execution.

Chunk Alignment with Dynamic Time Warping. Latency in asynchronous inference systems typically causes discontinuous jumps between consecutive action chunks. Existing methods [2, 5, 8] address this issue via action inpainting or adaptive conditional guidance, but depend on specific model architectures and introduce additional computational overhead. In contrast, we propose a computationally efficient, model-agnostic waypoint dropout strategy based on Dynamic Time Warping (DTW) [51]. DTW enables nonlinear temporal alignment and has been widely adopted from speech recognition [33, 52, 61] to motion retargeting [46]. By aligning the predicted trajectory with the recent execution history, we can identify an optimal entry index and drop preceding waypoints, thereby eliminating latency-induced discontinuities without modifying the training procedure or model architecture. Additional implementation details are provided in the supplementary material.

The proposed scheduler enables latency-aware deployment of dual-policy control by decoupling asynchronous inference from execution while maintaining smooth, force-consistent motion. Unified-frequency resampling, DTW-based chunk alignment, and acceleration-continuous planning together ensure stable execution without constraining policy architecture.

V. Experiments

Through experiments, we seek to answer the following questions: **(Q1)** Can ***Force Policy*** effectively handle different types of contact-rich tasks? **(Q2)** Does the ***Force Policy*** achieve superior force control compared to existing force-aware baselines? **(Q3)** Does our global-local design generalize better than the baselines? **(Q4)** Does our IF recovery method produce more accurate IF than prior approaches, and does improved interaction labeling enable the policy to learn more accurate contact behaviors? **(Q5)** Does the asynchronous scheduler reduce motion and force jerks, leading to smoother trajectories during deployment?

A. Setup

Platform. The robot platform comprises a Flexiv Rizon 4 arm equipped with a Flexiv GN-02 gripper and a 6-DoF force-torque sensor at the flange. Two Intel RealSense D415 RGB-D cameras are used to capture visual observations, serving as a global camera and a wrist-mounted camera, respectively.

Tasks. As shown in Fig. 4, we design three tasks from two major applications: polishing and insertion, including scenarios with continuously varying IFs and tasks that demand reactive control and accurate force regulation.

Data Collection. We use the arm-to-arm teleoperation toolkit with force feedback [48] to collect 50 high-quality demonstrations for each task. The system provides high-fidelity force feedback during teleoperation, enabling precise force regulation across tasks and task phases. Inspection of all demonstrations shows that the recorded force signals consistently fall within the required range for each task stage.

Baselines. For vision-only baselines, we consider two state-of-the-art policies: RISE-2 [24] and $\pi_{0.5}$ [6]. For force-aware baselines, we include the representative RDP [79], FoAR [29], ForceVLA [80], and TA-VLA [82] in evaluations.

Metrics. We report success rate as the primary metric for each task. For the "flip" phase in ***Push and Flip*** and the "plug in" phase in ***Plug in EV Charger***, we also count partial completion as a half success (0.5) per trial.

Evaluation. The π-based policies run on a server with an NVIDIA A800 GPU due to memory limitations, while others run on a workstation with an RTX 3090. Following [16, 78], all policies are evaluated consistently: test positions are randomly

TABLE II: **Evaluation Success Rates.** We report success rates evaluated in an accumulative, stage-wise manner. ***Force Policy*** achieves the highest success rates compared to vision-based and force-aware baselines among all tasks.

Policy	*Push and Flip*		*Plug in EV Charger*			*Scrape off Sticker* (Easy)			*Scrape off Sticker* (Hard)		
	push	flip	contact	match	plug in	contact	one-side	full off	contact	one-side	full off
RISE-2 [24]	**100.0%**	42.5%	**100.0%**	**90.0%**	0.0%	**100.0%**	80.0%	80.0%	**100.0%**	20.0%	10.0%
$\pi_{0.5}$ [6]	80.0%	52.5%	**100.0%**	30.0%	0.0%	70.0%	70.0%	70.0%	**100.0%**	40.0%	20.0%
RDP [79]	**100.0%**	57.5%	**100.0%**	70.0%	5.0%	**100.0%**	90.0%	70.0%	**100.0%**	30.0%	20.0%
FoAR [29]	**100.0%**	60.0%	**100.0%**	80.0%	10.0%	**100.0%**	70.0%	40.0%	**100.0%**	60.0%	20.0%
ForceVLA [80]	50.0%	30.0%	20.0%	0.0%	0.0%	10.0%	0.0%	0.0%	10.0%	0.0%	0.0%
TA-VLA [82]	85.0%	62.5%	80.0%	10.0%	0.0%	50.0%	50.0%	50.0%	40.0%	20.0%	10.0%
Force Policy (ours)	**100.0%**	**95.0%**	**100.0%**	**90.0%**	**65.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**90.0%**	**90.0%**

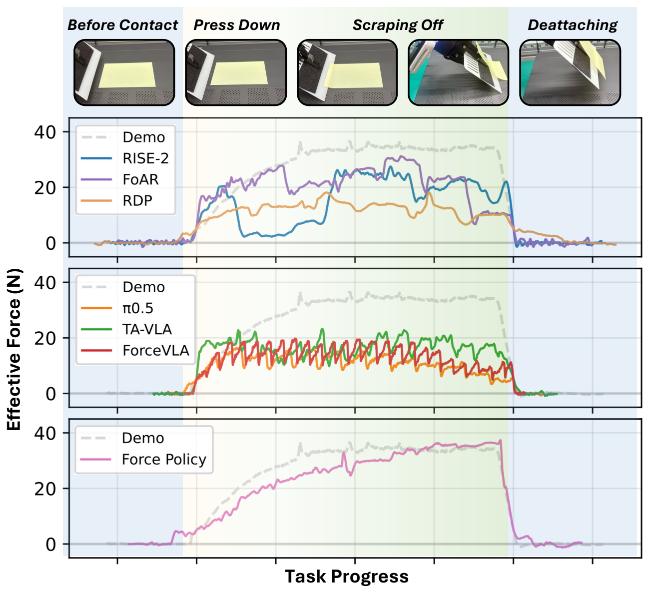

Fig. 5: Visualization of Effective Forces during Deployment and from Demonstrations on the *Scrape off Sticker* (Hard) Task. All baselines fail to replicate the force behavior from demonstrations, resulting in degraded performance. ***Force Policy*** closely imitates the effective force in demonstrations, achieving higher success rates.

generated beforehand, the workspace is identical, and metrics are counted over 20 trials for ***Push and Flip***, and 10 for others.

B. Results

Force Policy **exhibits significant effectiveness in handling a diverse range of contact-rich tasks (Q1).** As shown in Tab. II, ***Force Policy*** consistently outperforms both vision-only and force-aware baselines across a wide range of contact-rich tasks, with the largest gains appearing in phases that require precise and rapid force regulation, *e.g.*, during the insertion stage in the ***Plug in EV Charger*** task, and the scraping stage in the ***Scrape off Sticker*** task. The underlying insight is that *force feedback is genuinely useful, but only when integrated in a way that matches its intermittent, regime-dependent nature*.

Many prior designs misuse force in ways that reduce robustness. (1) Monolithic vision-force policies that simply concatenate force into observations, such as ForceVLA and TA-VLA, are prone to failure, because *force is noisy or uninformative during non-contact motion* and can pollute global decision-making [29], sometimes even underperforming their

TABLE III: **Force Regulation Evaluation on the *Push and Flip* Task.** *(Left)* Pushing the heavy object requires approximately 45N, while demonstrations apply about 15N pushing force to flip the target object; we then measure its pushed distance d as an indicator of force regulation. *(Right)* Statistics of the pushed distance for each method.

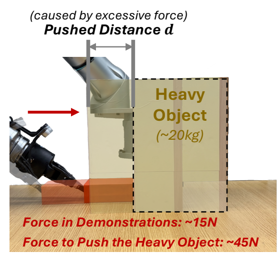

Policy	Avg. d (cm) ↓
RISE-2 [24]	0.86
$\pi_{0.5}$ [6]	**0.00**
RDP [79]	1.13
FoAR [29]	0.25
ForceVLA [80]	0.50
TA-VLA [82]	1.25
Force Policy (ours)	**0.00**

vision-only backbone. The effect is particularly pronounced in the evaluation results of ForceVLA. (2) Low-frequency force-aware pipelines such as FoAR *lack sufficient reactivity to handle fast contact transients*. (3) RDP employs a hierarchically-coupled policy design, in which low-level force-aware action decoding is directly conditioned on the high-level latent action chunk. When the high-level policy encodes a no-contact latent action but unexpected contact occurs during execution, the low-level tokenizer may map out-of-distribution contact force signals to erroneous actions. *This coupling limits the ability of low-level policy to correct wrong high-level latent action.*

In contrast to all baselines, our global-local vision-force design decouples high-level intent from low-level execution: the global vision module provides a stable long-horizon reference, *insulated from noisy force signals*, while the high-frequency local force policy *independently determines actions* based on this reference and real-time force feedback, enabling robust performance under sensor noise and fast contact transients.

Force Policy **achieves significantly superior force control compared to existing force-aware baselines (Q2).** As visualized in Fig. 5 and Tab. III, our method closely tracks the force profile of human demonstrations, maintaining the necessary effective force magnitude throughout the critical phases. On the contrary, baselines typically exhibit severe force oscillations, fail to exert sufficient downward force, or apply excessive force, resulting in task failure. This performance advantage stems from our adoption of a hybrid force-position control framework. Unlike baselines that regulate force implicitly or treat it as a sensory input, our approach explicitly controls interaction forces, enabling stable execution that faithfully

TABLE IV: Generalization Evaluation on the *Push and Flip* Task. We use unseen objects of different colors, geometries, and stiffnesses to evaluate the generalization ability of force-aware policies. Please refer to the supplementary material for detailed object comparisons.

Unseen Obj. / Policy						
RISE-2 [24]	3/5	0/5	0/5	3/5	0/5	0/5
$\pi_{0.5}$ [6]	2/5	2/5	1/5	3/5	**3/5**	1/5
RDP [79]	2/5	0/5	0/5	0/5	0/5	0/5
FoAR [29]	1/5	0/5	0/5	0/5	0/5	0/5
ForceVLA [80]	2/5	0/5	0/5	1/5	0/5	0/5
TA-VLA [82]	3/5	2/5	1/5	1/5	**3/5**	1/5
Force Policy (ours)	**5/5**	**5/5**	**4/5**	**4/5**	**3/5**	**2/5**

follows demonstrations, which is essential for contact-rich tasks requiring precise force control.

Force Policy* demonstrates decent generalization ability compared to baselines (Q3).** We aim to evaluate two aspects of generalization: visual generalization to *unseen objects*, and force-regulation generalization to *unseen geometries* during contact. The evaluation results of the policies on unseen objects with varying colors, geometries, and stiffnesses in the ***Push and Flip task are demonstrated in Tab. IV. Force Policy consistently achieves high success rates, whereas baselines frequently suffer catastrophic failure. Crucially, we must clarify that *although RISE-2, our global vision policy, often yields a 0% success rate on novel objects, it successfully navigates to the near-contact region in most trials*. Its failure mainly arises from relying on vision alone to handle the critical transition from free motion to contact, causing timing errors or improper contact establishment. *Our design successfully inherits this visual generalization capability from the global vision backbone to reliably guide the end-effector to the target vicinity*. Once in the near-contact region, the local force policy takes over, leveraging force feedback and hybrid force-position control to robustly establish contact and execute the actions. By decoupling visual reaching from physical interaction, our method effectively combines the broad generalization of visuomotor policies with the precision of local force control.

C. Ablations

Our IF recovery method achieves more accurate interaction frames from demonstrations than previous approaches, improving interaction quality (Q4). We evaluate several IF recovery approaches on the *Scrape off Sticker* task, including analytic [50], power-based [55], wrench-only [20, 45], twist-only, and our adaptive method. The z-axis of the ground-truth IF aligns with the world vertical axis in this task, so we measure angular error and count failures above 20°. The results are shown in Fig. 6. Analytic and power-based approaches exhibit larger errors, while twist-only and our method outperform wrench-only due to dissipative residuals like frictions dominating power in this task. Our adaptive method slightly improves over twist-only by capturing brief pressing periods before scraping, where structural residuals like scraper deformation dominate. We further evaluate the ***Force Policy*** trained with interaction frames labeled by wrench-only methods [45] on this task. It achieves only a 50% success rate, dropping from 90% with our IF recovery strategy.

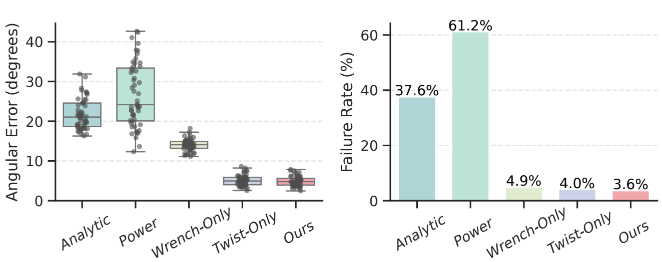

Fig. 6: IF Recovery Evaluation on the *Scrape off Sticker* Task. Angular error is computed between the recovered and ground-truth force control axes; a failure is counted if it exceeds 20°.

The semantic classification in our IF recovery method is robust in classifying different contact behaviors in the real world (Q4). The semantic classification accuracies for the three tasks are 100.0%, 92.0%, and 98.0%, respectively. We argue that ***Force Policy*** is not sensitive to these errors for the following two reasons: (1) Misclassified samples are infrequent, so their effect on policy learning is limited. (2) Even if an incorrect semantic prior produces a suboptimal IF and the policy learns to predict the suboptimal IF, the high-frequency local policy can still detect abnormal force signals and correct the IF online using the feedback.

The asynchronous scheduler substantially reduces jerks and yields smoother trajectories during deployment (Q5). We analyze both force and motion jerks and measure trajectory smoothness using spectral arc length (SPARC, ↑) [3] following [2]. As shown in Fig. 7, the temporal profiles indicate that the scheduler effectively attenuates signal discontinuities, resulting in smoother trajectories during the inference chunk alignment phase. Additional analysis of the asynchronous scheduler is provided in the supplementary material.

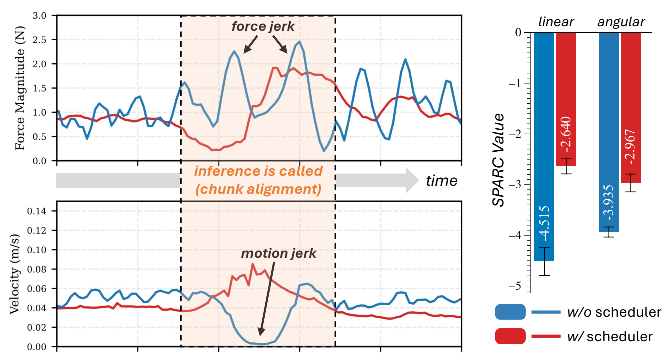

Fig. 7: Asynchronous Scheduler Evaluation on the *Push and Flip* Task. The scheduler effectively reduces both motion and force jerks *(left)*, resulting in more smoother executed trajectories *(right)*.

VI. Limitations and Future Works

IF Scope. We clarify that the present scope of our IF theory encompasses *stable-contact manipulation*, including *most frictional and compliant contact settings*. In these prevalent regimes, explicitly modeling structural residuals and continuous dissipative forces offers a distinct practical advantage over prior methods. Our formulation is grounded in *locally conservative* and *topologically invariant* assumptions. Consequently, highly non-smooth scenarios, such as fracture or non-conservative contact transitions, fall outside this scope and are left for future work. Multi-point and surface contact cases are discussed in the remark to Theorem 1 in the supplementary.

Evaluation Scope. Our current evaluation covers two broad classes of contact-rich tasks, *peg-in-hole* and *surface interaction*, spanning common contact-rich manipulation scenarios. In particular, ***Push and Flip*** involves a *continuously changing IF*, serving as a dynamic test of both IF recovery and IF prediction. Broader evaluations of complex contact-rich tasks is indeed valuable, and we leave it to future work.

Torque Control. Our current formulation focuses more on recovering interaction orientation for force control, rather than explicitly estimating the true contact point; fully modeling the contact structure would be necessary to support torque control.

Contact for High-Level Decision-Making. For some tasks, force/torque signals might help high-level decision-making, *e.g.*, briefly tugging to check if the connector is fully inserted; a promising direction is to extend the architecture to enable such force-driven judgments.

VII. Conclusion

Motivated by the biological division of labor where vision guides global decision-making and reaching while haptics governs local contact, we propose ***Force Policy***, a global-local vision-force framework that decouples the control hierarchy. The global vision policy decides "*where to act*" with strong generalization, while a high-frequency local controller decides "*how to act*" by regulating contact via hybrid force-position control. This design is enabled by a physically grounded interaction-frame formulation and a practical method to recover intrinsic frames from demonstrations. Empirically, we show (1) substantially improved robustness and precise force regulation on diverse contact-rich manipulation tasks by explicitly modeling contact structure and interaction forces, and (2) strong generalization across various object geometries and physical properties, driven by the complementary roles of the global vision policy that handles visual and geometric variation and the local force policy that handles contact dynamics. Overall, ***Force Policy*** bridges the generalization of visuomotor policies with the precision of classical force control, providing a scalable approach to contact-rich manipulation.

Acknowledgement

This work is supported by the New Generation Artificial Intelligence — National Science and Technology Major Project (No. 2025ZD0122901), the Science and Technology Major Project of Jiangsu Province (No. BG2024041), Shanghai Artificial Intelligence Laboratory, and XPLORER PRIZE grants.

We would like to thank Yiming Wang from Shanghai Jiao Tong University for insightful discussions on theoretical formulation; Zhipeng Zhang from Flexiv for his support on the arm-to-arm teleoperation system with force feedback; Wenbo Tang from Flexiv for discussions about the robot controller; Peishen Yan and Chunyu Xue from Shanghai Jiao Tong University for writing advice and proofreading; Lijia Yao from Noematrix for suggestions on figure design; Junchao Zhang from Noematrix for his help on evaluations; and Shangning Xia and Linhao Chen from Noematrix for discussions.

Author Contributions. *H. Fang led the project under the guidance of S. Wang and C. Lu. H. Fang, S. Tang, and M. Mei formulated the core idea of the project together. H. Fang, S. Tang, M. Mei, and S. Wang discussed and derived the theoretical framework of the interaction frame. H. Fang, S. Tang, and M. Mei designed the framework of Force Policy together. H. Fang and S. Tang implement the Force Policy and the asynchronous scheduler, respectively. C. Wang helped with the policy design and scheduler implementation. H. Fang, S. Tang and M. Mei designed the tasks together. M. Mei, Z. He, and H. Qin implemented the baselines. M. Mei, S. Tang, H. Qin, Z. He, J. Chen, Y. Feng, and H. Fang conducted the experiments. H. Fang, S. Tang, M. Mei, J. Chen wrote the paper. S. Wang, C. Lu, W. Liu and Z. He supervised the project and provided hardware support. This work was done while H. Fang was a research collaborator with Noematrix, and M. Mei was a research intern at Noematrix.*

References

[1] Malek Aburub et al. "Learning Diffusion Policies from Demonstrations for Compliant Contact-Rich Manipulation". In: *arXiv preprint arXiv:2410.19235* (2024).

[2] Nadun Ranawaka Arachchige et al. "SAIL: Faster-than-Demonstration Execution of Imitation Learning Policies". In: *Conference on Robot Learning*. Vol. 305. PMLR, 2025, pp. 721–749.

[3] Sivakumar Balasubramanian et al. "On the Analysis of Movement Smoothness". In: *Journal of Neuroengineering and Rehabilitation* 12.1 (2015), p. 112.

[4] Johan Bjorck et al. "GR00T N1: An Open Foundation Model for Generalist Humanoid Robots". In: *arXiv preprint arXiv:2503.14734* (2025).

[5] Kevin Black, Manuel Y Galliker, and Sergey Levine. "Real-Time Execution of Action Chunking Flow Policies". In: *Advances in Neural Information Processing Systems*. 2025.

[6] Kevin Black et al. "$\pi_{0.5}$: a Vision-Language-Action Model with Open-World Generalization". In: *Conference on Robot Learning*. Vol. 305. PMLR, 2025, pp. 17–40.

[7] Kevin Black et al. "π_0: A Vision-Language-Action Flow Model for General Robot Control". In: *Robotics: Science and Systems*. 2025.

[8] Kevin Black et al. "Training-Time Action Conditioning for Efficient Real-Time Chunking". In: *arXiv preprint arXiv:2512.05964* (2025).

[9] Herman Bruyninckx and Joris De Schutter. "Specification of Force-Controlled Actions in the "Task Frame Formalism" - a Synthesis". In: *IEEE Transactions on Robotics and Automation* 12.4 (1996), pp. 581–589.

[10] Herman Bruyninckx et al. "Kinematic Models for Model-Based Compliant Motion in the Presence of Uncertainty". In: *International Journal of Robotics Research* 14.5 (1995), pp. 465–482.

[11] Thanpimon Buamanee et al. "Bi-ACT: Bilateral Control-Based Imitation Learning via Action Chunking with Transformer". In: *IEEE International Conference on Advanced Intelligent Mechatronics*. IEEE. 2024, pp. 410–415.

[12] Jonas Buchli et al. "Learning Variable Impedance Control". In: *International Journal of Robotics Research* 30.7 (2011), pp. 820–833.

[13] Ivan Camponogara and Robert Volcic. "Integration of Haptics and Vision in Human Multisensory Grasping". In: *Cortex* 135 (2021), pp. 173–185.

[14] Tao Chen et al. "Vegetable Peeling: A Case Study in Constrained Dexterous Manipulation". In: *arXiv preprint arXiv:2407.07884* (2024).

[15] Cheng Chi et al. "Diffusion Policy: Visuomotor Policy Learning via Action Diffusion". In: *Robotics: Science and Systems*. 2023.

[16] Cheng Chi et al. "Universal Manipulation Interface: In-The-Wild Robot Teaching Without In-The-Wild Robots". In: *Robotics: Science and Systems*. 2024.

[17] Stefano Chiaverini and Lorenzo Sciavicco. "The Parallel Approach to Force/Position Control of Robotic Manipulators". In: *IEEE Transactions on Robotics and Automation* 9.4 (2002), pp. 361–373.

[18] Kyunghyun Cho et al. "Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation". In: *Empirical Methods in Natural Language Processing*. ACL, 2014, pp. 1724–1734.

[19] Naomi Chukwurah, Abiodun Sunday Adebayo, and Olanrewaju Oluwaseun Ajayi. "Sim-to-Real Transfer in Robotics: Addressing the Gap Between Simulation and Real-World Performance". In: *International Journal of Robotics and Simulation* 6.1 (2024), pp. 89–102.

[20] Adam Conkey and Tucker Hermans. "Learning Task Constraints from Demonstration for Hybrid Force/Position Control". In: *IEEE-RAS International Conference on Humanoid Robots*. IEEE. 2019, pp. 162–169.

[21] Joris De Schutter and Hendrik Van Brussel. "Compliant Robot Motion I. A Formalism for Specifying Compliant Motion Tasks". In: *International Journal of Robotics Research* 7.4 (1988), pp. 3–17.

[22] Google DeepMind. *Gemini 3 Pro*. Accessed January 2026.

[23] Hao-Shu Fang et al. "DEXOP: A Device for Robotic Transfer of Dexterous Human Manipulation". In: *arXiv preprint arXiv:2509.04441* (2025).

[24] Hongjie Fang et al. "AirExo-2: Scaling up Generalizable Robotic Imitation Learning with Low-Cost Exoskeletons". In: *Conference on Robot Learning*. Vol. 305. PMLR, 2025, pp. 198–220.

[25] Ruoxuan Feng et al. "Play to the Score: Stage-Guided Dynamic Multi-Sensory Fusion for Robotic Manipulation". In: *Conference on Robot Learning*. 2024.

[26] Omar A Gharbawie, Iwona Stepniewska, and Jon H Kaas. "Cortical Connections of Functional Zones in Posterior Parietal Cortex and Frontal Cortex Motor Regions in New World Monkeys". In: *Cerebral Cortex* 21.9 (2011), pp. 1981–2002.

[27] Adam B Goldring et al. "Functional Characterization of the Fronto-Parietal Reaching and Grasping Network: Reversible Deactivation of M1 and Areas 2, 5, and 7b in Awake Behaving Monkeys". In: *Journal of Neurophysiology* 127.5 (2022), pp. 1363–1387.

[28] Kaiming He et al. "Deep Residual Learning for Image Recognition". In: *IEEE Conference on Computer Vision and Pattern Recognition*. 2016, pp. 770–778.

[29] Zihao He et al. "FoAR: Force-Aware Reactive Policy for Contact-Rich Robotic Manipulation". In: *IEEE Robotics and Automation Letters* (2025).

[30] Neville Hogan. "Impedance Control: An Approach to Manipulation". In: *Journal of Dynamic Systems, Measurement, and Control* 107 (1985), pp. 1–24.

[31] Yifan Hou et al. "Adaptive Compliance Policy: Learning Approximate Compliance for Diffusion Guided Control". In: *IEEE International Conference on Robotics and Automation*. 2025, pp. 4829–4836.

[32] Binghao Huang et al. "3D-ViTac: Learning Fine-Grained Manipulation with Visuo-Tactile Sensing". In: *arXiv preprint arXiv:2410.24091* (2024).

[33] Abderrahmane Issam et al. "DTW-Align: Bridging the Modality Gap in End-to-End Speech Translation with Dynamic Time Warping Alignment". In: *Proceedings of the Tenth Conference on Machine Translation*. 2025, pp. 191–199.

[34] Zhenzhong Jia et al. "A Survey of Automated Threaded Fastening". In: *IEEE Transactions on Automation Science and Engineering* 16.1 (2018), pp. 298–310.

[35] Mrinal Kalakrishnan et al. "Learning Force Control Policies for Compliant Manipulation". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. 2011, pp. 4639–4644.

[36] 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: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE. 2024, pp. 12663–12670.

[37] Arvid QL Keemink, Herman Van der Kooij, and Arno HA Stienen. "Admittance control for physical human–robot interaction". In: *The International Journal of Robotics Research* 37.11 (2018), pp. 1421–1444.

[38] 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 (2003), pp. 43–53.

[39] Jens Kober, Michael Gienger, and Jochen J Steil. "Learning Movement Primitives for Force Interaction Tasks". In: *IEEE International Conference on Robotics and Automation*. IEEE. 2015, pp. 3192–3199.

[40] Klas Kronander and Aude Billard. "Stability Considerations for Variable Impedance Control". In: *IEEE Transactions on Robotics* 32.5 (2016), pp. 1298–1305.

[41] Sergey Levine, Nolan Wagener, and Pieter Abbeel. "Learning Contact-Rich Manipulation Skills With Guided Policy Search". In: *IEEE International Conference on Robotics and Automation*. 2015, pp. 156–163.

[42] Hao Li et al. "See, Hear, and Feel: Smart Sensory Fusion for Robotic Manipulation". In: *Conference on Robot Learning*. 2022, pp. 1368–1378.

[43] Xing Li, Manuel Baum, and Oliver Brock. "Augmentation Enables One-Shot Generalization in Learning from Demonstration for Contact-Rich Manipulation". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE. 2023, pp. 3656–3663.

[44] Xing Li and Oliver Brock. "Learning from Demonstration Based on Environmental Constraints". In: *IEEE Robotics and Automation Letters* 7.4 (2022), pp. 10938–10945.

[45] Wenhai Liu et al. "ForceMimic: Force-Centric Imitation Learning With Force-Motion Capture System for Contact-Rich Manipulation". In: *ICRA*. IEEE. 2025, pp. 1105–1112.

[46] Yangcen Liu et al. "ImMimic: Cross-Domain Imitation from Human Videos via Mapping and Interpolation". In: *arXiv preprint arXiv:2509.10952* (2025).

[47] Zeyi Liu et al. "ManiWAV: Learning Robot Manipulation from In-the-Wild Audio-Visual Data". In: *Conference on Robot Learning*. 2024.

[48] Flexiv Ltd. *Flexiv Teleoperation Development Kit (TDK)*. Accessed January 2026.

[49] Matthew T Mason. "Compliance and Force Control for Computer Controlled Manipulators". In: *IEEE Transactions on Systems, Man, and Cybernetics* 11.6 (2007), pp. 418–432.

[50] Ali Mousavi Mohammadi et al. "A Generic Task Model and Control Strategy to Support Learning, Robust Control, and Generalization of Contact-Rich Manipulation Tasks". In: *Robotics and Autonomous Systems* 197 (2026), p. 105270.

[51] Meinard Müller. "Dynamic Time Warping". In: *Information Retrieval for Music and Motion*. Springer, 2007, pp. 69–84.

[52] C. Myers, L. Rabiner, and A. Rosenberg. "Performance Tradeoffs in Dynamic Time Warping Algorithms for Isolated Word Recognition". In: *IEEE Transactions on Acoustics, Speech, and Signal Processing* 28.6 (1980), pp. 623–635.

[53] Michael Noseworthy et al. "FORGE: Force-Guided Exploration for Robust Contact-Rich Manipulation Under Uncertainty". In: *IEEE Robotics and Automation Letters* (2025).

[54] Alex H. G. Overbeek et al. "Versatile Kinematics-Based Constraint Identification Applied to Robot Task Reproduction". In: *Frontiers in Robotics and AI* 12 (2025).

[55] Alex Harm Gert-Jan Overbeek et al. "Identifying Physical Interactions in Contact-Based Robot Manipulation for Learning from Demonstration". In: *Advanced Robotics Research* (2025), e202500109.

[56] Chaoyi Pan et al. "Much Ado About Noising: Dispelling the Myths of Generative Robotic Control". In: *arXiv preprint arXiv:2512.01809* (2025).

[57] Ethan Perez et al. "FiLM: Visual Reasoning with a General Conditioning Layer". In: *AAAI conference on Artificial Intelligence*. 2018, pp. 3942–3951.

[58] Claudia Pérez-D'Arpino and Julie A Shah. "C-Learn: Learning Geometric Constraints from Demonstrations for Multi-Step Manipulation in Shared Autonomy". In: *IEEE International Conference on Robotics and Automation*. IEEE. 2017, pp. 4058–4065.

[59] Marc H Raibert and John J Craig. "Hybrid Position/Force Control of Manipulators". In: *Journal of dynamic systems, measurement, and control* 103.2 (1981), pp. 126–133.

[60] Branden Romero et al. "Eyesight Hand: Design of a Fully-Actuated Dexterous Robot Hand with Integrated Vision-based Tactile Sensors and Compliant Actuation". In: *arXiv preprint arXiv:2408.06265* (2024).

[61] H. Sakoe and S. Chiba. "Dynamic Programming Algorithm Optimization for Spoken Word Recognition". In: *IEEE Transactions on Acoustics, Speech, and Signal Processing* 26.1 (1978), pp. 43–49.

[62] Joohwan Seo et al. "EquiContact: A Hierarchical SE(3) Vision-To-Force Equivariant Policy for Spatially Generalizable Contact-Rich Tasks". In: *arXiv preprint arXiv:2507.10961* (2025).

[63] Homayoun Seraji. "Adaptive Admittance Control: An Approach to Explicit Force Control in Compliant Motion". In: *IEEE International Conference on Robotics and Automation*. IEEE. 1994, pp. 2705–2712.

[64] Guru Subramani, Michael Zinn, and Michael Gleicher. "Inferring Geometric Constraints in Human Demonstrations". In: *Conference on Robot Learning*. PMLR. 2018, pp. 223–236.

[65] Markku Suomalainen, Fares J Abu-Dakka, and Ville Kyrki. "Imitation Learning-Based Framework for Learning 6-D Linear Compliant Motions". In: *Autonomous Robots* 45.3 (2021), pp. 389–405.

[66] Markku Suomalainen, Yiannis Karayiannidis, and Ville Kyrki. "A Survey of Robot Manipulation in Contact". In: *Robotics and Autonomous Systems* 156 (2022), p. 104224.

[67] Markku Suomalainen and Ville Kyrki. "Learning Compliant Assembly Motions from Demonstration". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE. 2016, pp. 871–876.

[68] Markku Suomalainen and Ville Kyrki. "A Geometric Approach for Learning Compliant Motions from Demonstration". In: *IEEE-RAS International Conference on Humanoid Robotics*. IEEE. 2017, pp. 783–790.

[69] Markku Suomalainen et al. "Learning from Demonstration for Hydraulic Manipulators". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE. 2018, pp. 3579–3586.

[70] Wenbo Tang, Weiming Wang, and Shiquan Wang. "Partially Decoupled Impedance Motion Force Control Using Prioritized Inertia Shaping". In: *IEEE Robotics and Automation Letters* (2024).

[71] Ana Lucia Pais Ureche et al. "Task Parameterization Using Continuous Constraints Extracted from Human Demonstrations". In: *IEEE Transactions on Robotics* 31.6 (2015), pp. 1458–1471.

[72] Madhusudhan Venkadesan and Francisco J Valero-Cuevas. "Neural Control of Motion-to-Force Transitions With the Fingertip". In: *Journal of Neuroscience* 28.6 (2008), pp. 1366–1373.

[73] Luigi Villani and Joris De Schutter. "Force Control". In: *Springer Handbook of Robotics*. Springer, 2016, pp. 195–220.

[74] Chenxi Wang et al. "RISE: 3D Perception Makes Real-World Robot Imitation Simple and Effective". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. 2024, pp. 2870–2877.

[75] Renhao Wang et al. "The Sound of Simulation: Learning Multimodal Sim-to-Real Robot Policies with Generative Audio". In: *Conference on Robot Learning*. Vol. 305. PMLR, 2025, pp. 420–436.

[76] Daniel E Whitney. "Historical perspective and state of the art in robot force control". In: *The International Journal of Robotics Research* 6.1 (1987), pp. 3–14.

[77] Yansong Wu et al. "TacDiffusion: Force-Domain Diffusion Policy for Precise Tactile Manipulation". In: *IEEE International Conference on Robotics and Automation*. IEEE. 2025, pp. 11831–11837.

[78] Shangning Xia et al. "CAGE: Causal Attention Enables Data-Efficient Generalizable Robotic Manipulation". In: *IEEE International Conference on Robotics and Automation*. IEEE. 2025, pp. 13242–13249.

[79] Han Xue et al. "Reactive Diffusion Policy: Slow-Fast Visual-Tactile Policy Learning for Contact-Rich Manipulation". In: *Robotics: Science and Systems*. 2025.

[80] Jiawen Yu et al. "ForceVLA: Enhancing VLA Models with a Force-Aware MoE for Contact-Rich Manipulation". In: *Advances in Neural Information Processing Systems*. 2025.

[81] Kelin Yu et al. "MimicTouch: Leveraging Multi-Modal Human Tactile Demonstrations for Contact-Rich Manipulation". In: *Conference on Robot Learning*. 2024.

[82] Zongzheng Zhang et al. "Elucidating the Design Space of Torque-aware Vision-Language-Action Models". In: *Conference on Robot Learning*. Vol. 305. PMLR, 2025, pp. 4019–4037.

[83] Tony Z Zhao et al. "Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware". In: *Robotics: Science and Systems*. 2023.

[84] Bo Zhou et al. "Admittance Visuomotor Policy Learning for General-Purpose Contact-Rich Manipulations". In: *IEEE Transactions on Industrial Electronics* (2025).