

Minimalist Compliance Control

Haochen Shi* Songbo Hu* Yifan Hou Weizhuo Wang C. Karen Liu† Shuran Song†

*Equal contribution †Equal advising

Stanford University

<https://minimalist-compliance-control.github.io>




Figure 1: **Minimalist Compliance Control** (A) requires no force sensors or learning, estimating external wrenches $\hat{f}_{ext}$ directly from motor current or voltage signals using a motor torque model and Jacobians. These estimates drive a spring–mass–damper model to update task-space position references. (B) This minimalist approach generalizes across embodiment, from robot arm and dexterous hand to humanoid robot, and (C) remains plug-and-play with any policy such as VLM-based policy, imitation policy, and model-based policy, and across diverse tasks such as wiping, drawing, scooping, and in-hand manipulation.

Abstract—Compliance control is essential for safe physical interaction, yet its adoption is limited by hardware requirements such as force/torque sensors. While recent reinforcement learning approaches aim to bypass these constraints, they often suffer from sim-to-real gaps, lack safety guarantees, and add system complexity. We propose Minimalist Compliance Control, which enables compliant behavior using only motor current or voltage signals readily available in modern servos and quasi-direct-drive motors—without force sensors, current control, or learning. External wrenches are estimated from actuator signals and Jacobians and incorporated into a task-space admittance controller, preserving sufficient force measurement accuracy for stable and responsive compliance control. Our method is embodiment-agnostic and plug-and-play with diverse high-level planners. We validate our approach on a robot arm, a dexterous hand, and two humanoid robots across multiple contact-rich tasks, using vision-language models, imitation learning, and model-based planning. The results demonstrate robust, safe, and compliant interaction across embodiments and planning paradigms.

I. INTRODUCTION

Widespread adoption of compliance control is currently hindered by significant hardware barriers. Standard admittance and impedance control typically rely on expensive force/torque sensors or actuators with precise force feedback

capabilities that are not available on many widely used robot platforms [11, 32, 26, 31, 50, 21, 44, 40, 51, 9, 17, 16, 54, 15]. To address this, recent work has proposed leveraging Reinforcement Learning (RL) to learn compliance control policies [41, 52, 43, 29, 5, 19, 53, 30]. Typically, these approaches learn a position controller to mimic compliance behavior through reward shaping during training. However, there are two major limitations of RL-based compliance: First, these methods suffer from the sim-to-real gap when deployed in the real world, such as discrepancies in position tracking and stiffness accuracy. Consequently, the learned controllers often lack safety guarantees and can produce unexpected, large force spikes during physical interaction [52, 43, 29]. Moreover, RL-based pipelines increase the framework’s complexity and can be difficult to tune [30, 41].

To overcome hardware barriers while avoiding unnecessary complexity, a key observation is that motor current or PWM (pulse-width modulation, which is essentially voltage) inherently contains information for estimating external wrenches. In fact, early work in the 1990s leveraged motor current for torque estimation, but required carefully designed observer-based filtering to mitigate substantial measurement noise and

recover physically meaningful external forces [33, 12]. In contrast, modern actuator designs feature significantly improved current sensing and high-bandwidth current control loops, and have demonstrated that torque estimated from motor current closely matches ground-truth measurements [34, 24, 46].

Another key observation is that inaccurate stiffness realization and imprecise external wrench estimation affect compliance in different ways. Compliance control does not require highly accurate force magnitudes; instead, it primarily depends on correct force sign/direction and reasonable accuracy in the relevant frequency band to ensure stable and timely responses to disturbances. When these properties are preserved, approximate wrench estimates are sufficient to induce compliant behavior, and improved magnitude accuracy mainly enhances performance rather than safety [52]. In contrast, RL-based approaches do not explicitly guarantee either force sign consistency or frequency-domain behavior.

Therefore, we propose Minimalist Compliance Control for robots equipped with modern servos or quasi-direct-drive (QDD) motors (Fig. 1). Our method (1) calibrates the motor characteristics (e.g., torque constant), (2) estimates external wrenches from motor current or PWM, and (3) applies task-space admittance control using the estimated wrench. This approach provides three advantages that together democratize compliant behavior for a broad class of robotic systems:

- **No Force Sensors:** By estimating external forces and torques solely via motor current or PWM signals ubiquitous in modern actuators, we eliminate the dependency on direct force/torque/tactile sensing hardware, drastically reducing system cost and maintenance requirements. Notably, our approach does not even require current sensors or current-controlled motors; raw PWM signals alone are sufficient. Despite its simplicity, this model-based estimation preserves sufficient force measurement accuracy for stable and responsive compliance control. In contrast, RL-based methods provide no explicit guarantees on these, increasing the risk of unsafe behavior during real-world contact. Meanwhile, our approach retains additional advantages, including reactive behavior and the ability to respond to disturbances beyond end-effector wrenches.
- **Embodiment-Agnostic:** Grounded in task-space admittance control, our framework naturally generalizes across diverse morphologies, from fixed-base robot arms to dexterous hands and floating-base humanoids.
- **Plug-and-Play with Any Policy:** Our model-based approach remains compatible with a wide variety of high-level planning strategies and requires minimal tuning, ranging from vision-language models (VLMs), imitation learning, to model-based motion planning.

In summary, we present a Minimalist Compliance Control framework that enables compliant interaction without force sensors or learning, extending this capability to a wide range of robotic systems equipped with QDD motors and even high gear-ratio servos typically considered unsuitable for sensorless estimation. We validate our method through comprehensive

experiments on four distinct platforms—a robot arm, a dexterous hand, and two humanoid robots—across contact-rich tasks such as wiping, drawing, scooping, and in-hand manipulation. Moreover, our method can be rapidly deployed across new platforms without complex calibration—a key advantage of the “minimalist” approach. For example, we deploy our method on the Unitree G1 within 5 hours, without motor sysID—using manufacturer-provided torque constant.

II. RELATED WORKS

A. Robots and Force Torque Sensors

Many widely used robot platforms lack end-effector force/torque sensors or joint torque sensors. This is a limitation that spans diverse embodiments, including robot arms (e.g., ARX, ViperX, and ALOHA [1]), dexterous hands (e.g., LEAP [35], ORCA [10], and DexHand [47]), and humanoid robots (e.g., Unitree G1, Berkeley Humanoid Lite [7], and ToddlerBot [36, 22, 45]). While installing external 6-axis force/torque sensors offers a direct solution, industry-standard options like the ATI Mini-45 are often bulky, limited in robustness, and prohibitively expensive [37, 3, 8]. While alternative tactile technologies like optical [48], resistive [23], and magnetic [2] sensors exist, they are often limited by durability, signal drift, or susceptibility to interference. Moreover, they generally require complex calibration and primarily sense normal forces and rough estimates of shear forces. Another key observation is that these robotic systems are all equipped with modern servo and QDD motors, which motivates our approach to estimate external wrenches from motor current or PWM signals that are readily available on these platforms.

B. Model-based Compliance Control

Classical model-based strategies, such as task-space admittance and impedance control [11, 32, 26], provide principled frameworks for regulating contact forces via virtual spring-mass-damper dynamics. However, these approaches typically require either 6-axis force/torque sensors or accurate joint torque sensing and control. To relax these hardware requirements, prior work has explored sensorless force estimation using joint-level disturbance observers that separate externally induced torques from internally modeled dynamics, enabling force or motion control without dedicated force sensors [33, 12, 38, 25]. These methods, however, rely on accurate full-body dynamics models or specialized joint instrumentation (e.g., strain-gauge torque sensors), which are unavailable on the low-cost servo and QDD platforms we target. More recently, improved current sensing, high-bandwidth current control, and proper calibration have enabled accurate torque estimation directly from motor current, primarily in systems with low transmission distortion such as direct-drive and QDD motors [34, 24, 46, 14, 27]. Our work extends beyond low-transmission actuators and shows that servo motors with high-reduction transmissions (gear ratios exceeding 200:1), even without current/torque control, can be reliably used for torque estimation and subsequent admittance control.

C. RL-based Compliance Control

As another approach to address hardware limitations, recent work has explored reinforcement learning (RL) for compliant control, aiming to implicitly infer external wrenches from proprioceptive signals through a neural network [30, 43, 52]. These methods generally fall into two categories: task-specific policies, which train on augmented datasets around specific motion clips to absorb disturbances [30], and task-agnostic policies, which aim to generalize across tasks by mimicking spring-mass-damper dynamics [52, 43, 29]. However, the black-box nature of these policies often leads to sim-to-real discrepancies, resulting in a lack of safety guarantees and susceptibility to dangerous force spikes. In this work, we achieve reliable compliance control using explicit wrench estimation and classical model-based control, without sim-to-real transfer or black-box policies.

III. METHOD

We present our low-level compliance controller in Section III-A, III-B, III-C, III-D and high-level planning strategies in Section III-F, III-G, III-H. Each planning strategy can be used independently with our compliance controller. Fig. 2 summarizes the inputs and outputs of each policy.

A. Spring-Mass-Damper Model

In our compliance controller, each end-effector is modeled as a spring-mass-damper system that responds to commanded motion and external wrenches. The dynamics are governed by:

$$m\ddot{\mathbf{x}} = \mathbf{K}_p(\mathbf{x}_{\text{des}} - \mathbf{x}) + \mathbf{K}_d(\dot{\mathbf{x}}_{\text{des}} - \dot{\mathbf{x}}) + \mathbf{f}_{\text{cmd}} + \mathbf{f}_{\text{ext}}, \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^6$ denotes the end-effector pose with orientation represented as a rotation vector, m is the effective mass (set to 1 for simplicity), $\mathbf{K}_p$ and $\mathbf{K}_d$ are the stiffness and damping matrices, $\mathbf{x}_{\text{des}}$ and $\dot{\mathbf{x}}_{\text{des}}$ are the desired pose and velocity, and $\mathbf{f}_{\text{cmd}}$ and $\mathbf{f}_{\text{ext}}$ are the commanded and external wrenches.

B. Motor Torque Estimation

Motor torque sensing is not directly available in most servo and QDD motors. Therefore, we estimate motor torques from motor current or pulse-width modulation (PWM) signals that can be used to infer motor current. We sysID the motors following standard motor models described in [13, 28].

For motors with current sensors, we directly read the motor current I_w from the sensor. For motors without current sensors, we estimate the motor current using the following relations:

$$V_{\text{PWM}} = \text{PWM} \times V_{\text{bus}}, \quad (2)$$

$$V_{\text{emf}} = \frac{\dot{q}}{K_v}, \quad (3)$$

$$I_w = \frac{V_{\text{PWM}} - V_{\text{emf}}}{R_w}, \quad (4)$$

where I_w is the motor winding current, V_{PWM} is the PWM voltage, V_{emf} is the back EMF, R_w is the winding resistance, $\text{PWM} \in [0, 1]$ is the duty cycle, V_{bus} is a constant bus voltage, $\dot{q}$ is the motor velocity, and K_v is the velocity constant.

With Eq. (3), we calibrate K_v by sweeping the motor through different velocities under no-load conditions where $V_{\text{PWM}} \approx V_{\text{emf}}$. We estimate R_w using an external current sensor to measure I_w under varying PWM duty cycles, then solve for R_w with Eq. (4). The motor torque constant K_t is calibrated from the current-torque relationship:

$$K_t = \frac{\tau_{\text{load}}}{\eta I_w}, \quad (5)$$

where τ_{load} is the measured load torque, $\eta \in (0, 1]$ is the motor efficiency accounting for power conversion losses, and I_w is the motor winding current. We use manufacturer-provided torque constants K_t when available; otherwise, we calibrate K_t under different load and velocity conditions. This means that for motors equipped with current sensors and manufacturer-provided K_t , no additional sysID is required. Concretely, for QDD motors used in the ARX X5 arm and Unitree G1, we use manufacturer-provided K_t . For Dynamixel servos used in ToddlerBot and the LEAP Hand, which lack current sensing, we calibrate K_t ourselves using an open-source testbed [36].

To account for asymmetric power transmission in forward and backward drive, we apply direction-dependent gains:

$$d = \begin{cases} \text{sign}(\tau_w \dot{q}), & |\dot{q}| > \epsilon_{\text{vel}}, \\ d_{\text{prev}}, & |\dot{q}| \leq \epsilon_{\text{vel}}, \end{cases} \quad (6)$$

$$\tau_{\text{load}} = \begin{cases} \eta K_t I_w, & d > 0 \text{ (forward drive)}, \\ \eta^{-1} K_t I_w, & d \leq 0 \text{ (backward drive)}, \end{cases} \quad (7)$$

where d is the forward/backward drive state determined from the sign of power flow $\tau_w \dot{q}$, ϵ_{vel} is a velocity threshold to debounce, and τ_{load} is the output torque at the joint. d_{prev} is initialized to be 1. In forward drive ($d > 0$), the motor accelerates the load, and efficiency η accounts for power losses. In backward drive ($d \leq 0$), external forces back-drive the motor, and the inverse efficiency η^{-1} accounts for brake amplification through the transmission. Calibration is performed using a low-cost, open-source, 3D-printed motor testbed [36] and requires no specialized equipment. In practice, sysID takes less than 20 minutes per motor. All code is open-sourced. SysID takes under 20 minutes per motor model; the identified parameters (K_v, R_w, K_t, η) are shared across all joints of the same model, avoiding per-joint calibration. Manufacturing variability and simple testbeds can introduce 10–20% parameter errors, but our method remains robust. The controller depends on force direction and relative magnitude, not absolute values, making it resilient to such errors.

To isolate external torques, we compensate for the gravity by subtracting τ_{grav} from the measured torque:

$$\tau_{\text{ext}} = -(r \cdot \tau_{\text{load}} - \tau_{\text{grav}}), \quad (8)$$

where r is the gear ratio and τ_{grav} is the gravity-induced torque. We assume a quasi-static interaction regime, in which inertial and velocity-dependent terms, such as joint acceleration and Coriolis effects, are negligible compared to gravity and externally applied torques. This is a deliberate design choice for our minimalist goal: incorporating these terms




Figure 2: **Policy Inputs and Outputs.** We show that Minimalist Compliance Control is plug-and-play with a range of high-level policies that benefit from compliant interaction. We illustrate the inputs and outputs of (A) a VLM-based policy, (B) an imitation policy, and (C) a model-based policy. The math symbols refer to those in Eq. (1). (A) and (C) predict only the directions of the stiffness $\mathbf{K}_p$ and force command $\mathbf{f}_{\text{cmd}}$, while their magnitudes remain fixed. For all three, the damping matrix is set for critical damping as $\mathbf{K}_d = 2\mathbf{K}_p^{1/2}$, assuming an identity inertia matrix, where $(\cdot)^{1/2}$ denotes the matrix square root. We also visualize the VLM pipeline to predict 3D contact points and normals, which are interpolated to generate $\mathbf{x}_{\text{des}}$.

would require accurate full-body inertia models and joint-acceleration estimates, increasing the calibration burden and the risk of instability. Under this assumption, τ_{ext} provides a reliable estimate of the externally induced torque. Empirically, this assumption holds for most tasks that require compliance, where interactions are dominated by low-frequency, quasi-static contacts rather than high dynamic motions.

C. External Wrench Estimation

The external wrench is estimated from motor torques using a flexible formulation that supports either full-wrench recovery or selective, axis-aligned estimation. When full-wrench estimation is enabled, we recover the contact force and optionally the torque by solving a regularized least-squares problem:

$$\hat{\mathbf{f}}_{\text{ext}}^p = \arg \min_{\mathbf{f}} \|\mathbf{J}_p^T \mathbf{f} - \tau_{\text{ext}}\|^2 + \lambda \|\mathbf{f}\|^2, \quad (9)$$

where $\mathbf{f}_{\text{ext}}^p \in \mathbb{R}^3$ is the translational component of $\mathbf{f}_{\text{ext}}$, $\mathbf{J}_p \in \mathbb{R}^{3 \times n}$ is the translational Jacobian of the contact site, and λ is a regularization coefficient. Torque components can be recovered analogously using the rotational Jacobian $\mathbf{J}_r$.

For improved numerical conditioning and noise robustness, we optionally estimate only the physically relevant components—namely, the force along the contact normal. For a selected unit axis $\hat{\mathbf{u}} \in \mathbb{R}^3$, we estimate the corresponding axis-aligned force as

$$\hat{\mathbf{f}}_{\text{ext}}^p = \frac{(\hat{\mathbf{u}}^T \mathbf{J}_p) \tau_{\text{ext}}}{(\hat{\mathbf{u}}^T \mathbf{J}_p)(\hat{\mathbf{u}}^T \mathbf{J}_p)^T + \lambda} \hat{\mathbf{u}}. \quad (10)$$

Both formulations assume a single contact point per end-effector; we do not localize multiple contacts and instead map them to an equivalent end-effector wrench. This formulation reduces to a set of independent 1D regularized least-squares problems, avoiding ill-conditioned Jacobian inversions while preserving compliant behavior in the directions of interest.

Torque components along selected axes are estimated in the same manner using $\mathbf{J}_r$.

D. Admittance Control with Inverse Kinematics

Given the estimated external wrench $\hat{\mathbf{f}}_{\text{ext}}$, admittance control is realized by integrating the spring–mass–damper dynamics to update the task-space motion reference. At each control step k , the reference acceleration $\ddot{\mathbf{x}}_{\text{ref}}^k$ is computed from Eq. (1), and the system is advanced using a semi-implicit Euler scheme:

$$\dot{\mathbf{x}}_{\text{ref}}^{k+1} = \dot{\mathbf{x}}_{\text{ref}}^k + \Delta t \ddot{\mathbf{x}}_{\text{ref}}^k, \quad (11)$$

$$\mathbf{x}_{\text{ref}}^{k+1} = \mathbf{x}_{\text{ref}}^k + \Delta t \dot{\mathbf{x}}_{\text{ref}}^{k+1}. \quad (12)$$

Here, k denotes the discrete control step, Δt is the control period. The resulting task-space reference is mapped to joint-space via an inverse kinematics (IK) solver [49], yielding position targets tracked by joint controllers. The IK solver uses damped least squares to handle near-singular Jacobians.

E. Safety Analysis

The admittance dynamics in Eq. (1) enforce a structural force–displacement relation. Letting $\Delta \mathbf{x} = \mathbf{x} - \mathbf{x}_{\text{des}}$ denote the deviation from the desired pose, with $\mathbf{f}_{\text{cmd}} = \mathbf{0}$ the dynamics reduce to:

$$m\ddot{\Delta \mathbf{x}} + \mathbf{K}_d \dot{\Delta \mathbf{x}} + \mathbf{K}_p \Delta \mathbf{x} = \hat{\mathbf{f}}_{\text{ext}}, \quad (13)$$

where m , $\mathbf{K}_p$, and $\mathbf{K}_d$ are inherited from Eq. (1). The corresponding transfer function is

$$\mathbf{A}(j\omega) = -\omega^2 m \mathbf{I} + j\omega \mathbf{K}_d + \mathbf{K}_p, \quad (14)$$

which reduces to $\mathbf{A}(j\omega) \approx \mathbf{K}_p$ in the quasi-static band, giving $\|\mathbf{f}_{\text{ext}}\| \approx \|\mathbf{K}_p\| \|\Delta \mathbf{x}\|$.

With bounded multiplicative estimation error $\hat{\mathbf{f}}_{\text{ext}} = (\mathbf{I} + \mathbf{E})\mathbf{f}_{\text{ext}}$, $\|\mathbf{E}\| \leq \varepsilon < 1$, the worst-case interaction force satisfies:

$$\|\mathbf{f}_{\text{ext}}\| \leq \frac{\|\mathbf{K}_p\| \|\Delta \mathbf{x}\|}{1 - \varepsilon}, \quad (15)$$

preserving admittance-induced safety bound up to a factor $(1 - \epsilon)^{-1}$ when the estimate is accurate in the task frequency band.

F. VLM-based Planning

The first planning method uses vision foundation models and vision-language models to predict 3D contact points and normals, which are interpolated into a dense trajectory.

Reasoning in Pixel Space. Our pipeline first applies open-vocabulary segmentation using SAM3 [4] on rectified RGB images to localize the target object. For wiping tasks, candidate pixels are grid-sampled within the mask bounding box, while for drawing tasks, candidates are generated from predefined end-effector workspace rectangles. A vision-language model (VLM) is then prompted with the rectified image and the list of candidate pixel coordinates to output an ordered sequence of waypoints for each end-effector.

Lifting to 3D Space. Depth is estimated using Foundation Stereo [42] and lifted to 3D using camera intrinsics. The resulting 3D point cloud is transformed from the camera frame to the world frame using the head pose and camera extrinsics. We then interpolate smooth trajectories at 50 Hz with approach and retreat segments, pauses, separate free-space and contact speed limits. Contact forces $\mathbf{f}_{\text{cmd}}$ are applied along the estimated surface normals, while stiffness matrices $\mathbf{K}_p$ are computed from surface normals and fixed tangential and normal values, and damping matrices $\mathbf{K}_d$ are derived assuming critical damping. The pipeline is illustrated in Fig. 2.

G. Imitation Learning

Data Collection. We employ a teleoperation setup similar to [1]. A leader arm computes end-effector poses via forward kinematics in MuJoCo [39] and streams them to the follower, where the received pose is used to set the desired end-effector target $\mathbf{x}_{\text{des}}$. The follower then tracks this target using a low-level compliance controller during physical interaction.

Policy Deployment. We use a diffusion policy [6] with a conditional 1D UNet that predicts diffusion noise over action sequences, conditioned on visual features extracted by a ResNet-18 encoder [18]. At each control step, the policy takes the observed end-effector pose $\mathbf{x}$ as input and outputs the desired end-effector target $\mathbf{x}_{\text{des}}$, which is passed to our compliance controller to generate $\mathbf{x}_{\text{ref}}$ for execution.

H. Model-based Planning

Hybrid Force-Velocity Planning. When contact location is known, a robust compliance profile can be computed in closed-form using optimally-conditioned hybrid servoing (OCHS) [20]. At each step, OCHS takes in the desired system velocities and contact information to compute a hybrid force-velocity control. The force control action maintains the safe engagement of desired contacts, while the velocity control action guarantees the desired object/robot velocities. To implement the computed hybrid force-velocity action on our compliance controller, we assign high stiffness in the velocity-controlled directions, while setting low stiffness and




Figure 3: **Comparison with Force Sensor Readings.** Dashed lines ($\hat{f}$) denote the estimated forces, while solid lines (f) show ground-truth measurements from an ATI Mini45 sensor. The mean absolute error is 0.69 ± 0.73 N for ToddlerBot with servo motors (gear ratio $> 200:1$) and 1.05 ± 1.60 N for ARX arm with QDD motor (gear ratio $\approx 10:1$).

force offset in the force-controlled directions for each end-effector. The stiffness matrix for an end-effector with a desired Cartesian velocity $\mathbf{v}$ is computed as

$$\mathbf{K} = k_{\text{low}}\mathbf{I} + (k_{\text{high}} - k_{\text{low}}) \frac{\mathbf{v}\mathbf{v}^\top}{\|\mathbf{v}\|^2}. \quad (16)$$

To simplify implementation, the velocity component produced by OCHS is integrated over time to generate Cartesian position commands, which, together with the force and stiffness commands, are jointly executed by our compliance controller. Using this formulation, the robot can perform contact-rich manipulation tasks, such as in-hand rotation, while maintaining robust contact behavior.

IV. EXPERIMENTS

A. Hardware Setup

We validate our compliance control method on four robotic platforms with distinct morphologies: ARX X5 robot arm with QDD motors; Unitree G1 with QDD motors; ToddlerBot [36] with Dynamixel servos; and LEAP Hand [35] with Dynamixel servos. None of the platforms is equipped with force/torque sensors, and some even lack current sensing and control, making them ideal for evaluating our approach based solely on motor current or PWM feedback and position control.

B. External Wrench Estimation

In this experiment, we press an ATI Mini45 force/torque sensor against the end-effector while running our compliance controller, apply pushes along each Cartesian axis, and estimate the resulting three-axis contact forces. The estimated forces closely match the sensor measurements, remaining near zero in free space and accurately tracking contact events.




Figure 4: **Qualitative Comparison with Baselines.** In this experiment, ToddlerBot draws a heart on a whiteboard using a target trajectory generated by a VLM. All methods follow the same command, which specifies the desired end-effector position $\mathbf{x}_{\text{des}}$, velocity $\dot{\mathbf{x}}_{\text{des}}$, stiffness $\mathbf{K}_p$, damping $\mathbf{K}_d$, and commanded wrench $\mathbf{f}_{\text{cmd}}$. We visualize the 3D end-effector trajectories and real-world execution results for UniFP [52], FACET [43], our method without $\hat{\mathbf{f}}_{\text{ext}}$, and our full method.

Table I: **Quantitative Comparison with Baselines.** We report position tracking errors in millimeters and orientation tracking errors in radians, defined as the deviation of $\mathbf{x}$ from $\mathbf{x}_{\text{des}}$. We also report the humanoid’s root pitch as a proxy for contact force. Since the contact point is at an approximately fixed distance from the center of mass, the moment arm is nearly constant, making root pitch correlated with the applied force.

	Pos. Error (mm)	Ori. Error (rad)	Root Pitch (rad)
UniFP [52]	57.8 ± 30.2	0.147 ± 0.119	0.068 ± 0.020
FACET [43]	22.4 ± 11.0	0.151 ± 0.087	0.018 ± 0.008
Ours w/o $\hat{\mathbf{f}}_{\text{ext}}$	22.5 ± 9.6	0.082 ± 0.040	0.060 ± 0.038
Ours	15.9 ± 5.1	0.048 ± 0.043	0.029 ± 0.012

Quantitative results are reported in Fig. 3. Although friction is not explicitly modeled, accounting for the drive state d (forward/backward) is sufficient for accurate estimation, with gearbox efficiency η implicitly capturing friction and transmission losses from high gear reductions. To reduce noise, we apply an exponential moving average (EMA) filter with $\alpha = 0.9$ for ToddlerBot and $\alpha = 0.1$ for the ARX arm. A larger α yields more responsive but noisier estimates, highlighting the trade-off between responsiveness and smoothing introduced by the EMA filter. A transient spike in ToddlerBot’s z -axis force around $t = 5$ s arises from a temporarily ill-conditioned Jacobian that couples x - and z -axis estimates; in practice, this has negligible impact on task performance.

C. Compliance Control

Baseline Implementation. We compare our compliance controller against two RL baselines, UniFP [52] and FACET [43], the state of the art in learning-based compliance. We retain the original policy formulations with only minor

implementation-level adjustments. Both baselines are tuned on our hardware for over one week in the same MuJoCo infrastructure, sharing identical command sampling ranges (Tab. II); reward weights are tuned per method (Tab. III).

During training, end-effector pose, external and commanded interaction wrenches, stiffness, and damping are uniformly sampled within the robot’s reachable workspace. The ranges ensure kinematic feasibility and remain consistent with those trajectories encountered during evaluation (Fig. 4), avoiding infeasible or out-of-distribution commands. We also adjust the reward weights to emphasize stable and responsive behavior under contact. More details are described in Appendix A.

Evaluation Results. As shown in Fig. 4, the heart-shaped target trajectory is generated by the VLM-based policy and executed by all methods. UniFP and FACET struggle to track the position accurately while maintaining appropriate contact force with the whiteboard. The external wrench estimation in our method is also crucial for regulating contact force. Tab. I reports quantitative results on position and orientation tracking errors, as well as the humanoid’s root pitch, which correlates with contact force. FACET exhibits the smallest root pitch because it fails to apply sufficient contact force to the board.

Additional Attributes. Our controller runs at about 12 ms per step, with about 10 ms spent in the IK solver [49]. These timings are profiled on a Jetson Orin NX. Moreover, since motor currents respond to disturbances anywhere on the robot, the controller can react beyond end-effector contacts. Qualitative results are shown in the supplementary video.

D. ToddlerBot Results

We evaluate three high-level planners on the ToddlerBot humanoid platform to demonstrate that our compliance controller integrates seamlessly across planning paradigms. As shown in Fig. 5, we consider three tasks: drawing and wiping on a whiteboard with a VLM-based policy, placing an egg on bread with a spatula using an imitation policy, and rotating a ball using a model-based policy. These tasks require appropriate contact force, motivating the use of compliance control. Each task includes a position-control baseline, which fails due to loss of contact, insufficient force (failing to establish or maintain contact), or excessive force (increasing tangential friction and tracking errors). We observe that the quasi-static assumption does not degrade performance in these tasks.

VLM-based Policy. For drawing and wiping, we attach a sponge to ToddlerBot’s left hand and a pen to the right. Given a text prompt (e.g., “star”) and an initial stereo observation of the whiteboard, vision foundation models and VLMs predict a 3D contact trajectory, which our compliance controller executes while maintaining contact force. After drawing, the robot captures a new observation to infer contact points for wiping. The compliance controller compensates for sometimes inaccurate predictions of 3D contact points and normals through compliant interaction. In contrast, the position-control baseline either fails to establish contact or applies excessive force, resulting in poor tangential tracking. The end-to-end pipeline runs in 3–6 s, with about 90% of the latency from the VLM




Figure 5: ToddlerBot Results. We demonstrate that Minimalist Compliance Control works on a floating-base humanoid robot and integrates seamlessly as a plug-and-play module across diverse high-level planners. Tasks include drawing and wiping on a whiteboard with a VLM-based policy, placing an egg on bread with a spatula using an imitation policy, and rotating a ball with a model-based policy. Odd rows (Ours) present successful executions of our method, while even rows (Position Control) illustrate representative failure cases of the baseline. Our controller maintains appropriate contact forces and stable interaction, whereas the position-control baseline frequently loses contact, applies insufficient force, or exerts excessive force, leading to task failure. Insufficient force typically fails to establish or maintain contact, while excessive force increases tangential friction and leads to larger tangential tracking errors.



Figure 6: **LEAP Hand Results**. We show that Minimalist Compliance Control generalizes across embodiments, including dexterous manipulation. We evaluate two high-level strategies: a VLM-based policy for multi-finger drawing and wiping on a whiteboard, and model-based planning for in-hand reorientation of various objects. Both tasks demand precise contact force regulation. Each also includes a position-control baseline, which fails due to contact loss, insufficient force, or excessive force.

API call. Qualitative results are shown in Fig. 5. More details are described in Appendix B.

Imitation Policy. For the egg placement task, we attach a spatula to ToddlerBot’s right hand. To train diffusion policies, we collect 200 demonstrations with the compliance controller active, logging both $\mathbf{x}_{\text{des}}$ and $\mathbf{x}_{\text{ref}}$. For a fair comparison, the compliance-enabled policy uses $\mathbf{x}_{\text{des}}$ as the action output, whereas the position-control baseline uses $\mathbf{x}_{\text{ref}}$, which is the trajectory after compliance-based adjustment via the spring-mass-damper model. Because inference-time trajectories are

not strictly in-distribution, the compliance controller corrects local errors by reacting to estimated external wrenches, whereas a pure position controller cannot. The baseline often fails to stabilize the bread as a stopper or applies excessive force with the spatula, leading to poor tangential tracking. Our method achieves a success rate of $16/20 = 80\%$, compared to $8/20 = 40\%$ for the position-control baseline.

Model-based Policy. In the ball rotation task, we command a target angular velocity and follow a fixed sequence of rotation commands specifying the axis and which hand to

use. For horizontal-axis rotations, the robot uses a single hand (left or right), whereas vertical-axis rotations are performed bimanually. The planner also takes the object’s geometry and scale as input. Our method achieves an average rotation of 21° , while the baseline typically slips or rotates the ball by only 2° – 3° . Rotation is estimated from multi-view observations by fitting the ball contour, matching surface features, lifting them to 3D sphere points, and solving for the best-fit rotation.

E. LEAP Hand Results

We further evaluate our controller on the LEAP Hand to show its effectiveness in dexterous manipulation. We consider two tasks: whiteboard drawing and wiping using a VLM-based policy, and in-hand reorientation of various objects using a model-based policy (Fig. 6). Both require precise contact force regulation. Although the LEAP Hand’s mechanical design provides some passive compliance that a position controller can partially exploit, we show that incorporating active compliance control is still crucial for reliable performance.

VLM-based Policy. We use a setup similar to the Toddler-Bot experiments, but the LEAP Hand can draw with multiple fingers simultaneously. A third-person view stereo camera observes the whiteboard. A pen is mounted on the index and ring fingers, and a sponge on the middle finger. Given the prompt “Letter O and Letter K,” a VLM-based policy predicts the 3D trajectory, which the compliance controller executes reliably. In contrast, the position-control baseline either fails to establish contact or tracks poorly due to excessive force.

Model-based Policy. For in-hand reorientation, we command target linear or angular velocities and execute a predefined sequence of commands. The planner receives the object’s geometry and scale as input. Our method achieves an average translation of 3.6 cm and a rotation of 27° , and can sustain reorientation for more than three cycles of translation and rotation commands. In contrast, the position-control baseline typically drops the object within the first few commands.

V. CONCLUSION

In conclusion, Minimalist Compliance Control achieves reliable compliance using widely available actuator signals, without force sensors or learning, making this capability accessible to many robotic systems with modern servo and QDD motors and lowering the barriers to contact-rich manipulation. We validate the method on a robot arm, a dexterous hand, and two humanoid robots across diverse contact-rich tasks, demonstrating that it generalizes across embodiments and integrates with vision–language, imitation, and model-based policies to achieve more stable contact than position control.

Limitations. Our current approach does not account for non-backdrivable or self-locking actuators. Motors with significant stiction or backlash are also not explicitly modeled and would require additional system identification. The wrench estimator further assumes a single contact point per end-effector and does not localize multiple contacts on the same kinematic chain. Additional effects, including joint acceleration, Coriolis terms, and thermal dynamics, could improve accuracy if

modeled, but this introduces a trade-off between accuracy and complexity. In this work, we prioritize a minimalist design that provides sufficient accuracy for reliable performance in contact-rich tasks.

ACKNOWLEDGMENT

The authors would like to express their great gratitude to the Toddy Team for hardware support and the members of Stanford TML and REALab for helpful discussions. This work was supported in part by the NSF Award #2143601, #2037101, #2132519, #2153854, Sloan Fellowship, and Stanford Institute for Human-Centered AI. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the sponsors.

REFERENCES

- [1] Jorge Aldaco, Travis Armstrong, Robert Baruch, Jeff Bingham, Sanky Chan, Debidatta Dwibedi, Chelsea Finn, Pete Florence, Spencer Goodrich, Wayne Gramlich, Alexander Herzog, Jonathan Hoech, Thinh Nguyen, Ian Storz, Baruch Tabanpour, Jonathan Tompson, Ayzaan Wahid, Ted Wahrburg, Sichun Xu, Sergey Yaroshenko, and Tony Z Zhao. ALOHA 2: An Enhanced Low-Cost Hardware for Bimanual Teleoperation, 2024.
- [2] Raunaq Bhirangi, Venkatesh Pattabiraman, Enes Er-ciyes, Yifeng Cao, Tess Hellebrekers, and Lerrel Pinto. AnySkin: Plug-and-Play Skin Sensing for Robotic Touch. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 16563–16570, May 2025. doi: 10.1109/ICRA55743.2025.11128638.
- [3] Max Yiye Cao, Stephen Laws, and Ferdinando Rodriguez y Baena. Six-Axis Force/Torque Sensors for Robotics Applications: A Review. *IEEE Sensors Journal*, 21(24): 27238–27251, December 2021. ISSN 1558-1748. doi: 10.1109/JSEN.2021.3123638.
- [4] Nicolas Carion, Laura Gustafson, Yuan-Ting Hu, Shoubhik Debnath, Ronghang Hu, Didac Suris, Chaitanya Ryal, Kalyan Vasudev Alwala, Haitham Khedr, Andrew Huang, Jie Lei, Tengyu Ma, Baishan Guo, Arpit Kalla, Markus Marks, Joseph Greer, Meng Wang, Peize Sun, Roman Rädle, Triantafyllos Afouras, Effrosyni Mavroudi, Katherine Xu, Tsung-Han Wu, Yu Zhou, Liliane Momeni, Rishi Hazra, Shuangrui Ding, Sagar Vaze, Francois Porcher, Feng Li, Siyuan Li, Aishwarya Kamath, Ho Kei Cheng, Piotr Dollár, Nikhila Ravi, Kate Saenko, Pengchuan Zhang, and Christoph Feichtenhofer. SAM 3: Segment Anything with Concepts, November 2025.
- [5] Sirui Chen, Zi-ang Cao, Zhengyi Luo, Fernando Castañeda, Chenran Li, Tingwu Wang, Ye Yuan, Linxi “Jim” Fan, C. Karen Liu, and Yuke Zhu. CHIP: Adaptive Compliance for Humanoid Control through Hindsight Perturbation, December 2025.
- [6] Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin Burchfiel, and Shuran Song. Diffu-

- sion Policy: Visuomotor Policy Learning via Action Diffusion. In *Robotics: Science and Systems XIX*. Robotics: Science and Systems Foundation, July 2023. ISBN 978-0-9923747-9-2. doi: 10.15607/RSS.2023.XIX.026.
- [7] Yufeng Chi, Qiayuan Liao, Junfeng Long, Xiaoyu Huang, Sophia Shao, Borivoje Nikolic, Zhongyu Li, and Koushil Sreenath. Demonstrating Berkeley Humanoid Lite: An Open-source, Accessible, and Customizable 3D-printed Humanoid Robot, April 2025.
- [8] Hojung Choi, Jun En Low, Tae Myung Huh, Gabriela A. Uribe, Seongheon Hong, Kenneth A. W. Hoffman, Julia Di, Tony G. Chen, Andrew A. Stanley, and Mark R. Cutkosky. CoinFT: A Coin-Sized, Capacitive 6-Axis Force Torque Sensor for Robotic Applications. <https://arxiv.org/abs/2503.19225v2>, March 2025.
- [9] Hojung Choi, Yifan Hou, Chuer Pan, Seongheon Hong, Austin Patel, Xiaomeng Xu, Mark R. Cutkosky, and Shuran Song. In-the-Wild Compliant Manipulation with UMI-FT, January 2026.
- [10] Clemens C. Christoph, Maximilian Eberlein, Filippos Katsimalis, Arturo Roberti, Aristotelis Sympetheros, Michel R. Vogt, Davide Liconti, Chenyu Yang, Barnabas Gavin Cangan, Ronan J. Hinchet, and Robert K. Katzschmann. ORCA: An Open-Source, Reliable, Cost-Effective, Anthropomorphic Robotic Hand for Uninterrupted Dexterous Task Learning, September 2025.
- [11] J.J. Craig and M.H. Raibert. A systematic method of hybrid position/force control of a manipulator. In *COMPSAC 79. Proceedings. Computer Software and The IEEE Computer Society's Third International Applications Conference, 1979.*, pages 446–451, November 1979. doi: 10.1109/COMPSAC.1979.762539.
- [12] K.S. Eom, I.H. Suh, W.K. Chung, and S.-R. Oh. Disturbance observer based force control of robot manipulator without force sensor. In *Proceedings. 1998 IEEE International Conference on Robotics and Automation (Cat. No.98CH36146)*, volume 4, pages 3012–3017 vol.4, May 1998. doi: 10.1109/ROBOT.1998.680888.
- [13] José Carlos Gamazo-Real, Ernesto Vázquez-Sánchez, and Jaime Gómez-Gil. Position and Speed Control of Brushless DC Motors Using Sensorless Techniques and Application Trends. *Sensors (Basel, Switzerland)*, 10(7):6901–6947, July 2010. ISSN 1424-8220. doi: 10.3390/s100706901.
- [14] David V. Gealy, Stephen McKinley, Brent Yi, Philipp Wu, Phillip R. Downey, Greg Balke, Allan Zhao, Menglong Guo, Rachel Thomasson, Anthony Sinclair, Peter Cuellar, Zoe McCarthy, and Pieter Abbeel. Quasi-Direct Drive for Low-Cost Compliant Robotic Manipulation, April 2019.
- [15] Noah Geiger, Tamim Asfour, Neville Hogan, and Johannes Lachner. Diffusion-Based Impedance Learning for Contact-Rich Manipulation Tasks, 2025.
- [16] Sami Haddadin and Erfan Shahriari. Unified force-impedance control. *The International Journal of Robotics Research*, 43(13):2112–2141, November 2024. ISSN 0278-3649. doi: 10.1177/02783649241249194.
- [17] Sami Haddadin, Alessandro De Luca, and Alin Albu-Schäffer. Robot Collisions: A Survey on Detection, Isolation, and Identification. *IEEE Transactions on Robotics*, 33(6):1292–1312, December 2017. ISSN 1941-0468. doi: 10.1109/TRO.2017.2723903.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, Las Vegas, NV, USA, June 2016. IEEE. ISBN 978-1-4673-8851-1. doi: 10.1109/CVPR.2016.90.
- [19] Zewen He, Chenyuan Chen, Dilshod Azizov, and Yoshihiko Nakamura. CoTaP: Compliant Task Pipeline and Reinforcement Learning of Its Controller with Compliance Modulation, October 2025.
- [20] Yifan Hou and Matthew T. Mason. An Efficient Closed-Form Method for Optimal Hybrid Force-Velocity Control. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11125–11131, May 2021. doi: 10.1109/ICRA48506.2021.9561426.
- [21] Yifan Hou, Zeyi Liu, Cheng Chi, Eric Cousineau, Naveen Kuppaswamy, Siyuan Feng, Benjamin Burchfiel, and Shuran Song. Adaptive Compliance Policy: Learning Approximate Compliance for Diffusion Guided Control. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4829–4836, May 2025. doi: 10.1109/ICRA55743.2025.11128452.
- [22] Kaizhe Hu, Haochen Shi, Yao He, Weizhuo Wang, C. Karen Liu, and Shuran Song. Robot Trains Robot: Automatic Real-World Policy Adaptation and Learning for Humanoids, August 2025.
- [23] Binghao Huang, Yixuan Wang, Xinyi Yang, Yiyue Luo, and Yunzhu Li. 3D-ViTac: Learning Fine-Grained Manipulation with Visuo-Tactile Sensing. In *8th Annual Conference on Robot Learning*, September 2024.
- [24] Marco Hutter, Christian Gehring, Dominic Jud, Andreas Lauber, C. Dario Bellicoso, Vassilios Tsounis, Jemin Hwangbo, Karen Bodie, Peter Fankhauser, Michael Bloesch, Remo Diethelm, Samuel Bachmann, Amir Melzer, and Mark Hoepflinger. ANYmal - a highly mobile and dynamic quadrupedal robot. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 38–44, October 2016. doi: 10.1109/IROS.2016.7758092.
- [25] Maged Iskandar, Alin Albu-Schäffer, and Alexander Dietrich. Intrinsic sense of touch for intuitive physical human-robot interaction. *Science Robotics*, 9(93):eadn4008, August 2024. doi: 10.1126/scirobotics.adn4008.
- [26] O. Khatib. A unified approach for motion and force control of robot manipulators: The operational space formulation. *IEEE Journal on Robotics and Automation*, 3(1):43–53, February 1987. ISSN 2374-8710. doi: 10.1109/JRA.1987.1087068.
- [27] Seunghwan Lee, Phil Sik Chang, and Jehee Lee. Deep

- Compliant Control. In *ACM SIGGRAPH 2022 Conference Proceedings*, SIGGRAPH '22, pages 1–9, New York, NY, USA, July 2022. Association for Computing Machinery. ISBN 978-1-4503-9337-9. doi: 10.1145/3528233.3530719.
- [28] Ung Hee Lee, Tor Shepherd, Sangbae Kim, Avik De, Hao Su, Robert Gregg, Luke Mooney, and Elliott Rouse. How to Model Brushless Electric Motors for the Design of Lightweight Robotic Systems, September 2023.
- [29] Qingzhou Lu, Yao Feng, Baiyu Shi, Michael Pisen, Zhenan Bao, and C. Karen Liu. GentleHumanoid: Learning Upper-body Compliance for Contact-rich Human and Object Interaction, November 2025.
- [30] Gabriel B. Margolis, Michelle Wang, Nolan Fey, and Pulkit Agrawal. SoftMimic: Learning Compliant Whole-body Control from Examples, October 2025.
- [31] Roberto Martín-Martín, Michelle A. Lee, Rachel Gardner, Silvio Savarese, Jeannette Bohg, and Animesh Garg. Variable Impedance Control in End-Effector Space: An Action Space for Reinforcement Learning in Contact-Rich Tasks. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1010–1017, November 2019. doi: 10.1109/IROS40897.2019.8968201.
- [32] Matthew T. Mason. Compliance and Force Control for Computer Controlled Manipulators. *IEEE Transactions on Systems, Man, and Cybernetics*, 11(6):418–432, June 1981. ISSN 2168-2909. doi: 10.1109/TSMC.1981.4308708.
- [33] T. Murakami, F. Yu, and K. Ohnishi. Torque sensorless control in multidegree-of-freedom manipulator. *IEEE Transactions on Industrial Electronics*, 40(2):259–265, April 1993. ISSN 1557-9948. doi: 10.1109/41.222648.
- [34] Sangok Seok, Albert Wang, Meng Yee Chuah, David Otten, Jeffrey Lang, and Sangbae Kim. Design principles for highly efficient quadrupeds and implementation on the MIT Cheetah robot. In *2013 IEEE International Conference on Robotics and Automation*, pages 3307–3312, May 2013. doi: 10.1109/ICRA.2013.6631038.
- [35] Kenneth Shaw, Ananye Agarwal, and Deepak Pathak. LEAP Hand: Low-Cost, Efficient, and Anthropomorphic Hand for Robot Learning. In *Robotics: Science and Systems XIX*. Robotics: Science and Systems Foundation, July 2023. ISBN 978-0-9923747-9-2. doi: 10.15607/RSS.2023.XIX.089.
- [36] Haochen Shi, Weizhuo Wang, Shuran Song, and C. Karen Liu. ToddlerBot: Open-Source ML-Compatible Humanoid Platform for Loco-Manipulation, February 2025.
- [37] Stefano Stassi, Valentina Cauda, Giancarlo Canavese, and Candido Fabrizio Pirri. Flexible tactile sensing based on piezoresistive composites: A review. *Sensors (Basel, Switzerland)*, 14(3):5296–5332, March 2014. ISSN 1424-8220. doi: 10.3390/s140305296.
- [38] Andreas Stolt, Magnus Linderöth, Anders Robertsson, and Rolf Johansson. Force controlled robotic assembly without a force sensor. In *2012 IEEE International Conference on Robotics and Automation*, pages 1538–1543, May 2012. doi: 10.1109/ICRA.2012.6224837.
- [39] Emanuel Todorov, Tom Erez, and Yuval Tassa. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033, October 2012. doi: 10.1109/IROS.2012.6386109.
- [40] Changhao Wang, Xiang Zhang, Zhian Kuang, and Masayoshi Tomizuka. Safe Online Gain Optimization for Cartesian Space Variable Impedance Control. In *2022 IEEE 18th International Conference on Automation Science and Engineering (CASE)*, pages 751–757, August 2022. doi: 10.1109/CASE49997.2022.9926697.
- [41] Lai Wei, Xuanbin Peng, Ri-Zhao Qiu, Tianshu Huang, Xuxin Cheng, and Xiaolong Wang. HMC: Learning Heterogeneous Meta-Control for Contact-Rich Loco-Manipulation, November 2025.
- [42] Bowen Wen, Matthew Trepte, Joseph Aribido, Jan Kautz, Orazio Gallo, and Stan Birchfield. FoundationStereo: Zero-Shot Stereo Matching, April 2025.
- [43] Botian Xu, Haoyang Weng, Qingzhou Lu, Yang Gao, and Huazhe Xu. FACET: Force-Adaptive Control via Impedance Reference Tracking for Legged Robots. In *9th Annual Conference on Robot Learning*, September 2025.
- [44] Xiaomeng Xu, Yifan Hou, Zeyi Liu, and Shuran Song. Compliant Residual DAgger: Improving Real-World Contact-Rich Manipulation with Human Corrections. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, October 2025.
- [45] Tae Hoon Yang, Haochen Shi, Jiacheng Hu, Zhicong Zhang, Daniel Jiang, Weizhuo Wang, Yao He, Zhen Wu, Yuming Chen, Yifan Hou, Monroe Kennedy, Shuran Song, and C. Karen Liu. Locomotion Beyond Feet, January 2026.
- [46] Shuangyue Yu, Tzu-Hao Huang, Xiaolong Yang, Chunhai Jiao, Jianfu Yang, Yue Chen, Jingang Yi, and Hao Su. Quasi-Direct Drive Actuation for a Lightweight Hip Exoskeleton With High Backdrivability and High Bandwidth. *IEEE/ASME Transactions on Mechatronics*, 25(4):1794–1802, August 2020. ISSN 1941-014X. doi: 10.1109/TMECH.2020.2995134.
- [47] Jianbo Yuan, Haohua Zhu, Jing Dai, and Sheng Yi. Development of the Bioinspired Tendon-Driven DexHand 021 with Proprioceptive Compliance Control, November 2025.
- [48] Wenzhen Yuan, Siyuan Dong, and Edward H. Adelson. GelSight: High-Resolution Robot Tactile Sensors for Estimating Geometry and Force. *Sensors*, 17(12):2762, December 2017. ISSN 1424-8220. doi: 10.3390/s17122762.
- [49] Kevin Zakka. Mink: Python inverse kinematics based on MuJoCo, May 2025.
- [50] Xiang Zhang, Liting Sun, Zhian Kuang, and Masayoshi Tomizuka. Learning Variable Impedance Control via Inverse Reinforcement Learning for Force-Related Tasks. *IEEE Robotics and Automation Letters*, 6(2):2225–2232,

April 2021. ISSN 2377-3766. doi: 10.1109/LRA.2021.3061374.

- [51] Xiang Zhang, Changhao Wang, Lingfeng Sun, Zheng Wu, Xinghao Zhu, and Masayoshi Tomizuka. Efficient Sim-to-real Transfer of Contact-Rich Manipulation Skills with Online Admittance Residual Learning. In *7th Annual Conference on Robot Learning*, August 2023.
- [52] Peiyuan Zhi, Peiyang Li, Jianqin Yin, Baoxiong Jia, and Siyuan Huang. Learning a Unified Policy for Position and Force Control in Legged Loco-Manipulation. In *9th Annual Conference on Robot Learning*, September 2025.
- [53] Xiang Zhou, Xinyu Zhang, and Qingrui Zhang. HAC-LOCO: Learning Hierarchical Active Compliance Control for Quadruped Locomotion under Continuous External Disturbances, July 2025.
- [54] Adrian Zwiener, Christian Geckeler, and Andreas Zell. Contact Point Localization for Articulated Manipulators with Proprioceptive Sensors and Machine Learning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 323–329, Brisbane, Australia, May 2018. IEEE Press. doi: 10.1109/ICRA.2018.8462869.