IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

RoboMT: Human-like Compliance Control for Assembly via a Bilateral Robotic Teleoperation and Hybrid Mamba-Transformer Framework

Wang Rundong[1†], Cheng Yanchun[1†], Yuan Qilong[2], Prakash Alok[3], Francis EH Tay[1] and Marcelo H. Ang Jr[1]

***Abstract*—Robotic compliance control is critical for delicate tasks such as electronic connector assembly, where precise force regulation and adaptability are paramount. However, traditional methods often struggle with modeling inaccuracies and sensor noise. Inspired by human adaptability in complex assembly operations, we present RoboMT, a novel framework that integrates a Mamba algorithm with a Transformer architecture to achieve human-like compliance control. By leveraging a bilateral teleoperation platform, we collect extensive real-time force/torque and motion data to form a comprehensive dataset for training. Furthermore, RoboMT incorporates an Adaptive Action Chunk module and a Temporal Fusion module to ensure smooth and robust action prediction. Experimental results across four electronic assembly tasks show that RoboMT achieves superior success rates (62–98%) over baselines (29–98%), while maintaining stable force regulation around 2.5N, closely resembling human performance. During task transitions, RoboMT quickly stabilizes at 5N with minimal overshoot, avoiding the large force spikes (over 24N) seen in baselines. Additionally, RoboMT maintains an average inference speed of 55 ms per batch, balancing real-time responsiveness and control robustness. Overall, RoboMT presents a compelling pathway toward error-minimized, human-level compliance control, and generalization for real-world robotic assembly, setting a new benchmark for precision, adaptability, and robustness in robotic assembly.**

***Index Terms*—Compliance and Impedance Control, Model Learning for Control, Telerobotics and Teleoperation**

I. Introduction

ROBOTS have been widely adopted in contactless tasks such as welding [1], drone spraying [2], and ultraviolet disinfection [3], but contact-rich and precise assembly tasks based on conventional compliance control algorithms are still challenging for robots to handle. Conventional compliance control generally falls into passive and active categories. Passive compliance uses mechanical flexibility (e.g., compliant wrists or actuators) to accommodate contact forces but suffers from limited compatibility and accuracy [4]. Active compliance typically employs hybrid force/position control—where the unconstrained directions follow position control while the constrained directions rely on force control [5]—or impedance control, which models the robot as a mass-damper-spring system to regulate interaction forces [6]. For hybrid force/position control, continuous switching based on contact conditions can negatively affect system stability [7]. The impedance control method still depends on the manual parameter adjustment, which limits adaptability under rapidly changing conditions [8]. In this study, electronic component assembly tasks, characterized by tight assembly tolerances and the need for flexible force control, are selected as representative tasks.

We draw inspiration from human strategies in assembly. For instance, in assembling electronic components, operators sense and adjust contact forces at the hand, aligning component pose, and arm movement to reduce undesired forces. By continuously modulating arm stiffness, they maintain the required force control to complete tasks successfully [9]. Therefore, we analyze and capture human compliance in both the force and motion domains and propose an end-to-end paradigm that can adaptively adjust motion based on contact forces/torques for human-like compliant behavior.

A fundamental challenge is the absence of comprehensive datasets capturing intricate force-motion dynamics at scale. While prior efforts in embodied intelligence have predominantly focused on visual, language and motion, the critical dimension of physical interaction forces remains underexplored due to technical difficulties in acquiring synchronized force-pose measurements during contact-rich manipulation tasks [10]. To bridge this gap, we develop a bilateral teleoperation system that combines human expertise in adaptive force modulation with robotic sensing precision. As shown on the left side of Fig.1, the platform enables: (1) real-time haptic feedback for operators to perceive the contact forces, (2) incremental pose/force adjustments guided by human tactile intelligence, and (3) synchronous recording of 6 degrees of freedom (DOF) force-torque and motion. By collecting human-guided compliant interactions across diverse electronic component assembly tasks, we establish the long-horizon force-motion dataset, directly supporting human-like compliance control learning framework.

With high-frequency interaction data continuously generated via the bilateral teleoperation platform, we face significant challenges in processing long sequential data. Recent breakthroughs in large-scale models have inspired us to apply

Manuscript received: January, 15, 2025; Revised April, 29, 2025; Accepted May, 30, 2025.

This paper was recommended for publication by Editor Ki-Uk Kyung upon evaluation of the Associate Editor and Reviewers' comments.

†Equal contribution.

[1]Authors are with the Advanced Robotics Centre, College of Design and Engineering, National University of Singapore, 117608, Singapore. {rundong, Yanchun.Cheng}@u.nus.edu, {mpetayeh, mpeangh}@nus.edu.sg

[2]Yuan Qilong is with Singapore Institute of Technology, 138683, Singapore. yu0017ng@e.ntu.edu.sg

[3]Prakash Alok is with Singapore-MIT Alliance for Research and Technology, 138602, Singapore. alok.prakash@smart.mit.edu

Accompanying content at https://github.com/robomt/supplementary

Digital Object Identifier (DOI): see top of this page.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

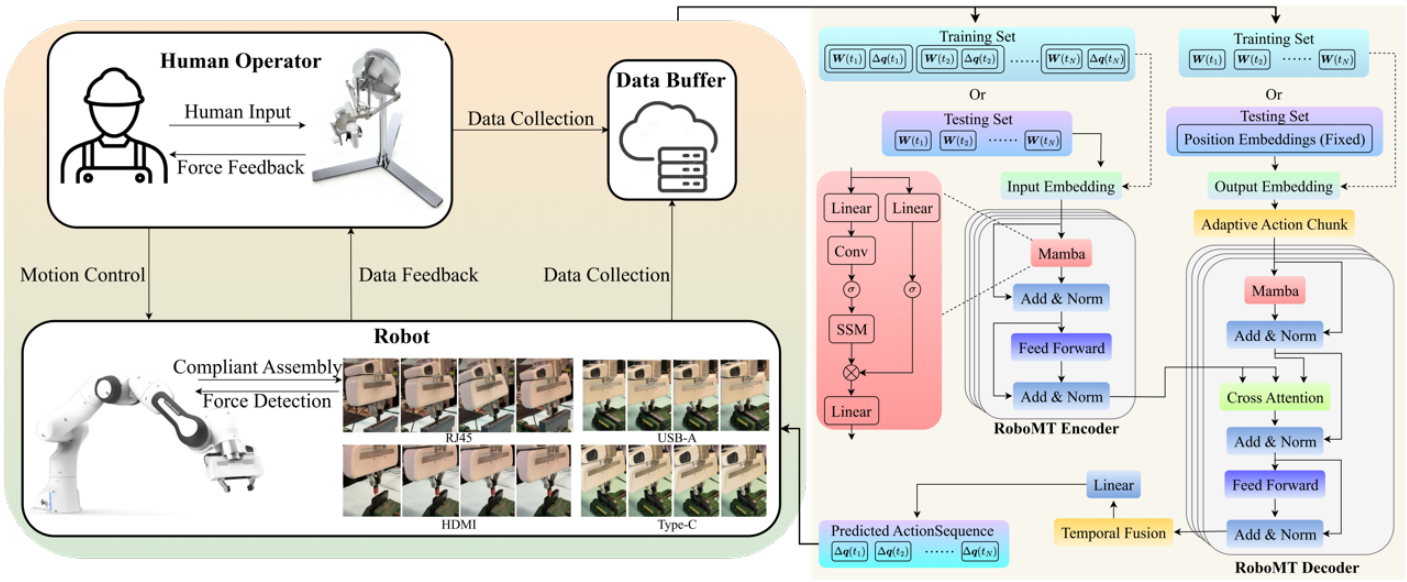

Fig. 1. Overview of the proposed framework for compliance control: The bilateral teleoperation platform on the left is designed to collect long-sequence force-motion data, which serves as the foundation for training the framework. On the right, the RoboMT architecture, combining the strengths of Mamba and Transformer, processes the collected data to achieve efficient global feature, ultimately enabling optimized compliance control.

similar approaches to robotic compliance control. One of the most representative methods for handling long sequences is the Transformer [11]. However, the quadratic computational complexity inherent in the Transformer's self-attention mechanism poses substantial challenges for continuous and high-frequency sensor data, demanding considerable resources and leading to efficiency constraints [12].

To address these concerns, we adapt the Mamba algorithm [13], which leverages the selective state-space model (SSM) to achieve linear computational scaling with sequence length. The algorithm's selective mechanism preserves global trend information while filtering high-frequency noise, enhancing robustness in dynamic interaction scenarios. However, Mamba lacks a native encoder-decoder architecture, limiting its flexibility for input-output sequence mapping compared to Transformers. We therefore propose RoboMT, a hybrid framework (Fig.1, right) that integrates Mamba into an encoder-decoder structure by replacing self-attention layers. On the one hand, Mamba's near linear-time complexity and reduced memory overhead allow more computing resources to be allocated to the cross-attention module, boosting the efficiency and robustness of information exchange between the encoder and the decoder. On the other hand, Mamba's strength in modeling long sequences and extracting global features integrates with the flexibility of Transformer in input-output mapping, achieving an appealing balance between performance and efficiency.

This letter presents an innovative robotic learning framework focused on compliance control and its main contributions are summarized as follows:

1) **Bilateral teleoperation platform and long-sequence force-motion dataset:** We build a bilateral teleoperation platform with real-time force feedback to the user, enabling precise and reliable fine motion control for delicate assembly tasks. This platform continuously records force-motion interactions, producing a comprehensive dataset reflecting operator compliance behaviors on diverse electronics assembly tasks, which is beneficial for training learning models for compliance control and embodied intelligence.
2) **A hybrid Mamba-Transformer framework:** We propose a learning framework-RoboMT-that achieves near-linear complexity in sequence modeling and efficient feature extraction. To the best of our knowledge, RoboMT is the first learning framework to process multiple DOF in an integrated manner for compliance control. In addition, we incorporate an Adaptive Action Chunk (AAC) module and an Temporal Fusion (TF) module for optimal action prediction. We plan to open-source the entire control framework to further support the robotics community.

II. Related Work

Achieving human-like compliance in robotic systems is critical for safe interaction, particularly in tasks that demand adaptability and precision. Traditional methods for estimating arm stiffness, such as electromyography-based sensing [14] and least-squares fitting algorithms [15], have been used to adjust robotic compliance accordingly. Although these approaches perform well in environments with larger tolerances, they often struggle in high-precision tasks due to limitations arising from modeling inaccuracies and measurement noise.

To address these challenges, recent research has shifted towards integrating intelligent algorithms into compliance control frameworks. Specifically, techniques such as reinforcement learning (RL) [16] and imitation learning (IL) [17] have been employed to cast impedance tuning as a policy learning or regression problem, enabling more adaptive and dynamic control strategies. Although RL can learn adaptive controllers

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

without expert demonstrations, its reliance on well-defined reward functions, safety assurances, and sample efficiency hinders direct real-world deployment [18]. In contrast, IL bypasses reward function engineering by directly learning state-to-action mappings from expert demonstrations, making it especially suitable for tasks where success criteria are hard to quantify. However, single-step IL often struggles with non-Markovian human behavior and long-horizon dependencies [19]. Although recurrent neural networks address some temporal aspects [20], they are limited in handling extended sequences, motivating the use of Transformer architectures for global temporal modeling and improved policy coherence. Despite progress in tackling robotic tasks, many studies underutilize force data, reducing applicability in force-sensitive tasks. For instance, [21] and [22] only consider joint trajectories, with [22] relying exclusively on simulation data that may not transfer well to real scenarios. Although [23] incorporates force sensing, it merely estimates environmental stiffness without deeper control integration. And [24] employs vibration to convey contact force, such qualitative feedback lacks the granularity required for more complex tasks. Additionally, certain works partition data with fixed patch sizes [25], [26] or truncate sequences to reduce computational load [27], potentially sacrificing long-range dependency modeling. Although Mamba-based approaches [28] achieve linear complexity, they often require specialized preprocessing for inputs of variable length, increasing the complexity of the system. Building on these insights, this study leverages both Mamba and Transformer to efficiently model force-motion data over extended horizons, offering robust compliance control in complex environments.

III. Preliminaries

We adapt position-based impedance control as the basic mathematical model to capture the relationship between external contact forces/torques and corresponding motion adjustments. The core idea of this control is to translate deviations in external forces/torques into motion correction signals, which the robot controller then executes to maintain compliant end-effector force regulation [29]. Formally, the robot–environment interaction is modeled by a unified six DOF impedance model that simultaneously addresses translational and rotational compliance. This hybrid position-orientation impedance relationship is expressed as:

$$\boldsymbol{M}\Delta\ddot{\boldsymbol{q}} + \boldsymbol{B}\Delta\dot{\boldsymbol{q}} + \boldsymbol{K}\Delta\boldsymbol{q} = \boldsymbol{W} \tag{1}$$

where $\Delta\boldsymbol{q} = \left[\Delta\boldsymbol{p}, \Delta\boldsymbol{\theta}\right]^T$ represents the generalized pose adjustment, comprising translational displacement $\Delta\boldsymbol{p} \in \mathbb{R}^3$ and orientation deviation $\Delta\boldsymbol{\theta} \in \mathbb{R}^3$. The impedance matrices $\boldsymbol{M}$, $\boldsymbol{B}$, and $\boldsymbol{K}$ are full 6×6 positive-definite matrices that respectively represent the generalized inertia, damping, and stiffness parameters without assuming decoupling between translational and rotational components. The external wrench $\boldsymbol{W} = \left[\boldsymbol{F}, \boldsymbol{\tau}\right]^T$ combines contact forces $\boldsymbol{F}$ and torques $\boldsymbol{\tau}$.

In precision assembly tasks, operators typically work at low velocities, making velocity and acceleration errors less critical. While existing studies predominantly simplify robotic motion as decoupled across individual DOF for analytical tractability, practical physical interactions inherently exhibit dynamic coupling — where forces/torques in one DoF induce motion perturbations in others. Hence, we emphasize an end-to-end mapping from force/torque to motion adjustments across 6 DOF, providing a more holistic and robust compliance control.

IV. System and Methods

A. Bilateral Teleoperation Platform

In this section, we develop a bilateral teleoperation platform that provides dual feedback in both the force and the motion domains, enabling operators to perceive and manipulate remote tactile information. Our hardware setup comprises a Sigma.7 leader device for force feedback to operators and motion input, and a Franka Emika Panda robot as the remote execution platform. Sigma.7 offers high-resolution force feedback with minimal interference [30], while the Franka robot's built-in sensors measure and report real-time forces [31]. This configuration ensures an accurate, bidirectional exchange of force and motion data, establishing a foundation for subsequent compliance control learning.

Software-wise, we integrate Force Dimension SDK and Franka ROS to exchange data between Sigma.7 and the Franka robot. The rosbag module is used to record contact forces, torques, and motion adjustments with synchronized timestamps, ensuring data integrity for subsequent visualization, analysis, and model training. In conclusion, the developed platform is shown in Fig.2.

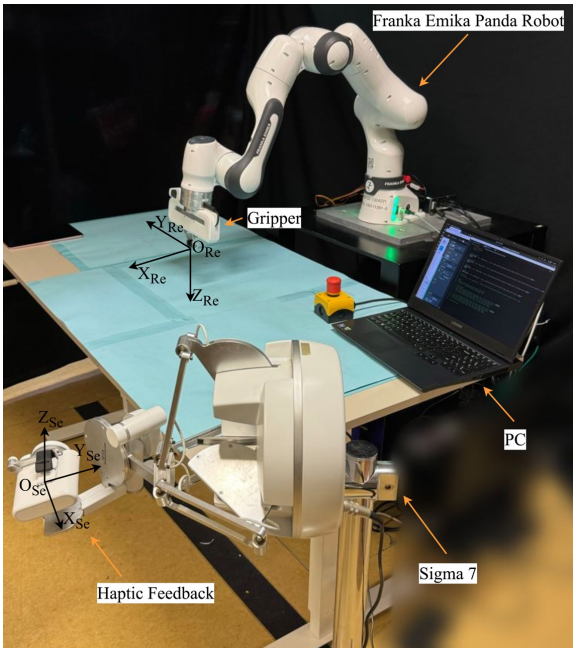

Fig. 2. Bilateral teleoperation platform. The system consists of a Franka Emika Panda robot with a parallel gripper, a Sigma.7 haptic interface for real-time force feedback, and a PC for motion control and data collection.

Due to the smaller workspace of Sigma.7 compared to the Franka robot, direct mapping of end-effector poses or joint coordinates can lead to poor workspace utilization and frequent recalibration. To address this, we scale the relative pose

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

changes from Sigma.7 proportionally by a factor, using end-effector pose difference mapping aligned with the compliance control principles in Section III, thereby minimizing post-processing overhead. We also consider the problem of singularities and the computational burden. Given that the spatial range required for electronic assembly tasks is relatively small, the robot is configured to effectively prevent singularities. Moreover, the redundancy of the Franka robot allows efficient inverse kinematics solutions without substantially increasing computational load, thus improving efficiency while preserving operational flexibility [32].

Our system forms a closed-loop control architecture: the operator senses the forces/torques of the Franka robot's end effector through Sigma.7 and adjusts the pose of Sigma.7, indirectly influencing the Franka robot's motion to reduce undesired contact forces or torques. At time step t, we record the force/torque vector $\boldsymbol{W}(t)$:

$$\begin{aligned}\boldsymbol{W}(t) &= \left[\boldsymbol{F}(t), \boldsymbol{\tau}(t)\right]^T \\ &= \left[f_x(t), f_y(t), f_z(t), \tau_x(t), \tau_y(t), \tau_z(t)\right]^T\end{aligned} \tag{2}$$

where $f_x(t)$, $f_y(t)$ and $f_z(t)$ represent the contact force components, and $\tau_x(t)$, $\tau_y(t)$ and $\tau_z(t)$ represent the torque components, all defined along or about the x, y and z axes of the Franka robot's end-effector frame (see Fig. 2).

And the corresponding motion adjustments $\Delta\boldsymbol{q}(t)$:

$$\begin{aligned}\Delta\boldsymbol{q}(t) &= \left[\Delta\boldsymbol{p}(t), \Delta\boldsymbol{\theta}(t)\right]^T \\ &= \left[\Delta x(t), \Delta y(t), \Delta z(t), \Delta\alpha(t), \Delta\beta(t), \Delta\gamma(t)\right]^T\end{aligned} \tag{3}$$

Here, $\Delta x(t)$, $\Delta y(t)$, and $\Delta z(t)$ denote the positional changes, and $\Delta\alpha(t)$, $\Delta\beta(t)$, and $\Delta\gamma(t)$ represent the incremental Euler angles, all defined along or about the x, y and z axes of the Sigma.7 end-effector frame (see Fig. 2).

For different types of assembly tasks, the system records a comprehensive dataset D, which consists of time steps t_i, contact forces/torques $\boldsymbol{W}\left(t_i\right)$, and the corresponding motion adjustments $\Delta\boldsymbol{q}\left(t_i\right)$.

$$\begin{aligned}D &= \{(t_i, \boldsymbol{W}\left(t_i\right), \Delta\boldsymbol{q}\left(t_i\right))\}_{i=1}^N \\ &= \{(t_i, \boldsymbol{f}(t_i), \boldsymbol{\tau}(t_i), \Delta\boldsymbol{x}(t_i), \Delta\boldsymbol{\theta}(t_i))\}_{i=1}^N\end{aligned} \tag{4}$$

B. RoboMT: A hybrid Mamba-Transformer Framework

We collect the dataset D where the robot's end-effector forces/torques serve as state variables and the corresponding motion adjustments as action variables. As shown in Fig. 1, during training, the forces, torques and motion adjustments sampled are embedded in a high-dimensional feature space and passed to the RoboMT encoder. The encoder comprises stacked layers, each featuring a Mamba sub-layer and a Feed-Forward Network (FFN) sub-layer. Within the Mamba sub-layer, inputs are linearly projected and then processed by an SSM and a convolution module. The SSM filters out irrelevant information while preserving essential global features, and the convolution module captures local spatio-temporal patterns. A nonlinear activation function further enhances the model's ability to capture complex motion relationships. The fused features are then projected linearly to form the output of this sub-layer. Subsequently, the FFN sub-layer performs nonlinear transformations on the features. Each sub-layer is followed by a residual connection and layer normalization to stabilize training and preserve deep-level information. After stacking multiple encoder layers, the output incorporates global information and provides high-quality feature representations for the subsequent decoder.

The RoboMT decoder predicts motion adjustments given the end-effector forces/torques. These are first passed through an embedding layer, and the AAC module determines the length of the output action sequence. In the cross-attention sub-layer, the decoder uses the output of the RoboMT encoder as key-value pairs, aligning predictions with current states and source sequences. Like the encoder, the decoder employs multi-layer Mamba and FFN sub-layers with residual connections and normalization. This design ensures that the decoder can fully exploit the information from both the input states and the previous actions at each decoding step. Finally, the decoder output layer will predict the action sequence, while the TF module refines the predicted target actions. The predicted actions are then compared to human demonstrations to compute the loss function. By iteratively updating the encoder and decoder parameters via backpropagation, RoboMT progressively improves its predictions.

For testing, we deploy RoboMT with the lowest validation loss in real-world scenarios. During compliance control, RoboMT processes end-effector forces/torques and outputs motion adjustments to the robot controller. By leveraging the learned relationships among forces, torques, and motion adjustments, RoboMT delivers precise, robust motion predictions that meet practical task demands.

Additionally, we introduce an optimized action-prediction modules: AAC and TF. The motivation for this optimization is twofold. First, the dataset collected in Section IV-A reflects the inherent stochasticity of human operators—diverse control strategies, incidental pauses during demonstrations, and the randomness of human behavior. Because human decision-making is intrinsically non-Markovian, it depends on prior experience and previously visited states. Second, even small inaccuracies in the predicted actions can propagate and cause large deviations in the subsequent system state. This compounding error problem drives the agent toward regions of the state space that diverge progressively from the demonstrated trajectories.

To address the above issues, previous research [24] explored the action chunking, where the model predicts the next k time steps of actions rather than just one, reducing short-horizon ambiguity. However, fixed k can be problematic: too small k fails to capture long-range dependencies, whereas large k risks overfitting [21]. We therefore propose the AAC model, which dynamically updates the chunk size k_t based on current task conditions. Specifically, we track the variations in contact force and torque, to assess environmental dynamics and adjust k_t accordingly:

$$k_{t+1} = k_t + \alpha \cdot \sigma(\boldsymbol{f}(t), \boldsymbol{\tau}(t)) \tag{5}$$

where k_t is the magnitude of the current action at step t; α is the adjustment step size, which controls the range of dynamic

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

changes in action magnitude; $\sigma(\boldsymbol{f}(t), \boldsymbol{\tau}(t))$ is a state function based on the contact force and torque variation, defined as:

$$\sigma(\boldsymbol{f}(t), \boldsymbol{\tau}(t)) = \begin{cases} +1, & \text{if } \|\boldsymbol{f}(t)\| + \|\boldsymbol{\tau}(t)\| < \epsilon, \\ -1, & \text{if } \|\boldsymbol{f}(t)\| + \|\boldsymbol{\tau}(t)\| \geq \epsilon. \end{cases} \tag{6}$$

where ϵ is the threshold distinguishing the environmental variations.

In addition, switching between the 'observation' and 'execution' modes may introduce sudden shifts in the predicted action sequence. To mitigate this, we incorporate the TF that queries the policy at each timestep, allowing the predicted actions from overlapping chunks to coincide. As a result, at time t, the model produces multiple action predictions. We employ Gaussian-based weighting to fuse these predictions, highlighting key time steps while attenuating the influence of others. The weight for each time step is given by:

$$w_i = \exp\left(-\frac{(i-c)^2}{2\sigma^2}\right) \tag{7}$$

where w_i is the weight of timestep; c is the central timestep; σ controls the weight spread.

The final action a_t is then computed by:

$$a_t = \frac{\sum_{i=t-k_t+1}^{t} w_i A_t[i]}{\sum_{i=t-k_t+1}^{t} w_i} \tag{8}$$

where $A_t[i]$ denotes the predicted action at time t.

C. Methodological and Architectural Innovations

The RoboMT framework draws inspiration from [21] and [22], yet both of these works exhibit multiple shortcomings that necessitate further optimization. First, they underutilize contact force feedback and primarily focus on positions, limiting their capability in intricate contact tasks. Moreover, [22] relies solely on simulated data, introducing a pronounced sim-to-real gap that hinders practical deployment. Furthermore, both studies adopt a conditional variational autoencoder (CVAE) to model multi-modal distributions, requiring a latent style variable z. However, real-world task conditions often diverge from the Gaussian assumption, resulting in unstable predictions. Estimating and reparameterizing z for long sequences imposes additional computational overhead.

To address these constraints, RoboMT leverages real-world task data collected by our developed bilateral teleoperation platform to comprehensively capture the coupling of contact forces, torques, and motion adjustments. Therefore, RoboMT enhances robustness and adapts more effectively to real-world conditions. Unlike CVAE-based approaches that explicitly model z, RoboMT uses Mamba with recursive mechanisms and state storage to implicitly capture global patterns in long sequences, avoiding the instability caused by mismatched distributions and maintaining near-linear complexity. This mechanism concentrates on critical features and filters out irrelevant fluctuations, enabling more stable imitation learning and compliant force control. Moreover, whereas [21] and [22] utilize positional embeddings, the recursive nature of Mamba inherently encodes spatiotemporal positions, thus simplifying the architecture and improving the processing efficiency of long sequence.

V. Experiments

A. Implementation Details

Four representative electronic components were chosen for this study (assembly stages are illustrated in Fig. 3). Assembly data for each component type was collected through our custom-developed bilateral teleoperation platform, yielding 270 demonstrations per component. The bilateral teleoperation and data acquisition were executed on a Lenovo Legion 5 laptop equipped with an Intel Core i7-11800H CPU and an NVIDIA GeForce RTX 3060 GPU.

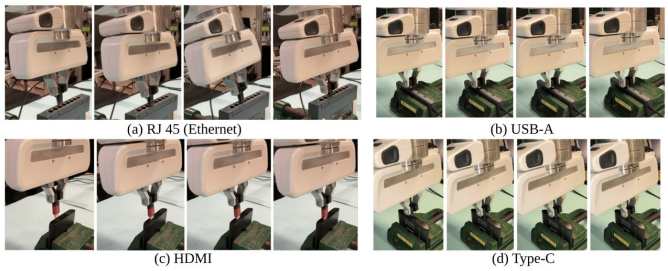

Fig. 3. Four typical electronic component assembly tasks. Four assembly tasks involving RJ45 (a), USB-A (b), HDMI (c), and Type-C (d), each with assembly tolerances of 0.1–0.2 mm. The components feature distinct geometries: RJ45 (polygonal), USB-A (rectangular), HDMI (trapezoidal), and Type-C (oval). The high precision of the assembly tolerances and irregular sizes require compliance control of pose and force. Specifically, the robot aligns the plug with the socket under medium stiffness and switches to low stiffness upon contact to ensure a smooth and secure connection. And we provide the video for each task in the supplementary materials.

Six volunteers (three males, three females; age 26–30) who have the professional experience in electronics assembly participated in data collection sessions. Before recording began, each operator spent 20 minutes familiarizing themselves with the teleoperation system. All participants provided written informed consent. During each trial the operator manipulated a Sigma.7 haptic device and, guided by the force/torque vector feedback, strove to complete the assembly as efficiently as possible while minimizing part–part collisions. Successful completion was defined by full component insertion.

The duration of demonstrations ranged from 7 to 10 seconds depending on component complexity, with sampling standardized at 50 Hz, generating approximately 94.5 to 135k temporal steps per component type. The acquired dataset was employed for RoboMT training using a computational cluster equipped with four NVIDIA L40S GPUs. The training protocol encompassed 800 epochs with three distinct random seeds to ensure statistical robustness.

To ensure fairness, all comparative experiments were conducted with identical hardware configurations and training protocols. The experimental results are averaged over 20 trials to ensure statistical reliability. Additional implementation details, including hyperparameter settings, are available in our accompanying GitHub repository.

B. Comparison of RoboMT and Baseline Models in Assembly

We evaluate RoboMT on various assembly tasks and compare it to five other models suitable for real world settings. These include Action Chunking with Transformers (ACT)

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

TABLE I
COMPARISON OF DIFFERENT MODELS UNDER VARIOUS CONFIGURATIONS.

Models	Inference time (ms/batch)			Minimum loss			Success rate (%): RJ45			Success rate (%): USB-A			Success rate (%): HDMI			Success rate (%): Type-C		
	w/o AC	w/ FAC	w/ AAC	w/o AC	w/ FAC	w/ AAC	w/o AC	w/ FAC	w/ AAC	w/o AC	w/ FAC	w/ AAC	w/o AC	w/ FAC	w/ AAC	w/o AC	w/ FAC	w/ AAC
ACT	163	205	227	0.1632	0.1508	0.1253	52	67	72	61	70	77	64	74	82	79	85	90
Comp-ACT	171	209	235	0.1357	0.1173	0.0895	59	72	81	65	72	82	76	88	94	87	93	**98**
Bi-ACT	183	215	232	0.1483	0.1245	0.0955	55	70	82	68	75	80	70	80	89	84	90	92
MaIL	**22**	**29**	**37**	0.1721	0.1515	0.1312	53	66	75	59	72	79	68	77	84	75	83	88
RoboS4T	29	33	43	0.3143	0.2372	0.2181	29	42	56	31	45	59	40	52	63	52	67	72
RoboMT(Ours)	35	41	55	**0.1205**	**0.0973**	**0.0774**	**62**	**78**	**85**	**70**	**81**	**88**	**81**	**90**	**95**	**88**	**95**	**98**

[21], Compliance Control via Action Chunking with Transformers (Comp-ACT) [24], Bilateral Control-Based Imitation Learning via Action Chunking with Transformer (Bi-ACT) [25], Mamba Imitation Learning (MaIL) [28], and RoboS4T.

ACT integrates a CVAE with a Transformer backbone to capture multimodal demonstration data and model high-dimensional spatiotemporal dependencies. Comp-ACT extends ACT to learn variable stiffness control from limited demonstrations, predicting future end-effector poses and stiffness in Cartesian space. Bi-ACT integrates bilateral control with ACT to enable imitation learning that captures both position and force information, achieving robust manipulation of objects with varying hardness and weight. MaIL leverages the Mamba architecture, to enhance imitation learning efficiency and generalization. To validate the impact of the selective state-space mechanism, we introduce RoboS4T, a variant of RoboMT in which Mamba is replaced with the original Structured State Space (S4) model while retaining the hybrid encoder-decoder framework, enabling a direct comparison of linear sequence modeling capabilities without selective information filtering.

Table I presents a comparative analysis of inference time (ms/batch), minimum loss, and success rates across four electronic assembly tasks (RJ45, USB-A, HDMI, and Type-C) under three configurations: without Action Chunk (w/o AC), with Fixed Action Chunk (w/ FAC), and with Adaptive Action Chunk (w/ AAC). To clearly evaluate the inherent advantages of our framework, we initially compare the performance of different methods under the same configuration of w/ AAC and TF. Detailed ablation studies for AAC and TF will be discussed in Section V-E.

As shown in Table I, RoboMT achieves a minimum loss (0.0774) and superior success rates of 85% (RJ45), 88% (USB-A), 95% (HDMI) and 98% (Type-C), outperforming all baselines. This significant improvement arises from RoboMT's modeling of the dynamic relationship between force/torque feedback and corresponding motion adjustments, enabling human-like compliance strategies learned from collected datasets reflecting human compliant assembly strategies. In contrast, baseline models exhibit critical limitations. ACT and MaIL disregard force/torque feedback entirely, preventing detection of contact force variations or dynamic compliance adjustment. This rigidity causes behavioral instability in critical phases (e.g., final assembly alignment), often resulting in task failure. While Comp-ACT and Bi-ACT incorporate force information, their implementations are flawed: Comp-ACT relies on vibration-based qualitative force estimation, introducing stiffness-level errors, whereas Bi-ACT focuses on isolated force control without integrating force-motion dynamics, and its limited consideration of both force and motion dimensions further restricts overall performance improvements compared to RoboMT. RoboS4T leverages the S4 architecture for faster inference (43 ms vs. RoboMT's 55 ms). However, it lacks Mamba's selective state mechanism, making it more susceptible to noise in high-frequency sensor sequences. This results in maximum losses (e.g., 0.2181 vs. RoboMT's 0.0774) and lowest success rates (e.g., 56-72%), highlighting the importance of Mamba's selective state-space mechanism.

It is noteworthy that MaIL, employing exclusively the Mamba framework, shows a significant inference-time advantage (lowest inference times of 37 ms) compared to Transformer-based methods (ACT, Comp-ACT, and Bi-ACT). It also slightly outperforms RoboS4T and RoboMT. Nevertheless, MaIL exhibits limited robustness in sequence processing, necessitating input preprocessing to handle sequential data [28]. Furthermore, the absence of an encoder-decoder architecture and selective information cross-attention restricts its performance in complex compliance control tasks. This trade-off highlights RoboMT's balanced design: it combines Mamba's efficiency with Transformer's sequence modeling strength, achieving both speed and precision.

In the w/ FAC configuration, all methods utilize an identical number of action chunks. RoboMT consistently outperforms baseline models, achieving the lowest loss (0.0973) and the highest success rates across all tasks (78% for RJ45, 81% for USB-A, 90% for HDMI and 95% for Type-C). In comparison, ACT, Comp-ACT, Bi-ACT, and MaIL show moderate success (66–93%) but suffer higher losses (0.1173–0.1515). RoboS4T exhibits substantial performance gaps, with significantly the highest loss (0.2372) and the lowest success rates (42–67%).

Under the w/o AC condition, RoboMT maintains superior performance with the lowest loss (0.1205) and the highest success rates (62% for RJ45, 70% for USB-A, 81% for HDMI and 88% for Type-C), further demonstrating its inherent modeling advantages even without additional adaptive components. Baseline methods suffer substantial performance drops, particularly RoboS4T. This deficiency forces RoboS4T to rely on AAC and TF to compensate for its inability to handle sensor noise and dynamic task variations. These observations underscore RoboMT's intrinsic capability in effectively leveraging the relationship between force/torque data and corresponding motion adjustments, thus ensuring robust and reliable task execution across diverse assembly scenarios.

Thus, RoboMT demonstrates a balanced advantage across inference efficiency, loss minimization, and robustness in task performance, positioning it as a highly effective solution for precise electronic assembly tasks.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

C. Axial Force Regulation in Electronics Assembly Task

One of the goals in compliant assembly is to minimize the axial force during the process. Using Type-C insertion as an example, Fig.4 shows the axial contact force profiles for RoboMT, Comp-ACT, and human demonstration, where RoboMT and Comp-ACT achieve the highest success rates.

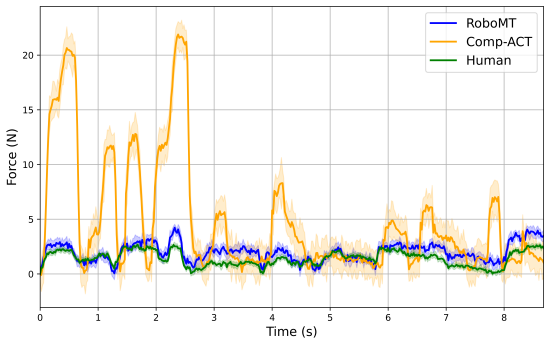

Fig. 4. The contact force measured during Type-C insertion. The solid line depicts the average contact force across multiple trials for each policy, while the shaded region represents the variability across trials, shown as one standard deviation above and below the mean contact force at each time.

Human operators (green) achieve near-identical force stability (average: 2.5N; shading: ±1.0N), while RoboMT (blue) exhibits a force trajectory that closely mimics human performance, maintaining stable oscillations between 0–5N with an average of 2–3N and minimal variance (shading width: ±1.2N), thereby highlighting RoboMT's success in replicating human-like force modulation. This contrasts sharply with Comp-ACT (orange), which shows erratic spikes exceeding 20N in early stages and persistent 5–7N fluctuations later, along with broader uncertainty regions (±4.5N). These results demonstrate RoboMT's ability to dynamically adapt to contact variations, akin to human proprioceptive adjustments, achieving biologically inspired adaptability.

D. Axial Force Regulation Across Assembly Task Transitions

Fig.5 presents the axial contact force profiles for a continuous assembly sequence in which the task switches from an RJ45 to an HDMI connector.

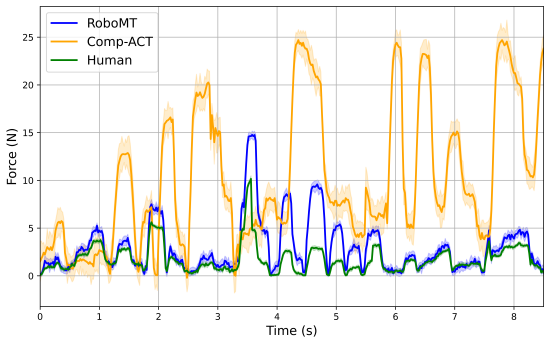

Fig. 5. Comparison of contact force during the transition between variable assembly tasks. The meanings of the solid line and shaded region are consistent with those in Fig. 4.

Throughout the sequence, RoboMT (blue) maintains forces largely within the 0–10 N safety envelope, exhibiting only brief transients when the assembly conditions change at 3.3s and 5s, after which it quickly re-converges. The visibly narrower band for RoboMT (shaded region: ±1.5N) attests to its higher consistency and compliance in force control. In contrast, Comp-ACT (orange) produces recurrent force peaks exceeding 20N—and at times approaching 25N—together with greater variability (±4.8N). We attribute this advantage to RoboMT's nearly linear-time inference scheme, which permits on-the-fly traversal of longer interaction histories and adaptive adjustment of control gains, thereby preserving human-like compliance even under abrupt changes in different environments. In contrast, Comp-ACT's latent variable cannot be updated online during extended tasks, leading to delayed compensation and excessive contact forces. Collectively, RoboMT achieves a force profile that not only more closely matches that of human operators (green), but also exhibits smaller fluctuations, highlighting its robustness and suitability for dynamic, high-precision electronic assembly.

E. Ablation Studies on the AAC and TF Modules

We conduct ablation studies to validate the necessity and generality of the proposed AAC module. As illustrated in Table I, integrating AAC consistently enhances performance across all evaluated models. For instance, in the RJ45 insertion task, RoboMT with AAC improves the success rate to 85%, significantly outperforming FAC (78%) and the baseline without AC (62%). The reduction in minimum loss (e.g., 0.0774 for RoboMT w/ AAC vs. 0.1205 for RoboMT w/o AC) further highlights AAC's ability to mitigate cumulative prediction errors. Moreover, even baseline models (ACT, Comp-ACT, Bi-ACT, MaIL), despite inherently limited or absent modeling of force-motion dynamics, benefit notably from AAC. Their success rates increase markedly (e.g., ACT from 67% to 72% for RJ45, MaIL from 59% to 79% for USB-A), accompanied by reductions in minimum losses. This suggests that AAC can generalize to diverse models, leading to improved robustness and higher precision in complex assembly scenarios.

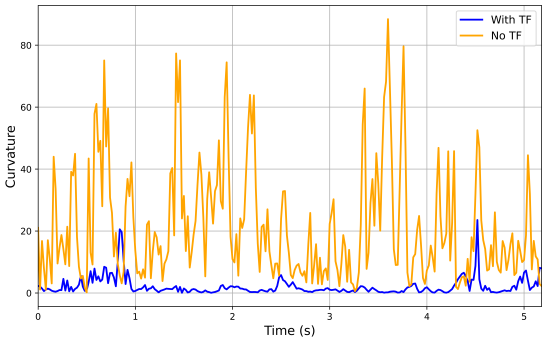

Fig. 6. Curvature profiles with and without TF over time.

Fig. 6 shows the curvature evolution over 5 seconds under two conditions: with TF (blue) and without TF (orange). Without TF, curvature frequently spikes above 40 and approaches 80, indicating large fluctuations. With TF, it stays mostly below 10, reflecting much more stable behavior. These results suggest that TF reduces curvature variations, improving the smoothness and reliability of the assembly.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

VI. Conclusion

In this letter, we present RoboMT, a novel learning framework that integrates a hybrid Mamba-Transformer architecture for enhanced compliance control. By constructing a force-interaction enabled teleoperation system, we collect the dataset capturing force/torque feedback and operator-driven motion adjustments, providing critical training dataset for robust policy learning. Furthermore, the proposed AAC module and TF module improve the stability and continuity of action predictions, ensuring smooth operation.

We perform experiments on four representative electronic assembly tasks. RoboMT achieves state-of-the-art performance by outperforming existing baselines in control robustness, success rates, and stability of force regulation, while maintaining low inference latency. These results demonstrate the superiority of RoboMT in achieving accurate, efficient, and human-like compliance behavior for assembly tasks.

A key limitation of the current method is its reliance solely on force feedback, without incorporating visual perception. This limits its performance in scenarios requiring visual guidance, such as handling large initial misalignment. In the future, we aim to integrate visual perception with force feedback and consider the optimal weighting between visual and force-motion information, further advancing the system's ability to meet the demands of automated assembly across a broader range of complex and dynamic scenarios.

References

[1] R. Suzuki, Y. Okada, Y. Yokota, T. Saijo, H. Eto, Y. Sakai, K. Murano, K. Ohno, K. Tadakuma, and S. Tadokoro, "Cooperative towing by multi-robot system that maintains welding cable in optimized shape," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11 783–11 790, 2022.

[2] A. S. Vempati, M. Kamel, N. Stilinovic, Q. Zhang, D. Reusser, I. Sa, J. Nieto, R. Siegwart, and P. Beardsley, "Paintcopter: An autonomous uav for spray painting on three-dimensional surfaces," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 2862–2869, 2018.

[3] J. Byun, J. Byun, J. Kang, I. Yi, J. Lee, K. Noh, J. Kim, Y. Choi, G. Chung, S. Oh, *et al.*, "Autonomous ultraviolet-c disinfection and wiping robot: Assessment in a hospital environment," *IEEE Robotics & Automation Magazine*, 2025.

[4] T. Kastritsi and A. Ajoudani, "A passive power-based control strategy for phri tasks with omni-directional robotic mobile platforms," *IEEE Robotics and Automation Letters*, vol. 9, no. 8, pp. 6959–6966, 2024.

[5] J. Li, H. Tang, M. Lv, X. Liao, P. Zhang, B. Zhao, P. K. Wong, and Y. Hu, "Force-position hybrid control for robot assisted thoracic-abdominal puncture with respiratory movement," *IEEE Robotics and Automation Letters*, vol. 9, no. 6, pp. 5262–5268, 2024.

[6] K. Karacan, R. J. Kirschner, H. Sadeghian, F. Wu, and S. Haddadin, "Tactile robot programming: Transferring task constraints into constraint-based unified force-impedance control," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024.

[7] G. Nava, Q. Sablé, M. Tognon, D. Pucci, and A. Franchi, "Direct force feedback control and online multi-task optimization for aerial manipulators," *IEEE Robotics and Automation Letters*, 2019.

[8] T. K. Best, C. G. Welker, E. J. Rouse, and R. D. Gregg, "Data-driven variable impedance control of a powered knee–ankle prosthesis for adaptive speed and incline walking," *IEEE Transactions on Robotics*, vol. 39, no. 3, pp. 2151–2169, 2023.

[9] D. J. Braun, V. Chalvet, T.-H. Chong, S. S. Apte, and N. Hogan, "Variable stiffness spring actuators for low-energy-cost human augmentation," *IEEE Transactions on Robotics*, 2019.

[10] W. Liu, J. Wang, Y. Wang, W. Wang, and C. Lu, "Forcemimic: Force-centric imitation learning with force-motion capture system for contact-rich manipulation," *arXiv preprint arXiv:2410.07554*, 2024.

[11] X. Chen, H. Peng, D. Wang, H. Lu, and H. Hu, "Seqtrack: Sequence to sequence learning for visual object tracking," in *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.

[12] F. Ilhan, G. Su, S. F. Tekin, T. Huang, S. Hu, and L. Liu, "Resource-efficient transformer pruning for finetuning of large models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 16 206–16 215.

[13] A. Gu and T. Dao, "Mamba: Linear-time sequence modeling with selective state spaces," *arXiv preprint arXiv:2312.00752*, 2023.

[14] Y. Zhuang, Y. Leng, J. Zhou, R. Song, L. Li, and S. W. Su, "Voluntary control of an ankle joint exoskeleton by able-bodied individuals and stroke survivors using emg-based admittance control scheme," *IEEE Transactions on Biomedical Engineering*, 2021.

[15] P. Atreya, H. Karnan, K. S. Sikand, X. Xiao, S. Rabiee, and J. Biswas, "High-speed accurate robot control using learned forward kinodynamics and non-linear least squares optimization," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 11 789–11 795.

[16] H. Zhang, G. Solak, G. J. G. Lahr, and A. Ajoudani, "Srl-vic: A variable stiffness-based safe reinforcement learning for contact-rich robotic tasks," *IEEE Robotics and Automation Letters*, 2024.

[17] G. N. S. Gubbala, M. Nagashima, H. Mori, Y. A. Seong, H. Sato, R. Niiyama, Y. Suga, and T. Ogata, "Augmenting compliance with motion generation through imitation learning using drop-stitch reinforced inflatable robot arm with rigid joints," *IEEE Robotics and Automation Letters*, vol. 9, no. 10, pp. 8595–8602, 2024.

[18] F. Pizarro Bejarano, L. Brunke, and A. P. Schoellig, "Safety filtering while training: Improving the performance and sample efficiency of reinforcement learning agents," *arXiv e-prints*, pp. arXiv–2410, 2024.

[19] C. Weaver, C. Tang, C. Hao, K. Kawamoto, M. Tomizuka, and W. Zhan, "Betail: Behavior transformer adversarial imitation learning from human racing gameplay," *IEEE Robotics and Automation Letters*, vol. 9, no. 8, pp. 7302–7309, 2024.

[20] A. N. Fathima, S. S. Ibrahim, and A. Khraisat, "Enhancing network traffic anomaly detection: Leveraging temporal correlation index in a hybrid framework," *IEEE Access*, 2024.

[21] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware," in *Proceedings of Robotics: Science and Systems*, 2023.

[22] J. Huang, Q. Shi, D. Xie, Y. Ma, X. Liu, C. Li, and X. Duan, "Visuomotor policy learning for task automation of surgical robot," *IEEE Transactions on Medical Robotics and Bionics*, 2024.

[23] M. Kobayashi, T. Buamanee, Y. Uranishi, and H. Takemura, "Ilbit: Imitation learning for robot using position and torque information based on bilateral control with transformer," *arXiv preprint arXiv:2401.16653*, 2024.

[24] T. Kamijo, C. C. Beltran-Hernandez, and M. Hamaya, "Learning variable compliance control from a few demonstrations for bimanual robot with haptic feedback teleoperation system," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 12 663–12 670.

[25] T. Buamanee, M. Kobayashi, Y. Uranishi, and H. Takemura, "Bi-act: Bilateral control-based imitation learning via action chunking with transformer," *arXiv preprint arXiv:2401.17698*, 2024.

[26] X. Xu, Y. Liang, B. Huang, Z. Lan, and K. Shu, "Integrating mamba and transformer for long-short range time series forecasting," *arXiv preprint arXiv:2404.14757*, 2024.

[27] T. Tsuji, "Mamba as a motion encoder for robotic imitation learning," *arXiv preprint arXiv:2409.02636*, 2024.

[28] X. Jia, Q. Wang, A. Donat, B. Xing, G. Li, H. Zhou, O. Celik, D. Blessing, R. Lioutikov, and G. Neumann, "Mail: Improving imitation learning with mamba," *arXiv preprint arXiv:2406.08234*, 2024.

[29] K. Samuel, K. Haninger, R. Oboe, S. Haddadin, and S. Oh, "A perturbation-robust framework for admittance control of robotic systems with high-stiffness contacts and heavy payload," *IEEE Robotics and Automation Letters*, vol. 9, no. 7, pp. 6432–6439, 2024.

[30] H. Su, Y. Schmirander, Z. Li, X. Zhou, G. Ferrigno, and E. De Momi, "Bilateral teleoperation control of a redundant manipulator with an rcm kinematic constraint," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 4477–4482.

[31] S. Haddadin, S. Parusel, L. Johannsmeier, S. Golz, S. Gabl, F. Walch, M. Sabaghian, C. Jähne, L. Hausperger, and S. Haddadin, "The franka emika robot: A reference platform for robotics research and education," *IEEE Robotics Automation Magazine*, 2022.

[32] S. Tittel, "Analytical solution for the inverse kinematics problem of the franka emika panda seven-dof light-weight robot arm," in *2021 20th International Conference on Advanced Robotics (ICAR)*. IEEE, 2021, pp. 1042–1047.