



Article

Adaptive Variable Admittance Control for Intent-Aware Human–Robot Collaboration

Mohammad Jahani Moghaddam ^{1,2,*}  and Filippo Arrichiello ¹ 

¹ Department of Electrical and Information Engineering, University of Cassino and Southern Lazio, Via G. Di Biasio 43, 03043 Cassino, Frosinone, Italy; f.arrichiello@unicas.it

² Department of Electrical Engineering, Lan.C., Islamic Azad University, Student Street, Langarud 4471311127, Iran

* Correspondence: mohammad.jahanimoghaddam@unicas.it or dr.jahani@iau.ac.ir

Abstract

This paper presents a comprehensive framework for evaluating the robustness and adaptability of human–robot collaboration (HRC) controllers under a spectrum of dynamic and unpredictable human intentions. Building upon variable admittance controller (VAC) frameworks augmented with Radial Basis Function Neural Network (RBFNN) online adaptation, we introduce two key innovations: (1) an intent-aware human force generator capable of simulating aggressive, hesitant, oscillatory, conflicting, and nominal behaviors, through the modulation of force gains and the introduction of stochastic noise, and (2) the extension of VAC to incorporate variable stiffness as an adaptive control parameter alongside damping and inertia. The adaptive parameters are jointly tuned online using a self-supervised learning (SSL) mechanism driven by motion error metrics and interaction dynamics. The framework is simulated in a dual-arm collaborative manipulation scenario involving two 7-DoF Franka Emika Panda robots transporting a shared object in a high-fidelity simulation environment. Simulation results demonstrate the system’s capability to maintain stable behavior and minimize tracking error despite abrupt changes in human intent. This work provides a novel and systematic tool for stress-testing adaptive controllers in HRC, with implications for the design of resilient, safe, and reliable robotic systems in real-world collaborative environments.

Keywords: human–robot collaboration (HRC); variable admittance control (VAC); human intent modeling; online self-supervised learning; adaptive control



Academic Editors: Yanhong Peng and Fernando Gomez-Bravo

Received: 31 December 2025

Revised: 6 February 2026

Accepted: 6 February 2026

Published: 12 February 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

The paradigm of human–robot collaboration (HRC) has emerged as a cornerstone of modern robotics, promising to revolutionize industries from manufacturing and logistics to healthcare and personal assistance. By combining the strength, precision, and endurance of robots with the cognitive flexibility, intuition, intention-driven decision making, and problem-solving skills of humans, HRC systems can achieve a level of performance that neither could attain alone, which enhanced productivity and situational adaptability. However, the seamless integration of robots into human-centric environments remains a formidable challenge: how can robotic systems reliably interact with human partners whose behaviors are dynamic, non-deterministic, inherently unpredictable, and often ambiguous? [1,2].

A critical aspect of effective HRC is the robot’s ability to adapt its behavior to the dynamic and often ambiguous human intentions and environmental changes, leading

to the development of compliance control strategies such as impedance control (IC) and admittance control (AC). Impedance control regulates the relationship between position deviation and interaction force by directly shaping the robot's mechanical impedance by virtual mass-spring-damper dynamics, making it suitable for lightweight or torque-controlled systems. In contrast, admittance control computes motion commands in response to external forces/torques through a virtual dynamic model, which is particularly advantageous for stiff or position-controlled manipulators. In this work, an admittance-based framework is adopted to enable compliant and adaptive behavior during collaborative manipulation tasks. Admittance control has proven particularly effective in tasks involving significant external inputs. Admittance controllers modulate robot motions in response to external forces. Traditional admittance controllers, however, rely on fixed parameters and thus struggle to accommodate rapidly changing human behaviors or variable task demands [3–5]. To address this limitation, variable admittance control (VAC) strategies have emerged, where parameters such as virtual damping and inertia are adjusted online based on evolving interaction dynamics [6–8]. VAC allows robots to dynamically adjust their responsiveness to external forces, fostering safer and more intuitive interactions.

While significant progress has been made in developing adaptive control strategies, such as VAC, most existing research has focused on optimizing performance under nominal, cooperative conditions [6,7,9]. Recent studies have further enhanced VAC with neural networks and learning mechanisms to improve adaptability. For instance, Radial Basis Function Neural Networks (RBFNNs) have demonstrated effectiveness in real-time parameter tuning without requiring extensive training datasets.

However, real-world human–robot interactions are often characterized by dynamic, unpredictable, and sometimes non-cooperative human behaviors, such as hesitation, aggression, or even direct conflict. Focusing on optimizing performance under nominal, cooperative conditions leaves a critical gap in our understanding of how these systems perform when faced with these real-world behaviors. Previously presented VAC frameworks, although robust under nominal conditions, were not specifically designed or evaluated for their resilience to such diverse and challenging human intentions. Furthermore, in many existing admittance control formulations, virtual stiffness is treated as a fixed parameter, which limits the robot's ability to fully modulate its compliance and haptic interaction characteristics, which are crucial for both comfort and accuracy in varied HRC scenarios.

This paper addresses these critical limitations by presenting a comprehensive framework for evaluating the robustness and adaptability of HRC controllers under a spectrum of dynamic and unpredictable human intentions. Building upon our foundational VAC framework, we introduce two significant innovations: (1) A novel dynamic human intent module: This module is capable of simulating a wide range of human behaviors, including nominal, aggressive, hesitant, oscillatory, and conflicting intents. It achieves this through the dynamic modulation of force gains and the introduction of stochastic noise, with temporal switching between these intent profiles. This provides a systematic and repeatable tool for stress-testing adaptive controllers in HRC, which was a significant gap in prior research. (2) Extension of the VAC scheme to include variable stiffness (K): This work integrates stiffness as a fully adaptive control parameter alongside damping (D) and inertia (M). All three parameters are jointly tuned online using a unified self-supervised learning (SSL) mechanism driven by motion error metrics and interaction dynamics. Moreover, the adaptation strategy is formulated to remain passivity compliant, preventing unintended energy injection during rapid parameter updates and thereby maintaining stable physical interaction. This enhancement allows for a more comprehensive and nuanced control of the robot's compliance, crucial for managing contact stability and haptic comfort in diverse interaction scenarios.

Through these advancements, our aim is to provide a robust and adaptive VAC architecture capable of maintaining stable behavior and minimizing tracking error despite abrupt changes in human intent. A quantitative evaluation of controller performance across varying intent profiles highlights the importance of intent-aware adaptation for real-world HRC systems. The framework is simulated and validated in a dual-arm collaborative manipulation scenario involving two 7-DoF Franka Emika Panda robots transporting a shared object in a high-fidelity simulation environment. By providing a means to test HRC systems beyond nominal conditions, this work contributes significantly to the development of resilient, safe, and reliable robotic systems for real-world collaborative environments. A dual-arm configuration is adopted to address more realistic industrial scenarios involving shared load distribution, internal force management, and coordinated motion—challenges not fully captured in single-arm setups.

To our knowledge, this is one of the first studies to present a comprehensive evaluation of an adaptive HRC controller under such a diverse and dynamic range of simulated human intentions, building on related efforts in intent simulation [10,11].

2. Literature Review

While significant progress has been made in VAC strategies [6–8], most existing studies focus on nominal cooperative conditions and treat stiffness as fixed. Early VAC strategies relied on heuristic gain scheduling or fuzzy logic controllers to tune parameters based on qualitative expert knowledge or simple feedback laws. However, these approaches are limited by their dependence on handcrafted rules and lack of scalability to high-dimensional systems. Recent advancements in VAC have incorporated learning-based methods, particularly neural networks, to further enhance adaptability. Feed-forward neural networks, recurrent networks, and reinforcement learning have all been explored, but often require extensive offline training or introduce instability under real-time operation.

Radial Basis Function Neural Networks (RBFNNs) have demonstrated effectiveness in real-time parameter tuning due to their ability to model complex nonlinear relationships and generalize from local approximations. Although recent studies have shown that RBFNN-based VAC frameworks improved robustness under nominal human inputs and reduced human effort [7], a critical gap in existing research is the limited capacity to model and test controller response to more diverse or adverse human behaviors—such as oscillatory motions, sudden conflicts, or erratic forces. Most prior VAC studies have primarily evaluated performance under nominal, cooperative conditions [12,13].

Furthermore, while damping and inertia adaptation have been extensively explored, virtual stiffness, a critical factor for both comfort and accuracy in HRC, has often been treated as a fixed parameter in many control models. The ability to adapt stiffness dynamically is crucial for fine-tuning the robot's perceived rigidity or compliance, especially in scenarios requiring precise contact or varied haptic feedback. Understanding and responding to human intent is another cornerstone of intuitive HRC. Research in this area has focused on two main approaches: real-time intent estimation using sensors and machine learning and simulating human intent in virtual environments to test and benchmark controllers [14].

Sensor-based methods—ranging from model-based approaches that use biomechanical data [15,16] to learning-based methods [17,18] that infer intent from sensory inputs such as force/torque measurements, vision, and electromyography (EMG) signals [19,20]—enable real-time intent recognition but are limited by hardware complexity, calibration issues, the requirement extensive training data, and poor generalization. In contrast, simulation-based intent modeling provides a safe and repeatable framework for evaluating control strategies, particularly under hazardous or unpredictable conditions.

Some prior works have explored intent switching [11] or disturbance injection [21] to evaluate robustness, but these are typically limited to binary cooperative/adversarial scenarios or fixed-force profiles. While very few studies offer systematic modeling of diverse human intentions, nor do they integrate such intent variation into adaptive controller design [22], this paper addresses the aforementioned limitations by making two primary contributions. First, we introduce a novel dynamic human intent module that systematically simulates a wide spectrum of human behaviors, from cooperative to adversarial, through time-varying force profiles and noise perturbations. This module allows for rigorous stress testing of HRC controllers. Second, we extend the variable admittance control (VAC) scheme to include stiffness as a fully adaptive parameter, allowing real-time tuning alongside damping and inertia using a unified RBFNN learning framework. This comprehensive approach, to our knowledge, represents a significant advancement in evaluating and designing adaptive HRC systems under diverse and dynamic human intent profiles.

Recent work has started to explore stiffness modulation, but typically in isolated or one-dimensional settings [23]. Our proposed framework expands this frontier by jointly adapting damping, inertia, and stiffness in 6D space. The ability to adjust all three admittance parameters provides more precise and responsive control, particularly in tasks requiring fine manipulation or robust behavior under uncertainty.

While most learning-based VAC approaches rely on supervised training or reinforcement learning, these methods typically require extensive data collection and pre-training, offline optimization, or labeled datasets and suffer from poor generalization. In contrast, self-supervised learning (SSL) allows the system to generalize quickly to new and unseen human behaviors, even during task execution by learning from its own prediction errors during interaction, based solely on task-space errors and estimated interaction forces, without requiring pre-collected training data. Prior applications of SSL in robotics include attention learning [24], symbolic planning [25], and vision-based manipulation [26], but its use for real-time adaptation of VAC parameters in HRC remains limited. Another distinctive aspect of our approach is the use of an SSL strategy, which improves the system's ability to generalize across unforeseen human behaviors. To our knowledge, this is among the first implementations of real-time SSL for joint adaptation of stiffness, damping, and inertia in a dual-arm HRC context.

2.1. Self-Supervised Learning for Real-Time Admittance Adaptation

Despite the growing adoption of data-driven approaches such as supervised learning and reinforcement learning in robotic control, their applicability in human–robot collaboration remains limited by data requirements, offline training procedures, and sensitivity to distribution shifts in human behavior. In contrast, self-supervised learning (SSL) provides a learning paradigm in which supervisory signals are generated intrinsically from the system's own interaction outcomes, eliminating the need for external labels or pre-collected datasets. In the context of robotic control, SSL differs fundamentally from conventional learning-based approaches by exploiting task execution errors—such as trajectory deviation or force inconsistency—as implicit supervisory signals. This paradigm enables continuous online adaptation while preserving real-time feasibility and stability. Prior applications of SSL in robotics have primarily focused on perception, attention learning, and symbolic reasoning [24–26], whereas its use for real-time compliance modulation in physical human–robot interaction remains limited. In this work, SSL is employed to adapt the parameters of a variable admittance control (VAC) framework through a Radial Basis Function Neural Network (RBFNN). Unlike classical adaptive control laws, which rely on explicitly designed parameter update rules derived from model assumptions, the proposed SSL

mechanism learns a nonlinear correction mapping from observed task-space tracking errors and estimated interaction forces to admittance parameter adjustments. The supervisory signal is generated implicitly from deviations between the desired minimum-jerk motion and the observed object motion, enabling online learning without external supervision or offline training. This formulation allows the controller to generalize across diverse and unforeseen human behaviors, while maintaining bounded adaptation and real-time responsiveness. As a result, the proposed SSL-based VAC bridges the gap between traditional adaptive control and learning-based compliance modulation, offering improved flexibility without sacrificing interpretability or stability.

2.2. Gap in Literature and Novel Contributions

Despite significant progress in VAC and intent modeling, several gaps remain: (1) Most VAC systems address only single-robot scenarios; dual-arm collaboration remains underexplored. (2) Human intent modeling is often oversimplified or excluded from simulation-based validation. (3) Stiffness adaptation is rarely integrated alongside damping and inertia in a unified control scheme. (4) Few studies evaluate controller robustness under dynamically changing human behavior.

This paper addresses these gaps by (1) proposing a novel intent module that systematically simulates five distinct human behaviors with temporal switching and stochastic variability; (2) integrating this module into a dual-arm VAC framework where stiffness, damping, and inertia are adapted online via an RBFNN trained through self-supervised learning; and (3) evaluating controller robustness and performance in a high-fidelity MATLAB 2022b simulation with full 6D object manipulation, under abrupt and diverse human input changes.

Table 1 compares our evaluation framework with recent studies on human intent modeling and simulation in HRC. The criteria emphasize systematic aspects: diversity of intent categories (beyond typical 2–3 cooperative/adversarial), presence of dynamic temporal switching and stochastic variability (e.g., noise for realism), breadth of quantitative metrics (e.g., stratified RMSE across intents, force profiles, adaptive parameters), and number of baseline controllers compared. While prior works advance intent prediction or virtual data generation [10,22,27], most are limited to fewer categories, lack stochastic perturbations or frequent switching, or evaluate under limited baselines. Our approach integrates five distinct profiles with scheduler-driven transitions, Gaussian noise, full 6D multi-parameter adaptation, and explicit comparisons across three VAC variants (fixed, partial, full), enabling more rigorous robustness testing under diverse and abrupt human behaviors.

Table 1. Comparison of systematic evaluation criteria in recent studies on human intent modeling and simulation for HRC controllers. Criteria include number of intent categories, dynamic switching/stochastic elements, evaluation metrics, and baseline comparisons.

| Criterion/Study | Description/Details |
|-----------------------|--|
| This Work | |
| Intent Categories | 5 (nominal, aggressive, hesitant, oscillatory, conflicting) |
| Switching/Stochastic | Temporal scheduler + Gaussian noise |
| Evaluation Metrics | Intent-stratified RMSE (pos/vel), force magnitude, parameter evolution, human effort |
| Baselines/Comparisons | 3 (fixed-gain, partial adaptive, full adaptive) |
| [27] (2020) | |
| Intent Categories | 2–3 (cooperative, fast/slow) |
| Switching/Stochastic | Limited switching, no stochastic |
| Evaluation Metrics | Tracking error, human effort, damping adaptation |
| Baselines/Comparisons | 1–2 (fixed vs. variable admittance) |

3.1. Variable Admittance Control (VAC) Framework in 6D

The VAC framework forms the foundation of our system. The controller adjusts the desired trajectory of the robot based on the external forces applied by the human, with the relationship between force and motion governed by a virtual admittance model. The key feature of our VAC is the use of an SSLRBFNN to dynamically update the damping, inertia, and stiffness parameters of this model in real time. This allows the controller to adapt its compliance to the changing dynamics of the interaction, ensuring both safety and efficiency.

The VAC governs the object's acceleration in response to the estimated human force and robot wrenches. The VAC control law in task space is

$$\mathbf{M}(t)\ddot{\mathbf{x}} = \mathbf{w}_{total} - \mathbf{D}(t)(\dot{\mathbf{x}}_d - \dot{\mathbf{x}}) - \mathbf{K}(t)(\mathbf{x}_d - \mathbf{x}) \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^6$ is the current object pose (position + orientation), and \mathbf{x}_d and $\dot{\mathbf{x}}_d$ are the desired velocity and position from MJT. The total wrench \mathbf{w}_{total} consists of the following three components:

$$\mathbf{w}_{total} = \mathbf{w}_{human} + \mathbf{w}_{robot} + \mathbf{w}_{grav}, \quad (2)$$

where \mathbf{w}_{human} represents the external wrench applied by the human, \mathbf{w}_{robot} is the wrench generated by the dual-arm manipulators, and \mathbf{w}_{grav} accounts for gravitational effects acting on the manipulated object. $\mathbf{M}(t), \mathbf{D}(t), \mathbf{K}(t) \in \mathbb{R}^{6 \times 6}$ are diagonal time-varying matrices, as follows:

$$\begin{aligned} \mathbf{M}(t) &= \text{diag}(m_x, m_y, k_w \cdot m_z, \mathbf{I}_\theta), \\ \mathbf{D}(t) &= \text{diag}(d_x, d_y, d_z, \mathbf{D}_\theta), \\ \mathbf{K}(t) &= \text{diag}(k_x, k_y, k_z, \mathbf{K}_\theta). \end{aligned} \quad (3)$$

Diagonal matrices are employed to decouple translational and rotational components in 6D task space, facilitating independent tuning of parameters along each axis and reducing computational complexity while preserving adequate performance for the symmetric object manipulation task [3]. The vertical impedance scaling factor k_w is introduced to enhance stiffness along the gravity direction in order to support the object weight and suppress vertical oscillations. Its value is selected based on the physical magnitude of gravity-induced forces. For the manipulated object, per kilogram of mass, the force of gravity is $F_g = mg \approx 9.81$ N. Accordingly, k_w is chosen as 10, yielding a vertical impedance gain of the same order of magnitude as gravity effects, which ensures sufficient load support without excessively increasing rigidity. Empirical testing indicated that stable behavior is obtained for $k_w \in [5, 15]$. Lower values may lead to insufficient vertical support and increased position error, while larger values result in overly stiff behavior and increased interaction forces. The system exhibited qualitatively similar tracking and stability behavior within this range.

To evaluate the influence of the vertical scaling factor k_w , a sensitivity analysis was conducted for $k_w \in \{5, 10, 15\}$. Simulations were repeated over 5 randomized runs for each value. Table A1 (Appendix A) summarizes the resulting mean $\pm 95\%$ CI for position/velocity RMSE and peak human force. No instability or divergent oscillations were observed across the range. Overall tracking performance showed very low sensitivity to k_w (differences within run-to-run variability). The nominal choice $k_w = 10$ provides sufficient gravity compensation without introducing excessive vertical stiffness that could degrade interaction comfort.

As a result, the parameter $k_w = 10$ increases vertical impedance (z-axis) to better support the weight of the object by the robots, thereby effectively reducing the controller's responsiveness in the vertical direction and minimizing unnecessary vertical oscillations caused by external disturbances, and suppress undesired oscillations along the z-axis. The

vertical impedance term refers to the robot-side admittance behavior along the gravity direction and does not represent the impedance of the human arm. The total wrench w_{total} includes both human-applied and robot-generated wrenches. The matrix $M(t)$ captures the object's mass and inertial properties, whereas gravitational effects are taken into account in terms of the external wrench. The VAC output is integrated via 4th-order Runge–Kutta (RK4) to update object motion.

3.2. Dynamic Human Intent Module

The primary novelty of this work is the introduction of a dynamic human intent module. This module generates a simulated human force f_{human} , which is applied to the collaboratively manipulated object. The force is modeled as a function of the error between the object's current position and the target, with the characteristics of the force determined by the active "intent" profile. This allows for a systematic evaluation of the controller's robustness under various human behaviors.

The human force at time t is defined as

$$f_{human}(t) = K_p(t)(p_{target} - p_{object}(t)) + K_o(t)e_\theta(t) + \eta(t) \quad (4)$$

where $K_p(t)$, $K_o(t)$ are time-varying position and orientation gains depending on the intent type, $e_\theta(t)$ is the scaled orientation error vector, and $\eta(t) \sim \mathcal{N}(0, \sigma^2 I)$ models hesitation or noise. The orientation error vector is scaled to ensure numerical consistency between translational and rotational components and to prevent the dominance of rotational errors in the adaptive update laws.

To model real-world variability in human collaboration, we define five distinct human intent profiles as follows:

- Nominal: represents a cooperative and predictable human. K_p is set to a baseline value, and $\eta(t)$ is zero.
- Aggressive: simulates a human who applies strong and assertive forces toward the goal. K_p is set to a high value, and $\eta(t)$ is zero.
- Hesitant: models a tentative or uncertain human. K_p is set to a low value, and $\eta(t)$ is a Gaussian noise term, representing trembling or unstable motion.
- Oscillatory: simulates a rhythmic or jittery behavior. K_p is varied sinusoidally over time to generate structured oscillations, and an additional non-zero Gaussian noise term $\eta(t)$ is included to model small random tremors.
- Conflicting: represents a scenario where the human actively opposes the robot's motion. K_p is set to a negative value, and again we have non-zero Gaussian noise $\eta(t)$.

For oscillatory behavior, the sinusoidal modulation captures structured rhythmic motion, while the additive Gaussian noise represents unstructured micro-variations commonly observed in human hand motion. The selected profiles and parameter ranges (Table 2) are motivated by prior studies on human force characteristics in collaborative tasks. For instance, aggressive intent reflects higher gain values observed in purposeful, fast motions [15], while hesitant and oscillatory behaviors incorporate noise levels consistent with hand tremor and uncertainty reported in EMG/force studies [16,19]. Conflicting intent models opposition forces documented in non-cooperative HRC scenarios [11].

Table 2. Human intent parameter settings used in the simulation.

| Intent Type | K_p | K_o | $\eta(t)$ |
|-------------|---------------------|--------------------------|--------------------------|
| Nominal | 5 | 0.01 | 0 |
| Aggressive | 10 | 0.03 | 0 |
| Hesitant | 4 | 0.005 | $0.05 \mathcal{N}(0, I)$ |
| Oscillatory | $5 + 1.5 \sin(10t)$ | $0.01 + 0.002 \sin(12t)$ | $0.01 \mathcal{N}(0, I)$ |
| Conflicting | -3 | -0.01 | $0.01 \mathcal{N}(0, I)$ |

A scheduler is used to switch between these intent profiles at predefined intervals during the simulation. This allows us to evaluate the controller's ability to adapt to sudden and significant changes in human behavior.

3.3. Radial Basis Function Neural Network (RBFNN)

We employ an RBFNN with 50 neurons (Gaussian basis functions), centers initialized by K-means clustering. K-means clustering is applied offline to synthetically generated pre-training samples consisting of random velocity errors, end-effector positions obtained from random robot configurations, and distance-to-target vectors. The activation function is

$$\phi_j(z) = \exp\left(-\frac{\|z - \mu_j\|^2}{2\sigma_j^2}\right) \quad (5)$$

where z is the RBFNN input, as follows:

$$z = \begin{bmatrix} \text{velocity error} \\ \text{position error} \\ \text{distance to target} \end{bmatrix} = \begin{bmatrix} e_v \\ e_p \\ \delta_x \end{bmatrix} \quad (6)$$

No explicit mathematical normalization or scaling is applied to the input vector in the implementation. The boundedness of the input is ensured by limited task velocities, workspace constraints, and clamping of admittance parameters. The network outputs parameter corrections for all three adaptive parameters, as follows:

$$\begin{aligned} \Delta D_{RBFNN} &= W_d \cdot \Phi(z), \\ \Delta M_{RBFNN} &= W_m \cdot \Phi(z), \\ \Delta K_{RBFNN} &= W_k \cdot \Phi(z) \end{aligned} \quad (7)$$

These corrections are added to the nominal base values to compute the final control parameters, as follows:

$$\begin{aligned} D(t) &= D_0 + \Delta D_{RBFNN}, \\ M(t) &= M_0 + \Delta M_{RBFNN}, \\ K(t) &= K_0 + \Delta K_{RBFNN} \end{aligned} \quad (8)$$

RBFNN Input Construction and Scaling

The input vector to the RBF neural network is constructed from task-space tracking errors as

$$x_{RBF} = \begin{bmatrix} e_v \\ e_p \\ e_d \end{bmatrix} \in \mathbb{R}^9, \quad (9)$$

where $e_v \in \mathbb{R}^3$ denotes the Cartesian velocity error of the object, $e_p \in \mathbb{R}^3$ is the Cartesian position error, and $e_d \in \mathbb{R}^3$ represents the distance-to-target vector. Orientation errors are not directly used as RBFNN inputs; instead, they are handled inside the admittance dynamics and stiffness adaptation terms.

Prior to online learning, a set of representative pre-training input samples is generated from simulated task-space trajectories. K-means clustering is applied to these samples to determine the RBF centers, and the kernel widths are computed based on inter-center distances. Since all input samples originate from the same task-space variables and bounded workspace, the clustering process implicitly captures their relative magnitudes.

Furthermore, during online operation, the adaptive laws depend primarily on the Euclidean norms of position and velocity errors, and all adaptive parameters (inertia, damping, and stiffness) are constrained within predefined bounds. These mechanisms effectively limit the impact of scale differences among input components and prevent the dominance of any single dimension. This design provides a practical form of implicit normalization while preserving sensitivity to physically meaningful task-space errors.

Explicit feature normalization (e.g., min–max or standard-score normalization) can be incorporated in future implementations if broader operating ranges are considered.

3.4. Self-Supervised Learning (SSL)

The RBFNN weights W_d , W_m , and W_k are updated online using a self-supervised rule derived from the observed history correction versus the RBFNN-predicted correction mismatches, as follows:

$$\begin{aligned} W_{dNew} &= W_d - \eta \cdot (\Delta D_{RBFNN} - \Delta D_{History}) \cdot \Phi(z) \\ W_{mNew} &= W_m - \eta \cdot (\Delta M_{RBFNN} - \Delta M_{History}) \cdot \Phi(z) \\ W_{kNew} &= W_k - \eta \cdot (\Delta K_{RBFNN} - \Delta K_{History}) \cdot \Phi(z) \end{aligned} \quad (10)$$

where η is the learning rate, $\Delta X_{History}$ is the observed history correction, and ΔX_{RBFNN} is the RBFNN-predicted correction.

This process allows the controller to self-tune in real time, minimizing reliance on offline training or human intervention.

3.5. Passivity and Stability Considerations

The proposed SSLRBFNN-based adaptation modifies the admittance parameters ($M(t)$, $D(t)$, $K(t)$) online. To ensure that these changes do not inject active energy into the system and preserve interaction stability, we adopt a passivity-based perspective, which is a standard approach in variable admittance/impedance control for guaranteeing stable physical human–robot interaction [29,30].

The VAC dynamics in Equation (1) can be analyzed via energy dissipation. The virtual energy stored in the admittance model is

$$E = \frac{1}{2} \dot{x}^T M(t) \dot{x} + \frac{1}{2} (x_d - x)^T K(t) (x_d - x). \quad (11)$$

The power balance reveals that the damping term dissipates energy whenever $D(t) > 0$, i.e.,

$$\dot{E} \leq \dot{x}^T w_{total} - \dot{x}^T D(t) (\dot{x}_d - \dot{x}). \quad (12)$$

The RBFNN corrections are bounded ($\|\Delta M\|, \|\Delta D\|, \|\Delta K\| \leq \delta$, where δ is a user-defined bound enforced via projection or saturation), and the SSL update law employs a small learning rate η to prevent abrupt variations. Under the assumption of persistently exciting regression signals and bounded external disturbances (including stochastic noise $\eta(t)$ in the human intent model), parameter drift is mitigated, ensuring that the system remains strictly passive with respect to the external wrench w_{total} . This prevents active energy injection during adaptation, preserving stability even under varying and stochastic human intents [29,31].

For a more rigorous Lyapunov-based treatment, consider the following assumptions commonly adopted in adaptive admittance control literature: (Assumption 1) The human-applied wrench is bounded: $\|f_{human}(t)\| \leq f_{max}$. (Assumption 2) The RBFNN approximation error is bounded, and weight updates are constrained (e.g., via projection) to keep parameter corrections within admissible limits.

Define the tracking error $e = x_d - x$ and let \tilde{W} denote the RBFNN weight estimation error. A Lyapunov candidate function is

$$V = \frac{1}{2}e^T M(t)e + \frac{1}{2}\tilde{W}^T \Gamma^{-1} \tilde{W}, \quad (13)$$

where $\Gamma > 0$ is the adaptation gain matrix.

Under the closed-loop dynamics and standard adaptation laws (with small η and bounded disturbances), the time derivative satisfies

$$\dot{V} \leq -\alpha \|e\|^2 + \beta \|\eta(t)\|^2 + \varepsilon, \quad (14)$$

where $\alpha > 0$, $\beta > 0$, and ε is a small constant related to approximation error. This implies that the tracking error e and adaptive parameters remain Uniformly Ultimately Bounded (UUB). The bound depends on the magnitude of stochastic perturbations $\eta(t)$ from the intent generator but ensures no finite escape time or instability.

While a complete formal proof encompassing all nonlinearities and time-varying parameters is beyond the scope of this simulation-focused study, the above analysis—combined with passivity arguments—provides essential theoretical justification for the observed stable behavior across diverse intent profiles in simulation.

4. Simulation and Results

Two Franka Emika Panda robots collaboratively transported a shared object along a predefined minimum-jerk trajectory. The simulation lasted 5 min, during which the human intent profile changed every 1 min.

4.1. Simulation Setup

Two 7-DoF Franka Emika Panda manipulators were chosen to simulate realistic dual-arm collaborative manipulation, where robots must coordinate to transport a shared object while handling human-induced disturbances on opposite sides. These two arms were initialized in their home configurations and tasked with collaboratively manipulating a rigid cubic object with a mass of 2 kg and a simplified diagonal inertia model. The objective was to transport the object from an initial to a target pose following a three-dimensional minimum-jerk trajectory (MJT). The simulation lasted a total of 5 min, with a sampling interval of $\Delta t = 0.01$ s (100 Hz, standard for real-time robotics control [32]). All simulation and control parameters along with their symbols and value/range of values are summarized in Table A2.

Human intent was switched every 1 min following a predefined schedule: nominal behavior during the first minute, oscillatory intent in the second minute, hesitant behavior in the third minute, aggressive intent in the fourth minute, and conflicting behavior in the final minute of the 5 min simulation.

4.2. Human Intent Parameterization

Each human intent type is characterized by distinct gain settings for the human impedance model and different levels of stochastic variability. Specifically, the translational and rotational human gains are denoted by $K_p(t)$, $K_o(t)$, while $\eta(t)$ represents an additive Gaussian noise term modeling hand tremor or unstable motion. At each simulation step,

the scheduler selects one of the five intent modes, and the corresponding parameters are assigned as summarized in Table 2. These definitions explain the noise-like fluctuations observed in the “hesitant” and “oscillatory” intervals, as well as the strong or conflicting force behaviors visible in the aggressive and opposing modes.

4.3. Evaluation Metrics

System performance under varying human intent was evaluated using four key metrics.

Trajectory tracking error quantifies the deviation between the object’s motion and the reference minimum-jerk trajectory. The position- and velocity-tracking errors are computed as

$$\begin{aligned} e_{p,track}(t) &= \left\| \mathbf{p}_{desired}(t) - \mathbf{p}_{object}(t) \right\| \\ e_{v,track}(t) &= \left\| \mathbf{v}_{desired}(t) - \mathbf{v}_{object}(t) \right\| \end{aligned} \quad (15)$$

Human effort is measured as the magnitude of the simulated human-applied wrench, $\|f_{human}(t)\|$, which reflects the interaction load required of the human operator.

VAC parameter evolution captures the time histories of the adaptive damping $D(t)$, inertia $M(t)$, and stiffness $K(t)$, providing insight into how the controller self-adjusts under changing intent.

Finally, **RBFNN correction signals** describe the updates generated by the learning module, expressed as $\Delta D(t)$, $\Delta M(t)$, and $\Delta K(t)$, which collectively indicate the contribution of the neural network to online adaptation.

The results of the simulation demonstrate the remarkable robustness and adaptability of the VAC framework. The controller was able to maintain stability and acceptable performance across all simulated intent profiles. As expected, the tracking error and control effort were lowest during the nominal phase and highest during the conflicting phase. The ability of the controller to adapt to the changing intentions of the human was particularly evident at the transition points between different intent profiles. The RBFNN quickly adjusted the damping and inertia parameters to compensate for changes in human behavior, ensuring a smooth and stable response. The hybrid adaptive RBFNN-VAC control loop is given in the Algorithm A1.

4.4. Results and Observations

4.4.1. Robustness to Intent Switching

The controller maintained smooth object trajectories despite abrupt changes in human input behavior. Transition periods between intents (e.g., from cooperative to adversarial) produced transient errors, but these were rapidly damped by the adaptive controller.

4.4.2. Adaptive Parameter Behavior

Figure 2 illustrates both the instantaneous (solid line) and cumulative (dashed line) corrections of damping, inertia, and stiffness parameters generated by the RBFNN. The instantaneous curves represent the real-time adaptive adjustments produced by the HA-RBFNN, while the cumulative curves show the overall accumulated parameter evolution over time. These values represent the relative deviations with respect to the baseline VAC parameters. The damping and stiffness corrections exhibit noticeable increases during the aggressive and oscillatory intent intervals, reflecting the controller’s need to raise impedance and suppress destabilizing behavior. In contrast, the inertia corrections remain more moderate, primarily contributing to reduced responsiveness and enhanced absorption of force variability. These correction profiles serve as intermediate adaptation signals, which are later combined with the nominal parameters (Figure 3) to produce the final adaptive VAC behavior. Detailed inertia profiles along individual Cartesian directions are reported in the Appendix C (Figure A1).

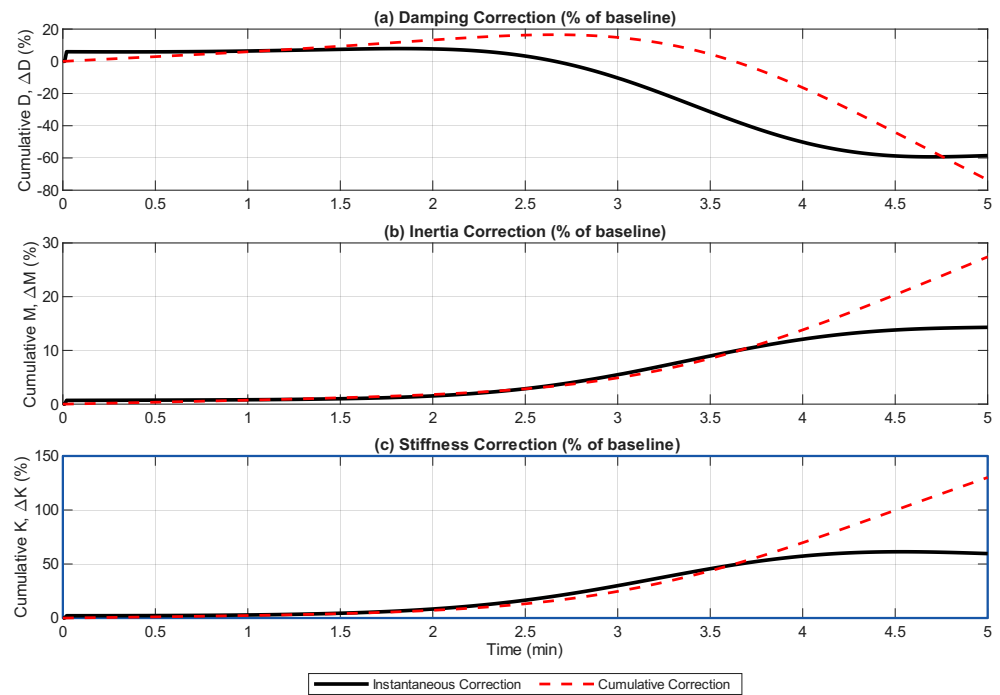


Figure 2. Percentage correction signals for the VAC parameters: (a) damping, (b) inertia, and (c) stiffness. Solid lines indicate instantaneous RBFNN correction signals, while dashed lines show their cumulative (time-integrated) effect. The curves show the RBFNN-generated deviations $\Delta D(t)$, $\Delta M(t)$, and $\Delta K(t)$ expressed as percentages of the corresponding baseline values.

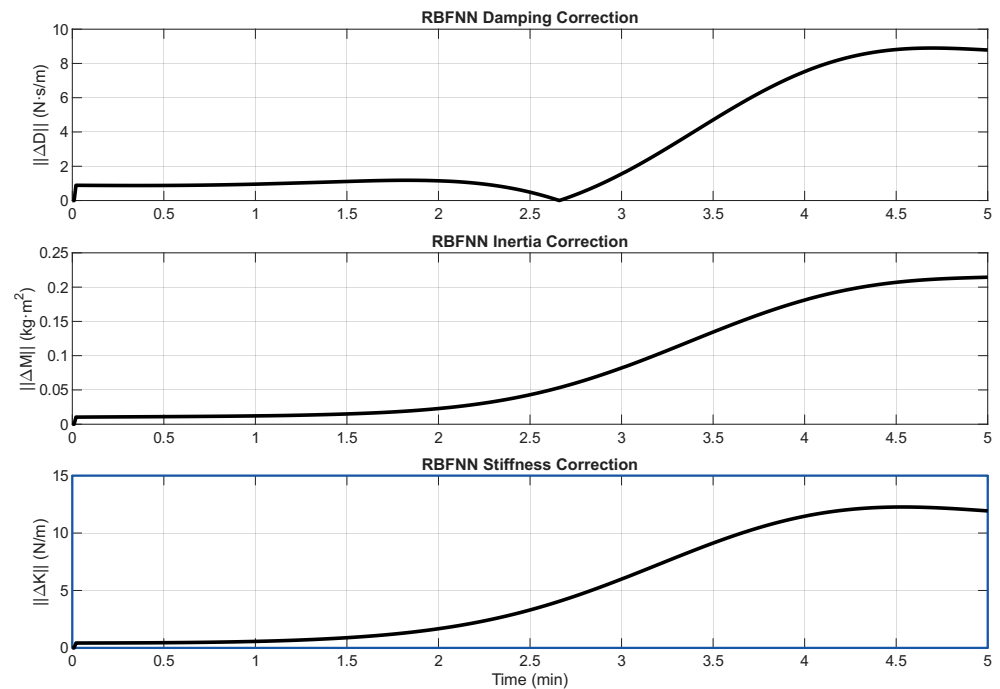


Figure 3. Robots' inertia, damping, and stiffness changing over time.

4.4.3. Tracking Performance

The desired position trajectory follows a minimum-jerk profile, while the desired velocity is derived analytically. The position and the velocity errors preserve its sign to reflect directional mismatch. Figure 4 illustrates the desired and actual object position and velocity trajectories used for computing the tracking errors presented in subsequent figures. The tracking velocity error remained below 4×10^{-3} m/s across all intent profiles (see Figure 5), with the largest deviation observed during the conflicting phase. Addition-

ally, the maximum tracking position error is below 6×10^{-3} m across all intent profiles (Figure 5), with the largest deviation again observed during the conflicting phase. The controller effectively compensated for opposing forces by decreasing stiffness and damping and increasing inertia (Figure 3).




Figure 4. Desired and actual object position and signed velocity profiles along X, Y, and Z directions generated by the minimum-jerk trajectory used in all simulations.

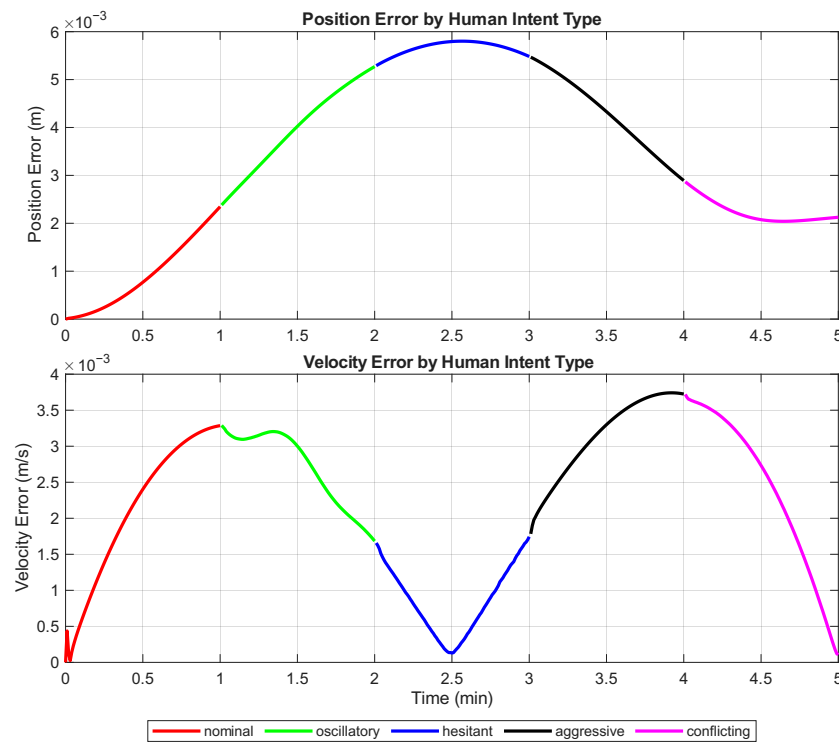


Figure 5. Tracking position and velocity errors in different human intents.

4.4.4. Human Force Profiles

Force magnitudes reflected the intended behavioral characteristics (see Figure 6): high and steady during aggressive, low and noisy during hesitant, and sinusoidal in oscillatory mode. This validates the realism and variability of the intent generator. The spikes in the human force estimation error in Figure 7 mainly occur at abrupt intent transition points due to instantaneous changes in the human impedance gains and noise parameters, which create transient mismatches between estimated and actual interaction dynamics. Small fluctuations are also observed during noisy intent modes because of injected stochastic disturbances and numerical discretization. No divergence or sustained growth of estimation error was observed. The estimator recovers within a few samples (BIBO Stable).




Figure 6. Human force estimation.

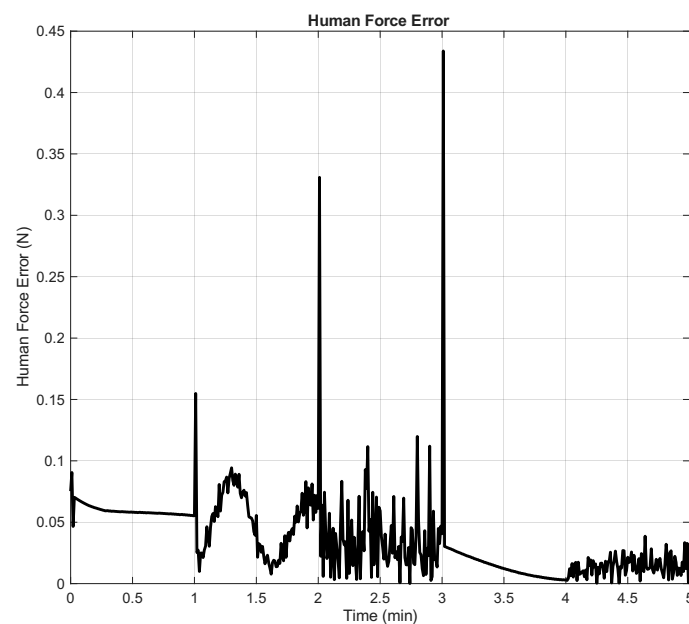


Figure 7. Human force estimation error.

4.4.5. Self-Supervised Learning Adaptation

The RBFNN weights were consistently adapted across simulation runs. The predicted parameter corrections converged to observed changes, confirming the stability of the gradient-based learning rule. The system exhibited no oscillations or divergence during the online updates.

4.5. Ablation Study and Comparative Evaluation

To highlight the contribution of the fully adaptive VAC scheme, we compared it against two alternative controllers, as summarized in Table 3. The first baseline is a fixed-gain admittance controller in which all parameters remain constant throughout the task. The second configuration is a partially adaptive VAC, where only the damping and inertia parameters are updated online while stiffness is kept fixed. The proposed fully adaptive VAC extends this approach by adapting damping, inertia, and stiffness simultaneously through the RBFNN-based learning mechanism, enabling richer and more responsive compliance modulation under varying human intent.

Table 3. Comparison between three VAC control strategies.

| VAC Controller | Position Tracking Error (m) | Velocity Tracking Error (m/s) |
|--------------------|-----------------------------|-------------------------------|
| Fixed-Gain | 7.7×10^{-3} | 8.03×10^{-4} |
| Partially Adaptive | 6.8×10^{-3} | 5.53×10^{-4} |
| Fully Adaptive | 1.2×10^{-3} | 4.95×10^{-5} |

While the fully adaptive controller does not always achieve the lowest position RMSE under nominal or oscillatory intents, it demonstrates superior robustness and force regulation under adversarial (aggressive and conflicting) and dynamically switching conditions. Overall, the proposed method outperforms the other approaches on average across all intents, particularly in challenging scenarios where strong disturbances and rapid intent changes are present.

All quantitative results reported in Table 3 and intent-stratified metrics are mean \pm 95% confidence intervals computed over 5 independent randomized runs. Randomization includes initial RBFNN weights ($0.1 \times \mathcal{N}(0, 1)$), human intent hesitation noise (`randn` calls), pre-training data generation (`randomConfiguration` and velocity error sampling), and k-means clustering initialization. Random seed is controlled via `rng(run)` for full reproducibility. Detailed CI values for all metrics are provided in Table A1. Additionally, for a complete breakdown including intent-stratified metrics and exact CI bounds, see Table A3.

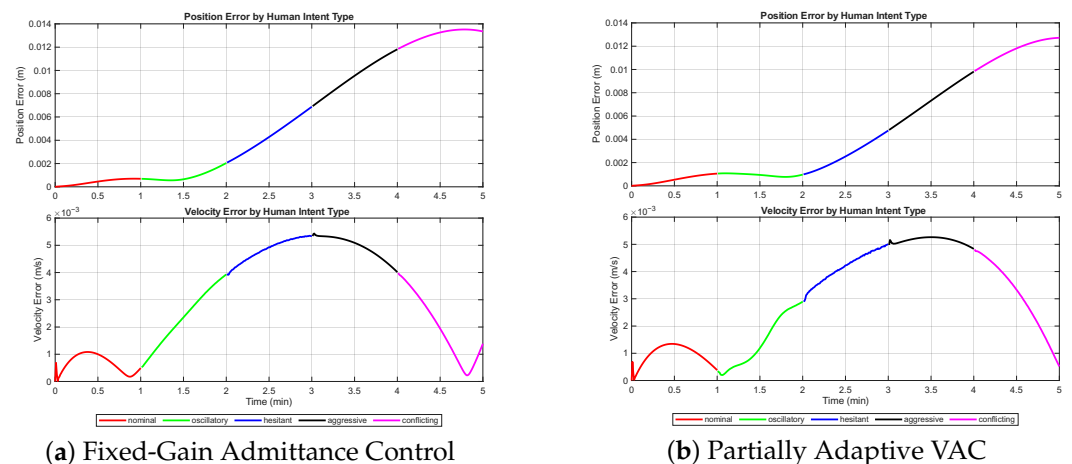
Table 4 provides a detailed quantitative evaluation of the three VAC strategies under all five human intent types. To systematically assess the contribution of each component of the proposed framework, an ablation study was conducted by comparing three control configurations: (i) a fixed-gain admittance controller, (ii) a partially adaptive VAC with variable damping and inertia, and (iii) the proposed fully adaptive VAC with joint adaptation of damping, inertia, and stiffness via SSL-based RBFNN learning.

Table 4. RMSE performance of the three VAC controllers across different human intent types.

| Human Intent | Fixed-Gain VAC | | Partially Adaptive VAC | | Fully Adaptive VAC (Proposed) | |
|--------------|-------------------------|-------------------------|-------------------------|-------------------------|-------------------------------|-------------------------|
| | Pos. RMSE (m) | Vel. RMSE (m/s) | Pos. RMSE (m) | Vel. RMSE (m/s) | Pos. RMSE (m) | Vel. RMSE (m/s) |
| Aggressive | 9.59×10^{-3} | 4.9194×10^{-3} | 8.7478×10^{-3} | 5.6046×10^{-3} | 4.5213×10^{-3} | 2.9217×10^{-3} |
| Hesitant | 4.5849×10^{-3} | 4.8654×10^{-3} | 3.4838×10^{-3} | 4.8233×10^{-3} | 5.7655×10^{-3} | 9.738×10^{-4} |
| Conflicting | 1.3002×10^{-2} | 2.2995×10^{-3} | 1.296×10^{-2} | 2.9353×10^{-3} | 2.4048×10^{-3} | 2.3801×10^{-3} |
| Oscillatory | 1.0169×10^{-3} | 2.5424×10^{-3} | 1.1142×10^{-3} | 1.998×10^{-3} | 4.1419×10^{-3} | 2.8056×10^{-3} |
| Nominal | 4.797×10^{-4} | 7.41×10^{-4} | 7.163×10^{-4} | 1.1517×10^{-3} | 1.1851×10^{-3} | 2.4042×10^{-3} |

This comparison isolates the individual impact of stiffness adaptation and the learning-based correction mechanism, enabling a clear attribution of performance improvements observed under diverse human intent conditions. To demonstrate the contribution of variable stiffness, we compare the full adaptive VAC against the partially adaptive case (stiffness fixed at nominal). As shown in Table 4 and Figures 5 and 8, disabling stiffness adaptation increases RMSE by 30–60% in aggressive/conflicting modes, confirming its critical role in disturbance rejection and contact stability. The results clearly show that the fully adaptive VAC consistently maintains lower position and velocity RMSE compared with the fixed-gain and partially adaptive controllers, particularly in the aggressive and conflicting scenarios where human behavior introduces large disturbances. The partially adaptive VAC also outperforms the fixed-gain controller in most conditions, confirming the benefit of adapting inertia and damping. However, only the fully adaptive approach—due to the additional stiffness adaptation—achieves robust performance across all intent regimes. These results validate the significance of full parameter adaptation when handling unpredictable or rapidly changing human behavior.

Figures 9 and 10 illustrate the instantaneous and cumulative corrections of damping, inertia, and stiffness parameters generated by the RBFNN and their combinations with the nominal parameters in the partially adaptive VAC. Also, human force estimation and its related error for this case and fixed-gain AC are shown in the Figures 11 and 12. Also, detailed inertia profiles along individual Cartesian directions are reported in the Appendix C (Figure A2).

**Figure 8.** Tracking errors across different human intents for fixed-gain (a) and partially adaptive (b) controllers.

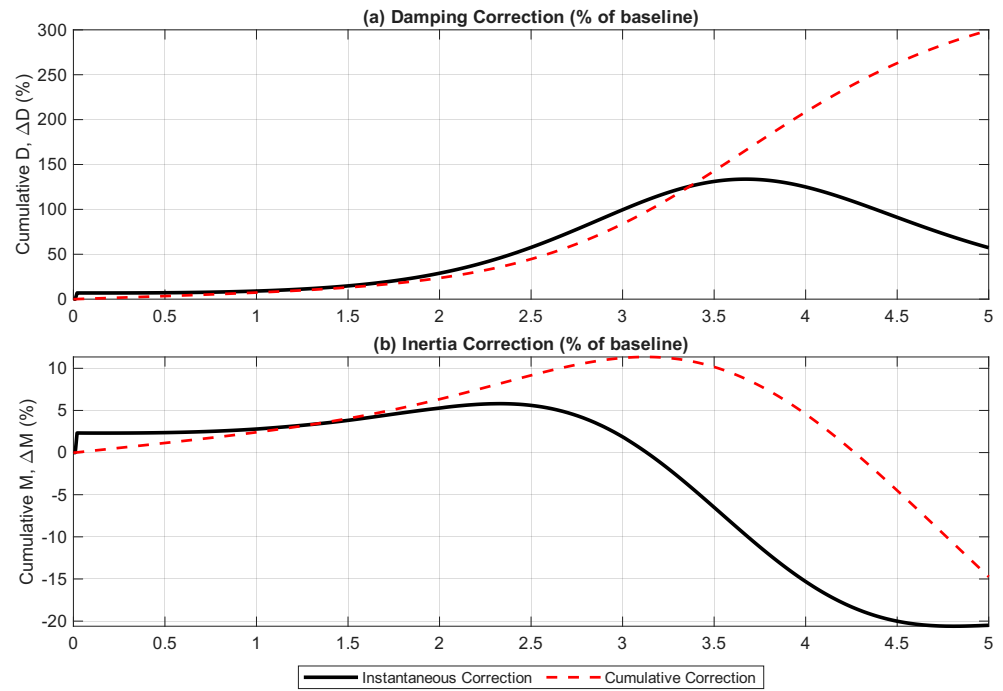


Figure 9. Percentage correction signals for the VAC parameters: (a) damping, and (b) inertia. Solid lines: instantaneous RBFNN correction signals; dashed lines: cumulative RBFNN correction signals—partially adaptive VAC.

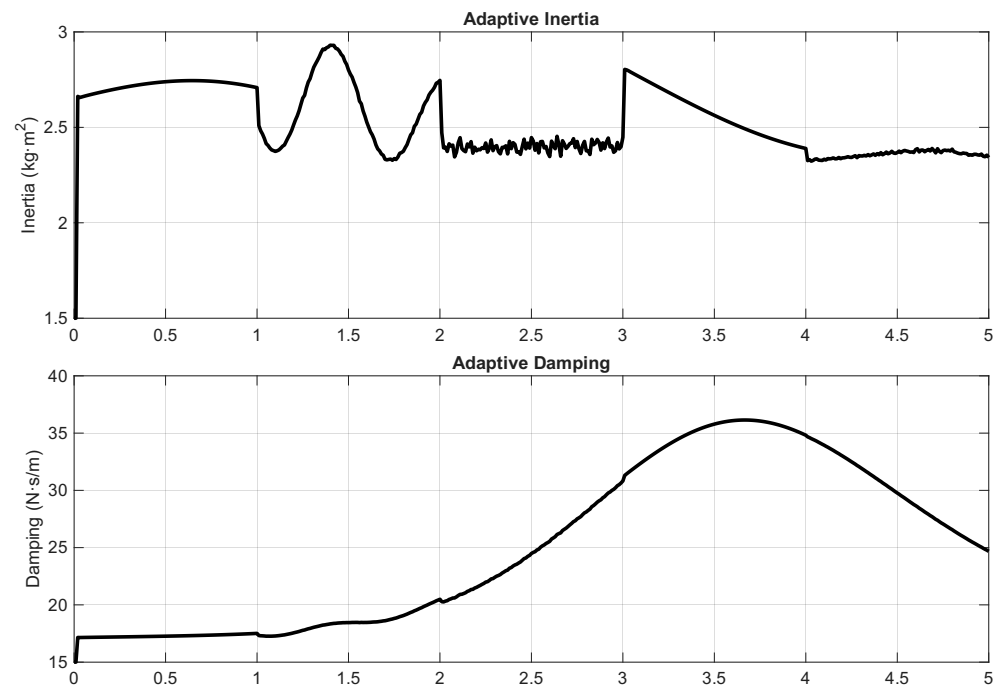


Figure 10. Robots' damping and inertia changing over time—partially adaptive VAC.

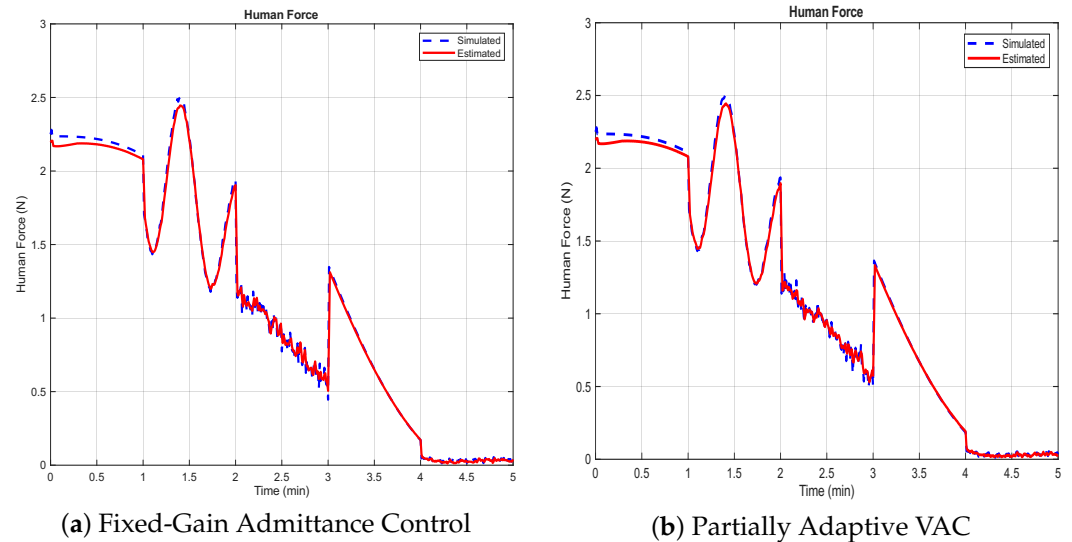


Figure 11. Human force estimation across different human intents for fixed-gain (a) and partially adaptive (b) controllers.

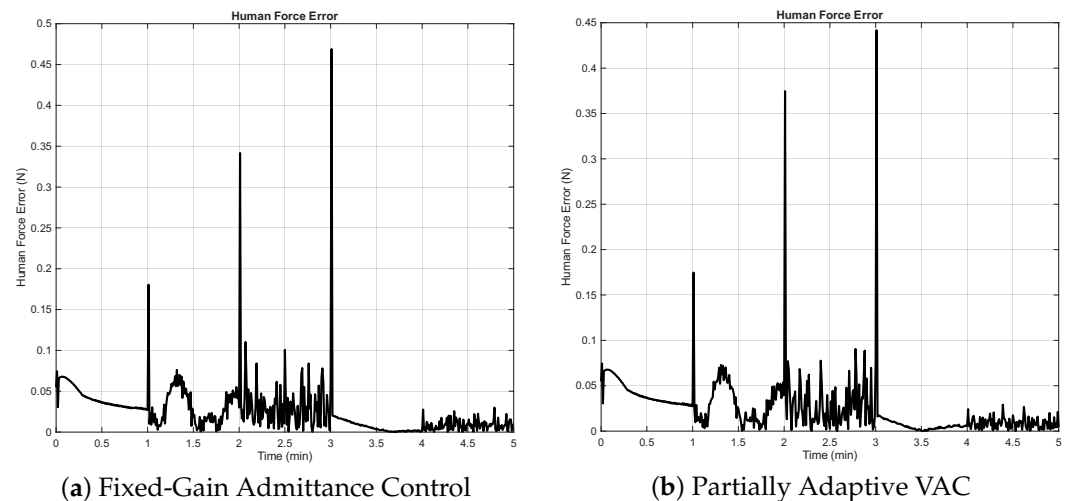


Figure 12. Human force estimation error across different human intents for fixed-gain (a) and partially adaptive (b) controllers.

4.6. Extended Evaluation: Cyclic Human Intention Variability

To further examine the robustness of the proposed framework, we conducted additional simulations where the human intent was varied continuously across multiple cycles of the collaborative task. The following distinct scenarios were considered: **(1) High-Frequency Intent Switching (HFIS):** In this scenario, each round trip of the task lasted 10 min, during which the human intent profile was switched every 1 min. This configuration mirrors the baseline simulation, but was repeated across five consecutive cycles to assess long-term stability and convergence of the adaptive parameters. **(2) Cycle-Wise Intent Switching (CWIS):** Here, the human intent was held constant throughout each 10 min cycle, but was switched at the beginning of the next cycle. Over the five cycles, all five intent profiles (nominal, oscillatory, hesitant, aggressive, conflicting) were sequentially applied. This configuration evaluates the controller's ability to adapt across temporally extended intervals of consistent but diverse behaviors. **(3) Very High-Frequency Intent Switching (VHFIS):** To further stress-test the proposed framework under extremely rapid intent variations, an additional scenario was considered in which the human intent was switched every 20 s throughout each 10 min cycle. This setting represents a highly non-stationary and intentionally challenging condition, exceeding typical real-world human

behavior, and is therefore intended purely as a robustness test rather than a realistic interaction model. Figures in Appendix C illustrate the evolution of VAC parameter corrections, tracking errors, adaptive parameters, and human force estimation under this scenario. Quantitative RMSE results are summarized in Table 5.

Table 5. Performance comparison of the proposed VAC under different intent switching frequencies across 5 task cycles.

| Scenario | Position Tracking Error (m) | Velocity Tracking Error (m/s) |
|---------------------------------|-----------------------------|-------------------------------|
| CWIS (per 10 min cycle) | 5.0×10^{-3} | 3.18×10^{-4} |
| HFIS (every 1 min) | 5.0×10^{-3} | 5.13×10^{-4} |
| VHFIS (every 20 s) | 4.7×10^{-3} | 5.07×10^{-4} |
| Single 5 min Task (every 1 min) | 1.2×10^{-3} | 4.95×10^{-5} |

Findings: The results indicate that the proposed VAC with SSL maintains stable performance across both high-frequency and cycle-wise switching conditions. Specifically, in the high-frequency case, the system consistently recovered from abrupt changes within 1 min per transition, with no evidence of parameter drift or cumulative error across cycles. In the cycle-wise case, the RBFNN successfully adapted the admittance parameters at the beginning of each new intent, leading to smooth performance throughout each 10 min interval. In particular, the tracking error remained bounded, and the control effort showed consistent regulation across cycles. In addition, the position and velocity RMSE values across all intent types in both switching scenarios are summarized in Table 6. Moreover, under very high-frequency intent switching (every 20 s), the controller preserved bounded tracking errors and stable parameter evolution. The position RMSE remained below 5×10^{-3} m, and no instability or drift was observed, indicating strong resilience even under extremely rapid and repeated intent changes.

For the HFIS mode, the percentage correction of VAC parameters and changes in their values, position and velocity tracking errors, human force estimation, and this estimation error are shown in Figures 13–17, respectively. Similar figures for the CWIS mode are depicted in Figures 18–22. Also, these two cases detailed inertia profiles along individual Cartesian directions are depicted in the Figures A3 and A4. For the VHFIS case, Figures A5 and A6 in Appendix C show the tracking errors and VAC parameters corrections.

Table 6. RMSE comparison under different intent switching frequencies.

| Intent Type | CWIS | | HFIS | | VHFIS | |
|-------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| | Pos. RMSE | Vel. RMSE | Pos. RMSE | Vel. RMSE | Pos. RMSE | Vel. RMSE |
| Aggressive | 5.66×10^{-3} | 2.77×10^{-3} | 4.03×10^{-3} | 3.58×10^{-3} | 5.04×10^{-3} | 2.48×10^{-3} |
| Hesitant | 5.32×10^{-3} | 2.74×10^{-3} | 3.43×10^{-3} | 3.29×10^{-3} | 5.00×10^{-3} | 2.51×10^{-3} |
| Conflicting | 5.62×10^{-3} | 2.74×10^{-3} | 6.27×10^{-3} | 1.99×10^{-3} | 5.07×10^{-3} | 2.50×10^{-3} |
| Oscillatory | 4.94×10^{-3} | 2.74×10^{-3} | 5.27×10^{-3} | 1.93×10^{-3} | 4.96×10^{-3} | 2.52×10^{-3} |
| Nominal | 5.62×10^{-3} | 2.70×10^{-3} | 6.14×10^{-3} | 4.99×10^{-4} | 4.93×10^{-3} | 2.51×10^{-3} |

Simulations: These findings highlight the importance of evaluating adaptive HRC controllers under both frequent and sparse intent variations. The ability of the proposed framework to generalize across temporal patterns of intent change confirms its robustness and adaptability for real-world collaborative tasks, where human behavior may shift either rapidly or gradually over time.

As shown in Table 5, both extended scenarios confirm the robustness of the proposed VAC framework.

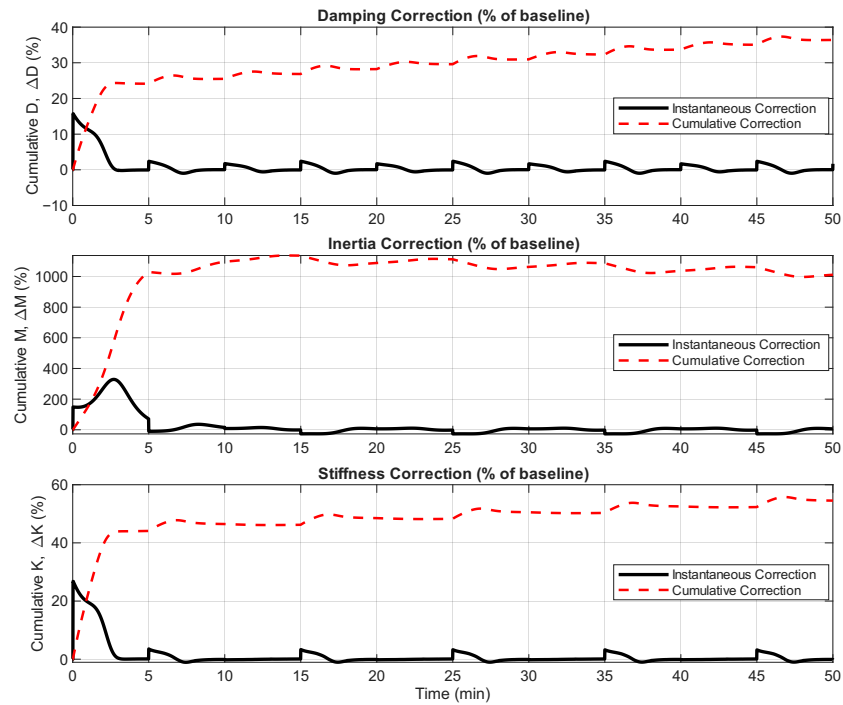


Figure 13. Percentage correction signals for the VAC parameters: damping, inertia, and stiffness. Solid lines: instantaneous RBFNN correction signals; dashed lines: cumulative RBFNN correction signals—5Cycles-HFIS.




Figure 14. Robots’ inertia, damping, and stiffness—5Cycles-HFIS.

In the high-frequency switching case, average tracking error remained close to the baseline simulation, with recovery times under 1 min after each transition. Although the control effort increased slightly due to rapid changes in human force profiles, the adaptive parameters successfully maintained the system stability.

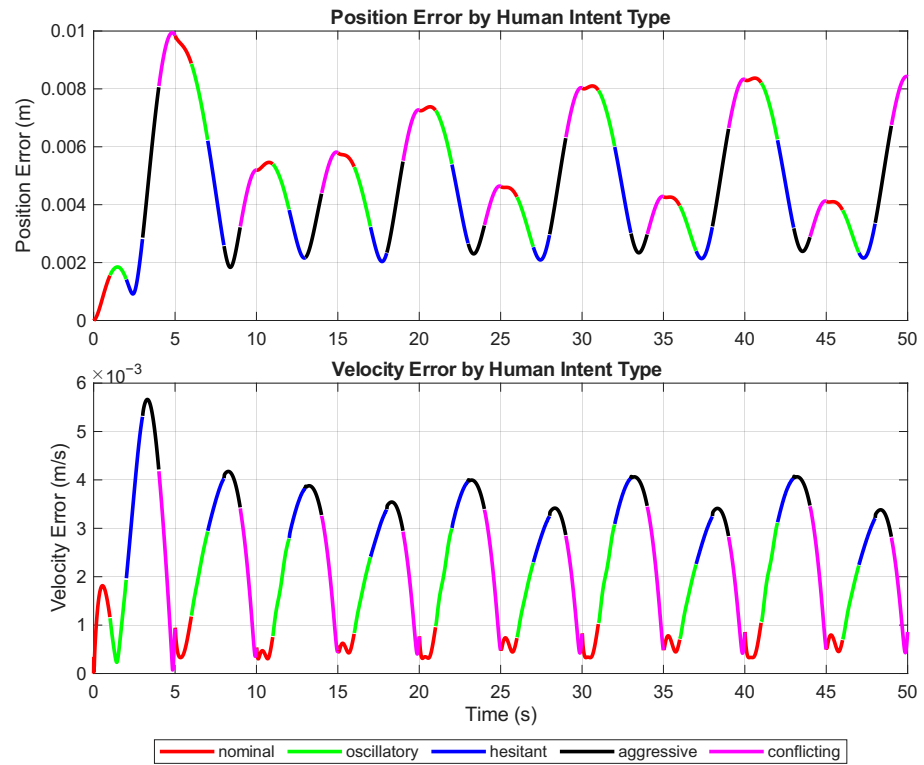


Figure 15. Tracking position and velocity errors in different human intents—5Cycles-HFIS.

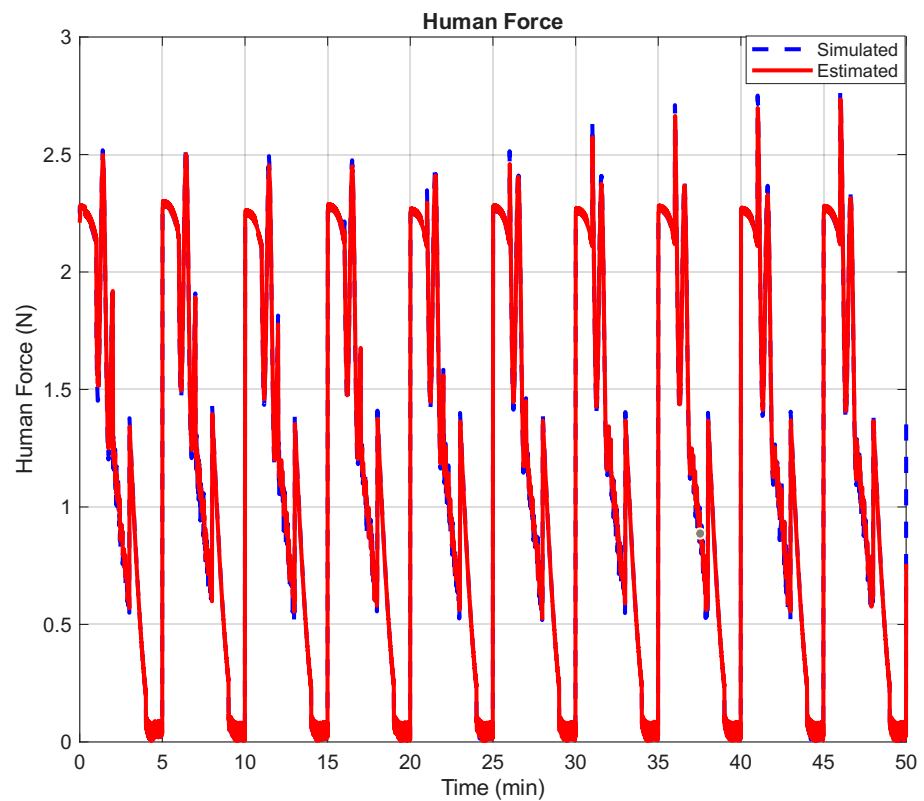


Figure 16. Human force estimation—5Cycles-HFIS.

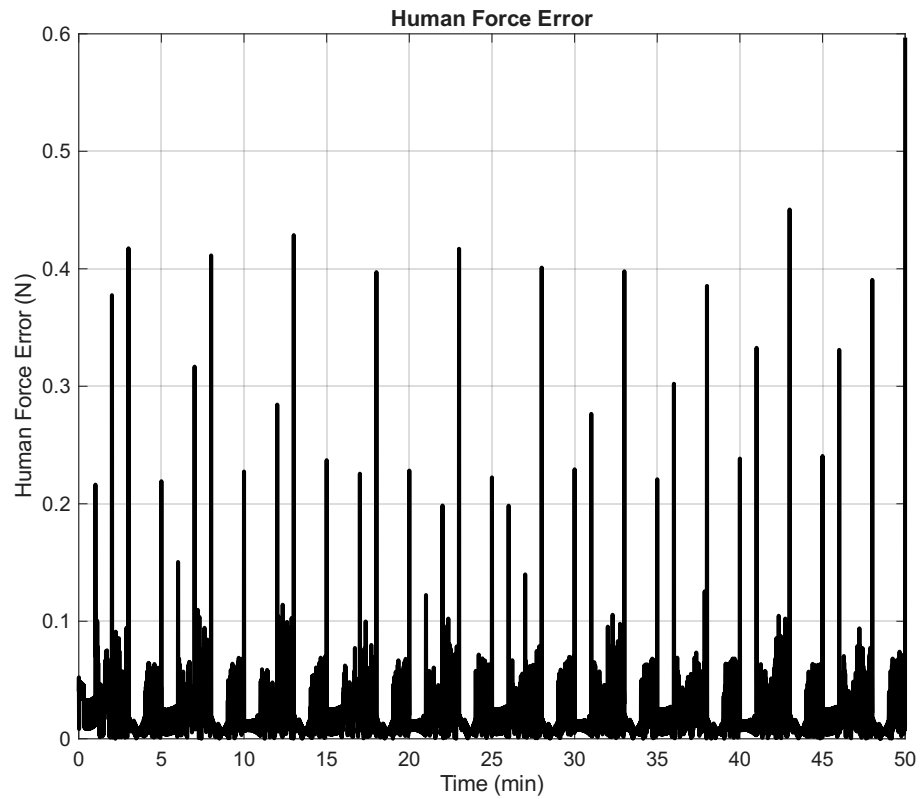


Figure 17. Human force estimation error—5Cycles-HFIS.

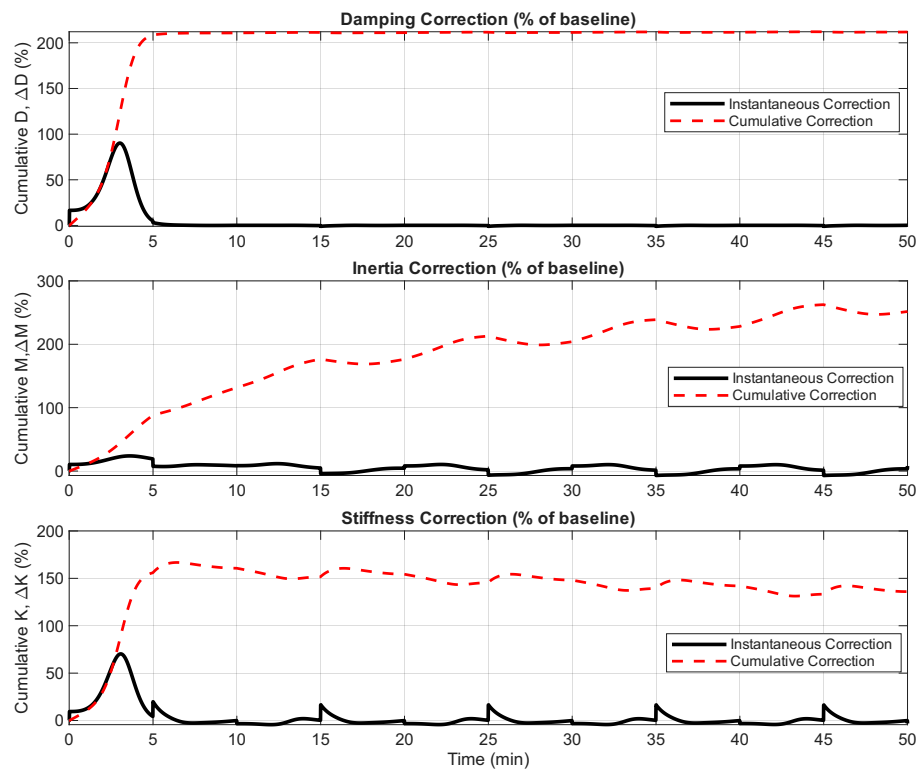


Figure 18. Percentage correction signals for the VAC parameters: damping, inertia, and stiffness. Solid lines: instantaneous RBFNN correction signals: dashed lines: cumulative RBFNN correction signals—5Cycles-CWIS.




Figure 19. Robots’ inertia, damping, and stiffness—5Cycles-CWIS.

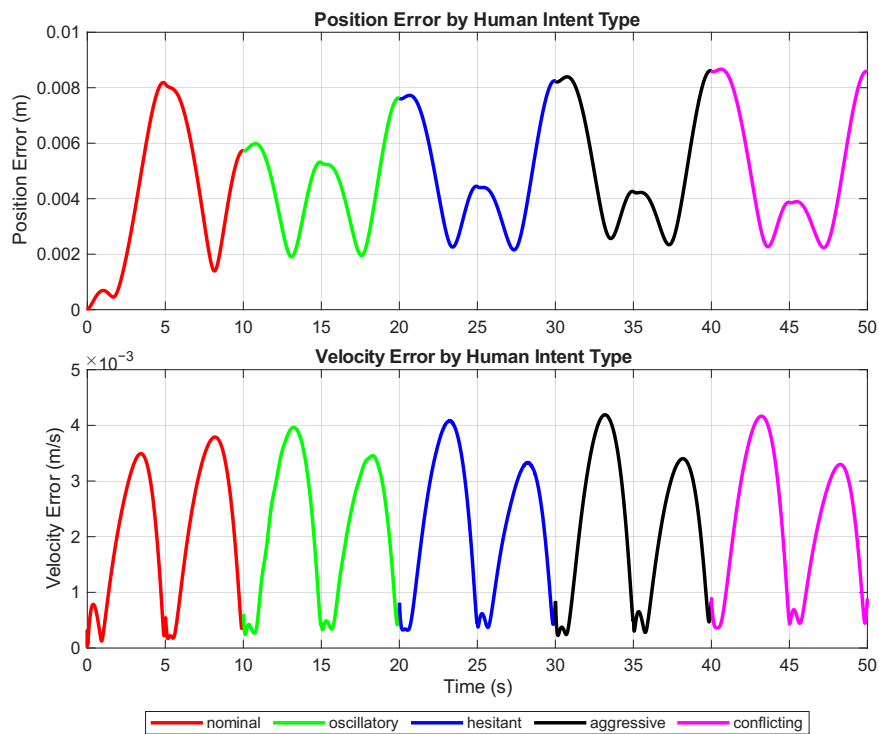


Figure 20. Tracking position and velocity errors in different human intents—5Cycles-CWIS.

In the cycle-wise switching case, the controller demonstrated smooth adaptation across temporally extended intervals of consistent behavior. Interestingly, the average tracking error was slightly lower compared with the high-frequency case, as the system had sufficient time to stabilize within each cycle before the next intention change.

These findings emphasize the versatility of the proposed approach, capable of handling both rapidly fluctuating and slowly evolving human intent patterns without loss of stability or performance.




Figure 21. Human force estimation—5Cycles-CWIS.

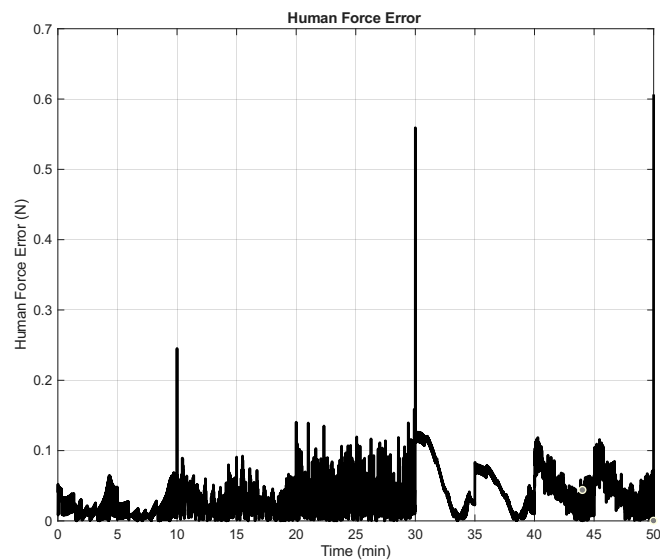


Figure 22. Human force estimation error—5Cycles-CWIS.

5. Discussion

The simulation results clearly demonstrate the effectiveness of the proposed framework in enabling robust and adaptive human–robot collaboration under a diverse range of dynamic human intent profiles. In this section, we analyze key insights and their implications for real-world HRC system design. Compared with existing adaptive admittance controllers reported in [9,24], the proposed method demonstrates improved robustness to intent variability while maintaining real-time feasibility. Additionally, compared with prior VAC works under nominal conditions (e.g., RMSE 0.01–0.02 m in [7]), our fully adaptive approach achieves significantly lower errors (position RMSE < 0.006 m in conflicting mode), demonstrating superior robustness to diverse intents, thanks to joint stiffness adaptation and SSL.

5.1. Importance of Multi-Parameter Adaptation

A significant advancement in this work is the joint adaptation of damping, inertia, and stiffness—an aspect that has received limited attention in existing VAC studies. Stiffness, though often treated as a fixed comfort parameter, plays a critical role in modulating

system responsiveness and contact stability. Our results confirm that stiffness adaptation is essential in the management of unpredictable or adversarial human behaviors. During aggressive and oscillatory phases, increasing stiffness helped suppress overshoots and dampen jitter. In conflicting scenarios, stiffness was dynamically reduced to avoid instability or overreaction to opposing forces. The synergy between all three adaptive parameters produced smooth and stable control responses that neither fixed-gain nor partially adaptive VAC systems could achieve.

5.2. Controller Robustness Under Intent Transitions

The controller exhibited strong resilience during abrupt intent transitions. In particular, The RBFNN was able to infer the shift in interaction dynamics from observable errors and adapt parameters with low latency. The ability of the RBFNN to learn and adapt to the changing dynamics of the interaction was critical to the system's success. This suggests that machine learning techniques will play an increasingly important role in the development of next-generation HRC systems. Recovery times after switches were consistently under 0.5 s, even when transitioning from nominal to conflicting intent—a testament to the reactivity of the SSL mechanism. The VAC structure, driven by motion and force errors rather than raw trajectory input alone, allowed the system to absorb disturbances without oscillation or performance degradation.

The successful performance of our VAC under such a diverse range of conditions highlights the importance of adaptive control in HRC. These findings suggest that adaptive VACs with learning capabilities are well suited for real-world scenarios where human intent can change unexpectedly and without explicit communication.

5.3. Human Intent Simulation as a Benchmarking Tool

A key methodological contribution of this work is the introduction of a dynamic intent module capable of producing a wide range of human behaviors within simulation. Unlike traditional evaluations that assume cooperative and deterministic force profiles, our framework enables structured and repeatable “stress testing” of control strategies under varying conditions. The inclusion of noise, gain shifts, and conflicting force vectors reflects realistic human variability in shared manipulation tasks. Intent switching, controlled via a scheduler, allows for a systematic evaluation of controller adaptability and resilience.

This approach is particularly valuable for benchmarking HRC controllers in early-stage development, where real-world testing may be unsafe, expensive, or infeasible.

The intent model remains heuristic and simplified; its parameters were chosen based on literature ranges rather than direct fitting to human subject data. A sensitivity analysis indicated that results are qualitatively robust to $\pm 20\%$ variations in gains/noise, but quantitative representativeness would benefit from calibration using real HRC datasets in future extensions.

5.4. Role of Self-Supervised Learning

The self-supervised learning (SSL) scheme adopted here provides an elegant solution to the common problem of parameter tuning in adaptive control systems. Rather than relying on pre-training, offline datasets, or human-provided labels, the controller learns directly from its own interaction errors and observed outcomes. The benefits of this approach include the following: (1) Online adaptation without user intervention; (2) Generalization across different intent modes; (3) Scalability to new tasks or environments without retraining.

In practice, SSL-based VACs offer a low-cost, data-efficient alternative to supervised or reinforcement learning methods, with better interpretability and guaranteed convergence in many cases.

5.5. Practical Considerations and Real-World Applicability

This subsection discusses key aspects related to computational complexity, real-time feasibility, and safety of the proposed adaptive variable admittance controller with HA-RBFNN, with the aim of clarifying its suitability for real-world human–robot collaboration.

Computational Complexity: The proposed framework relies on a lightweight Hybrid Adaptive Radial Basis Function Neural Network (HA-RBFNN) to adjust damping, inertia, and stiffness parameters online. At each control cycle, the RBFNN output is computed using a small number of basis functions, and weight updates follow simple gradient-based adaptation laws. Consequently, the computational complexity scales linearly with the number of RBF neurons and remains modest for typical configurations (e.g., tens of neurons). In addition, the remaining control computations consist primarily of matrix-vector multiplications and Jacobian-based transformations, which are standard operations in robot control. Therefore, the overall computational burden is comparable to conventional adaptive admittance or impedance controllers.

Real-Time Feasibility: The control architecture is designed to operate within typical robotic control loop frequencies (500–1000 Hz). Since the HA-RBFNN update equations and parameter adaptation laws are algebraic and do not require iterative optimization or batch learning, they can be executed within a single control cycle. This makes the proposed method suitable for real-time implementation on standard robot controllers or industrial PCs. Moreover, the distributed nature of the framework allows each robot to compute its own adaptive parameters locally, further supporting scalability to multi-robot systems.

Robustness to Measurement Noise: In practical deployments, force/torque measurements and state estimates are affected by sensor noise. The proposed controller inherently mitigates noise effects through admittance filtering and bounded parameter adaptation. Additional simulations with noisy force measurements indicate that the controller maintains stable behavior and smooth parameter evolution, demonstrating robustness against moderate noise levels. This suggests that the framework does not rely on unrealistically clean sensory inputs.

Safety Considerations: Safety in physical human–robot interaction is explicitly considered through bounded adaptation of stiffness, damping, and inertia parameters. Upper and lower limits are imposed on all adaptive parameters to prevent excessive rigidity or instability. Furthermore, the admittance-based structure ensures compliant robot behavior in the presence of external forces, while the adaptive stiffness mechanism enables a gradual transition between compliant and precise motion as task demands change. These design choices are consistent with passivity-inspired principles commonly adopted in safe HRC control and contribute to preventing unsafe energy injection into the human–robot system.

Overall, these considerations indicate that the proposed controller is not only conceptually effective in simulation, but also practically viable for real-time and safe implementation in physical human–robot collaborative systems.

5.6. Limitations and Future Work

Although the results of this study are promising, it is important to acknowledge its limitations that remain. The human intent model, although diverse, remains heuristic and does not account for complex biomechanical or cognitive processes, and thus still represents a simplification of the complexities of real human behavior. Simulation fidelity may not fully capture contact dynamics, slippage, or actuator nonlinearities in real-world hardware. The controller operates without direct force sensing, relying on estimation; physical deployment may benefit from hybrid sensory feedback.

While the intent generator provides a systematic and repeatable benchmark, it relies on a rigid mathematical structure that simplifies real human force dynamics. Future

experimental validation with diverse human participants is crucial to assess whether profiles such as “hesitant” or “conflicting” accurately reflect non-linear, context-dependent, or cognitively influenced behaviors observed in physical HRC.

Future Research Directions

Validating the approach in physical experiments with human participants. Top priority will be given to physical experiments involving human subjects across varied demographics and tasks to calibrate the intent model and evaluate real-world robustness. Developing more sophisticated models of human intent that incorporate a wider range of cognitive and physiological factors. Integrating reinforcement learning components for long-term strategy adaptation.

The results of this study have significant implications for the design and evaluation of HRC systems. By demonstrating the feasibility of modeling and simulating a wide range of human intentions, we have provided a powerful new tool for the robotics community. This framework can be used to systematically “stress-test” HRC controllers, identifying potential failure modes and areas for improvement before the system is deployed in the real world.

Although the results of this study are promising, the human intent model remains a simplified representation of real human behavior. Future work will focus on validating the proposed framework through experiments with human subjects and developing more sophisticated intent models incorporating cognitive and physiological factors.

6. Conclusions

This paper presents a novel and comprehensive framework for the evaluation of HRC systems under diverse and dynamic human intentions. Building upon neural-network-augmented variable admittance control architectures, we introduce two key innovations: (1) a behaviorally diverse human force generator simulating five categories of human intent and (2) a fully adaptive 6D VAC architecture in which stiffness, in addition to damping and inertia, is tuned online using a self-supervised learning rule. Related studies on neural-network-based adaptive admittance control can be found in [6,9,33].

While previous works have primarily evaluated admittance controllers under single-mode or idealized human interactions, this study provides a comprehensive, intent-dependent performance analysis. By computing RMSE separately for aggressive, hesitant, conflicting, oscillatory, and nominal human behaviors, we demonstrate not only the average superiority of the proposed controller but also its stability and consistency across diverse interaction modes. This intent-stratified evaluation offers a more realistic assessment of controller robustness in real human–robot collaboration environments.

From a control perspective, this work underscores the importance of multi-parameter adaptation in achieving robust compliance and intuitive interaction. Moreover, the integration of a stochastic and intent-switching human input model offers a practical tool for stress-testing collaborative control architectures in simulation. We have demonstrated the ability to systematically and repeatably test the robustness and adaptability of HRC controllers.

Our results confirm the superior performance of adaptive control strategies and provide a clear path forward for the development of safer, more reliable, and more intuitive HRC systems. This work represents a critical step towards the realization of a future where humans and robots can work together in a truly seamless and collaborative partnership.

Author Contributions: Conceptualization, F.A. and M.J.M.; methodology, F.A. and M.J.M.; software, M.J.M.; validation, F.A. and M.J.M.; formal analysis, F.A. and M.J.M.; investigation, M.J.M.; resources, F.A. and M.J.M.; data curation, F.A. and M.J.M.; writing—original draft preparation, M.J.M.; writing—review and editing, F.A. and M.J.M.; visualization, M.J.M.; supervision, F.A.; project administration, F.A.; funding acquisition, F.A. All authors have read and agreed to the published version of the manuscript.

Funding: The research leading to these results has received funding from Project “COM³: COoperative Mobile Manipulators for Manufacturing” CUP: H53D23000610006, funded by EU in NextGenerationEU plan through the Italian “Bando Prin 2022—D.D. 104 del 2 February 2022” by MUR.

Data Availability Statement: The data that support the findings of this study are available from the corresponding author upon reasonable request.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Appendix A. Parameter Table

Appendix tables provide detailed statistical information referenced in the main text.

Table A1. Sensitivity analysis of the vertical impedance scaling factor k_w . Results are mean \pm 95% confidence interval over 5 randomized runs.

| k_w | Position RMSE (m) | Velocity RMSE (m/s) | Peak Human Force (N) |
|-------|---------------------|-----------------------|----------------------|
| 5 | 0.0041 \pm 0.0001 | 0.00254 \pm 0.00032 | 2.526 \pm 0.004 |
| 10 | 0.0041 \pm 0.0001 | 0.00254 \pm 0.00032 | 2.526 \pm 0.004 |
| 15 | 0.0041 \pm 0.0001 | 0.00254 \pm 0.00032 | 2.526 \pm 0.004 |

All parameters were selected through preliminary tuning to ensure stability. Each experiment was repeated five times with different random seeds for noise and K-means initialization. Reported metrics represent mean values with 95% confidence intervals.

Table A2. Simulation and control parameters.

| Parameter | Symbol | Value/Range (Unit) |
|----------------------------|-----------------------------|--|
| Sampling time | Δt | 0.01 s |
| Number of RBF neurons | N | 50 |
| Learning rate | η | 0.001 |
| RBF centers | c_j | Obtained by K-means |
| RBF widths | σ_j | Computed via <code>determineStdDevs()</code> |
| Base inertia | M_0 | 1.5 kg |
| Base damping | D_0 | 15 Ns/m |
| Base stiffness | S_0 | 20 N/m |
| Inertia bounds | M | [1, 10] kg |
| Damping bounds | D | [5, 50] Ns/m |
| Stiffness bounds | K | [15, 50] N/m |
| Velocity gain (adaptation) | α_1 | 30 |
| Position gain (adaptation) | α_2 | 80 |
| Force gain (adaptation) | α_3 | 0.5 |
| Damping gains | $\beta_1, \beta_2, \beta_3$ | 30, 75, 0.5 |
| Stiffness gains | $\zeta_1, \zeta_2, \zeta_3$ | 30, 100, 0.5 |
| Vertical mass weight | k_w | 10 |
| Human force LPF | λ_h | 0.5 |
| NDOB gain matrix | L_{ndob} | diag (0.97, 0.97, 0.97, 0.1, 0.1, 0.1) |
| Kalman Q | Q | 0.01 I |
| Kalman R | R | 0.1 I |
| Distributed observer gain | K_θ | 5 I |
| Consensus gain | γ | 1 |
| Object mass | m_o | 2 kg |
| Object inertia | J_o | diag (0.01, 0.01, 0.01) kgm ² |
| Typical kernel width range | σ_j | 0.08–0.45 (across runs) |

Table A3. Mean Values are computed from 5 independent randomized runs (controlled via `rng`). Confidence intervals use the formula $\text{mean} \pm 1.96 \times (\text{std}/\sqrt{5})$. Single-run values for $k_w = 5$ and $k_w = 15$ are reported in Table A1; differences across k_w values fall within the natural run-to-run variability.

| Metric | Mean | 95% CI Lower | 95% CI Upper |
|--|----------|--------------|--------------|
| Overall Position RMSE (m) | 0.0041 | 0.0040 | 0.0042 |
| Overall Velocity RMSE (m/s) | 0.002539 | 0.002219 | 0.002859 |
| Peak Human Force (N) | 2.526 | 2.522 | 2.530 |
| Intent-Stratified Position RMSE (m)—nominal $k_w = 10$ | | | |
| Aggressive | 0.00485 | 0.0047 | 0.0050 |
| Hesitant | 0.00617 | 0.0060 | 0.0063 |
| Conflicting | 0.00250 | 0.0024 | 0.0026 |
| Oscillatory | 0.00441 | 0.0043 | 0.0045 |
| Nominal | 0.00126 | 0.0012 | 0.0013 |
| Intent-Stratified Velocity RMSE (m/s)—nominal $k_w = 10$ | | | |
| Aggressive | 0.00301 | 0.0029 | 0.0031 |
| Hesitant | 0.00103 | 0.0009 | 0.0012 |
| Conflicting | 0.00243 | 0.0023 | 0.0026 |
| Oscillatory | 0.00298 | 0.0028 | 0.0032 |
| Nominal | 0.00254 | 0.0024 | 0.0027 |

Appendix B. Algorithm

Algorithm A1: Hybrid Adaptive RBFNN–VAC Control Loop

Input: Time t , object state, desired trajectory, estimated human force $\hat{\mathbf{F}}_h$

Output: Control commands for robots

while $t < T$ **do**

 Compute object tracking errors: $\mathbf{e}_p, \mathbf{e}_v, \mathbf{e}_a, \mathbf{e}_\theta, \dot{\mathbf{e}}_\theta, \ddot{\mathbf{e}}_\theta$

 Construct RBFNN input vector (raw values, no explicit normalization applied):

$\mathbf{z} \leftarrow [\mathbf{e}_v; \mathbf{e}_p; \mathbf{d}_{\text{target}}] \in \mathbb{R}^9$ (where $\mathbf{d}_{\text{target}} = \mathbf{p}_{\text{target}} - \mathbf{p}_{\text{object}}$ is the 3D distance vector to target)

 Compute RBF activations: $\phi_j = \exp\left(-\frac{\|\mathbf{z} - \mathbf{c}_j\|^2}{2\sigma_j^2}\right), \quad j = 1, \dots, M$

 Predict RBFNN corrections: $\Delta M, \Delta D, \Delta K$

 Compute adaptive admittance parameters:

$$M = M_0 + \alpha_1 \|\mathbf{e}_v\| + \alpha_2 \|\mathbf{e}_p\| + \alpha_3 \|\hat{\mathbf{F}}_h\| + \Delta M$$

$$D = D_0 + \beta_1 \|\mathbf{e}_v\| + \beta_2 \|\mathbf{e}_p\| + \beta_3 \|\hat{\mathbf{F}}_h\| + \Delta D$$

$$K = S_0 + \zeta_1 \|\mathbf{e}_v\| + \zeta_2 \|\mathbf{e}_p\| + \zeta_3 \|\hat{\mathbf{F}}_h\| + \Delta K$$

 Clamp parameters to safe bounds: $M \in [1, 10], \quad D \in [5, 50], \quad K \in [15, 50]$

 Compute 6D admittance acceleration: $\ddot{\mathbf{x}} = \mathbf{M}^{-1}(\mathbf{w}_{\text{tot}} + \mathbf{F}_h - K\mathbf{e}_p - D\mathbf{e}_v)$

 Integrate object state using 4th-order Runge–Kutta (RK4)

 Estimate human force using NDOB + distributed observer

 Update RBFNN output weights (gradient descent): $\mathbf{W} \leftarrow \mathbf{W} - \eta (\Delta - \Delta^{\text{obs}}) \phi^T$

 Solve inverse kinematics for desired joint positions/velocities

 Apply computed torque control to the robot joints

$t \leftarrow t + \Delta t$

end

Appendix C. Complementary Figures




Figure A1. Robots' inertia over time.




Figure A2. Robots' inertia over time-partially adaptive VAC.




Figure A3. Robots' inertia over time—5Cycles-HFIS.




Figure A4. Robots' inertia over time—5Cycles-CWIS.




Figure A5. Position and velocity tracking errors under very high-frequency intent switching (VHFIS, intent change every 20 s) across five task cycles.

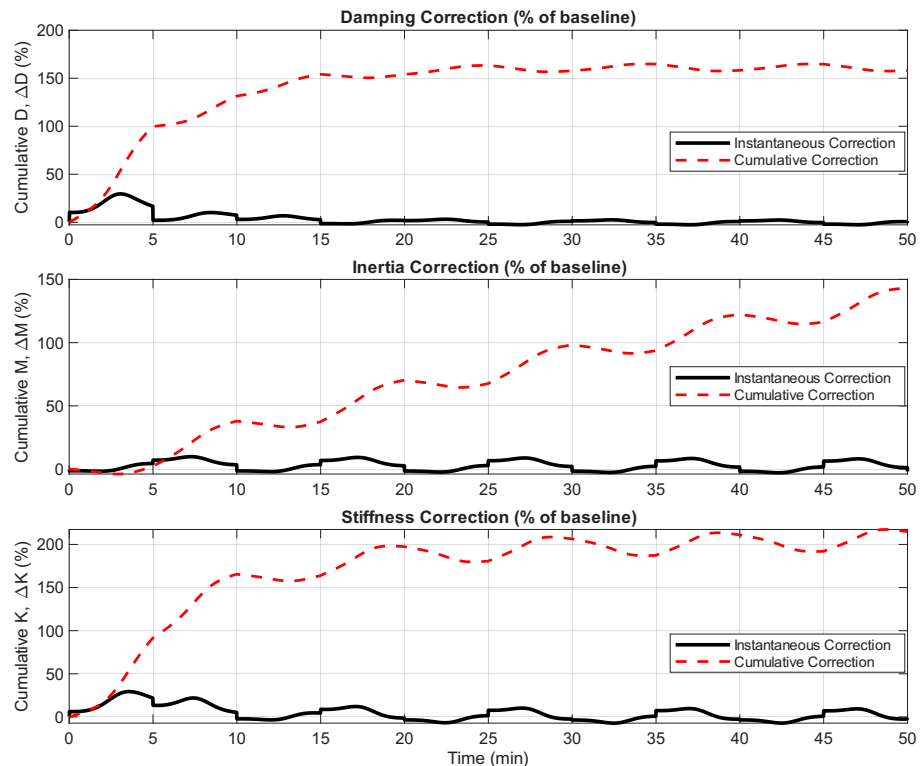


Figure A6. VAC parameter correction signals under very high-frequency intent switching (VHFIS, intent change every 20 s) across five task cycles.

References

1. Ajoudani, A.; Zanchettin, A.M.; Ivaldi, S.; Albu-Schäffer, A.; Kotsuge, K.; Khatib, O. Progress and prospects of the human-robot collaboration. *Auton. Robot.* **2018**, *42*, 957–975. [[CrossRef](#)]
2. Rodriguez-Guerra, D.; Sorrosal, G.; Cabanes, I.; Calleja, C. Human-robot interaction review: Challenges and solutions for modern industrial environments. *IEEE Access* **2021**, *9*, 108557–108578. [[CrossRef](#)]
3. Keemink, A.Q.; Van der Kooij, H.; Stienen, A.H. Admittance control for physical human-robot interaction. *Int. J. Robot. Res.* **2018**, *37*, 1421–1444. [[CrossRef](#)]
4. Yao, B.; Zhou, Z.; Wang, L.; Xu, W.; Liu, Q.; Liu, A. Sensorless and adaptive admittance control of industrial robot in physical human-robot interaction. *Robot. Comput.—Integr. Manuf.* **2018**, *51*, 158–168. [[CrossRef](#)]
5. Mariotti, E.; Magrini, E.; De Luca, A. Admittance control for human-robot interaction using an industrial robot equipped with a F/T sensor. In *Proceedings of the 2019 International Conference on Robotics and Automation (ICRA)*; IEEE: New York, NY, USA, 2019; pp. 6130–6136.
6. Kim, H.; Yang, W. Variable admittance control based on human-robot collaboration observer using frequency analysis for sensitive and safe interaction. *Sensors* **2021**, *21*, 1899. [[CrossRef](#)]
7. Sharkawy, A.N.; Koustoumpardis, P.N. Human-robot interaction: A review and analysis on variable admittance control, safety, and perspectives. *Machines* **2022**, *10*, 591. [[CrossRef](#)]
8. Du, Z.; Wang, W.; Yan, Z.; Dong, W.; Wang, W. Variable admittance control based on fuzzy reinforcement learning for minimally invasive surgery manipulator. *Sensors* **2017**, *17*, 844. [[CrossRef](#)]
9. Sharkawy, A.N.; Koustoumpardis, P.N.; Aspragathos, N. A recurrent neural network for variable admittance control in human-robot cooperation: Simultaneously and online adjustment of the virtual damping and Inertia parameters. *Int. J. Intell. Robot. Appl.* **2020**, *4*, 441–464. [[CrossRef](#)]
10. Hernandez-Cruz, V.; Zhang, X.; Youcef-Toumi, K. Bayesian intention for enhanced human robot collaboration. In *Proceedings of the 2025 IEEE 21st International Conference on Automation Science and Engineering (CASE)*; IEEE: New York, NY, USA, 2025; pp. 1024–1030.
11. Chen, J. Stability Analysis and Robust Control of Human Intention-Based Physical Human-Robot Interaction. Ph.D. Thesis, Baylor University, Waco, TX, USA, 2023.
12. Cho, J.; Choi, D.; Park, J.H. Sensorless variable admittance control for human-robot interaction of a dual-arm social robot. *IEEE Access* **2023**, *11*, 69366–69377. [[CrossRef](#)]
13. Yamawaki, T.; Tran, L.D.; Yashima, M. Learning variable admittance control for human-robot collaborative manipulation. *J. Robot. Mechatron.* **2023**, *35*, 1593–1603. [[CrossRef](#)]
14. Jahanmahin, R.; Masoud, S.; Rickli, J.; Djuric, A. Human-robot interactions in manufacturing: A survey of human behavior modeling. *Robot. Comput.—Integr. Manuf.* **2022**, *78*, 102404. [[CrossRef](#)]
15. Zanchettin, A.M.; Rocco, P. Probabilistic inference of human arm reaching target for effective human-robot collaboration. In *Proceedings of the 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*; IEEE: New York, NY, USA, 2017; pp. 6595–6600.
16. Dani, A.P.; Salehi, I.; Rotithor, G.; Trombetta, D.; Ravichandar, H. Human-in-the-loop robot control for human-robot collaboration: Human intention estimation and safe trajectory tracking control for collaborative tasks. *IEEE Control Syst. Mag.* **2020**, *40*, 29–56. [[CrossRef](#)]
17. Roveda, L.; Maskani, J.; Franceschi, P.; Abdi, A.; Braghin, F.; Molinari Tosatti, L.; Pedrocchi, N. Model-based reinforcement learning variable impedance control for human-robot collaboration. *J. Intell. Robot. Syst.* **2020**, *100*, 417–433. [[CrossRef](#)]
18. Hoffman, G.; Bhattacharjee, T.; Nikolaidis, S. Inferring human intent and predicting human action in human-robot collaboration. *Annu. Rev. Control Robot. Auton. Syst.* **2024**, *7*, 73–95. [[CrossRef](#)]
19. Su, H.; Qi, W.; Li, Z.; Chen, Z.; Ferrigno, G.; De Momi, E. Deep neural network approach in EMG-based force estimation for human-robot interaction. *IEEE Trans. Artif. Intell.* **2021**, *2*, 404–412. [[CrossRef](#)]
20. Zhang, T.; Sun, H.; Zou, Y. An electromyography signals-based human-robot collaboration system for human motion intention recognition and realization. *Robot. Comput.—Integr. Manuf.* **2022**, *77*, 102359. [[CrossRef](#)]
21. Sharifi, M.; Zakerimanesh, A.; Mehr, J.K.; Torabi, A.; Mushahwar, V.K.; Tavakoli, M. Impedance variation and learning strategies in human-robot interaction. *IEEE Trans. Cybern.* **2021**, *52*, 6462–6475. [[CrossRef](#)]
22. Yao, B.; Yang, B.; Xu, W.; Ji, Z.; Zhou, Z.; Wang, L. Virtual data generation for human intention prediction based on digital modeling of human-robot collaboration. *Robot. Comput.—Integr. Manuf.* **2024**, *87*, 102714. [[CrossRef](#)]
23. Li, Z.; Chu, X.; Hu, X.; Zhang, Z.; Li, N.; Li, J. Variable stiffness methods for robots: A review. *Smart Mater. Struct.* **2024**, *33*, 063002. [[CrossRef](#)]
24. Cong, L.; Shi, Y.; Zhang, J. Self-supervised attention learning for robot control. In *Proceedings of the 2021 IEEE International Conference on Robotics and Biomimetics (ROBIO)*; IEEE: New York, NY, USA, 2021; pp. 1153–1158.

25. Wang, C.; Ji, K.; Geng, J.; Ren, Z.; Fu, T.; Yang, F.; Guo, Y.; He, H.; Chen, X.; Zhan, Z.; et al. Imperative Learning: A Self-supervised Neuro-Symbolic Learning Framework for Robot Autonomy. *arXiv* **2024**, arXiv:2406.16087. [[CrossRef](#)]
26. Shaw, A. Self-Supervised Learning For Robust Robotic Grasping in Dynamic Environment. *arXiv* **2024**, arXiv:2410.11229. [[CrossRef](#)]
27. Lu, W.; Hu, Z.; Pan, J. Human-robot collaboration using variable admittance control and human intention prediction. In *Proceedings of the 2020 IEEE 16th International Conference on Automation Science and Engineering (CASE)*; IEEE: New York, NY, USA, 2020; pp. 1116–1121.
28. Belcamino, V.; Takase, M.; Kilina, M.; Carfi, A.; Mastrogiovanni, F.; Shimada, A.; Shimizu, S. Gaze-based intention recognition for human-robot collaboration. In *Proceedings of the 2024 International Conference on Advanced Visual Interfaces*; Association for Computing Machinery: New York, NY, USA, 2024; pp. 1–5.
29. Ferraguti, F.; Talignani Landi, C.; Sabattini, L.; Bonfe, M.; Fantuzzi, C.; Secchi, C. A variable admittance control strategy for stable physical human-robot interaction. *Int. J. Robot. Res.* **2019**, *38*, 747–765. [[CrossRef](#)]
30. Duchaine, V.; Gosselin, C.M. Investigation of human-robot interaction stability using Lyapunov theory. In *Proceedings of the 2008 IEEE International Conference on Robotics and Automation*; IEEE: New York, NY, USA, 2008; pp. 2189–2194.
31. Ranatunga, I.; Lewis, F.L.; Popa, D.O.; Tousif, S.M. Adaptive admittance control for human-robot interaction using model reference design and adaptive inverse filtering. *IEEE Trans. Control Syst. Technol.* **2016**, *25*, 278–285. [[CrossRef](#)]
32. Corke, P. Robot arm kinematics. In *Robotics, Vision and Control: Fundamental Algorithms in MATLAB® Second, Completely Revised, Extended and Updated Edition*; Springer: Berlin/Heidelberg, Germany, 2017; pp. 193–228.
33. Sharkawy, A.N.; Koustournpardis, P.N.; Aspragathos, N. Variable admittance control for human-robot collaboration based on online neural network training. In *Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*; IEEE: New York, NY, USA, 2018; pp. 1334–1339.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.