Thesis for the Degree of Master of Engineering

Building Blocks for Full Upper-Limb Haptics: Kinesthetic Exoskeleton and Pneumatic Glove

Minwoo Lee

Department of Software Convergence
Graduate School
Kyung Hee University
Seoul, Korea

February, 2026

Building Blocks for Full Upper-Limb Haptics: Kinesthetic Exoskeleton and Pneumatic Glove

Minwoo Lee

Department of Software Convergence
Graduate School
Kyung Hee University
Seoul, Korea

February, 2026

Building Blocks for Full Upper-Limb Haptics: Kinesthetic Exoskeleton and Pneumatic Glove

by

Minwoo Lee

Advised by

Prof. Seungjae Oh

Submitted to the Department of Software Convergence
and the Faculty of the Graduate School of
Kyung Hee University in partial fullfillment
of the requirements for degree of
Master of Engineering

Dissertation Committee:

Chairperson Hyoseok Hwang

Sanghyun Kim

Seungjae Oh

Contents

List of Figures

Abstract

Building Blocks for Full Upper-Limb Haptics: Kinesthetic Exoskeleton and Pneumatic Glove

Minwoo Lee
Master of Engineering
Graduate School of Kyung Hee University
Advised by Prof. Seungjae Oh

The human upper limb constitutes the primary channel for physical interaction, yet modern Human-Computer Interaction (HCI) interfaces largely fail to utilize its full kinesthetic potential. This limitation severs the link between digital intelligence and physical experience, hindering effective robotic embodiment in applications such as imitation learning and telepresence. This thesis proposes a modular haptic architecture comprising two complementary systems: a kinesthetic exoskeleton for the proximal arm and a bidirectional pneumatic glove for the distal hand.

First, the upper-limb exoskeleton is designed to prioritize biomechanical compatibility and intrinsic transparency. The exoskeleton design achieves biomechanical compatibility through a double-parallelogram remote-center-of-motion mechanism for the shoulder, while quasi-direct-drive actuators provide intrinsic transparency via low-ratio, high-efficiency transmissions. A real-time controller implements active inverse-dynamics compensation, establishing a transparent baseline upon which virtual forces are rendered. A redundancy-exploiting control strategy is developed to distribute virtual Cartesian forces across joint torques, enabling posture-dependent kinesthetic rendering. A preliminary musculoskeletal simulation using OpenSim demonstrates that this approach evokes muscle activation patterns comparable to those caused by a real external load, validating the exoskeleton's ability to render convincing force sensations.

Second, the pneumatic glove addresses the bandwidth limitations of soft actuators through a custom dual-reservoir electro-pneumatic architecture capable of active pressurization and vacuuming. A cascaded control loop incorporating feedforward hysteresis compensation is implemented to achieve high-bandwidth pose tracking. A user study demonstrates that an AI agent can communicate with users via physical hand pose intervention, with participants reporting improved intuitiveness and satisfaction compared to visual-only cues.

Together, these systems establish the foundational building blocks for a comprehensive upper-limb haptic interface, effectively bridging the gap between human motor intent and machine embodiment.

Keywords: Kinesthetic Haptics, Upper-Limb Exoskeleton, Pneumatic Haptic Glove, Posture-Dependent Haptic Rendering, Backdrivability

1 Introduction

The human upper limb is the principal tool through which people reach, manipulate objects, and physically interact with their surroundings. Its high degrees of freedom (DoFs), dexterous end-effector, and rich proprioceptive feedback make it the dominant channel for purposeful physical behavior. This centrality is evident from clinical literature [7], where upper-limb impairments resulting from neurological or physical injury are a major focus of treatment and rehabilitation due to their critical role in activities of daily living (ADLs).

1.1 Status of Upper Limb in Human–Computer Interaction (HCI)

In the course of expanding human capability through machine computation, the primary mission of modern human–computer interaction has been to utilize and better interpret the myriad actions the upper limb can produce—clicking, typing, swiping, pointing, gesturing, and so forth [8].

However, despite this importance, interactions involving the upper limb are still predominantly unidirectional in terms of information transfer. Widely used interface modalities—such as keyboard and mouse, touchscreen, and more recently VR controller and body-tracking systems—can capture and digitize human motion with high fidelity; yet, they provide at best localized haptic cues as feedback, which are insufficient to convey physical contexts involving the entire arm.

While interaction is often reduced to what feels like little more than buzzing and clicking at the surface, our innate, lifelong-developed ability to process haptic sensation, dynamics, and physical causality remains largely detached from the interaction context. As a result, HCI has evolved around an imbalanced exchange of information, fostering a heavy reliance on audiovisual channels. Consequently, the human capacity for information transmission and reception—now vastly outpaced by that of silicon processors—approaches saturation, posing the risk of rendering humans slow, indecisive, and error-prone components in performance-critical systems.

This observation motivates a research question—or rather, an engineering imperative:

We need better interfaces that leverage the physical capabilities of the human body, thereby maximizing the bandwidth of information we can handle.

1.2 Kinesthetic Haptics

While it is worth noting visionary perspectives on human–computer "integration" [9] and brain–computer interfaces point toward deep coupling between humans and machines, approaches that are closer to practical realization largely emerge from the field of haptics. In the haptics literature, interaction modalities are commonly classified according to the primary sensory channels they address [10]. Cutaneous haptics concerns sensations arising at the skin, including contact, pressure, vibration, and texture. In contrast, kinesthesia refers to the perception of limb motion, position, and force mediated by the musculoskeletal and proprioceptive systems. Haptic systems that target this modality constitute the subdomain of kinesthetic haptics, enabling the representation of physical dynamics, constraints, and forceful interaction beyond localized contact cues.

As stated, this thesis emphasizes the importance of making better use of kinesthetic information in human–computer interaction.

The significance of kinesthesia to HCI follows directly from the physics of interaction. All physical bodies obey the equations of motion, and interaction with the environment is fundamentally mediated by force. In the macroscopic world, interaction almost always involves contacts, through which forces are exchanged between the body and external objects. Unless a body accelerates, forces remain locally balanced (i.e., $\sum \vec{F} = 0$) typically by muscle contractions in the human body. Such forces propagate through the musculoskeletal system via chains of action and reaction, until they are ultimately counteracted by external supports, such as the ground or a seated surface.

Concurrently, the human sensorimotor system continuously processes these physical interactions. Mechanoreceptors generate proprioceptive and kinesthetic signals that are transmitted to the central nervous system [10]. While some of these signals are processed through low-level reflex pathways, a higher-level neural processor—the brain—integrates sensory input with prior knowledge to produce voluntary, goal-directed motor commands.

1.3 Limitations of Existing Kinesthetic Interfaces

Existing kinesthetic haptic interfaces span diverse mechanical architectures, each targeting different trade-offs between workspace, force fidelity, wearability, and complexity. This section surveys representative approaches and examines their limitations in supporting comprehensive, whole–upper-limb kinesthetic interaction.

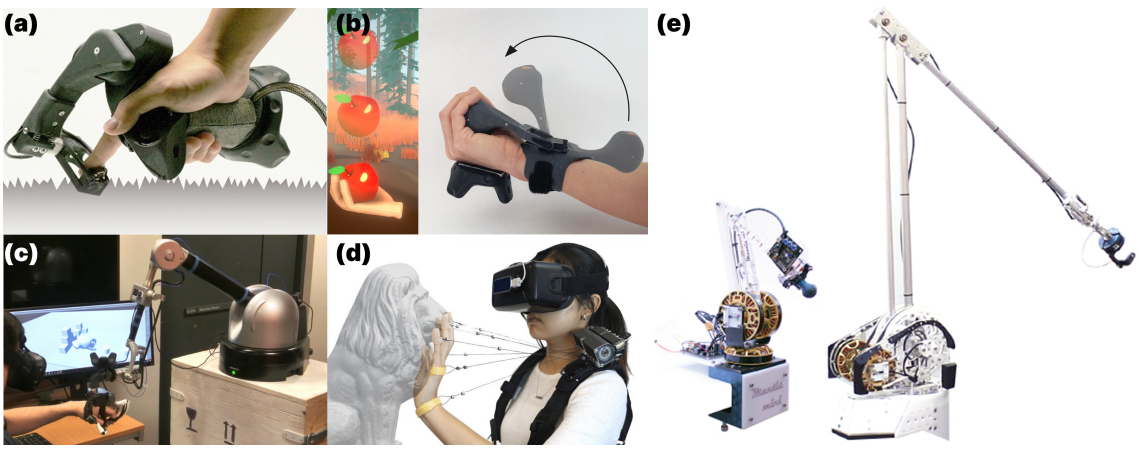

Figure 1: Examples of interface from existing kinesthetic methodologies: (a) CLAW [1] (b) Haptic PIVOT [2] (c) Docking Haptics [3] (d) Reel Feel [4] (e) Mantis [5]

As handheld controllers remain dominant in consumer VR, recent work has focused on expanding their in-hand haptic expressiveness. For example, “Grabity” [11] creates weight illusions through mechanical brake and asymmetric skin deformation; “CLAW” [1] augments a conventional VR controller with an actuated index-finger mechanism with motion-synchronized vibration; and “Haptic PIVOT” [2] swings a wrist-worn handle into the palm on demand. Despite these advances, handheld controllers remain fundamentally localized interfaces confined to the hand and, at most, the wrist.

Hybrid topologies and novel actuation mechanisms have been proposed to overcome the traditional trade-off between workspace and force fidelity. The "Docking Haptics" concept [3] combines wearable interfaces with grounded robots, allowing users to "dock" with physical manipulators only within designated interaction volumes. Pneumatic systems such as "AirPush" [12] generate multi-dimensional forces without rigid grounding by ejecting pressurized air jets at the fingertip. String-based architectures maximize workspace-to-weight ratio: "Wireality" [13] uses passive spring-loaded retractors with solenoid locks to arrest motion and simulate rigid objects, while "Reel Feel" [4] introduces active motor actuation to render continuous forces and compliance. However, because cables can transmit tensile forces, such systems can impose only unilateral (pulling) constraints.

Grounded 6-DoF haptic interfaces such as the "PHANTOM Omni" [14] and "sigma.7" [15] set benchmarks for high-fidelity force feedback, while "Mantis" [5] offers a low-cost, scalable alternative using brushless motors with admittance control. Nevertheless, because these devices are grounded manipulators, their reachable workspace is limited by link geometry and installation constraints, and enlarging the workspace typically comes with increased size, mass, and inertia—factors that reduce portability and can degrade transparency.

Electric Muscle Stimulation (EMS) can induce kinesthetic effects by electrically activating the user's muscles, offering a lightweight route to body-scale feedback that is not constrained by an external device's geometric workspace. In practice, however, EMS closes the control loop through time-varying physiological dynamics (e.g., electrode placement, skin impedance, fatigue, and inter-user response), which often necessitates extensive manual calibration for comfort and repeatability [16]. This variability complicates both precise control and the establishment of rigorous, standardized metrics required for reproducible force sensations [17].

Across these methodologies, a consistent pattern emerges: contributions have rarely formed into universally applicable solutions, mainly because the unpredictable, personalized nature of human perception and response mechanisms demands tailored design choices. From a technical perspective, grounded interfaces achieve high force fidelity but restrict workspace; wearable and cable-driven systems expand reachable volume but often sacrifice bidirectional force capability or user compliance; and hybrid systems mitigate trade-offs at the cost of added complexity or constrained interaction scenarios. Consequently, no existing approach simultaneously provides wide workspace coverage, distributed whole-arm force interaction, high transparency, practical wearability, and versatility.

These limitations suggest that achieving full upper-limb kinesthetic coverage requires reconsidering the mechanical form factor. Importantly, the Cartesian forces encountered at the hand are ultimately supported by joint torques distributed across the shoulder–elbow chain; for a kinematic chain, this relationship is captured by $\boldsymbol{\tau} = \mathbf{J}(\mathbf{q})^{\top}\mathbf{f}$. Therefore, a wearable robot capable of applying controlled torques at multiple joints offers a general and versatile route to rendering a broad class of kinesthetic interactions throughout the user's reachable workspace.

Interfaces whose kinematic chains run parallel to the human arm—namely, exoskeletons—provide such a path toward whole-arm kinesthetics. They can distribute forces across multiple joints while preserving kinematic compatibility, motivating the exploration of upper-limb exoskeletons paired with hand interfaces for comprehensive haptic embodiment.

2 Kinesthetic Exoskeleton Design for Full Upper-Limb Coverage

This chapter motivates the upper-limb exoskeleton as a form factor for multi-joint kinesthetic interaction, presents a lightweight prototype emphasizing intrinsic transparency, and introduces a joint-torque control method for rendering virtual Cartesian forces. Crucially, we leverage the arm's kinematic redundancy to diversify perceived sensation: the same virtual force can be distributed across joints differently depending on posture, engaging different muscle groups and enabling selective kinesthetic rendering.

2.1 Rationale for the Choice of Exoskeleton

Selecting a mechanical form factor for a comprehensive upper-limb haptic interface involves competing requirements: workspace coverage consistent with human arm mobility, the ability to apply forces across multiple body segments, minimal interference with voluntary motion, and practical deployability. Among candidate architectures—including grounded serial manipulators [18], cable-driven systems [4], and other wearable devices [19]—an exoskeleton is uniquely suited to these requirements because its kinematic chain runs parallel to the human limb. This parallel structure enables multiple attachment points distributed across the upper arm, forearm, and hand, forming the basis for whole-arm kinesthetic interaction.

2.1.1 Virtual Force Rendering Task

Virtual force rendering refers to the methodology of physically presenting, through a haptic interface, forces that do not exist in the user's local environment but are computed from external causes—for example, contact forces encountered by a remotely operated robot or interaction forces generated in a virtual environment [20]. Conventional approaches typically assume a single interaction point and therefore rely on 6-DoF endpoint devices held in the hand. While effective for localized interaction, such devices implicitly collapse the human arm into a single contact point.

The human upper limb, however, is a multi-joint structure with kinematic redundancy beyond six degrees of freedom. Consequently, whole-arm kinesthetic feedback must account for how forces distribute across multiple joints rather than being perceived solely at the hand. This observation motivates haptic interfaces that can engage the arm at multiple locations simultaneously.

Prior work aiming to distribute kinesthetic effects beyond a single contact point has frequently adopted Electric Muscle Stimulation (EMS), largely because electrodes can be placed relatively easily at multiple locations on the body to influence different joints. Some systems augment EMS with passive mechanical constraints (e.g., brakes or ratchets) to bound motion and stabilize the perceived output [21].

This thesis hypothesizes that an actively motorized haptic interface with a multi-joint mechanical structure—namely, an upper-limb exoskeleton—can more directly realize quantitative virtual force rendering. By applying measurable joint torques through a rigid kinematic structure, such a system enables deterministic control while explicitly exploiting the arm's redundant degrees of freedom.

2.1.2 Advantages of the Exoskeleton Form Factor

The distributed physical contact afforded by an exoskeleton fundamentally distinguishes it from endpoint 6-DoF interfaces, which constrain interaction to a single contact location (often via a stylus-type handle). Whereas single-point devices can impose only one endpoint constraint, an exoskeleton can apply and sense interaction torques across multiple joints simultaneously. This capability is essential for rendering whole-arm kinesthetic conditions, such as simulating the inertia of a held object or selectively resisting shoulder motion while permitting elbow movement.

The parallel kinematic structure of an exoskeleton also provides intrinsic pose sensing: joint encoders directly measure limb configuration without the occlusion, latency, and drift commonly associated with vision-based tracking. This property enables high-bandwidth proprioceptive feedback for closed-loop haptic control. Moreover, exoskeletons avoid the inherent workspace–force trade-off of grounded manipulators. Enlarging workspace in a grounded robot typically requires longer and heavier links, whereas an exoskeleton moves with the user and therefore naturally matches the wearer's range of motion without external fixtures.

In this work, we therefore target two goals: (1) to implement a simplified, wearable arm exoskeleton suitable as a kinesthetic haptic interface rather than a full rehabilitation robot, and (2) to formulate a control method that renders virtual Cartesian forces while leveraging the arm's redundancy to diversify how the same virtual force is perceived across joints and muscles.

2.2 Material: Upper Limb Exoskeleton

2.2.1 Biomechanical Compatibility

In wearable kinesthetic interfaces, biomechanical compatibility does not necessarily require strict anatomical replication. Although one could, in principle, align an exoskeleton's degrees of freedom (DoFs), centers of motion, and joint axes with those of the human arm, such a one-to-one correspondence is often impractical. Human upper-limb motion is instead shaped by neuromotor coordination strategies in which multiple joints are recruited together as functional units—often

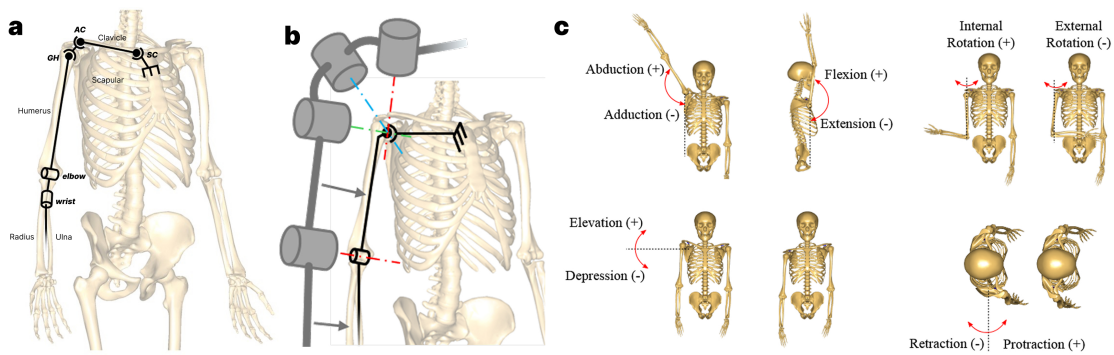

Figure 2: Human upper-limb kinematics: (a) commonly used simplified kinematic definitions; (b) further reduced representation used for exoskeleton alignment; and (c) representative motions highlighting the contribution of the shoulder girdle [6].

described as motor (or muscle) synergies [22]. From this perspective, an upper-limb exoskeleton may target a reduced set of representative DoFs that captures the dominant components of whole-arm motion while preserving functional compatibility.

A widely adopted abstraction is the 4-DoF arm model (see Figure 2b), in which the glenohumeral (GH) joint is approximated as a 3-DoF spherical joint rigidly attached to the thorax and the elbow as a 1-DoF revolute joint [23, 6]. While such reduction is convenient and mechanically economical, they effectively treat the sternoclavicular (SC) and acromioclavicular (AC) joints as rigid. Consequently, shoulder-girdle motions such as elevation, protraction/retraction, and scapular rotation are only indirectly represented, which can lead to kinematic mismatch during high-elevation or large-amplitude movements. Increasing the number of anatomically aligned DoFs can alleviate these mismatches, but doing so comes at the cost of additional mass, larger linkages, and higher mechanical complexity, all of which directly undermine wearability and intrinsic transparency [24, 25]. For wearable haptic interfaces, this trade-off motivates designs that prioritize functional compatibility over strict anatomical fidelity.

An alternative data-driven kinematic synthesis may adopted, in which joint placements are optimized to best reproduce observed human motions under explicit design constraints, such as self-collision avoidance, keep-out regions around the torso, and inertia limits. Formal formulations for joint-placement optimization and co-optimization of tracking configurations on $SE(3)$ can substantially improve coverage and reduce relative motion at attachment interfaces [26, 27]. Even though such optimization is regarded as future work and synthesis process was not integrated into the current hardware iteration, the detailed mathematical formulation and algorithms provided in Appendix A. In this work, available real human motion datasets [28] used primarily to inform target workspace coverage rather than as direct optimization objectives, and the current prototype based on 4-DoF model was designed heuristically.

From a mechanical standpoint, the shoulder being reduced as a single 3-DoF GH joint is the most challenging joint to interface in an upper-limb exoskeleton. Its effective center of rotation lies inside the body and migrates with scapular motion, making external alignment inherently difficult. To accommodate this, remote-center-of-motion (RCM) mechanisms are commonly employed. Among them, double-parallelogram linkages offer a compact and structurally stiff solution, while the inclusion of passive joints can further improve load transmission and robustness [29, 30].

2.2.2 Actuator Choices for Intrinsic Transparency

A wearable exoskeleton intended for continuous physical interaction with a human must exhibit compliant and transparent behavior, that is during free motion, the device should neither impede voluntary movement nor introduce parasitic forces perceptible to the wearer. Unlike industrial manipulator where external disturbances are deliberately rejected through high-ratio transmissions and position control, a human–robot interaction (HRI) system must interpret the human's motion itself as the primary input rather than as a disturbance [31].

This requirement becomes more apparent for kinesthetic interfaces. In contrast to collaborative robots that are designed to safely tolerate occasional contact, an exoskeleton must support continuous, bidirectional force exchange while remaining perceptually "invisible" whenever no virtual force is commanded. Achieving such transparency is therefore a core design objective at

the actuator level.

Design Strategies for Transparency Two broad strategies have been explored to realize transparency. The first relies on explicit force sensing at the human–device interface, closing a feedback loop that regulates actuator torque to track a desired interaction force, typically zero during free motion [32]. Series elastic actuators (SEAs) follow this approach by inserting a compliant element between the motor and the load, whose deflection provides a direct measurement of transmitted torque [33, 34]. When combined with accurate feedforward compensation of gravity and inertia, SEAs can achieve high levels of dynamic transparency, as demonstrated in upper-limb rehabilitation platforms such as ANYexo [33].

The second strategy pursues intrinsic transparency by minimizing the mechanical impedance of the transmission itself. If an actuator is sufficiently backdrivable, the motor can be rotated from the load side with little resistance, allowing motor current to serve as a reliable proxy for output torque. This enables proprioceptive force control without dedicated force-torque sensors [35, 36]. By eliminating the compliant element required in SEAs, intrinsic transparency increases control bandwidth and reduces mechanical complexity—both critical advantages for a lightweight, wearable exoskeleton.

Backdrivability and Quasi-Direct-Drive Actuation Backdrivability is governed by the combined effects of gear ratio, transmission efficiency, reflected rotor inertia, and friction [37, 38]. High-ratio gearboxes amplify output torque but also amplify the torque required to backdrive the joint, often rendering the system perceptually stiff. Conversely, low-ratio or direct-drive configurations preserve backdrivability at the expense of requiring motors with higher torque density.

Recent advances in motor design have shifted this trade-off. As articulated by Seok et al. [35], actuators built around high–torque-density brushless DC (BLDC) motors paired with low-ratio, high-efficiency transmissions can simultaneously deliver meaningful force output and high transparency. This design paradigm, commonly referred to as quasi-direct drive (QDD), has been successfully applied in dynamically demanding systems such as legged robots, whose force and bandwidth requirements exceed those of most haptic interfaces [39].

Actuator Implementation in the Present Exoskeleton The present exoskeleton adopts this intrinsic-transparency philosophy while accommodating the practical constraints of a wearable device. Performance targets were derived from upper-limb biomechanics. Typical elbow angular velocities during activities of daily living are on the order of 40 RPM, while highly dynamic motions such as tennis serves can reach several hundred RPM [40]. For haptic interaction, however, the primary requirement is not extreme speed but sufficient torque and bandwidth to render virtual forces without perceptible delay. A representative scenario rendering a 1 kgf virtual force at the palm approximately 0.3 m from the elbow corresponds to a joint torque of roughly 3 Nm.

Based on these considerations, a custom actuator was developed for the elbow joint by pairing a high–torque-density BLDC motor (maxon ECX FLAT 32L) with a compact strain-wave reducer (Harmonic Drive CSF-11-2UH-ULW). Although the resulting 50:1 reduction ratio exceeds that of typical QDD designs, the selected gearbox offers low backdrive torque for its class and favorable efficiency, helping preserve intrinsic transparency relative to conventional high-ratio transmissions. Figure 3 illustrates the integrated actuator module, and Table 1 summarizes the resulting performance.

For the remaining joints, commercially available QDD actuator modules (CubeMars AK60-6) were employed to reduce development overhead while maintaining consistent transparency characteristics across the arm. Their specifications are listed in Table 2.

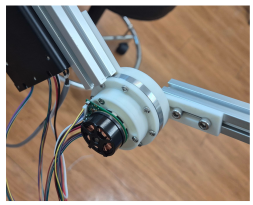

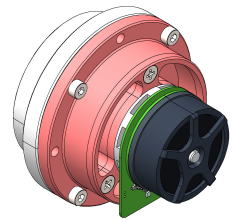

NO	ITEM NAME	MTRL	QTY	NOTE
1	HD Body	-	1	CSF-11-050-2UH-ULW
2	HD Wave Generator	-	1	CSF-11-050-2UH-ULW
3	Housing Body	A6061	1	
4	Housing Mount	A6061	1	
5	Hub Spacer	BRASS	1	
6	BLDC	-	1	Maxon ECX FLAT 32 L
7	Housing Endcap	A6061	1	
8	Hub-Housing Bearing	SUJ2	2	C-E6703
9	Shim Ring Spacer	SUS304	1	
10	Hub Endcap	A6061	1	
11	Hub Body	A6061	1	
12	Pan M2.5x12	SUS304	4	
13	Set Screw M2.5x6	SCM435	2	SSHC-ST-M2_5-6
14	Flat M2x6	SUS304	7	
15	Internal Hex M3x12	SUS304	4	
16	Flat M3x10	SUS304	3	
17	Flat M3x5	SUS304	3	

Figure 3: Designed elbow joint actuator module assembly: brushless DC motor integrated with low-ratio harmonic reducer fit to QDD definition.

Table 1: Elbow actuator component specifications and output performance.

Component	**Parameter**	**Rated**	**Peak**
Motor (ECX FLAT 32 L 48V)	Output torque (mNm)	106	488
	Speed (RPM)	10,700	13,800
	Mass (g)	71.2	
Gearbox (CSF-11-2UH-ULW)	Reduction ratio	50:1	
	Backdrive torque (Nm)	0.55	
	Mass (g)	150	
Combined output	Output torque (Nm)	3.5†	17†
	Output speed (RPM)	70†	170†
	Total mass (g)	∼300	

† Limited by gearbox intermittent load capacity.

Proprioceptive Control and Transparency in Operation The use of backdrivable actuators enables a proprioceptive control paradigm in which motor current provides an estimate of joint torque and encoders measure joint position, eliminating the need for dedicated force–torque sensors [35, 36]. In the context of an exoskeleton, this approach is particularly advantageous: force sensors placed at the human–device interface add mass, compliance, and calibration complexity precisely where high stiffness and bandwidth are desired [32].

Table 2: CubeMars AK60-6 actuator specifications.

Parameter	Rated	Peak
Output torque (Nm)	3	9
Output speed (RPM)	233	—
Current (A)	3.8	13.1
Reduction ratio	6:1	
Backdrive torque (Nm)	0.2	
Mass (g)	368	

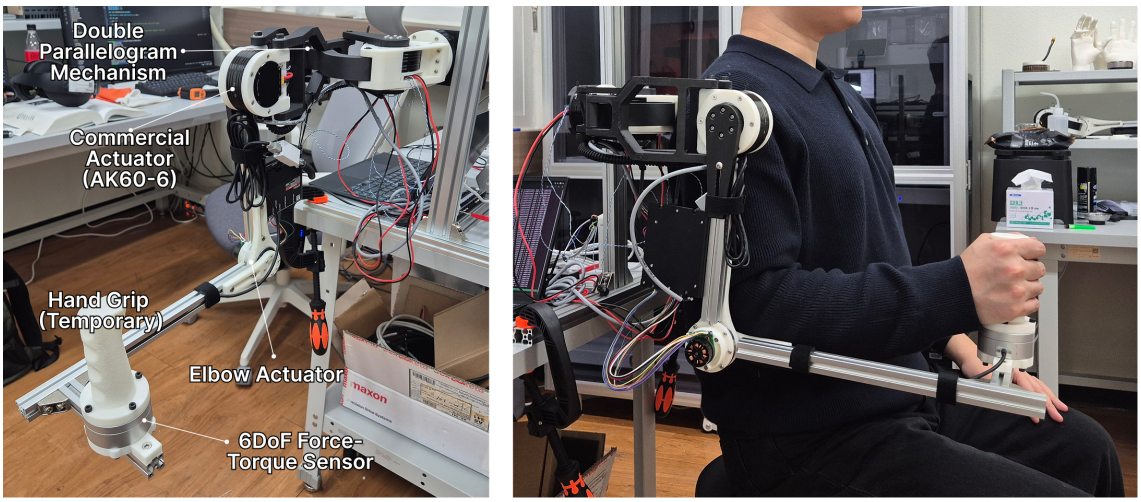

Figure 4: (Left) Overview of the manufactured upper-limb exoskeleton prototype. (Right) Snapshot of the operating setup with a hand-grip interface at the end-effector.

In operation, intrinsic transparency is achieved by actively compensating the exoskeleton's own dynamics—primarily gravity and inertia—so that, in the absence of commanded virtual forces, the device presents a near-zero-impedance behavior to the user. This transparent baseline allows virtual forces to be introduced additively and perceptually cleanly when required, forming the foundation for the redundancy-aware virtual force rendering methods described in the following sections.

2.2.3 Prototype Fabrication

The upper-limb exoskeleton prototype was fabricated to balance kinematic coverage, intrinsic transparency (backdrivability), and manufacturability, while remaining compact for laboratory operation. To achieve a lightweight yet stiff structure, load-bearing parts were produced via 3D printing using engineering-grade polymer (PA6-CF, eSun), and long-span members were reinforced with standard aluminum extrusion frames.

Owing to the backdrivable actuator selection, the prototype remains inherently compliant to the user's motion even when unpowered. The total mass is below 3 kg, primarily enabled by the extensive use of 3D-printed components and modular aluminum extrusions. Figure 4 shows the fabricated device and an example operating setup with a hand-grip interface at the end-effector.

2.2.4 Software Framework and Low-level Torque Control

To support high-bandwidth kinesthetic interaction, the control system is implemented on a real-time Linux environment and executes a 1 kHz torque-control loop. The software stack is built upon the ROS 2 architecture, leveraging its modular abstractions to decouple hardware-specific communication from control logic and to facilitate rapid iteration [41]. In particular, we adopt the `ros2_control` framework to standardize the real-time control pipeline and to expose the exoskeleton as a hardware interface that can be reused across controllers [42].

Hardware Abstraction and Real-time Communication At the lowest level, actuator communication is handled via the CAN 2.0B bus using the Linux SocketCAN API. A dedicated Hardware Abstraction Layer (HAL) is implemented by inheriting the `ros2_control` `hardware_interface` class, providing deterministic message handling and a unified joint-level API (state readback and torque command write) to the upper-level controller.

Model-based Compensation for Active Transparency While intrinsic backdrivability reduces mechanical impedance, perceptual transparency during free motion further benefits from actively compensating the exoskeleton's own dynamics. The real-time controller therefore computes model-based torques using the `pinocchio` library for efficient rigid-body dynamics [43]. The full joint-space dynamics are expressed as

$$M(q)\ddot{q} + C(q,\dot{q}) + g(q) = \tau,$$

where $M(q)$ is the inertia matrix, $C(q,\dot{q})$ represents Coriolis and centrifugal effects, and $g(q)$ is the gravity vector.

To realize a transparent baseline, we define the commanded torque as the sum of a dynamics-compensation term and an interaction (virtual) torque:

$$\tau = \tau_{\text{comp}} + \tau_{\text{virtual}}.$$

In free motion, τ_{virtual} is set to zero, and the controller applies

$$\tau_{\text{comp}} = C(q,\dot{q}) + g(q),$$

which offsets gravitational and velocity-dependent torques that would otherwise be felt by the wearer. This "zero-force rendering" mode reduces parasitic impedance and reserves actuator authority for the additive rendering of interaction forces.

Interface to Virtual Force Rendering The virtual torque τ_{virtual} is generated by the higher-level virtual force rendering module (described in the following section), which maps desired Cartesian interaction forces into joint torques and distributes them across the arm's redundant degrees of freedom.

2.3 Methodology: Multi-Joint Virtual Force Rendering

In conventional robotic control, kinematic redundancy is often treated as a solution space for optimizing secondary criteria such as singularity avoidance, torque minimization, or obstacle avoidance. In physical human–robot interaction (pHRI) and haptic rendering, redundancy plays a different role: it becomes a controllable distribution mechanism for how the same task-level interaction is perceived across joints and muscle groups.

This interaction mode remains under-explored in many practical interfaces, because a large fraction of haptic devices intentionally collapse interaction to a single contact point or simplify the mechanism to avoid redundant actuation. In contrast, the present work explicitly retains whole-arm actuation and uses redundancy as a capability to transfer selective kinesthetic information.

2.3.1 Formulating Cartesian Wrench to Joint Torque

The fundamental mapping between a desired Cartesian interaction and the required joint actuation is derived using the Principle of Virtual Work. Let $\delta\mathcal{W}$ denote virtual work. In the Cartesian domain, it can be expressed as the inner product of the Cartesian wrench $\mathcal{F} \in \mathbb{R}^6$ (forces and moments) and an infinitesimal virtual displacement δx:

$$\delta\mathcal{W} = \mathcal{F}^T \delta x.$$

Equivalently, in joint space, virtual work is expressed by the joint torque vector $\boldsymbol{\tau} \in \mathbb{R}^n$ and a virtual joint displacement δq:

$$\delta\mathcal{W} = \boldsymbol{\tau}^T \delta q.$$

The differential kinematics relates these virtual displacements through the geometric Jacobian $J(q)$:

$$\delta x = J(q)\, \delta q.$$

Substituting and equating the two expressions yields:

$$\mathcal{F}^T J(q)\, \delta q = \boldsymbol{\tau}^T \delta q,$$

and since this must hold for any admissible δq, the static wrench-to-torque mapping follows:

$$\boldsymbol{\tau} = J^T(q)\, \mathcal{F}.$$

In this thesis, we define the resulting joint torque as the virtual rendering torque

$$\tau_{\text{virtual}} \triangleq J^T(q)\, \mathcal{F},$$

which is provided as the joint-torque reference to the low-level torque controller (Section 2.2.4).

Embedding into the torque-controlled exoskeleton dynamics Using the joint-space equation of motion

$$M(q)\ddot{q} + C(q, \dot{q}) + g(q) = \tau_{\text{comp}} + \tau_{\text{virtual}},$$

the closed-loop behavior becomes

$$M(q)\ddot{q} \approx \tau_{\text{virtual}}.$$

In practice, during quasi-static rendering the user counteracts this tendency to accelerate, resulting in $\ddot{q} \approx 0$ while interaction forces are transmitted through the attachment interfaces. This is the primary regime considered in the present evaluation; dynamic rendering extensions are discussed as future work.

Posture dependence and quasi-static perception model A key implication of (2.3.1) is that the torque distribution depends on posture through $J(q)$: the same Cartesian wrench can induce different joint-torque patterns at different configurations. In the proposed rendering mode, the controller generates an additive virtual torque τ_{virtual} from a desired Cartesian wrench via (2.3.1). We assume the user maintains a fixed arm posture against the device, so any virtual wrench $\mathcal{F}$ is counteracted by user effort rather than causing motion. Under this quasi-static interaction, the user is assumed to maintain a posture against the rendered wrench; i.e., the user provides counteracting joint torques such that net acceleration remains small. In this regime, the virtual wrench is primarily perceived through the support effort required to equilibrate τ_{virtual}, rather than through large involuntary motion.

2.3.2 Redundancy-Exploiting Rendering for Diversified Kinesthetic Perception

The redundancy of the upper limb enables the proposed virtual-force rendering to be extended beyond a single "force output" into a family of rendering modes that diversify how the same task-level virtual wrench is perceived. Concretely, we consider two complementary mechanisms: (i) posture-driven diversification under quasi-static support, and (ii) energy-based dynamic rendering when nonzero accelerations are allowed.

Mode I: Posture-driven diversification under quasi-static support Even for an identical virtual wrench magnitude, changing the arm posture alters the Jacobian $J(q)$ and therefore changes the joint-torque distribution $\tau_{\text{virtual}} = J^T(q)\mathcal{F}$. Under quasi-static rendering ($\ddot{q} \approx 0$), the user is assumed to maintain the posture by providing counteracting torques through the attachment interfaces. In this regime, the virtual interaction is primarily perceived through the support effort required to hold the configuration, and posture becomes an explicit handle to shape where the effort is felt.

An intuitive analogy is weight training: holding the same dumbbell weight while extending the arm forward versus laterally produces distinct joint-torque patterns and muscle engagement. Similarly, by intentionally guiding the user toward a posture consistent with a targeted muscle-engagement objective, the same virtual force can be rendered as different kinesthetic "modes" without changing the task-level wrench.

Mode II: Energy-aware dynamic rendering (future work) The quasi-static formulation above explains how a virtual wrench can be perceived through support effort when the user maintains posture. Extending the same idea to dynamic interaction ($\ddot{q} \neq 0$), however, requires explicit consideration of energy exchange. In haptic interaction, the human impedance and the environment dynamics are generally unknown; therefore, a promising direction is to adopt passivity-inspired, energy-aware rendering schemes that regulate how much net energy the device can inject into the coupled human–robot system.

One representative concept is the virtual energy tank [44]. Intuitively, the tank stores an energy budget associated with the intended rendering action (e.g., the virtual work that would have been performed if the interaction force physically existed along the user's motion). During operation, the controller monitors the actuator-side mechanical power and modulates the rendering command so that the net energy delivered through the interface remains within the available budget. When the budget becomes low, the rendering action is smoothly reduced (e.g., by scaling the commanded virtual torque), which can improve robustness against destabilizing energy injection during fast motion or contact transitions. Incorporating such energy-aware constraints while retaining redundancy-based torque distribution such as shaping where the effort is felt across joints is a natural extension of the present framework and is left as future work.

2.3.3 Human Musculoskeletal Simulation Study

To validate whether the proposed quasi-static virtual force rendering induces a musculoskeletal response comparable to that of a physically applied external load, we conducted a static upper-limb simulation study and compared muscle-force patterns between (i) a baseline case with an actual external force applied at the hand and (ii) a rendering case where the same interaction is produced through the exoskeleton and transmitted to the user through the attachment interfaces.

Simulation environment We used the OpenSim musculoskeletal simulation suite [45] and an upper-limb model that includes anatomically detailed shoulder-girdle mechanics [46]. A representative arm posture was prescribed, and the analysis focused on quasi-static force support (i.e., negligible joint accelerations).

Protocol: baseline vs. rendering cases Two simulation conditions were constructed under the same posture:

1. **Baseline (physical load):** a constant 10 N external force was applied at the distal hand (end-effector) to represent a real interaction force.

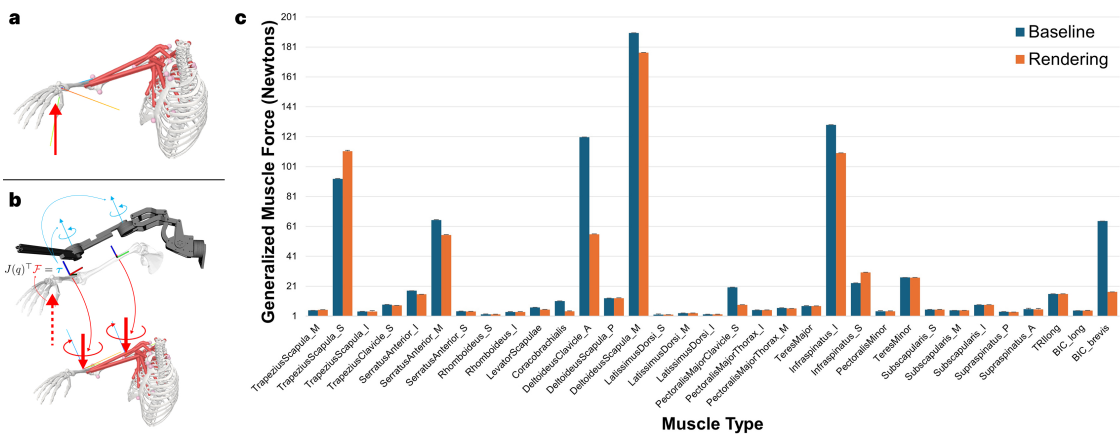

Figure 5: Musculoskeletal simulation study: (a) Baseline: 10 N hand load. (b) Exoskeleton-mediated rendering via upper-arm/forearm attachments. (c) Generalized muscle force comparison (baseline vs rendering).

2. **Rendering (exoskeleton-mediated):** the exoskeleton generated the virtual rendering torque corresponding to the same task-level interaction, and the load was transmitted to the human model through the two attachment interfaces located on the upper arm and forearm. This condition represents the intended quasi-static rendering mode, where the user supports the interaction through the attachment points.

Metric We compared the generalized muscle force across all modeled muscles. The comparison was primarily qualitative (pattern-level similarity and outliers), focusing on whether major contributors and their relative magnitudes were preserved between the two conditions.

Results and interpretation Figure 5 shows the generalized muscle forces for the baseline and rendering conditions. Overall, the rendering condition reproduced a similar muscle-force pattern to the physical-load baseline, with most muscles exhibiting comparable activation magnitudes. Deviations were mainly observed in a limited subset of muscles, which tended to be anatomically close to the exoskeleton attachment regions. This suggests that differences in load transmission paths (hand-applied vs. interface-applied) and local stabilization requirements around the attachments can alter the recruitment of nearby muscles, even when the task-level force is matched.

2.4 Limitation and Future Resolution

This work demonstrates a kinesthetic exoskeleton interface and a quasi-static virtual-force rendering formulation, together with a preliminary musculoskeletal simulation study. While the results support the feasibility of the approach, several limitations remain and directly motivate future work.

- Prototype still relies on heuristic mechanical design, which constrains effective workspace and requires attachment-point relative motion to remain small for reliable rendering. Future work should use optimization-driven kinematic design to reduce interface slip and expand usable range.
- Rendering framework is currently centered on quasi-static force equilibrium and does not address dynamic interaction where torque must vary along trajectories; extending it toward energy-aware control is a natural next step.
- Evaluation is limited to simulation, so the immediate plan is to conduct user studies with the prototype to assess perceived equivalence and diversification effects and to guide the next hardware/control iteration.

Concluding Remarks In summary, we presented a lightweight wearable upper-limb exoskeleton prototype with high dynamic transparency, and discussed a quasi-static virtual force rendering scheme in which the perceived load can be diversified via posture-dependent torque distribution. A preliminary musculoskeletal simulation further suggests that exoskeleton-mediated rendering can elicit muscle-force patterns comparable to those induced by a physically applied external load, providing an initial indication that quantitatively controlled haptic sensations are attainable. Collectively, these results connect joint-space control authority to task-space virtual interactions, supporting the potential of arm exoskeletons as general-purpose haptic interfaces for extending HCI toward spatial and physical embodiment.

3 Haptic Glove with Bidirectional Pneumatics for Hand Pose Control

This section is an adaptation based on the previous work: "AI Speaks with Hands: A Bidirectionally Actuated Pneumatic Glove for Pose-Based Human-AI Communication" by Minwoo Lee, Sungjoon Yoon, Seongmin Yun, Yunseo Do, Seungjae Oh.

With the arm exoskeleton enabling kinesthetic force feedback at the shoulder and elbow, we next address the distal upper limb. This chapter presents a pneumatic haptic glove that extends the interface to the user's fingers, providing pose-level feedback that complements the arm's force-level rendering.

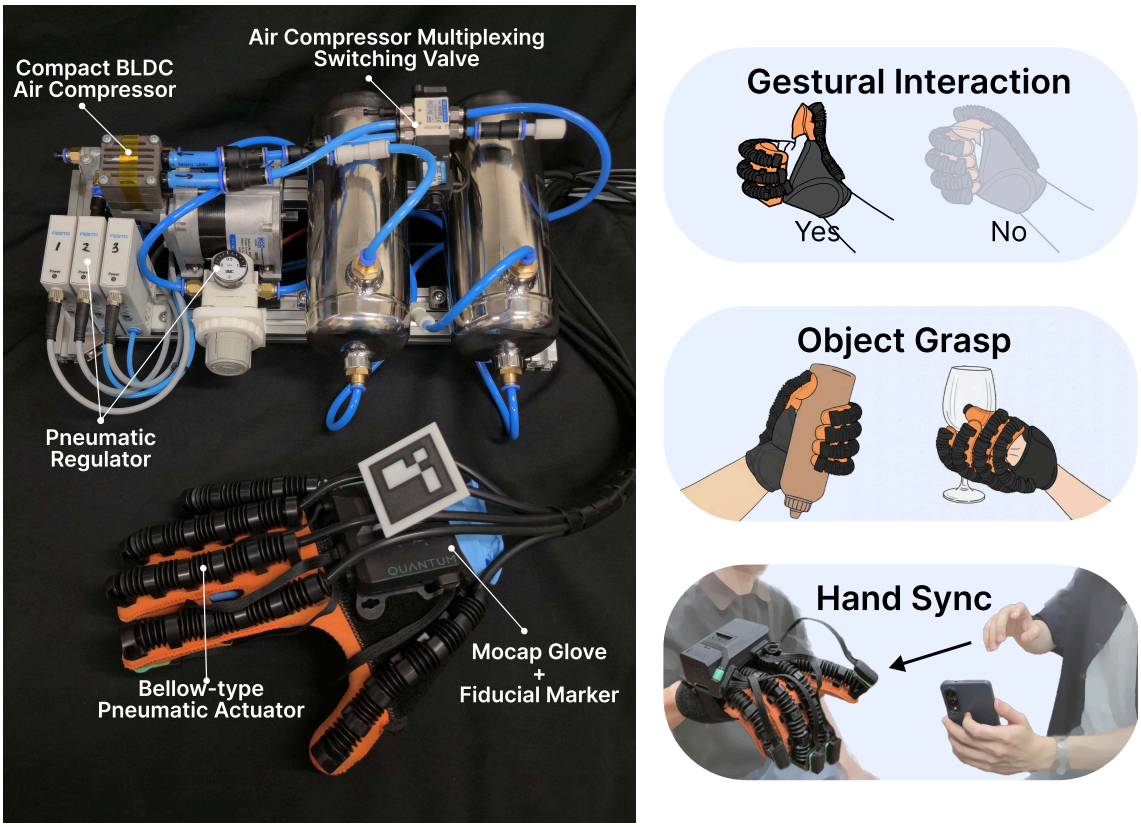

Figure 6: Technical overview and interaction scenarios of proposed glove system.

3.1 Introduction

Recent advances in large language models (LLMs) and AI companion have dramatically changed the way in which humans interact and communicate with computers. While the development of vision-language models (VLMs) expanding input modalities to include visual image data, enabling deeper contextual understanding, interaction with the user and AI companion remains highly limited to textural or auditory input and output, which limits the intuitiveness and expressiveness of

interaction. These constraints raise critical research questions in scenarios that demand nuanced, embodied, and context-aware communication.

One of the solutions is accounting additional channel of communication from novel interaction modalities, and in this study we focuses on taking hand gestures as the one. Through hand-based communications, AI agents can express their subtle intentions via a user's hand and understand the user's intentions by observing their hand poses.

Soft pneumatic gloves, as capable conveyor of gestural contexts, fit right into the role of method, while having decent characteristics such as mechanical simplicity, intrinsic safety, and user comfort [47, 48]. However, they have not been regarded as suitable for fine and responsive tasks, as these advantages come at the cost of reduced motion precision and actuation speed due to the inherent mechanical compliance of pneumatic bellows [49]. To circumvent this, researchers have explored the use of multiple pneumatic actuators assigned to movements in different directions to increase the degree of freedom in motion and enhance control bandwidth for motion responsiveness [50, 51, 52, 53]. Introducing negative air pressure allows flexion and extension with a single actuator; however, only a few studies address this system due to technical difficulty in pneumatic system design and control [54, 55, 49]. Moreover, many studies still focus on pneumatic actuation mechanisms alone [55, 49], rather than holistic devices or systems.

Among the studies focusing on glove-type pneumatic devices, open-loop control can suffer from low motion precision and resolution [52, 54]. Some studies have introduced closed-loop control using sensors such as pressure sensors [47], QR codes [53], force sensors [51], and bending sensors [51, 50]; however, these presented low-to-moderate motion bandwidth (0.1-4 Hz) due to low sensor performance. Chen et al. [50] achieved the fastest motion bandwidth (up to 4 Hz), but it requires specialized optical-waveguide sensor for bending detection.

In this study, we developed a dedicated haptic glove powered by pneumatic system, with closed-loop hand pose control architecture, and a dedicated pneumatic circuit that actively inflates and deflates pneumatic actuators, enabling bidirectional control capability creating responsive output motion. Equipped with each single positive and negative pressure reservoir, air pressure can be managed with stable and sustained manner while having single compact air compressor unit installed. The result, pneumatic powerhouse, fits into a small enclosure suitable for tabletop use.

Being able to control the pose of hands with the proposed system, context-rich gestural communication scenarios could be realized and evaluated. While a key technical challenge presents in the system is how to assist—or occasionally inhibit—fine-grained human hand movements in a fast and reliable manner, system was validated with an empirical user study focusing on two interaction scenarios where AI agents respond and intervene through the pneumatic gestural communication system, to determine the factor affecting user experience.

3.2 Hardware Design of Pneumatic System and Glove Interface

Glove Interface For a wearable glove form factor, we modify a commercial pneumatic glove for rehabilitation (ML-115A). The glove is equipped with five bellow-type pneumatic actuators, one for each finger, providing 1 DoF flexion and extension. Moreover, we integrate a hand motion capture device (Quantum Metaglove, Manus) to the glove by installing the mount structure of its sensing components for each digit on the pad side of the glove. This modification enables fast and robust acquisition of hand skeleton positions (>120 Hz), which is not feasible with typical

vision-based approaches that could not result robust estimation due to the the distinct appearance of the glove and occlusion. We use a 3d-printed arUco marker onto the glove, enabling global hand position tracking useful for interaction design. Furthermore, we use a mobile computing system paired with the mocap device (Quantum Bodypack, Manus) to stream hand tracking data to the control system.

Electro-pneumatic Control & Regulation The pneumatic glove is operated by pneumatic actuation and control systems. We use electro-pneumatic regulator valves (VEAB Proportional Regulator, Festo) that control target pressure values in the range of $+100$ to -100 kPa, directly supplying air to or removing air from the pneumatic bellows. Our early prototypes, which relied solely on conventional positive-pressure-only regulators, suffered from slow finger extension because the bellows could only vent passively. Even recent pneumatic gloves use more than one actuator to achieve bi-directional motion if only positive pressure actuation is available [51, 50, 52, 53]. In summary, incorporating negative pneumatic actuation enables rapid and responsive extension with a single bellow actuator (see Figure 9).

Pneumatic Circuit Design We design a pneumatic circuit with a dual-reservoir system for stable system responses. For its compact pneumatic source, we employ a BLDC-driven rocking-piston compressor (30RNS-ED, KCC) capable of supplying both positive and negative pressures ($+650$ to -65 kPa). This single pressure source is electronically controlled to supply either the positive or negative pressure reservoir, which is dynamically switched by a 5-way/2-mode solenoid valve (KS105S, KCC).

The positive pressure reservoir can store a sufficient amount of compressed air (over $+600$ kPa), which is then decompressed via a manual pressure regulator to match the maximum operating pressure requirement of the regulator ($+200$ kPa). The reservoir is capable of keeping the operational pressure level for extended time of operation, varying over the frequency and intensity of actuation yet sufficient for supporting a short burst of gestures equivalent of 10 repeated full extension/flexion.

Conversely for the negative pressure source, because vacuum cannot be stored like compressed air—it represents the absence of pressure rather than its accumulation—the reservoir functions primarily as a buffer. It absorbs sudden inflows when bellows depressurize, while the continuously running pump maintains the vacuum level for rapid finger extension. And simultaneously, entering air must be immediately pushed out to maintain vacuum pressure level. Therefore, the solenoid valve normally connects the air compressor to negative pressure side making air out, and sparsely switches to the positive pressure reservoir when sensed pressure level drops below the operation range. From this strategy, the maximum negative pressure (-65 kPa) is maintained in the reservoir for the most time.

Additionally, pneumatic check valves functioning like diodes in electric circuits are installed between the reservoirs and the pneumatic source to minimize air flow leakage while not charging, prevent air compressor from being forced to continuously operate. Overall, the system is compact enough for tabletop use yet supports reliable, sustained operation.

Micro Controller Unit The MCU (Pico W, Raspberry Pi) controls the operation of all electro-pneumatic parts and coordinates communication between the pneumatic system and its closed-control loop controller. It receives digital set points from the control loop controller and generates logic-level signals, which are then amplified by a custom DAC-amp circuit to the analog voltage levels (0-10 V) required by the regulator valves.

3.3 Control Software Framework

Middleware Our middleware aggregates, processes, and transmits hand pose information, working between high-level applications and low-level embedded systems. We design it to support hand skeleton data from different APIs or sources. The middleware generates pressure commands with a closed-loop controller and facilitates communication with the microprocessor.

Hand Pose Tracking & Analysis The pose tracking module received a 15-DOF finger joint position set (i.e., three joints per finger) from the mocap device via its proprietary C++ SDK. Joint composition and indices remain consistent, compatible across any combination of other hand tracking systems providing joint positions, opening rooms for future integration with computer vision-based systems like Google Mediapipe, or virtual reality frameworks such as OpenXR, and Meta Quest SDK. Since our glove is capable to impose only lumped flexion/extension for each finger, result of synergetic motion consist with multiple sub-finger joints. To simplifying the representation, values from the sum of flexion angles for all joints, the metric known as composite finger flexion(CFF), is used. This representation is interchangeably used in the system to compute either target poses from interaction context (any hand motion that are synthesized, recorded, synced with external hand tracking module being imposed to user) or measured poses (feedback data from the mocap device) to compute the 5 DoF per-finger pose errors.

2-Stage Cascaded Closed-Loop Control Dynamics of pneumatic, soft robotic systems require a control scheme with robust feedback and error compensation for handling its nonlinearity. We also have the human user wearing the glove as the core of the system, timely and responsive operation is critical for ensuring pleasant interaction experience. With those requirement in mind, from given target finger flexion angle θ_{tar}, the pressure command p_{cmd} is derived from the error between θ_{tar} and acquired feedback θ using standard PID control law:

$$p_{\text{cmd}} = f^{-1}(\theta_{\text{tar}}) + K_p(\theta_{\text{tar}} - \theta) + K_i \int (\theta_{\text{tar}} - \theta)dt + K_d(\dot{\theta}_{\text{tar}} - \dot{\theta})$$

K_p, K_i and K_d are gains of the feedback controller while $f^{-1}(\theta)$ is the feedforward term which is calculated from the inverse of system response $f : p \rightarrow \theta$ which will be elaborated after. The output value of the positional closed-loop, p_{cmd} is feed into another closed loop PD control based law [55] :

$$u = \phi^{-1}(K_d\dot{p}_{cmd} + p_{cmd} + K_p(p - p_{\text{cmd}}))$$

In here, ϕ^{-1} is system-specific inverse relationship compensating of piezoelectric-introduced nonlinearity of pneumatic valve.

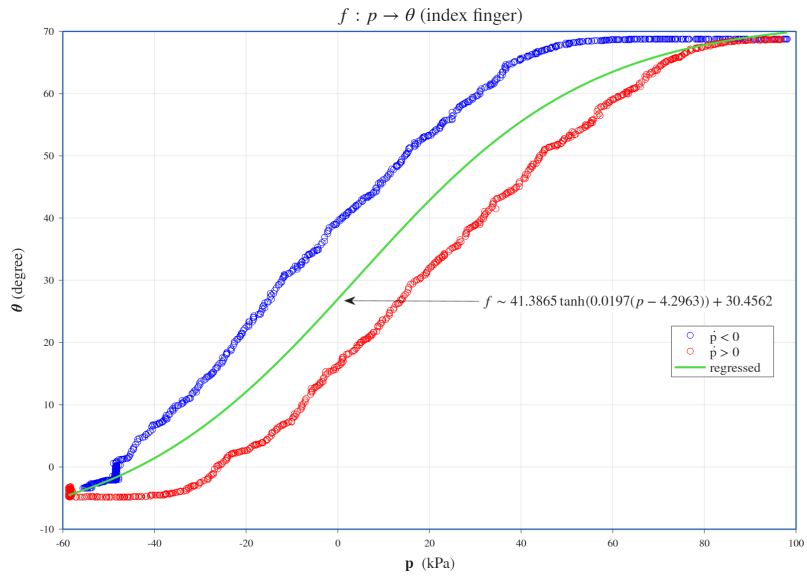

Figure 7: Directional hysteresis of index finger flexion angle over input increasing (Blue, air flowing into), and decreasing (Red, air blowing out) input pressure. A hyperbolic tangent regressed between (Green) is used to formulate the feed-forward term.

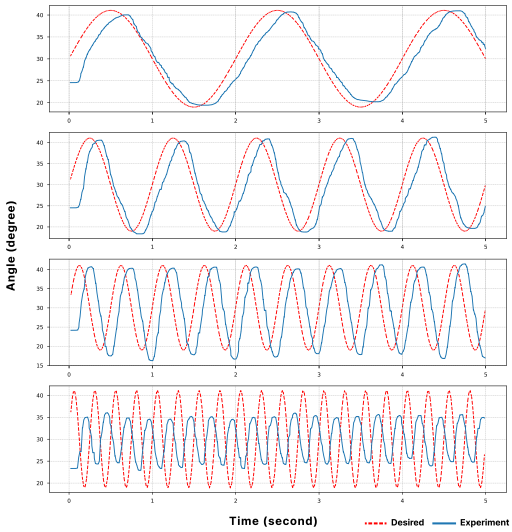

Figure 8: Plots of index finger control responses for 0.5, 1, 2, and 4 Hz.

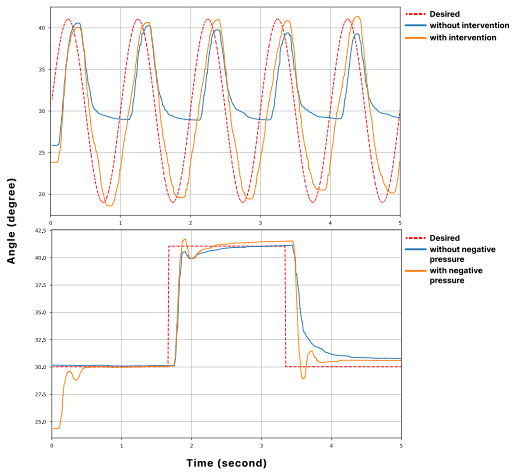

Figure 9: Plots of index finger control responses, with and without the utilization of negative pressure on pneumatic control.

Performance Evaluation Upon system characterization process that identifying aforementioned $f : p \rightarrow \theta$ from the sample pairs of pressure and flexion angle, we observed severe hysteresis depending on the direction of operation. Figure 7 illustrates this: pressurizing the bellow (blue curve)

follows a different flexion trajectory than depressurizing (red curve). To linearize this response, we identified a feedforward term $f^{-1}(\theta)$ fitted as a hyperbolic tangent through the hysteresis midpoint (green curve). This feedforward component, combined with PID feedback, reduces residual tracking error and improves response bandwidth.

Overall, with the optimal feedback gains of PID control determined empirically, the system exhibits an end-to-end finger flexion command to actual motion latency of 140 ms as from the Figure 9 measured from the initiation of target pose update to the point at which the tracking error stabilized within ±5%.

3.4 User Study

In this study, we explore the user experience of glove operation in two potential use scenarios. The first scenario involves gestural communication between LLM and user in a binary decision situation. Depending on the polarity of the LLM's response, the appropriate gesture is commanded; for example, a thumbs-up gesture for a positive response and a fist gesture for a negative response. Here, users may choose to comply with or reject the given constraint, whose intensity is naturally conveyed as haptic feedback (see Figure 10). Their decision can be explicitly confirmed by, for instance, pressing a key, and the final pose is recorded by the mocap system.

The second scenario demonstrates agent motion assistance and intervention during ADL. When users interact with everyday objects, the agent intervenes to assist with grasping objects or inhibit touching certain objects. We select a predefined set of hand-object configurations and then encode these configurations with ArUco markers [56]. A grasp pose is selected by checking the proximity between the markers attached to the real world object and one on the hand (see Figure 11). The selected pose is safely reproduced by the glove according to additional decision rules, including priority, conflict resolution, and timeout constraints.

All features in the experimental scenarios were fully implemented and functional, but we deactivated the feedforward term and give explicit pressure command predetermined per gesture to ensure stable and simplified operation. This user study was approved by the Institutional Review Board at the authors' institution.

Experimental Design & Condition For both scenarios, the independent variable was *Gesture Intervention*. The first scenario consisted of two levels that differed in the presence or absence of gesture responses (text-only and text+pose conditions). The second scenario also consisted of two conditions with and without physical interventions (no-intervention and with-intervention).

We designed a set of user experience questionnaire for each scenario, modifying questions from previous studies [57, 58, 59]. Participants completed three common questions: agency [57, 58], intuitiveness, and satisfaction. We additionally asked about novelty in the first scenario and matched expectation [59, 60] in the second scenario. All questions were rated on a 7-point Likert scale ranging from -3 (strongly disagree) to 3 (strongly agree).

The experiment used a one-factor within-subjects design for each scenario. We counterbalanced the presentation order in each scenario with a balanced Latin square. The entire process of completing the two interaction scenarios took approximately 50 minutes.

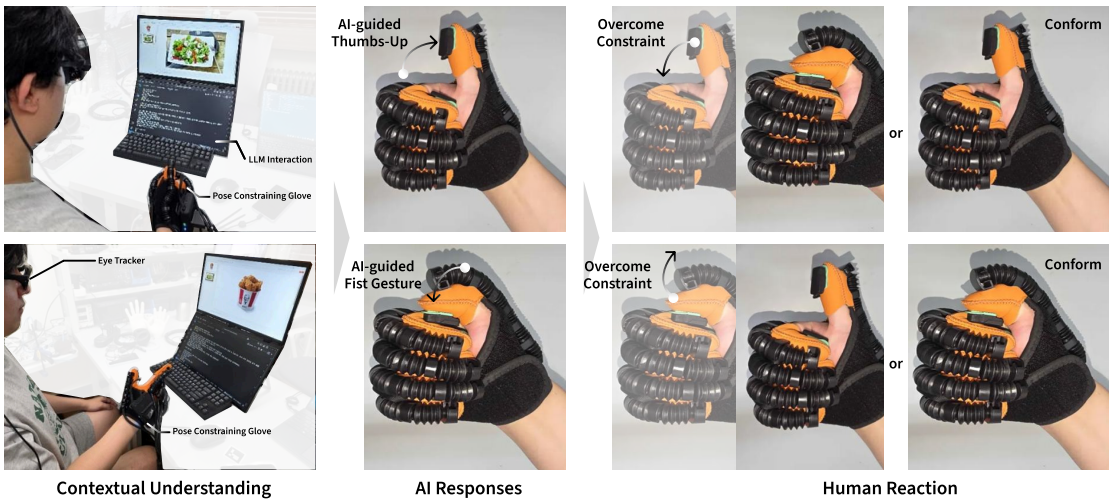

Figure 10: Overview of experiment 1. Poses transmitted from AI agents are illustrated. Then, there are basically two options for users, either overcome or conform to the constraint. The gradient-arrows indicate start (white) and end (black) of a motion.

Participants We recruited 16 right-handed participants (12 male, 4 female; age M = 23.1 years and SD = 1.8 years) from the authors' institution. Each participant was paid USD 10. Eleven out of sixteen participants reported using LLMs everyday, and the remaining five reported at least three to four times per week. Before the experiment, participants were informed of the procedure via a written document and signed a consent form.

3.4.1 Scenario 1: Positive-Negative Gesture Communication

Procedure Participants performed a daily task of selecting diet-appropriate foods with the assistance of an AI agent (Figure 10). Participants wore the eye tracker and the pose glove. At the start of each trial, they were shown a photo of either a bucket of fried chicken or a salad, selected at random. For convenience, we pre-typed the text message, "I'm on a diet. Should I eat the food in this photo?" Participants pressed the enter key to confirm their food selection and submit the text query. The AI agent (GPT-4o) responded in real time via the ChatGPT response API, following our JSON query (in our supplementary materials). The AI agent responded with both textual answers and a gesture via the haptic glove. Following the agent's response, participants could also indicate their agreement or disagreement using either a hand gesture or text, depending on the condition. Specifically, participants could decide to comply with or resist the pose constraint imposed by the glove. They then repeated the trial with the other food and were asked to comply if they had previously resisted, and vice versa. After completing two trials, they answered the questionnaire and took a one-minute break. All participants completed both the text-only and text-and-pose conditions. A semi-structured interview was conducted after the completion of the first scenario twice.

Figure 11: Overview of experiment 2. Depending on motion trajectory toward the frying pan, different poses are assigned to assist or inhibit subsequent actions.

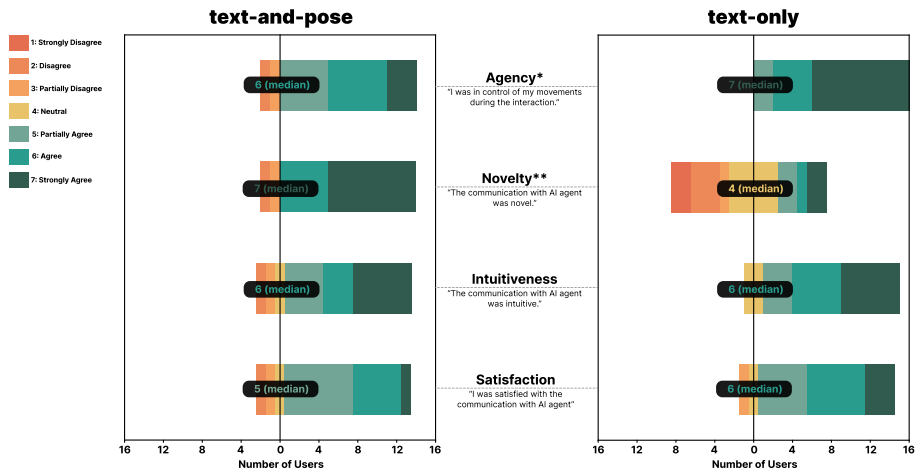

Figure 12: Results of experiment 1 with the distribution of participants' responses on the 7-point Likert scale for Agency, Novelty, Intuitiveness, and Satisfaction. Statistical significance is indicated by asterisks. (*: $p < .05$; **: $< .01$; ***: $p < .001$)

Results and Discussion We performed Wilcoxon signed rank tests with a significance level of 0.05. As expected, we observed a significant effect of physical intervention on agency ($W = 3.00, P = .022$). The median agency scores with and without gestural responses were 6 and 7, respectively. The results showed that the presence of intervention significantly increased perceived novelty ($W = 66.00, P = .003$), while no significant differences were found in intuitiveness ($W = 6.50, P = .461$) or satisfaction ($W = 20.50, P = .138$).

Participants reported that the system slightly limited their sense of free will but considered this AI-human communication modality to be novel (Figure 12). We analyzed interview responses to understand the reasons behind the lack of significance in intuitiveness and satisfaction. A few participants mentioned a preference for text outputs due to simplicity and clarity. Some participants noted that the system would be much more informative if it could perform more than two gestures. There were also comments that the gesture-constraining force was a bit stronger than expected and

that they had to wear too many devices.

3.4.2 Scenario 2: Hand-Object Interaction

Procedure Participants were equipped with the glove. We informed them that they would be interacting with everyday objects with the assistance of another AI agent (Figure 11). Participants were asked to reach their gloved hand toward a frying pan and a wine glass. We attached three fiducial markers to the pan: one on the handle, one on the base, and one representing hot oily foods. They repeated with or without the marker representing hot oily foods in the pan. For example, all fingers are extended to signal caution when they approaches a pan with a hot food marker. For the wine glass, a marker was place on the top to assist a grasping pose suitable for holding the stem. They experienced as much as they want and then answered the questionnaire. After finishing the trial, they had a one-minute break. All participants completed the trials with the two conditions of with and without intervention. We used pre-recorded hand poses in the user study to ensure stable and safe operation.

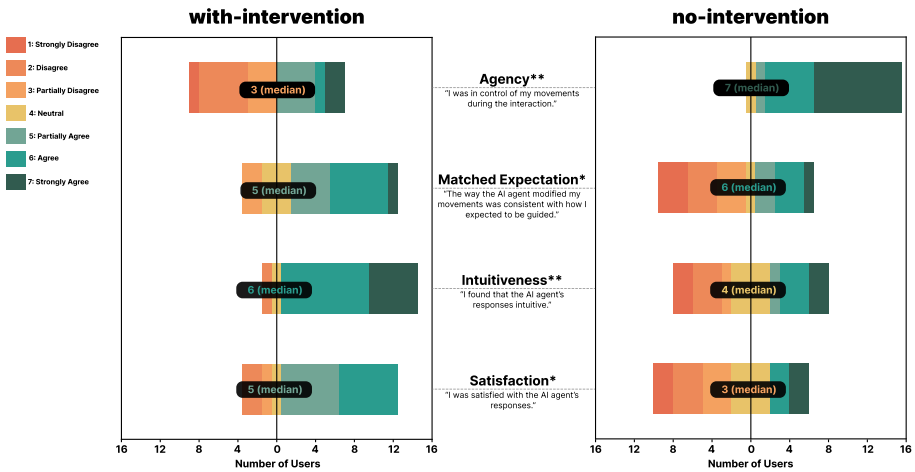

Figure 13: Results of experiment 2 with the distribution of participants' responses on the 7-point Likert scale for Agency, Matched Expectation, Intuitiveness, and Satisfaction. Statistical significance is indicated by asterisks.(*: $p < .05$; **: $< .01$; ***: $p < .001$)

Results and Discussion In line with the previous scenario, we observed a significant effect of physical intervention on agency ($W = 0.00, P = .002$), and the median agency scores with and without intervention were 3 and 7, respectively. The results showed that the presence of intervention significantly increased all other measures: matched expectation ($W = 98.00, P = .032$), intuitiveness ($W = 74.50, P = .006$), and satisfaction ($W = 74.50, P = .044$).

Interestingly, participants reported a greater loss of agency in the second scenario than the first, but they found the intervention to be intuitive and satisfactory (Figure 13). In the interview, all participants responded that the intervention was informative and beneficial, but some of them again reported that constraining forces in our experimental setting were too strong. Interestingly,

some participants highlighted that it would be much more useful if these physical interventions took place based on the inference of safety contexts, e.g., for blind users, firefighters, or children.

3.5 Discussion

3.5.1 Designing AI-Human Gestural Interactions

Our findings highlight several important design considerations.

First, our results support the common understanding that gesture-based communication is essential, but it can be confusing when used alone. According to participants' subjective responses, they expect capabilities that are difficult to achieve with conventional methods, such as enhanced safety and agile responses. Similar to how gestures complement speech in human-human communication, AI gestures in our system are best received when used to augment—not replace—conventional text and voice outputs.

Second, scenario-sensitive intervention emerged as a key requirement. While the system's physical feedback was perceived as even more forceful in the second scenario, our participants found such interventions satisfying. Moreover, they expressed a desire for more adaptive behaviors based on the inference of user intent and context. This points to the need for intelligent decision-making layers that determine *when* and *how* to intervene based on situational relevance.

Third, participants' qualitative feedback emphasized that the perceived strength of the constraint strongly influenced their experience. In future studies, we recommend to carefully calibrating the degree of intervention in conjunction with context and its importance.

3.5.2 Limitations and Future Work

Despite the promising results, a few limitations remain. Current bellow-type actuators on the gloves are still bulky and mechanically rigid, which reduces user comfort during prolonged use. Although we were able to produce a desktop-scale, fully functional prototype, miniaturization and silent operation remain important challenges. In future iterations, replacing the bellows with softer actuators and adopting more compact compressed-air cartridges could further enhance both wearability and comfort during extended use.

3.6 Conclusion

By introducing pneumatic glove and its context-aware operation, this work advances a new interaction paradigm—to our knowledge—that lets AI agents convey semantically grounded, context-aware information through physical hand pose intervention.

Our results demonstrate that AI-mediated pose assistance can enhance the subjective experience of novelty and intuitiveness. In summary, this study makes the following contributions: (1) We propose a novel AI interaction paradigm that leverages hand pose recognition and assistance as a bi-directional communication channel between a user and an AI agent. (2) We validate the technical feasibility of a pneumatic hand pose constraining hardware platform with a software framework that supports AI-human gestural communication with vision-enabled AI agents.

4 Discussion

This section synthesizes the contributions of the exoskeleton and glove, highlights design trade-offs and lessons learned, and reflects on the broader implications for embodied human–AI interaction.

4.1 Synthesis of Exoskeleton and Glove

The kinesthetic exoskeleton and pneumatic glove were developed as separate systems, yet they can be interpreted as a synergistic "upper-limb haptic stack" when considered together.

Kinesthetic Exoskeleton Provides joint-space coverage at the shoulder and elbow, featuring:

- Biomechanical compatibility via a remote-center-of-motion mechanism
- Intrinsic transparency through backdrivable quasi-direct-drive actuators
- A virtual force rendering framework enabling posture-dependent modulation of muscle engagement

Pneumatic Glove Provides local, finger-level pose control, featuring:

- Bidirectional pressure actuation via dual reservoirs
- Motion-capture integration for closed-loop feedback
- Sub-200 ms response latency enabling real-time gestural AI communication

From a robotics perspective, this division of labor mirrors the common pairing of an arm manipulator with a pneumatic gripper affixed at the end-effector. However, in the present work, the "manipulator" is not a remote robot but a body-worn exoskeleton, and the "gripper" is not merely a grasping tool but a gestural channel through which an AI agent can write and read subtle motor patterns on the user's own hand. The exoskeleton provides global, joint-space control and proprioceptive sensing of the limb configuration, while the glove provides high-gain, contact-proximal modulation of finger pose and perceived compliance.

Although a fully integrated prototype was not realized within the scope of this thesis, the two subsystems were intentionally designed with compatible mechanical assumptions and software abstractions. One practical perspective is that the exoskeleton can serve as a rigid mounting point and force ground for other wearable interfaces such as the proposed glove, potentially offloading pneumatic hardware from the hand and improving wearability. More broadly, the exoskeleton can be viewed as a proximal infrastructure for future wearable HCI: a body-conformal, force-grounded frame that offers (i) repeatable attachment geometry and alignment, (ii) a mechanically reliable load path for transmitting forces without relying on soft tissue alone, and (iii) a shared platform for routing power, communication, and actuation resources. This "backbone" perspective shifts the design burden away from distal segments—where added mass and bulk most strongly penalize dexterity—toward a proximal structure that can host heavier components (e.g., reservoirs, pumps, batteries, compute modules) while leaving the hand and fingers lightweight. As a result, diverse add-on interfaces such as tactile displays, instrumented cuffs, motion-capture markers, or auxiliary end-effectors can be integrated in a modular manner while preserving wearability and calibration consistency across sessions. On the software side, both systems expose clear interfaces—joint-space torque commands for the arm and hand-pose targets for the glove—that can be coordinated by a higher-level controller, such as a teleoperation, shared autonomy, or imitation learning.

In combination, these characteristics indicate a path toward covering most of the upper-limb functional space: reaching and orienting the arm while simultaneously shaping grasp posture and gestural nuance. Importantly, the exoskeleton's ability to render posture-dependent joint torques provides a proximal context in which distal finger-level interventions can be perceived not as isolated cues, but as components of a coherent whole-arm interaction. Together, the two devices outline a holistic approach to upper-limb haptics spanning from the shoulder to the fingertips, while remaining modular enough to support incremental integration and future extensions.

4.2 Meeting the Demand for Capable Human–Machine Interfaces

Recent breakthroughs in data-driven methods have opened a new frontier in robotics centered on training large-scale probabilistic robot policies that transfer and deploy human skills to robots. A major paradigm is imitation learning, also known as learning from demonstration (LfD), which critically depends on high-quality human demonstrations that guide a robot through a task. Vision-based motion capture and teleoperation remain the dominant means of collecting such demonstrations and converting them into time-sequenced robot actions.

However, in many systems the human operator receives little or no haptic feedback, resulting in a physically open-loop process that can degrade pose synchronization, weaken contact perception, and reduce demonstration quality. This limitation has fueled increasing interest in human–machine interfaces that both capture human motion and close the loop with physical feedback.

One representative example is the *ALOHA* system introduced by Zhao et al. [61], which uses two identical robot arms in a master–follower configuration: as the demonstrator moves the master arm, the follower reproduces the motion while recording kinematics and onboard camera views. A key contribution of that work is the degree to which the hardware platform has been open-sourced and packaged into accessible configurations, enabling follow-up work using consistent system assumptions. This trend continued with *Mobile ALOHA*, which integrates a moving base [62]. Beyond arm-level interfaces, Xu et al. [63] proposed the hand-worn mechanical glove *DexUMI*, providing rich sensing of hand pose. Pushing anthropomorphism further, Fang et al. [64], Zhong et al. [65, 66] developed sophisticated upper-limb exoskeleton interfaces intended to replicate a wide range of human motion for teleoperating humanoid robots. Yet, these systems largely remain at the kinematic motion capture level, with limited bilateral force feedback, though their shared focus reflects a broader anthropomorphism trend, both the interface hardware and its robotic counterpart increasingly mirror human kinematics, whether as manipulators or humanoids.

A natural next step is to extend these kinematic interfaces into *haptic* ones. A full upper-limb haptic system that combines a kinesthetic exoskeleton (for distributed shoulder–elbow torques rendering Cartesian contact forces) with a hand-level glove interface (for grasp posture and fingertip-level constraints) offers a promising channel for both data acquisition and teleoperation. Such a system can capture demonstrations that include not only "where the arm moved," but also "what forces were exchanged," enabling higher-fidelity demonstrations for contact-rich tasks and more intuitive teleoperation through closed-loop kinesthetic feedback.

4.3 Robotic Embodiment of Intelligence and of Ourselves

A central motivation of this thesis has been the emerging paradigm of *robotic embodiment*: humans and AI agents acting through physical bodies that are not co-located with them, yet are tightly coupled to their sensorimotor or decision-making loops. In most present-day systems, this coupling is dominated by vision and kinematics: the human provides joint trajectories for a remote robot, or an AI policy maps camera streams to motor commands, with little or no bilateral physical feedback. The upper limb is reduced to an input device, and the robot to a distant effector.

The exoskeleton and glove developed here suggest a different framing. Rather than treating the human as a mere demonstrator and the robot as a separate embodiment, the two devices form a composite *avatar* that can be shared. When the exoskeleton backdrives freely and the glove modulates hand pose with controllable stiffness, the human can experience remote forces and constraints as if they were acting directly on their own arm and hand. Conversely, an AI policy controlling the same system can imprint its decisions as physically perceivable guidance or resistance, creating a closed loop in which human and artificial intelligence can negotiate control over the same limb.

In this sense, the thesis addresses two intertwined notions of embodiment:

- **Embodiment of Intelligence.** Data-driven robot controllers, such as imitation learning policies, need physical instantiations to reveal their capabilities and limitations. An upper-limb haptic interface that can both capture rich motion data and replay force conditions offers a higher-fidelity substrate for training and evaluating such policies than kinematics-only teleoperation. It enables demonstrations that include not just "where the hand went" but also "what it felt like" to interact with objects and environments.

- **Embodiment of Ourselves.** For the human user, the same interface can act as a vehicle for projecting agency into remote or virtual spaces. High-bandwidth bidirectional coupling of arm and hand makes it possible to manipulate tools, surfaces, and avatars at a distance while retaining a sense of bodily ownership over those actions. As communication and work increasingly span physical locations, such surrogate embodiments may become as routine as video calls are today.

Ultimately, high-fidelity upper-limb haptics should be viewed not only as a technical refinement but as an enabling step toward physically grounded human–AI interaction. If we are to build a saddle for riding the ever-accelerating technological advance, high-fidelity haptic feedback for the upper limb is more than a technical refinement; it is a necessary step by which we can truly benefit from our creations, rather than be replaced by them. When the barrier between local and remote dissolves, and our physical constraints become secondary to our digital reach, one may say that our existence has well transitioned into another mode—what I have called *Homo ubiquitarius*: a transient category preceding whatever comes next soon.

5 Conclusions

This thesis developed modular building blocks for full upper-limb haptics by combining a kinesthetic exoskeleton for the shoulder–elbow chain with a bidirectionally actuated pneumatic glove for the hand. Rather than delivering a single, fully integrated device, the work pursued three complementary goals: (1) establishing design principles and concrete implementations for biomechanically compatible and intrinsically transparent exoskeleton joints; (2) realizing a compact, closed-loop pneumatic glove capable of responsive, bidirectional hand pose control; and (3) providing task-grounded evidence—via simulation and user studies—that these components can support interaction paradigms in which both humans and AI agents act through the upper limb.

The exoskeleton's primary contribution is a posture-dependent virtual force rendering framework implemented on a mechanically transparent, backdrivable platform. Preliminary musculoskeletal simulations supported that the proposed rendering strategy can elicit muscle-activation trends comparable to those induced by representative physical loads under matched postural conditions. The pneumatic glove demonstrated that physical hand-pose interventions can serve as a communication channel from an AI agent to a human user. In the conducted user studies, the proposed approach improved perceived intuitiveness and alignment with user expectations compared to visual-only feedback within the evaluated interaction tasks.

Although a fully integrated prototype combining the exoskeleton and glove remains future work, their complementary capabilities—joint-space force rendering and finger-level pose control—establish practical building blocks toward full upper-limb haptic interfaces. More broadly, the results suggest that physically grounded upper-limb interfaces can expand how humans and artificial agents share control, guidance, and intent through embodied interaction.

Appendix A Optimized Kinematic Synthesis of Robotic System

A.1 Modern Kinematics Description Method: Product of Exponentials (POE)

Among numerous established description methods for interpreting multiple-body dynamics problems of robotics, Lie group theory has established itself as the modern rule of thumb. Unlike using geometry-describing parameters from Denavit-Hartenberg (DH) convention, which often suffers from artifacts such as singularities and gimbal locks, the Product of Exponentials (POE) formula instead allows robust representation of robot kinematics free from those problems [67, 68, 69].

In this framework, the configuration of a robot is defined on the 'Special Euclidian Group' $SE(3)$. A critical distinction must be made between reference frame $SE(3)$ based on, that is whether it is from global reference frame or the local frames attached to individual joints and frames. These spaces form distinct manifold of the configuration space, and element from one cannot be directly related with the one from the other $SE(3)$ even if it represent the same robot motion, due to the difference of describing reference frame.

To navigate between these spaces, we thus utilize the Adjoint map (Ad). The Adjoint operation maps an element from the $SE(3)$ group of one frame to the other, effectively providing the offset transform required to describe a frame with respect to reference frame coordinate. Furthermore, the velocity of these bodies is represented by twists $\hat{\xi}$, elements of another Lie group vector space $\mathfrak{se}(3)$, which is the tangent space spanned at the $SE(3)$ point on the manifold. One powerful advantage of Lie theory is that, there are well established mapping between curved manifold and linear tangent space back-and-forth. Mapping that projects element of linear tangent spaces $\mathfrak{se}(3)$ back onto the curved $SE(3)$ manifold is handled via the exponential map $\exp(\cdot)$, and logarithmic map $\log(\cdot)$ maps item on $SE(3)$ manifold to the designated tangent space.

POE formula representing robot kinematics itself is utilizing all those context to construct subsequent product of $\exp(\hat{\xi}_i q_i) \in SE(3)$ which is a transform by the displacement of moving joint over time. In here, $\hat{\xi}_i$ being i-th joint screw axis defined as an twist of world frame, which multiplied by joint angle q_i representing joint frame's transform done by rotating q_i radian around the axis $\hat{\xi}$.

Finally with POE, $SE(3)$ placement of tool frame starting from initial, 'zero-configuration' ${}^{o}M_{\text{tool},0}$, moved by configuration q is expressed as:

$$ {}^{o}M_{\text{tool}} = \exp(\hat{\xi}_1 q_1)\exp(\hat{\xi}_2 q_2)...\exp(\hat{\xi}_n q_n){}^{o}M_{\text{tool},0} \tag{1} $$

This theoretical foundation will further allows for robust gradient-based optimization.

A.2 Optimization of Joint Placements for Maximized Human Kinematics Compatibility

Following Lie-group-based kinematic synthesis methods [27, 26], the formulation process assumes the exoskeleton is modeled as a serial kinematic chain with n revolute joints, described using $SE(3)$ POE formulation. Each joint is assumed to be revolute type, rotating about its own local z-axis, thus its local joint screw $\xi_{local} = [v^\top \ \omega^\top]^\top = [0\ 0\ 0\ 0\ 0\ 1]^\top$. Let $\rho \in \mathbb{R}^{6n}$ denote the structural parameters of the robot, where each block $\rho_i \in \mathbb{R}^6$ is the vector form of a twist in the

parent frame specifying the static, local placement of the joint i. Then the term of each joint in (1) can be expressed as:

$$\exp(\hat{\xi}_i q_i) = \mathrm{Ad}_{W\leftarrow P_i} \cdot {}^{parent}M_i \cdot \exp(\hat{\xi}_{local} q_i) = \mathrm{Ad}_{W\leftarrow P_i} \cdot \exp(\hat{\rho}_i) \cdot \exp(\hat{\xi}_{local} q_i) \tag{2}$$

with $\hat{\cdot}$ the standard se(3) "hat" operator, $\mathbb{R} \to \mathfrak{se}(3)$, and $\mathrm{Ad}_{W\leftarrow P_i}$ is the adjoint operator mapping a local placement in the parent frame P_i into an equivalent $SE(3)$ item expressed in the World frame. Notice that Given joint angles $q \in \mathbb{R}^n$, the POE formula yields the tool pose in the World frame that depends on q and ρ:

$$^{o}M_{\text{tool}}(q;\rho) \in SE(3), \tag{3}$$

constructed as the ordered product of (2) as described in Section A.1.

For each human sample $k \in \{1,\dots,m\}$ we have the target tool pose $M_{\text{target},k} \in SE(3)$ expressed in the World frame. We define the pose error as an element of $SE(3)$,

$$M_{e,k}(q_k,\rho) \;=\; M_{\text{target},k}\, {}^{o}M_{\text{tool}}(q_k;\rho)^{-1}, \tag{4}$$

and its logarithm as a twist in the World frame:

$$e_k(q_k,\rho) \;=\; \mathrm{Log}(M_{e,k}(q_k,\rho))^{\vee} \in \mathbb{R}^6 \tag{5}$$

The overall cost function aggregates the squared twist norms:

$$F(\rho,\{q_k\}) \;=\; \frac{1}{2}\sum_{k=1}^{m} \big\|e_k(q_k,\rho)\big\|^2 \tag{6}$$

Lie-group LM formulation for structural parameters The structural parameters ρ live in the tangent spaces of each joint's parent frame $SE(3)$, while the pose errors e_k live in the tangent space of the World frame. At each iteration we consider a perturbation $\delta\rho_i \in \mathbb{R}^6$ applied in the parent frame:

$${}^{parent}M_i^{\text{new}} \;=\; {}^{parent}M_i \,\exp\big(\widehat{\delta\rho_i}\big) \quad\Longleftrightarrow\quad \rho_i^{\text{new}} \;=\; \hat{\rho}_i + J_{log}\delta\rho_i \tag{7}$$

$J_{\log}$ is Jacobian of the logarithm map operation of $\mathfrak{se}(3)$ to $\mathbb{R}^6$, linear vector space of error updated by LM optimization. With the treatment of perturbation as derived from the Baker–Campbell–Hausdorff (BCH) formula [67], (7) reduces to the familiar additive update $\rho_i^{\text{new}} \approx \rho_i + \delta\rho_i$ for small $\|\delta\rho_i\|$, but remains well-defined and singularity-free $SE(3)$ for larger motions.

To use Levenberg–Marquardt (LM) algorithm for optimization, we require the Jacobian that maps structural perturbations $\delta\rho$ to pose errors δe [26]. For a given sample k, a perturbation of i-th joint placement $\delta\rho_i$ induces a perturbation of the tool pose, which can be represented as a body twist in the World frame. Using adjoint mappings between frames, the contribution of joint perturbations to the error variation can be written as

$$\delta e_k \;\approx\; J_{\rho,k}\,\delta\rho, \quad J_{\rho,k} = \Big[J_{\rho_1,k}^{\top}\; J_{\rho_2,k}^{\top}\; \dots\; J_{\rho_n,k}^{\top}\Big]^{\top} \;\in\; \mathbb{R}^{6\times 6n}, \tag{8}$$

where each 6×6 block of joint i is given by

$$J_{\rho_i,k} \;=\; J_{\log}\big(e_k\big)\,\mathrm{ad}_{W\leftarrow P_i}, \tag{9}$$

Intuitively, $\mathrm{ad}_{W \leftarrow P_i}$ transports the effect of a structural perturbation twist at joint i to the task-space error twist at World frame, while $J_{\log}$ accounts for the curvature of the logarithm map.

Stacking all samples yields the global residual vector

$$r(\rho, \{q_k\}) = \begin{bmatrix} e_1^\top & e_2^\top & \dots & e_m^\top \end{bmatrix}^\top \in \mathbb{R}^{6m}, \tag{10}$$

and the global Jacobian $J_\rho \in \mathbb{R}^{6m \times 6n}$. The LM step for the structural parameters solves the damped normal equations

$$\left(J_\rho^\top J_\rho + \lambda I\right) \delta\rho \;=\; -J_\rho^\top r, \tag{11}$$

with damping factor $\lambda \geq 0$. From this, the choice Gauss-Newton method and gradient descent is adaptively made throughout the iterations. The solution $\delta\rho$ is a stacked vector of local twists $\delta\rho_i$, which are then applied via the Lie-group update (7). In this way the linear algebra of LM operates purely in Euclidean parameter space, while all geometric operations (composition, parallel transport, logarithms and exponentials) are consistently handled on $SE(3)$, avoiding singularities and ensuring continuity of the kinematic chain throughout the optimization.

Algorithm 1 Simultaneous Optimization: Joint Placements & Trajectory Tracking

Require: Target human poses $\{M_{\text{target},k}\}$, damping schedule λ

$\rho \leftarrow \rho_0$ (random or heuristic initialization)
Initialize $\{q_k\}$ by inverse kinematics for each $M_{\text{target},k}$ given ρ
repeat
  for $k = 1$ to m **do**
    Compute tool pose ${}^{o}M_{\text{tool}}(q_k; \rho)$ via POE
    $M_{e,k} \leftarrow M_{\text{target},k}\, {}^{o}M_{\text{tool}}(q_k; \rho)^{-1}$
    $e_k \leftarrow \mathrm{Log}(M_{e,k})^\vee$
    Compute Jacobian blocks $J_{q,k}(q_k, \rho)$ and $J_{\rho,k}(q_k, \rho)$ using adjoint mappings and $J_{\log}$
    Update q_k by a Gauss–Newton step using $J_{q,k}$ (inner IK refinement)
  end for
  Stack residuals $r \leftarrow [e_1^\top, \dots, e_m^\top]^\top$
  Stack structural Jacobian $J_\rho \leftarrow [J_{\rho,1}^\top, \dots, J_{\rho,m}^\top]^\top$
  Solve $(J_\rho^\top J_\rho + \lambda I)\, \delta\rho = -J_\rho^\top r$ for $\delta\rho$
  for $i = 1$ to n **do**
    Update joint placement: ${}^{parent}M_i \leftarrow \exp(\widehat{\delta\rho_i})\, {}^{parent}M_i$
    Reparametrize: $\rho_i \leftarrow \mathrm{Log}({}^{parent}M_i)^\vee$
  end for
  Adjust λ according to LM gain ratio (decrease on successful step, increase otherwise)
until convergence of $F(\rho, \{q_k\})$ or maximum iterations reached

Simultaneous optimization with tracking configurations The joint configurations $\{q_k\}$ are optimized jointly with the structural parameters. Conceptually, $F(\rho, \{q_k\})$ is minimized over both sets of variables. In practice, we alternate between updating q_k and ρ in a Gauss–Newton / LM fashion:

1. For fixed ρ, solve an inverse-kinematics (IK) problem for each sample k to update q_k so as to reduce $\|e_k(q_k, \rho)\|$. This uses the standard manipulator Jacobian with respect to q_k.

2. For the updated $\{q_k\}$, assemble r and J_ρ, solve (11) for $\delta\rho$, and update ρ via (7).

This "simultaneous" viewpoint guarantees that the structural design (joint placements) and the tracking configuration (joint angles along the human trajectory) are consistently co-optimized. The resulting exoskeleton achieves a workspace that closely matches the human motion data while minimizing geometric misalignment, thereby suppressing parasitic interaction forces at the attachment points and improving comfort and safety.

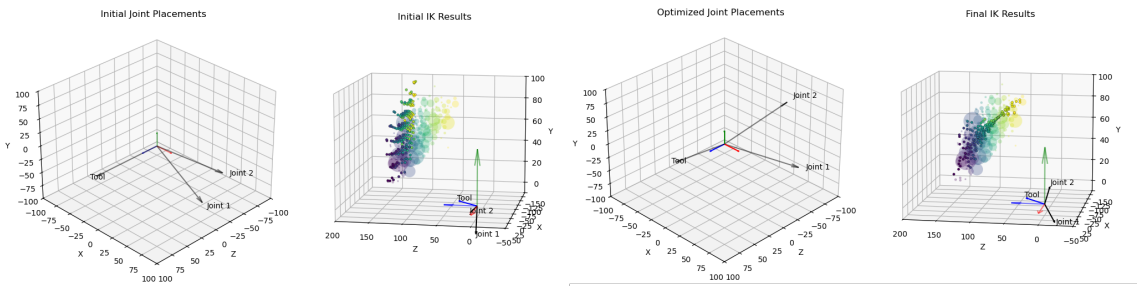

Figure 14: Example of simultaneous optimization framework in use, applied for two revolute joint covering 3 DoF SC joint. Left pair: Heuristically selected initial joint placements and its kinematic coverage; Right pair: Same kind of plot with optimized result. Shaded spheres are each indicating target pose sample location and its size displaying density, with IK solution points with matching color of its target.

References

[1] Inrak Choi, Eyal Ofek, Hrvoje Benko, Mike Sinclair, and Christian Holz. CLAW: A Multifunctional Handheld Haptic Controller for Grasping, Touching, and Triggering in Virtual Reality. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–13, Montreal QC Canada, 2018. ACM. ISBN 978-1-4503-5620-6. doi: 10.1145/3173574.3174228.

[2] Robert Kovacs, Eyal Ofek, Mar Gonzalez Franco, Alexa Fay Siu, Sebastian Marwecki, Christian Holz, and Mike Sinclair. Haptic PIVOT: On-Demand Handhelds in VR. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, UIST '20, pages 1046–1059, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 978-1-4503-7514-6. doi: 10.1145/3379337.3415854.

[3] Anthony Steed, Sebastian Friston, Vijay Pawar, and David Swapp. Docking Haptics: Extending the Reach of Haptics by Dynamic Combinations of Grounded and Worn Devices. In *26th ACM Symposium on Virtual Reality Software and Technology*, pages 1–11, Virtual Event Canada, November 2020. ACM. ISBN 978-1-4503-7619-8. doi: 10.1145/3385956.3418943.

[4] Nathan DeVrio and Chris Harrison. Reel Feel: Rich Haptic XR Experiences Using an Active, Worn, Multi-String Device. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–20, Yokohama Japan, April 2025. ACM. ISBN 979-8-4007-1394-1. doi: 10.1145/3706598.3713615.

[5] Gareth Barnaby and Anne Roudaut. Mantis: A scalable, lightweight and accessible architecture to build multiform force feedback systems. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, UIST '19, pages 937–948, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 978-1-4503-6816-2. doi: 10.1145/3332165.3347909.

[6] Simon Søndergaard Christensen. User-Centered Modelling and Design of Assistive Exoskeletons. 2023. doi: 10.54337/AAU528218520.

[7] Manuel Cardona and Fernando E. Serrano. *Rehabilitation Robotics: Kinematics, Dynamics, and Control Techniques*. Springer Nature Switzerland, Cham, 2025. ISBN 978-3-031-83654-1 978-3-031-83655-8. doi: 10.1007/978-3-031-83655-8.

[8] I. Scott MacKenzie. *Human-Computer Interaction: An Empirical Research Perspective*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1st edition, 2013. ISBN 978-0-12-405865-1.

[9] Pedro Lopes. What if the "I" in HCI stands for Integration? In *Adjunct Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology*, pages 1–5, Busan Republic of Korea, September 2025. ACM. ISBN 979-8-4007-2036-9. doi: 10.1145/3746058.3762829.

[10] Heather Culbertson, Samuel B. Schorr, and Allison M. Okamura. Haptics: The Present and Future of Artificial Touch Sensation. *Annual Review of Control, Robotics, and Autonomous Systems*, 1(1):385–409, May 2018. ISSN 2573-5144, 2573-5144. doi: 10.1146/annurev-control-060117-105043.

[11] Inrak Choi, Heather Culbertson, Mark R. Miller, Alex Olwal, and Sean Follmer. Grabity: A Wearable Haptic Interface for Simulating Weight and Grasping in Virtual Reality. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, pages 119–130, Québec City QC Canada, October 2017. ACM. ISBN 978-1-4503-4981-9. doi: 10.1145/3126594.3126599.

[12] Yuxin Ma, Tianze Xie, Peng Zhang, Hwan Kim, and Seungwoo Je. AirPush: A Pneumatic Wearable Haptic Device Providing Multi-Dimensional Force Feedback on a Fingertip. *Conference on Human Factors in Computing Systems - Proceedings*, 2024. doi: 10.1145/3613904.3642536.

[13] Cathy Fang, Yang Zhang, Matthew Dworman, and Chris Harrison. Wireality: Enabling Complex Tangible Geometries in Virtual Reality with Worn Multi-String Haptics. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–10, Honolulu HI USA, April 2020. ACM. ISBN 978-1-4503-6708-0. doi: 10.1145/3313831.3376470.

[14] Alejandro Jarillo Silva, Omar A. Domínguez-Ramírez, Vicente Parra-Vega, and Patricio Ordaz. PHANToM OMNI haptic device: Kinematic and manipulability. *2009 Electronics, Robotics and Automotive Mechanics Conference (CERMA)*, 2009. doi: 10.1109/CERMA.2009.55.

[15] Andreas Tobergte, Patrick Helmer, Ulrich Hagn, P. Rouiller, Sophie Thielmann, S. Grange, Alin Albu-Schäffer, François Conti, and Gerd Hirzinger. The sigma.7 haptic interface for MiroSurge: A new bi-manual surgical console. *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2011. doi: 10.1109/IROS.2011.6094433.

[16] Elisa Galofaro, Erika D'Antonio, Nicola Lotti, and Lorenzo Masia. Rendering immersive haptic force feedback via neuromuscular electrical stimulation. *Sensors*, 22(14):5069, 2022.

[17] Sarah Faltaous, Marion Koelle, and Stefan Schneegass. From perception to action: A review and taxonomy on electrical muscle stimulation in hci. In *Proceedings of the 21st International Conference on Mobile and Ubiquitous Multimedia*, MUM '22, page 159–171, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450398206. doi: 10.1145/3568444.3568460. URL `https://doi.org/10.1145/3568444.3568460`.

[18] Choongin Lee, Taeyoon Lee, Jae-Kyung Min, Albert Wang, SungPyo Lee, Jaesung Oh, Chang-Woo Park, and Keunjun Choi. A Proprioceptive Haptic Device Design for Teaching Bimanual Manipulation. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 8862–8868, May 2022. doi: 10.1109/ICRA46639.2022.9811694.

[19] Zhijun Li, Xiang Li, Qinjian Li, Hang Su, Zhen Kan, and Wei He. Human-in-the-Loop Control of Soft Exosuits Using Impedance Learning on Different Terrains. *IEEE Transactions on Robotics*, 38(5):2979–2993, 2022. ISSN 1941-0468. doi: 10.1109/TRO.2022.3160052.

[20] Kenneth Salisbury, Francois Conti, and Federico Barbagli. Haptic rendering: Introductory concepts. *IEEE computer graphics and applications*, 24(2):24–32, 2004. ISSN 1558-1756. doi: 10.1109/MCG.2004.1274058.

[21] Romain Nith, Shan-Yuan Teng, Pengyu Li, Yujie Tao, and Pedro Lopes. DextrEMS: Increasing Dexterity in Electrical Muscle Stimulation by Combining it with Brakes. In *The 34th Annual ACM Symposium on User Interface Software and Technology*, pages 414–430, Virtual Event USA, October 2021. ACM. ISBN 978-1-4503-8635-7. doi: 10.1145/3472749.3474759.

[22] Cristiano Alessandro, Ioannis Delis, Francesco Nori, Stefano Panzeri, and Bastien Berret. Muscle synergies in neuroscience and robotics: From input-space to task-space perspectives. *Frontiers in Computational Neuroscience*, 7, 2013. ISSN 1662-5188. doi: 10.3389/fncom.2013.00043.

[23] Muhammad Ahsan Gull, Mikkel Thoegersen, Stefan Hein Bengtson, Mostafa Mohammadi, Lotte N S Andreasen Struijk, Thomas B Moeslund, Thomas Bak, and Shaoping Bai. A 4-DOF Upper Limb Exoskeleton for Physical Assistance: Design, Modeling, Control and Performance Evaluation. *Applied Sciences*, 11(13):5865, 2021. doi: 10.3390/app11135865.

[24] Michael D. Matthews. *Exoskeletons: State-of-the-Art, Design Challenges, and Future Directions*, pages 78–98. Oxford University PressNew York, 1 edition, May 2020. ISBN 978-0-19-087047-8 978-0-19-751093-3. doi: 10.1093/oso/9780190870478.003.0005.

[25] Yves Zimmermann, Michael Sommerhalder, Peter Wolf, Robert Riener, and Marco Hutter. ANYexo 2.0: A fully actuated upper-limb exoskeleton for manipulation and joint-oriented training in all stages of rehabilitation. *IEEE TRANSACTIONS ON ROBOTICS*, 39(3):2131–2150, 2023. ISSN 1941-0468. doi: 10.1109/TRO.2022.3226890.

[26] Shouhei Shirafuji and Keiichiro Shimamura. Kinematic Synthesis of a Serial Manipulator Using Gradient-Based Optimization on Lie Groups. *IEEE Robotics and Automation Letters*, 2025. doi: 10.1109/LRA.2025.3534064.

[27] Jaewoon Kwon, Seung-Hyun Kim, and Frank C. Park. Physically Consistent Lie Group Mesh Models for Robot Design and Motion Co-Optimization. *IEEE Robotics and Automation Letters*, 2022. doi: 10.1109/LRA.2022.3189806.

[28] Florent Moissenet, Pierre Puchaud, Alexandre Naaim, Nicolas Holzer, and Mickaël Begon. Spartacus: A review and aggregation of reference datasets reporting the normal shoulder girdle kinematics during uniplanar humerus motions. *Journal of Biomechanics*, 189:112642, August 2025. ISSN 00219290. doi: 10.1016/j.jbiomech.2025.112642.

[29] Simon Christensen and Shaoping Bai. Kinematic analysis and design of a novel shoulder exoskeleton using a double parallelogram linkage. *Journal of Mechanisms and Robotics*, 10 (4):41008, 2018. ISSN 1942-4302, 1942-4310. doi: 10.1115/1.4040132.

[30] Wei Qian, Junbei Liao, Linjun Lu, Letian Ai, Miao Li, Xiao-Hui Xiao, and Zhao Guo. CURER: A Lightweight Cable-Driven Compliant Upper Limb Rehabilitation Exoskeleton Robot. *IEEE/ASME transactions on mechatronics*, 28(3):1730–1741, June 2023. doi: 10.1109/tmech.2022.3224423.

[31] Bruno Siciliano and Oussama Khatib. *Springer Handbook of Robotics*. Springer, Berlin, 2008. ISBN 978-3-540-23957-4.

[32] Emek Barış Küçüktabak, Yue Wen, Sangjoon J. Kim, Matthew R. Short, Daniel Ludvig, Levi Hargrove, Eric J. Perreault, Kevin M. Lynch, José L Jose L. Pons, Emek Bars Kucuktabak, Yue Wen, Sangjoon J. Kim, Matthew R. Short, Daniel Ludvig, Levi Hargrove, Eric J. Perreault, Kevin M. Lynch, and José L Jose L. Pons. Haptic Transparency and Interaction Force Control for a Lower Limb Exoskeleton. *IEEE Transactions on Robotics*, 40:1842–1859, 2024. ISSN 19410468. doi: 10.1109/TRO.2024.3359541.

[33] Yves Zimmermann, Alessandro Forino, Robert Riener, and Marco Hutter. ANYexo: A versatile and dynamic upper-limb rehabilitation robot. *IEEE Robotics and Automation Letters*, 4(4):3649–3656, 2019. ISSN 2377-3766, 2377-3774. doi: 10.1109/LRA.2019.2926958.

[34] Gabriel Aguirre-Ollinger and Haoyong Yu. Lower-Limb Exoskeleton With Variable-Structure Series Elastic Actuators: Phase-Synchronized Force Control for Gait Asymmetry

Correction. *IEEE Transactions on Robotics*, 37(3):763–779, 2021. ISSN 1941-0468. doi: 10.1109/TRO.2020.3034017.

[35] Sangok Seok, Albert Wang, David Otten, and Sangbae Kim. Actuator design for high force proprioceptive control in fast legged locomotion. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1970–1975, October 2012. doi: 10.1109/IROS.2012.6386252.

[36] Patrick M Wensing, Albert Wang, Sangok Seok, David Otten, Jeffrey Lang, and Sangbae Kim. Proprioceptive actuator design in the mit cheetah: Impact mitigation and high-bandwidth physical interaction for dynamic legged robots. *Ieee transactions on robotics*, 33 (3):509–522, 2017.

[37] Pablo Lopez Garcia, Elias Saerens, Stein Crispel, Anand Varadharajan, Dirk Lefeber, and Tom Verstraten. Factors influencing actuator's backdrivability in human-centered robotics. *MATEC Web of Conferences*, 366:01002, 2022. ISSN 2261-236X. doi: 10.1051/MATECCONF/202236601002.

[38] Albert Wang and Sangbae Kim. Directional efficiency in geared transmissions: Characterization of backdrivability towards improved proprioceptive control. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1055–1062. IEEE, 2015. ISBN 1-4799-6923-0.

[39] Benjamin G. Katz. *A Low Cost Modular Actuator for Dynamic Robots*. Thesis, Massachusetts Institute of Technology, 2018.

[40] Aïda M. Valevicius, Quinn A. Boser, Ewen B. Lavoie, Craig S. Chapman, Patrick M. Pilarski, Jacqueline S. Hebert, and Albert H. Vette. Characterization of normative angular joint kinematics during two functional upper limb tasks. *Gait and Posture*, 69(July 2018):176–186, 2019. ISSN 18792219. doi: 10.1016/j.gaitpost.2019.01.037.

[41] Steven Macenski, Tully Foote, Brian Gerkey, Chris Lalancette, and William Woodall. Robot Operating System 2: Design, architecture, and uses in the wild. *Science Robotics*, 7(66), 2022. doi: 10.1126/scirobotics.abm6074.

[42] ros2_control framework. URL `https://control.ros.org/rolling/doc/acknowledgements/acknowledgements.html`.

[43] Justin Carpentier, Guilhem Saurel, Gabriele Buondonno, Joseph Mirabel, Florent Lamiraux, Olivier Stasse, and Nicolas Mansard. The pinocchio c++ library – a fast and flexible implementation of rigid body dynamics algorithms and their analytical derivatives. In *IEEE International Symposium on System Integrations (SII)*, 2019.

[44] Federico Califano, Ramy Rashad, Cristian Secchi, Stefano Stramigioli, Federico Califano, Ramy Rashad, Cristian Secchi, and Stefano Stramigioli. On the use of energy tanks for robotic systems. *Springer proceedings in advanced robotics*, 2023. doi: 10.1007/978-3-031-22731-8_13.

[45] Scott L. Delp, Frank C. Anderson, Allison S. Arnold, Peter Loan, Ayman Habib, Chand T. John, Eran Guendelman, and Darryl G. Thelen. OpenSim: Open-source software to create and analyze dynamic simulations of movement. *IEEE Transactions on Biomedical Engineering*, 2007. doi: 10.1109/TBME.2007.901024.

[46] Ajay Seth, Meilin Dong, Ricardo Matias, and Scott Delp. Muscle Contributions to Upper-Extremity Movement and Work From a Musculoskeletal Model of the Human Shoulder. *Frontiers in Neurorobotics*, 13:90, November 2019. ISSN 1662-5218. doi: 10.3389/fnbot.2019.00090.

[47] Panagiotis Polygerinos, Zheng Wang, Kevin C. Galloway, Robert J. Wood, and Conor J. Walsh. Soft robotic glove for combined assistance and at-home rehabilitation. *Robotics and Autonomous Systems*, 73:135–143, 2015. ISSN 09218890. doi: 10.1016/j.robot.2014.08.014. URL `http://dx.doi.org/10.1016/j.robot.2014.08.014`.

[48] Sung Sik Yun, Brian Byunghyun Kang, and Kyu Jin Cho. Exo-glove PM: An easily customizable modularized pneumatic assistive glove. *IEEE Robotics and Automation Letters*, 2 (3):1725–1732, 2017. ISSN 23773766. doi: 10.1109/LRA.2017.2678545.

[49] Brandon J. Caasenbrood, Alexander Y. Pogromsky, and Henk Nijmeijer. Sorotoki: A Matlab Toolkit for Design, Modeling, and Control of Soft Robots. *IEEE Access*, 12(December 2023): 17604–17638, 2024. ISSN 21693536. doi: 10.1109/ACCESS.2024.3357351.

[50] Wenbin Chen, Caihua Xiong, Chenlong Liu, Peimin Li, and Yonghua Chen. Fabrication and Dynamic Modeling of Bidirectional Bending Soft Actuator Integrated with Optical Waveguide Curvature Sensor. *Soft Robotics*, 6(4):495–506, 2019. ISSN 21695180. doi: 10.1089/soro.2018.0061.

[51] Boran Wang, Andrew McDaid, Kean C. Aw, and Morteza Biglari-Abhari. Design and development of a skinny bidirectional soft glove for post-stroke hand rehabilitation. In *Intelligent Systems Conference*, volume 2018-Janua, pages 979–987, 2017. ISBN 9781509064359. doi: 10.1109/IntelliSys.2017.8324248.

[52] Jianwei Lai, Aiguo Song, Jiajin Wang, Ye Lu, Ting Wu, Huijun Li, Baoguo Xu, and Xiangshan Wei. A Novel Soft Glove Utilizing Honeycomb Pneumatic Actuators (HPAs) for Assisting Activities of Daily Living. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 31:3223–3233, 2023. ISSN 15580210. doi: 10.1109/TNSRE.2023.3302612.

[53] Jianwei Lai, Aiguo Song, Ke Shi, Qinjie Ji, Ye Lu, and Huijun Li. Design and Evaluation of a Bidirectional Soft Glove for Hand Rehabilitation-Assistance Tasks. *IEEE Transactions on Medical Robotics and Bionics*, 5(3):730–740, 2023. ISSN 25763202. doi: 10.1109/TMRB.2023.3292414.

[54] Debin Hu, Jinhua Zhang, Yuhan Yang, Qiuyang Li, Dahai Li, and Jun Hong. A novel soft robotic glove with positive-negative pneumatic actuator for hand rehabilitation. In *IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, volume 2020-July, pages 1840–1847, 2020. ISBN 9781728167947. doi: 10.1109/AIM43001.2020.9158826.

[55] Brandon J. Caasenbrood, Femke E. Van Beek, Hoang Khanh Chu, and Irene A. Kuling. A Desktop-sized Platform for Real-time Control Applications of Pneumatic Soft Robots. *2022 IEEE 5th International Conference on Soft Robotics, RoboSoft 2022*, pages 217–223, 2022. doi: 10.1109/RoboSoft54090.2022.9762137.

[56] OpenCV: Detection of ArUco Markers. URL `https://docs.opencv.org/4.x/d5/dae/tutorial{_}aruco{_}detection.html`.

[57] Kyungyeon Lee, Daniel S Yang, Kriti Singh, and Jun Nishida. Hapticus: Exploring the Effects of Haptic Feedback and its Customization on Motor Skill Learning: Tactile, Haptic, and Somatosensory Approaches. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–20, Yokohama Japan, April 2025. ACM. ISBN 979-8-4007-1394-1. doi: 10.1145/3706598.3713821.

[58] Jun Nishida, Yudai Tanaka, Romain Nith, and Pedro Lopes. DigituSync: A Dual-User Passive Exoskeleton Glove That Adaptively Shares Hand Gestures. In *Proceedings of ACM Symposium on User Interface Software and Technology*, volume 1. Association for Computing Machinery, 2022. ISBN 9781450393201. doi: 10.1145/3526113.3545630.

[59] David Wagmann and Marie Muehlhaus. Move with Style ! Enhancing Avatar Embodiment in Virtual Reality through Proprioceptive Motion Feedback. In *Proceedings of ACM Symposium on User Interface Software and Technology*, 2025. ISBN 9798400720376. doi: 10.1145/3746059.3747649.

[60] Adam Tapal, Ela Oren, Reuven Dar, and Baruch Eitam. The sense of agency scale: A measure of consciously perceived control over one's mind, body, and the immediate environment. *Frontiers in Psychology*, 8(SEP):1–11, 2017. ISSN 16641078. doi: 10.3389/fpsyg.2017.01552.

[61] Tony Z. Zhao, Vikas Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *Robotics: Science and Systems*, 2023. doi: 10.48550/ARXIV.2304.13705.

[62] Zipeng Fu, Tony Z Zhao, and Chelsea Finn. Mobile ALOHA: Learning Bimanual Mobile Manipulation with Low-Cost Whole-Body Teleoperation, 2024.

[63] Mengda Xu, Han Zhang, Yifan Hou, Zhenjia Xu, Linxi Fan, and Shuran Song. DexUMI: Using human hand as the universal manipulation interface for dexterous manipulation. *arXiv.org*, 2025. doi: 10.48550/ARXIV.2505.21864.

[64] Hongjie Fang, Hao-Shu Fang, Yiming Wang, Jieji Ren, Jingjing Chen, Ruo Zhang, Weiming Wang, and Cewu Lu. AirExo: Low-cost exoskeletons for learning whole-arm manipulation in the wild. *IEEE International Conference on Robotics and Automation*, 2024. doi: 10.1109/ICRA57147.2024.10610799.

[65] Rui Zhong, Chuang Cheng, Junpeng Xu, Yantong Wei, Ce Guo, Daoxun Zhang, Wei Dai, and Huimin Lu. NuExo: A wearable exoskeleton covering all upper limb ROM for outdoor data

collection and teleoperation of humanoid robots. *arXiv.org*, 2025. doi: 10.48550/ARXIV.2503.10554.

[66] Rui Zhong, Yizhe Sun, Junjie Wen, Jinming Li, Chuang Cheng, Wei Dai, Zhiwen Zeng, Huimin Lu, Yichen Zhu, and Yi Xu. HumanoidExo: Scalable whole-body humanoid manipulation via wearable exoskeleton. 2025.

[67] Richard M. Murray, Zexiang Li, and S. Shankar Sastry. *A Mathematical Introduction to Robotic Manipulation*. CRC Press, 1 edition, December 2017. ISBN 978-1-315-13637-0. doi: 10.1201/9781315136370.

[68] Kevin M Lynch and Frank C Park. *Modern robotics*. Cambridge University Press, 2017.

[69] Joan Solà, Jeremie Deray, and Dinesh Atchuthan. A micro Lie theory for state estimation in robotics, December 2021.