

Document Version

Final published version

Licence

CC BY

Citation (APA)

Ratschat, A. L., Abbink, D. A., Ribbers, G. M., & Marchal-Crespo, L. (2026). Partial Arm Weight Support Does Not Hamper the Weight Perception of Haptically-Rendered Virtual Objects: A Psychophysics Study. *IEEE Transactions on Haptics*, 19(3), 687-700. <https://doi.org/10.1109/TOH.2026.3680656>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Partial Arm Weight Support Does Not Hamper the Weight Perception of Haptically-Rendered Virtual Objects: A Psychophysics Study

Alexandre L. Ratschat¹, Graduate Student Member, IEEE, David A. Abbink², Senior Member, IEEE, Gerard M. Ribbers³, and Laura Marchal-Crespo¹, Member, IEEE

Abstract—Arm weight support (AWS) from upper-limb exoskeletons aids training by reducing fatigue and increasing the movement workspace, especially in people with acquired brain injury (ABI). Simultaneously, task-specific training could be enabled via haptic rendering, simulating the interaction forces with tangible virtual objects. However, how AWS affects somatosensory perception, such as the perception of the weight of haptically rendered virtual objects, is unclear. We therefore conducted a psychophysics experiment with 40 participants investigating the effect of partial AWS on virtual object weight perception during dynamic lifting. Participants performed the lifting task under two conditions: with and without AWS, where AWS compensated for 50% of their arm weight. An upper-limb exoskeleton equipped with a haptic hand module provided both AWS and haptic rendering of virtual objects. The virtual objects and a representation of the participant's arm and hand were visualized using immersive virtual reality. Weight perception was assessed using just-noticeable differences (JNDs). To evaluate the experimental setup, participants' estimated elbow torques were analyzed. Additionally, subjective measures of task load, motivation, and user experience with the exoskeleton were collected to explore potential secondary factors influencing weight perception. We found that partial AWS effectively reduced elbow load during dynamic lifting compared to no-AWS. The JND was not significantly different between conditions, suggesting that AWS does not hamper weight discrimination. However, participants with a larger self-reported sense of control over their movement and those who reported that the robot followed their movements more closely were associated with better weight perception. Our findings suggest that including haptic rendering into exoskeletons that provide partial AWS could be a viable solution to provide task-specific training.

Index Terms—Force rendering, haptic rendering, robotics, perception and psychophysics.

I. BACKGROUND

FOR the past decades, upper-limb exoskeletons have been employed to support users during motor task execution, such as during the re-training of functional movements after an acquired brain injury. Upper-limb exoskeletons can assist in performing functional arm movements, e.g., those needed to perform activities of daily living (ADL), by providing adjustable arm weight support (AWS) to partially or fully counter the gravitational forces people experience from the weight of their arms [1], [2]. AWS has been shown to increase arm movement workspace and decrease abnormal shoulder-elbow coupling in people after stroke [3]. In addition, AWS limits muscle fatigue [2], [4], promoting high-intensity training [5], crucial to enhance learning [6].

Learners usually train with AWS while interacting with virtual objects, and the training environment is then visualized within a virtual reality (VR) setting, creating opportunities to train simulated ADL while playing motivating serious games [7]. However, this interaction usually lacks somatosensory information, which originates from the interaction with physical objects [8]. This somatosensory information is crucial for fine motor control [9] and successfully manipulating objects, such as stable object grasping, which is influenced by the object's physical characteristics [10]. Indeed, the absence or impairment of somatosensation hampers the learning of new tasks [11] and hinders stroke recovery [12]. Thus, integrating somatosensory information in robot-aided training may be vital to maximize (re)learning [8], in line with the specificity of practice hypothesis [13].

A versatile method for delivering somatosensory information during robotic training involves simulating the interaction forces from tangible virtual objects through haptic rendering [8], [14]. Several exoskeletons with haptic rendering capabilities have been developed. For instance, Galofaro et al. developed a bi-manual exoskeleton setup that simulates object pickup using the god-object algorithm [15]. Özen et al. utilized the ARMin exoskeleton to render the dynamics of a virtual pendulum via its equation of motion [14]. While both approaches successfully delivered the intended haptic rendering, the sensation was limited to the arm's proximal joints. Building on Galofaro et al.'s

Received 10 November 2025; revised 29 January 2026 and 24 March 2026; accepted 30 March 2026. Date of publication 3 April 2026; date of current version 30 September 2026. This work was supported by the Dutch Research Council NWO, VIDI, under Grant 18934. This article was recommended for publication by Associate Editor S. Kuroki and Editor-in-Chief K. Drewing upon evaluation of the reviewers' comments. (Corresponding author: Alexandre L. Ratschat.)

Alexandre L. Ratschat and Laura Marchal-Crespo are with the Department of Cognitive Robotics, Delft University of Technology, 2628 CD Delft, The Netherlands, and also with the Department of Rehabilitation Medicine, Erasmus MC, 3015 GD Rotterdam, The Netherlands (e-mail: a.l.ratschat@tudelft.nl).

David A. Abbink is with the Department of Cognitive Robotics, Delft University of Technology, 2628 CD Delft, The Netherlands, and also with the Department of Sustainable Design Engineering, Delft University of Technology, 2628 CD Delft, The Netherlands.

Gerard M. Ribbers is with the Department of Rehabilitation Medicine, Erasmus MC, 3015 GD Rotterdam, The Netherlands.

This article has supplementary downloadable material available at <https://doi.org/10.1109/TOH.2026.3680656>, provided by the authors.

Digital Object Identifier 10.1109/TOH.2026.3680656

setup, Buongiorno et al. integrated wrist and hand exoskeletons to render the interaction of a hand-held virtual stick with virtual walls [16]. Yet, they did not simulate the grasping of the stick itself. Rätz et al. designed an end-effector hand module mounted on a grounded end-effector delta device that leveraged the *Bullet* physics simulation engine to provide unprecedented levels of kinesthetic haptic rendering during hand-object interaction [17]. Transferring this method to upper-limb exoskeletons would extend the rich and versatile kinesthetic haptic rendering from the hand to the entire arm while leveraging the exoskeleton's ability to provide targeted assistance at each individual joint.

However, robotic assistance, such as AWS, may alter the perception of haptic stimuli, as assistive and haptic forces are usually provided by the same actuators [18], potentially limiting the effectiveness of haptic rendering in robot-aided training. In a recent study, Özen et al. showed that haptic rendering enhances motor learning and skill transfer during a dynamic sensorimotor task. However, adding AWS on top of the haptic rendering hampers learning [14]. The authors attributed the limited learning to the participants' reliance on assistance and the formation of internal models that failed to account for the dynamics of the uncompensated arm, which impeded skill transfer to unassisted tasks. Yet, another explanation is that the AWS corrupted the perception of the task's dynamics, in line with the Guidance Hypothesis [19]. Indeed, the effect of AWS on the perception of virtual objects' physical characteristics remains mainly unexplored. Since AWS provides support against gravitational forces, it can potentially alter the perception of a virtual object's weight. This is quite a limitation as object weight mediates the initial phases of object interaction and the forces applied for grasping and lifting objects [20].

Weight perception has been extensively studied for both real [21] and virtual objects [22]. Of relevance is Weber's Law, which states that the just-noticeable difference (JND)—i.e., the smallest detectable difference between two stimuli—is proportional to the intensity of a standard stimulus, which serves as the baseline or reference against which changes are compared [23]. In the context of weight perception, the standard stimulus might represent a known weight, while the JND reflects the smallest discernible weight difference with respect to this weight. The proportionality between the JND and the standard stimulus is quantified by the Weber fraction, which is considered to be almost constant, except for very small standard stimuli [24]. Weber's law has been employed to predict the effect of artificially increasing arm weight on weight perception. For example, Gregory and Ross investigated how artificially modifying the arm weight via a 500 g cuff attached to the forearm affects weight discrimination thresholds during the dynamic lifting of objects [25]. They found that adding the weighted cuff hampered weight perception, increasing the JND by a factor of 1.2. However, this increase was less than the predicted factor of 6 by Weber's law. Yet, the effect of alleviating arm weight on weight perception remains unclear.

Extending beyond the lifting of real objects, researchers have also explored weight perception for haptically rendered virtual objects. The majority of these previous studies focused on comparing the perception of weight when rendered via kinesthetic feedback—i.e., feedback relating to the movement,

position, and forces acting on the body and its parts—versus tactile cues through, e.g., skin stretch displays. For example, van Beek et al. [26] found Weber fractions for weight perception of approximately 30% to 40% in the kinesthetic condition, 30% to 35% in the tactile condition, and 20% to 30% in the combined condition. Similarly, Giachritsis et al. found Weber fractions of 16.1% for the weight perception of virtual objects simulated by a haptic interface for the index finger and thumb [27].

While these previous studies provide valuable insights into weight perception, they do not address how AWS—i.e., reducing the physical load of the arm—might influence the weight perception of haptically rendered virtual objects. Furthermore, they do not account for the potential effect of secondary factors on weight perception. For instance, muscular weakness, whether due to fatigue or neurological disorders, can increase the perceived heaviness of an object [21]. Additionally, increased cognitive load can lead to attentional lapses—i.e., instances where participants guess their responses rather than relying on the stimuli presented—negatively affecting perception measurements [28]. Another factor that may influence weight perception is the participants' feeling of controlling their own actions, i.e., their perceived agency, as higher agency is linked to improved detection accuracy [29]. Therefore, there is a need to investigate how AWS influences the perception of haptically rendered weight while addressing the role of secondary factors that may shape psychophysical outcomes.

In this study, we address this gap by investigating the effect of AWS on weight perception during dynamic lifting of haptically rendered virtual objects. We used an upper-limb exoskeleton to compensate for 50% of participants' arm weight. We exploited the *Bullet* physics simulation engine to haptically render the interaction forces between participants' hand and arms with the experimental environment, consisting of the lifted virtual object and a virtual table. Participants visualized the experimental environment, including a virtual representation of the user's arm and hand, from a first-person perspective using a commercial Head-Mounted Display (HMD). Using a within-subject counterbalanced experimental design, 40 unimpaired participants completed a psychophysics evaluation under two conditions: with AWS and without. The participants' weight perception under each condition was measured through the JND. Here, the JND reflects perceptual sensitivity in a constrained sensorimotor task rather than an isolated sensory threshold. To account for the influence of secondary factors on weight perception and potential interactions with AWS, we standardized the lifting task and arm posture, and used self-reported questionnaires to assess task load, intrinsic motivation, and user interaction experience with the exoskeleton.

We hypothesized that AWS improves weight perception compared to no AWS. This hypothesis is motivated by Weber's law and the assumption that with the same standard stimulus and psychophysical procedure, participants generate joint torques of smaller magnitude when lifting their arms with AWS than without AWS. While Weber's law is traditionally formulated in terms of changes in the task-relevant stimulus ($dS = K \cdot S$, where K is a constant [23]), weight perception during dynamic lifting depends on multiple concurrent cues. We therefore use the overall physical load (as reflected in the required joint torques)




Fig. 1. (a) Experimental setup with a mock participant in the exoskeleton, wearing the head-mounted display (HMD) and holding a gaming controller for user input. The VR trackers on top of the exoskeleton were used to place the exoskeleton in the virtual environment. Three force/torque (F/T) sensors were located under the cuffs and hand module. (b) Kinematic structure of the exoskeleton, including the seven active DOFs ($\theta_{e,1-7}$), and three passive DOFs with their respective angles and lengths ($\theta_{e,sh}$, l_{ua} , l_{la}). Note that in the schematic $\theta_{e,sh} = 0$. The kinematic structure of the hand module is simplified from [32]. (c) Schematic drawing of the *Average AWS* method from [2]. From the estimated AWS forces at the CoM (f_{CoM}), forces at the cuffs (f_{cuff}) resulting in equivalent torques at the segment origin (τ_{ua} , τ_{la}) are calculated. (d) Physics simulation environment in *Bullet* with the simulated arm, cylinder, and table. The angles of the active DOFs and the segment lengths of the upper and lower arm are indicated by $\theta_{s,1-7}$, l_{ua} , and l_{la} , respectively. (e) Side view of the virtual environment in *Unity* with the arm avatar, cylinder, table, screen, lift height indicator, and cylinder and elbow markers. (f) Zoomed-in point of view of the participants during the lifting task, with the cylinder top above the lift height indicator. The screen has three areas, with Area 1 showing the stage of the experiment (Lift 1, Lift 2, or Decision), Area 2 showing feedback after a decision was made (Weight 1/2 selected: Correct/Wrong), and Area 3 showing the total number of correct decisions on the left and the total number of comparisons on the right during each condition.

as a proxy for the effective stimulus context. We expect that reducing this load through AWS (i.e., a smaller effective S) may lower the JND (dS). We had no hypothesis on how secondary factors might affect weight perception or its interaction with AWS.

II. METHODS

A. Participants

A total of 42 unimpaired adults from the student and staff population of TU Delft, the Netherlands, were recruited for the experiment. Eligibility criteria included being right-handed, having no impairments to the right upper limb, body weight under 120 kg, no pacemaker, and arm length requirements of 26.5 cm to 31 cm for the lower arm and 23.5 cm to 40.5 cm for the upper arm. These inclusion criteria were defined based on safety measures and anthropomorphic measurements required to comfortably use the upper-limb exoskeleton employed in the study in its right arm configuration (see Section II-B). We restricted participation to right-handed individuals, as verified by the 10-item Edinburgh Handedness Inventory (EHI) [30], to ensure consistency in dominant-hand use. All participants provided written informed consent before participation. The study was approved by the TU Delft Human Research Ethics

Committee (approval numbers 4521 and 4633) and adhered to the ethical standards outlined in the Declaration of Helsinki.

B. Upper-Limb Exoskeleton

We utilized the ARMin upper-limb exoskeleton [31] augmented with the PRIDE hand module [32] as a haptic and assistive interface (Fig. 1(a)). The system includes seven active DOFs (Fig. 1(b)): shoulder abduction/adduction ($\theta_{e,1}$), elevation ($\theta_{e,2}$), internal/external rotation ($\theta_{e,3}$), elbow flexion ($\theta_{e,4}$), forearm supination/pronation ($\theta_{e,5}$), wrist flexion/extension ($\theta_{e,6}$), and collective finger flexion/extension, from the index to the little finger, ($\theta_{e,7}$). Additionally, there are three passive DOFs for user-specific adjustments, namely shoulder angle ($\theta_{e,sh}$), upper-arm length (l_{ua}), and lower-arm length (l_{la}). In this study, we set the wrist flexion/extension and the shoulder angle to 0° , yielding overall six active and two passive DOFs. For each participant, the upper and lower arm lengths were adjusted and remained constant for the duration of the experiment. Further, the hand module has four handles which can be exchanged to account for different hand sizes (small, small-medium, medium-large, large) [32].

All active DOFs ($\theta_{e,1-7}$) are equipped with optical incremental encoders (Maxon Encoder MR 1000 impulses per rotation,

Maxon Group, Switzerland) to accurately measure the exoskeleton joint angles. Additionally, all active DOFs, excluding the finger flexion/extension, are equipped with redundant analog sensors, such as string potentiometers, hollow shaft potentiometers, and hall sensors. The passive DOF angles and positions are measured with string potentiometers and hollow shaft potentiometers. The exoskeleton includes three six-DOF force/torque sensors (Mini45, ATI Industrial Automation, USA) positioned beneath the upper and lower arm cuffs and the hand module to measure the interaction forces and torques between the user and exoskeleton.

Velocity-based disturbance observers, using the interaction forces and torques as input, were employed at all joints at the shoulder, elbow, and forearm (θ_{e1-5}) to allow the exoskeleton to transparently follow the user's self-generated movement [31], [33]. The wrist flexion/extension was fixed in a neutral position ($\theta_{e,6} = 0$) using a stiff position controller (stiffness 100 Nm/rad, damping 20 Nms/rad). The hand module did not require compensation, as it is inherently transparent by design [32]. To compute the final output torques for the exoskeleton joints, the haptic rendering torques from the physics simulation (see Section II-D) were combined with the control torques generated by the disturbance observers. For the finger flexion/extension joint of the hand module, the haptic rendering torques were applied directly. The exoskeleton control, including disturbance observers, position controller, hand module control, and arm weight support (see Section II-C), was executed in Simulink Realtime R2017b (MathWorks, USA) at 3 kHz.

C. Arm Weight Support

We provided personalized AWS to the participants via the *Average AWS* method described in [2]. This method models the arm as two rigid segments: the upper arm (*ua*) and the lower arm with the hand (*la*). Each segment's weight and center of mass (CoM) are estimated using anthropometric data [34], the specific participant body weight measured before the experiment, and robot segment lengths measured by the exoskeleton through the sensors located in the passive DOFs.

For each segment, the vertical support force that compensates for the gravitational force at the human arm CoM is calculated as $f_{CoM} = [0, 0, gm]$ in the world coordinate frame (see Fig. 1(c)). Here, g is the Earth's gravitational acceleration and m is the estimated mass of each respective human arm segment. As the user interacts with the exoskeleton through the upper and lower arm cuffs, the CoM support forces are transformed into forces at the cuffs following equation:

$$f_{cuff} = f_{CoM} l_{CoM} / l_{cuff}, \quad (1)$$

where l_{CoM} and l_{cuff} are the distances from the segment start to the CoM and cuff, respectively (see Fig. 1(c)). This transformation results in equivalent torques $\tau_{aws,ua}$ and $\tau_{aws,la}$ at the intersection of the rotation axes (q_e) of the shoulder ($q_{e,1-3}$) and elbow ($q_{e,4-5}$), respectively. The AWS torques applied to the exoskeleton joints are then calculated as:

$$\tau_{aws} = J_{cuff,ua} f_{cuff,ua} + J_{cuff,la} f_{cuff,la}, \quad (2)$$

where $J_{cuff,ua}$ and $J_{cuff,la}$ are the Jacobians at the upper and lower arm cuffs, respectively. In this study, τ_{aws} was scaled by a factor of 0.5, providing 50% AWS. This scaling helped prevent support forces from exceeding the gravitational forces acting on the user's arm, which could result from approximations in the arm segment weights and CoM locations.

D. Haptic Weight Rendering

1) *Physics Simulation*: We developed a simulation of a virtual arm (*simulated arm*) and virtual objects (*simulated objects*) to calculate the interaction forces between them and render those forces using the exoskeleton. We modeled the interaction of the simulated arm with simulated objects using the *Bullet Real-Time Physics Simulation* (Version 3.2.5, bulletphysics.org, USA) due to its versatility in modifying the simulation environment and high performance [17], [35].

To represent the user's arm in the physics simulation following the kinematic chain of the ARMin exoskeleton, we connected capsule objects for the upper arm, lower arm, and hand (Fig. 1(d)). The simulated upper and lower arm had six active DOFs, namely $\theta_{s,1}-\theta_{s,6}$, which correspond to those from the exoskeleton $\theta_{e,1}-\theta_{e,6}$. The hand was modeled with capsule colliders for each finger and thumb segment, representing the metacarpals and phalanges. The thumb was fixed at 65° palmar abduction to facilitate power grasping (see Fig. 1(d)). Each finger included three joints for the metacarpophalangeal and the two interphalangeal joints, resulting in a total of 12 active DOFs for the simulated hand. Here, the angles of the phalangeal joints were set to 1/3 of the total finger flexion/extension angle, $\theta_{s,7}$, to approximate interjoint coupling. The total finger flexion/extension angle $\theta_{s,7}$ was defined as the flexion angle of the most distal finger joint with respect to the index finger metacarpal, where $\theta_{s,7} = 0$ indicates fully extended fingers (Fig. 1(d)). The passive DOFs, l_{ua} and l_{la} , were initialized based on the user's segment lengths as measured by the exoskeleton.

The physics simulation included a virtual table and an object of cylindrical shape. The table provided a stable surface for participants to rest their elbows during the lifting task to help reduce shoulder fatigue, promote consistent movement patterns between participants, and protect their legs from accidental exoskeleton collisions. The table height was adjusted to the participant's upper-arm length, ensuring a shoulder elevation of $\theta_{s,2} = -35^\circ$ when the elbow was at rest on the virtual table. The cylinder (height: 15 cm, diameter: 7 cm) was used as the object to be lifted. Its mass could be varied throughout the experiment to assess weight perception. We disabled gravity in *Bullet* to prevent the simulated arm from experiencing gravitational forces. The weight of the cylinder was, therefore, rendered through a gravitational force, $f_{g,cylinder} = [0, 0, -m_c g]$, applied to the cylinder's CoM at each simulation step, where m_c is the cylinder's mass.

Bullet's unitless dynamics parameters of the simulated upper and lower arm, hand, cylinder, and table were set empirically for a naturalistic experience. In particular, the lateral friction, spinning friction, rolling friction, and linear damping were set to 2×10^{-1} , 1×10^{-3} , 2×10^{-5} , 0 for the upper

and lower arm, 8×10^{-1} , 5×10^{-3} , 1×10^{-4} , 1×10^{-1} for the hand, 4 , 5×10^{-3} , 1×10^{-4} , 1×10^{-1} for the cylinder, and 1.3 , 5×10^{-3} , 1×10^{-4} , 1×10^{-1} for the table. Additionally, we set the contact stiffness and damping of the hand to 8×10^3 and 8×10^2 , respectively.

2) *Bilateral Exoskeleton-Physics Simulation Controller*: A bilateral control scheme was employed between the exoskeleton and the physics simulation with two objectives: (i) The simulated arm should follow the movements of the real arm in the exoskeleton (*Follower Controller*), and (ii) the interaction forces and torques between the simulated arm and objects should be projected back to the real arm through the exoskeleton (*Leader Controller*).

For the *Follower Controller*, we chose a proportional-derivative (PD) controller based on the joint positions and velocities of the exoskeleton. The 7×1 vector of generalized coordinates of the exoskeleton is $q_e = [\theta_{e,1}, \theta_{e,2}, \theta_{e,3}, \theta_{e,4}, \theta_{e,5}, \theta_{e,6}, \theta_{e,7}]^T$. The finger flexion/extension entries in q_e , i.e., $\theta_{e,7}$, were replaced with the phalangeal joint values, calculated as one-third of the original flexion/extension angles, yielding the extended 18×1 vector $q_{e,ext} = [\theta_{e,1}, \theta_{e,2}, \dots, \theta_{e,6}, \theta_{e,index,1}, \theta_{e,index,2}, \theta_{e,index,3}, \theta_{e,middle,1}, \dots, \theta_{e,little,3}]^T$. The exoskeleton joint velocities, $\dot{q}_e$ and $\dot{q}_{e,ext}$, were computed via Euler-backwards differentiation from the exoskeleton joint angles. The obtained joint velocities were first-order low-pass filtered to reduce noise, with cutoff frequencies of 10 Hz for the shoulder abduction/adduction and 30 Hz for all other joints. The 18×1 vectors containing the simulated arm's joint positions q_s and velocities $\dot{q}_s$ were obtained directly from the physics simulation, following the same structure as the extended generalized coordinates of the exoskeleton.

The *Follower Controller* yields a vector τ_f of virtual torques applied at each joint of the simulated arm, enforcing it to follow the movements of the real arm in the exoskeleton:

$$\tau_f = \tilde{M}(K_{p,f}(q_{e,ext} - q_s) + K_{d,f}(\dot{q}_{e,ext} - \dot{q}_s)) + h_s. \quad (3)$$

Here, $\tilde{M}$ represents the simulated arm's mass matrix with only the diagonal elements obtained from the physics simulation. Using a diagonal mass matrix ensured that forces and torques from off-diagonal elements did not undesirably affect other DOFs during hand-object interactions. Examples of $\tilde{M}$ for different arm lengths are provided in Section I of the Supplementary Material. The diagonal proportional gain matrix $K_{p,f}$ contains the gains (6×10^3 , 6×10^3 , 6×10^3 , 1×10^4 , 1×10^4 , 1×10^4) s^{-2} for all arm joints ($\theta_{s,1} - \theta_{s,6}$) and $9 \times 10^3 s^{-2}$ for each finger joint. The diagonal damping gain matrix $K_{d,f}$ contains the critical damping gains, where the individual elements were calculated as:

$$k_{d,f,i} = 2\sqrt{k_{p,f,i}} - k_{p,f,i}T, \quad (4)$$

for the proportional gain $k_{p,f,i}$ of each joint i ($i = 1, 2, \dots, 18$) and the time increment $T = 1$ ms. Finally, h_s in (3) represents velocity-dependent forces and torques within the simulation, including Coriolis and centrifugal forces, as well as any modeled damping effects. These forces and torques are computed by evaluating the inverse dynamics of the simulated system under

the assumption of zero accelerations. A detailed explanation of how (3) and (4) are derived can be found in [35].

The *Leader Controller* also followed a PD controller scheme defined as:

$$\tau_l = K_{p,l}(q_{e,ext} - q_s) - K_{d,l}(\dot{q}_{e,ext} - \dot{q}_s). \quad (5)$$

Here, τ_l represents the vector of exoskeleton joint torques with extended DOFs for the finger joints, corresponding to $q_{e,ext}$. These exoskeleton joint torques result in the user feeling the interactions of the virtual arm with the virtual object. For the arm excluding the hand ($\theta_{s,1} - \theta_{s,6}$), the diagonal proportional gain matrix $K_{p,l}$ contains the gains (1×10^3 , 2×10^3 , 1×10^3 , 4×10^4 , 2×10^4 , 0) Nm/rad. To reduce the dimensionality at the hand back to one DOF for finger flexion/extension, we consider the normal force at the fingertips instead of the torques at the individual finger joints. This was achieved by calculating the proportional gains of each metacarpophalangeal joint with:

$$k_{p,l,i} = k_{p,l,finger} J_i J_i, \quad (6)$$

with $k_{p,l,finger} = 5 \times 10^2$ N/m, and J_i the Jacobian for each finger metacarpophalangeal joint $i = 7, 10, 13, 16$ with respect to the normal force at the fingertip.

The diagonal damping gain matrix $K_{d,l}$ in (5) contains the gains (2.7×10^1 , 5.15×10^1 , 2.8 , 8.3 , 4.5×10^{-1} , 0) Nms/rad for all arm joints ($\theta_{s,1} - \theta_{s,6}$). The damping gains for the metacarpophalangeal joints were calculated with $k_{d,l,finger} = 11.0$ Ns/m and

$$k_{d,l,i} = k_{d,l,finger} J_i J_i, \quad (7)$$

for $i = 7, 10, 13, 16$. The proportional and damping gains for all interphalangeal joints were set to zero. All controller gains were empirically tuned across the full range of experimental motions and loads to ensure stable and safe operation.

Finally, the total normal fingertip force was calculated as the sum of the individual finger forces $f_{l,mcp,i}$, which in turn were calculated from transforming the metacarpophalangeal torques $\tau_{l,mcp,i}$ back to forces using the Jacobians:

$$f_{l,mcp,i} = \tau_{l,mcp,i} / J_i. \quad (8)$$

For a detailed description of how the finger force was transformed into motor torques, please refer to [32].

The physics simulation setup, the bilateral controller, and the software with the experimental protocol were programmed in C++ (Version C++20) on a desktop computer (OS: Ubuntu 22.04.4 LTS, Kernel: 5.15.0-1067-realtime, RAM: 32GB, CPU: AMD Ryzen threadripper 2950X 16-core 3.5 GHz, GPU: NVIDIA GeForce RTX 2080 Ti 11GB). The experimental environment with the physics simulation ran at 1 kHz and communicated bilaterally with the exoskeleton control software via the User Datagram Protocol (UDP) at 1 kHz.

E. Virtual Environment

We developed a Virtual Environment (VE) to provide participants with a first-person perspective visualization of their arms and the experimental setup, i.e., the object and table, using a

head-mounted display (Varjo XR-3, Varjo, Finland). We developed the VE in the Unity game engine (Version 2022.2.15f1, Unity Technologies, USA) and ran it with a desktop computer (OS: Windows 10 Enterprise, RAM: 32GB, CPU: AMD Ryzen threadripper 2920X 12-core 3.5 GHz, GPU: NVIDIA GeForce RTX 4070 Ti 12GB). Two HTC VIVE 3.0 trackers (HTC Corporation, Taiwan) mounted on top of the exoskeleton (see Fig. 1(a)) and two SteamVR 2.0 base stations (HTC Corporation, Taiwan) were used to align the coordinate frames of the VE and the physics simulation, as well as to track the position of the HMD within the VE [36].

The virtual arm (see Fig. 1(e)–(f)) was modeled using the Unity asset *Adventure Character* (Version 1.0, Maksim Bugri-mov). The avatar was animated using the data from the physics simulation, including simulation hand and arm joint positions ($\theta_{s,1-7}$) and segment lengths (l_{ua} , l_{la}) sent via UDP at 200 Hz. The VE included a virtual cylinder and a table, mirroring the dimensions and positions of the simulated counterparts in the physics engine. The virtual cylinder followed the movements of the simulated one, sent through UDP, while the virtual table was fixed in the same position as the simulated one.

Two markers, visualized as grey disks, were displayed on the table to indicate the desired elbow and cylinder positions to participants. The position of the elbow marker was defined by the desired shoulder abduction/adduction ($\theta_1 = -50^\circ$) and elevation ($\theta_2 = -35^\circ$) and the upper-arm length l_{ua} , and therefore, was specific for each participant. The cylinder marker was placed on the table along the user's sagittal plane (i.e., along x) at a distance of $l_{la} + l_h$ from the elbow marker (Fig. 1(e)), corresponding roughly to the length from the elbow to the grasped object, where $l_h = 0.1$ m is the distance from the wrist to the cylinder's center line.

To assist participants during the task, the VE included two additional elements: a lift-height indicator and a virtual screen. The lift-height indicator was visualized as a yellow semi-transparent plane positioned above the cylinder marker at a height of 19 cm with respect to the table, i.e., 4 cm above the top of the cylinder. The virtual screen, located in front of the participants behind the cylinder marker, displayed messages related to the experimental protocol (see Section II-F). The virtual environment and avatar were kept identical between conditions within each participant.

F. Experimental Protocol

The experiment consisted of a single session of approximately two hours and was conducted by a research engineer. The experiment followed a within-subject design, where participants underwent two conditions—*Without AWS* and *With AWS*. The order of the conditions was counterbalanced, i.e., half the participants started with either condition. The participants were informed about the condition order before beginning the experiment to provide a rationale for the experienced AWS forces. The experimental protocol is visualized in Fig. 2.

Upon arrival, the participants' information regarding their age and sex was collected. A commercial digital scale was used to measure their body weight. In addition, they filled in the 10-item version of the EHI [30] to corroborate their




Fig. 2. Experimental protocol with pre-experiment preparation, two conditions—Without and With AWS—in counterbalanced order, and the post-experiment interview. The entire experiment took approximately two hours. Abbreviations: Edinburgh Handedness Inventory (EHI), Raw Task Load Index (RTLX), Intrinsic Motivation Inventory (IMI), Human-Robot Interaction Questionnaire (HRIQ).

right-handedness. Their hand size from the wrist to the middle finger tip was measured, and the corresponding handle size for the hand module was installed. Following, the participants received an introduction to the experimental task and setup. Participants were then set up in the exoskeleton by carefully aligning the exoskeleton links to the participants' arms.

Following, the participants donned the HMD and the first condition was tested. Each condition evaluation consisted of three main parts. First, we measured the interaction forces and torques between the participants and exoskeleton during static lifting for different cylinder masses (see Section II-F1). This measurement allowed us to evaluate the elbow load participants experienced from the gravitational forces acting on the cylinder. Second, we measured the participants' weight perception during dynamic lifting (see Section II-F2). In the third part, we asked participants to complete three questionnaires (see Section II-G2).

After completing all the measurements and questionnaires for both conditions, participants could report any adverse events, i.e., any undesired experience, like dizziness or pain, from the interaction with the exoskeleton and the VR experience. All demographic data, questionnaires, and adverse events were collected using the Qualtrics XM survey tool (Qualtrics, USA). Finally, a short semi-structured interview was conducted to gather their overall experience during the experiment, the interaction with the exoskeleton and the virtual objects, the VR experience, and whether it was easier to distinguish weight in either of the conditions.

1) *Static Lift Measurement*: We performed a first static lift measurement to validate our experimental setup, i.e., evaluate the elbow load participants experience for different cylinder masses, with and without AWS. For the static lift measurement, participants received instructions to grasp the virtual cylinder in a power grasp and hold it steadily for 48 s, with the top of the cylinder slightly above the lift height indicator. They were instructed to rest their elbow at all times on the elbow marker on the virtual table. These instructions were read to the participants for each condition. After the participants were in the correct position and indicated that they were ready, the measurement started. During the measurement, the cylinder mass was increased from 0 kg to 1.75 kg and then decreased back to 0 kg in steps of 0.25 kg for a duration of 3 s each. Participants could see a countdown on the virtual screen in front of them with the remaining total measurement time. The exoskeleton and simulation active and passive joint positions, the interaction torques and forces generated by the simulation, the human-robot interaction forces and torques measured by the force/torque sensors, the AWS torques, and the virtual cylinder position and rotation were recorded at 100 Hz.

2) *Dynamic Lift Weight Perception Measurement*: The weight perception measurement for the dynamic lifting of virtual objects followed a two-alternative forced-choice (2-AFC) psychophysical procedure [37]. Participants were asked to lift the virtual cylinder twice and decide whether the cylinder was heavier the first or the second time. Again, participants were instructed to rest their elbows on the elbow marker while lifting. For each lift, they had to grasp the cylinder in a power grasp and lift it until the cylinder's top slightly penetrated the lift height indicator. After remaining above the indicator for 0.5 s, the indicator disappeared, prompting the participant to gently place the cylinder on the table. Between lifts, they were asked to open their hand to release the cylinder. We then changed the cylinder mass and the lift height indicator reappeared for the next lift. When the second lift was completed, participants were asked to indicate whether they felt the first or second cylinder was heavier using a gaming controller (Xbox 360 wired controller, Microsoft, USA). They always had to answer, even if unsure.

To practice the lifting procedure, we included a familiarization phase of 20 comparisons (i.e., 40 lifts), during which the participants could ask questions and received instructions from the experimenter about the experimental procedure. Participants were presented with a standard stimulus of 1 kg and a comparison stimulus of either 0.6 kg or 1.4 kg. Each comparison stimulus was presented ten times in a randomized order, and the order of presenting the standard and comparison stimuli was pseudo-randomized, so each was presented first equally often.

After the familiarization, participants were equipped with noise-canceling headphones playing white noise to reduce any auditory bias. They then commenced with the weight perception measurement. For each comparison, they were presented with a standard stimulus of 1 kg and a comparison stimulus. This standard stimulus was chosen based on a balance of factors, including the exoskeleton's maximum joint torques, its transparency, and the physical load on the participants. The comparison stimulus was based on a transformed and weighted up/down

method [37], [38]. The method followed a 3-down/1-up rule, i.e., after three correct decisions on the heavier stimulus, the comparison stimulus moved closer to the standard stimulus, and for one wrong decision, the comparison stimulus moved further away. The step size moving further away was set to 0.05 kg and the step size moving closer to $0.968 \cdot 0.05$ kg, yielding a targeted correct-response ratio of 79.28% [38]. The use of this specific step size ratio makes the method's correct-response ratio independent of the starting point and relative step size, as well as increasing its accuracy [38]. Further, the chosen step size is well below both the expected JND for a 1 kg standard (~ 0.09 kg to 0.13 kg reported for real-object lifting [21]). To allow the procedure to converge faster to the desired correct-response level, a 1-down/1-up rule was used until the first reversal [37], where a reversal is the point where the comparison stimulus switches from moving closer to moving further from the standard stimulus and vice versa. Two of these so-called staircases were intertwined, with one descending staircase starting at 1.4 kg following the described up/down method and an ascending staircase starting from 0.6 kg in a mirrored fashion (see Fig. 3(a)).

The initial comparison stimuli were selected based on pilot experiments, which indicated they were easily distinguishable from the standard stimulus. The comparison stimuli were bounded to be within [1.75, 1.0] kg for the descending staircase and [0.25, 1.0] kg for the ascending staircase. These limits ensured the required joint torques stayed within the exoskeleton's operational range and kept the rendered torques for lighter weights sufficiently above any unintended artifacts arising from the system's transparency. The boundaries followed the *carry-on* condition within the up/down method [39], meaning that whenever the rules called for off-limits values, the stimulus was presented at the boundary level. The order of presentation of the standard and comparison stimulus, as well as the order of the descending and ascending comparison stimuli, were pseudo-randomized. Each staircase was terminated after 12 reversals and the measurement was automatically terminated with both staircases. After every 25 comparisons, we enforced a break of 60 s to mitigate fatigue. During this time, participants were encouraged to rest and stretch their necks and received positive feedback to maintain motivation. During the dynamic lift measurement, we collected the same data from the exoskeleton and simulation at the same frequency as for the static lift measurement.

G. Outcome Metrics

All outcome metrics were processed in Python (v. 3.12).

1) *Just-Noticeable Difference for Weight Perception*: Our main outcome metric of weight perception was the JND for each condition. We calculated the JND by fitting a psychometric curve to the data from the up/down method and evaluating it at the method's correct response level of 79.28%, which represents the asymptotic percentage of correct responses achieved by the staircase procedure over an infinite number of trials [40]. We used a logistic function as the psychometric curve for the 2-AFC discrimination task:

$$\psi(s) = \lambda + \frac{1 - 2\lambda}{1 + e^{-\beta(s-s_0)}}, \quad (9)$$




Fig. 3. (a) Example of the weighted up/down transformed response rule with intertwined descending and ascending staircase of one participant during the Without AWS condition. (b) Corresponding logistic fit to the responses in (a). For each comparison, the value on the ordinate denotes whether the participant indicated that the comparison stimulus was heavier than the standard stimulus (1 if the participant decided the comparison stimulus was heavier, 0 otherwise). Abbreviation: Just-Noticeable Difference (JND).

where $\psi(s)$ denotes the probability that a participant decides that the comparison stimulus s is heavier than the standard stimulus. The parameter s_0 is the Point of Subjective Equality (PSE), representing the stimulus level at which participants have a 50% probability of perceiving the comparison stimulus as heavier than the standard stimulus. We fixed the lapse rate $\lambda = 0.025$ and bounded the scale $\beta \in [0.1, 100]$ and PSE $s_0 \in [0.25, 1.75]$. Fixing the lapse rate to a small non-zero value avoids biases found when fixing the lapse rate to zero [41]. To fit the logistic function to our data, we used the non-linear least squares *curve_fit* function of the SciPy optimization and root finding package with initial guesses $\beta_0 = 3.0$ and $s_0 = 1.0$. Finally, the JND was obtained by calculating:

$$JND = \frac{\psi^{-1}(0.7928) - \psi^{-1}(1 - 0.7928)}{2}, \quad (10)$$

where ψ^{-1} is the inverse of the logistic function in (9) (an example is provided in Fig. 3(b)).

2) *Questionnaires and Participant Statements*: We asked participants to fill in three questionnaires after completing the

dynamic lift measurements for each condition to account for secondary factors that may affect their weight perception and potential interaction effects with AWS.

The first questionnaire assessed the participants' perceived physical and mental workload via the Raw Task Load Index (RTLX) [42]. The RTLX was answered using a visual analog scale yielding values from 0–20, with 0 representing low and 20 high scores. The second questionnaire investigated the participants' motivation via the Intrinsic Motivation Inventory (IMI), where we used the Effort/Importance, Interest/Enjoyment, and Perceived Competence subscales [43]. The last questionnaire assessed how participants perceived the transparency of their interaction with the exoskeleton, which we call the Human-Robot Interaction Questionnaire (HRIQ) [33], with subscales for Comfort, Low Resistance, Compliance, Precision, Stabilization, and Agency. Comfort reflects how comfortable the participants' movement in the exoskeleton felt. Low Resistance assesses the extent to which participants perceived the exoskeleton as offering minimal resistance to their movements. Compliance measures how well the exoskeleton followed the participants' intended movements. Precision captures the accuracy with which participants felt that they could perform the tasks using the device. Stabilization evaluates the perception of the robot's ability to support and steady the participants' movements. Lastly, Agency refers to the participants' perceived control over their movements when using the exoskeleton. The IMI and HRIQ were answered using seven-item Likert scales, with 1 indicating low and 7 high scores. The three questionnaires were always presented in the same order; however, the order of the questions within each questionnaire was randomized. The complete survey can be found in the supplementary material.

The scores for each questionnaire were averaged over all subscales to yield a single outcome. For the IMI and HRIQ, there could be several questions per subscale, which were also averaged to obtain a score for the subscale. Some scores were reversed before averaging, according to the questionnaire instructions. Participant statements from the reported adverse events and semi-structured interviews were summarized, and recurring statements were paraphrased and used to complement our discussion.

3) *Experimental Setup Evaluation*: To test whether participants experienced joint torques of smaller magnitude with AWS than without AWS during the experiment, we compared the torques at the elbow during static and dynamic lift measurements between conditions. Since participants rest their elbows on the virtual table while lifting, we focused on the torques at the elbow, calculated as:

$$\tau_{elbow} = \tau_{la} + \tau_h + r_{elbow,la} \times f_{la} + r_{elbow,h} \times f_h, \quad (11)$$

where τ_{la} and f_{la} are the recorded interaction torques and forces at the lower arm, and τ_h and f_h at the hand. The position vectors between the elbow and the force/torque sensors at the lower arm and hand are denoted by $r_{elbow,la}$ and $r_{elbow,h}$, respectively. We focused the evaluation on the $\tau_{elbow,y}$ component of the elbow torque, as y is parallel to the lower arm's axis of rotation while lifting. Note that as interaction forces increase in z during lift onset, elbow torques decrease around y . Also, the reported

elbow torque is an estimated external elbow moment demand; we did not directly measure participants' internal biological joint torques.

We measured the interaction forces and torques during the 3 s static lifting of each virtual cylinder mass as described in Section II-F1. We discarded the first 0.5 s of each measurement period to account for dynamic fluctuations due to mass changes. The resulting elbow torques were averaged for each cylinder mass per participant and per condition.

For the dynamic lift measurement, we first identified all individual lifts as the time interval where the cylinder top was at least 0.5 s above the lift height indicator. This interval was then extended by 0.5 s before and after to include the movement onset and cylinder lowering. We chose the elbow torque minima in the y-axis as a measure for the largest load participants experience during dynamic lifts as elbow torques decrease during lift onset and the extrema are unaffected by varying lift durations. We also calculated the elbow torque range for each lift, i.e., the difference between the torque maximum and minimum within a lift interval, as a measure for the experienced dynamic load. Note that the advantage of calculating this range is that it is unaffected by the torque level before lift onset. An exemplary visualization of the cylinder position, velocity, and elbow torques during a lift with and without AWS can be found in the supplementary material.

H. Statistical Analysis

To evaluate our hypothesis and account for secondary factors, we defined a Linear Mixed Model (LMM) with the *Condition* (With and Without AWS) and questionnaire subscales as the independent variables and the log-transformed JND as the dependent variable. The JND was log-transformed to meet the residual normality assumption in the LMM. To reduce the model complexity, prevent overfitting, and ensure model stability, we compared reduced candidate models with fewer secondary factors as independent variables using the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC). The model with the smallest AIC and BIC was chosen, indicating a good balance between model fit and complexity. The chosen model contained the Physical and Temporal Demand subscales of the RTLX (PD and TD), the Effort/Importance subscale of the IMI (EI), and the Agency and Compliance subscales of the HRIQ (AG and CP). The participant ID was included as a random effect:

$$\text{LMM1. } \log(JND) \sim \text{Condition} \times (PD + TD + EI + AG + CP) + (1|ID).$$

In order to evaluate not only differences between conditions but also practical equivalence, we performed Two One-Sided Tests (TOST) equivalence testing on the LMM-estimated *Condition* contrast at the mean levels of the normalized covariates. This was done within prespecified margins of $\pm 10\%$ and $\pm 20\%$. While clinically meaningful change thresholds have been established for some motor outcomes and sensory measures, no comparable clinically accepted threshold is currently available for weight-discrimination JNDs in this context. Therefore, instead, the selected bounds were intended as stricter and more permissive statistical margins.

To evaluate our experimental setup, we tested the effect of the condition and virtual cylinder mass on the mean elbow torque during static lifting. Additionally, we tested the effect of the condition, lift number, and virtual cylinder mass on the elbow torques during dynamic lifting. Here, we included the lift number to monitor potential changes in elbow torques throughout the experiment. We included the data of all participant lifts in the analysis. Note that the number of lifts varied per participant, as the termination of the psychophysical procedure depended on the number of reversals (see Section II-F2). Also, the experienced weights could differ between conditions because the weight was determined by the adaptive psychophysical procedure, which adjusted the weights based on each participant's performance. We defined three LMMs, which included the participant ID as a random effect:

$$\text{LMM2. } \tau_{\text{elbow},y,\text{static},\text{mean}} \sim \text{Condition} \times \text{Mass} + (1|ID)$$

$$\text{LMM3. } \tau_{\text{elbow},y,\text{dynamic},\text{min}} \sim \text{Condition} \times (\text{LiftNumber} + \text{Mass}) + (1|ID)$$

$$\text{LMM4. } \tau_{\text{elbow},y,\text{dynamic},\text{range}} \sim \text{Condition} \times (\text{LiftNumber} + \text{Mass}) + (1|ID).$$

For all models, the lift number, cylinder mass, and questionnaire data were centered and normalized to improve model interpretability and reduce multicollinearity among predictors. We used the *lmer* function of the *lme4* package (Version 1.1-35.5) in R (Version 4.4.1). Model assumptions, i.e., normality and independence of the residuals and homoscedasticity, were assessed visually using QQ plots, residuals vs. fitted plots, and boxplots of residuals by participant ID. All *p*-values were adjusted using the Bonferroni correction for multiple comparisons. The significance level was set to $\alpha < 0.05$.

III. RESULTS

Out of the 42 recruited participants, 41 completed the study. One participant struggled with using the gaming controller, repeatedly pressing the opposite buttons from those intended, probably due to inexperience or trouble seeing the controller while wearing the HMD. This led to the participant terminating the study before completion. We did not include the dataset of another participant in the analysis due to the participant's hand frequently slipping out of the hand module, which was used to grasp the virtual objects. A total of 40 participants were included in the analysis, who were between 22 and 35 years old (median age 26 years) and weighed between 49.4 and 111.7 kg (median weight 73.8 kg), with 30 males and 10 females. The EHI confirmed right-handedness for 39 participants, with one participant categorized as ambidextrous (EHI Laterality Index median 100, interquartile range 15.98). The questionnaire results, along with their subscales, are presented in the supplementary material.

A. Just-Noticeable Difference for Weight Perception

We included 80 data points in the analysis of the JND, i.e., one data point per condition and participant (LMM1). Results from the LMM1, i.e., estimates (β), standard errors (SE), and Bonferroni-corrected *p*-values for the main and interaction effects are summarized in Table I.

TABLE I
RESULTS FROM THE LINEAR MIXED MODEL FOR THE log-Transformed
JND, INCLUDING THE ESTIMATES, STANDARD ERRORS, AND
BONFERRONI-CORRECTED P -VALUES. ALL INCLUDED QUESTIONNAIRE
SCORES WERE CENTERED AND NORMALIZED.

Term	Estimate	Std. Error	p -value
Intercept	-1.511	0.053	< 0.001
Condition Without AWS	0.081	0.058	1.000
Physical Demand (PD)	-0.089	0.054	1.000
Temporal Demand (TD)	0.124	0.052	0.234
Effort Importance (EI)	0.037	0.050	1.000
Agency (AG)	-0.185	0.044	0.002
Compliance (CP)	0.029	0.042	1.000
Condition Without AWS:PD	0.083	0.072	1.000
Condition Without AWS:TD	-0.133	0.060	0.393
Condition Without AWS:EI	-0.158	0.057	0.094
Condition Without AWS:AG	0.150	0.075	0.600
Condition Without AWS:CP	-0.201	0.064	0.037

We did not observe a significant main effect of arm weight support on the JND (With AWS: 0.253 ± 0.113 kg, Without AWS: 0.241 ± 0.084 kg; Table I). The estimated effect size was also small (Without/With AWS ratio = 1.085, 95% confidence interval (CI) [0.969, 1.214]). We found a main effect of Agency ($p < 0.01$), with increases in self-reported Agency associated with a 16.9% reduction in JND per unit increase in centered and normalized Agency (Fig. 4(a)).

We also found an interaction effect between Condition and Compliance on JND ($p < 0.05$). Specifically, in the Without AWS condition, higher Compliance was related to an 18.2% reduction in JND per unit increase in centered and normalized Compliance score (Fig. 4(b)). The equivalence testing indicated that the AWS–no-AWS contrast (90% CI: $[-16.4\%, 1.7\%]$) was equivalent within the $\pm 20\%$ margin ($p_{\text{low}} = 0.0096$, $p_{\text{high}} < 0.001$), whereas equivalence within the stricter $\pm 10\%$ margin was not established ($p_{\text{low}} = 0.341$, $p_{\text{high}} = 0.0021$).

B. Evaluation of the Experimental Setup

For the mean elbow torque during static lifting (LMM2), we included 640 data points in the analysis, i.e., 40 participants $\times$ 8 cylinder masses. As expected, overall, participants lifting with AWS experienced significantly higher mean elbow torques (-0.08 ± 1.17 Nm) compared to lifting without AWS (-2.01 ± 1.0 Nm; $\beta = 1.92$, $SE = 0.15$, $p < 0.001$; Fig. 4(c)). This means that, as expected, participants experienced, on average, smaller elbow torque magnitudes with AWS than without AWS. We also found a main effect of the cylinder mass on the joint torque. In particular, for each unit increase in the normalized and centered cylinder mass, participants experienced significantly decreased mean elbow torques by 0.96 Nm ($SE = 0.027$, $p < 0.001$). Importantly, we did not find a significant interaction effect between the condition and the cylinder mass ($p = 0.435$).

For the elbow torques during the dynamic lifting (LMM3 and LMM4), we compared 18044 lifts (40 participants $\times$ 2 conditions $\times$ total number of comparisons per condition per participant $\times$ 2 lifts per comparison). Overall, participants exhibited significantly higher elbow torque minima while lifting with AWS (-1.32 ± 1.04 Nm) compared to lifting without

AWS (-3.26 ± 0.95 Nm); $\beta = 1.98$, $SE = 0.0082$, $p < 0.001$; Fig. 4(d)). Thus, on average, participants experienced smaller magnitudes of elbow torque minima lifting with AWS compared to without AWS. Further, the elbow torque minima significantly decreased by 0.057 Nm for each unit increase in the normalized and centered lift number ($SE = 0.0060$, $p < 0.001$). Finally, the elbow torque minima decreased significantly for each unit increase in the normalized and centered cylinder mass by 0.29 Nm ($SE = 0.0059$, $p < 0.001$). We did not find statistically significant interactions between the condition and lift number ($p = 1.000$) and condition and cylinder mass ($p = 1.000$) on the elbow torque minima.

For the elbow torque range, representing the magnitude of torque change in each lift, we observed significantly reduced elbow torque ranges in participants when lifting with AWS (3.41 ± 1.30 Nm) compared to lifting without AWS (3.80 ± 1.47 Nm; $\beta = 0.38$, $SE = 0.012$, $p < 0.001$; Fig. 4(e)). The torque range significantly increased by 0.24 Nm for each unit increase in the normalized and centered lift number ($SE = 0.0091$, $p < 0.001$). Further, the torque range was significantly increased by 0.082 Nm for unit increases in the normalized and centered mass ($SE = 0.0090$, $p < 0.001$). Additionally, a significant interaction between the condition and lift number ($p < 0.001$) shows that the effect of lift number on the elbow torque range is less pronounced when AWS is used ($\beta = -0.097$, $SE = 0.013$). We found no statistically significant interaction between the condition and cylinder mass ($p = 1.000$) on the elbow torque range.

IV. DISCUSSION

A. Arm Weight Support Does Not Significantly Affect Weight Perception of Haptically Rendered Virtual Objects

To put our findings into context with current literature, we use the Weber fraction by dividing the JND by the standard stimulus. This allows comparison with other studies that used different standard stimuli. Overall, the Weber fractions we observed for dynamic weight perception with arm weight support are $25.3 \pm 11.3\%$ and without arm weight support $24.1 \pm 8.4\%$. When compared to Weber fractions reported in dynamic lifting of real objects (~ 9 – 13% [21]), the Weber fractions found in our study are approximately twice as large. These observed larger Weber fractions could partially be attributed to the experimental setup primarily delivering kinesthetic feedback, with minimal tactile feedback. Indeed, studies with real weights and virtual weights rendered with robotic setups have shown that combining kinesthetic and tactile feedback improves weight perception, compared to only receiving one type of feedback [26], [44]. For example, van Beek et al. observed Weber fractions for static weight perception of 20–30% when providing combined kinesthetic and tactile feedback compared to 30–40% with kinesthetic feedback alone [26]. The difference between the Weber fractions reported by van Beek et al. in the only kinesthetic condition, i.e., the one that better resembles our experimental setup, compared to ours, might be explained by the difference in the lifting task, i.e., they performed static lifting, compared to our dynamic lifting experiment. Overall, given differences in haptic information, dynamics, and virtual-reality-mediated interaction, our results




Fig. 4. Resulting JNDs vs. self-reported Agency (a) and self-reported Compliance (b) per Condition. We added lines visualizing the predicted back-transformed JNDs for each level of Compliance and Condition to illustrate the significant interaction effect. The marginal distributions are shown for each condition. (c) Mean elbow torques around y computed from the static lift measurement of the virtual cylinder with the eight different masses. The individual results, means, and standard deviations are represented by markers, lines, and ribbons, respectively. (d) Elbow torque minima around y computed from the dynamic lift measurements at different cylinder masses. (e) Elbow torque range around y computed from the dynamic lift measurement at different cylinder masses. One data point from subfigure (e) was removed to increase the visibility (torque 16.08 Nm, mass 0.67 kg, condition with AWS).

cannot be directly compared with real-object lifting or virtual lifting with different lifting apparatuses.

Contrary to our hypothesis, we found no statistically significant difference in weight perception when lifting haptically rendered virtual objects with or without arm weight support. Further, our results indicate that the difference between conditions fell within a $\pm 20\%$ margin, and that arm weight support did not worsen weight perception by 10% or more. Motivated by Weber's law, we expected better weight discrimination with arm weight support as it reduces the elbow torque magnitudes. For example, a rendered mass of 1 kg resulted in a mean elbow torque minima of -1.32 Nm with arm weight support and -3.26 Nm without. Using Weber's law ($dS = K \cdot S$; see Section I) as an indirect heuristic, this difference in torque may predict a difference in JNDs, with values for the arm weight support condition expected to be approximately 2.5 times lower than those for the non-AWS condition. However, this prediction assumes a constant Weber fraction K across reference weights and,

thus, across the tested conditions [21]. Yet, the Weber fraction is known to increase drastically for very small stimuli during weight lifting [24]. Engen observed that the Weber fraction for the weight perception of sequentially lifted real objects was close to constant at 3% for standard stimuli from 50 g to 600 g, but nearly doubled when the standard stimulus was at 25 g [24]. Thus, small elbow torques, resulting from small weight forces multiplied by the lower arm length, may also lead to larger Weber fractions. As shown in Fig. 4(d), the distribution of elbow torque minima during lifting with arm weight support includes considerably smaller elbow torque magnitudes, compared to those for lifting without AWS. This suggests that the lower torque levels observed with arm weight support may result in higher Weber fractions, thereby reducing the expected differences in JNDs between conditions. Yet, additional analysis of the effect of elbow torque on weight discrimination performance does not seem to support this explanation (see Supplementary Material, Section III).

B. Higher Compliance Without Arm Weight Support and Higher Agency Were Associated With Better Weight Perception

We observed a positive association between participants' sense of control over their movements, i.e., their agency, and their weight perception. It is possible that participants who felt less in control of their movements in the exoskeleton shifted their attention to managing their movements, which may have impaired their ability to perceive weight differences. While some participants reported during the interview that they found it easier to control their movements in one condition over the other, we found no evidence that agency was consistently higher in either condition, nor did we find an interaction effect between agency and condition on JND. Yet, it is unclear whether greater agency improves weight perception or whether better weight perception leads to a greater sense of agency.

However, we observed a significant interaction between self-reported compliance and condition on JND. Agency and compliance both relate to participants' perception of control during interaction with the exoskeleton—agency reflects the sense of voluntarily controlling movements, while compliance indicates how well the exoskeleton follows the participants' intended motions (i.e., perceived tracking performance). In particular, a decrease in compliance was related to worse weight perception in the condition without arm weight support. The positive relationship between compliance and weight perception may follow the same argumentation as the positive relationship between agency and weight perception, i.e., participants who feel more in control over their movements or feel like the exoskeleton follows their movements better may be able to focus better on the weight they are perceiving. Why this positive relationship between compliance and weight rendering only appears in the condition without arm weight support may be explained by the higher elbow torque ranges participants experienced in the condition without arm weight support compared to the condition with arm weight support. This difference may be attributed to the participants needing to generate greater elbow torques to overcome the resistance of the exoskeleton and initiate movement in the absence of support, despite our effort to enhance the transparency of the exoskeleton. In contrast, with arm weight support, assistive torques were continuously provided, helping to counteract the exoskeleton's resistance and reducing the initial effort required to start a lift. However, as compliance is only one measure that contributes to self-reported transparency via the HRIQ, and since we did not see a significant main effect of condition on weight perception, we argue that the observed interaction does not compromise the primary outcomes of our study.

C. The Experimental Setup is Well-Suited to Evaluate Weight Perception of Virtual Objects

The results from the static and dynamic lift measurements suggest that our experimental setup is well-suited to assess weight perception of haptically rendered virtual objects with and without arm weight support. Participants experienced increasing elbow torque magnitudes as cylinder mass increased, while the presence of arm weight support reduced some of the elbow torque magnitudes from the combined weight of the

arm and cylinder. The relationship between the gravitational forces acting on the virtual objects and the experienced elbow torque magnitudes became particularly clear during static lifting, where increased object mass resulted in larger elbow torque magnitudes. This relationship was also apparent in the elbow torque minima during dynamic lifts; however, with additional variability in elbow torque due to inertial forces acting on the cylinder originating from different accelerations between lifts. These findings align with participants' reports, where none of the participants reported not being able to perceive both the weight of the cylinder and the support provided by AWS. Importantly, we did not find significant interaction effects of condition and mass on elbow torque magnitudes.

D. Implications for Robot-Aided Training

Our finding that providing arm weight support does not hamper the weight perception of haptically rendered virtual objects may have relevant implications for robot-aided training, especially for the rehabilitation of individuals with acquired brain injury. Adjustable arm weight support in upper-limb exoskeletons enhances rehabilitation by increasing arm movement workspace and reducing muscle fatigue [2], [3], thereby enabling high-intensity training crucial for improving functional outcomes [45]. Additionally, haptic rendering provides essential somatosensory feedback by simulating interaction forces with virtual objects and supporting fine motor control [8], [9]. Our findings suggest that users may benefit from the combination of AWS and haptic rendering through upper-limb exoskeletons without compromising the perception of essential somatosensory information. Yet, while preserving accurate perception may be necessary for effective robot-assisted training based on the Guidance-Hypothesis, it is not sufficient to rule out other Guidance-Hypothesis-related risks, such as over-reliance on the provided physical assistance. Importantly, establishing how accurate perception affects functional outcomes in the rehabilitation of individuals with an acquired brain injury remains an open question.

Finally, our findings on the positive relationship between agency and compliance with better weight perception highlight the need to develop robotic systems that promote a sense of control over voluntary movements. This could be achieved in the future through the use of controllers that attempt to promote agency while also optimally assisting users, for example, with model predictive controllers [46].

E. Study Limitations

This study has several limitations. First, we focused only on weight discrimination across two levels of arm weight support, 0% and 50%, and one standard stimulus to keep the study within a reasonable duration. However, in clinical contexts, higher assistance levels (e.g., 80–100% AWS) are also common. Notably, the tested arm weight support level resulted in rather small elbow torque loads, which may not fully capture the potential impact of arm weight support on weight perception. Further, we tested only unimpaired participants. We acknowledge that our findings may therefore not generalize to higher-support regimes

in individuals with acquired brain injury who experience altered somatosensation and muscle weakness.

Another limitation is the choice to rest the elbow on a virtual table and restrict wrist movements. We choose to limit shoulder and wrist movements to limit fatigue and promote a consistent lifting strategy across participants and conditions, in line with previous weight-lifting studies [26], [47]. However, even with these constraints, we cannot fully disentangle differences in weight perception between conditions from concurrent changes in sensorimotor strategy, movement kinematics, or interaction forces. We do not expect substantial differences in weight perception between constrained and unconstrained lifting, since real-object lifting studies report comparable Weber fractions (0.09 – 0.13 [47] and 0.08 – 0.12 [48], respectively).

A further limitation is that the experiment was conducted with the ARMin exoskeleton, PRIDE hand module, and its specific control architecture; therefore, the perceptual effects of AWS may differ for devices with different mechanical and control properties. While the transparency of ARMin has already been well established without AWS [33], we did not explicitly evaluate transparency in the with-AWS condition. The hardware employed also limited the possibility of including left-handed participants. Finally, we did not directly measure participants' joint torques. Incorporating additional measurements, for instance, via electromyography, could improve estimates of joint torques.

V. CONCLUSION

This study investigated the effects of adjustable arm weight support on weight perception during the dynamic lifting of haptically rendered virtual objects. To that end, we developed a sophisticated experimental setup to provide arm weight support and haptic rendering to participants' arms and hands through an upper-limb exoskeleton. Participants visualized the experimental environment with immersive virtual reality. While arm weight support effectively reduced elbow load during dynamic lifting, it did not significantly influence weight perception. Further, we found a positive relationship between agency and better weight perception and between the system-perceived compliance and weight perception (only in the condition without arm weight support), underscoring the importance of designing robotic systems that enhance the feeling of being in control of self-generated movements. These findings provide valuable insights into integrating arm weight support and haptic rendering in robot-aided training environments, which may inform more effective and engaging training programs.

V. DATA AVAILABILITY

The dataset supporting this work is available on the Zenodo repository: <https://doi.org/10.5281/zenodo.17600763>.

ACKNOWLEDGMENT

The authors thank Prof. Raphael Rätz for the support with the Haptic Rendering Environment and Dr. Stefano Dalla Gasperina for his contribution to developing the exoskeleton control. Furthermore, the authors also thank Alberto Garzás Villar for

his advice on the statistical analysis. Finally, the authors thank Alex van den Berg for the insightful conversations and testing of the experimental setup, and Irene Beck for proofreading the manuscript.

REFERENCES

- [1] T. G. Sugar et al., "Design and control of RUPERT: A device for robotic upper extremity repetitive therapy," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 15, no. 3, pp. 336–346, Sep. 2007.
- [2] F. Just, Ö. Özen, S. Tortora, V. Klamroth-Marganska, R. Riener, and G. Rauter, "Human arm weight compensation in rehabilitation robotics: Efficacy of three distinct methods," *J. NeuroEng. Rehabil.*, vol. 17, 2020, Art. no. 13.
- [3] M. Ellis, T. Sukal-Moulton, and J. Dewald, "Impairment-based 3-D robotic intervention improves upper extremity work area in chronic stroke: Targeting abnormal joint torque coupling with progressive shoulder abduction loading," *IEEE Trans. Robot.*, vol. 25, no. 3, pp. 549–555, Jun. 2009.
- [4] F. Guo, L. Liu, M. Li, and R. L. Greene, "The effect of arm support and control modes on muscle fatigue, eye fatigue, and posture changes in mobile game use among young adults," *Hum. Factors Ergonom. Manuf. Service Industries*, vol. 31, pp. 496–505, 2021.
- [5] D. J. Reinkensmeyer and S. J. Housman, "If I can't do it once, why do it a hundred times?": Connecting volition to movement selection in a virtual environment motivates people to exercise the arm after stroke," in *Proc. 2007 Virtual Rehabil.*, 2007, pp. 44–48.
- [6] E. Basalp, P. Wolf, and L. Marchal-Crespo, "Haptic training: Which types facilitate (re)learning of which motor task and for whom? Answers by a review," *IEEE Trans. Haptics*, vol. 14, no. 4, pp. 722–739, Oct.–Dec. 2021.
- [7] V. Klamroth-Marganska et al., "Three-dimensional, task-specific robot therapy of the arm after stroke: A multicentre, parallel-group randomised trial," *Lancet Neurol.*, vol. 13, pp. 159–166, 2014.
- [8] R. Gassert and V. Dietz, "Rehabilitation robots for the treatment of sensorimotor deficits: A neurophysiological perspective," *J. NeuroEng. Rehabil.*, vol. 15, 2018, Art. no. 46.
- [9] F. Danion, J. S. Diamond, and J. R. Flanagan, "The role of haptic feedback when manipulating nonrigid objects," *J. Neurophysiol.*, vol. 107, pp. 433–441, 2012.
- [10] R. Johansson, "How is grasping modified by somatosensory input?," in *Motor Control: Concepts Issues*, vol. 38, pp. 331–355, 1991.
- [11] E. Vidoni, N. Acerra, E. Dao, S. Meehan, and L. Boyd, "Role of the primary somatosensory cortex in motor learning: An rTMS study," *Neurobiol. Learn. Memory*, vol. 93, pp. 532–539, 2010.
- [12] S. Meyer, A. H. Karttunen, V. Thijs, H. Feys, and G. Verheyden, "How do somatosensory deficits in the arm and hand relate to upper limb impairment, activity, and participation problems after stroke? a systematic review," *Phys. Ther.*, vol. 94, pp. 1220–1231, 2014.
- [13] L. Proteau, "Chapter 4 on the specificity of learning and the role of visual information for movement control," in *Vision and Motor Control (Advances in Psychology Series)*, L. Proteau and D. Elliott, Eds. Amsterdam, The Netherlands: North-Holland, 1992, vol. 85, pp. 67–103.
- [14] Ö. Özen, K. A. Buetler, and L. Marchal-Crespo, "Towards functional robotic training: Motor learning of dynamic tasks is enhanced by haptic rendering but hampered by arm weight support," *J. NeuroEngineering Rehabil.*, vol. 19, 2022, Art. no. 19.
- [15] E. Galofaro, E. D'Antonio, N. Lotti, F. Patane, M. Casadio, and L. Masia, "Bimanual motor strategies and handedness role in human-robot haptic interaction," *IEEE Trans. Haptics*, vol. 16, no. 2, pp. 296–310, Apr.–Jun. 2023.
- [16] D. Buongiorno, E. Sotgiu, D. Leonardis, S. Marcheschi, M. Solazzi, and A. Frisoli, "WRES: A novel 3 DoF WRist ExoSkeleton with tendon-driven differential transmission for neuro-rehabilitation and teleoperation," *IEEE Robot. Automat. Lett.*, vol. 3, no. 3, pp. 2152–2159, Jul. 2018.
- [17] R. Rätz, F. Conti, I. Thaler, R. M. Müri, and L. Marchal-Crespo, "Enhancing stroke rehabilitation with whole-hand haptic rendering: Development and clinical usability evaluation of a novel upper-limb rehabilitation device," *J. NeuroEngineering Rehabil.*, vol. 21, pp. 1–28, 2024.
- [18] D. Powell and M. K. O'Malley, "The task-dependent efficacy of shared-control haptic guidance paradigms," *IEEE Trans. Haptics*, vol. 5, no. 3, pp. 208–219, Third Quarter 2012.
- [19] R. A. Schmidt and R. A. Bjork, "New conceptualizations of practice: Common principles in three paradigms suggest new concepts for training," *Psychol. Sci.*, vol. 3, pp. 207–217, 1992.

- [20] R. S. Johansson and G. Westling, "Roles of glabrous skin receptors and sensorimotor memory in automatic control of precision grip when lifting rougher or more slippery objects," *Exp. Brain Res.*, vol. 56, pp. 550–564, 1984.
- [21] L. A. Jones, "Perception of force and weight: Theory and research," *Psychol. Bull.*, vol. 100, pp. 29–42, 1986.
- [22] R. Martín-Rodríguez, A. L. Ratschat, L. Marchal-Crespo, and Y. Vardar, "Tactile weight rendering: A review for researchers and developers," *IEEE Trans. Haptics*, vol. 18, no. 1, pp. 93–109, Jan.–Mar. 2025.
- [23] G. A. Gescheider, *Psychophysics: The Fundamentals*. New York, NY, USA: Psychology Press, 1997.
- [24] T. Engen, "Psychophysics: Discrimination and detection," *Woodworth Schlosberg's Exp. Psychol.*, pp. 11–46, 1971.
- [25] R. L. Gregory and H. E. Ross, "Arm weight, adaptation, and weight discrimination," *Perceptual Motor Skills*, vol. 24, pp. 1127–1130, 1967.
- [26] F. E. V. Beek, R. J. King, C. Brown, M. D. Luca, and S. Keller, "Static weight perception through skin stretch and kinesthetic information: Detection thresholds, JNDs, and PSEs," *IEEE Trans. Haptics*, vol. 14, no. 1, pp. 20–31, Jan.–Mar. 2021.
- [27] C. D. Giachritsis, P. Garcia-Robledo, J. Barrio, A. M. Wing, and M. Ferre, "Unimanual, bimanual and bilateral weight perception of virtual objects in the master finger 2 environment," in *Proc. IEEE Int. Workshop Robot Hum. Interactive Commun.*, 2010, pp. 513–519.
- [28] C. Manning, P. R. Jones, T. M. Dekker, and E. Pellicano, "Psychophysics with children: Investigating the effects of attentional lapses on threshold estimates," *Attention, Percept., Psychophys.*, vol. 80, pp. 1311–1324, 2018.
- [29] W. Wen and H. Imamizu, "The sense of agency in perception, behaviour and human-machine interactions," *Nature Rev. Psychol.*, vol. 1, no. 4, pp. 211–222, 2022.
- [30] R. Oldfield, "The assessment and analysis of handedness: The Edinburgh inventory," *Neuropsychologia*, vol. 9, pp. 97–113, 1971.
- [31] F. Just et al., "Exoskeleton transparency: Feed-forward compensation vs. disturbance observer," *At-Automatisierungstechnik*, vol. 66, pp. 1014–1026, 2018.
- [32] R. Rätz, F. Conti, R. M. Müri, and L. Marchal-Crespo, "A novel clinical-driven design for robotic hand rehabilitation: Combining sensory training, effortless setup, and large range of motion in a palmar device," *Front. Neurobot.*, vol. 15, 2021, Art. no. 748196.
- [33] S. D. Gasperina, A. L. Ratschat, and L. Marchal-Crespo, "Quantitative and qualitative evaluation of exoskeleton transparency controllers for upper-limb neurorehabilitation," in *Proc. Int. Conf. Rehabil. Robot.*, 2023, pp. 1–6.
- [34] D. A. Winter, *Biomechanics and Motor Control of Human Movement*, 3rd ed. Hoboken, NJ, USA: Wiley, 2006.
- [35] R. Rätz and L. Marchal-Crespo, "Physics engine-based whole-hand haptic rendering for sensorimotor neurorehabilitation," in *Proc. IEEE World Haptics Conf.*, 2023, pp. 279–285.
- [36] A. Ratschat, T. M. Lomba, S. D. Gasperina, and L. Marchal-Crespo, "Development and validation of a kinematically accurate upper-limb exoskeleton digital twin for stroke rehabilitation," in *Proc. IEEE Int. Conf. Rehabil. Robot.*, 2023, pp. 1–6.
- [37] F. Kingdom and N. Prins, *Psychophysics*. London, U.K.: Elsevier, 2016.
- [38] M. A. García-Pérez, "A cautionary note on the use of the adaptive up-down method," *J. Acoust. Soc. Amer.*, vol. 130, pp. 2098–2107, 2011.
- [39] M. A. García-Pérez, "Forced-choice staircases with fixed step sizes: Asymptotic and small-sample properties," *Vis. Res.*, vol. 38, pp. 1861–1881, 1998.
- [40] M. A. García-Pérez and R. Alcalá-Quintana, "Sampling plans for fitting the psychometric function," *Spanish J. Psychol.*, vol. 8, pp. 256–289, 2005.
- [41] N. Prins, "The psychometric function: The lapse rate revisited," *J. Vis.*, vol. 12, pp. 25–25, 2012.
- [42] S. G. Hart, "NASA-task load index (NASA-TLX); 20 years later," *Proc. Hum. Factors Ergonom. Soc. Annu. Meeting*, vol. 50, pp. 904–908, 2006.
- [43] E. McAuley, T. Duncan, and V. V. Tammien, "Psychometric properties of the intrinsic motivation inventory in a competitive sport setting: A confirmatory factor analysis," *Res. Quart. Exercise Sport*, vol. 60, pp. 48–58, 1989.
- [44] C. Giachritsis, R. Wright, and A. Wing, "The contribution of proprioceptive and cutaneous cues in weight perception: Early evidence for maximum-likelihood integration," in *Haptics: Generating and Perceiving Tangible Sensations* (Lecture Notes in Computer Science), vol. 6191, A.M.L. Kappers, J.B.F. van Erp, W.M. Bergmann Tiest, F.C.T. van der Helm, Eds. Berlin, Heidelberg: Springer, 2010, pp. 11–16.
- [45] G. Kwakkel et al., "Effects of augmented exercise therapy time after stroke: A meta-analysis," *Stroke*, vol. 35, pp. 2529–2536, 2004.

- [46] Ö. Özen, K. A. Buetler, and L. Marchal-Crespo, "Promoting motor variability during robotic assistance enhances motor learning of dynamic tasks," *Front. Neurosci.*, vol. 14, 2021.
- [47] E. E. Brodie and H. E. Ross, "Sensorimotor mechanisms in weight discrimination," *Percept. Psychophys.*, vol. 36, pp. 477–481, 1984.
- [48] H. E. Ross and E. E. Brodie, "Weber fractions for weight and mass as a function of stimulus intensity," *Quart. J. Exp. Psychol. Sect. A*, vol. 39, pp. 77–88, 1987.



Alexandre L. Ratschat (Graduate Student Member, IEEE) received the B.Sc. and M.Sc. degrees in mechanical engineering from ETH Zurich, Zürich, Switzerland. He is currently working toward the Ph.D. degree with the Delft University of Technology, Delft, The Netherlands. From 2019 to 2020 he was with F&P Robotics AG, Switzerland as a Software Engineer. He is also affiliated with the Department of Rehabilitation Medicine at Erasmus MC, Rotterdam, The Netherlands. His research interests include leveraging robotic devices and haptics for upper-limb rehabilitation and motor learning of people with acquired brain injury.



David A. Abbink (Senior Member, IEEE) received the M.Sc. and Ph.D. degrees in mechanical engineering from the Delft University of Technology, Delft, The Netherlands, in 2002 and 2006, respectively. He is currently a Full Professor of haptic human-robot interaction with the Delft University of Technology, with a joint appointment with the Faculty of Mechanical Engineering and the Faculty of Industrial Design Engineering. He is also Leading FRAM, a transdisciplinary research and innovation center shaping the future of physical work. He has authored or coauthored extensively on human-robot interaction, haptics and sensorimotor control. He founded the Delft Haptics Lab, and was an Associate Editor for IEEE TRANSACTION ON HAPTICS.



Gerard M. Ribbers received the Ph.D. degree from Erasmus MC, Rotterdam, The Netherlands, in 2001. He was appointed Special Professor of neurorehabilitation with Erasmus MC, in 2012. He is currently the Head of the Department of Rehabilitation Medicine with Erasmus MC. He trained in rehabilitation medicine with Sint Maartenskliniek, Nijmegen from 1991 to 1995 and joined the Rijndam Rehabilitation Center in 1995, where, for 24 years, he led a clinical department for patients with severe cognitive and behavioral regulation disorders due to acquired brain injury. His clinical and research interests include early, integrated rehabilitation care and recovery after stroke and traumatic brain injury.



Laura Marchal-Crespo (Member, IEEE) received the M.Sc. and Ph.D. degrees from the University of California at Irvine, Irvine, CA, USA. She was a Postdoctoral Researcher with ETH Zurich and an Assistant Professor with the University of Bern, Bern, Switzerland. She is currently an Associate Professor with the Delft University of Technology, Delft, The Netherlands, since 2020, where she led the Human-Robot Interaction Group. She is also an Honorary Associate Professor with the Erasmus Medical Center, Rotterdam, The Netherlands. Her research interests include human-machine interfaces and biological learning, and, specifically, the use of robotics and virtual reality to aid people to rehabilitate after neurologic injuries.