

# User-customizable Shared Control for Robot Teleoperation via Virtual Reality

Rui Luo<sup>1\*</sup>, Mark Zolotas<sup>1\*</sup>, Drake Moore<sup>1</sup>, and Taşkın Padır<sup>1,2</sup>

**Abstract**—Shared control can ease and enhance a human operator’s ability to teleoperate robots, particularly for intricate tasks demanding fine control over multiple degrees of freedom. However, the arbitration process dictating how much autonomous assistance to administer in shared control can confuse novice operators and impede their understanding of the robot’s behavior. To overcome these adverse side-effects, we propose a novel formulation of shared control that enables operators to tailor the arbitration to their unique capabilities and preferences. Unlike prior approaches to “customizable” shared control where users could indirectly modify the latent parameters of the arbitration function by issuing a feedback command, we instead make these parameters *observable* and *directly* editable via a virtual reality (VR) interface. We present our user-customizable shared control method for a teleoperation task in  $SE(3)$ , known as the buzz wire game. A user study is conducted with participants teleoperating a robotic arm in VR to complete the game. The experiment spanned two weeks per subject to investigate longitudinal trends. Our findings reveal that users allowed to interactively tune the arbitration parameters across trials generalize well to adaptations in the task, exhibiting improvements in precision and fluency over direct teleoperation and conventional shared control.

## I. INTRODUCTION

Teleoperation is the foundation behind many robotic applications. These applications range from preventing human presence in hazardous environments to providing health-care services through remote patient care and robot-assisted surgery. Nevertheless, remotely controlling a complex robot warrants a certain level of operator skill and domain expertise in order to successfully complete challenging teleoperation tasks [1]. Without this baseline, a novice operator may be exposed to detrimental levels of physical and cognitive workload [2], potentially resulting in catastrophic task failures [3]. A common means of relieving this burden on the operator is to employ an autonomous controller that continuously assists the human user by sharing control over the robot [4].

While shared control alleviates excess workload exerted on operators, it may also give rise to new issues that impact task performance and impede dexterity training [5]. One prominent example of such an issue is *model misalignment*. This phenomenon occurs when the division in control between




Fig. 1. A user wearing a virtual reality (VR) headset teleoperates the Kinova Gen 3 robotic arm to play the buzz wire game. The VR controller acts as a handle in the virtual world, where its simulated effects also have a direct consequence on the physical counterpart played by the robot.

the human and autonomy creates a misunderstanding in how the operator expects the robot to behave [6]. In direct teleoperation, model misalignment may arise whenever a robot has higher degrees-of-freedom (DoF) than the user’s control interface [7]. Despite this interface asymmetry in DoF, a user with adequate training should still gradually build a mental model of the robot’s behavior. In contrast, the extra layer of arbitration introduced in shared control exacerbates the user’s difficulty in comprehending the system. A communication medium for user feedback is therefore essential to how an operator understands the arbitration [4].

Aside from equipping operators with suitable feedback, another integral aspect of shared control is *user customizability*. Over a long-term interaction, operators will experience variations in their capabilities, preferences, and environments [8]. Shared control must account for these changes and adapt accordingly, especially in terms of how control authority is assigned [9]. Prior works have developed promising frameworks for adaptive shared control by enabling users to interactively customize the parameters that characterize the arbitration function [10], [11] or suggest corrections for the resulting robot behavior [12], [13]. However, these works rely entirely on the robot’s legible state to convey the effects of a user’s modifications to the arbitration procedure.

In this paper, we establish user-customizable shared control by *directly* communicating the internal arbitration parameters to an operator for refinement. This transparency allows users to better understand the arbitration process and thus effectively modify it for personalized outcomes. We ground this idea in a teleoperation task where operators must control

\*These authors contributed equally to this work.

This material is based upon work supported by the National Science Foundation under Award No. 1928654.

<sup>1</sup>Institute for Experiential Robotics (IER), Northeastern University, Boston, Massachusetts, USA. {luo.rui, m.zolotas, moore.dr, t.padir}@northeastern.edu

<sup>2</sup>Taşkın Padır holds concurrent appointments as a Professor of Electrical and Computer Engineering at Northeastern University and as an Amazon Scholar. This paper describes work performed at Northeastern University and is not associated with Amazon.

a 7-DoF robotic arm in virtual reality (VR) to complete the buzz wire game (see Fig. 1). The buzz wire game is a suitable testbed for fine robot control, as it requires both translational and rotational motion, significant eye-hand coordination, and sustained operator focus [14], [15]. Moreover, the game is frequently utilized for skill training in domains that would benefit from robot teleoperation, such as surgery [16]. To assess the quality and longevity of the proposed shared control method, we conducted a user study to analyze how participant performance evolves over repeated interactions.

The key contributions of this paper are as follows:

- A novel mathematical framework that factors in user feedback to formalize user-customizable shared control;
- A comprehensive demonstration of this framework within the context of a teleoperation task in SE(3), involving a real 7-DoF robotic arm and a VR interface as the bidirectional communication channel;
- Results from a longitudinal user study with 12 subjects teleoperating the robotic arm to perform the buzz wire game over four sessions spread across two weeks.

## II. RELATED WORK

Sensory feedback can be supplied across multiple modalities in teleoperation. For example, haptic devices are widely used in bilateral teleoperation to simultaneously enhance an operator’s task efficiency and encourage growth in skill [2], [4]. Other modalities, *e.g.*, audio, force, and visual feedback, have enjoyed similar success [17]–[19]. In recent years, VR headsets have proven advantageous in complex manipulation tasks where high-DoF robots are remotely controlled [20]–[22]. Furthermore, VR interfaces offer operators the opportunity to modify the inner workings of the teleoperated robot, which is an asset we exploit in this work.

In shared control teleoperation, the challenge is to also effectively “blend” user inputs with an autonomous controller’s outputs. This blending or arbitration process is typically described by parameters that determine how control is allocated. Arbitration parameters may be fixed, *e.g.*, to favor the human leading, or dynamically updated based on heuristics [4], such as safety, fluency, and confidence [23]–[25]. These parameters can also be personalized to the user by balancing out their capability to independently complete a task with their need for assistance [26], [27]. However, striking the desired balance in autonomy is an open problem [9], as heuristically tuning arbitration parameters may not be suitable for all users – even when the tuning is personalized.

Therefore, it may be more appropriate to empower users with the ability to *customize* the arbitration process. While the literature on customizable shared control is sparse, some previous works have successfully granted this ability via user feedback [10], [11] or “corrections” [12], [13]. Hence, users can haptically or verbally correct arbitration parameters on a trial-by-trial basis, *e.g.*, requesting “less assistance” for the next interaction. Although effective, these works relied solely on the robot’s appearance and movement to communicate the underlying autonomous assistance, which is often insufficient at explaining arbitration to users [27], [28]. Instead, our paper

opts for greater transparency by externalizing the factors that influence arbitration and making them editable to operators.

## III. USER-CUSTOMIZABLE SHARED CONTROL

In the following, we propose a mathematical framework for shared control that enables a user to dictate how the robot autonomy should provide teleoperation assistance.

### A. Problem Setup

The shared control system considered in this work is driven by the following interaction data. At time  $t$ , a human operator’s measured state,  $\mathbf{x}_h(t) \in \mathbb{R}^{n_h}$  (*e.g.*, joystick deflections), is fed into a control interface to produce control commands,  $\mathbf{u}_h(t) \in \mathbb{R}^{m_h}$ , for direct robot teleoperation. Similarly, a robot’s state and the autonomy’s assistive commands are  $\mathbf{x}_r(t) \in \mathbb{R}^{n_r}$  (*e.g.*, joint angles) and  $\mathbf{u}_r(t) \in \mathbb{R}^{m_r}$ , respectively. The output commands of the shared control system,  $\mathbf{u}_{sc}(t) = \beta_{\theta}(\mathbf{u}_r(t), \mathbf{u}_h(t))$ , adhere to an arbitration function,  $\beta(\cdot, \cdot)$ , parameterized by  $\theta \in \mathbb{R}^L$ . For instance, a linear blending scheme can be written as:

$$\beta_{\theta}(\mathbf{u}_r(t), \mathbf{u}_h(t)) = (1 - \alpha_{\theta})\mathbf{u}_r(t) + \alpha_{\theta}\mathbf{u}_h(t), \quad (1)$$

where  $\alpha_{\theta} \in [0, 1]$  is the blending variable [10]. While  $\alpha$  is usually a scalar, we will adopt a matrix representation for more granular authority over the system behavior. Let  $\mathbf{A}_{\theta} = \text{diag}(\alpha_1, \dots, \alpha_{m_r})$  be a positive definite diagonal arbitration matrix, with  $\forall i, \alpha_i \in [0, 1]$ , such that linear blending is:

$$\beta_{\theta}(\mathbf{u}_r(t), \mathbf{u}_h(t)) = (\mathbf{I} - \mathbf{A}_{\theta})\mathbf{u}_r(t) + \mathbf{A}_{\theta}\mathbf{u}_h(t), \quad (2)$$

where  $\mathbf{I}$  is an identity matrix of appropriate dimensions.

Selecting the parameters,  $\theta$ , of an arbitration function,  $\beta_{\theta}(\cdot, \cdot)$ , is critical in shaping the robot’s assistive behavior. Instead of manually selecting a fixed  $\theta$ , an “optimal” selection is often derived by solving an optimization problem with a sensible objective,  $\Gamma$ , like minimizing human effort. The issue with this approach is that  $\Gamma$  may not be the true cost function of the shared control,  $\Gamma^*$ , nor would a tractable solution necessarily exist [10].

### B. Feedback-informed User Optimization

To circumvent the intractable problem of identifying  $\Gamma^*$  and its “optimal” parameters, we present a method that allows users to customize the arbitration function by configuring their preferred  $\theta$ . The core concept of our method is to establish a *bidirectional* communication channel between the human and robot. We formalize this channel as a vector-valued function  $\psi(\cdot, \cdot)$  (dropping  $t$  for brevity):

$$\theta \leftarrow \psi(\mathbf{v}_r, v_h), \quad (3)$$

where the robot administers feedback,  $\mathbf{v}_r$ , to the operator about the shared control, while the user performs parameter updates via an interface action,  $v_h$ . For example, parameter updates could be obtained from discrete verbal commands made by the operator,  $v_h \in \mathbb{Z}^+$ , *e.g.*, “more/less assistance”. Most previous efforts at user customization of  $\theta$  have solely depended on the robot’s legible state,  $\mathbf{v}_r = \mathbf{x}_r$ , for operator feedback [10], [11].

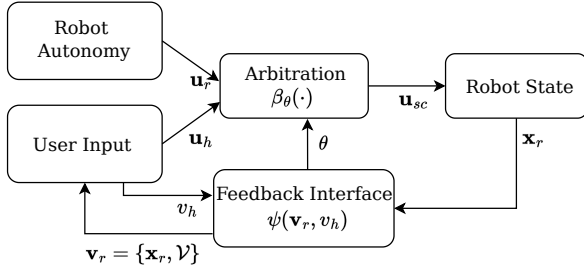


Fig. 2. The user-customizable shared control architecture.

Using this channel formulation to obtain parameters,  $\theta$ , we represent the arbitration matrix,  $\mathbf{A}_\theta$ , as follows:

$$\mathbf{A}_\theta = \text{diag}(\boldsymbol{\alpha}) = \text{diag}(\alpha_1, \dots, \alpha_{m_r}), \quad (4)$$

$$\text{where } \alpha_j = \sum_{l=1}^L w_l \psi_l(\mathbf{v}_r, v_h), \quad j = 1, \dots, m_r$$

Each scalar blending quantity,  $\alpha_j$ , is governed by potentially more than one parameter,  $\theta_l$ , output as an element of the channel,  $\psi_l(\mathbf{v}_r, v_h)$ , depending on the selection weight,  $w_l$ . Our overall framework architecture is shown in Fig. 2.

Since the feedback variables,  $\mathbf{v}_r$  and  $v_h$ , dramatically impact how users perceive the shared control, we recommend two requirements for generating these variables. First, the robot should communicate more than its external state,  $\mathbf{x}_r$ , to help operators tune  $\theta$ . In Section IV, we demonstrate how a set of human-interpretable cues,  $\mathcal{V}$  (*i.e.*, visualizations in VR), can be incorporated into  $\mathbf{v}_r = \{\mathbf{x}_r, \mathcal{V}\}$ . Second, an operator's parameter updates should have a proportionate effect on robot behavior. This proportionality can be achieved by letting the operator submit actions,  $v_h \in [0, 1]$  (*e.g.*, using a graphical slider in VR), over the same constrained range as the arbitration parameters,  $\forall l, \theta_l \in [0, 1]$ .

#### IV. PROPOSED FRAMEWORK INSTANTIATION

To illustrate the operational principles of the proposed framework, this section provides an instantiation for assistive teleoperation in the scope of the buzz wire game.

##### A. Testbed: Buzz Wire Game

In the buzz wire game, the operator's goal is to carefully guide a loop handle in  $SE(3)$  from start to end while minimizing collisions with the wire. The game is a well-known assessment environment for fine motor skill and eye-hand coordination [14], [15]. Moreover, 3D buzz wire games are advantageous when evaluating dexterity tasks that benefit from stereoscopic depth perception [14]. As precise motor control, eye-hand coordination, and 3D spatial awareness are critical capabilities for teleoperation in general, we deem this game to be a pertinent testbed. Immersive VR versions of the game have also been validated as an effective simulation for operator training, *e.g.*, in stroke rehabilitation [29].

Fig. 1 demonstrates our buzz wire game replica for robot teleoperation. The layout is comprised of a Kinova Gen 3 arm with a loop handle extension at its end-effector to complete the physical game, as well as a person wearing an HTC Vive headset to perform the same task in VR with

a handheld controller. The virtual scene is an environment model reconstruction of the buzz wire game and robot, which has been shown to be less cognitively demanding and more usable in VR contexts than point cloud or image renderings [30]. An electric circuit board on the physical setup also imitates the game's "buzz" feature by detecting collisions. Collisions are then relayed in VR via flashing red lights and vibrations in the handheld input device.

##### B. Shared Assistance in Teleoperation

In teleoperation, a common strategy is to control the robot's end-effector pose,  $\mathbf{x}_r \in \mathbb{R}^6$ , by sending desired twist commands,  $\dot{\mathbf{x}}_d \in \mathbb{R}^6$ , via admittance control:

$$\dot{\mathbf{x}}_d = \mathbf{M}^{-1} \int (\mathbf{k} \odot \mathbf{e} + \mathbf{d} \odot \dot{\mathbf{e}} + \mathbf{f}_{\text{ext}}) dt, \quad (5)$$

where  $\mathbf{M} \in \mathbb{R}^{6 \times 6}$ ,  $\mathbf{k} \in \mathbb{R}^6$ , and  $\mathbf{d} \in \mathbb{R}^6$  are mass, stiffness, and damping terms, respectively, while  $\mathbf{f}_{\text{ext}} \in \mathbb{R}^6$  is the externally applied force,  $\mathbf{e} \in \mathbb{R}^6$  is the pose error vector, and  $\odot$  denotes component-wise multiplication. In our buzz wire setup,  $\mathbf{e}$  is the difference between the handheld VR controller's pose,  $\mathbf{x}_h \in \mathbb{R}^6$ , and the robot's state,  $\mathbf{x}_r \in \mathbb{R}^6$ .

Robot autonomy can then be integrated into Eq. (5) as an assistive force by reformulating the linear blending in Eq. (2) to create an admittance control policy for shared control:

$$\mathbf{u}_{sc} = \dot{\mathbf{x}}_d = (\mathbf{I} - \mathbf{A}_\theta) \mathbf{u}_r + \mathbf{A}_\theta \mathbf{u}_h, \quad (6)$$

$$\mathbf{u}_h = \mathbf{M}^{-1} \int (\mathbf{k} \odot \mathbf{e} + \mathbf{d} \odot \dot{\mathbf{e}}) dt, \quad (7)$$

$$\mathbf{u}_r = \mathbf{M}^{-1} \int \mathbf{w}_a dt. \quad (8)$$

Here,  $\mathbf{f}_{\text{ext}}$  is replaced by a wrench,  $\mathbf{w}_a \in \mathbb{R}^6$ , that represents autonomous assistance in translational and rotational space.

##### C. Real-time Assistive Wrench

Generating real-time autonomous assistance for teleoperation is vital in ensuring reactive and smooth robot motion. To fulfil this requirement, we efficiently compute the assistive wrench,  $\mathbf{w}_a$ , using potential fields, where targets produce attractive vectors and obstacles yield repulsive vectors. Attractive and repulsive vectors are typically calculated between a pair of control and environment points [31]. In our setting, we define *eight* control points as the vertices of an octagon that encloses the circular end-effector. The environment points are acquired in real-time given point cloud data detected using two depth cameras.

Environment points within a pre-defined range of the control points compose a neighborhood set,  $\mathcal{C} = \mathcal{C}_{\text{att}} \sqcup \mathcal{C}_{\text{rep}}$ , where  $|\mathcal{C}| = N$ . The  $\mathcal{C}_{\text{att}}$  subset contains attractive points, whereas  $\mathcal{C}_{\text{rep}}$  contains repulsive points. The potential forces between an environment point,  $\mathbf{p}_k \in \mathcal{C}$ , and a control point,  $\mathbf{p}_i \in \mathbb{R}^3$ , are computed based on a logistic function:

$$\mathbf{f}_{i,k} = \frac{f_{\text{max}}}{1 + \exp \left[ -\lambda_{i,k} \left( \frac{\|\mathbf{d}\|}{\rho} - d_{i,k} \right) \right]} \hat{\mathbf{d}}_{i,k}, \quad (9)$$

where each function-shaping parameter  $\lambda_{i,k} \in \{\lambda_{\text{rep}}, \lambda_{\text{att}}\}$ ,  $d_{i,k} \in \{d_{\text{rep}}, d_{\text{att}}\}$ , and  $\hat{\mathbf{d}}_{i,k} = \pm \frac{\mathbf{d}_{i,k}}{\|\mathbf{d}_{i,k}\|}$  has two possible

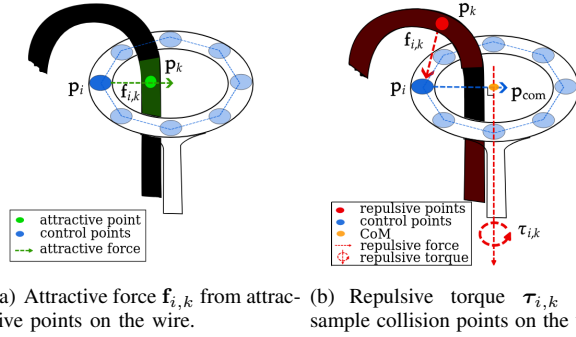


Fig. 3. Illustration of how the attractive and repulsive forces are computed given a control point  $\mathbf{p}_i$  (dark blue) and environment point  $\mathbf{p}_k$  (green/red).

values depending on whether the relationship between  $\mathbf{p}_k$  and  $\mathbf{p}_i$  is repulsive or attractive. Hence, forces  $\mathbf{f}_{i,k}$  are either repulsive or attractive, with  $f_{\max}$  determining the maximum magnitude of the generated force and  $\rho$  adjusting the change rate in force given distance.

As the buzz wire game requires 6-DoF teleoperation, we compute the assistive torque  $\boldsymbol{\tau}_{i,k}$  for rotational motion as:

$$\boldsymbol{\tau}_{i,k} = \mathbf{f}_{i,k} \times (\mathbf{p}_{\text{com}} - \mathbf{p}_i), \quad (10)$$

where  $\mathbf{p}_{\text{com}}$  is the center of the mass of the robot’s end-effector. Fig. 3 illustrates an example of how attractive forces and repulsive torques are calculated for the buzz wire game from a pair of control and environment points ( $\mathbf{p}_i, \mathbf{p}_k$ ).

The resultant assistive wrench,  $\mathbf{w}_a = [\mathbf{f}_{\text{net}}^T \ \boldsymbol{\tau}_{\text{net}}^T]^T$ , is derived by aggregating a force component,  $\mathbf{f}_{\text{net}} \in \mathbb{R}^3$ , and a torque component,  $\boldsymbol{\tau}_{\text{net}} \in \mathbb{R}^3$ , each calculated as:

$$\mathbf{f}_{\text{net}} = \sum_i \frac{\mathbf{f}_i}{\|\mathbf{f}_i\|} \cdot \max_k \|\mathbf{f}_{i,k}\|, \quad (11)$$

$$\boldsymbol{\tau}_{\text{net}} = \sum_i \frac{\boldsymbol{\tau}_i}{\|\boldsymbol{\tau}_i\|} \cdot \max_k \|\boldsymbol{\tau}_{i,k}\|, \quad (12)$$

where  $\mathbf{f}_i = \sum_k \mathbf{f}_{i,k}$  and  $\boldsymbol{\tau}_i = \sum_k \boldsymbol{\tau}_{i,k}$ .

#### D. User-customizable Arbitration

A VR interface acts as the basis for bidirectional communication,  $\psi(\mathbf{v}_r, v_h)$ . We created a spider chart interface, depicted in Fig. 4, to augment the robot’s feedback,  $\mathbf{v}_r$ , with a variety of human-readable indicators,  $\mathcal{V}$ . Before introducing these indicators, we emphasize that the spider chart is *only visible* when a user wishes to edit the shared control. Hence, no extra cognitive workload is induced during a game trial.

The  $\mathcal{V}$  set consists of three arbitration factors displayed in VR: *speed*, *depth assistance*, and *turnability*. *Speed* regulates translational velocity along axes of motion relevant to the buzz wire game (e.g., horizontally along the wire). *Depth assistance* governs the margin between the end-effector and wire along the depth axis, with the aim of assisting line-of-sight occlusions. *Turnability* adjusts the rotational speed of the end-effector. Users can therefore examine the current parameters,  $\boldsymbol{\theta}$ , via the spider chart, as well as supply edits per factor,  $v_h$ , in increments/decrements (+/- buttons in Fig. 4). Although we defined these three factors for the buzz wire game, regulating motion in relevant DoF is widely used in

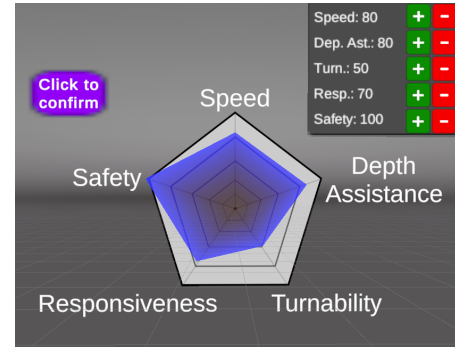


Fig. 4. The customizable spider chart interface in virtual reality. The operator can modify any of the five factors in steps of 5 before confirming their choice. Each factor is scaled from  $[0.1, 1]$  to  $[10, 100]$ .

shared control [9] (e.g., virtual fixture methods), and maintaining a clear line-of-sight is crucial for robot teleoperation in general [18]. We thus believe these arbitration factors can be easily adapted to many different robot teleoperation tasks.

To apply user edits  $v_h$  to the arbitration, we express  $\boldsymbol{\alpha} \in \mathbb{R}^6$  from Eq. (4) as a linear combination:

$$\boldsymbol{\alpha} = \mathbf{W}\boldsymbol{\theta} = \mathbf{W}\psi(\mathbf{v}_r, v_h) = \mathbf{W} \begin{bmatrix} v_h^{\text{speed}} \\ 1 - v_h^{\text{depth}} \\ v_h^{\text{turn}} \end{bmatrix}, \quad (13)$$

$$\text{where } \mathbf{W} = \begin{bmatrix} 1 & 0 & 1 & 0 & 0.4 & 0.2 \\ 0 & 1 & 0 & 0.2 & 0 & 0 \\ 0 & 0 & 0 & 0.8 & 0.6 & 0.8 \end{bmatrix}^T.$$

The coefficient matrix,  $\mathbf{W} \in \mathbb{R}^{6 \times 3}$ , is structured in a manner where each editable facet,  $v_h$ , directly maps to a parameter with a semantically similar effect on arbitration. The weighting of elements was chosen empirically to fit the task. For example,  $v_h^{\text{turn}}$  is masked by  $\mathbf{W}$  to contribute to an 80% change in arbitration on rotation about the z-axis (yaw), while  $v_h^{\text{speed}}$  contributes to the remaining 20%. Here, the term *turnability* intuitively represents more influence over robot rotation than translational speed. The weighting criteria for *depth assistance* uses the complement of  $v_h^{\text{depth}}$ , as “assistance” implies more contribution from the autonomy.

We also expose two task-independent factors through  $\mathcal{V}$ : *safety* and *responsiveness*. These factors adjust  $f_{\max}$  in Eq. (9) and  $\mathbf{k}$  in Eq. (7), respectively, akin to a variable impedance controller. Using *safety* to tune the force magnitude and *responsiveness* to adjust the stiffness are generalizable to any admittance control scheme. Each of the five configurable factors in  $\mathcal{V}$  is constrained in the range  $[0.1, 1]$  to prevent a “full autonomy” configuration.

## V. EXPERIMENT

To gauge the effectiveness of the proposed framework, we conducted a between-subjects study involving 12 subjects (3 female; aged 20-34). Subjects were asked to complete the buzz wire game by teleoperating a robotic arm in VR. All participants provided written consent prior to data collection.

### A. Experiment Protocol

Participants were allocated into one of three groups, each designed to evaluate a different control strategy: (1) direct




Fig. 5. New wire testbed for the last session to assess skill generalization.

teleoperation (**teleop**), which is implemented by setting  $\mathbf{A}_\theta = \mathbf{I}$  in Eq. (6); (2) shared control with a heuristics-based arbitration function (**sc**), as outlined in Section V-B; and (3) our user-customizable shared control method (**sc\_user**). Participants were unaware of the distinctions between each group. Subjects in **sc** and **sc\_user** were provided with identical written explanations of the five factors from Section IV-D before the experiment. For instance, “Increases/Decreases in this parameter mean more/less robot assistance in avoiding obstacles” explained *safety*. It is also worth noting that an equal distribution of the subject pool reported no background in robotics, which was preserved across control strategies.

Each group consisted of four participants and every subject carried out five consecutive game trials in a single session using their assigned control strategy. This procedure was repeated for a total of four sessions. Subjects were also offered five minutes before their first session to become familiar with the VR setup while not playing the game. This included a brief overview of the spider chart shown in Fig. 4 for the **sc** and **sc\_user** groups. To evaluate the learning rates of operators over an extended period, a minimum of one workday separated each session, resulting in an experiment spanning two weeks. Additionally, the last session involved a distinct wire testbed, as portrayed in Fig. 5. This alteration aimed to assess whether any teleoperation skills gained are transferable to a related, yet different task. Subjects also answered post-session surveys on their user experience.

After each trial, participants could spend a minute operating the arm without playing the game. For the **teleop** group, this post-trial time could be used as an opportunity to improve their understanding of direct robot teleoperation. Participants in **sc** could instead examine the automatically updated arbitration factors (see Section V-B) through the spider chart and test out the resulting control. Only the **sc\_user** group could use this post-trial time to edit the arbitration factors displayed on the spider chart and explore different configurations. User-selected arbitration settings persisted across trials unless altered. During game trials, the spider chart was *not* visible and arbitration settings were fixed.

### B. Heuristics-based Shared Control

We implemented a heuristics-based assist-as-needed controller as the reference method, **sc**, for our experiment. The arbitration parameters,  $\theta$ , in this controller are iteratively updated according to the operator’s performance after each trial. We linearly updated the parameters with a change rate

akin to the multi-trial update function from [26]:

$$\theta_{t+1} = (1 + \chi_t)\theta_t, \quad (14)$$

where  $t$  is the trial index and  $\chi_t \in \mathbb{R}^3$  is the change rate vector, which is updated using the two most recent trials:

$$\chi_t^i = \chi_{\text{nom}} \frac{\bar{r}_t^i - r_d^i}{r_d^i} \left( \frac{|\bar{r}_t^i - r_d^i|}{|\bar{r}_{t-1}^i - r_d^i|} \right)^{\text{sign}(r_d^i - \bar{r}_t^i)}, \quad (15)$$

where  $r_d^i$  is the  $i^{\text{th}}$  component of the desired error vector  $\mathbf{r}_d$  (pre-calculated from an expert demonstration),  $\bar{r}_t^i$  is the  $i^{\text{th}}$  component of the average error vector  $\bar{\mathbf{r}}_t$  during trial  $t$ , and  $\chi_{\text{nom}}$  is the nominal change rate, which determines the maximum change rate within one update step.

Distance from the end-effector to the closest point along the wire is a suitable performance metric to denote safety in the buzz wire game. This error is conveniently contained in the assistive wrench,  $\mathbf{w}_a$ , such that  $\bar{\mathbf{r}}_t$  can be expressed as:

$$\bar{\mathbf{r}}_t = \mathbf{W}^\dagger \bar{\mathbf{w}}_{a,t}, \quad (16)$$

where  $\bar{\mathbf{w}}_{a,t} \in \mathbb{R}^6$  is the average wrench from trial  $t$  and  $\mathbf{W}^\dagger$  is the pseudo-inverse of the coefficient matrix in Eq. (13).

### C. Evaluation Metrics

Performance evaluation for the buzz wire game primarily revolves around two measures: successful task completion time and number of collisions. A trial is deemed successful and collisions are counted when the participant teleoperates the robot’s end-effector to the game end (a pre-defined location along the wire) without encountering a fatal failure, *i.e.*, when collision forces exceed a threshold. Aside from task time and collisions, we also examine a user’s *jerk* (the mean squared measure from [32]) in handling the VR controller to capture smoothness or fluency. Moreover, we explore the learning trajectories of users, *i.e.*, each user’s change in “skill” metrics relative to their first session. Lastly, we investigate a user’s perceptions of the different control modes in terms of *skill*, *success*, and *difficulty*.

Therefore, the hypotheses for this study are as follows:

- **H1:** Subjects will report higher ratings on skill, success, and difficulty for **sc\_user** than **teleop** and **sc**.
- **H2:** There will be faster task success times and less collisions for **sc\_user** than **teleop** and **sc**.
- **H3:** Participants in **sc\_user** will exhibit better teleoperation performance relative to their ‘baseline’ session when compared against users from **teleop** and **sc**.

## VI. RESULTS

In the following, we report on our results and test our hypotheses. When evaluating the learning trajectories of participants, we term the first session as the ‘baseline’, the combined second and third sessions as ‘training’, and the last session as ‘transfer’. Unless otherwise stated, mixed ANOVAs were performed to analyze the effects of control modes (between-subjects) and sessions (within-subjects) on the relevant evaluation metrics. The ‘baseline’ session was left out of the ANOVAs, leaving only the combined ‘training’ sessions to compare against the ‘transfer’ condition.

TABLE I  
SUBJECTIVE RESULTS ACROSS SESSIONS (1 = “VERY LOW” OR “FAILURE”, 10 = “VERY HIGH” OR “PERFECT”)

Sessions	Mode	Skill $\uparrow$	Success $\uparrow$	Difficulty $\downarrow$
1 (baseline)	sc	<b>7.5 <math>\pm</math> 1.0</b>	<b>7.5 <math>\pm</math> 1.3</b>	<b>4.8 <math>\pm</math> 2.2</b>
	sc_user	6.3 $\pm$ 0.5	6.0 $\pm$ 1.8	6.5 $\pm$ 1.7
	teleop	6.8 $\pm$ 1.0	<b>7.5 <math>\pm</math> 1.0</b>	5.5 $\pm$ 2.4
2 & 3 (training)	sc	6.1 $\pm$ 3.9	6.0 $\pm$ 1.6	5.9 $\pm$ 1.8
	sc_user	7.1 $\pm$ 3.8	7.4 $\pm$ 1.4	6.1 $\pm$ 1.4
	teleop	<b>7.6 <math>\pm</math> 3.6</b>	<b>7.5 <math>\pm</math> 1.4</b>	<b>5.8 <math>\pm</math> 2.4</b>
4 (transfer)	sc	6.0 $\pm$ 0.8	6.8 $\pm$ 1.0	<b>5.0 <math>\pm</math> 1.8</b>
	sc_user	<b>7.3 <math>\pm</math> 0.5</b>	<b>7.0 <math>\pm</math> 0.8</b>	5.8 $\pm$ 2.1
	teleop	6.5 $\pm$ 2.4	6.5 $\pm$ 2.4	7.0 $\pm$ 3.6

TABLE II  
SUBJECT RESPONSES TO OVERALL EXPERIMENT QUESTIONS (1 = “STRONGLY DISAGREE”, 5 = “STRONGLY AGREE”)

Question	Teleop	SC	SC-User
I found the spider chart <i>helpful</i>	-	3.5 $\pm$ 1.0	<b>4.3 <math>\pm</math> 1.0</b>
Most would learn this robot <i>quickly</i>	4.0 $\pm$ 0.8	3.8 $\pm$ 0.5	<b>4.5 <math>\pm</math> 0.6</b>
The last session was <i>more difficult</i>	3.3 $\pm$ 1.7	<b>2.5 <math>\pm</math> 0.6</b>	3.0 $\pm$ 1.2


### A. Subjective Results

Subject ratings to post-session questions on *skill* (“How would you rate your skill?”), *success* (“How successful were you?”), and *difficulty* (“How hard did you have to work?”) are shown in Table I. For the ‘baseline’ session, **sc** subjects reported better ratings than the other groups, especially in terms of difficulty, with a mean difference of 1.7 compared to **sc\_user**. There is a tendency toward more favorable responses for **teleop** and **sc\_user** in the ‘training’ sessions, with **teleop** receiving the best ratings across all categories. Nevertheless, **sc\_user** has the superior ‘transfer’ according to perceived skill (7.3 $\pm$ 0.5) and success (7.0 $\pm$ 0.8). The **teleop** users found the new wire testbed more difficult to complete (7.0  $\pm$  3.6), with a mean difference of 2.0 from **sc** and 1.2 from **sc\_user**. Despite these patterns, no main or interaction effects were revealed from mixed ANOVAs performed on each question category and so we cannot substantiate **H1**.


Survey questions on the entire experiment are summarized in Table II. Indeed, the last session is voted as marginally more difficult by **teleop** participants. Subjects in the **sc\_user** group also viewed operating the robot as quicker to learn than **sc** and even **teleop**. Interestingly, **sc\_user** participants considered the spider chart to be more helpful (4.3 $\pm$ 1.0) than the **sc** chart (3.5  $\pm$  1.0), which had autonomously computed values and could not be edited. All **sc\_user** individuals “found the VR interface easy to edit” (5.0, not in table).

### B. Quantitative Results

Table III summarizes the quantitative results of the study. In the first ‘baseline’ session, the **sc** group had the least number of collisions (1.74  $\pm$  2.6), whereas **teleop** had the fastest time-to-success (35.8s  $\pm$  9.8s). On the other hand, **sc\_user** subjects exhibited notable progress in ‘training’ sessions, as there was a major reduction in average number of collisions, yielding the lowest collision count (1.50  $\pm$  1.3). The **teleop** group continued to maintain the fastest success



(a) Time-to-success per session



(b) Time-to-success per trial

Fig. 6. Time-to-success demonstrated for participants across (a) sessions and (b) trials. Gray shaded area represents the session with a new test bed.


TABLE III  
QUANTITATIVE RESULTS ACROSS SESSIONS

Sessions	Mode	Collisions	Time (s)	No. Fails
1 (baseline)	sc	<b>1.74 <math>\pm</math> 2.6</b>	51.4 $\pm$ 24.8	1
	sc_user	2.74 $\pm$ 1.9	38.4 $\pm$ 8.5	1
	teleop	2.10 $\pm$ 1.8	<b>35.8 <math>\pm</math> 9.8</b>	0
2 & 3 (training)	sc	1.60 $\pm$ 2.0	53.1 $\pm$ 12.8	6
	sc_user	<b>1.50 <math>\pm</math> 1.3</b>	35.0 $\pm$ 8.5	0
	teleop	1.54 $\pm$ 1.4	<b>29.3 <math>\pm</math> 8.0</b>	1
4 (transfer)	sc	2.26 $\pm$ 1.9	43.2 $\pm$ 5.4	1
	sc_user	<b>1.45 <math>\pm</math> 1.1</b>	<b>25.5 <math>\pm</math> 6.3</b>	0
	teleop	1.70 $\pm$ 1.4	28.5 $\pm$ 10.3	0


times (29.3s  $\pm$  8.0s). Although in the ‘transfer’ session, **sc\_user** participants attained the swiftest completion times (25.5s  $\pm$  6.3s) and the fewest collisions (1.45  $\pm$  1.1).

Average time-to-success is presented per session and trial in Figs. 6(a) and 6(b), respectively. A mixed ANOVA run on completion times found no interaction effect between modes and sessions ( $F(2, 9) = 1.17, p = 0.35$ ). Though a main effect was found for both mode ( $p < 0.01$ ) and session ( $p < 0.05$ ). Post-hoc pairwise comparisons validated that the **sc** times were slower than those for both **teleop** ( $p < 0.01$ ) and **sc\_user** ( $p < 0.001$ ). A paired post-hoc test also suggests that completion times were faster in the final session ( $p < 0.05$ ), hinting at the ‘transfer’ task being a quicker game to navigate. These results only partially support **H2**, as **teleop** and **sc\_user** obtained statistically similar times ( $p = 0.81$ ). We still highlight the trend in **sc\_user** operators displaying faster task times compared to the **teleop** group when transitioning to the new testbed. Smaller standard deviations also bolster this trend, as shown in Fig. 6(b).

Comparing a subject’s teleoperation performance between



(a) Collisions vs baseline per session



(b) Control interface jerk vs baseline per session

Fig. 7. Differences in collision occurrences (a) and control interface jerk (b) relative to the subject’s average baseline performance in the first session.

their ‘baseline’ session and subsequent sessions is informative of their learning trajectory. Figs. 7(a) and 7(b) illustrate these trajectories for collision frequency and controller jerk, respectively, which we use to quantify operator “skill”. A mixed ANOVA for collisions found no interaction or main effects in modes and session. Yet a one-way ANOVA to observe if the control mode in the ‘transfer’ session influenced collisions relative to the subject’s ‘baseline’ revealed a significant effect ( $F(2, 56) = 5.47, p < 0.01$ ). Post-hoc comparisons showed that the only significant effect was **sc.user** incurring fewer relative collisions than **sc** ( $p < 0.01$ ). Similarly, a mixed ANOVA found no interactions or main differences on control interface jerk, but a one-way ANOVA found a significant effect of mode in the final session ( $F(2, 57) = 3.52, p < 0.05$ ). A post-hoc test indicated that **sc.user** issued less jerky commands than **teleop** ( $p < 0.05$ ). Overall, these findings offer partial support for **H3**.

For the analyses discussed above, we did not observe any significant effects of prior robotics experience when isolating control strategies. Even one-way ANOVAs on the ‘baseline’ session (where significance would be most likely) found no differences, *e.g.*, task times for subjects in **teleop** with robotics background were not faster than those without ( $p = 0.066$ ). Nonetheless, we cannot draw conclusions about these results given the small number of participants.

## VII. DISCUSSION

A few notable findings emerged from the buzz wire experiment. First, while the subjective and *absolute* results were comparable for **sc.user** and **teleop** during ‘training’ sessions, **sc.user** operators outperformed the other control modes in measures of *relative* improvement. For example, **sc.user** subjects had consistently less collisions and produced less jerky




Fig. 8. User-customized vs heuristics-based arbitration parameters at the conclusion of the experiment.

inputs than their ‘baseline’ session, which were statistically significant differences over **teleop** and **sc** in the ‘transfer’ condition. Indeed, **sc.user** obtained markedly better results across all subjective and quantitative (absolute and relative) metrics in the new game. Altogether, these results reinforce the benefits of offering users a means of customizing shared control, particularly when there are variations in the task.

Another observation is the inferior performance of the **sc** policy. Comparing the distributions of arbitration parameters between **sc.user** and **sc** reveals higher variance in the **sc.user** group’s final  $\theta$  selection, as depicted in Fig. 8. The invariance in **sc** is due to the performance heuristic  $\bar{r}_t$  from Eq. (16) being solely dependent on the distance between the end-effector and wire. An alternative update routine could prioritize factors that **sc.user** participants gravitated toward, such as *turnability* and *responsiveness*. However, any heuristics-based method for automatically adapting  $\theta$  would still converge to a similar parameter distribution across users, without accounting for individual preferences. Shared control policies learned from large demonstration datasets would also suffer from the lack of personalization to unseen users.

Despite our promising preliminary findings, there are areas for improvement. One important limitation is the broader applicability of our results due to the constrained task and environment. Future research should explore more complex telemanipulation tasks, *e.g.*, those involving humanoids [1], and include a larger, more diverse participant pool with varied backgrounds and skill levels in robot teleoperation. Another imperative requirement going forward will be to ensure that users have a clear understanding of the implications in visualization signals,  $\mathcal{V}$ , and interface actions,  $v_h$ . Regardless of our best efforts to make  $\mathcal{V}$  and  $v_h$  intuitive, there was one occasion where an **sc.user** subject misunderstood the meaning of a parameter, leading to a deterioration in performance. Simplifying the selection of arbitration factors may alleviate this issue, especially since certain factors, *e.g.*, *turnability*, were redundant as all **sc.user** subjects converged on a similar setting (see Fig. 8).

Finally, we suggest ways to extend our user-customizable formulation to make it more widely applicable. Instead of using various task-specific factors, like *speed*, a single scalar  $\alpha \in [0, 1]$  for “autonomous assistance” could be easily tuned in VR, *e.g.*, with a graphical slider. Modern frameworks for shared control, such as those based on reinforcement learning [33], could then let users calibrate  $\alpha$  *online*, rather than relying on a fixed or learned value that they cannot see or refine. Alternatively, latent actions learned from expert

demonstrations could be made user-adjustable [13], [34]. The main challenges here would be in communicating the latent dimensions through  $\mathcal{V}$  and exposing user interface actions to modify their values, as these latent actions may not represent a human-interpretable factor or be easy to edit via  $v_h$ . We defer investigation into these ideas for future work.

## VIII. CONCLUSIONS

In this paper, we introduced a mathematical shared control framework for users to customize the arbitration process. We presented an instantiation of this framework for a teleoperation task in SE(3) and conducted a longitudinal study spanning two weeks per subject to evaluate the proposed method. Our user-customizable shared control method enhanced teleoperation outcomes, reducing collision frequency and input jerk, while also accommodating the baseline skill levels of different users. Furthermore, operators transferred to generalizations in the teleoperation task better when they could edit the inner workings of the shared control. We believe these results hold promise for long-term adaptability to diverse variations in users and tasks. Future research will extend the scale of the user study, explore a wider array of tasks, and consider more flexible representations of the arbitration parameters for communication and user editing.

## REFERENCES

- [1] K. Darvish, L. Penco, J. Ramos, R. Cisneros, J. Pratt, E. Yoshida, S. Ivaldi, and D. Pucci, "Teleoperation of humanoid robots: A survey," *IEEE Transactions on Robotics*, 2023.
- [2] M. Selvaggio, M. Cognetti, S. Nikolaidis, S. Ivaldi, and B. Siciliano, "Autonomy in physical human-robot interaction: A brief survey," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 7989–7996, 2021.
- [3] R. Luo, C. Wang, C. Keil, D. Nguyen, H. Mayne, S. Alt, E. Schwarm, E. Mendoza, T. Padir, and J. P. Whitney, "Team Northeastern's Approach to ANA XPRIZE Avatar Final Testing: A Holistic Approach to Telepresence and Lessons Learned," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2023, pp. 7054–7060.
- [4] D. P. Losey, C. G. McDonald, E. Battaglia, and M. K. O'Malley, "A Review of Intent Detection, Arbitration, and Communication Aspects of Shared Control for Physical Human–Robot Interaction," *Applied Mechanics Reviews*, vol. 70, no. 1, 2018.
- [5] M. K. O'Malley, A. Gupta, M. Gen, and Y. Li, "Shared Control in Haptic Systems for Performance Enhancement and Training," *Journal of Dynamic Systems, Measurement, and Control*, vol. 128, no. 1, pp. 75–85, 2005.
- [6] M. Zolotas and Y. Demiris, "Towards Explainable Shared Control using Augmented Reality," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2019, pp. 3020–3026.
- [7] S. Musić and S. Hirche, "Control sharing in human-robot team interaction," *Annual Reviews in Control*, vol. 44, pp. 342–354, 2017.
- [8] B. D. Argall, "Autonomy in rehabilitation robotics: An intersection," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, no. 1, pp. 441–463, 2018.
- [9] G. Li, Q. Li, C. Yang, Y. Su, Z. Yuan, and X. Wu, "The classification and new trends of shared control strategies in telerobotic systems: A survey," *IEEE Transactions on Haptics*, vol. 16, no. 2, pp. 118–133, 2023.
- [10] D. Gopinath, S. Jain, and B. D. Argall, "Human-in-the-loop optimization of shared autonomy in assistive robotics," *IEEE Robotics and Automation Letters*, vol. 2, no. 1, pp. 247–254, 2017.
- [11] C. X. Miller, T. Gebrekristos, M. Young, E. Montague, and B. Argall, "An analysis of human-robot information streams to inform dynamic autonomy allocation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2021, pp. 1872–1878.
- [12] M. Hagenow, E. Senft, R. Radwin, M. Gleicher, B. Mutlu, and M. Zinn, "Corrective Shared Autonomy for Addressing Task Variability," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3720–3727, 2021.
- [13] Y. Cui, S. Karamcheti, R. Palleti, N. Shivakumar, P. Liang, and D. Sadigh, "No, to the right: Online language corrections for robotic manipulation via shared autonomy," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2023, pp. 93–101.
- [14] J. C. A. Read, S. F. Begum, A. McDonald, and J. Trowbridge, "The Binocular Advantage in Visuomotor Tasks Involving Tools," *i-Perception*, vol. 4, no. 2, pp. 101–110, 2013.
- [15] F. Budini, M. M. Lowery, M. Hutchinson, D. Bradley, L. Conroy, and G. De Vito, "Dexterity training improves manual precision in patients affected by essential tremor," *Archives of Physical Medicine and Rehabilitation*, vol. 95, no. 4, pp. 705–710, 2014.
- [16] H. Schreuder, C. van den Berg, E. Hazebroek, R. Verheijen, and M. Schijven, "Laparoscopic skills training using inexpensive box trainers: which exercises to choose when constructing a validated training course," *BJOG: An International Journal of Obstetrics & Gynaecology*, vol. 118, no. 13, pp. 1576–1584, 2011.
- [17] M. J. A. Zeestraten, I. Havoutis, and S. Calinon, "Programming by demonstration for shared control with an application in teleoperation," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1848–1855, 2018.
- [18] D. Nicolis, *A General Framework for Shared Control in Robot Teleoperation with Force and Visual Feedback*. Springer International Publishing, 2020, pp. 119–131.
- [19] K. Fitzsimons, A. Kalinowska, J. P. Dewald, and T. D. Murphey, "Task-based hybrid shared control for training through forceful interaction," *The International Journal of Robotics Research*, vol. 39, no. 9, pp. 1138–1154, 2020.
- [20] R. Hetrick, N. Amerson, B. Kim, E. Rosen, E. J. d. Visser, and E. Phillips, "Comparing Virtual Reality Interfaces for the Teleoperation of Robots," in *Systems and Information Engineering Design Symposium*, 2020, pp. 1–7.
- [21] M. Zolotas, M. Wonsick, P. Long, and T. Padir, "Motion Polytopes in Virtual Reality for Shared Control in Remote Manipulation Applications," *Frontiers in Robotics and AI*, vol. 8, p. 286, 2021.
- [22] E. Babaian, D. Yang, M. Karimi, X. Xu, S. Ayvasik, and E. Steinbach, "Skill-cpd: Real-time skill refinement for shared autonomy in manipulator teleoperation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2022, pp. 6189–6196.
- [23] Q. Li, W. Chen, and J. Wang, "Dynamic shared control for human–wheelchair cooperation," in *IEEE International Conference on Robotics and Automation*, 2011, pp. 4278–4283.
- [24] A. D. Dragan and S. S. Srinivasa, "A policy-blending formalism for shared control," *The International Journal of Robotics Research*, vol. 32, no. 7, pp. 790–805, 2013.
- [25] M. Zurek, A. Bobu, D. S. Brown, and A. D. Dragan, "Situational confidence assistance for lifelong shared autonomy," in *IEEE International Conference on Robotics and Automation*, 2021, pp. 2783–2789.
- [26] A. U. Pehlivan, D. P. Losey, and M. K. O'Malley, "Minimal assist-as-needed controller for upper limb robotic rehabilitation," *IEEE Transactions on Robotics*, vol. 32, no. 1, pp. 113–124, 2016.
- [27] S. Jain and B. Argall, "Probabilistic Human Intent Recognition for Shared Autonomy in Assistive Robotics," *ACM Transactions on Human-Robot Interaction*, vol. 9, no. 1, 2019.
- [28] C. Brooks and D. Szafir, "Visualization of intended assistance for acceptance of shared control," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020, pp. 11 425–11 430.
- [29] C. G. Christou, D. Michael-Grigoriou, and D. Sokratous, "Virtual Buzzwire: Assessment of a Prototype VR Game for Stroke Rehabilitation," in *IEEE Conference on Virtual Reality and 3D User Interfaces*, 2018, pp. 531–532.
- [30] M. Wonsick, T. Keleştemur, S. Alt, and T. Padir, "Telemanipulation via virtual reality interfaces with enhanced environment models," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2021, pp. 2999–3004.
- [31] F. Flacco, T. Kröger, A. De Luca, and O. Khatib, "A depth space approach to human-robot collision avoidance," in *IEEE international conference on robotics and automation*, 2012, pp. 338–345.
- [32] N. Hogan and D. Sternad, "Sensitivity of smoothness measures to movement duration, amplitude, and arrests," *Journal of Motor Behavior*, vol. 41, no. 6, pp. 529–534, 2009.
- [33] S. Reddy, A. Dragan, and S. Levine, "Shared autonomy via deep reinforcement learning," in *Proceedings of Robotics: Science and Systems*, 2018.
- [34] H. J. Jeon, D. Losey, and D. Sadigh, "Shared autonomy with learned latent actions," in *Robotics: Science and Systems*, 2020.