Bimanual Telemanipulation with Force and Haptic Feedback through an Anthropomorphic Avatar System

Christian Lenz*, Sven Behnke

Institute for Computer Science VI, Autonomous Intelligent Systems, University of Bonn, Friedrich-Hirzebruch-Allee 8, 53115 Bonn, Germany

Abstract

Robotic teleoperation is a key technology for a wide variety of applications. It allows sending robots instead of humans in remote, possibly dangerous locations while still using the human brain with its enormous knowledge and creativity, especially for solving unexpected problems. A main challenge in teleoperation consists of providing enough feedback to the human operator for situation awareness and thus create full immersion, as well as offering the operator suitable control interfaces to achieve efficient and robust task fulfillment. We present a bimanual telemanipulation system consisting of an anthropomorphic avatar robot and an operator station providing force and haptic feedback to the human operator. The avatar arms are controlled in Cartesian space with a direct mapping of the operator movements. The measured forces and torques on the avatar side are haptically displayed to the operator. We developed a predictive avatar model for limit avoidance which runs on the operator side, ensuring low latency. The system was successfully evaluated during the ANA Avatar XPRIZE competition semifinals. In addition, we performed in lab experiments and carried out a small user study with mostly untrained operators.

Keywords: force feedback control, teleoperation, dual arm manipulation, human robot interaction, real-time control, haptic interface

1. Introduction

Teleoperation is a very powerful method to control robots. It enables humans to explore remote locations and to interact there with objects and persons without being physically present. Although state-of-the-art methods for autonomous control are improving rapidly, the experience and instincts of humans, especially for solving unpredictable problems is unparalleled so far. The current COVID-19 pandemic is a great example of scenarios where remote work is highly desirable. Further possible applications for teleoperation include disaster response where humans can operate remotely and use critical situationsaving skills without risking their lives as well as maintenance and healthcare to allow experts operating in remote locations for manipulation tasks without the need of travel. Robotic teleoperation is a popular research area which is advanced by multiple robotic competitions like the DARPA Robotics Challenge [1] and RoboCup Rescue [2]. These events are a great opportunity to benchmark and evaluate different highly integrated and complex systems in standardized test scenarios under comparable conditions. Our team NimbRo participates in the ANA

Email addresses: lenz@ais.uni-bonn.de (Christian Lenz), behnke@cs.uni-bonn.de (Sven Behnke)

Avatar XPRIZE Competition¹ with the goal to advance the state of the art of such robotic telemanipulation and telepresence systems.

In addition to immersive visualization of the remote location, one important aspect is telemanipulation which enables the operator to physically interact with the remote environment. This capability is critical for many applications—without it, we are constrained to mere telepresence.

In this work, we present a humanoid bimanual telemanipulation system built from off-the-shelf components, which allows a human operator to interact and manipulate in remote locations (see Fig. 1). Our contributions include:

- 1. Integrating a bimanual robotic avatar and an upperbody operator exoskeleton for Cartesian telemanipulation,
- 2. an arm and hand controller with force and haptic feedback,
- 3. a model-based arm movement prediction to haptically display position and velocity limitations of the remote avatar in real time,
- 4. an oscillation observer module to detect and suppress oscillations introduced in the force-feedback control loop, and

 $^{^{*}}$ Corresponding author

¹https://www.xprize.org/prizes/avatar



Figure 1: Our bimanual haptic telemanipulation system: Human operator in operator station (right) inspecting an object during the ANA Avatar XPRIZE Competition semifinals through a remote anthropomorphic avatar robot (left).

5. subsystem evaluation in lab experiments, a user study, as well as our participation at the ANA Avatar XPRIZE competition semifinals.

2. Related Work

Teleoperation is a widely investigated research area. A leading device (in our context called the Operator Station, see Section 3), often with haptic feedback is used to control a following device (Avatar Robot) in a remote location. The DARPA Robotics Challenge (DRC) 2015 [1] required the development of mobile telemanipulation systems. Several research groups, such as DRC-HUBO [3], CHIMP [4], RoboSimian [5], and our own entry Momaro [6] presented teleoperation systems with impressive manipulation capabilities. The focus was on completing as many manipulation and locomotion tasks as possible using a team of trained operators. Thus, some hardware and software components were highly specialized towards solving predefined tasks. In addition, the robots were not required to communicate or interact with other humans in the remote location and thus did not feature respective capabilities. In contrast, our developed avatar solution was designed for interaction with humans in the remote location and the operator interface is designed to give intuitive control over the robot to a single, possibly untrained operator. Presence of the operator in the remote location is prioritized over specific task solution skills.

Passivity control constitutes a large research field in the context of teleoperation. Uncertainties of the operator's input dynamics, as well as the remote environment are factors which introduce potential instability in control loops. Many different passivity control methods tackle the stability problem, e.g. [7, 8, 9, 10, 11, 12]. These control schemes use the concept of passivity which is a sufficient condition to obtain a stable control system. Control systems are considered passive if and only if the energy flowing into the system exceeds the energy flowing out at any time. Conveniently, if all subsystems are passive the entire system is guaranteed to be passive as well.

Ryu and Preusche [11] use a time-domain passivity control approach to ensure stable teleoperation, handling time delay of up to 120 ms. A passivity observer is used to monitor the energy transferred from the operator to the avatar system and vice versa. The passivity controller actively dampens the system to ensure passivity and thus stability of the system. Our approach uses a similar observer and dampening approach to ensure stable teleoperation control loops. One drawback of energy-based time-domain passivity controllers is the occurrence of position drift, which is handled by [13, 14] and improved in a 1 DOF (degree of freedom) teleoperation setup by Coelho et al. [15]. The design of our control architecture assures position drift-free teleoperation by commanding goal poses for each avatar hand. Small position derivations can occur due to motion execution. However, these are negligible small for our application (see Section 5.3.2). Overall, passivity control can suffer from distortion of the displayed environment [16].

In presence of large communication time delays (in the order of seconds) between operator and operator station (e.g., in space teleoperation missions), predictive control methods can improve task performance [17]. The remote robot tries to anticipate human control commands to complete partially transmitted instructions. Hauser [18] presents a prediction and planning method to assist teleoperation without assuming a finite set of candidate tasks. Nevertheless, this method is limited to a finite set of task types and needs a large number of training data to overcome this limitation. Our control approach is designed to limit the operator only within the underlying hardware constrains, without constraining the task solution itself.

Some recent approaches use teleoperation interfaces which only send commands to the robot without providing any force or haptic feedback to the operator [19, 20]. The advantage of such systems is clearly the low weight of the capture devices which hinder the operator only marginally. The downside is missing force or haptic feedback, especially for tasks that cannot be solved with visual feedback alone, such as difficult peg-in-hole tasks.

Other methods use custom-developed operator interfaces including force and haptic feedback. Klamt et al. [21] developed a centaur-like platform for teleoperation in disaster response scenarios. The proposed teleoperation interfaces uses actuators located at the base of the device and metallic tendons for torque transmission to the actuated joints. This approach benefits from light-weight moving parts with low inertia resulting in an easily backdrivable system. However, the used metallic tendons introduce some compliant behavior which need to be considered. In our approach, we use off-the-shelf robotic arms with actuators located inside the joints. We utilize an external FT-Sensor to actively follow the operator's arm movement. Guanyang et al. [22] use two haptic devices (3DoF rotational and 3DoF translational device) to control a robot's end-effector. Both devices can display contact forces with high stiffness due to mechanical design. However, limited workspace and thus the requirement of a motion mapping between operator station and controlled robot as well as using two devices to control full 6D motion are some drawbacks compared to our approach. In Abi-Farrajl et al. [23] the operator station is comparable to our approach, but the focus there is placed on haptic feedback for balance control of the bipedal humanoid avatar robot.

Wearable haptic feedback devices [24] overcome the workspace constraints generated by stationary devices but are limited to displaying contact since they cannot create any force towards the operator. Other research projects focus on controlling a teleoperation system under time delays [25] or with two different kinematic chains on the avatar side [26].

In contrast to the highlighted related research, our approach focuses on off-the-shelf components which allow for easy replication and maintenance. Furthermore, the used robotic arms are replaceable with any other appropriate actuators with different kinematic chains, since the whole communication between the systems uses only the 6D end-effector pose.

3. Hardware Setup

The developed robotic teleoperation system consists of an operator station and an avatar robot, as shown in Fig. 1. The operator station allows the operator to control the avatar from a remote location. It includes two robotic arms, hand exoskeleton, 3D Rudder foot paddle, and a head mounted display with additional sensors. The avatar robot is designed to interact with humans and in



Figure 2: Operator (left) and avatar (right) arm with used hardware components. For simplicity, only the right arm is shown. The axes depict the common hand frame which is used for control commands and feedback.

human-made indoor environments and, thus, features an anthropomorphic upper body mounted on a mobile base.

The operator station and the avatar robot are controlled with a standard desktop computer (Intel i9-9900K @ 3.60 GHz, NVidia RTX 2080) each. The communication between these computers is achieved by a single Gigabit Ethernet connection. We successfully tested the system with artificial delay of up to 30 ms in both directions. Thus, our system allows operating the avatar from a distant location. On the software side, the Robot Operating System (ROS) framework is used. Both, the operator station and the avatar robot run their own roscore. We use NimbRo Network² for any communication between both roscores. The hardware design of the operator station and avatar robot is described in the following.

3.1. Avatar Robot

The avatar robot's anthropomorphic upper body mimics the human arm configuration using two 7 DoF Franka Emika Panda arms, mounted in slightly V-shaped angle. The shoulder height of 110 cm above the floor allows convenient manipulation of objects on a table, as well interaction with both sitting and standing persons. The shoulder width of under 90 cm enables easy navigation through standard doors.

The Panda arms have a sufficient payload of 3 kg and a maximal reach of 855 mm. The extra degree of freedom gives some flexibility in the elbow position. While the arm measures joint torques in each arm joint, we mounted additional OnRobot HEX-E 6-Axis force/torque sensors at the wrists for more accurate force and torque measurements close to the robotic hands, since this is the default location of contact with the robot's environment (see Fig. 2). The avatar robot is equipped with two anthropomorphic hands. A 20 DoF Schunk SVH hand is mounted on the right side. The nine actuated DoF provide very dexterous manipulation capabilities. The left arm features a 5 DoF

²https://github.com/AIS-Bonn/nimbro_network



Figure 3: Control System overview. For simplicity, only the right side is depicted. The left side is controlled similarly, besides a different Hand Controller for the different Schunk hands.

Schunk SIH hand for simpler but more force-requiring manipulation tasks. Both hand types thus complement each other.

The avatar's head is equipped with two RGB cameras, a microphone, and a small screen displaying the animated face of the operator [27]. It is attached to the upper body using a 6 DoF UFactory xArm for free head movement. In addition, two wide-angle RGB cameras are capturing the robot's vicinity for situation awareness during locomotion. Further details on the VR remote visualization system are provided in [28]. The anthropomorphic upper body has been mounted on a movable base, which allows omnidirectional movement.

3.2. Operator Station

The operator controls the avatar through the Operator Station from a comfortable sitting pose. The human hand movement is captured with a similar setup as already described for the avatar robot: Two Panda arms are equipped with an OnRobot HEX-E force/torque sensor and connected to the operator hand using a SenseGlove haptic interaction device. The Panda arms thus serve dual purposes: They provide precise 6D human hand pose measurements for avatar control, as well as the possibility to induce force feedback measured by the Avatar onto the human wrists. The operator-side force/torque sensor is used to measure the slightest operator hand movement to assist the operator in moving their arm, reducing the felt mass and friction to a minimum.

The SenseGlove haptic interaction device features 20 DoF finger joint position measurements (four per finger) and a 1 DoF haptic feedback channel per finger (i.e., when activated the human feels resistance, which prevents further finger closing movement).

For visual and audio communication, the operator is wearing a head mounted display equipped with eye trackers, audio headset, and a camera viewing the lower face part (for more details see [29] and [30]). The avatar locomotion can be controlled using a 3D Rudder foot paddle device.

The Panda arms feature built-in safety measures and will stop immediately if force, torque, or velocity limits are exceeded. This ensures safe human-robot interactions both on the operator and the avatar side.

4. Force Feedback Controller

The control architecture for the force feedback teleoperation system consists of two arm and two hand controllers (one for each the operator and the avatar side). For the right and the left arm, each controller pair is running separately. The hand controller for the right and left hand are slightly different since different robotic hands are used. An overview of the control architecture is shown in Fig. 3. The arm controllers run with an update rate of 1 kHz and the force-torque sensor measurements are captured with 500 Hz. The force-torque measurements are smoothed using a sensor-sided low-pass filter with a cut-of frequency of 15 Hz.

Since the robot arms are attached from outside to the operator's wrists (see Fig. 3), the kinematic chains of avatar and operator station differ, and thus, a joint-by-joint mapping of the operator and avatar arm is not possible. Consequently, the developed control concept does not rely on similar kinematic chains. Instead, a common control frame is defined in the middle of the palm of both the human and robotic hands, i.e., all necessary command and feedback data are transformed such that they refer to this frame before being transmitted. The controllers for the operator and avatar arms and both hands are described in the following subsections.

4.1. Operator Arm Controller

The operator arm controller commands joint torques to the Panda arm and reads the current operator hand



Figure 4: Unintended lower arm contact: Typical situation in which the avatar's lower arm establishes contact with the environment, which cannot be measured by the force-torque sensor. Panda arm torques are used to provide the operator with appropriate force feedback.



Figure 5: Operator arm torque command (see Eq. (1)) is scaled using α to reduce oscillations when getting close to joint position or velocity limits. The scalar decreases linearly if the distance to a joint position limit d_p exceeds the threshold t_p . Velocity limits are handled analogously.

pose to generate the commanded hand pose sent to the avatar robot. The goal is to generate a weightless feeling for the operator while moving the arm—if no force feedback is displayed. Even though the Panda arm has a convenient teach-mode using the internal gravity compensation when zero torques are commanded, the weightless feeling can be further improved by using precise external force-torque measurements. Any contact established by the avatar robot with its hands and lower arms is haptically displayed to the operator. Since the teleoperation system has no information about the operator's intention, contact should not be avoided but displayed to the operator to keep the human in control of the situation.

For simplicity, just one arm is mentioned in the following, since the left and right arms are controlled equally.

4.1.1. Torque Controller

Let us denote with $\tau_o \in \mathbb{R}^7$ the commanded joint torques for a particular time step. Then

$$\tau_o = \alpha \tau_{cmd} + \beta \tau_f + \tau_{lo} + \tau_{la} + \tau_{no} + \tau_{co} \tag{1}$$

describes the used torque components (command, force feedback, operator limit avoidance, avatar limit avoidance,

null-space, and Coriolis) which will be explained in the following. Note that the gravity compensation is not considered here, since it is done by the Franka Control Interface (FCI) itself.

The commanded joint torques τ_{cmd} to move the Panda arm based on the force/torque sensor measurements and are defined as

$$\tau_{cmd} = J^T F \tag{2}$$

with J being the body Jacobian relative to the hand frame and $F \in \mathbb{R}^6$ denoting the measured 3D forces and 3D torques. The scalars $\alpha \in \mathbb{R}^7$ are computed by the predictive limit avoidance module (see Section 4.1.2). Note that F has to be corrected taking sensor bias and attached endeffector weight into account, as well as transformed into the common hand control frame (see Section 4.3).

The term τ_f denotes the force feedback induced by the avatar-side force-torque sensor and Panda arm torque measurements. The scalar β is computed by the oscillation observer module (see Section 4.1.3) to prevent possible oscillations in the feedback loop. The force-torque measurements are used as the primary feedback source. They are already bias-corrected and correctly transformed, therefore Eq. (2) can be directly applied analogously to compute the induced joint torques τ_{sensor} .

In some situations, especially when performing manipulation tasks on a table, the avatar establishes contact between the lower arm and the environment (for example the table, see Fig. 4), which are not visible from the operator's view pose. Since the force-torque sensor and hand are above the table, this type of contact cannot be measured by the force-torque sensor. Here, we must use the joint torque measurements of the Panda arm to give the operator feedback about the contact. The Franka API provides estimated Cartesian forces f_{panda} at the end-effector. We use

$$f_{diff} = f_{panda} - f_{sensor} \tag{3}$$

to calculate the forces in the end-effector frame which cannot be measured using the force-torque sensor, where \hat{f}_{panda} and \hat{f}_{sensor} are the respective low-pass filtered end-effector forces. Finally, we transform the calculated forces into the wrist frame and compute

$$\tau_f = \tau_{sensor} + \tau_{diff},\tag{4}$$

where $\tau_{diff} \in \mathbb{R}^6$ is $f_{diff} \in \mathbb{R}^3$ extended with $t_0 \in \mathbb{R}^3$, since we ignore torque measurements here.

4.1.2. Predictive Limit Avoidance

Humans can achieve high speeds moving their arm, which can exceed the Panda joint velocity limits (up to $150^{\circ}/s$). To prevent the operator from exceeding joint position or velocity limits of the Panda arm, the term $\tau_{lo} \in \mathbb{R}^7$ is introduced to apply torques pushing the arm away from those limits. For a single joint *i*, the torque to avoid its position limit is defined as



Figure 6: Oscillation observer: Frequency analysis (left) of force feedback measurements from the avatar robot's right arm while placing a vase on a table (right). The initial contact between vase and table is marked. The graph shows the amplitude of the 4th frequency (ca. 5.7 Hz) computed using a sliding DFT for each axis over time.



Figure 7: Arm workspace evaluation. Left: Initial arm setup, similar to the avatar side. Right: Optimized mounting pose. Turquoise (reachable) and red (not reachable) arrows depict the captured human left-hand poses. The coordinate axes depict the operator sitting pose.

$$\tau_{lo-position}^{i} = \begin{cases} \gamma_{p}(\frac{1}{d_{p}^{i}} - \frac{1}{t_{p}}), & d_{p}^{i} < t_{p} \\ 0, & \text{else} \end{cases}$$
(5)

with γ_p being a constant scalar, d_p^i being the distance for joint *i* to its closer position limit, and $t_p = 10^{\circ}$ being a threshold how close a joint must be at a limit to activate this behavior. $\tau_{lo-velocity}^i$ is calculated analogously with $t_v = 40^{\circ}/sec$. Together, τ_{lo} is defined as

$$\tau_{lo} = \tau_{lo-position} + \tau_{lo-velocity}.$$
 (6)

The torques τ_{lo} exhibit hyperbolical growth when getting closer to respective limits. Since the operator-side forcetorque sensor will measure the generated limit avoidance torques, the arm can end up oscillating, especially being close to one or multiple position limits. Thus, the torques τ_{cmd} , which are influenced by the force-torque sensor, are scaled per joint by α , which is defined as

$$\alpha = \max(0, \min(1, 2\min(\frac{d_p}{t_p}, \frac{d_v}{t_v}) - 1)).$$
(7)

The scalar α is designed to decrease linearly and reach zero when the limit is approached halfway after activating the limit avoidance (see Fig. 5). This reduces the commanded torques τ_{cmd} enough when approaching a position or velocity limit and prevents the oscillation.

As already mentioned, the operator station and avatar robot have different kinematic arm chains. Therefore, avoiding position and velocity limits on the operator side does not guarantee limit avoidance on the avatar side. Calculating the joint torques preventing joint limits on the avatar robot in a similar way is not beneficial, since the feedback information would arrive with high latency (mainly because of the delay generated by motion execution on the avatar side). To overcome this issue, we use a model of the avatar inside the operator arm controller to predict the avatar arm movement for the next time step and calculate the needed joint torques to prevent joint limit violations in advance.

The current operator hand pose is used as the desired goal pose for the avatar arm in the common gripper frame. We estimate the avatar arm joint configuration reaching this goal pose using inverse kinematics (IK). The latest received avatar arm joint positions are used to initialize the IK solver. The current joint velocities are approximated by a low-pass filtered version of the joint position first derivatives. Having estimated the joint positions and velocities, we can apply the same avoidance strategy as described above (see Eq. (6)). Finally, the resulting joint torques can be transformed into the common 6D hand frame using the pseudoinverse of the Jacobian transpose $(J_A^T)^+$ and back to joint torques for the operator arm with J_O^T . This results in

$$\tau_{la} = J_O^T (J_A^T)^+ \tau_{la-model}.$$
(8)

The remaining two torque components from Eq. (1) are τ_{no} and τ_{co} . The term τ_{no} is a null-space optimization term which pulls the elbow towards a defined convenient pose in the null-space of the Jacobian. The result is a more human-like elbow pose to maximize the oper-

ator workspace by pushing the arm away from singularities. The last torque component is the Coriolis term τ_{co} obtained by the Panda model.

4.1.3. Oscillation Observer

Our telemanipulation control system as described above is designed as simple as possible, but yet powerful to give the operator control over the robot along with necessary force feedback. The downside is, that our controller cannot promise any stability, resulting in oscillations when certain contact forces are measured. We address this issue by observing the force feedback channel, detecting any critical oscillations, and reducing the feedback introduced into the system.

The sliding Discrete Fourier-Transformation (DFT) [31, 32] method is used to analyze the force feedback channel per axis. Since we did not measure isolated oscillations in the torque channels, analyzing the 3D forces is sufficient. The force measurements are sampled with around 1 kHz. We experimentally investigated the oscillation frequency to obtain suitable DFT parameters and the observed frequency band. In the experiments, we provoked vibrations in the system by placing a vase harshly on a table. The used sliding DFT has resolution of 512 and uses the Hanning window to minimize spectral leakage. We observe correlation between the generated oscillation and the amplitude of the 4th frequency (ca. 5.7 Hz), is depicted in Fig. 6.

In each time step, the force measurements of all three axes are analyzed using separate sliding DFTs. Next, the combined result is computed using the Euclidean norm over all axes amplitudes. The resulting value v is clamped using experimentally obtained parameters (min = 163 and max = 500) and scaled to be within the interval [0, 1]. Finally, the scalar $\beta = 1 - v$ is used to scale the force feedback provided to the operator (see Section 4.1.1). The change rate of β is limited to gradually remove and restore the force feedback, s.t. the observer needs 0.8 s to fully remove the feedback in case of detected oscillations and 1.7 s to reach the normal force feedback control status again. This gentle oscillation elimination behavior is as little noticeable for the operator as possible. Note that Fig. 6 shows the oscillation without feedback reduction. The computation of β and the effect oscillation damping are evaluated in Section 5.1.

The avatar arm controller commands the Panda arm of the avatar robot to follow the commanded 6D pose by sending joint torques $\tau_a \in \mathbb{R}^7$ to the Franka Control Interface. The commanded torque is defined as

$$\tau_a = \tau_{cmd} + \tau_{init} + \tau_{na} + \tau_{ca}, \tag{9}$$

where $\tau_{cmd} \in \mathbb{R}^7$ and $\tau_{init} \in \mathbb{R}^7$ are calculated to reach the goal pose during operation and initialization, respectively (see below). The components τ_{na} and τ_{ca} are the null-space optimization and Coriolis terms similar to τ_{no} and τ_{co} , as described in Section 4.1. The convenient elbow poses used for the null-space optimization is defined such that the elbows are slightly stretched out. This generates a more human-like arm configuration and keeps the arm away from singularities in the elbow joints. Other singularities do not occur due to the nature of the arm kinematic in the usable workspace and other joint position limits. Therefore, no special singularity handling is needed here.

The goal torques τ_{cmd} and τ_{init} are generated using a Cartesian impedance controller that emulates a springdamper system. Its equilibrium point is the 6D goal pose commanded by the operator station in the common hand frame:

$$\tau_{cmd} = J^T (-S\Delta p - D(J\dot{q})), \tag{10}$$

where J denotes the zero Jacobian, $S \in \mathbb{R}^{6\times 6}$ and $D \in \mathbb{R}^{6\times 6}$ denote the stiffness and damping matrix, $\Delta p \in \mathbb{R}^{6}$ is the error in translation and rotation between the current and goal end-effector pose, and $\dot{q} \in \mathbb{R}^{7}$ denotes the current joint velocities. τ_{init} is only used for a safe initialization procedure (see below) and generated similarly using the current end-effector pose. The stiffness and damping parameters symbols and their values are empirically tuned to achieve some compliance while still closely following the operator command.

When no goal pose command is received, the controller keeps commanding the current arm pose to remain in a safe state. This happens when the operator station is not active or if a communication breakdown occurs (no operator command within the last 100 ms). After receiving a command, the controller performs an initialization procedure which fades linearly between the current and the new received goal pose. This prevents the robot from generating high torques to suddenly reach the new, possibly distant pose. This initialization process takes about 3 s.

The Panda arm stops immediately when excessive forces are measured, for example when there is unintended contact that exceeds force/torque thresholds. This feature is necessary to operate in a safe way. After notification of the human operator, the avatar arm controller can restart the arm automatically. After performing the initialization procedure, normal teleoperation can be resumed.

4.2. Hand Control

The operator finger movements are captured using two SenseGlove haptic interaction devices. Four separate finger joint measurements are provided per finger. Since the Schunk SVH and SIH robotic hands on the avatar have nine and five actuated joints, respectively, only the corresponding joint measurements are selected and linearly mapped to the avatar hands. While this mapping does not precisely replicate hand postures – this is impossible anyways due to the different kinematic structure – it gives the operator full control over all hand DoFs.

Both hands provide feedback in the form of motor currents, which is used to provide per-finger haptic feedback to the operator. The SenseGlove brake system is switched on or off depending on a pre-defined current threshold.



Figure 8: Predictive avatar model: Measured joint position for the first joint of the right avatar arm during a grasping motion (green) and predicted joint position for predictive limit avoidance (blue). Both measurements are captured on the operator side. Communication between both systems and motion execution generate a delay of up to 200 ms (Δt), which is compensated by the predictive model.

4.3. Force-Torque Sensor Calibration

Different end-effectors (SenseGloves, Schunk SIH, Schunk SVH hand, and corresponding 3D printed mounting adapters) are mounted on each of the four involved force-torque sensors. In addition, sensor bias results in barely usable raw sensor data. Thus, each sensor is calibrated separately to compensate these effects. To this end, 20 data samples from different sensor poses are collected. Each sample includes the gravity vector in the sensor frame and the mean of 100 sensor measurements from a static pose. A standard least squares solver [33] is used to estimate the force-torque sensor parameters, i.e., the force and torque bias and the mass and center of mass of all attached components. The same parameters including the additional mass and center of mass transformation resulting by the force-torque sensor itself is used to configure the built-in gravity compensation of the Panda arms. The calibration is performed once after hardware changes at the end-effectors or if the bias drift is too large. This method does not compensate for bias drift during usage, but is sufficient for our application.

5. Evaluation

In addition to our participation at the ANA Avatar XPRIZE Competition semifinals, we performed multiple experiments along with a small user study to evaluate the developed teleoperation system in our lab environment.

5.1. Quantitative Experiments

In a first experiment, we evaluated the operator arm workspace. 2,959 different 6D left hand poses were captured from a sitting person performing typical arm motions with a VR tracker on their wrist. In addition to hand poses with a fully extended arm, most of the poses are directly in front of the person, likely to be performed during manipulation tasks. First, the initial arm mounting pose (motivated by the avatar configuration) of the operator arm was evaluated. Each captured hand pose was marked as reachable if an inverse kinematic solution for the



Figure 9: Operator arm movement: Force in z-direction (in the direction of the human palm) needed to move the arm in the same repetitive motion with our operator arm controller running (top) and using only the Panda built-in gravity compensation (bottom).

Table 1: Operator arm workspace analysis

Mounting Pose	Reached	Missed	Reached $[\%]$
Initial	1,795	$1,\!164$	60.6%
Optimized	$2,\!848$	111	96.3%

arm was found. In a second step, different arm mounting poses were sampled to find an optimal pose, maximizing the number of reachable hand poses (see Fig. 7). Table 1 reports quantitative results. The resulting arm mounting pose drastically increases the overlap (from 60.6% to 96.3%) between the human operator's and the avatar's arm workspace, but requires a more complicated mounting setup.

In a second experiment, we evaluated the predictive limit avoidance module. Avatar arm position and velocity limits are haptically displayed via joint forces to the operator. Since measured joint positions and velocities are afflicted with latency generated by network communication (<1 ms) and motion execution using the Cartesian impedance controller, which can reach up to 200 ms (see Section 4.1.1), the operator control predicts the avatar arm joint configuration. Fig. 8 shows the measured and predicted joint position of the first right arm joint. The prediction compensates the delays, which allows for instantaneous feedback of the avatar arm limits to the operator.

In a third experiment, we investigated the forces and torques required to move the operator station arm, since this directly affects operator fatigue. We measured the forces and torques applied to the arm by reading the force-torque sensor measurements. The arm was moved in a comparable manner once with only the Panda gravity compensation enabled (i.e., $\tau_{cmd} = 0$ see Eq. (1)) and a second



Figure 10: Oscillation observer module. Placing a vase onto a table results in the depicted amplitude response for 5.7 Hz for each Cartesian axis ((a) and (b)). We performed the experiment twice with comparable executions. First, the oscillation observer was active ((a) and (c)). In the second execution, the observer was inactive ((b) and (d)). The corresponding joint torques commanded to the operator Panda arm are plotted exemplary for the 7th joint (c) and (d). The force feedback is scaled with β (c) which is computed from the oscillation observer. Note that β shown in (d) was computed but not used during execution.

time with our arm force controller running. In Fig. 9, the forces in the direction of one exemplary axis are shown. The results demonstrate the advantage of using an external force-torque sensor to generate a more unencumbered feeling for the operator while using the system.

In the last experiment, we analyzed the oscillation observer module (see Section 4.1.3), which observes the avatar force feedback channel to detect and prevent oscillations in the closed-loop controller. We used the avatar system to place a metal vase harshly on a table (see Fig. 6). The executed joint torques on the operator side along with the frequency responses of the force-feedback measurements are shown in Fig. 10. Although the oscillations are not eliminated immediately (c), reducing the feedback gain yields the expected oscillation suppression. The oscillation is detected by analyzing the frequency response shown in (a) and (b). The gain β is reduced with the maximum allowed rate to a value of 0.1 within 0.8 s. Next, the operator does not feel the contact forces as strong, the oscillation stops (shown in the exemplary torque command for the 7th operator arm joint (c)), and the gain is increased again, which brings back the force feedback for the operator. In conclusion, the oscillation observer module is a necessary part of our simple force-feedback control loop to ensure safe operation.

5.2. User Study

Our goal was to create an immersive and intuitive feeling for operation at remote locations using our system. Since humans have their very own preferences and subjective feelings of how good or intuitive certain control mechanisms perform, we carried out a user study with untrained operators, comparing different telemanipulation approaches. Due to the COVID-19 pandemic, we were limited to immediate colleagues as subjects, which severely constrained the scope of our study. Although all participants had a rough idea of the system, they controlled it for the first time during this study.

A total of five participants were asked to perform a bimanual peg-in-hole manipulation task. First, two different objects had to be grasped: a small aluminum bar and a 3D printed part with a hole. Afterwards, the bar should be inserted into the hole (see Fig. 11). The avatar robot was already placed in front of a table and both objects were within the avatar's workspace. Participants controlled the robot using the operator station located within the same room. Only the HMD with the avatar's perspective was used for visual feedback. The task was challenging due to very little friction between the finger and objects and tight tolerances, which required precise insertion alignment.

Each participant performed the task three times with the following control modes:

- 1. Operator station with force feedback enabled,
- 2. Operator station with force feedback disabled, and
- 3. VR controllers.



Figure 11: User study with untrained operator: Both objects had to be grasped and the bar had to be inserted into the hole.

Table 2: User study success rates and timings.

Telemanipulation mode	Success	Completion time [s]		
		Mean	StdDev	
1) Exoskeleton with feedback	4/5	119.0	117.1	
2) Exoskeleton w/o feedback	5/5	123.0	88.4	
3) VR controllers	3/5	126.3	25.8	

In the first control mode, all system components were active as described in this article. Any force and haptic feedback were disabled for the second control mode, i.e. the operator had to rely on visual feedback only. In the third control mode two HTC Vive VR controllers were used as input devices. As long as the trigger button was pressed, the corresponding avatar arm followed the controller movement. A different button was programmed to toggle between a defined closed and open hand pose.

A maximum of 5 min were granted to solve the task before it was marked as a failure. An object dropped outside the reachable workspace resulted in a failed trial. Objects dropped onto the table within the workspace of the avatar could be grasped again with no penalty. The participants were allowed to test each control mode about 1 min before starting the measured test.

Table 2 reports the quantitative results of the user study. The time needed to successfully solve the task is quite similar over the different telemanipulation modes. From these experiments, we realized that the completion time was highly influenced by external factors, such as losing the object due to not enough finger friction or different grasping and object handling solutions, which are unrelated to the used operator interface. In addition, humans can easily compensate missing force feedback using visual feedback. All three unsuccessful trials failed due to reaching the maximum experiment time of 5 min. To generate a more meaningful statement based on performance scores, more test samples are needed.

In addition to these quantitative measurements, we asked the participants to answer a short questionnaire about each telemanipulation mode with answers from the 1-7 Likert scale (see Fig. 12). The results show that especially the feeling of handling the objects and intuitive finger control was subjectively much better using the Operator Station. Enabling the force and haptic feedback gives the highest advantage when picking up the objects

Table 3: Avatar Arm safety stops					
	Scenario 1	Scenario 2	Scenario 3		
Day 1	Exceeded torque limits	Exceeded vel. limits	None		
Day 2	None	None	Software failure		

from the table. This can be explained by the additional feedback indicating contact between the hand and the table which cannot be perceived visually due to occlusions. All participants reported to feel safe and comfortable using the system. Although the experiment time was limited, this suggests non-excessive cognitive load on the operator.

Overall, the user study showed that our developed system is intuitive to use for untrained operators. Even though the force and haptic feedback did not increase the success rate of solving the task, it increases the immersive feeling as shown by the questionnaire.

5.3. ANA Avatar XPRIZE Competition

Our avatar system was evaluated by independent judges during the ANA Avatar XPRIZE competition semifinals over two days³. At each competition day, a first judge acted as the operator who performed 18 tasks in three predefined scenarios together with the second judge acting as the so called "recipient". The recipient was sitting with the avatar robot at a table in a different room about 100 m away from the operator control room. Communication of any kind between both judges was only possible through the avatar system. Both judges were trained to get comfortable with our system during the first hour of a trial. In the second hour, the judges had to solve the tasks without any instructions or support from our team. Thus, the judges were no experts but slightly more trained compared to the completely untrained operators in our user study (see Section 5.2). Fig. 13 shows the three scenarios: Solving a jigsaw puzzle, celebrating a business deal, and exploring an artifact from a historical exhibition with the robot's senses. The enabled force feedback (Control Mode 1, see Section 5.2) enabled the operator to feel some texture of the artifact. The judges evaluated the avatar system with a major focus on the ability to convey human senses, actions, and presence in the remote location in real time. At both competition days, the same three scenarios were tested by different judges. The better score per scenario (max. 30 points per scenario) counted towards the final score. Additional 10 points were given based on a video submitted prior to the semifinal showing the system in action in our lab⁴. Our team NimbRo archived an almost perfect score of 99 out of 100 points, which placed us first in the semifinal.

³Video material about the competition can be found here: http: //ais.uni-bonn.de/nimbro/AVATAR/

 $^{^4}$ https://www.youtube.com/watch?v=yGwJIDBMolk



Figure 12: Qualitative results of our user questionnaire. We show the median, lower and upper quartile (includes interquartile range), lower and upper fence, outliers (marked with \bullet) as well as the average value (marked with \times), for each aspect as recorded in our questionnaire.



Figure 13: ANA Avatar XPRIZE Semifinal Scenarios (left) and the used objects (right). The three scenarios were: Solving a jigsaw puzzle (top), celebrating a business deal (middle), and exploring an artifact (bottom).

We analyzed the recorded data during the official semifinal test runs to obtain more insight on the technical performance of our system. In the following, we report some results.

5.3.1. Safety System

One important aspect of our avatar robot is to operate safely next to and in cooperation with humans. Therefore, the avatar arm controller stops the Panda arm immediately when unexpected high forces/torques are measured (see Section 4.1.1). This behavior was activated three times during the performance of the six official scenarios (see Table 3).

At Day 1, both arm stops were triggered by exceeded torque and velocity limits. In the first case, operator and recipient performed a powerful 'high five' gesture, which resulted in a rapid increase of torque applied to the robot, exceeding the joint torque limits. The predictive limit avoidance module was not able to prevent the operator executing such high contact forces, since the operator arm torque limits were reached as well. In the second scenario on Day 1, the operator waved to the recipient. The different kinematic chains of the operator and avatar arm demand high joint accelerations and speeds on the avatar side for relative low operator arm speeds and accelerations. The limit predictive module tried to slow down the operator waving speed, but humans can overcome these applied feedback forces to always give the operator control for safety reasons. Here, the operator learned immediately continuing waving in a slightly slower and less accelerationneeding manner. At the second competition day, the right avatar arm stopped once while not actively being moved. After some investigation, we found a software failure which led to a very unlikely race condition. This was easily fixed after the competition.

In all cases, the avatar and operator robot continued to work in a safe manner, which is necessary for humanrobot interactions. Our intuitive feedback to the operator



Figure 14: Workspace analysis for all scenarios (solving a puzzle, celebrating a business deal, exploring an artifact) during the ANA Avatar XPRIZE Competition semifinal. Top: Hand positions of the avatar (blue) and operator (magenta) depicted in the common frame. Bottom: Operator VR perspective.



Figure 15: Avatar arm latency analysis. The graph shows the resulting mean translation error, comparing the current avatar hand position with an operator command shifted into the future. The mean is calculated over both arms and all six official competition scenarios. The minimum can be found at 44 ms, which is the estimated round-trip latency in our system — including network and motion execution.

gave instantaneous situation awareness and allowed the operator to react and learn immediately to continue the ongoing tasks.

5.3.2. Arm Controller Accuracy

After optimizing the usable workspace for the operator (see Section 5.1), Fig. 14 shows the used workspace for all three scenarios at the second competition day. The required workspaces during the competition tasks were relatively low, such that we had no issues with our system. We used the same data to analyze the spatial and temporal accuracy of our arm controller, i.e., how precise and fast the avatar arms follow the operator commands during an evaluation task.

The measured operator and avatar hand positions were captured with 1000 Hz over the duration of the six official competition scenarios. We then calculated the translation error between the corresponding target and goal position for each arm, resulting in a 6.6 mm mean translation error. In addition, we investigated the estimated delay in our arm controller including network communication, controller runtime, and motion execution. We compared the measured hand position of the avatar with the up to 100 ms delayed operator's arm control position and calculated the mean translation error as explained above. Fig. 15 shows the mean translation error over all competition tasks and both arms for a given temporal offset. The minimum translation error is archived using a temporal shift of 44 ms, resulting in a mean translation error of

Table 4: Translation error [mm] at ANA Avatar XPRIZE semifinals.

	Day 1			Day 2		
Scenario	1	2	3	1	2	3
$\begin{array}{c} \text{Left Arm} \\ 44\text{ms shift}^1 \end{array}$	$6.4 \\ 5.8$	$5.7 \\ 5.3$	7.9 7.6	$4.9 \\ 4.6$	$\begin{array}{c} 3.8\\ 3.6\end{array}$	4.8 4.5
$\begin{array}{c} \text{Right Arm} \\ 44 \text{ms shift}^1 \end{array}$	$13.2 \\ 12.5$	$9.5 \\ 9.2$	$6.4 \\ 6.0$	$\begin{array}{c} 6.3 \\ 6.3 \end{array}$	$5.5 \\ 5.2$	$5.1 \\ 4.8$

 $\overline{1 \text{ Delay between operator station and avatar (see Fig. 15)}}$

6.3 mm. The results with and without temporal shift for both arms and all tasks are reported in Table 4. Noticeable is the comparably large error for the right arm during Scenario 1 and 2 at the first competition day. This can be explained by the executed safety stops (see Section 5.3.1). It takes about two seconds to safely fade the avatar arm to the commanded pose. This results in possibly large position errors. Overall, the observed errors are rather small and not noticeable for the operator and could potentially be decreased by an even more accurate calibration on a whole-system level. The estimated round-trip execution latency of 44 ms could be considered for future system improvement.

6. Discussion & Conclusion

This work presented a bimanual telemanipulation system consisting of an exoskeleton-based operator control station and an anthropomorphic avatar robot. Both components communicate using our force and haptic feedback controller, which allows safe and intuitive teleoperation for both the operator and persons directly interacting with the avatar. The control method is agnostic to the kinematic parameters and uses only a common Cartesian hand frame for commands and feedback. Using the predictive limit avoidance avatar model, arm limits for both the operator and avatar side can be force-displayed to the operator with low latency. The oscillation observer and damper modules detect and suppress oscillations in the feedback control loop by reducing the force feedback gains temporarily. Additional force-torque sensors measurements are used to generate a weightless feeling for the operator while moving the arms without establishing contact on the avatar side.

We evaluated the system using a user study with untrained operators as well as in lab experiments. In addition, the system performed very well at the ANA Avatar XPRIZE Competition Semifinals, scoring 99 out of 100 points. This demonstrates the intuitiveness and reliability of our system and its control methods.

Acknowledgments

This work has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy,

EXC-2070 - 390732324 - PhenoRob.

References

- E. Krotkov, D. Hackett, L. Jackel, M. Perschbacher, J. Pippine, J. Strauss, G. Pratt, C. Orlowski, The DARPA Robotics Challenge finals: Results and perspectives, Journal of Field Robotics 34 (2017) 229–240.
- [2] H. Kitano, S. Tadokoro, I. Noda, H. Matsubara, T. Takahashi, A. Shinjou, S. Shimada, RoboCup Rescue: Search and rescue in large-scale disasters as a domain for autonomous agents research, in: International Conference on Systems, Man, and Cybernetics (SMC), volume 6, 1999, pp. 739–743.
- [3] P. Oh, K. Sohn, G. Jang, Y. Jun, B.-K. Cho, Technical overview of team DRC-Hubo@UNLV's approach to the 2015 DARPA Robotics Challenge finals, Journal of Field Robotics 34 (2017) 874–896.
- [4] A. Stentz, H. Herman, A. Kelly, E. Meyhofer, G. C. Haynes, D. Stager, B. Zajac, J. A. Bagnell, J. Brindza, C. Dellin, et al., CHIMP, the CMU highly intelligent mobile platform, Journal of Field Robotics 32 (2015) 209–228.
- [5] S. Karumanchi, K. Edelberg, I. Baldwin, J. Nash, J. Reid, C. Bergh, J. Leichty, K. Carpenter, M. Shekels, M. Gildner, et al., Team RoboSimian: semi-autonomous mobile manipulation at the 2015 DARPA Robotics Challenge finals, Journal of Field Robotics 34 (2017) 305–332.
- [6] M. Schwarz, T. Rodehutskors, D. Droeschel, M. Beul, M. Schreiber, N. Araslanov, I. Ivanov, C. Lenz, J. Razlaw, S. Schüller, D. Schwarz, A. Topalidou-Kyniazopoulou, S. Behnke, NimbRo Rescue: Solving disaster-response tasks with the mobile manipulation robot Momaro, Journal of Field Robotics 34 (2017) 400–425.
- [7] D. Lee, K. Huang, Passive-set-position-modulation framework for interactive robotic systems, IEEE Transactions on Robotics 26 (2010) 354–369.
- [8] R. J. Anderson, M. W. Spong, Bilateral control of teleoperators with time delay, in: IEEE International Conference on Systems, Man, and Cybernetics (SMC), volume 1, 1988, pp. 131–138.
- B. Hannaford, J.-H. Ryu, Time-domain passivity control of haptic interfaces, IEEE Transactions on Robotics and Automation 18 (2002) 1–10.
- [10] P. Berestesky, N. Chopra, M. W. Spong, Discrete time passivity in bilateral teleoperation over the internet, in: IEEE International Conference on Robotics and Automation (ICRA), volume 5, 2004, pp. 4557–4564.
- [11] J.-H. Ryu, C. Preusche, Stable bilateral control of teleoperators under time-varying communication delay: Time domain passivity approach, in: IEEE International Conference on Robotics and Automation (ICRA), 2007, pp. 3508–3513.
- [12] S. Hirche, M. Buss, Human-oriented control for haptic teleoperation, Proceedings of the IEEE 100 (2012) 623–647.
- [13] V. Chawda, M. K. O'Malley, Position synchronization in bilateral teleoperation under time-varying communication delays, IEEE/ASME Transactions on Mechatronics 20 (2014) 245–253.
- [14] J. Artigas, J.-H. Ryu, C. Preusche, Position drift compensation in time domain passivity based teleoperation, in: 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems, IEEE, 2010, pp. 4250–4256.
- [15] A. Coelho, H. Singh, T. Muskardin, R. Balachandran, K. Kondak, Smoother position-drift compensation for time domain passivity approach based teleoperation, in: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2018, pp. 5525–5532.
- [16] X. Xu, Haptic communication for time-delayed teleoperation, Ph.D. thesis, Technische Universität München, 2017.

- [17] R. R. Burridge, K. A. Hambuchen, Using prediction to enhance remote robot supervision across time delay, in: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2009, pp. 5628–5634.
- [18] K. Hauser, Recognition, prediction, and planning for assisted teleoperation of freeform tasks, Autonomous Robots 35 (2013) 241–254.
- [19] Y. Wu, P. Balatti, M. Lorenzini, F. Zhao, W. Kim, A. Ajoudani, A teleoperation interface for loco-manipulation control of mobile collaborative robotic assistant, IEEE Robotics and Automation Letters (RAL) 4 (2019) 3593–3600.
- [20] D. Rakita, B. Mutlu, M. Gleicher, A motion retargeting method for effective mimicry-based teleoperation of robot arms, in: ACM/IEEE International Conference on Human-Robot Interaction (HRI), 2017, pp. 361–370.
- [21] T. Klamt, M. Schwarz, C. Lenz, L. Baccelliere, D. Buongiorno, T. Cichon, A. DiGuardo, D. Droeschel, M. Gabardi, M. Kamedula, N. Kashiri, A. Laurenzi, D. Leonardis, L. Muratore, D. Pavlichenko, A. S. Periyasamy, D. Rodriguez, M. Solazzi, A. Frisoli, M. Gustmann, J. Roßmann, U. Süss, N. G. Tsagarakis, S. Behnke, Remote mobile manipulation with the Centauro robot: Full-body telepresence and autonomous operator assistance, Journal of Field Robotics 37 (2020) 889–919.
- [22] L. Guanyang, G. Xuda, L. Lingzhi, W. Yan, Haptic based teleoperation with master-slave motion mapping and haptic rendering for space exploration, Chinese Journal of Aeronautics 32 (2019) 723–736.
- [23] F. Abi-Farrajl, B. Henze, A. Werner, M. Panzirsch, C. Ott, M. A. Roa, Humanoid teleoperation using task-relevant haptic feedback, in: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2018, pp. 5010–5017.
- [24] J. Bimbo, C. Pacchierotti, M. Aggravi, N. Tsagarakis, D. Prattichizzo, Teleoperation in cluttered environments using wearable haptic feedback, in: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2017, pp. 3401–3408.
- [25] G. Niemeyer, J.-J. E. Slotine, Telemanipulation with time delays, The International Journal of Robotics Research 23 (2004) 873–890.
- [26] F. Porcini, D. Chiaradia, S. Marcheschi, M. Solazzi, A. Frisoli, Evaluation of an exoskeleton-based bimanual teleoperation architecture with independently passivated slave devices, in: IEEE International Conference on Robotics and Automation (ICRA), 2020, pp. 10205–10211.
- [27] A. Rochow, M. Schwarz, M. Schreiber, S. Behnke, VR facial animation for immersive telepresence avatars, in: 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE, 2022, pp. 2167–2174.
- [28] M. Schwarz, S. Behnke, Low-latency immersive 6D televisualization with spherical rendering, in: IEEE-RAS 20th International Conference on Humanoid Robots (Humanoids), 2021.
- [29] M. Schwarz, C. Lenz, A. Rochow, M. Schreiber, S. Behnke, NimbRo Avatar: Interactive immersive telepresence with forcefeedback telemanipulation, in: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2021.
- [30] M. Schwarz, S. Behnke, Low-latency immersive 6d televisualization with spherical rendering, in: 2020 IEEE-RAS 20th International Conference on Humanoid Robots (Humanoids), IEEE, 2021, pp. 320–325.
- [31] E. Jacobsen, R. Lyons, The sliding DFT, IEEE Signal Processing Magazine 20 (2003) 74–80.
- [32] E. Jacobsen, R. Lyons, An update to the sliding dft, IEEE Signal Processing Magazine 21 (2004) 110–111.
- [33] D. Kubus, T. Kroger, F. M. Wahl, On-line rigid object recognition and pose estimation based on inertial parameters, in: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2007, pp. 1402–1408.