Comparing Virtual Reality Interfaces for the Teleoperation of Robots

Rebecca Hetrick, Nicholas Amerson, Boyoung Kim, Eric Rosen [1], Ewart J. de Visser, and Elizabeth Phillips
United States Air Force Academy, Brown University [1]
{C20rebecca.hetrick, C20nicholas.amerson, boyoung.kim.kr.ctr, eric_rosen[1], ewart.devisser.nl.ctr, elizabeth.phillips}
@usafa.edu,@brown.edu[1]

***Abstract*—Whether exploring a defunct nuclear reactor, defusing a bomb, delivering medicine to quarantined patients, repairing the International Space Station from the outside, or providing dexterous manipulation for those with motor impairments, robots have the ability to be in places where humans cannot go, can augment the capabilities of humans, and improve quality of life and work. Since even the most advanced robots have difficulty completing tasks that require grasping and manipulation, human teleoperation is often a practical alternative for these types of tasks. By importing the dexterity, expertise, and wealth of background knowledge of a human operator, robots can leverage the skills of their human teammates without requiring humans to be physically present. However, existing robot teleoperation interfaces often rely on 2D methods to view and interact with the 3D world, which is cumbersome for human operators. Virtual reality interfaces may be suitable for resolving problems with traditional teleoperation interfaces (e.g., perspective adjustment, action specification). The goal of this research was to investigate the efficacy of using two different Virtual Reality interfaces—positional control, similar to waypoint navigation, and trajectory control, similar to click and drag—for remotely controlling a Baxter robot to complete a variety of dexterous manipulation tasks. The results of this study will help us to develop control interfaces that allow for more intuitive robot manipulation and ultimately, better distal collaborations between humans and robots.**

***Index Terms*—Virtual reality, human-robot interaction, teleoperation, interface design**

I. Introduction

Many of the application domains for robotic teleoperation involve humans controlling robots at a distance. Whether navigating a defunct nuclear reactor, defusing a bomb, repairing the International Space Station from the outside, delivering medicine in quarantined environments, or providing dexterous manipulation for those with motor impairments, robots have the ability to be in places where humans cannot go, can augment the capabilities of humans, and can improve quality of work and of life. Due to challenges in robot perception, planning, and control, even the most advanced robots have difficulty completing grasping and manipulation tasks. Thus, human teleoperation is often a practical alternative to autonomous manipulation, because it leverages the dexterity, expertise, and wealth of background knowledge of human operators without requiring them to be physically present. However, in order to perform manipulation tasks well, human operators need high-fidelity control over the robot's actuators and an accurate and rich visualization of the robot's environment.

State-of-the-art teleoperation systems require the operator to both manage their view of the scene and, separately, command the robot's actuators. Existing interfaces generally rely on computer monitors to display sensor data, and joysticks or keyboards to actuate the robot. These approaches force humans to rely on 2D methods to view and interact with a 3D world, which often makes using them cumbersome and workload-intensive for human operators [1, 2]. Consequently, there is a need for alternative methods of interfacing with and controlling robotic systems. Recent renewed interest in, and lowered cost of virtual reality (VR) devices have raised the possibility of using VR as an interface for robot teleoperation. VR promises a user experience that is both immersive and detailed, coupled with complete freedom of viewpoint (i.e., 3D) and provides a natural method of controlling robot action (e.g., mapping human arm movement to robot arm movement for instance [3]). VR systems may, therefore, be suitable for resolving many problems associated with state-of-the-art robot teleoperation interfaces.

II. Background

Researchers have begun investigating the use of consumer-grade AR/VR systems for robotic control. Whitney et al. [4] showed that non-expert users were more efficient in teleoperating a robot to complete a number of dexterous manipulation tasks using VR than more traditional keyboard and monitor interfaces. Teleoperating a robot with a VR interface resulted in faster completion times, lower cognitive workload, and higher usability and user satisfaction scores than teleoperating the robot via more traditional keyboard and monitor interfaces. Lipton [5] developed a VR-robot teleoperation system that employed a homunculus model of the robot, in which the user was virtually embedded in a "control room" inside the robot's "mind." In an informal evaluation, the researchers asked users to control the robot to engage in a number of manufacturing and pick and place tasks with objects of different shapes and compliance. Via the homunculus VR model, users successfully picked up each item, transferred that item between robot hands, and finally placed each item in a bin.

Consumer grade mixed reality displays are also being used to support human-robot interaction. In a study by Rosen et al [6], researchers investigated the use of mixed reality displays (i.e., the Microsoft HoloLens) to communicate a robot's planned motion while working with humans in collocated spaces. Communicating intended motion in such a way could help reduce risk for accidents in many industrial settings. Participants who used the HoloLens to view the robot's planned motion made faster and more accurate judgements about collisions than participants using a 2D display. Research by Gadre et al [7] extended the use of the Microsoft HoloLens to allow users to not only view information from the robot, but also use the device as a visual programming interface, using the interface to command the robot to perform and revise tasks. The majority of participants found visual programming using the mixed reality headset to be more intuitive and less cognitively demanding than using a more traditional syntax method of programming the robot. They concluded that using consumer-grade mixed reality displays shows great promise for making robot programming accessible to a variety of end users.

Because several researchers have demonstrated the promise of using consumer grade AR/VR hardware for both robot communication *to users*, and robot control *by users*, there is great value in continuing to explore the use of consumer grade AR/VR hardware to improve human-robot interaction. Thus, it is worthwhile to investigate the conditions under which different VR interface designs may be more effective than others. Doing so will allow for the creation of control paradigms that work best under a variety of different use cases—teleoperation for fine manipulation tasks may be improved by a different type of VR interface than teleoperation for gross motor tasks, for instance. The purpose of this study was to evaluate the effectiveness, usability, and resulting workload associated with using two distinct VR robotic teleoperation interfaces for controlling a robot to perform a variety of dexterous manipulation tasks.

III. Method

A. Participants

Supported by a power analysis conducted using G*Power [8], we planned to recruit 52 participants for this study. However, disruptions to in-person data collection due to the spread of COVID-19 resulted in only 12 participants (11 males, 1 female) completing the study. All participants were volunteers recruited from the United States Air Force Academy, with ages ranging from 18 to 28 ($M = 20.82, SD = 3.28$). Additionally, we asked participants to report their prior experience working with VR, as well as their prior experience working with robots from 0 (No experience) to 10 (A lot of experience). On average participants reported little prior experience working with VR ($M = 3.17, SD = 2.89$). However, four participants reported a 5 or higher on the scale used to indicate prior VR experience. One of those four participants reported "A lot of experience" working with VR. Participants also reported low prior experience working with robots, ($M = 2.42, SD = 2.15$). Only two participants reported above a 5 on the scale used to report prior experience working with robots. No participants reported symptoms of VR sickness while completing or shortly after the study. Participants received course credit in return for their participation. All materials associated with the study were reviewed and approved by the United States Air Force Academy's Institutional Review Board and Survey Control Office.

B. System Overview: Baxter Robot, HTC Vive and ROS Reality Bridge

1) Baxter robot: Baxter, a ROS enabled robot created by Rethink Robotics for industrial pick and place tasks, was used as the robotic platform for this experiment. Baxter is equipped with two 7 DOF arms, each with a parallel gripper capable of picking up and replacing objects. Baxter is also equipped with two RGB eye-in-hand cameras near its wrists, that point down toward its end effectors. We outfitted our Baxter robot with a Microsoft Kinect v2, mounted to the top of the robot's head. The Kinect was used to build a 3D point-cloud of the physical experimental environment to be displayed to users in VR.

2) HTC Vive: We used the HTC Vive headset and controllers as the Virtual Reality head and handsets for the user. On the Vive controllers, we used the side and trigger buttons as binary inputs for completing the experimental tasks. Using either, or alternating between, both side buttons on a given controller was used to move the corresponding Baxter arm. The trigger button on the controller was used to open and close Baxter's corresponding gripper. The controllers were mapped bilaterally onto Baxter's two arms, such that participants could control both of Baxter's arms independently using the two Vive controllers.

3) ROS Reality Bridge: ROS Reality Bridge is an open source, over-the-Internet system that allows any ROS-enabled robot, like Baxter, to communicate with any Unity-compatible virtual or augmented reality headset (with or without handsets) via Unity game engine. It allows users to bilaterally view and control robots over-the-Internet using consumer-grade VR and AR hardware. ROS Reality has served as the technical basis for the VR research in Whitney et al. [1] and for the AR research in Rosen et al. [6].

Specifically, the ROS Reality Bridge system is composed of an an HTC Vive connected to a computer running the Unity game engine. Unity builds a local copy of the robot based on its URDF (Unified Robot Description Format) with a custom-made URDF parser. Unity connects to a ROS network over the Internet via a Rosbridge WebSocket connection. The pose and imagery of the robot's wrist cameras are compressed, bundled, and sent via this WebSocket connection, as well as the color and depth image of the Kinect v2 mounted on top of the robot's head. The color and depth image are built into a point-cloud in Unity via a custom shader. The robot transform data is converted to Unity coordinates and updates the virtual robot's pose. When the user holds down the Vive's handset side controller buttons, the pose(s) of the user's controller(s) are converted to the ROS coordinate frame and sent back to the robot, which uses an inverse kinematics solver to move the robot's end effectors to the specified pose(s). Full system details can be found in [1] with open source code and documentation available at https://github.com/h2r/ros_reality_bridge.

4) Virtual environment: The environment represented in VR consisted of three primary informational components: 1) a 3D model of the robot, obtained by importing a description of the robot in URDF format which was continually updating with the TF topic from the robot's ROS network; 2) a 3D point-cloud of the scene, obtained by a manually calibrated Kinect v2 sensor mounted on top of the robot's head; and 3) a display of Baxter's two wrist cameras (i.e., two 1280 x 800 pixel RGB cameras, downscaled to 400 x 600, showing a live image of the environment immediately forward of the robot's gripper).

In the virtual environment, participants were able to choose to assume either a robocentric or egocentric position, or switch between them at any time by walking around the virtual space, or around the virtual model of the robot. In a robocentric model, the human and the robot share a virtual space, but are not necessarily superimposed on one another. In an egocentric model, the person can step into the virtual model of the robot, thereby superimposing themselves on the robot and assuming the pose and first-person perspective of the robot as if they "were the robot."

5) Physical environment: The physical laboratory space was equipped with the Baxter robot, a 5 x 3 ft. table placed in front of the Baxter, a computer workstation for participants to complete our subjective measures and for experimenters to record our objective measures, and an approximately 10 x 10 area of floor space to allow participants to physically walk around the virtual robot depicted in the Vive headset. Experimenters also surrounded the table with a small temporary wall made with cardboard to help keep objects used in the experimental tasks from falling onto the floor.

C. Experimental tasks

Participants teleoperated the Baxter robot to complete four tasks; two gross motor movement tasks and two fine motor movement tasks.

1) Gross motor tasks: : The first gross motor task was a keyboard press. A QWERTY keyboard was placed in front of Baxter and the participants' goal was to use the VR headset and controllers to teleoprate Baxter to complete 5 presses (i.e., trials) of the "Space Bar" without missing.

The second gross motor task was controlling Baxter to perform a drumming rudiment—a paradiddle-diddle. We placed a dry erase marker in each of Baxter's grippers with large sheets of paper attached to the table underneath each of the robot's end effectors. Participants' goal was to complete the paradiddle-diddle pattern (i.e., right, left, right, right, left, left) while aiming for two small dots drawn on each piece of paper. Participants were given one attempt to complete this task.

2) Fine motor tasks.: The first fine motor movement task was the YCB (Yale-CMU-Berkeley) cup stacking task [9]. This task involved stacking 4 cups that gradually increase in size inside one another. Participants were given 3 attempts to direct Baxter to stack the smaller cups into the larger cups.

The last fine motor movement task was a pouring task where participants were asked to control Baxter to pour small plastic beads into a glass places inside of a bowl placed on top of a plate. Participants used the VR system to control the Baxter robot in pouring the beads out of a container and into the cup. Participants were given three attempts to try to pour as many beads into the cup as possible (i.e., by not spilling beads into the bowl, or onto the plate).

D. Subjective measures

1) Biographical data questionnaire: Participants were asked to fill out a biographical data questionnaire that asked them for biographical information regarding age, gender, handedness, prior experience with VR, and prior experience with robots.

2) Cognitive workload NASA-TLX: Participants also completed the NASA-TLX, a subjective measure of cognitive workload most frequently used in order to study human-machine or human-interface interaction. The participants provided ratings of their workload on six subscales: mental demand, physical demand, temporal demand, effort, frustration, and performance. The initial five subscales range from 0 (Low) to 100 (High). The sub-scale for performance ranges from 0 (Perfect) to 100 (Failure). Overall workload scores were derived by computing an average of participant ratings on the six subscales, after reversing scores on the performance subscale.

3) System Usability Scale (SUS): The SUS assesses perceptions of overall system usability by asking participants to rate ten statements on 5-point Likert scales ranging from 1 "strongly disagree" to 5 "strongly agree." The statements cover different perceptions of the system, such as the participant's perceived likelihood to use the system in the future, the frequency with which they would like to use it, as well as perceived complexity, consistency, and cumbersomeness. After reversing negatively worded items, scores on each item on the SUS are combined and then multiplied by 2.5 to provide an overall usability score that can range from 0 to 100.

E. Objective measures

For all tasks, time to completion across attempts of the task were recorded along with a measure of precision and a measure of accuracy for each attempt. Accuracy was defined as how closely the participant could control the robot to complete each task (e.g., reach a commanded position), and precision was defined as how well participants could control the robot to repeat each task. A detailed description of the measures of precision and accuracy for each of the gross motor and fine motor tasks are given below.

1) Gross motor task 1: Keyboard press: Accuracy was measured as the average distance in inches from the center of the space bar (marked with a white square) to each key pressed across 5 attempts. Precision was measured as the standard deviation of the distance between each key press across 5 attempts.

2) Gross motor task 2: Drum Rudiment: Accuracy was measured as the average distance away from the center of the target (large dot on the on the page). Precision was measured as the standard deviation of each mark created by the participants away from the center of the target.

3) Fine motor task 1: Cup stacking: Accuracy was measured on of a scale ranging between 1 and 0. A participant would receive a 1 if the smaller cup(s) were successfully placed into the larger cup. They would receive a 0.5 if the cup went in diagonally, but stayed in, and they would receive a 0 if the cup failed to go into the other cup. The average of these scores across attempts was used to calculate task accuracy. Precision was measured as the standard deviation of these attempts.

4) Fine motor task 2: Pouring: For the pouring task, accuracy was measured as the proportion of beads that the participant was successfully able to control Baxter to pour into the glass alone. Precision was measured as the average difference across trials in the amount of beads in each of the glass, bowl, and plate (i.e., change across trials in each).

F. Design

This study used a 2 (Trajectory control vs. Positional control) x 2 (Fine motor tasks vs. Gross motor tasks) mixed between-within subjects design. Each participant was pseudo-randomly assigned to control the Baxter robot using one of the two control interfaces. Participants in each interface condition completed both the gross motor and fine motor tasks. Descriptions of the control interface conditions are given below.

1) Trajectory control: The actions of the trajectory control interface condition somewhat mimicked that of clicking and dragging a computer mouse. When the user pressed the HTC Vive side controller button and moved the controller to their desired position, the Baxter robot's corresponding arm would follow the relative trajectory of the controller's movement.

2) Positional control: The positional control interface condition somewhat resembled waypoint navigation commonly used with mobile robots. When the HTC Vive side controller button was pressed, the system recorded the location of the controller as a coordinate in 3-dimensional space. Then, as long as the controller button remained pressed, the Baxter robot's arm would move until it reached that 3-D coordinate in space and then would stop. Once the button had been released, the system forgot the last recorded location and the button needed to be re-pressed to record another location for Baxter's arms to move.

G. Procedure

Upon arriving in the lab, participants were provided with informed consent information. Once consent was provided, participants were asked to put on the HTC Vive headset, to adjust it for comfort, and to orient themselves in the virtual world. Participants were then instructed that they could choose which orientation of the robot they would like (i.e., egocentric, robotcentric, or switch between the two). All participants assumed the egocentric model of the robot. They were then given a brief tutorial on how to use the HTC Vive controllers to teleoperate the robot in accordance with their interface control condition. Then, participants completed both gross motor manipulation tasks, followed by both fine motor manipulation tasks. Participants were asked to complete the SUS and the NASA-TLX after each pair of tasks. Once complete, participants were thanked for their time and granted course credit for their participation. The study took approximately one hour to complete.

IV. Results

A. Accuracy and Precision

A series of independent samples t-tests were conducted to test for statistically significant differences in accuracy scores and in precision scores between the two VR robot control conditions across each of the manipulation tasks.

1) Gross motor: Keyboard press: For the keyboard press task, the difference between accuracy scores in the positional and trajectory control conditions was statistically significant, $t(10) = 4.402, p = 0.001$, Cohen's $d = 2.38$, but violated the equal variance assumption of the statistical test, equal variance not assumed $t(10) = 3.878, p = 0.011$. Participants in the positional control condition had better accuracy scores, (i.e., shorter average distances in inches from the target) than participants in the trajectory control condition, $M = 1.11, SD = 0.31$ and $M = 2.40, SD = 0.70$ respectively. Participants also completed the keyboard press trials faster (in seconds) in the positional VR control interface condition $M = 141.43, SD = 65.30$ than in the trajectory interface condition $M = 233.60, SD = 88.30$. However, the differences in completion times between conditions only approached statistical significance $t(10) = 2.089, p = 0.063$, Cohen's $d = 1.19$. Between the two control conditions, scores on the measure of precision for the keyboard press task were not statistically significant from one another $t(10) = 0.166, p = 0.872$, Cohen's $d = 0.10$, indicating that participants in both conditions were roughly equally precise in controlling the robot to complete the key presses.

2) Gross motor: Drum rudiment: Because completing the drum rudiment task required participants to control both of Baxter's arms, accuracy and precision scores were computed for each. There was a significant difference in accuracy scores between the VR robot control conditions for the right arm. When controlling the right arm of Baxter, participants who used the positional control interface condition were more accurate than participants who used the trajectory control interface, $M = 1.67, SD = 0.79$ and $M = 2.80, SD = 0.88$ respectively $t(10) = 2.353, p = 0.040$, Cohen's $d = 1.35$. Participants in the positional VR control condition were also statistically significantly faster, $M = 83.43, SD = 22.63$, in completing the drum rudiment task than participants in the trajectory control condition, $M = 269.00, SD = 188.35$, $t(10) = 2.632, p = 0.025$, Cohen's $d = 1.35$. Accuracy scores between conditions for the left arm were not statistically significantly different from one another, and neither were precision scores between conditions for both the left and right arms, all $p's > 0.05$

3) Fine motor: Cup stacking: Cup stacking using the positional control interface was more accurate $t(10) = 2.757, p = 0.020$, Cohen's $d = 1.51$, more precise $t(10) = 2.406, p = 0.037$, Cohen's $d = 1.39$, and faster $t(10) = 2.831, p = 0.018$, Cohen's $d = 1.49$, equal variance not assumed for task completion time $t(10) = 2.831, p = .018$ than cup stacking using the trajectory control interface.

TABLE I
TABLE OF MEANS AND STANDARD DEVIATIONS FOR MEASURES OF ACCURACY, PRECISION, AND COMPLETION TIMES FOR THE CUP STACKING TASK.

Measure	**Condition**	**Mean**	**SD**	**N**
Accuracy	Positional	0.93	0.13	7
	Trajectory	0.60	0.28	5
Precision	Positional	0.12	0.23	7
	Trajectory	0.46	0.26	5
Task completion time (sec.)	Positional	180.29	37.76	7
	Trajectory	467.00	269.53	5

4) Fine motor: Pouring: Finally, for the pouring task, there were no significant differences between VR robot control conditions on the measures of accuracy, $t(10) = 1.491, p = 0.167$ precision, $t(10) = 0.143, p = 0.889$ or task completion times $t(10) = 0.455, p = 0.641$.

B. NASA-TLX: Cognitive workload

A 2 x 2 mixed between-within subjects ANCOVA was conducted to assess the impact of two different VR robot control interface conditions (trajectory vs. positional control) on participants' cognitive workload scores across two task types (gross motor vs. fine motor tasks) while controlling for prior VR experience. The interaction between control condition and task type was not statistically significant $Wilks'\Lambda = 0.950,\ F(1, 9) = 0.474, p = 0.508, \eta^2 partial = 0.050$. There was a statistically significant main effect for task type, however. Participants reported significantly higher workload scores after completing the fine motor tasks ($M = 52.100, SD = 8.014$) than after completing the gross motor tasks ($M = 38.614, SD = 10.388$), $Wilks'\Lambda = 0.386$, $F(1, 9) = 14.305, p = 0.004, \eta^2 partial = 0.614$

C. SUS: Usability

A second 2 x 2 mixed between-within subjects ANCOVA was conducted to assess the impact of two different VR robot control interface conditions (trajectory vs. positional control) on participants' perceived usability of those interfaces across the task types (gross motor vs. fine motor tasks) while controlling for prior VR experience. The interaction between control interface and task type was not statistically significant $Wilks'\Lambda = 0.988,\ F(1, 9) = 0.106, p = 0.753, \eta^2 partial = 0.012$, nor were the main effects for task type or control conditions, all $p's > 0.05$.

However, participants in the positional control interface condition ($M = 47.18, SD = 11.05$) reported higher subjective usability scores than participants in the trajectory control interface condition ($M = 38.25, SD = 10.94$) across both the gross motor and fine motor tasks, although the differences were not statistically significant.

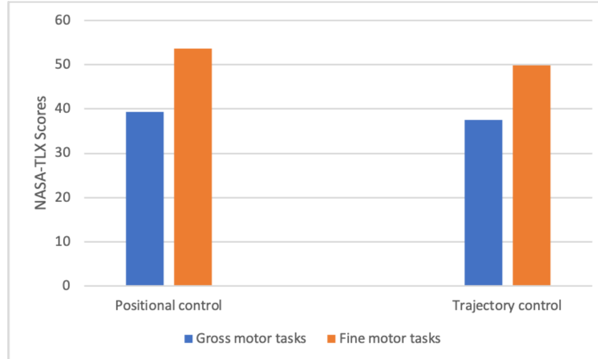

Fig. 1. Cognitive workload scores in interface conditions by task type. Interaction was not statistically significant

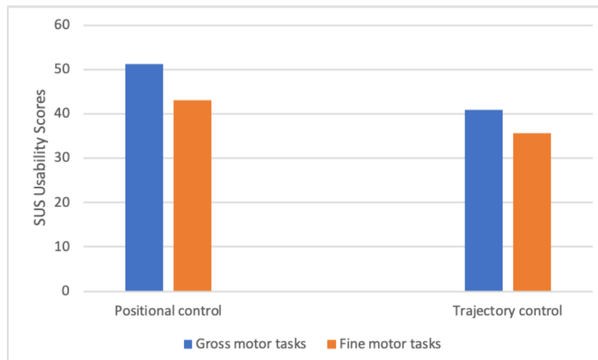

Fig. 2. System Usability Scores in interface conditions by task type. Interaction was not statistically significant.

V. DISCUSSION

With the rapid proliferation of consumer-grade VR systems at accessible price points, smaller form factors, and with better on-board computing, there has been increasing interest in using these systems to solve many of the problems associated with robot teleoperation, manipulation, and grasping. The purpose of this study was to test the effectiveness, usability, and cognitive workload associated with using two different VR robotic control interfaces for completing a variety of gross and fine robot manipulation tasks.

A clear pattern of differences between control interfaces on the subjective measures of performance did not emerge. There were no statistically significant differences in usability or workload scores between the interface conditions. However, participants did report that completing the fine motor tasks was cognitively more demanding than completing the gross motor tasks using both of the control interfaces. It may be important, then, to design VR robot control systems that can support users—perhaps by changing the magnitude in gain in mapping human movement to robot movement—when manipulation tasks change from gross to fine, or when increasingly precise teleoperation of the robot is needed. Similar to an adaptive automation paradigm [10, 11, 12], doing so may be beneficial for aiding humans in managing workload and maintaining situation awareness, both of which are significant and known challenges for robot teleoperation [2].

Although clear differences on the subjective measures were not found, stronger differences between conditions emerged for the objective measures of performance. For novice users of our system, teleoperating the Baxter robot using the positional control interface was faster and more accurate for both of the gross motor tasks (Keyboard press, and right arm of Paradiddle-diddle task), as well as for the fine motor cup stacking task. One potential reason that the difference in accuracy between control conditions for controlling Baxter's left hand in the paradiddle-diddle task did not emerge was likely because the vast majority (10 out of 11) of our participants were right as opposed to left-hand dominant. Cup stacking was also more precise when using the positional control interface as compared to the trajectory control interface.

Even though disruptions to in-person data collection due to the spread COVID-19 resulted in low power for our statistical analyses, taken together, these results suggest that positional control paradigms may provide many practical benefits for novice users controlling robots in manipulation tasks at a distance. Providing novice users with easy, intuitive means to control robots has many practical implications. With the proliferation of robots working in manufacturing, driving, pick and place and other work settings, there is an increasing need for collaborative robots (COBOTs) that can work in tandem with humans [13]. If VR interfaces can provide a simple means to program and control robots, we can leverage the expertise of humans to build better, more intuitive robotic partners. For instance, [14] noted that often the only thing a robot needs to perform well is a little assistance from a human. Using VR technologies, domain experts (who may or may not be robotics experts) may be able to virtually provide robots with the domain expertise needed to overcome software and hardware problems that [currently] limit the ability of robots to complete tasks autonomously—reconciling imagery, picking up deformable objects, or driving the last mile, for instance. VR interfaces show promise for allowing a variety of end users to successfully work with robots.

VI. Conclusion

Deft robot manipulation can save and improve many lives. Our study compared two different VR robot teleoperation interfaces for performing dexterous manipulation tasks with a Baxter robot. Results revealed that a positional control paradigm (i.e., similar to waypoint control) was more beneficial for allowing novice users to control the robot across manipulation task types. The results have important implications for designing better human-robot collaborations.

VII. Acknowledgment

This work is supported by the Air Force Office of Scientific Research Grant 16RT0881f. Special thanks are also extended to Col. Chris McClernon, Associate Dean for Research and the staff of the Research Office (DFRO) at the United States Air Force Academy for supporting this project.

References

[1] David Whitney et al. "Comparing robot grasping teleoperation across desktop and virtual reality with ROS reality". In: *Robotics Research*. Springer, 2020, pp. 335–350.

[2] Jennifer L Burke et al. "Moonlight in Miami: Field study of human-robot interaction in the context of an urban search and rescue disaster response training exercise". In: *Human–Computer Interaction* 19.1-2 (2004), pp. 85–116.

[3] Shannon KT Bailey et al. "Development of Gesture-based Commands for Natural User Interfaces". In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*. Vol. 61. 1. SAGE Publications Sage CA: Los Angeles, CA. 2017, pp. 1466–1467.

[4] D. Whitney et al. "Comparing Robot Grasping Teleoperation across Desktop and Virtual Reality with ROS Reality". In: ().

[5] Jeffrey I Lipton, Aidan J Fay, and Daniela Rus. "Baxter's Homunculus: Virtual Reality Spaces for Teleoperation in Manufacturing". In: *IEEE Robotics and Automation Letters* 3.1 (2018), pp. 179–186.

[6] E. Rosen et al. "Communicating robot arm motion intent through mixed reality head-mounted displays". In: *International Symposium On Robotics Research*. 2017.

[7] Samir Yitzhak Gadre et al. "End-user robot programming using mixed reality". In: *2019 International Conference on Robotics and Automation (ICRA)*. IEEE. 2019, pp. 2707–2713.

[8] Franz Faul et al. "G* Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences". In: *Behavior research methods* 39.2 (2007), pp. 175–191.

[9] Berk Calli et al. "The ycb object and model set: Towards common benchmarks for manipulation research". In: *2015 international conference on advanced robotics (ICAR)*. IEEE. 2015, pp. 510–517.

[10] Mark W Scerbo. "Adaptive automation". In: *Neuroergonomics: The brain at work* (2008), pp. 239–252.

[11] Toshiyuki Inagaki et al. "Adaptive automation: Sharing and trading of control". In: *Handbook of cognitive task design* 8 (2003), pp. 147–169.

[12] Evan A Byrne and Raja Parasuraman. "Psychophysiology and adaptive automation". In: *Biological psychology* 42.3 (1996), pp. 249–268.

[13] J Edward Colgate et al. "Cobots: Robots for collaboration with human operators". In: (1996).

[14] T. Fong, I. Nourbakhsh, and K. Dautenhahn. "A survey of socially interactive robots". In: *Robotics and autonomous systems* 42.3 (2003), pp. 143–166.