An Effective and Intuitive Control Interface for Remote Robot Teleoperation with Complete Haptic Feedback

Cody Glover
DePauw University
johnglover_2009@depauw.edu

Brian Russell
Univerisity of Arkansas
brussell@uark.edu

Allison White
Lehigh University
asw309@lehigh.edu

Matthew Miller
Iowa State University
mamille@iastate.edu

Alexander Stoytchev
Iowa State University
alexs@iastate.edu

ABSTRACT

This paper addresses the development and usability testing of a system to control a remote robotic arm using an identical local arm as a user interface. It allows a user to control each joint of the remote robot by physically manipulating the corresponding joint of the local one. Additionally, it is capable of delivering realistic haptic feedback through the local arm, allowing a user to "feel" the remote environment. Since the arms are identical, the feedback can reproduce the exact forces that are acting on the remote robot. The system was implemented using two state of the art Barrett Whole Arm Manipulators (WAMs), and several experiments were performed to demonstrate its usability. The tests showed that the haptic feedback improves user performance, but can also make a task seem more difficult.

1. INTRODUCTION

Robots have demonstrated a unique capacity to take on roles too difficult, too dangerous, or too undesirable for human workers. Industries have employed robots in hazardous work environments for decades. NASA has also used a wide range of robotic devices for space exploration. Urban search and rescue teams have also deployed robots in hope of locating survivors [10]. In each of these situations, robots have served their purpose in places where humans cannot follow.

There is still a limit, however, on the complexity of the tasks that robots can perform autonomously. Therefore, it is sometimes vital that a human be able to control the robot from a distance. This has led to the development of teleoperated robots, which can be controlled remotely by human operators. A teleoperated robot is able to benefit from a human's perception, judgment, and adaptability, while the operator has the safety and convenience of a remote location. Unfortunately, many of the weaknesses of the current robot control interfaces are amplified when the user is in a remote location.

Various user interfaces, including joysticks, teach pendants and virtual reality interfaces, have been used to control robotic systems. Many of these interfaces attempt to present all necessary information for control of the robot to the operator through some sort of visual feedback. Interpreting this information places a large cognitive load on the user, making the robot much more difficult to operate [2]. Furthermore, many of the simpler interfaces do not allow the user to control the individual joints of the robot. An interface that only allows a user to control the end effector of a robot will frequently cause it to move in an undesirable manner. Individual joints might move at unsafe speeds and can collide with other objects or people.

A way to ease the cognitive load is to register any forces met by the remote robot as haptic feedback. Haptic feedback is a mechanical force that is applied through a control device and then perceived by a user. The control device can be made to respond as though it is interacting with a virtual environment - resisting motion when it encounters a virtual object. This can be used to give a user a sense of touch, including factors like weight, resistance, and friction.

Generating haptic feedback can require a great deal of information, especially when representing a realistic or complex environment. In many cases, the information about the location, shape, and stiffness of objects can be pre-programmed [3, 4, 6]. This is useful when interacting with an entirely virtual environment, or an unchanging physical one. However, when using a haptic interface to interact with a remote real-world scene, this is simply impossible.

We propose an intuitive system for realtime teleoperation of a remote robotic arm with realistic haptic feedback using an identical robotic arm as the local control interface. The use of an identical local arm as a control interface allows a user to exercise precise control over each joint. Moreover, the haptic feedback can represent precisely the forces acting on the remote arm. A haptic interface device with a different geometry than the remote robot could never completely capture the forces acting upon it. So the use of an identical arm allows more rich and accurate haptic feedback than can be achieved otherwise. This paper demonstrates that the addition of this haptic feedback results in a significant improvement in a user's ability to perform teleoperated tasks.

2. RELATED WORK

The history of teleoperation of remote devices began with the creation of the first mechanically controlled master-slave systems in the mid 1940s. In 1954 a way to mechanically separate the master from the slave was developed [7, 8]. In the mid 1980s research began to focus on finding the optimal method for controlling a robot. By the 1990s, transparency, or the ability of a system to provide the user with the sense of actually being present in a remote environment, had become an area of interest [10].

Zhou et al. successfully used a Barrett WAM in a unilateral master-slave setup to control the end effector of a Titan II slave robot [14]. This is not the only research that suggests the WAM as an ideal platform for a system of the type we propose. Heinzmann and Zelinsky recognize the WAM's potential as a safe, human friendly robot [9].

Other research has focused on what can be accomplished using haptic feedback. Turner et al. used a bilateral haptic glove in a series of experiments as an anthropomorphic controller for a remote robot hand [12, 13]. Ansar et al. designed a system to make a real object into a virtual model through the use of a head-mounted virtual reality display and haptic feedback felt through a robot arm [1]. Cotin et al. proposed a surgical use of haptic feedback [4]. They designed a system that allows a surgeon to practice a procedure by using a robot on a pre-built haptic model. Cavusoglu et al. propose a system that has the potential to be used to perform telesurgery [3]. Their system uses haptic feedback to enhance teleoperation, but the feedback must still be generated by the computer rather than determined directly from the remote environment.

In contrast to this previous work, our system provides an intuitive interface that allows precise control of each joint of a remote robotic arm. Additionally, the haptic feedback is generated directly from the remote environment, and represented accurately to the user. None of these previous systems offer the level of control, or the richness of haptic feedback that our system provides.

3. ROBOT TELEOPERATION WITH HAPTIC FEEDBACK

Haptic feedback essentially consists of a mechanical force being applied to a user in some way. Therefore, it must be presented to a user through some kind of mechanical interface. The capabilities and limitations of that interface strongly affect the quality and usefulness of the haptic information. As mentioned before, a haptic interface with a different geometry than the controlled robot cannot completely represent the forces acting upon it. Those forces must be mapped onto the haptic device in some way, usually reducing the amount of information available to the user. The only way to completely represent the forces acting on a remote robot is to use a haptic interface with precisely the same geometry.

There are several advantages to using an identical robotic arm as a control interface for remote teleoperation, and these advantages are fully realized when the robotic arms are backdrivable. A backdrivable robot is one whose joints can be manipulated through the application of external force while the robot is in operation. That is, the motors can both apply and absorb torque simultaneously. A user can control a backdrivable robot by physically pushing on it. This is a very natural and intuitive way to interact with a robot. With virtually no training, a user can immediately begin to manipulate a robotic arm however he or she sees fit.

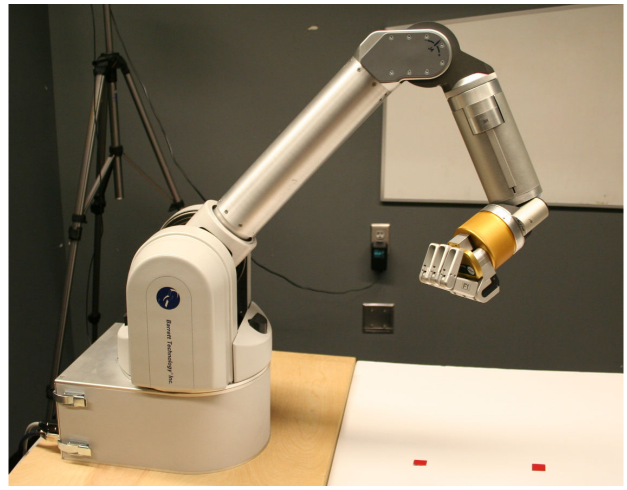

Figure 1: The Barrett Whole Arm Manipulator (WAM). It is a cable driven 7-dof arm that is completely backdrivable. We used a pair of WAMs in all of our experiments.

Therefore, given two identical backdrivable robotic arms we can construct an intuitive system for remote control. The remote arm can be made to mirror the configuration of the local arm, and then the local arm can be used as a control interface. A user can physically manipulate the local arm to control the remote one. If the arms are additionally capable of delivering haptic feedback, then the user can be made to feel the remote environment just as the remote robot does. The feedback can be instantaneous and complete, enabling the user to interact intuitively and proficiently with the remote environment.

4. ROBOTIC PLATFORM

Implementation of a teleoperation system with haptic feedback requires very specific capabilities to be present in the underlying hardware. As mentioned, the robot arms must be backdrivable, and capable of delivering realistic haptic feedback. We used two Barrett Whole Arm Manipulators (WAMs), which have exactly these capabilities (Figure 1).

The Barrett WAM is a human-scaled robotic arm with seven degrees of freedom. It utilizes a cable drive system, which generates less resistance than a gear mesh. This gives the WAM human-like dexterity and inherent backdrivability. Force can be applied to any part of the arm, and it will comply. The cable drive system is controlled by PUCKs - high performance miniature motor controllers connected via a CAN bus running at 500 Hz. This allows the arm to respond quickly to user input, or to apply torques to provide haptic feedback [11]. Additionally, the WAM is able to apply "gravity compensation" torques during real time operation. The robot applies the appropriate torque to each joint to negate its own weight, and thus remains suspended in place. A user can move the robot as if it is weightless, which allows for easy and natural control.

These capabilities are precisely the requirements for remote teleoperation with haptic feedback as we have described. We used two Barrett WAM arms to implement our algorithm, and to perform all of our experiments.

5. TELEOPERATION ALGORITHM

In order for a local arm to control a remote arm, any force applied to the local arm should be transmitted and applied to the remote arm. Conversely, in order for the local arm to respond appropriately when the remote arm encounters resistance, any forces applied to the remote arm should be transmitted and applied to the local arm. So the behavior of the 'controlling' and 'controlled' arms is actually identical, making those roles fundamentally ambiguous.

We developed an algorithm for remote teleoperation of a robotic arm with haptic feedback based on this observation. Initially, both robots are oriented in exactly the same position. Then, during each control cycle, each robot executes the control loop given in Algorithm 1.

When two robot arms execute this algorithm, they both try to move towards the average of their current positions (see Figure 2). This allows either arm to be the 'master' or 'slave.' When one arm is moved, the two arms become misaligned. The arm that was moved torques back towards the position of the one that did not move, and the arm that did not move torques towards the position of the one that did. The user who moves the arm feels extra resistance, as if the arm has the inertia of the two arms combined. The remote arm follows behind the moved arm, trying to reach the midpoint between the two. If the arms respond quickly, and with enough force, they appear to mirror each other.

If a user moves one of the arms, and the other arm hits an obstacle, the user will 'feel' the remote obstacle. More specifically, the remote arm will stop, and press against the obstacle. As the local arm is pulled further along the trajectory, it will torque back towards the joint position of the stuck remote arm. This feedback gives the tactile illusion that the local arm has hit an invisible obstacle.

Algorithm 1 Robot Control Loop

1: **while** true **do**
2: Read local joint angles $J_l = l_1, l_2 ... l_n$
3: Receive remote joint angles $J_r = r_1, r_2 ... r_n$ from remote robot over the network
4: Calculate average joint angle vector $A \Leftarrow (J_l + J_r)/2$
5: Use a PID control algorithm to calculate the additional torque that should be applied to each local joint in order to reach the average position A: $T_a \Leftarrow PID(A, J_l)$
6: Transmit the current joint angles of this robot to the remote robot via a network connection: $Transmit(J_l)$
7: Add the additional torque values T_a to the torques $T = t_1, t_2 ... t_n$ that are to be applied to each joint: $T \Leftarrow T + T_a$
8: Apply the torques T to each joint: $ApplyTorques(T)$
9: **end while**

The actual behavior of the arms is highly dependent on the control algorithm used in step 5 of Algorithm 1. It is specifically sensitive to how quickly the torques are applied - that is, the rate at which the torque increases as the deviation between the positions of the arms grows. We used standard PID controllers packaged with the API of the robot arms. The result was firm and realistic haptic feedback, without any noticeable lag or 'softness' between the arms. We performed several experiments to test the quality of alignment between the arms under several different circumstances. Those experiments and their results are addressed under the section Precision Testing (Section 6).

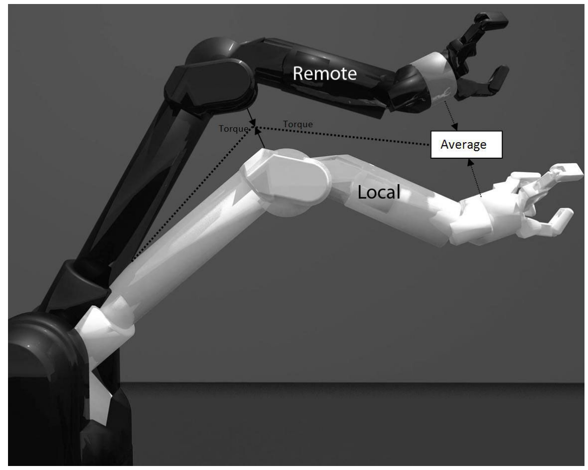

Figure 2: Visualization of the control system at work. The deviation between the two arms has been exaggerated for clarity. Both arms try to pull towards the average between their positions, simultaneously causing them to mirror each other, and transmit all appropriate haptic information.

The two WAM arms were configured to run their control cycle at 500 Hz. We used the UDP protocol to broadcast the joint angles of each robot to its counterpart. We were able to broadcast joint angles at about 5000 Hz, which was more than fast enough to keep up with the update cycle of the WAMs.

We ran the WAM arms with gravity compensation turned on in all of our experiments. Step 7 of our algorithm then consisted of adding the joint torques needed for gravity compensation T to the additional torques T_a calculated in step 5. A user manipulating one of the arms would feel as if it was weightless, but had twice the inertia of a single arm.

6. PRECISION TESTING EXPERIMENTS

In order to test the accuracy of the mirroring of the two WAM arms, we ran three separate tests. Each test lasted for one minute, and during each test all joint angles were recorded for both robots during each control cycle. Since the WAM updates at 500 hertz, 30,000 data points were recorded for each joint during each test.

Test 1: The arms were positioned in a neutral, stationary pose and not moved during the test.

Test 2: An experimenter constantly moved one of the arms through random positions, while the other arm was free to mirror the first arm.

Test 3: An experimenter held one arm in place, while another experimenter applied firm pressure to each joint of the other arm.

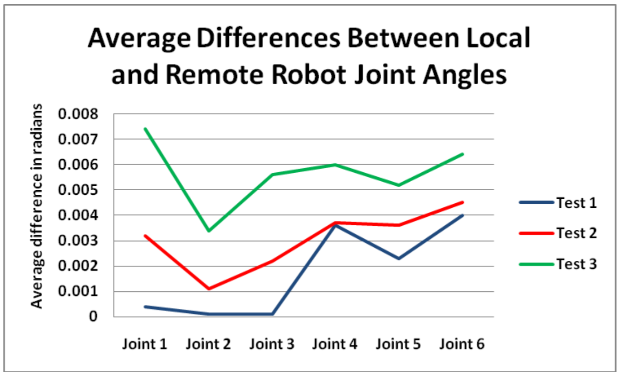

Figure 3: The figure shows the average difference of the joint angles between the remote and local robot arms (in radians). Each data point for each joint represents the average deviation between the two arms over 30,000 samples (one minute).

The joint angle difference between the two arms was calculated for each joint at each timestep in each of the three tests. For each test and for each joint the mean average error was calculated. This gives a measure of the average error between the joints when the arms are still, when they are moved freely, and when one has encountered an obstacle. The results are summarized in Figure 3.

As expected, Test 1 showed the lowest error. The error for each joint increased in test 2 and test 3. However, even in Test 3, in which one arm was held still and firm pressure was applied to the other, the error was very low. In no case did the average error climb above $1/100^{th}$ of a radian. This demonstrates that the arms mirror each other precisely, even when they encounter obstacles and provide strong haptic feedback.

Anecdotally, we observed that the arms would surpass their torque limits and automatically shut off before they would allow a high degree of error between the joints. The users generally experienced the arms to be connected "solidly," with little or no perceived error between them.

7. USABILITY TESTING EXPERIMENTS

We performed several usability tests to establish the usefulness of this interface. We specifically focused on the value of each mode of sensory feedback for performing certain tasks. The goal was to verify that the interface could be used by an untrained participant to perform teleoperation tasks, and to evaluate whether the haptic feedback was helpful in doing so. We also questioned the users as to their perception of the value of each mode of sensory feedback.

As approved by the IRB (08-275, July 9, 2008), thirty participants were randomly assigned to one of three groups. Each group was given a different form of feedback: haptic feedback only, visual and haptic feedback, or visual feedback only. In the haptic feedback group, participants could feel obstacles encountered by the remote arm, but they had no visual information about the remote environment around the remote arm. The haptic and visual feedback group was shown a live video stream recorded by a camera located in the remote environment, in addition to being able to "feel" the remote environment through haptic feedback. The visual feedback group had the same camera setup, but the remote robot was controlled by a unilateral master-slave system (one robot controlling the other with no feedback). The experimental setup is illustrated in Figures 4 & 5.

Figure 4: Remote arm being controlled by the user in the top left corner. During experiments, a screen prevented the user from seeing the remote arm. The camera on the right provided visual feedback.

Figure 5: The view from the users point of view. The remote arm is visible on the television screen.

The ten participants assigned to each condition completed three tasks. In all conditions, the experimental setup consisted of two Barrett WAM robot arms, two Ubuntu Linux PCs controlling the robots, two video cameras, and a television set (see Figures 4 & 5). One of the cameras was capable of displaying the environment around the remote arm on the television to give the participants visual feedback. The other recorded the participant's actions. During the tests, participants were not able to see the remote robot, except through the video feed if their group received visual feedback. Participants were briefed on the capabilities of the robot arms and their configuration.

Upon completion of the tasks, each participant was given a post task survey. This survey asked participants to rank the usefulness of each type of sensory feedback for each task on a scale from one to ten, ten being the highest and one being the lowest. It also asked the participants to rate the difficulty of each task, with ten being the most difficult and one being the least. The three tasks and the results obtained for each are described below.

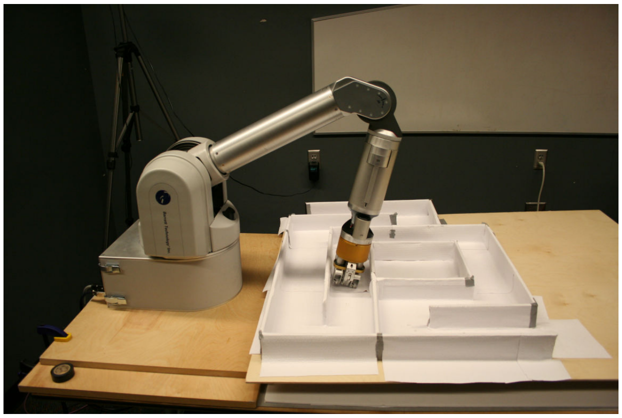

Figure 6: The remote Barrett WAM arm as a user guides it through the maze.

7.1 Task 1: Guiding the robot through a maze

The ability to navigate an arm through a complex environment is important in many real-life applications. Therefore, task 1 required the participants to guide the remote arm through a maze. The maze used in this test was constructed from wood, and its walls were covered with drawer liner to protect the robot. The maze itself measured approximately two feet by three feet with four inch wide paths. Certain areas of the maze were not accessible unless the participant was manipulating all six joints of the robot. In order to navigate the maze the end effector had to be placed down inside the maze walls, and then moved parallel to the floor (see Figure 6).

We evaluated participants based on the time they took to complete the maze task and the number of errors they made, where an error was defined as moving the robot out of the maze, or over a wall, or pushing the robot past its safety limits. When a participant made an error, we paused the timer and returned the arm back to the point where the error occurred. The completion times were recorded for all participants, along with any errors they made while traversing the maze. After five minutes, participants were offered a choice to continue or to stop. At any point after that, they were permitted to stop if they believed that they would be unable to make further progress. Stopping an attempt before reaching the end of the maze was counted as a failure. The results of the experiment are shown in Table 1.

Both groups with visual feedback had times significantly faster than those who only received haptic feedback. The haptic feedback only group was the only one in which participants became frustrated enough with the arm that they chose to give up rather than continue.

Table 1: The average completion time, standard deviation and success rate of all three groups of participants on the maze task.

Group	Average Completion Time (sec)	Std. Dev	Success Rate
Haptics and Vision	82.40	27.31	100%
Haptics Only	107.80	25.39	100%
Vision Only	459.66	216.17	66%

Table 2: The average rating of task difficulty and the benefits of visual and haptic feedback by each group of users. All results were based on a 10-point scale, 10 being the highest and 1 being the lowest.

Maze Task	Haptics Only	Haptics and Vision	Vision Only
Task Difficulty	8.11	7.3	5.0
Benefits of Vision	9.77	7.2	7.6
Benefits of Haptics	7.88	9.0	8.7

However, adding haptic feedback to visual feedback was clearly beneficial. A one-tailed t-test of the average completion time of the haptics and vision group and the visual feedback only group revealed a significant difference, with less than a .04 chance of error with 18 degrees of freedom.

Several general trends became apparent over the course of the experiment. Users in the haptic feedback group were less likely to change the location where they were holding the robot throughout the test. In many situations the arm could not be guided through the maze without simultaneously manipulating multiple joints. This problem occurred most often where the arm entered the straight-aways either closest or farthest from its base. Several participants from the haptic feedback group commented that while approaching these corners, they had reached positions where the arm would not go any farther in the direction they desired. This often led to extensive backtracking.

Participants in the groups with visual feedback avoided these problems. They were much quicker to switch their grasp of the robot when they knew where it had to go. The only instances of backtracking in these conditions occurred in the area closest to the base of the robot. This area was blocked from view by the base of the robot, so their behavior began to resemble that of the haptic group.

It should also be noted that teleoperation of the remote arm was much smoother in the conditions with haptic feedback. Often when a participant was only given visual feedback, he or she would not stop moving the primary arm when the remote arm hit an obstruction. As the difference in the positions of the arms grew, the remote arm applied higher amounts of torque to realign the arms. This resulted in the remote arm applying a great deal of force against the unyielding walls of the maze. Once the arm did get free of the obstacle, it would move at high speeds, which frequently ended in collisions with the walls of the maze. These violent collisions were not observed in the groups with haptic feedback.

The post task surveys revealed an interesting anomaly (see Table 2). The group that received both haptic and visual feedback rated the task as significantly more difficult than the group that received only visual feedback. The group that received only visual feedback perceived the task as easier, even though they performed more poorly. This is the first indication of a repeating theme among these experiments. The benefit of haptic feedback is often not perceived or is under-appreciated by the users.

In this case, the haptic feedback made the difficulty of navigating the remote robotic arm through the maze more tangible. This allowed the users to more fully appreciate the difficulty of the task, but also to perform it more effectively. The users who only received visual feedback could see the remote arm, but could not feel how hard it scraped the walls of the maze, or how fast it moved when it broke free from an obstacle. We believe that this dissociation gave the users the impression that the task was easier than it was. So haptic feedback was useful both in completing the task, and in aiding the understanding of the task's difficulty.

7.2 Task 2: Stacking Rings on a Peg

Task 2 was designed to test participants' ability to manipulate objects with the remote arm. The rings and peg used for ring stacking were part of the Fisher-Price Rock-a-Stack children's toy. The rings were loaded manually into the robot's hand by the researchers. Participants were then asked to orient the arm in such a way that the ring would be placed on the peg when the ring was released from the hand (see Figure 7). Because parts of the Barrett hand are not backdrivable, it was necessary for the participant to notify the experimenters when he or she wished to release the ring. Participants were timed to see how quickly they could orient the arm to the desired location. If the released ring fell on the peg, the task was counted as a success. The times along with the success rates of each test were recorded. The results are included in Table 3.

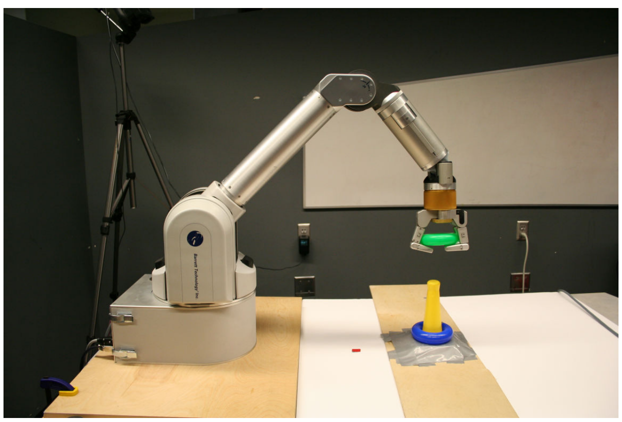

Figure 7: Image of the remote arm under the control of the primary arm during the ring stacking task. The camera used to capture the video can be seen in the upper left hand corner of the image.

The group which received both haptic and visual feedback had a significantly higher success rate than either of the other two groups (see Table 4). However, the time taken to complete the task, regardless of success or failure, was slightly lower for the group that received only visual feedback. The group that received only haptic feedback was by far the slowest and performed the worst at the task. Difficulties with the hand invalidated results from one participant in each condition.

Both haptic and visual feedback seem to be necessary for optimal performance of this task. A major hindrance to using visual feedback from one camera is the lack of depth perception. The two dimensional view made it difficult for some participants to accurately judge the position of the ring in relation to the peg. Haptic feedback provides a solution to this problem. It allows participants to sense when the two are in contact.

Table 3: The average success rate and time in which the two rings were stacked by users during the study.

Test Condition	Ring 1 Success Rate	Ring 2 Success Rate	Ring 1 Time	Ring 2 Time
Haptics & Vision	77.8%	77.8%	24.3	14.4
Haptics	33.3%	22.2%	49.7	29.5
Vision	44.4%	55.6%	18.0	14.6

Table 4: The average rating of task difficulty and the benefits of visual and haptic feedback by each group of users. Ratings were made on a scale from 1 to 10, with 1 being the least and 10 being the most.

Ring Task	Haptics Only	Haptics and Vision	Vision Only
Task Difficulty	7.44	6.1	5.66
Benefits of Vision	10.00	7.7	5.62
Benefits of Haptics	7.33	8.1	9.22

It is interesting, however, that the group that received only visual feedback tended to perform the task faster than the group that also received haptic feedback. Participants in the group with visual feedback had to line up the ring as best they could based on the video feed. Participants in the group with visual and haptic feedback would frequently start the same way, by lining up the ring visually. But then they would often use the haptic feedback to fine tune their position. They would "feel around" until the ring was lined up precisely. This additional step took extra time, but significantly improved the success rate of the participants.

The post task surveys showed that the group that received only visual feedback found this task the easiest, despite their inferior performance. This demonstrates, once again, that the addition of haptic feedback can make a task appear more difficult, while simultaneously increasing a user's proficiency at it.

In general, the groups that were deprived of a sensory modality showed a preference for the feedback they did not receive. That is, the group that received only visual feedback thought that haptic feedback would have been significantly more beneficial than the visual feedback was, and the group that received only haptic feedback thought that visual feedback would have been more beneficial. The group that received both rated the two modalities as roughly equivalent.

This demonstrates that this task requires both forms of perception to complete effectively. As mentioned before, the users tended to use visual feedback to line up the arm in roughly the right position, and haptic feedback to make fine adjustments. Users without visual feedback had a very hard time locating the peg at all, and users without haptic feedback had difficulty positioning the ring precisely over it.

We suggest that this is typical of remote object manipulation tasks. Gross movements can be performed quickly and easily by using visual feedback. However, small, precise movements are more difficult, and are much easier when a user can "feel" what they are doing.

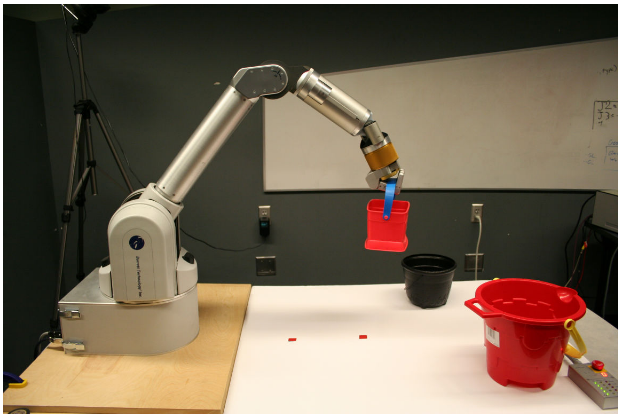

Figure 8: Image of the remote site while users try to compare the weight of three buckets.

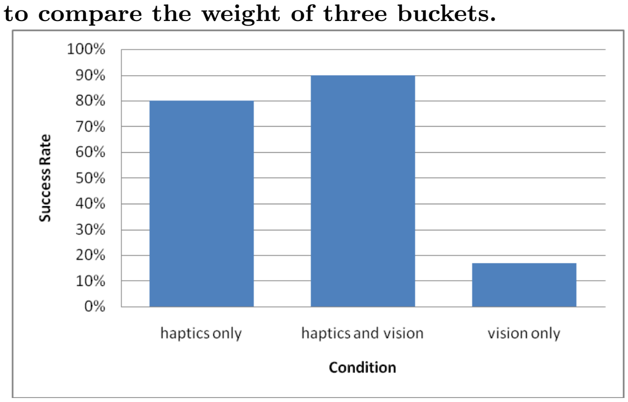

Figure 9: Success rates for comparing and ordering the three weighted buckets by participants in the three conditions.

7.3 Task 3: Ordering Buckets by Weight

In order to test the ability of humans to sense weight through haptic feedback, we created a weighted bucket test. This experiment was also designed to evaluate the sensitivity and quality of the haptic feedback of our teleoperation system. Three buckets were used, each of a different size, shape and weight. The largest bucket weighed .875 pounds and held no additional weight. The medium-sized bucket together with weights placed inside it, weighed a total of 2.5 pounds. The smallest bucket and its weights totaled 1.5 pounds. Each bucket was manually loaded into the robot's hand by the experimenters. Participants were asked to lift each bucket and rank the buckets in order of their weight (see Figure 8). Assigning all three buckets the correct rank was counted as a success. The completion times were not recorded for this test and subjects were allowed to lift each bucket as many times as they wanted. The results are included in Figure 9.

As expected, participants did significantly better at this task when they had haptic feedback. The participants did slightly better when they also received visual feedback, and quite poorly with only visual feedback. However, the 30% success rate for participants in the visual feedback group is higher than initially expected. But these findings are consistent with research showing that people can judge the weight of an object from visual cues alone [6, 5]. Observations showed that while many of the participants attempted to judge the weight of the buckets by simply picking them up, the three successful participants in that group used different techniques. These techniques included swinging each bucket, bouncing the bucket on the hand, and hitting the buckets against each other. The visual cues provided by these techniques can partially compensate for the lack of haptic feedback. It was also noted that members of the visual feedback group seemed to forget that they did not have haptic feedback during this task.

Once again the group that received only visual feedback rated the task as easier than the group that received both visual and haptic feedback (see Table 5). But group that received only haptic feedback thought the task was easier still. This is a very surprising result. The addition of visual feedback caused participants to perceive this task as more difficult, instead of less. One possible explanation is that the buckets' weights did not correspond directly to their size. The addition of visual feedback may have been confusing in this instance. However, despite this confusion, the group that received both forms of feedback still performed the best. So once again the apparent difficulty of the task was not related to how well the participants were able to perform it.

Table 5: The average rating of task difficulty and the benefits of visual and haptic feedback by each group of users. All results were based on 10.

Bucket Task	Haptics Only	Haptics and Vision	Vision Only
Task Difficulty	5.22	6.9	5.8
Benefits of Vision	5.00	8.0	9.4
Benefits of Haptics	8.33	5.4	8.0

An even more surprising result was that the participants in the group that received both visual and haptic feedback rated the usefulness of haptic feedback as very low. This suggests that users were either unaware of the haptic feedback, or became oblivious to its benefit when they also received visual feedback. In this instance, their perception and use of the haptic feedback seemed to be transparent. They believed that the visual feedback was sufficient to perform the task, but their conclusions about the weights of the buckets were clearly heavily informed by the haptic feedback. One can imagine using the remote robot arm to shake a bucket, and perceiving it as heavier than the last one, but attributing this perception to the visual cues of the bucket instead of the haptic feedback from the arm. This effect may have been amplified by the fact that the difference in weights between the buckets was very small - just over a pound. It is difficult to consciously feel that difference without concentrating on the weight. It is easier to obtain a holistic perception of "heavier" and "lighter" without specifically attributing the conclusion to a sensory modality.

Finally, participants in groups that were deprived of one of the two forms of feedback showed a preference for the feedback they did receive. This is in direct contrast to the last experiment. Here, those users who received only haptic feedback realized that they could easily tell the difference between the buckets if they concentrated on their weight. Those who received only visual feedback thought (incorrectly) that they could easily tell the weights of the buckets apart using vision alone.

8. DISCUSSION AND CONCLUSIONS

The current state of robotic teleoperation does not provide the level of control necessary for safe, proficient operation of a remote robot. The addition of haptic feedback effectively improves the proficiency of users at teleoperation tasks. This improves human-robot interaction; as well as the safety and overall control of remote environments.

We have demonstrated a simple intuitive system capable of exercising precise control over a remote robot while simultaneously delivering accurate haptic feedback. Moreover, the level of control over the robot and the realism of the haptic feedback in this system is greater than any standard teleoperation system. The use of an identical, backdrivable robotic arm as a user interface allowed for complete, intuitive control over every joint of the remote arm, and completely realistic haptic feedback.

We have shown that adding this haptic feedback significantly improved the proficiency of users at remote teleoperation tasks. The visual and haptic feedback group was the most successful in all three experiments. This suggests that including haptic feedback in this interface significantly improves an operator's performance. However, haptic feedback by itself was insufficient to replace visual feedback for successful and timely task completion.

In addition to the success rate of the haptics and vision group, users showed a consistent preference to have visual feedback when performing a task. This was true even in tasks where visual feedback might be unnecessary, *i.e.*, determining the weight of buckets.

Even though haptic feedback universally improved the performance of the participants, it also consistently increased the participants' estimation of the difficulty of the task. This seems to suggest that the benefit of haptic feedback goes largely unnoticed. Or, alternatively, that when an operator only receives visual feedback, they do not fully appreciate the difficulty of a task.

Even if an operator cannot perceive the difficulty of a task, that does not mean the task isn't difficult. During the visual feedback trials the remote arm would often hit obstacles or snap free of them only to knock into something else. In real world applications, e.g., telesurgery or manufacturing, these events could be detrimental, or even life threatening. These situations occurred because the user received no visual cues of the extreme situation of the remote arm. This lack of information led to undesirable and dangerous operation, and poor performance in teleoperation tasks. This is not only hazardous to people but to the equipment as well.

Haptic feedback clearly improves the safety and proficiency of robotic teleoperation. It has the potential to improve the proficiency of tasks already performed through teleoperation, and to allow robots to perform new tasks that require more detailed and specialized control.

Our system provides an extremely intuitive control interface for a robotic arm, with an unprecedented level of realistic haptic feedback. The system is simple to implement, requires little training to use, and enables a user to perform complex tasks remotely. The possible applications of this are wide ranging and beyond the scope of this work. However, these capabilities are certainly advantageous, and could provide solutions to many real world problems.

9. ACKNOWLEDGMENTS

This research was performed at Iowa State University as part of a Research Experience for Undergraduates (REU) internship sponsored by NSF (IIS-0552522). We would like to thank the Human Computer Interaction (HCI) Graduate Program at ISU and the Program for Women in Science and Engineering for hosting and sponsoring the REU program during Summer 2008. We would also like to thank Iowa State students Jivko Sinapov and Peter Wong for their help with the application development.

10. REFERENCES

[1] A. Ansar, D. Rodrigues, J. P. Desai, K. Daniilidis, V. Kumar, and M. F. M. Campos. Visual and haptic collaborative tele-presence. *Computers & Graphics*, 25(5):789–798, 2001.

[2] M. Baker, R. Casey, B. Keyes, and H. A. Yanco. Improved interfaces for human-robot interaction in urban search and rescue. *Systems, Man and Cybernetics, 2004 IEEE International Conference on*, 3:2960–2965 vol.3, Oct. 2004.

[3] M. C. Cavusoglu, A. Sherman, and F. Tendick. Design of bilateral teleoperation controllers for haptic exploration and telemanipulation of soft environments. *Robotics and Automation, IEEE Transactions on*, 18(4):641–647, Aug 2002.

[4] S. Cotin, H. Delingette, and N. Ayache. A hybrid elastic model for real-time cutting, deformations, and force feedback for surgery training and simulation. *The Visual Computer*, 16(7):437–452, 2000.

[5] A. F. de C. Hamilton, D. W. Joyce, J. R. Flanagan, C. D. Frith, and D. M. Wolpert. Kinematic cues in perceptual weight judgement and their origins in box lifting. *Psychological Research*, 71(1):13–21, 2007.

[6] M. Fleury, C. Bard, N. Teasdale, J. Paillard, J. Cole, Y. Lajoie, and Y. Lamarre. Weight judgement: the discrimination capacity of a deafferented subject. *Brain*, 118(5):1149–1156, 1995.

[7] R. Goertz and W. Thompson. Electronically controlled manipulator. *Nucleonics*, 12(11):46, Nov 1954.

[8] R. C. Goertz. Mechanical master-slave manipulator. *Nucleonics*, 12(11):45–46, Nov 1954.

[9] J. Heinzmann and A. Zelinsky. A safe-control paradigm for human-robot interaction. *Journal of Intelligent and Robotic Systems*, 1999.

[10] P. F. Hokayem and M. W. Spong. Bilateral teleoperation: An historical survey.

[11] B. Rooks. The harmonious robot. *Industrial Robot*, 33(2):125–130, 2006.

[12] M. Turner, D. Gomez, M. Tremblay, and M. R. C. and. Preliminary tests of an arm-grounded haptic feedback device in telemanipulation. In *Proc. ASME WAM*, volume DSC-64, pages 145–149, 1998.

[13] M. L. Turner, R. P. Findley, W. B. Griffen, M. R. Kutkosky, and D. H. Gomez. Development and testing of a telemanipulation system with arm and hand motion. In *ASME IMECE Symposium on Haptic Interfaces*, 2000.

[14] R. Zhou, W. Hamel, A. Hariharan, and M. Noakes. Using the wam as a master controller. In *Proceedings of the ANS 2006 International Joint Topical Meeting*, 2006.