

Weston B. Griffin
GE Global Research
Niskayuna, New York
wgriffin@stanfordalumni.org

William R. Provancher
Department of Mechanical
Engineering
University of Utah
Salt Lake City, Utah

Mark R. Cutkosky
Department of Mechanical
Engineering
Stanford University
Stanford, California

Feedback Strategies for Telemanipulation with Shared Control of Object Handling Forces

Abstract

Shared control represents a middle ground between supervisory control and traditional bilateral control in which the remote system can exert control over some aspects of the task while the human operator maintains access to low-level forces and motions. In the case of dexterous telemanipulation, a natural approach is to share control of the object handling forces while giving the human operator direct access to remote tactile and force information at the slave fingertips. We describe a set of experiments designed to determine whether shared control can improve the ability of an operator to handle objects delicately and to determine what combinations of force, visual, and audio feedback provide the best level of performance and operator sense of presence. The results demonstrate the benefits of shared control and the need to carefully choose the types and methods of direct and indirect feedback.

I Introduction

The work described herein is part of a project to enhance the dexterity and sensitivity of dexterous telemanipulation (i.e., manipulation in which fine forces and motions are imparted with the fingertips). Shared control provides a framework for extending the capabilities of an immersive telemanipulation system.

A shared control telemanipulation system combines some of the autonomy of supervised systems (Sheridan, 1992) with the telepresence found in direct master-slave bilateral systems (Figure 1). There are advantages to having the robot hand take over force regulation and object manipulation when the task is sufficiently well defined. By providing local control of forces, stiffness, and fine hand motions, the robot allows the human supervisor to focus on the task itself, concentrating on the desired motions and behavior of the grasped object or tool. Time delays and limitations in the accuracy of haptic feedback through the master become less detrimental because commands from the master are supplemented by local control to prevent unwanted slips or object damage. However, there is some concern that the operator's sense of presence will be reduced as the slave system takes more control over the interaction. Moreover, there is an indication from the physiology literature (Westling & Johansson, 1984) that without appropriate tactile cues, the human grasp force will gradually relax, leading to a divergence between conditions at the master and slave. Shared control seeks to avoid these problems by keeping the operator in direct contact with tactile and force information sensed at the slave fingertips.

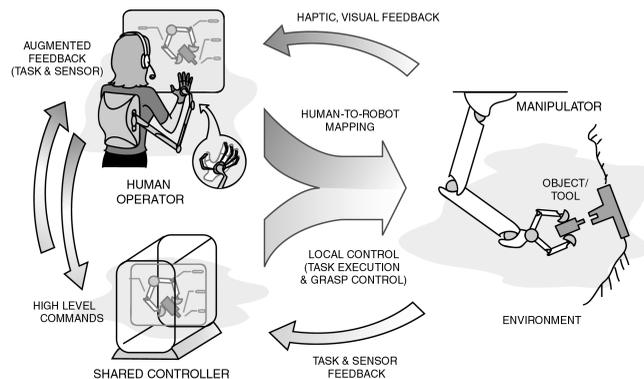


Figure 1. Shared control concept for dexterous telemanipulation.

In the experiment discussed in this paper, the operator and the robotic hand share control of the grasp force when handling an object. To prevent accidental drops, the robot can intervene and assume control over the internal force. However, the operator maintains the ability to override the dexterous controller to release the object, or grasp it tightly, if desired.

Several conditions were tested to investigate some of the issues associated with shared control, including the effects of robot task intervention, indirect feedback, and force feedback. In particular, we wished to determine whether task performance was improved if the operator was informed when robot intervention occurred. Additionally, because control was shared over parameters used for direct haptic feedback, different methods of force feedback were tested.

2 Previous Work

In cases where time delay is not a significant problem, shared control offers advantages over strict supervisory control by providing the operator with direct access to forces and motions at the slave. Many previous investigations have implemented a mix of hierarchical and shared control for telemanipulation.

Initial investigations focused on developing a framework for task-level sharing of motion trajectories for systems with moderate time delays (Hayati & Venkataraman, 1989; Oda et al., 1999). Other work focused on modify-

ing the impedance of slave-manipulators based on teleoperator commands and local sensor information (Hannaford, Wood, McAfee, & Zak, 1991; Backes, 1992).

Toward dexterous telemanipulation, Michelman and Allen (1994) applied the concept of shared control to the Utah/MIT dexterous hand. Their system focused on defining and sequencing primitives for operations such as grasping an object and inserting a peg in a hole. Researchers at NASA have developed Robonaut, a humanoid robotic system with a dexterous hand and telepresence interface. The control architecture is based on distributed control nodes that combine low-level intelligence for reflexive actions and high level commands for grasp configuration (Diftler, Culbert, Ambrose, Platt, & Buethmann, 2003). While these dexterous systems demonstrate the application of a shared control framework, none specifically addresses the issues of haptic feedback at the fingertip level.

Of particular relevance to our experiment, Hannaford et al. (1991) evaluated a six-axis generalized teleoperation system with arm/hand force feedback. Along with evaluation of the force feedback, a case was tested in which control was shared with the robot (utilizing local force/torque sensing) during a peg-hole insertion task. In this task, the operator controlled the end-effector position while task-space orientation control was shared with the robot. The authors observed a reduction in task completion time and sum-of-squared forces with the addition of shared control.

3 Experimental Setup

The dexterous telemanipulation system includes an arm-mounted master, a controller, and a slave consisting of an industrial robot and a two-fingered hand. Although small time delays must be considered when tuning the system, the treatment of large delays is beyond the scope of this paper. In addition, the operator has direct visual and aural feedback from the slave, located across the room from the master.

3.1 Master System

Human finger motions are recorded by an instrumented glove (Immersion CyberGlove) and wrist mo-

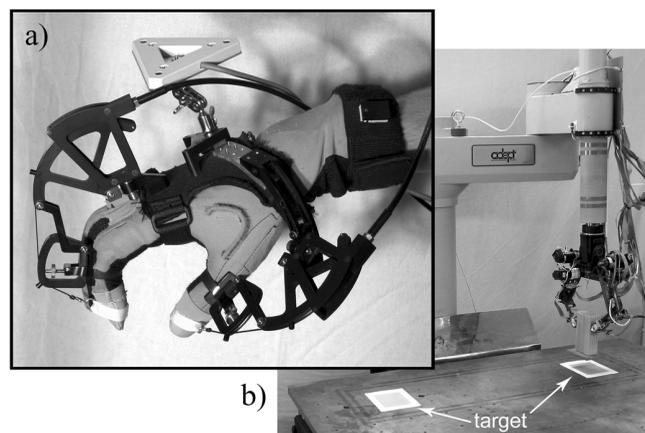


Figure 2. (a) Master interface with instrumented glove, force feedback exoskeleton, and ultrasonic wrist tracker. (b) Slave system with dexterous robotic hand mounted to industrial robotic arm.

tions are recorded using a six-DOF ultrasonic tracking system (Logitech Head Tracker). Calibration software developed in previous work (Griffin, Findley, Turner & Cutkosky, 2000) allows the telemanipulation controller to estimate the intended motions of a virtual object grasped between the operator's thumb and index finger. The glove and tracker signals are sampled at 200 Hz and 50 Hz, respectively, and smoothed to generate motion commands for the slave (Turner, Findley, Griffin, & Cutkosky, 2000; Griffin, 2003).

A cable-driven exoskeleton (Immersion CyberGrasp) provides a single degree of freedom of force feedback at each fingertip (Figure 2a). Additional feedback channels, including audio tones and visible LEDs on the robot hand, display state information computed by the telemanipulation controller.

3.2 Slave System

The slave system has a two-fingered hand with two degrees of freedom per finger. The hand is designed specifically for use with force and tactile sensors and therefore has a smooth back-driveable cable transmission to minimize friction and vibrations (Figure 2b). For the experiments in this paper, the fingertips were

equipped with sensors to measure normal and tangential contact forces and contact locations.¹

The robotic hand is mounted to an industrial five-axis SCARA robotic arm. Special-purpose software allows real-time control of the robot trajectory at 63 Hz update rates via ethernet.²

3.3 Control Framework

To accommodate the various communication rates and priorities associated with different system components, the controller was developed on a real-time operating system (QNX) with a multiprocess structure implemented on two networked nodes.

The control laws for the slave manipulator lay the foundation for implementing a shared control telemanipulation system. We are interested in both independent control of the fingers (for exploration) and coordinated control for manipulating grasped objects. Following the approach of Hyde (1995) and Schneider and Cannon (1989), the dynamics are computed using Khatib's (1987) operational space formulation, with Hogan's (1985) impedance control to specify the behavior with respect to disturbances. The impedance control loop runs at 1 kHz and the hand has a closed-loop bandwidth of approximately 10 Hz for small motions of the grasped object.

The transition between independent and coordinated manipulation is based on signals from force and tactile sensors, following the approach of Hyde (1995). For manipulation of a grasped object, the controller must first map the operator's finger motions to the corresponding motions of a virtual object held in the hand. The controller then computes the motions required of the (non-anthropomorphic) robot fingers, including rolling kinematics, to achieve a geometrically similar motion of the actual object (Griffin et al., 2000; Griffin, 2003).

Figure 3 illustrates the control system diagram for the cooperative control method. Separate control laws are

¹http://bdml.stanford.edu/touch/tele_projects/res_tacsens.htm

²Adept Technologies, Inc. V+ 12.4 OS with Enhanced Trajectory Control and Advanced Servo Library.

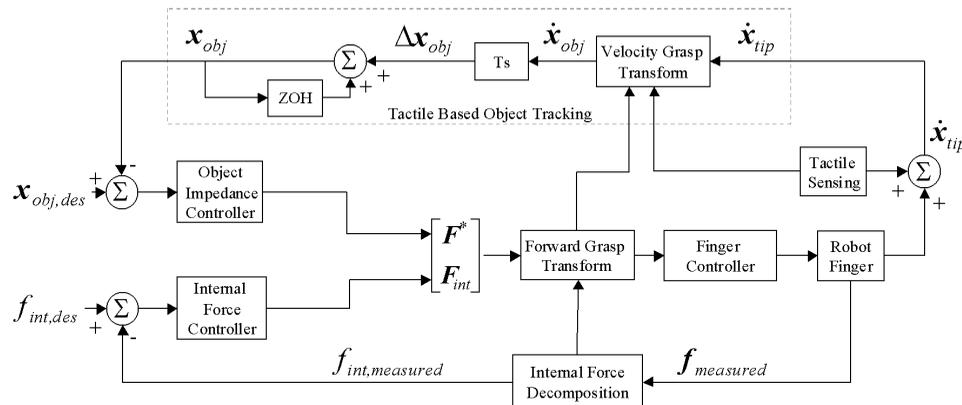


Figure 3. Control diagram for object impedance cooperative control framework.

used for controlling the (external) impedance forces and the (internal) grasp force on the object. The impedance forces are generated based on the difference between the current (modeled) object position and the operator's mapped virtual object position. The desired object positions are equilibrium points for a virtual spring and damper attached to the object; desired orientations are equilibrium points for a virtual torsional spring-damper. The internal force applied to the object is regulated based on the operator's desired internal force and on the actual internal force as measured by force sensors in the robot hand.

The internal and external forces are concatenated and multiplied by the forward grasp transform to determine the appropriate fingertip forces (Mason & Salisbury, 1985). The forward grasp transform is formulated and updated using finger/object contact geometry and provides a mathematical construct to determine the necessary fingertip forces required to create a net force on a grasped object (Mason & Salisbury, 1985). The operational space formulation is then used to generate appropriate torques for each finger link based on the desired fingertip forces, accounting for the dynamic properties of the finger mechanisms.

The upper portion of control diagram (Figure 3) illustrates the object-tracking algorithm (partially based on the approach of Maekawa, Tanie, and Komoriya [1995]) and is used to track and update the modeled object location. Using tactile sensors at the fingertips to

detect contact location, the control method sums differential object motions computed by multiplication of the velocity grasp transform and fingertip contact point velocities. The velocity grasp transform is derived directly from the forward grasp transform. The advantage of this contact tracking method lies in its ability to manipulate unknown objects, thus removing the need for having an object model.

Utilizing force and tactile sensors and the kinematics of contact, the cooperative control algorithm can stably manipulate an object through arbitrary small translations and rotations (i.e., until regrasping becomes necessary).

4 Experiment Description

Using the telemanipulation system, operators were instructed to pick up and carry an object across the workspace and set it down on a designated target (Figure 2b). The operators were asked to treat the object as fragile and thus to use a minimum grasp force, while taking care not to let the object drop. The object was a 0.2 kg wood block and was moved between targets separated by 0.7 m.

To assist the operator, the controller must have an estimate of the minimum grasp force. When grasping objects with their own hands, humans readily identify the minimum force required to prevent slipping and

generally maintain a safety margin of 10–30% (Westling & Johansson, 1984). For the robot, we estimate the minimum internal force based on a priori friction estimates and continuous measurements of the normal and tangential forces at the robot fingertips. Recall that in a multifingered grasp, the contact forces can be decomposed into f_{ext} , which balance the object weight, inertial forces, and contact with the environment, and f_{int} , which produce no net resultant and can be adjusted independently to prevent slipping (Yoshikawa and Nagai, 1991). For a two-fingered grasp on a block that is held approximately level, the minimum internal force becomes:

$$f_{int,min} = \max\left(\frac{f_{tan,i}}{\mu_i}\right) \quad (1)$$

where μ_i is the static friction estimate for the i th finger and $f_{tan,i}$ is the tangential force component. With this information, the robot is able to regulate the internal grasp force, using a PI control law, independent of external forces on the object.

We are interested in evaluating shared control as compared to unassisted telemanipulation and in determining what kinds of feedback are most useful to the operator. Accordingly, we examined various cases in which one or more of the following conditions applied.

- *Direct bilateral control (baseline case).* The desired grasp force from the operator (expressed as a reduction in the virtual distance between the fingertips, following the impedance control formulation [Hogan, 1985]) is used directly to control the grasp force on the object. The magnitude of the measured grasp force at each finger is fed directly back to the operator via CyberGrasp.
- *Robot assisted control.* When the desired grasp (internal) force drops below 110% of the minimum force in Eq. (1), the robot controller intervenes to maintain it at 110%, until the operator releases the object (desired grasp force < 0). When robot intervention is active, the force fed back to the user can be either the magnitude of the *actual* grasp force measured at the fingertips or proportional to the *commanded* grasp force.

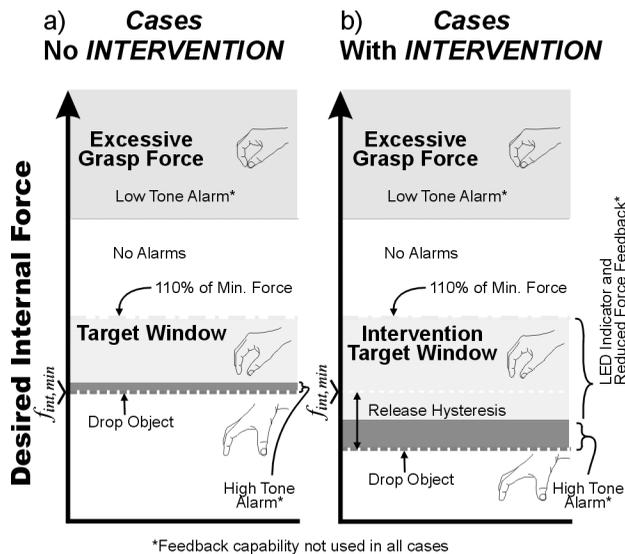


Figure 4. Comparison plots of feedback to operator for cases with and without intervention.

- *Visual indicators.* When robot intervention is enabled, LEDs on the robot hand are illuminated when the controller is actively maintaining the internal force at 110% of $f_{int,min}$.
- *High frequency audio tone.* A 500 Hz warning is sounded when the object is in danger of slipping. For unassisted telemanipulation, the tone is sounded whenever the desired grasp force approaches $f_{int,min}$ (at 101% of the minimum grasp force). If intervention is active, an audio tone is emitted when the desired grasp force approaches zero, which would trigger the robot to release the object.
- *Low frequency audio tone.* To discourage the user from squeezing the object too hard, a 50 Hz tone can be emitted when the desired grasp force exceeds 170% of the minimum required.

The audio and LED feedback channels create a target window in which the user can remain for safe, gentle object handling. The target windows for manipulation with and without robot intervention are depicted in Figure 4. Note that the target window with robot intervention (Figure 4b) is wider because the desired grasp

Table 1. Tested Effects for Experimental Cases

Case	Aid			Reduced force feedback
	Audio alarms	Robot intervention	LED indicator	
Case 1	No	No	No	No
Case 2	Yes	No	No	No
Case 3	No	Yes	No	No
Case 4	Yes	Yes	Yes	No
Case 5	No	Yes	No	Yes
Case 6	Yes	Yes	Yes	Yes
Case 7	No	Yes	Yes	Yes

Note: Audio alarms are sounded when subject force is too low or too high. Robot intervention occurs when the force commanded by the subject is too low. The LED indicator lights up during cases of robot intervention. Reduced force feedback occurs during robot intervention when the force commanded by the subject is too low.

force can drop below $f_{\text{int,min}}$ without adverse consequences.

The various cases tested, and the combinations of feedback associated with each case, are listed in Table 1.

5 Procedure

Eleven subjects, eight males and three females, were recruited for the experiments. Each subject was required to complete the experiment in two sessions. The first session was used to calibrate and adjust the human-to-robot mapping parameters (see Griffin et al., 2000) and familiarize the operator with controlling the robotic arm and hand with visual and force feedback, and required approximately one hour to complete.

The second session occurred two to four days after the first. The subjects were refamiliarized with the system and initially instructed to perform the pick-carry-and-place task under the control case, Case 1. Prior to testing, each case was explained using a graphic similar to Figure 4.

Subjects were asked to transport the block from target to target using the minimum force necessary without dropping it. The task was marked a failure if the block was dropped or not placed on the target. Subjects were told that trial completion time was not considered in evaluating performance.

Each subject was asked to complete four trials for each of the seven different cases. The case order for each subject was randomized. At the start of each new case, subjects were given a single practice trial prior to the four test trials.

During the test trials, the following data were recorded at 200 Hz by the computer: time, the measured internal force on the object, the operator's commanded internal force, the calculated minimum internal force per Eq. (1), the intervention state, and the states of any audio alarms and LED indicators. The experimenters manually recorded failures. On average, subjects took one and a half hours to complete the test trials in the second session.

6 Results

The results of the tests were examined to determine which combinations of control and which forms of feedback improved task performance in terms of minimizing the internal force and reducing failures.

Plotting key variables for each case immediately reveals some trends. Figure 5 shows typical trials of one subject for Cases 1, 2, and 6. For every case, the measured internal force closely tracks the operator's desired (commanded) internal force, unless intervention is active. For Case 1 (direct telemanipulation), the calculated minimum internal force is also shown. This force is approximately 1.7 N, with minor variations due to vertical accelerations of the block. Most of the time, the operator's commanded force is considerably above the minimum.

For Case 2 (alarms only), the effects of adding the audio tones are seen. Initially, the operator utilizes an excessive grasp force, which is gradually relaxed until the low frequency tone is no longer heard. The operator then uses the high frequency tone to maintain the grasp force above the minimum requirement (high frequency tone triggered at 101% of $f_{\text{int,min}}$).

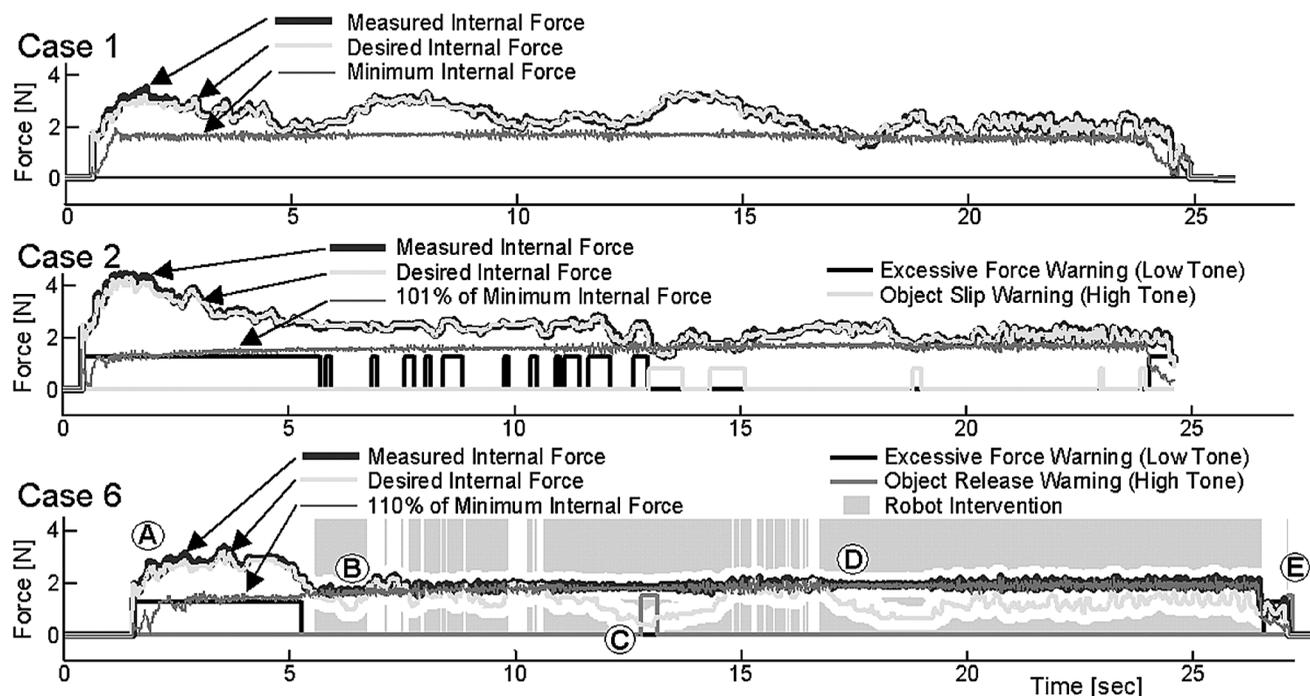


Figure 5. Typical subject data recorded during a single task trial for Cases 1, 2, and 6.

For Case 6 (robot intervention with alarms and LEDs) a trace is plotted corresponding to the 110% threshold at which the robot assumes control of the grasp force. Not surprisingly, the measured internal force tracks this value closely. As in Case 2, the operator initially applies an excessive force (marker A) and then reduces the force, allowing the robot to assume internal force control (B). However, the operator slowly continues to relax (consistent with predictions from Westling and Johansson [1984]) until the high frequency tone alarm sounds (C) to warn that the object may be released. The operator adjusts the grasp force and lets the robot continue to intervene (D) until the object is released at the target (E).

6.1 Objective Data Analysis

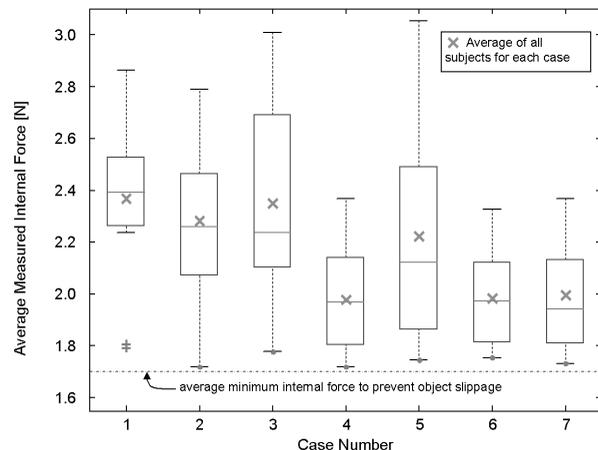
The objective data analysis is based primarily on the measured internal force applied to the object. Since the goal is to handle the object gently, the measured internal force is a logical performance metric. The mean

and standard deviation of the force for each case over all subjects is given in Table 2.

Figure 6 shows a boxplot of the average measured internal force for each case based on an average of successful trials for each subject. There is clearly a reduction in the measured internal force when comparing all cases to the control case (Case 1). In particular, Cases 4, 6, and 7 have a distribution that is much lower than Case 1. It should be pointed out that in theory it is possible to complete the task under Cases 1 and 2 with a lower force than cases with intervention because, during intervention, the lowest force the robot can apply is 110% of the minimum. To determine whether differences among the cases are statistically significant, an analysis of variance (ANOVA) was performed. A single factor, balanced ANOVA test with seven fixed effects was first run on the average measured internal force of each subject and for each case (yielding 11 data points per case with 76 DOF). The null hypothesis was that the different cases had no effect on the measured inter-

Table 2. Mean and Standard Deviation of Measured Internal Force of All Subjects for Each Case

Case	1	2	3	4	5	6	7
Average Force (N)	2.3	2.2	2.3	1.9	2.2	1.9	2.0
SD of Force (N)	0.3	0.3	0.4	0.2	0.4	0.1	0.2

**Figure 6.** Boxplot showing medians, quartiles, and outliers of subjects' averaged measured internal force applied to the object.

nal force. The ANOVA test results in a p value of .003 [$F(6,70) = 3.71$]. Thus we can conclude that at least two cases have a statistically different mean.

To determine specifically which cases are different, we must apply a multi-comparison procedure. For this we apply Dunnett's method, which is designed for the comparison of several effects to a control effect (Case 1) while limiting the possibility of a Type I error to the desired significance level ($\alpha = 0.05$) (Devore & Farnum, 1999). Applying this method, we can state with 95% confidence that Cases 4, 6, and 7 have a mean different than Case 1. From the averages in Table 2, we see a reduction of internal force on the order of 15% for these cases.

Even though subjects were informed that task completion time was not a factor, the task time may reveal information about the mental or physical difficulty associated with completing the task under the various conditions. An ANOVA test was performed using the 11

subjects' averaged trial times (excluding failures) for each case. The analysis resulted in a p value of .82 [$F(6,70) = 0.48$], indicating that the mean task completion times for the cases are not statistically different. While there is no improvement in task time for cases with shared control, there is importantly no increase in task time either.

Figure 7 shows the total number of failures for each case. With eleven subjects and four trials per case, each case was attempted 44 times. The numbers of failures for Cases 5 and 6 are the lowest and, interestingly, Case 7, which is identical to Case 5 except for the addition of LEDs, had the highest number of failures. Also note that no single subject contributed a disproportionate number of failures.

6.2 Discussion of Objective Analysis Results

Taking into account the task goals and the objective performance criteria based on measured internal force, task completion time, and number of failures, it is clear that the addition of a dexterous shared controller to a traditional bilateral telemanipulation system can enhance an operator's performance during a typical telemanipulation task.

A comparison of the measured internal force for Case 1 versus Case 2 indicates that warning the operator of a possible failure through audio feedback may be helpful. However, during preliminary testing we found that alarms could cause significant confusion for the operator if the activation levels were set too close relative to each other. Part of the advantage of adding robot intervention is that it allows us to make the target window wider, separating the conditions associated with the high and low audio tones.

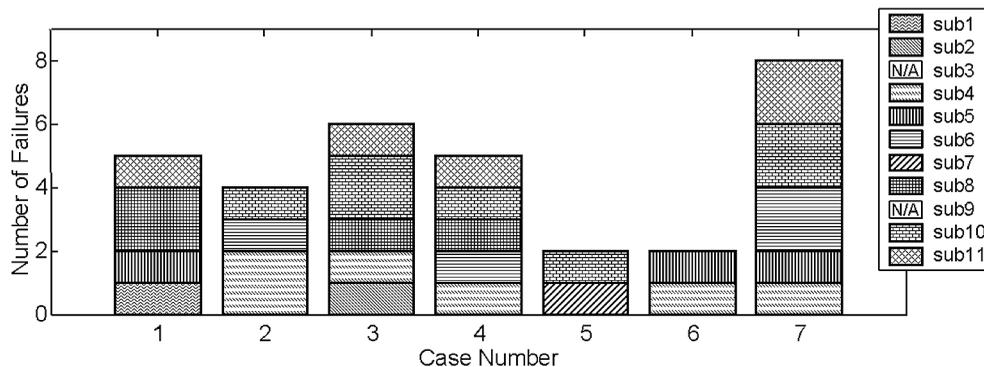


Figure 7. Total number of failures for each case for all subjects.

As anticipated, we found that robot intervention could improve task performance. However, the presence and type of direct and indirect feedback had a marked effect. The cases in which we informed the operator that the intervention was occurring (Cases 4, 6, and 7) had lower forces than the control (Case 1), whereas in Cases 3 and 5 (no indication of intervention), the average internal force was similar to Case 1. Moreover, if we examine the number of failures, we find that simply informing the operator that intervention was occurring, using LEDs as a visual indicator, was not adequate. The number of failures in Case 7 was four times that of Case 6, indicating that the audio alarms, particularly the high frequency tone alarm, reduced the number of failures.

The effects of two different approaches to force feedback are isolated in a comparison of Case 3 to Case 5. Neither of these cases had audio alarms or LEDs. In Case 3, the actual grasp force, based on measurements at the robot fingertips, is fed back to the operator. This force remains nearly constant when intervention is active. In Case 5, the force relayed to the operator is based upon the operator's commanded (desired) internal force. As seen in Table 2, the subjects used slightly less internal force and experienced three times fewer failures in Case 5 (see Figure 7). While the 6% force reduction is not statistically significant, the reduction in dropped objects indicates that it is useful to feed back forces proportional to the operators' commanded force, even as the robot holds the actual grasp force constant. In accord with the operator's expectations, the reduc-

tion in displayed force as the grasp is opened provides a haptic cue useful for manipulation.

Based on the objective data analysis and the performance criteria, Case 6, which combines robot intervention, audio alarms, LED indicators, and reduced force feedback, provides the best overall performance compared to the bilateral control case. Further analysis investigating the possible effects of fatigue and learning is presented in the appendix. However, based on this analysis, no evidence of learning and fatigue was found in our study.

6.3 Subjective Data Analysis

In addition to having the subjects complete the specified task to test the shared control system, a post-experiment questionnaire was administered to obtain qualitative data on the cases tested. Because the human operator is an integral part of the shared telemanipulation system, the expressed preference of the operator is an important parameter in assessing the overall effectiveness of a given case.

Subjects were asked to rank each case tested based on preference. The ranking was on a scale from negative two to positive two, corresponding to "disliked" and "preferred," respectively, with zero being "indifferent." Figure 8 shows the average ranking score for each case, for all subjects, with the errors bars indicating two standard deviations. From the figure, we can see that Cases 6 and 7 were most preferred. A simple *t*-test analysis was

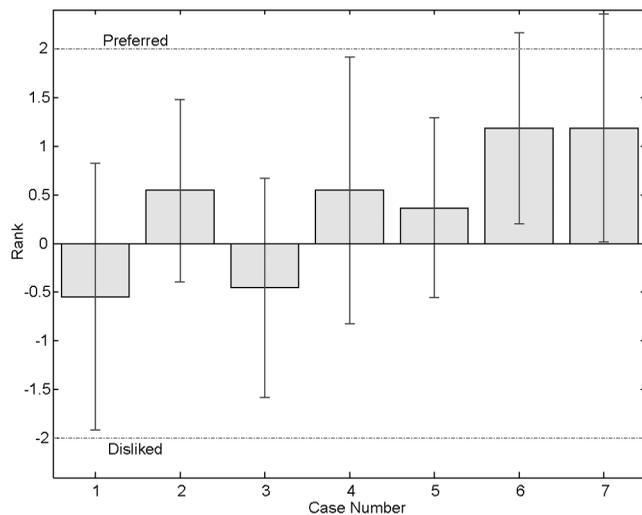


Figure 8. Average ranking of each case in terms of subjects' expressed preference (error bars represent two standard deviations).

performed to determine which cases had a mean preference ranking significantly greater than zero. At the 95% confidence level (p value $< .05$), Cases 2, 6, and 7 have a ranking greater than zero, indicating that subjects preferred these cases. While clearly some subjects preferred Case 2, the average ranking is much lower than for Cases 6 and 7 (0.55 as compared to 1.18 and 1.18, respectively). Interestingly, the two cases that subjects disliked the most were the same cases that had the highest applied internal force on the object, that is, Cases 1 and 3.

Figure 9 shows the percent difference in mean internal force compared to Case 1, and Figure 10 shows the average residuals based on order of task completion. Figures 9 and 10 will be discussed further in the appendix.

7 Conclusions

The performance of subjects during a shared-control fragile object handling task was evaluated through an objective data analysis. Utilizing results from a statistical analysis of each subject's measured internal force for each case and the number of failures for each

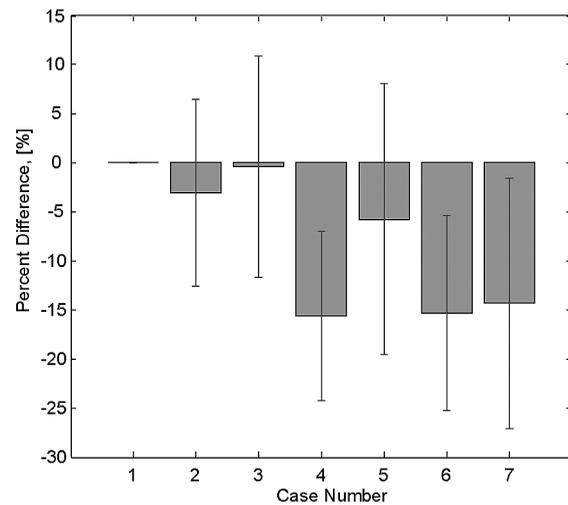


Figure 9. Average of each subject's percent difference in mean internal force for each case as compared to Case 1 (error bars represent two standard deviations).

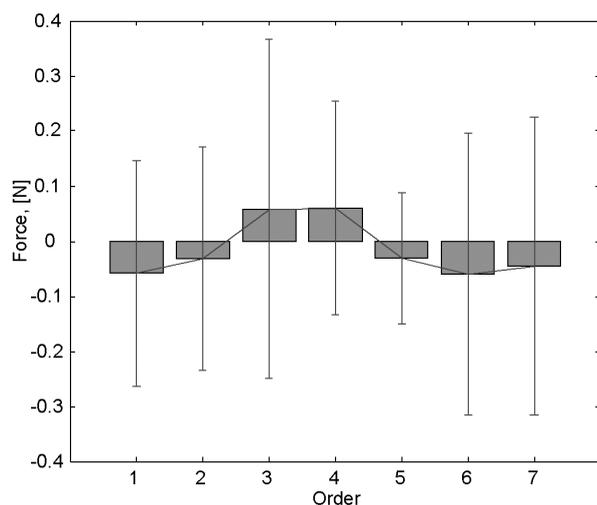


Figure 10. Average model residuals for all subjects based on order of task completion (error bars represent two standard deviations).

of the given cases, several conclusions can be drawn. We found that the addition of a shared control framework could improve an operator's ability to delicately and securely handle an object compared to direct telemanipulation. In comparing the performance of all seven cases, we further found that it is necessary to:

- inform the operator when the intervention is active (in other words, it is necessary to let the operator know that the robot has assumed control);
- inform the operator of impending state changes (in particular, inform the operator that the robot may release an object in its grasp if the operator's commands continue to diverge from the robot's commands); and
- feed back forces to the operator based on the operator's commanded force rather than feeding back the actual forces as measured by the robot during robot intervention.

Acknowledgments

We thank Dr. H. Rauch for his suggestions regarding the model analysis to evaluate possible learning and fatigue effects.

References

- Backes, P. G. (1992). Multi-sensor based impedance control for task execution. *Proceedings of the IEEE International Conference on Robotics and Automation*, 2, 1245–1250.
- Devore, J. L., & Farnum, N. R. (1999). *Applied statistics for engineers and scientists*. Pacific Grove, CA: Duxbury Press.
- Diftler, M. A., Culbert, C. J., Ambrose, R. O., Platt, R., Jr., & Buethmann, W. J. (2003). Evolution of the NASA/DARPA robonaut control system. *Proceedings of the IEEE International Conference on Robotics and Automation*, 2, 2543–2548.
- Griffin, W. B. (2003). *Shared control for dexterous telemanipulation with haptic feedback*. Ph.D. Thesis, Stanford University, Department of Mechanical Engineering.
- Griffin, W. B., Findley, R. P., Turner, M. L., & Cutkosky, M. R. (2000). Calibration and mapping of a human hand for dexterous telemanipulation. *Proceedings of the ASME International Mechanical Engineering Congress and Exposition, Dynamic Systems and Controls*, 69, 1145–1152.
- Hannaford, B., Wood, L., McAfee, D. A., & Zak, H. (1991). Performance evaluation of a six-axis generalized force-reflecting teleoperator. *IEEE Transactions of Systems, Man, and Cybernetics*, 21, 620–633.
- Hayati, S., & Venkataraman, S. T. (1989). Design and implementation of a robot control system with traded and shared control capability. *Proceedings of the IEEE International Conference on Robotics and Automation*, 3, 1310–1315.
- Hogan, N. (1985). Impedance control—An approach to manipulation. 1. Theory. *Journal of Dynamics Systems, Measurement, and Control—Transactions of the ASME*, 107, 8–16.
- Hyde, J. M. (1995). A phase management framework for event-driven dextrous manipulation. Ph.D. Thesis, Stanford University, Department of Mechanical Engineering.
- Khatib, O. (1987). Unified approach for motion and force control of robot manipulators: The operational space formulation. *IEEE Journal of Robotics and Automation*, 3, 43–53.
- Maekawa, H., Tanie, K., & Komoriya, K. (1995). Tactile sensor based manipulation of an unknown object by a multifingered hand with rolling contact. *Proceedings of the IEEE International Conference on Robotics and Automation*, 1, 743–750.
- Mason, M. T., & Salisbury, J. K. (1985). *Robot hand and the mechanics of manipulation*. Cambridge, MA: MIT Press.
- Michelman, P., & Allen, P. (1994). Shared autonomy in a robot hand teleoperation system. *Proceedings of the IEEE/RSJ/GI International Conference of Intelligent Robotics and Systems*, 1, 253–259.
- Oda, M., Inaba, N., Takano, Y., Nishida, S., Kayashi, M., & Sugano, Y. (1999). Onboard local compensation on Ets-space robot teleoperation. *Proceedings of the IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, 701–706.
- Schneider, S., & Cannon, R. H., Jr. (1989). Object impedance control for cooperative manipulation: Theory and experimental results. *Proceedings of the IEEE International Conference on Robotics and Automation*, 2, 383–394.
- Sheridan, T. B. (1992). *Telerobotic, automation, and human supervisory control*. Cambridge, MA: MIT Press.
- Turner, M. L., Findley, R. P., Griffin, W. B., & Cutkosky, M. R. (2000). Development and testing of a telemanipulation system with arm and hand motion. *Proceedings of the ASME International Mechanical Engineering Congress and Exposition, Dynamic Systems and Controls*, 69, 1057–1063.
- Westling, G., & Johansson, R. S. (1984). Factors influencing the force control during precision grip. *Experimental Brain Research*, 53, 277–284.
- Yoshikawa, T., and Nagai, K. (1991). Manipulating and grasping forces in manipulation by multifingered robot hands. *IEEE Transactions of Robotics and Automation*, 7, 67–77.

Appendix: Learning and Fatigue Effects

To address the question of whether learning or fatigue could have influenced the results, a model was developed to describe the observed data. In this way, a small set of parameters was used to develop a model equation and then the model was used to estimate the actual data. Using the model estimates, we can analyze the residuals of the model prediction (as compared to the observed data) to see how “good” the model is and if there are any additional parameters, perhaps arising from learning or fatigue, that should be added.

To investigate the differences in performance for each of the cases, an initial model was developed that describes the data with a single parameter representing each case. Thus, case-to-case comparisons can easily be made with a comparison of the describing parameter (assuming the model accurately describes the actual data).

There are many possible ways to describe the differences in the observed case data. We chose to use the averaged percent difference in measured internal force for each case as compared to Case 1. The percent difference, based on each subject’s mean internal force for each of the noncontrol cases, is computed as follows:

$$\phi_{(i,j)} = \frac{\bar{f}_{\text{int}(i,j)} - \bar{f}_{\text{int}(1,j)}}{\bar{f}_{\text{int}(1,j)}} \quad (2)$$

where $\bar{f}_{\text{int}(i,j)}$ is the mean measured internal force for case i ($i = 1, \dots, 7$) and subject j ($j = 1, \dots, 11$) and $\phi_{(i,j)}$ is the percent difference for Cases 2–7 compared to Case 1 for each subject ($\phi_{(1,j)}$ will always equal zero). The percent difference parameter was specifically chosen to help normalize the data and minimize the effects of subject-to-subject variability in the model parameters. Each subject’s percent difference for a given case is then averaged to form a single parameter, that describes the given case in a meaningful way:

$$\Phi_i = \sum_{j=1}^N \frac{\phi_{(i,j)}}{N} \quad (3)$$

where N is the total number of subjects. The averaged percent difference parameters can then be multiplied by

each subject’s Case 1 mean internal force to formulate an estimate of observed data. Thus, our single parameter based model equation is:

$$\bar{f}_{\text{int,estimate}(i,j)} = (1 + \Phi_i) \cdot \bar{f}_{\text{int}(1,j)} \quad (4)$$

for $i = 1, \dots, 7$.

To compare the cases, we analyze the mean percent difference, our model parameter Φ_i (and the standard deviation of Φ_i), for all subjects for each case. From Figure 9, the averaged percent difference in internal force for Case 6 is roughly -15% . If we assume the data are normally distributed, the empirical rule that 95% of the data falls within two standard deviations of the mean tells us that the Case 6 mean is significantly lower than zero, thus the percent difference in force is significantly less than Case 1 (this also holds for Cases 4 and 7), which matches our earlier statistical results.

To determine how well the model describes the data, we can perform an analysis of the model residuals. The residuals are based on the difference between the observed data (mean internal force) for each subject for each case, $\bar{f}_{\text{int}(i,j)}$, and the estimated value for the mean internal force for each subject for each case computed using Eq. (4):

$$\text{residual}_{(i,j)} = \bar{f}_{\text{int}(i,j)} - \bar{f}_{\text{int,estimate}(i,j)} \quad (5)$$

If all the residuals equal zero, then the model perfectly describes the data set. For this model, the mean of all the residuals is -0.01 N with a standard deviation of 0.24 N, both fairly small values indicating a good model. However, investigating potential sources of trends in the residual data can improve the model and create a better description of the observed data.

To determine if learning or fatigue plays a significant role in the observed data, we can group the residuals according to task order performed by each subject. Figure 10 shows a plot of the means of the residuals for all cases that occurred first, second, third, and so on. Also shown are error bars of two standard deviations. Clearly, the standard deviations are much larger than the mean residuals values, thus we can confidently state that learning or fatigue effects were not significant.