

Haptic Constraints for a Bimanual Command Selection and Location Task

Thomas N. Smyth*

Arthur E. Kirkpatrick†

School of Computing Science
Faculty of Applied Science
Simon Fraser University
Burnaby, BC, Canada

ABSTRACT

Haptic technology has the potential to make many everyday computing tasks faster, more precise, and more enjoyable. The common interaction task of command selection and location seems a suitable candidate for such improvement. Previous bimanual interaction techniques for this task have enjoyed relative success, but well designed unimanual techniques have proven faster. This work investigates the addition of haptic feedback to a bimanual technique for this task. We compared performance of 12 participants on two implementations of a bimanual interaction technique for this task. One implementation provided only graphical feedback, while the other additionally provided haptic feedback for the non-dominant hand. The haptic condition produced faster performance. However, the improvement was due more to faster movement times for the non-dominant hand rather than simultaneous movement. We discuss our results in the context of the motor control literature, and review the implications of our findings for future haptic interface designs.

Keywords: haptic interaction, bimanual interaction, tool palette

1 INTRODUCTION

Although haptic feedback is new to computer interfaces, it has long been an essential part of automobile controls. When shifting gears with a standard transmission, one hand controls the car's direction, while the other hand changes the gear. The path of the gear shift features a series of channels, and the driver uses the feel of these channels to change gears without removing the gaze from the road. With practice, the coordinated use of the two hands in a shifting motion becomes fluid. The workload for the gesture becomes low enough that it can be performed while navigating through traffic and conversing with a passenger.

Typical computer interfaces, which lack haptic display, are limited by comparison with a car's interface. All movements must be visually monitored, because the controls provide no haptic feedback about the current state. Imagine driving a car whose standard transmission was controlled by a typical (non-haptic) joystick: You would have to take your eyes off the road every time you changed gears. This is a particularly bad idea when travelling at 100 km/hr. Instead, the haptic feedback of gear shifts is designed so that a practiced driver can shift without looking at the control.

In this study, we investigate using a haptic interface to provide interaction with graphical user interfaces in a manner similar to shifting a car's gears. We focus on a common interaction task called *command selection and location*, wherein users must select both a

command and a location for the command to be applied. For example, to flood fill an enclosed region of a drawing, the user must specify both the command (flood fill) and the location (a point within the region to be filled). Note that shifting a car's gears may also be considered an instance of this task: The steering wheel controls the car's location, and the gear shift performs command selection. In computer applications, the command selection and location task is currently implemented by the *tool palette* interaction technique found in many popular design applications.

Although easily learned, the tool palette technique arguably does not support high levels of skilled performance, because the user must constantly shift both visual attention and the cursor from the current work area to the palette and back. An appealing alternative is to assign the two subtasks to separate hands: One hand can point to the location, while the other specifies the command. This introduces the potential for parallel activity, improving performance. Bier et al. [3] proposed an interaction technique dubbed *Toolglass* that does exactly that. Kabbash, Buxton, and Sellen [10] showed that Toolglass consistently outperformed conventional unimanual techniques in a command selection and location task.

However, in a later study, Guimbretière, Martin, and Winograd [7] tested the Toolglass technique against two unimanual techniques that support higher performance for the task than the conventional tool palette. Both unimanual techniques performed substantially better than Toolglass.

The sub-optimal performance of the Toolglass technique found by Guimbretière et al. may have been due to the difficulty of controlling both hands in an asymmetric task such as command selection and location. Users must either monitor both hands with their visual sense, attending to two distinct areas of the screen, or they must rely on their demonstrably inaccurate sense of proprioception to guide one hand, while visually monitoring the other. In both cases, the actions cannot be grouped into a single unit, simultaneity is sacrificed, and performance suffers.

Haptic interface technology offers a third alternative, whereby the position of one of the hands is monitored using haptic feedback. The problematic load on the visual system is then relieved, and vague proprioception gives way to a more concrete perception of force. Users may therefore be able to more easily organize the movements of the two hands as a single unit, resulting in greater simultaneity.

Unfortunately, this alternative is not well studied. As described in section 5, previous efforts to add haptic feedback to interaction with GUIs have concentrated mainly on the dominant hand, and have usually consisted of direct haptic augmentation of existing graphical interfaces. Our goal in this work is to investigate if and how users can make use of haptic feedback in a bimanual interaction technique.

We compared two implementations of an interaction technique for a simple command selection and location task: one implementation provided simple graphical feedback for both hands, while the other additionally provided force feedback for the non-dominant hand.

*e-mail: tsmyth@cs.sfu.ca

†e-mail: ted@cs.sfu.ca

2 METHODS

2.1 Task

At the start of each trial, a small circle containing a number from 1 to 8 appeared at a random location on the screen. Participants were required to select the appropriate number with their non-dominant hand using a Phantom haptic device (command selection), and to point to and click within the circle with their dominant hand (location selection), using a standard mouse. We used numbered commands rather than more common types such as colours or font styles, because the ordering of numbers is well-practiced and therefore participants could more quickly learn the mapping of numbers to octants. This allowed us to evaluate skilled performance on the task with less training than would have been required if the commands had no intrinsic ordering, such as colours.

The dominant hand mouse movement was the same for all conditions, while the non-dominant hand Phantom movement varied between three conditions.

The working volume of the Phantom was divided into eight octants about the x , y , and z axes. Octants were assigned the numbers 1 to 8. To select a given number, participants placed the Phantom's stylus tip (hereafter referred to as the probe) anywhere within the corresponding octant.

The experiment compared three variations of this number selection task. The first condition was named force feedback (F), in which the user's movements were restricted to the interior of a virtual cube rendered by the Phantom. The cube had a width of 2 cm, and was centered about the origin of the space described above.




Figure 1: The graphical Phantom position display provided in the GA and GR conditions. The upper marker indicated the x and y coordinates of the Phantom, while the lower marker indicated its z coordinate.

In the second condition, the haptic rendering of the cube was removed. Instead, participants were issued visual feedback as to their position within the cube with a small display as shown in Figure 1. The upper marker indicated their x and y coordinates, while the lower marker indicated their z coordinate. If participants moved outside the cube in any direction, the corresponding marker stopped at the edge of the indicator, just as a conventional mouse pointer stops at the edge of the screen. The background colour of the indicator also changed from dark gray to light gray when the probe crossed the plane $z = 0$. We refer to this condition as graphical-absolute, and denote it GA. This condition emphasizes the accuracy of the user's sense of proprioception, their knowledge of the absolute location of their limb in space. The more accurately they know where their limb is located, the more easily they can select a command without looking at the cursor location for that limb.

In both the F and GA conditions, the origin of the octant system was fixed to the origin of the Phantom's working volume. In the third condition, called graphical-relative (GR), the octant system origin shifted whenever participants moved the probe against the bounds of the cube. In other words, participants dragged the cube around with the probe. This is similar to the behaviour of a

conventional mouse, whereby the origin of the screen coordinate system shifts with respect to the physical table coordinate system whenever the pointer moves against the edge of the screen. This design emphasizes the accuracy of the user's sense of kinesthesia, their knowledge of the distance and direction of their latest movement. The more accurately they can plan a movement from their current resting location, the more easily they can select a command without looking at the cursor.

No graphical feedback was issued in F, and octants were not labeled in GA or GR, in order to bring about skilled performance in participants with as little practice as possible. Pilot studies in which octants were numbered, and in which graphical feedback was issued in all conditions, showed that participants would tend to use the graphical display as a crutch instead of learning the numberings and/or becoming habituated with the haptic feedback, both of which are otherwise easily achieved with little effort.

We acknowledge that there may exist non-haptic and/or unimanual interaction techniques which support better performance for the particular task used in this study. As mentioned in the previous section, our techniques were not designed primarily for optimal performance, but to investigate the effects of the addition of haptic feedback to a bimanual technique.

2.2 Experimental design and procedure

Dependent measures. A mixed design consisting of one within-subjects effect (force feedback vs. graphical feedback) and one between-subjects effect (relative vs. absolute) was chosen. Half the participants were tested under the F and GA conditions, and the other half under the F and GR conditions. The order of the F and G conditions was counterbalanced.

Participants. Participants were 12 right-handed volunteers, (10 male, 2 female), aged 23 to 42 years, with a median age of 25. All were computing science graduate students at Anonymous Institution. Participants were recruited by email solicitation, and were rewarded for their participation with baked goods.

Protocol. Participants first read a written instruction sheet summarizing the experiment, and completed a consent form and background questionnaire.

Next, participants were introduced to the Phantom. Only two reported having used such a device more than once or twice. A simple demo program which renders a frictionless sphere was run, and participants were asked to poke at the sphere with the stylus. Subjects were encouraged not to be afraid of breaking the Phantom, despite its fragile appearance.

Once apparently comfortable with the Phantom, verbal instructions were issued to participants. An oversized physical model was used to acquaint participants with the numbering of the octants.

Participants completed one practice block plus 5 timed blocks of 25 trials for each condition. Practice trials were monitored by the experimenter, and any abnormal or dangerous behaviour, such as pushing too hard on the Phantom stylus, was pointed out to participants at this time. Optional 20 second breaks were allowed between trials. Time data for the first trial of each block was discarded to allow for participants' readjustment to the apparatus after taking a break.

NASA TLX workload questionnaires were completed after each condition, and an open-ended subjective questionnaire was completed at the end. Sessions lasted approximately 35 to 45 minutes.

2.3 Apparatus

The haptic device used was a Phantom Premium 1.0, driven by a dual Intel Xeon 3.06GHz system with 2GB of RAM at a haptic refresh rate of 1000Hz. The mouse was a Microsoft IntelliMouse Optical 1.1A. In the Windows Control Panel, the mouse speed was set to 5 on a scale from 0 to 10, and mouse acceleration was turned on.

Mouse and Phantom path data were recorded at a rate of 100Hz. The experiment's interface was developed using Microsoft Visual C++ .NET, and ran on Microsoft Windows XP. No other applications were running during the experiment.

3 RESULTS

Analysis of the data revealed no significant differences for the GA and GR conditions for any dependent variable. In addition, observation of participants revealed no apparent difference in strategy or usage habits. Therefore, in the remaining discussion, we group the two conditions into a single graphical feedback condition called G.

3.1 Indications of skilled performance

Error counts and changes in block times indicated that the participants quickly reached skilled performance. Means of completion times for the first six blocks performed by each participant were compared to assess performance improvement over time. A plot of the means across each block for both conditions is shown in Figure 2. A clear improvement is seen between block 1 (the practice block) and block 2 (the first timed block). No appreciable improvement subsequently occurs. This indicates that participants became skilled during the initial 25 practice trials in both conditions. Consequently, we used the 125 trials in blocks 2-6 in our analysis of the time results.




Figure 2: Mean completion times per block for the first six blocks performed by each participant. Error bars show \pm one standard deviation. While a clear improvement is seen between block 1 (the practice block) and block 2 (the first timed block), no appreciable improvement is seen thereafter.

Error counts were low. Two types of errors were recorded. *Number selection errors* occurred when participants clicked within the circle, but had selected the wrong number. *Targeting errors* occurred when participants failed to click within the circle, regardless of the currently selected number. The following analysis does not include errors occurring during practice trials.

Mean error percentages were 4.2% (SD = 3.7%) for number selection and 3.5% (SD = 3.6%) for targeting for the F condition, and 7.4% (SD = 5.7%) and 2.5% (SD = 2.8%) for the G condition. Overall, these indicate that the users had little problem recalling the correct octant to select a number, although other factors appear to have increased the difficulty in the G condition.

All but 3 participants performed more number selection errors in the G condition, while all but 2 performed more targeting errors in the F condition. However, differences in number selection errors were more pronounced.

The differences in errors between the F and G conditions were small, at 3.92 ± 5.28 (95% confidence, $p = .13$) number selection errors and 1.16 ± 1.48 targeting errors (95% confidence, $p = .11$). Also, no appreciable differences were observed between mean error counts for the first (non-practice) and last blocks performed by participants in either condition.

These data suggest performance was consistent across all blocks. For the remainder of this analysis, we assume this was the case.

3.2 Completion time

Figure 3 shows box plots of each subject's completion times for the F and G conditions. Only two participants (2 and 10) did not exhibit faster median performance in the F condition. Additionally, the upper quartile, upper whisker, and upper extremum were lower in the F condition for all participants.




Figure 3: Box plots of completion times for all participants. The upper plot is for the G condition, and the lower is for the F condition. Only participants 2 and 10 exhibited faster median times in the G condition.

Paired t -tests performed on the means of the log completion times over each participant revealed that the F condition was significantly faster than G ($t(11) = 5.94, p < .001$). The G condition was 25% slower than the F condition (95% CI = (15%, 36%)).

3.3 Workload

Workload data was collected and computed using the NASA TLX instrument [8]. A paired t -test revealed marginally significant workload differences between F and G conditions ($t(11) = 1.80, p = .099$).

Table 1: Subjective workloads for each participant for F and G conditions. Only participants 1 and 10 reported markedly higher workloads for the F condition.

Condition	1	2	3	4	5	6
G	2.7	10.7	15.9	8.2	11.9	11.5
F	6.1	11.2	14.0	3.9	11.7	9.4

Condition	7	8	9	10	11	12
G	13.5	13.7	7.8	3.6	14.7	6.5
F	8.2	9.4	5.1	7.8	7.3	5.5

Table 1 shows the calculated workloads for each participant. Notably, the only participants who reported markedly higher workloads for the F condition (participants 1 and 10) were the only two participants who complained of arm fatigue in their subjective questionnaires. This suggests that their reported difference in workload may have been due to discomfort with the device itself rather than the interaction technique.

3.4 Path data




Figure 4: Mouse and Phantom speeds for trial 2 of block 6 in the F condition by participant 5. The mouse clearly completed its movement before the Phantom began.

Based on the recorded path data, the speeds of both the mouse and the Phantom were calculated using the difference in Cartesian distance from one timestep to the next. Those speeds were then plotted for each trial from a sample of blocks from each participant. Speeds for the Phantom (in cm/ms) were scaled by a factor of 10 for the purposes of comparison to mouse speeds (in pixels/ms). A typical example of such a plot is given in Figure 4.

The first sharp spike in Figure 4 depicts a fast mouse movement at the start of the trial. At the tail end of that spike, a slightly wider hump shows a movement of the Phantom. Some small, presumably corrective movements of the mouse then come near the end of the Phantom’s hump, and the trial ends. The mouse and Phantom movements exhibit little simultaneity. Figure 5 is similar to Figure 4, except for the two distinct maxima in the Phantom’s path, which likely represent a corrective movement following an error. Here also, little simultaneity is evident.

Comparison of the times of peak speeds for the mouse and Phantom supported the observation that the mouse moved first in most




Figure 5: Mouse and Phantom speeds for trial 11 of block 6 in the F condition by participant 5. Two spikes are evident in the path of the Phantom, likely indicating an error correction.




Figure 6: Mouse and Phantom speeds for trial 9 of block 6 in the G condition by participant 5. The Phantom’s spike is considerably wider than in Figure 4, indicating that the Phantom moved a greater distance.

trials. Paired t -tests revealed significant differences in times of peak speeds for the mouse and Phantom for both the G condition ($t(11) = 5.39, p < 0.001, 95\% \text{ CI } (380.8\text{ms}, 906.2\text{ms})$) and the F condition ($t(11) = 11.23, p < 0.001, 95\% \text{ CI } (530.7\text{ms}, 789.4\text{ms})$).

To derive a measure of the degree of simultaneity between the movements of the mouse and the Phantom, we estimated the similarity of their speed functions. Specifically, we computed

$$\frac{f \cdot g}{\|f\| \|g\|}$$

where f and g are the speed functions of the mouse and Phantom, treated as vectors. This resulted in a value between 0 and 1, where 0 implies no simultaneity, and 1 implies total simultaneity. Simultaneity as estimated by this measure was generally low, at an average of 0.27 for the G condition and 0.22 for the F condition.

Figure 6 shows a typical trial in the G condition. The same sharp mouse movement is evident at the start of the trial. The following Phantom movement is also present, but takes place over a longer

period. This suggests that participants moved the Phantom at similar speeds but over longer distances in this condition.

The data support this suggestion. A comparison of the mean distance moved by the Phantom in each condition revealed significantly greater distances in the G condition than in the F condition ($t(11) = 2.23, p = .047$).

Despite the trends described above, a non-trivial amount of individual difference was observed across participants. In particular, the amount of simultaneity was noticeably greater for at least three participants. However, increased simultaneity did not appear to affect completion times – these three participants exhibited near-average performance in both conditions, with the exception of one who was slower than the average for the G condition.

3.5 Qualitative Results

In subjective questionnaires, 11 out of 12 participants said that they preferred the F condition, claiming that it provided better awareness of their current position in the octant space (especially in the z direction), allowed them to focus their attention more on clicking the circle, and required less arm movement than the G condition. One participant preferred the F condition “since it restricts the range of motion. The force feedback frees the eyes from tracking the location of the [Phantom]”. This opinion was common.

Most participants were very receptive to the Phantom device, and quickly became accustomed to moving within the cube. However, a small number of participants claimed that the cube was “too small”. This was puzzling, since we assumed participants would determine their position in the cube by pushing against the walls. In this case, the size of the cube should have little effect on the perceived utility of the feedback. One possibility is that those participants were relying in part on proprioception to ascertain their position in the cube, though this seems unlikely.

Several participants reported having not made much use of the graphical feedback in the GA/GR condition. One participant commented, “For the most part I did not use the [graphical Phantom position display] at all, but I was using the [background color] changes from the corner of my eye.” Another reported that rather than shift his gaze continually from the graphical Phantom position display to the target, he ended up making “wider, more deliberate gestures.” Overall, it appears most participants checked the graphical Phantom position display from time to time, but not on every trial.

4 DISCUSSION

4.1 Reasons for performance difference

As shown, in the G conditions, participants wound up moving the Phantom further distances at similar speeds, naturally leading to longer completion times. Lacking the backstop provided by the F condition, users were forced into a trade-off: if they wished to save the time required to shift their gaze to the graphical Phantom position display, they had to move the Phantom a greater distance to ensure they had selected the proper number.

The F condition may also have reduced the time spent in the final phase of movement, verification that the intended target has been reached [19]. While the graphical Phantom position display did provide such confirmation, participants had to shift their gaze to see it, whereas the haptic sensation of the corners of the cube was available immediately. Thus the movement verification phase was likely shorter for the F condition.

In our estimation, it is these two factors that were responsible for the improved performance in the F condition.

4.2 Limited use of graphical feedback

As stated, participants reported making only limited use of the graphical Phantom position display in the G condition. This seems to agree with our finding that participants moved the Phantom further in those conditions. If the display had been attended to on every trial, participants would not presumably have needed to make “wider, more deliberate gestures” in order to ensure that the correct number was selected. Further study involving eye tracking data would be required to conclusively confirm this suspicion.

4.3 Movement simultaneity and ordering

Analysis of the path data showed that participants moved the mouse before the Phantom in most trials, and that they usually did not move the Phantom until the mouse movement was almost complete. This was not expected. Given that the Phantom’s task must be completed before that of the mouse, the natural order would seem to be Phantom-first. Instead, the tendency in most cases was to first make a sharp movement with the mouse toward the target, then move the Phantom to the specified octant, and finally move the mouse the remaining distance to the target and click.

Another unexpected result was the relatively small amount of simultaneity exhibited by participants in any of the conditions. This quandary is further compounded by the total lack of comments concerning simultaneity by participants in their subjective questionnaires.

In the motor control literature, two competing theories exist which may explain these results. The first theory, called the single channel hypothesis, posits that while humans can perceive multiple stimuli in parallel, we can only decide on and plan one action at a time in response to them [18]. This bottleneck could account for participant’s inability to move both devices at the same time.

The natural counterexample to the single channel hypothesis is the fact that humans make simultaneous movements frequently in the course of everyday life – the activity of car driving mentioned at the start of this paper is one such example. While some possible explanations for this fact exist [11], [9], another leading theory of attention asserts that humans actually possess multiple pools of attentional resources which can be drawn upon simultaneously [18].

This theory could also account for our results. If the time taken by participants to map presented numbers to octants were comparable to the time taken plan and execute the mouse movement, then the mouse movement might have been complete by the time the Phantom started moving, even though both processes were taking place in parallel.

5 RELATED WORK

Previous work has typically proceeded by adding haptics to traditional WIMP interaction techniques, and most has involved only the dominant hand.

Akamatsu, MacKenzie, and Hasbroucq [1] used a mouse that provided tactile feedback when users had entered a target. They found that users required less time to verify that they had reached the target, but that overall pointing time was not significantly improved. Engel, Goossens, and Haakma [5] reported similar findings with a modified trackball. Oakley, McGee, Brewster, and Gray [17] implemented several haptic pointing enhancements. They found that errors were reduced, but pointing time was not significantly improved by any of the enhancements.

Dennerlein, Martin, and Hasser [4] found that force display improved performance on a steering task, moving a cursor down a ‘tunnel’ to a target. For this task, where the path is more restricted than general pointing, force constraints improved performance times by 52%.

The above research dealt with pointing and steering tasks with only one potential target. Oakley, Adams, Brewster, and Gray [15] argued that results from studies with single targets may not generalize well to actual interfaces, which feature many possible targets. In this case, forces associated with unintended targets can disturb the movement of the user, reducing performance. Oakley, Brewster, and Gray [16] addressed this by modulating the magnitude of the haptic effects according to the speed of the user's movements to a menu item. They reported a significant reduction in error rates with no significant difference in speed, compared to a visual-only technique.

In another example of selecting from multiple targets, Komerska and Ware [13] added haptic effects to 2D menus in a 3D virtual reality environment. The effects included constraining the cursor to the plane containing the 2D menu, constraining the cursor within the menu boundaries, and snapping the cursor to the centre of a menu item. They found that the haptic features only improved performance 0–4%. Their participants did indicate a subjective preference for the snap-to effect, though.

Other researchers have added haptic effects to more complex techniques. Miller and Zeleznik [14] added haptic effects to GUI features, such as window borders, buttons, and checkboxes. In a similar vein, Komerska and Ware [12] extended the 3D interaction techniques of their GeoZUI3D. In both projects, forces were added to either pull the user's pointer towards a target or keep the pointer on a target once acquired. Neither of these papers reports empirical evaluations of their designs.

Two groups have developed bimanual haptic interaction techniques. Bernstein, Lawrence, and Pao [2] developed a bimanual interaction technique featuring haptic feedback to the dominant hand. The haptic feedback provided contact cues and snap-to-grid effects for a 3D object editor. No results for these effects were reported.

Grosjean, Burkhardt, Coquillart, and Richard [6] added vibrotactile feedback to a technique for selecting one of 27 commands in a virtual reality environment. Users felt a vibration in their non-dominant hand every time they crossed a threshold from one command to another. Unfortunately, the tactile feedback produced slower performance than the same technique with none.

6 CONCLUSION AND FUTURE WORK

In future, we plan develop richer haptic interaction techniques for more complex and realistic tasks. This work bears several suggestions for those designs. First, simultaneous bimanual movement should not be taken for granted. Second, the potential for haptic extents to shorten movement distance and provide confirmatory cues should be exploited.

We would also like to examine if simultaneity of movement would increase given significantly more practice. The high degree of simultaneity practiced by drivers of standard transmission automobiles suggests that it might. Finally, we plan to employ eye tracking technology to study user gaze habits during use of haptic bimanual interaction techniques such as the ones in this study.

7 ACKNOWLEDGEMENTS

This work was funded in part through an NSERC Julie Payette post-graduate scholarship, an NSERC Discovery grant, and a CFI New Opportunities grant.

REFERENCES

- [1] M. Akamatsu, I. S. MacKenzie, and T. Hasbrouq. A comparison of tactile, auditory, and visual feedback in a pointing task using a mouse-type device. *Ergonomics*, 38:816–827, 1995.
- [2] N. L. Bernstein, D. A. Lawrence, and L. Y. Pao. Design of a uniaxially-tuated bimanual haptic interface. In *Proceedings of IEEE HAPTICS 2003*, pages 310–317, 2003.
- [3] E. A. Bier, M. C. Stone, K. Pier, W. Buxton, and T. D. DeRose. Tool-glass and magic lenses: the see-through interface. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques*, pages 73–80. ACM Press, 1993.
- [4] J. T. Dennerlein, D. B. Martin, and C. Hassler. Force-feedback improves performance for steering and combined steering-targeting tasks. In *Proceedings of ACM CHI 2000*, pages 423–429, 2000.
- [5] F. L. Engel, P. Goossens, and R. Haakma. Improved efficiency through I-and E-feedback: A trackball with contextual force feedback. *International Journal of Human-Computer Studies*, 41(6):949–974, 1994.
- [6] J. Grosjean, J. M. Burkhardt, S. Coquillart, and P. Richard. Evaluation of the Command and Control Cube. In *Proceedings of ICMI 2002*, pages 473–478. IEEE Computer Society, 2002.
- [7] F. Guimbertière, A. Martin, and T. Winograd. Benefits of merging command selection and direct manipulation. *ACM Transactions on Computer-Human Interaction*, 12(3):460–476, 2005.
- [8] S. G. Hart and L. E. Staveland. Development of NASA-TLX (task load index): Results of empirical and theoretical research. In P. Hancock and N. Meshkati, editors, *Human Mental Workload*, pages 139–183. North Holland B.V., Amsterdam, 1988.
- [9] H. Heuer, R. A. Schmidt, and D. Ghodsian. Generalized motor programs for rapid bimanual tasks: A two-level multiplicative-rate model. *Biological Cybernetics*, 73:343–356, 1995.
- [10] P. Kabbash, W. Buxton, and A. Sellen. Two-handed input in a compound task. In *Proceedings of ACM CHI 1994*, pages 417–423, New York, NY, USA, 1994. ACM Press.
- [11] J. A. Kelso, D. L. Southard, and D. Goodman. On the coordination of two-handed movements. *Journal of Experimental Psychology: Human Perception and Performance*, 5(2):229–38, 1979.
- [12] R. Komerska and C. Ware. Haptic task constraints for 3D interaction. In *Proceedings of IEEE HAPTICS 2003*, pages 270–277, 2003.
- [13] R. Komerska and C. Ware. A study of haptic linear and pie menus in a 3D fish tank VR environment. In *Proceedings of IEEE HAPTICS 2004*, pages 224–231, 2004.
- [14] T. Miller and R. Zeleznik. An insidious haptic invasion: Adding force feedback to the X desktop. In *Proceedings of ACM UIST 1998*, pages 59–64, 1998.
- [15] I. Oakley, A. Adams, S. A. Brewster, and P. D. Gray. Guidelines for the design of haptic widgets. In *Proceedings of British HCI 2002*. British Computer Society, 2002. 16 pages.
- [16] I. Oakley, S. A. Brewster, and P. D. Gray. Solving multi-target haptic problems in menu interaction. In *Extended Abstracts of ACM CHI 2001*, pages 357–358, 2001.
- [17] I. Oakley, M.R. McGee, S.A. Brewster, and P.D. Gray. Putting the feel in 'look and feel'. In *Proceedings of ACM CHI 2000*, pages 415–422, 2000.
- [18] R. Schmidt and T. D. Lee. *Motor Control and Learning*, pages 64–67. Human Kinetics, Champaign, IL, USA, 3rd edition, 1999.
- [19] N. Walker, D. E. Meyer, and J. B. Smelcer. Spatial and temporal characteristics of rapid cursor-positioning movements with electromechanical mice in human-computer interaction. *Human Factors*, 35(3):431–458, 1993.