Oakley, I. and McGee, M.R. and Brewster, S.A. and Gray, P.D. (2000) Putting the feel in 'look and feel'. In, *Conference on Human Factors in Computing Systems, 1-5 April 2000,* pages pp. 415-422, Hague, Netherlands.

http://eprints.gla.ac.uk/3204/

Putting the Feel in 'Look and Feel'

Ian Oakley, Marilyn Rose McGee, Stephen Brewster and Philip Gray

Glasgow Interactive Systems Group

Department of Computing Science

University of Glasgow, Glasgow, G12 8QQ, UK

Email: {io, mcgeemr, stephen, pdg}@dcs.gla.ac.uk Web: www.dcs.gla.ac.uk/~stephen

ABSTRACT

Haptic devices are now commercially available and thus touch has become a potentially realistic solution to a variety of interaction design challenges. We report on an investigation of the use of touch as a way of reducing visual overload in the conventional desktop. In a two-phase study, we investigated the use of the PHANToM haptic device as a means of interacting with a conventional graphical user interface. The first experiment compared the effects of four different haptic augmentations on usability in a simple targeting task. The second experiment involved a more ecologically-oriented searching and scrolling task. Results indicated that the haptic effects did not improve users performance in terms of task completion time. However, the number of errors made was significantly reduced. Subjective workload measures showed that participants perceived many aspects of workload as significantly less with haptics. The results are described and the implications for the use of haptics in user interface design are discussed.

Keywords

Haptics, force feedback, multimodal interaction.

INTRODUCTION

Desktop interfaces are becoming increasingly complex, and with this added complexity, problems are beginning to emerge. One such problem is information overload, where so much information is presented graphically that it becomes difficult to attend to all relevant parts [4]. Presenting information in other sensory modalities has the potential to lessen this problem. Attempts have been made to overcome information overload using non-speech sound during interactions such as button clicking and scrolling [3, 5] but there have been no convincing empirical attempts to reduce overload by using haptic (or force feedback) technology. This new technology allows users to *feel* their interfaces and has the potential to radically change the way we use computers in the future. We will be able to use our powerful sense of touch as an alternative mechanism to send and receive information in computer interfaces.

Augmenting graphical user interfaces (GUIs) with haptic feedback is not a new idea. In 1994 Akamatsu and Sate [1] developed a haptic mouse with the ability to produce what they termed 'tactile feedback', the ability to vibrate a user's fingertip, and 'force feedback', a simple software controllable friction effect. Using this device they showed significantly decreased completion times in a targeting task offset by slightly increased error rates. Engel *et al.* [7] found improved speed and error rates in a generalised targeting task using a modified trackball with directional two degrees of freedom force feedback.

The devices used in these early studies have now been superseded. More advanced devices such as the Pantograph (Haptic Technologies Inc.), the FEELit mouse (Immersion Corp.), and the PHANToM (SensAble Technologies Inc.) have been developed. These devices have all been used to augment desktop interfaces. Ramstein *et al*. [11] used the Pantograph to demonstrate performance increases in desk-top interactions but provided little empirical evidence to support their claims. The FEELit mouse is a commercial product that offers users a haptically-enhanced desktop but there has been little evaluation of this device published [14]. Finally, the PHANToM has been used to create a haptically enhanced XWindows desktop [10]. No formal evaluation of this enhancement can be found in the literature.

The pace of technological advancement in this field is rapid, both in terms of the hardware produced and the software developed. Current projects to 'haptify' the desktop are not constrained to use the haptic effects described by Akamutsu and Engel. However, as technology has advanced there has been no corresponding progress in its evaluation. This disparity has led to a situation where there are no formal guidelines regarding what feedback is appropriate in different situations. This, along with evidence that shows arbitrary combinations of information presented to different senses is ineffective [12, 13], leads to the conclusion that empirical evaluation of modern haptic augmentations of the desktop is urgently required if much time and effort is not to be wasted. We might even end up with haptically-enhanced interfaces that are in fact harder to use than standard ones and haptics may become just a gimmick, rather than the key improvement in interaction technology that we believe it to be.

Haptic Terminology

Many different terms with many different definitions are used throughout the literature to describe haptic interaction. One reason for this is that the area is in its infancy. To rectify this problem we propose a set of haptic definitions that should prove useful for further research in this area.

The word ‘haptic’ has grown in popularity with the advent of touch in computing. We define the human haptic system to consist of the entire sensory, motor and cognitive components of the body-brain system. It is therefore closest to our understood meaning of proprioceptive (see Table 1). We define haptics therefore to be anything relating to the sense of touch. Under this umbrella term, however, fall several significant distinctions. Most important of these is the division between cutaneous and kinesthetic information (see Table 1). There is some overlap between these two categories; critically both can convey the sensation of contact with an object. The distinction becomes important however when we attempt to describe the emerging technology. In brief, a haptic device provides position input like a mouse but also stimulates the sense of touch by applying output to the user in the form of forces. Tactile devices affect the skin surface by stretching it or pulling it, for example. Force feedback devices affect the finger, hand, or body position and movement. Using these definitions (summarised in Table 1), devices can be categorised and understood by the sensory system that they primarily affect.

Term	Definition
Haptic	Relating to the sense of touch.
Proprioceptive	Relating to sensory information about the state of the body (including cutaneous, kinesthetic, and vestibular sensations).
Vestibular	Pertaining to the perception of head position, acceleration, and deceleration.
Kinesthetic	Meaning the feeling of motion. Relating to sensations originating in muscles, tendons and joints.
Cutaneous	Pertaining to the skin itself or the skin as a sense organ. Includes sensation of pressure, temperature, and pain.
Tactile	Pertaining to the cutaneous sense but more specifically the sensation of pressure rather than temperature or pain.
Force Feedback	Relating to the mechanical production of information sensed by the human kinesthetic system.

Table 1: Definitions of Terminology.

EXPERIMENTAL OVERVIEW

This paper describes two experiments that empirically test the use of haptics to augment targeting in the standard GUI. It is force feedback, and not tactile feedback that is evaluated in this work. Experiment 1 compared user performance with haptically-enhanced buttons using four different haptic effects in a simple targeting task. Experiment 2 involved a more ecologically oriented task in which participants searched for and selected targets using haptic scrolling. We hypothesise that in both experiments haptics will have a positive effect on performance.

Neither of the experiments described is concerned with the influence of haptic distracters; both investigate haptic augmentation when there is guaranteed to be a clear path to target. The decision to adopt this approach reflects the preliminary nature of empirical research in this field.

Device and Software

The device used in both experiments is the PHANToM 1.0 (see Figure 1). It is a force feedback device (provides kinesthetic information as defined in Table 1) which, in the experiments, acted as a cursor control device in place of the traditional mouse.

Optical sensors detect changes in the configuration of the PHANToM. The device uses mechanical actuators to apply forces back to the user calculated from this positional information and the stored algorithmic models of the objects with which the user is currently interacting. To operate the device users hold a stylus.

The graphical interface was generated using standard (MFC) widgets and these performed in exactly the same way as standard widgets. The workspace was a box 160 mm wide x 160 mm high x 2 mm deep. The haptic effects were present only on the back wall of the workspace.

Haptic Effects

Four haptic effects were used in the experiments. These built on and added to the effects used in previous studies. The effects were all aimed at improving targeting and reducing problems of mis-hitting or slipping off interface widgets. The effects used were:

Texture: Texturing a button in a texture-less, flat workspace is a potential way of haptically signifying that the cursor is positioned over an interesting object. The texture implemented here formed a set of concentric circles 7.5 mm apart and centred around the middle of the target. The texture was created by vector rotation (force perturbation) [15] and the maximum rotation applied was 12°. A visual representation is shown in Figure 2. This texture pattern was

Figure 1: The Phantom 3D force feedback device from SensAble Technologies. The stylus shown has a button that can be used for performing the mouse clicks in the experiments reported.

used because it was felt that it would maximise the possibility that users would encounter ridges irrespective of the direction they began from or travelled in.

Figure 2: Diagram of the geometry of haptic texture effect.

Friction: The friction effect damped a user's velocity. Haptically-enhanced interfaces that use a friction effect are common in previous literature. This is partly because they can be produced with simple hardware – for instance with an electromagnet placed in the base of a mouse [1, 2] – and partly because it seems advantageous to provide feedback that causes a user to stop when over an interesting target. The friction effect used here was realistically modelled with both a static and a dynamic component. The static component restricted users to a point until they attained an escape velocity. The dynamic component attempted to slow them whilst they were in motion.

Figure 3: Diagram of the geometry of haptic recess effect.

Recess: The recess effect was a hole in the back of the workspace, with a depth of 2 mm and edges sloped at 45°. This effect also features strongly in previous literature [10, 11]. A diagram of the geometry of a recess is presented in Figure 3. A recess could potentially provide useful feedback by the simple fact that to leave it, the wall at the edge must be climbed. This may make it harder to accidentally slip-off a button (a problem noted by Brewster *et al.* [5]).

Gravity Well: The gravity well was a 'snap-to' effect. When users moved over a button a constant force of 0.5 N was applied that pushed them towards the button's centre. This force tapered off around the very centre so that the user could rest in the centre. The gravity well promised the same benefits as the recess – a reduction in errors through the simple mechanism of preventing a user from accidentally slipping off a button.

General Measures Used in the Experiments

In order to get a full range of quantitative and qualitative results, time, error rates, and subjective workload measures were used in both of the experiments. The subjective workload measurement was a modified version of the NASA Task Load Index (TLX) [8]. NASA reduce workload to six factors: mental demand, physical demand, time pressure, effort expended, performance level achieved, and frustration experienced. We added a seventh factor: fatigue. One potential problem with force feedback devices is the physical strain placed on the user. By adding this factor it would be possible to find out if haptic effects caused any additional perceived fatigue. Participants filled-in workload charts after each condition in both experiments.

EXPERIMENT 1

In the first experiment the haptic effects were compared to investigate which was the most effective. To do this we added each of the haptic effects to standard graphical buttons. This allowed us to investigate targeting (moving the cursor to the button) and mis-hitting errors (slipping-off the button when trying to press it).

Hypotheses

Experiment 1 was an exploratory experiment – we wanted to investigate the differences between the different haptic effects and a control condition. Therefore, the experimental hypotheses were that differences would occur in task completion time, number of errors and in the subjective data gathered. We predicted that the gravity well and recess would provide the largest reduction in errors, time and workload as they provided feedback that was highly appropriate to a simple targeting task.

Participants

There were sixteen participants. Four were female and twelve were male. All were between the ages of eighteen and thirty. Most were computing students from the University of Glasgow. All were regular and fluent computer users. Three users were left-handed and one was dyslexic. None had anything more than trivial previous exposure to the PHANToM.

Design

The experiment followed a within-subjects repeated-measures design. Each participant underwent each of the four haptic conditions, each encompassing one of the effects described above, and a control condition. The control condition used the PHANToM device but no haptic effects were applied – in essence the device worked like a normal mouse. The order of the presentation of the conditions was counterbalanced to evenly distribute the effects of practice and fatigue. Participants were randomly allocated to conditions. Training was given in each condition in a session immediately prior to the experiment. Each condition in the training session constituted 60 button presses and in the experimental session 120 presses. The experiment's duration was typically 45 minutes.

Task

A simple button pressing/targeting task was used. This task was chosen because it featured prominently in the previous literature [1, 2, 7] and also because it is a very elementary operation – it is both simple to perform and also perhaps the most fundamental cursor operation.

Two factors were engineered into the task to make it more suitable for haptic augmentation. Firstly, it was felt that participants should experience some visual distraction. This is not an unlikely circumstance in the typical operation of a GUI, particularly in the case of expert users. They concentrate on some central task and interact with graphical widgets in the periphery of their attention [4]. Secondly, in this atmosphere of visual distraction, we assert that the haptic feedback will only really prove useful if the task encompasses some repetitive motion. Without such motions the haptic task would rapidly dissolve into exhaustively searching the entire workspace for some haptically distinct area. This is clearly an inefficient strategy when compared to visually scanning the screen. Repetitive motions are also common in desktop interactions (moving to menu bars, clicking buttons, etc.).

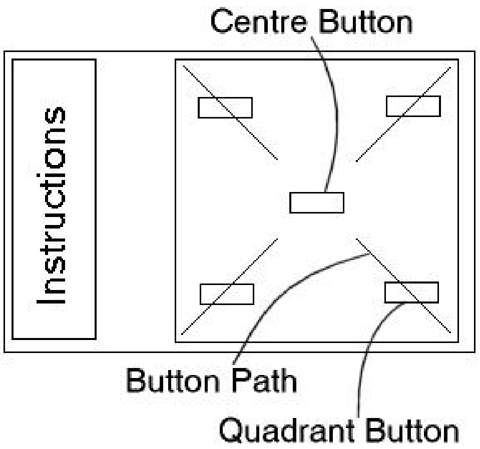

Figure 4: The interface used in Experiment 1.

To encompass these two factors two windows were placed on the screen at all times (see Figure 4). One, the instruction window, occupied the left-hand side of the screen and contained instructions as to the next target to seek. The other, larger, window occupied the centre and right-hand side of the screen and contained the targets in the form of five buttons. One button was always positioned in the centre, the other four were positioned one in each quadrant of the window, on the diagonals of the window. The position along the diagonals was changed in the course of the experiment, but each button remained in a single quadrant of the window throughout. This meant that each button remained in the same direction relative to the centre of the window at all times. The buttons moved along the diagonals to prevent users employing a purely mechanical repetition. To ensure users moved along only a few trajectories to reach each of the buttons, every second button press was the centre button. The buttons were labelled in accordance with their positions on screen, for instance “top right” or “bottom left”. The instruction window indicated the next target button, on successfully pressing the named button, a new name was presented.

Measures

Data were gathered from all button presses in the experiment. The performance measures were (a) mean time per trial (secs.), (b) mean number of errors, and (c) subjective workload ratings. Times were measured at four stages: time to find target button; time to move onto target button; time to press target button; and time to move off target button. Errors were measured as when a participant moved over a button but failed to press it. There were two categories: the first was where the user simply slid over the button, arguably as a part of the normal targeting process. The second, more serious error is known as a ‘slip-off’ [4]. This occurs when a user presses the mouse down over a button but moves off it before releasing the mouse, thus not selecting it. The feedback for this is the same as for a successful mouse click. An error of this type can go unnoticed for some time and cause considerable confusion.

Results from Experiment 1

The error data are presented in Figures 5 and 6. Results were analysed using ANOVA tests. Significant effects were found when comparing the mean scores for each haptic effect for both slide over ($F_{4,15}$ = 48.487, p<0.001) and slip off ($F_{4,15}$ = 20.81, p<0.001) errors. Order effects for both slide over ($F_{4,15}$ = 0.152, p=0.961) and slip off ($F_{4,15}$ = 0.123, p=0.974) errors were not found.

	Gravity	Recess	Friction	Texture	Control
Gravity	---------	Not sig	p<0.001	p<0.001	p<0.002
Recess	---------	---------	p<0.01	p<0.003	Not sig
Friction	---------	---------	---------	p<0.04	Not sig
Texture	---------	---------	---------	---------	p<0.016

Table 2: Analysis of slip-off errors in Experiment 1.

	Gravity	Recess	Friction	Texture	Control
Gravity	---------	Not sig	p<0.01	p<0.01	p<0.01
Recess	---------	---------	p<0.01	p<0.01	p<0.01
Friction	---------	---------	---------	p<0.01	Not sig
Texture	---------	---------	---------	---------	p<0.01

Table 3: Analysis of slide-over errors in Experiment 1.

A summary of the results revealed by *post-hoc* analysis of the means (using Bonferroni confidence interval adjustments) is shown in Tables 2 and 3. The most dramatic results were that participants in the gravity condition made significantly fewer errors of both sorts than in the control and that the converse was true of the texture condition – it caused significantly more errors than the control.

Analysis of the temporal data was less conclusive; the total time taken to complete a trial was strongly biased by the number of errors made in each condition. It was felt that this invalidated it as a measure – it would merely be a reflection of the number of errors in each condition. Instead, the total time on a button during a successful trial was analysed (see Figure 7). An ANOVA revealed significant differences between effects. Subsequent pair-wise comparisons (using Bonferroni adjustments) revealed that gravity was significantly slower than recess (p<0.05). It is also worth

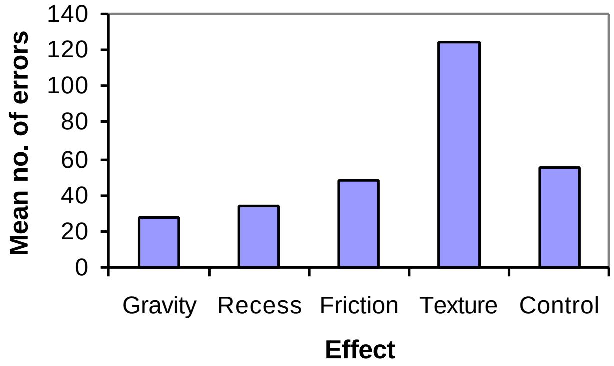

Figure 5: Slide over errors in Experiment 1.

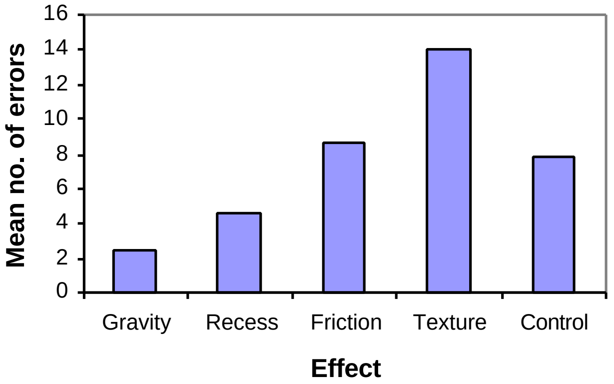

Figure 6: Slip-off errors in Experiment 1.

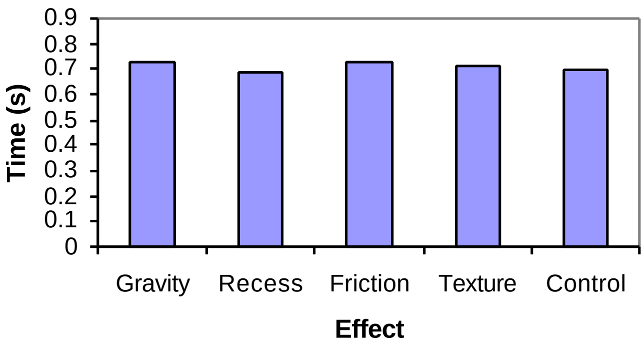

Figure 7: Total time on button in Experiment 1.

noting that the difference between the best and worse performing effects was only 42 ms, a very short time. No order effects were found in this temporal analysis ($F_{4,15}$ = 0.913, p=0.462).

To validate analysing time and errors separately we ran a Pearson correlation. The timing results did not correlate with the slide over (r=0.0, p<1.0) or slip off (r=0.019, p<0.976) errors. The two error results strongly correlated with one another (r=0.938, p<0.018).

Figure 8 shows the TLX workload scores (scored out of 20). The texture condition was significantly worse than the control across the whole board of measures. The gravity

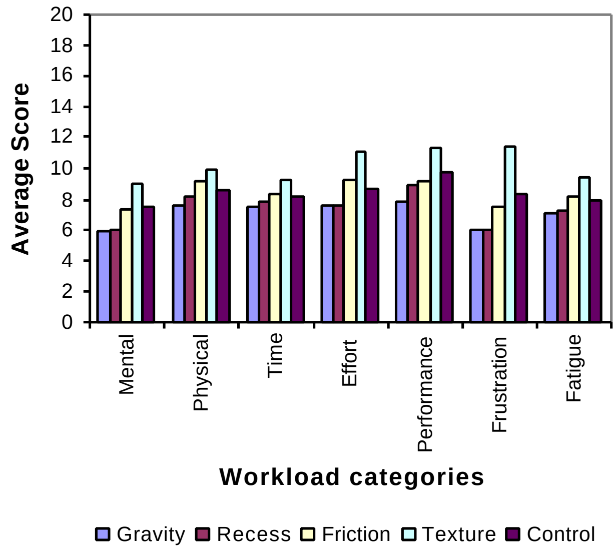

Figure 8: Workload results from Experiment 1.

condition consistently reduced workload and, in particular, achieved a significantly better score than the control in the performance level achieved category (p<0.018).

EXPERIMENT 2

This experiment simulated a more realistic task where reading was accompanied by scrolling through a document, selecting from the document, and returning to the scroll bar whilst still visually attending to the material being read. When users are required to scroll through a document it is the material in it that is of interest and not the scroll bar. Users want to concentrate on reading the material but often find themselves forced to move their visual attention to the scroll bar to ensure that the cursor is positioned appropriately to operate it. The time taken to make these frequent shifts in visual attention, and the frustration experienced by the need to do so, reduce the usability of the scroll bar. Problems associated with scrolling have been addressed previously [e.g. 4, 16]. Reducing these problems using force feedback technology has not yet been empirically evaluated.

Hypotheses

It was hypothesised that when the scroll bar was haptically-enhanced, the participants would (a) take significantly less time to complete the task; (b) move on and off the scroll bar significantly less; and (c) perceive the workload during the task as significantly less.

Participants

Twenty new participants were used: one was female and the remaining nineteen male. All were between the ages of seventeen and twenty-seven. Most participants were first-year computing science students from the University of Glasgow. All were regular and fluent computer users. All users were right-handed. Participants had nothing more than trivial previous exposure to the PHANToM.

Design

The experiment again used a within-subjects repeated-measures design. Each participant underwent both a visual-only condition (visual) and a visual and haptic condition (haptic). The visual condition used a standard graphical scroll bar only. In the haptic condition, this scroll bar was overlaid with haptic effects (recess and gravity well were chosen as these were the most effective in Experiment 1). The up and down arrow buttons used gravity wells. These acted as a haptic indication that the user was in the appropriate place to press the button successfully. The rest of the scrolling area used a recess effect that allowed the user to 'fall into' the slider area. Therefore, the haptic feedback allowed the user to reserve his/her visual attention for the primary task, as being over the widget was indicated through touch. The order of the presentation of the conditions was counterbalanced to evenly distribute the effects of practice and fatigue. Training was given to each participant in each condition prior to the experiment.

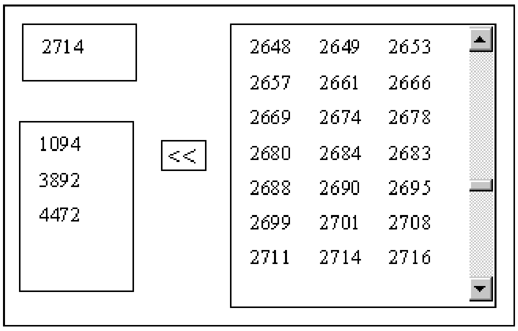

Figure 9: The interface used in Experiment 2. The top left window is the instruction window, the bottom left is the target window, the large window to the right is the data window and in the centre is the send button.

Procedure

Figure 9 shows the interface to the task. Participants had to read a four-digit numerical code from the instruction window. They then had to scroll vertically through a large file of codes (presented in the data window) to find the target code, highlight the code (either by double clicking on it or dragging across it), and press a button to send this code to the target window. The widgets operated as in standard desktop applications. The data window contained the same list of 2000 randomly generated but numerically ordered codes in each condition. Forty codes had to be entered in each condition. The list was formatted such that there were three columns of codes, simulating a standard document read from left to right and from top to bottom. The highlight operation was included to force the user off the scroll bar. This ensured repeated targeting of the scroll bar. The experiment's duration was typically 40 minutes.

Measures

The performance measures were (a) mean time per trial (secs.), (b) mean number of movements on/off scroll bar (including all required movements), and (c) workload ratings. Time was measured from when the user activated the send button at the end of the previous trial until the send button was activated at the end of the current trial. Subjective ratings were collected as before.

Results from Experiment 2

Timing results: Table 4 shows the timing and movement on/off scroll bar results. Paired T-tests established that haptic feedback did not significantly reduce the average trial time as predicted (T_{19} =0 .46, p< 0.32).

Mean Trial Time (secs.)		No. times on/off scroll bar	
Visual	Haptic	Visual	Haptic
11.7251	11.9668	107	97
SD=2.77	SD=2.84	SD=25	SD=22

Table 4: Timing and movement results from Experiment 2.

Movement on/off scroll bar: Paired T-tests showed that participants in the haptic condition moved on and off the scroll bar area significantly less than in the visual condition (T_{19} = 2.37, p< 0.05).

Workload Results: Figure 10 shows the workload scores. Paired T-tests were carried out on the visual versus haptic conditions for each of the categories. Mental demand was not significantly less in the haptic condition as expected. Both the effort and frustration ratings were significantly reduced in the haptic condition (Effort: T_{19} = 2.80, p<0.01, Frustration: T_{19} = 2.04, p<0.05). There was no significant difference in fatigue experienced. The hypothesis that the haptic condition would reduce workload is therefore confirmed in part.

GENERAL DISCUSSION

The timing results from the two studies indicate that the haptic effects added to the buttons and scroll bar did not reduce the time taken for either task, as hypothesised. There were also no real differences between the effects – only 42 ms between the best and worst effects (recess and gravity) in Experiment 1. The explicit separation of the error data from

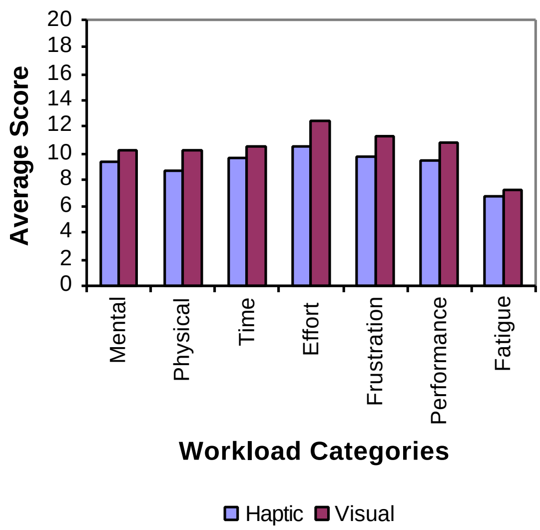

Figure 10: Workload results from Experiment 2.

the timing data is no doubt a contributing factor to the lack of temporal variations across conditions. However, we suggest that one potential reason for the lack of time reduction is that, in all of the effects used, participants had to exert more force to overcome the haptic effects. In the control condition they could just slide over the interface with no obstacles, in the haptic conditions they had to climb out of recesses, overcome gravity forces applied, etc. For participants to produce the forces required to do this could have taken them more time.

Further work is needed on the haptic effects themselves and the types of desktop tasks that would benefit most from them. It may have been that the haptic effects chosen were inappropriate either for reducing time or for the tasks chosen for these experiments. Other previous work has claimed a significant reduction in performance times [10, 14]. The present work suggests that things are no so clear-cut and care must be taken when using haptics to try to reduce performance times.

The error results were more conclusive. Experiment 1 showed a significant reduction in the number of errors produced across the different haptic conditions (where gravity and recess caused the fewest errors and texture the most). Gravity and recess were the most effective for targeting tasks (which are important for using many standard GUI widgets, for example hitting a button, selecting a menu item or dragging the scrollbar thumb) in the sense that they made it very hard to slip off a target once on it; participants could not just knock the pointer off the target, they had to make an explicit movement to leave. Texture only indicated that the cursor was over a target, and did not constrain users to the target, which was one of the reasons it was less effective in this case. Texture also had the problem that it could potentially perturb users' movements, making it hard for them to stay on target. This resulted from the kinesthetic force feedback device used here. We use cutaneous stimulation to feel much of the richness of fine-grained texture in the real world [9]. A kinesthetic device can only simulate gross textures, requiring larger forces, which then make it harder for users to move precisely. Texture is much more suitable to production by tactile devices such as the *Tractile* from Campbell *et al*. [6]. The PHANToM, on the other hand, is very effective at simulating gravity and recess effects as these require movement and so are kinesthetic tasks. There are no devices, as yet, which combine both tactile and kinesthetic force feedback.

Haptic devices are now reaching the desktop. For example, the FEELit Mouse [14] adds low cost haptic effects to the standard graphical interface. Our results show that interface designers must be aware of the facilities of the devices they are using in order to generate haptic effects that will improve usability. This might seem obvious, but this area is in its infancy and new devices are appearing all the time, each having different functionality to the last.

The movement results from Experiment 2 showed a significant reduction in the number of times a participant moved on/off the scroll bar in the haptic condition. This showed that the haptic recess aided participants in remaining on target, demonstrating that haptics can provide a significant practical benefit for interaction. The haptic groove placed over the scroll bar allowed users to scroll up and down without slipping off. They could do this without looking at the bar as once the cursor was in the groove it would stay there. To move out of the recess they had to lift off the scroll bar and it was difficult to do this by mistake as it required a conscious effort.

The subjective workload measures taken across both experiments are important. Papers concerning other haptically-enhanced desktops have not presented any such data. In developing multimodal interfaces (ones that use multiple sensory modalities) it is very important to consider what effects they have on users' workload. Users may perform tasks well and quickly and yet find them frustrating and requiring more effort to complete than they would expect. This dissociation between behavioral measures and subjective experience has been addressed in studies of workload. Hart and Wickens [8] suggest that cognitive resources are required for a task and there is a finite amount of these. As a task becomes more difficult, the same level of performance can only be achieved by the investment of more resources. Just measuring time or error rates does not give the whole picture of the usability of a haptic device. Workload is particularly important in this area as we know little yet of the effects on cognitive/attentional resources of using such devices.

Experiment 1 showed that the different effects had markedly different levels of workload. Gravity well and recess came out best, indicating that they were effective at reducing error rates and decreasing workload. This suggests that they are very robust and can be successfully used in haptic interfaces of the type described here. Texture came out the worst in terms of workload, suggesting that, in general, it is hard to do effectively with the device used here. Experiment 2 showed the effect of haptics in a more realistic situation. In this case there was a significant reduction in effort and frustration – the fact that it was easy to stay on the scroll bar due to the recess effect made the task much less effortful (the reduction in the number of movements on/off the scroll bar confirms this). We had expected that this might also lead to reductions in other categories (e.g. mental demand) but these showed no significant reductions. This suggests that we need further studies of workload to learn more about the affect of haptics in desktop interactions.

One other area that we investigated was fatigue. Using a device that requires the user to apply force could cause fatigue. It is important to investigate this if force feedback devices are to be used in desktop situations (where people might use the interfaces for long periods of time). Results from Experiment 1 showed that gravity and recess effects did not cause any more fatigue than the control condition. On the other hand, texture caused significantly more fatigue than the control. This is likely to be for the reasons as discussed above – to simulate texture with a kinesthetic device required larger forces to be applied and these, in turn, required the users to exert larger forces to overcome them. Experiment 2 again showed no increase in fatigue with the use of gravity well and recess effects. This research shows that appropriate haptic effects used correctly may have no impact on fatigue, but used incorrectly may significantly increase it. This is only a first step in investigating this problem and further work is needed to ensure that we can design haptic interfaces to avoid fatigue

CONCLUSIONS

Our research has shown that haptics may have some benefits in graphical user interfaces. Reductions in the number of errors made and subjective workload experienced can be gained. We have also shown that the haptic effects used must be matched to the capabilities of the device – trying to simulate effects not supported by the device in use can have serious negative effects on all aspects of usability. As technology progresses it is easy to focus on what benefits new equipment may afford whilst forgetting to measure the benefits actually produced. Recent work on haptically-enhanced desktops has been firmly orientated towards implementation and the experiments described here begin to redress the balance. Our empirical findings provide a firm foundation for future researchers to build on and some basic principles for developers to use.

ACKNOWLEDGEMENTS

This research was supported under EPSRC project GR/L79212 and EPSRC studentship 98700418. Thanks must also go to the SHEFC REVELATION Project, SensAble Technologies and Virtual Presence Ltd.

REFERENCES

1. Akamatsu, M. & Sate, S. (1994). A multi-modal mouse with tactile and force feedback. *International Journal of Human-Computer Studies* (40), 443-453.
2. Akamatsu, M. & MacKenzie, S. (1996). Movement Characteristics using a Mouse with Tactile and Force Feedback. *International Journal of Human Computer Studies* (45), 483-493.
3. Beaudouin-Lafon, M. & Conversy, S. (1996). Auditory illusions for audio feedback. In *ACM CHI'96 Conference Companion* (Vancouver, Canada) ACM Press, Addison-Wesley, 299-300.
4. Brewster, S.A. (1997). Using Non-Speech Sound to Overcome Information Overload. *Displays*, 17,179-189.
5. Brewster, S.A. (1998). The design of sonically-enhanced widgets. *Interacting with Computers*, 11(2), 211-235.
6. Campbell, C.S., Zhai, S., May, K.W. & Maglo, P. (1999). What you feel must be what you see: adding tactile feedback to the trackpoint. In *IFIP Interact'99*, (Edinburgh, UK), IOS Press, 383-390.
7. Engel, F.L., Goossens, P. & Haakma, R. (1994). Improved Efficiency through I- and E-Feedback: A Trackball with Contextual Force Feedback. *International Journal of Human-Computer Studies*, 41(6), 949–974.
8. Hart, S.G. & Wickens, C. (1990). Workload assessment and prediction. MANPRINT, an approach to systems integration, 257-296, Van Nostrand Reinhold.
9. Lederman, S.J., Summers, C. & Klatzky, R.L. (1996). Cognitive salience of haptic object properties: Role of modality-encoding bias. *Perception*, 25, 983-998.
10. Miller, T. & Zeleznik, R. (1998). An Insidious Haptic Invasion: Adding Force Feedback to the X Desktop. In *ACM UIST'98*, (San Francisco, CA) ACM Press, 59-64.
11. Ramstein, C. (1995). A Multimodal User Interface System with Force Feedback and Physical Models. In *IFIP Interact'95*, (Lillehammer, Norway) Chapman & Hall, 157-162.
12. Ramstein, C. & Hayward, V. (1994). The Pantograph: A Large Workspace Haptic Device for Multi-Modal Human-Computer Interaction. In *Summary Proceedings of ACM CHI'94*, (Boston, MA) ACM Press, Addison-Wesley, 57-58.
13. Ramstein, C., Martial, O., Dufresne, A., Carignan, M., Chassé, P. & Mabilleau, P. (1996). Touching and hearing GUIs: Design issues for the PC-access system. In *Proceedings of ACM Assets'96*, (Vancouver, Canada), ACM Press, 2-10.
14. Rosenberg, L.B. (1997). FEELit mouse: Adding a realistic sense of FEEL to the computing experience. http://www.force-feedback.com/feelit
15. Srinivasan, M.A. & Basdogan, C. (1997). Haptics in Virtual Environments: Taxonomy, Research Status, and Challenges. *Computers and Graphics*, 21(4), 393-404.
16. Zhai, S., Smith, B.A., & Selker, T. (1997). Improving Browsing Performance: A study of Four Input Devices for Scrolling and Pointing Tasks. In *Proceedings of Interact 97*, (Sydney, Aus) Chapman & Hall, 286-293.