Comparing haptic effects in a GUI

Ian Oakley
Department of Computing Science
University of Glasgow
Glasgow, G12 8QQ, UK
Email : io@dcs.gla.ac.uk

Abstract

Graphical user interfaces have been augmented with haptic feedback on several occasions, usually by either adding simple effects such as friction or vibration [1] or by mapping various parts of the environment to different height values, the transition between these areas being signified by slopes [5]. However the sophistication of haptic hardware has increased greatly in recent years and we believe that research in this area has not fully utilized the higher fidelity of feedback now available. We also believe that it is important to experimentally compare the different potential mechanisms for haptic feedback, attempting to identify feedback appropriate to different situations.

This paper describes an initial attempt to compare and contrast the use of a variety of haptic effects in a target selection task. Buttons were augmented with friction, a simple texture, a "snap to" gravity-like effect or formed the shape of a recessed button. Experimental evaluation of these effects indicates that haptic feedback can lead to significantly improved or reduced performance, yielding the conclusion that further evaluation of desktop haptic interfaces is required.

Introduction

Augmenting graphical user interfaces (GUI) with haptic feedback is not a new idea. In the early nineties Akamutsu [1, 2] developed a haptic mouse with the ability to produce what he termed tactile feedback, the ability to vibrate a user's fingertip, and force feedback, a simple software controllable friction effect. Using this device he showed significantly decreased completion times in a targeting task offset by slightly increased error rates. Engel et al [3] found improved speed and error rates in a generalized targeting task using a modified trackball with directional two degree of freedom force feedback.

The devices used in these early studies have now been vastly improved. More advanced devices such as the pantograph [7], the FEELit mouse [8] and the PHANToM [4] have been developed. These devices have all been used to augment a desktop interface. Ramstein et al [6] used the pantograph to demonstrate performance increases in desktop interactions but provided little empirical evidence to support their claims. The FEELit mouse is a commercial product aimed at offering customers a haptically enhanced desktop. However , there has been little concrete evaluation of this device. Finally the PHANToM has been used to create a haptically enhanced X Desktop [5]. No formal evaluation of this enhancement has yet taken place.

The pace of technological advancement in this field has forged ahead in leaps and bounds, both in terms of the hardware produced and the software developed. Current projects to "haptify" the desktop are not constrained to use the primitive haptic effects described by Akamutsu [1, 2] and Engel [3]. However as the 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, combined with the fact that there is strong evidence that arbitrary combinations of information presented in different modalities is ineffective [3], leads to the conclusion that empirical evaluation of modern haptic augmentations to the desktop is urgently required.

This paper describes an experiment that attempts to compare user performance in a haptically modified desktop using a PHANToM and four different haptic effects in a simple targeting task. The effects used were gravity well, friction, recess, and texture.

Methods

System

The haptic device used was a PHANToM 1.0. The hardware used was an Intel P II running at 300 MHz and under Windows NT 4, SP 3. All software was written in C++ using Microsoft Visual Studio version 5. The haptics code was created using the GHOST API supplied by SensAble Technologies. Subjects manipulated the PHANToM using the standard pen surrogate

Software

The software architecture used in this experiment was very simple. Separate but coincident representations were maintained for both the graphical and haptic displays. As the PHANToM's state changes, this information is extracted from the haptic model and mapped to a corresponding mouse event. This causes the appropriate event in the GUI. As the interface changes during the course of the experiment the two representations are updated manually. It is not suggested that an architecture such as this is suitable for more general use, nor that it would scale up to a realistic, real world, system.

The haptic representation was created from a management class and individual haptic effect classes. The management class held the positions and sizes of each haptic effect in the model. It also contained pointers to haptic effect classes which actually implemented the feedback presented to the user. The effect classes operated on a local coordinate frame created for them by the management class. This architecture allows the simple specification of different haptic augmentations independent of any large scale geometry and also allows ultimate flexibility in the positioning of haptic effects. There are no constraints as to what haptic effects appear in what locations. For instance using this architecture it is simple to create two buttons that are augmented using different effects but are otherwise identical.

Effects

Texture

Texturing a button in a texture-less, flat workspace is a potential way of haptically signifying that the cursor is positioned over some interesting object. The texture used in this experiment formed a set of concentric circles centered around the middle of the target. This is pictured in figure 1. This texture was used because it was felt that it would maximize the possibility that a user would encounter ridges irrespective of the direction they began from or traveled in.

Friction

The friction effect damped a user's velocity.

Figure 1. Diagram of Texture

Haptified GUIs 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.

Recess

The recess effect was a hole in the back of the workspace. This also features strongly in previous literature (eg Miller, Ramstein). A diagram of the geometry of a recess is presented in figure 2. A recess could potentially provide useful feedback by the simple fact that to leave a recess, the wall at the edge must be climbed. This may make it harder to accidentally fall off a button.

Figure 2. Diagram of Recess

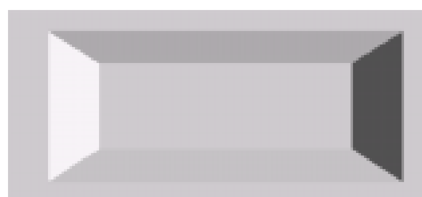

Gravity Well

The gravity well was a snap to effect. When a user moved over a button a constant force that pushed them towards the button center was applied. This force tapered off around the very center of the button to create an area of softer force there. The gravity well promises the same benefits that the recess does – a reduction in errors through the simple mechanism of preventing a user from accidentally leaving a button.

Task

The task studied here was a simple targeting task – subjects were required to seek to and press buttons. These buttons appeared on a large window occupying the center and right hand side of the screen. There were five buttons in total. All remained constantly on screen. One was positioned in the center, the others occupied a quadrant of the window, on the diagonals of the window. The buttons were labeled in accordance with their positions on screen, for instance “top right” or “bottom left”. Instructions as to which button to press next were presented in a window to the left of the screen. When a user pressed the button currently named in this window another name was displayed and the next trial began. As a new trial began the position of the four non-central buttons changed. They remained in their appropriate quadrants and on the diagonals of the window, merely changed position along this diagonal. Every second trial was the central button.

This task was felt to provide an element of visual distraction by presenting the instructions spatially separate from the experimental stimuli. It was also felt that sufficient repetition was produced by the consistent placing of the buttons to facilitate the successful use of haptic feedback. In a task with no repetition of motion, haptically locating a button would become a clearly inefficient exhaustive search.

Subjects

There were sixteen subjects. Four were female and twelve male. All were between the ages of eighteen and thirty. Most were computing students. All were regular and fluent computer users. Three users were left-handed and one was dyslexic. No subject had anything more than trivial previous exposure to the PHANToM.

Experimental Design

The experiment followed a within subjects repeated measures design. Each subject underwent each of the four haptic conditions and a control condition. The order subjects experienced the conditions was varied systematically. Training was given to each subject in each condition in a practice session prior to the experiment.

Measures

Extensive measurements were taken for each subject in the experiment. Each trial consisted of four stages - seeking to the required button, moving on to it, pressing it and moving off it. Timing information was kept for each of these stages. All errors when a subject moves over a button but fails to press it were also recorded.

Results

The error data is presented in graph 1. Using a one-way repeated measures ANOVA significant effects were found. Post hoc pair wise comparisons yielded significant effects between all haptic conditions at $p<0.001$ except for those between gravity and recess and the control condition and texture which were significant at $p<0.05$.

Analysis of the temporal data was less conclusive. Analysis of the total time taken to complete a trial was strongly biased by the number of errors experienced 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. To find differences in purely temporal performance, data for the time spent over a button for every successful trial was used. This data is presented in graph 2. A second one way repeated measures ANOVA on this data revealed significant differences. Pair wise comparisons isolated these as recess being significantly quicker than both gravity and friction ($p<0.01$) and the control condition being just better ($p<0.05$) than friction.

Graph 1. Error Data

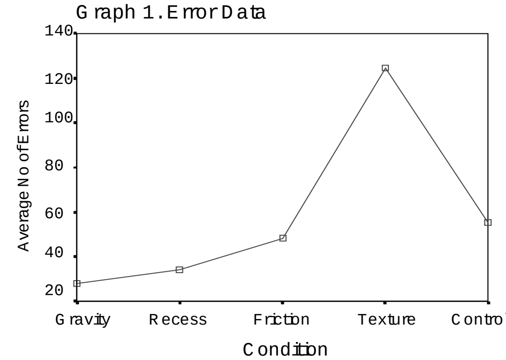

Graph 2. Total Time on Button

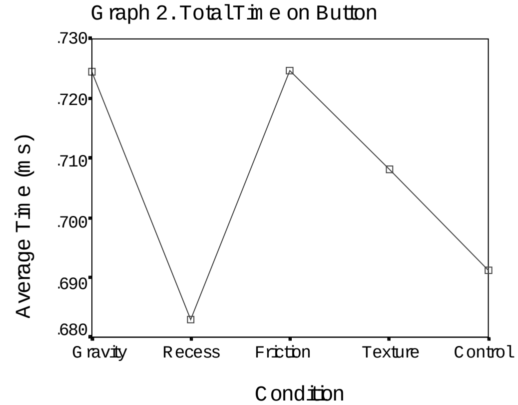

Discussion

The results discussed here are preliminary; further analysis of the results is in progress. A first look seems to nominate recess as the optimal haptic effect for targeting tasks. It has a low error rate and fast performance time. However Akamutsu [1] in a similar haptically enhanced targeting experiment measures temporal performance only in terms of the time it takes a user to reach and press a button. He dismisses the time it takes to move off a button after a click. An analysis of the temporal results from this experiment on this basis indicates that recess may not convey optimal performance benefits. This data is shown in graph 3.

Large relative improvements in both gravity and friction are found. The slow performance that these effects have in the overall time is a result of the delay they engender on the time it takes to leave a button. This suggests a trade off between best performance in targeting a button which is provided by the gravity well in both errors and temporal performance and the best performance for the entire interaction with a button, which is instead provided by the recess.

It is also worth noting that the differences in temporal performance, while significant, are not large. The difference between the fastest performance and the slowest is under 50 ms. Recent non-evaluatory studies [5,8] have observationally claimed time performance doubled through haptic augmentation of widgets. In light of these results this seems unlikely. More substantial and important are the error results. These conclusively demonstrate the need for the evaluation of haptic feedback in desktop user interfaces. The texture effect is shown to be highly disruptive, experiencing twice the number of errors that occurred in the control and demonstrating that not all feedback is helpful. The gravity effect experienced approximately half the number of errors as the control condition, showing substantial improvement.

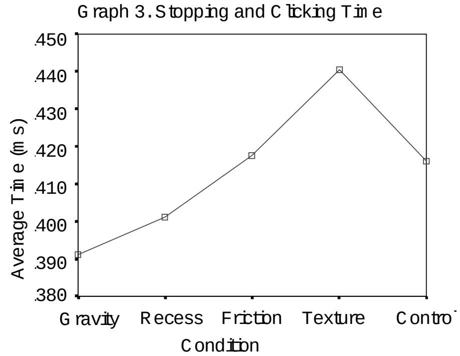

While this experiment can not come close to providing a full framework or set of guidelines for implementing haptic interfaces, it does identify the need for a body of empirical work that can do this. It shows that both positive and negative feedback can be generated and that there are substantial performance differences across various haptic effects.

Acknowledgements

The author is supported under ESPRC project GR/C 79212. Thanks to SHEFC Revelation Project, SensAble Technologies and Virtual Presence.

Bibliography

1. Akamutsu, A and Mackenzie, S. I. 1996, Movement characteristics using a mouse with tactile and force feedback, *International Journal of Human Computer Studies*, 44, 483-493.

2. Akamutsu, A. and Sata, S. 1994, A multi-modal mouse with tactile and force feedback, *International Journal of Human Computer Studies*, 40, 443-453.

3. Engel, F. L., Goossens, P. H. and Haakma, R. 1994, Improved efficiency through I- and E-feedback: a trackball with contextual force feedback, *International Journal of Human Computer Studies*, 41, 949-974.

4. Massie, T. H. 1996, Initial Haptic Explorations with the PHANToM: Virtual Touch through Pointer Interaction. *Masters Thesis, Massachusetts Institute of Technology*.

5. Miller, T. and Zeleznik, R. 1998, An Insideous Haptic Invasion: Adding Force Feedback to the X Desktop. *In Proceedings of UIST '98*.

6. Ramstein, C. 1994, MUIS : A Multi-modal User Interface System with Force Feedback and Physical Models.

7. Ramstein, C. and Hayward, V. 1994, The Pantograph: A Large Scale Workspace Haptic Device for Multi-modal Interaction, *Proceedings of CHI '94, Interactive Experience*.

8. Rosenburg, L. B. 1997, FEELit mouse: Adding a realistic sense of FEEL to the computing experience. *http://www.force-feedback.com/feelit/white-paper.html*