

# Web-Based Multimodal Graphs for Visually Impaired People

---

Wai Yu, Douglas Reid, and Stephen Brewster

## Abstract

This paper describes the development and evaluation of Web-based multimodal graphs designed for visually impaired and blind people. The information in the graphs is conveyed to visually impaired people through haptic and audio channels. The motivation of this work is to address problems faced by visually impaired people in accessing graphical information on the Internet, particularly the common types of graphs for data visualization. In our work, line graphs, bar charts and pie charts are accessible through a force feedback device, the Logitech WingMan Force Feedback Mouse. Pre-recorded sound files are used to represent graph contents to users. In order to test the usability of the developed Web graphs, an evaluation was conducted with bar charts as the experimental platform. The results showed that the participants could successfully use the haptic and audio features to extract information from the Web graphs.

## 1 Introduction

The work described in this paper is part of an ongoing research project 'Multivis' which aims to develop a multimodal data visualization system for visually impaired and blind people. Common data visualizations such as line graphs, bar charts and pie charts, etc, are presented to users through haptic and audio channels. To achieve this multimodal representation, we use force feedback devices, spatialized non-speech sound and synthesized speech. This work introduces our attempt at rendering various graph types on the World Wide Web with the aim of providing universal access.

The advent of World Wide Web (WWW) has changed many people's way of life. Not only do sighted people benefit from the growth of the Web but so do visually impaired people. With the help of synthesized speech and Braille display technology, visually impaired people can access a wealth of information and services such as news, travel timetables, online shopping and banking, etc. (RNIB, 2001). However, for these assistive technologies to work properly, Web pages must be appropriately designed and written in valid HyperText Mark-up Language (HTML). Moreover, due to the Graphical User Interface (GUI) orientated design,

Web pages often contain invisible tables, columnar text and graphical icons which are potential barriers to the assistive technologies mentioned above (Nguyen, 2001). Furthermore, pictures, diagrams and graphs, which are often found on Web pages, cannot be represented easily by synthesized speech or Braille. It is desirable to provide means for visually impaired people to access these kinds of information on the Web.

To address the inaccessibility of the GUI for visually impaired people, force feedback devices have been developed (Mynatt and Weber, 1994, Ramestein *et al*, 1996). Therefore, icons, window frames and other standard GUI features can be felt by visually impaired people through the device. Besides tackling GUI problems, other research has been done in the haptic representation of graphical contents on Web pages. Blind people's perception of virtual 3D objects and ability of using online education tools have been investigated (Hardwick *et al*, 1998, Wies *et al*, 2001). The commonly used force feedback devices are expensive 3D position input and force output devices. Our approach is to use a low-cost (~£60) Logitech WingMan Force Feedback Mouse (Figure 1) to relay information. The aim of this work is to create universal accessible Web pages which contain various types of graphs.



**Figure 1.** Logitech WingMan Force Feedback Mouse.

This paper describes the design and implementation of the Web-based graphs. Most effort has been put into the making of haptic and audio features in order to meet the needs of visually impaired people. Details of the evaluation, which has been conducted to assess the usability of Web graphs, are also given in this paper.

## 2. Web Graphs Development

Multiple senses modalities were used to present Web graphs. Besides the graphical display, haptic effects and audio feedback are incorporated into the Web pages. The haptic effect enables users to touch and interact with the graphical objects whereas audio provides additional information about the graphs. JavaScript is used for inserting the haptic effects into the Web pages as well as triggering the audio feedback based on users' mouse activities. The format of the Web pages and coordination of haptic and audio features are controlled in Cascading Style Sheets.

## 2.1. Web Page Format

Cascading Style Sheets (CSS) have been proposed by the World Wide Web Consortium (W3C) to control the style of appearance of a Web page (W3C, 2001). In this approach, the content of a Web page can be separated from its presentation. It is a simple, declarative language that allows Web developers to control the style and structure of documents written in HTML or XML (Extensible Markup Language). Values such as font, size, colour and positioning can be assigned to a document tag. The word ‘cascading’ in CSS refers to what occurs when several sources of style information are available for control of the elements on a page. Style information is passed down from a higher level style sheet until it is overridden by a style command with more weight. The CSS language is simple and easy to implement, and most importantly using it fulfills the requirements stated in the W3C Web page design guidelines (W3C, 1999).

Using CSS also has advantages when assistive technologies, i.e. screen readers, refreshable Braille displays, are used. The content can be accessed directly without confusions from style controls embedded in HTML. CSS also gives the flexibility required for multimodal Web pages. Due to the architecture of integrating multiple modalities into one Web page, two layers of images are needed. They can be implemented and controlled easily in style sheets. Flexible control and consistent appearance of the Web pages have been achieved.

## 2.2. Haptic graph implementation


The haptic effect, which makes the Web graphs tangible, is constructed by using the Immersion Web ActiveX control and Web Plug-in. These allow the communication between JavaScript within the source script of a Web page and a force feedback device to produce the desired sensations on the page. The plug-ins need to be installed on a user’s computer for the force effect to be activated. The force feedback device used in this work is the Logitech WingMan Force Feedback mouse. It is an absolute positioning device and provides 2-degree-of-freedom force feedback. The amount of force and the size of the workspace provided by the mouse are limited, but it is affordable to most blind individuals.

The haptic features of the Web graphs are constructed by integrating appropriate primitive haptic objects supported by the Immersion Web Plug-in. The general procedure to create the haptic effect involves three basic steps. Firstly, the appropriate type of haptic effect has to be decided. Secondly, parameters of the haptic effect such as location, size and force response have to be set to match with the modeling object. Finally, the events that control the start and stop of the haptic effect have to be defined. In general, these events are based on mouse movements.

To render different types of graphs which are commonly used, such as line graphs, bar charts and pie charts, dissimilar implementation is required. The methods that we used in this work are based on the primitive haptic objects supported by the Immersion Web Plug-in. The most used effect is enclosure which constrains the mouse movements in a pre-defined area (both rectangular and

elliptical). If users wish to exit from the enclosed area, they just need to push the mouse harder against the resisting force. The rectangle effect is used to construct the bar charts and line graphs whilst combination of ellipse and rectangle effects is used for pie charts.

A bar chart is relatively easy to implement due to its regularity. A rectangular enclosure effect is used to model the bars. The coordinates of each bar on the graph are measured and subsequently used to define the location and size of the effects. Both the inside and outside walls of the effect are set as ‘touchable’ so that users can feel the boundary of the bar. If only one side of the wall is enabled, force feedback only applies to the mouse when it travels towards the same side of the wall (Figure 2). Therefore, both sides of the wall are enabled to give a realistic wall simulation.




**Figure 2.** (a) A virtual wall with both sides enabled, (b) one side of the wall is disabled and the mouse can move across without any restrictions.

To render pie charts, both rectangular and elliptical enclosure effects have to be used. The ellipse effect is used to form the circumference of the pie. Only the inner wall of the ellipse is enabled so that users can feel the internal area of the pie. This is because the most interesting part of pie charts is the divisions of the pie rather than its overall size. The boundary of different portions of the pie is represented by a very narrow rectangular enclosure effect. Only the top and bottom walls of the rectangular effect are active and the gap between them is kept to a minimum distance. Therefore, moving between the top and bottom walls is like following a straight line. The position of the rectangular effect has to be precisely positioned so that the line effect starts from the centre of the pie to the edge. Figure 3a shows this transition, a rectangle enclosure is rotated and translated into the desired position. The number of rectangular effects required depends on the number of pie divisions.

Line graphs are more difficult to model using the primitive haptic objects. This is due to the potential complexity of the line or curve shapes. In order to construct a smooth curved line, techniques proposed by Fritz and Barner (1999) are more appropriate. However, it is not feasible to implement those techniques based on the functions provided in the plug-ins. Therefore, the rectangular enclosure effect is used to simulate simple straight-line approximations. The same technique used in the pie charts to form pie divisions is used here to represent line segments. Multiple straight lines are rotated and positioned on the graph so that the end points

of each line segments are jointed together (Figure 3b). The problem with this kind of formation is that smooth transition cannot be achieved when the line changes direction sharply.



**Figure 3.** (a) Pie chart arrangement, (b) Line graph arrangement.

## 2.3 Audio Implementation

Web-based graphs have an advantage over their traditional paper-based counterparts in that they have the capacity to include interactive sound to relay further information to the user. Our previous work has shown that users' performance has significantly improved when audio is introduced into a haptic graph (Yu and Brewster, 2002). In this work, redundant sounds were added to give additional feedback on the numerical information presented on these graphs.

Different forms of audio exist on the Internet including Wave, Musical Instrument Digital Interface (MIDI), Real Audio, MP3, etc. Each audio format has its own advantages and disadvantages. The criteria for choosing suitable audio format in this work are based on the popularity and bandwidth requirements. Popularity means whether the audio format has universal support so that it can be accessed in any Web browser. Not all Web users have the capacity or the desire to download and use plug-ins. Moreover, the downloading and installation of an audio plug-in can be a fairly daunting task even for the computer literate (RNIB, 2001, WCAG, 2001).

The limited bandwidth of standard analogue modems means Web developers can barely afford a fully featured graphical interface, even before audio content is considered (Beggs and Thede, 2001). The combination of intensive visual/audio pages, and slow modem speeds can result in slow loading of images, and audio frequency dropping or skipping ahead. Bandwidth limitations can also reduce the quality of the sound.


The suitable candidates for the audio feedback are MIDI and Wave. MIDI offers the smallest file size per minute of any audio format as a MIDI file only contains a simple set of instructions to how a synthesizer or sound engine should play back the file. Therefore, the download time is short and it is platform independent. However, the sound quality depends on the sound card installed on users' computer. The playback sound may sound different as each manufacturer may use a slightly different technique to sample and reproduce sounds.

Wave is a common sound format. It has an excellent browser compatibility rating with most browsers carrying a ready installed plug-in to play Wave files.

The format is also easy to place onto a standard Web server for downloading, and is ideally suited for simple Web sound effects such as button rollovers. The sound reproduced on users' computer will be more or less the same as the sound recorded on the developer's computer. However, Wave does not offer elegant compression and is therefore not suited for playback over limited bandwidth conditions. As the number of audio notes that playback on the Web graphs is fairly small, the limitation of Wave is not really a problem in the audio representation. Therefore, Wave format is used in the current development to present audio feedback.

Currently, audio feedback is only implemented on bar charts. Audio notes are played when the mouse cursor moves over a bar. A JavaScript method is used to perform this task. As the 'mouse over' event that triggers a JavaScript function is linked to one image only, multiple images have to be included on a Web page for a series of audio notes to be playback interactively. A separate layer of image is therefore implemented on top the haptic image. The arrangement of different layers of images is managed in the style sheet. The audio images are slightly smaller than the haptic images and overlay them. Sound used on the haptic graphs is recorded from a MIDI synthesizer. The height of the bars is mapped proportionally to the MIDI notes. Therefore, a tall bar will generate a high note and vice versa. By moving the mouse cursor across all the bars, the changing pitches will give users an idea about the bar values (Ramloll *et al*, 2001).

### 3. Evaluation



**Figure 4.** Two sample graphs used in the experiment: (a) Condition 1, (b) Condition 2.

An evaluation was conducted to assess how the Web-based graphs could be used. Bar charts were selected for the evaluation as they have been developed further in terms of haptic and audio features. In the experiment, the effect of two different bar chart arrangements was also investigated. In the first condition, bars have a gap between them (Figure 4a). In the second condition, all the bars are flush together (Figure 4b). The main testing purpose is to see which one is easier for users to navigate and understand the graph content. Therefore, the experiment was designed to compare their individual effectiveness in transmitting information to the user. The major investigation issues include how the different graph arrangements affect

users' navigation through the graphs, how the individual bars were compared against each other, and how much workload they placed on the users.

### 3.1 Experiment Setup

This experiment consisted of a two condition within subjects counter-balanced design. Six bar charts were constructed for each experiment condition. The data used to generate the bar charts were taken from the Website of the U. K. Department of Health (DOH, 2001). They were the statistics of hospitals in England from 1993 to 2000. There were seven bars on each chart and arranged according to the condition, either having a gap in between (Condition 1) or flush together (Condition 2).

Twelve sighted people took part in our experiment and they were students of the Department of Computing Science at the University of Glasgow. We did not use blind people in the experiment because our previous experimental results as well as other researchers' work indicate that there is no significant difference between blind and sighted people's performance in using the novel haptic interface (Ramstein *et al*, 1996, Yu *et al*, 2001). During the experiment, participants did not have any visual feedback on the graphs and had to work on the Logitech WingMan Force Feedback mouse and wore a pair of headphones.

In order to test whether participants could successfully use the haptic and audio features on the bar charts to extract useful information, four questions were asked.


- Q 1. What is the general trend of the graph?
- Q 2. Which bar represents the lowest value?
- Q 3. Which bar represents the highest value?
- Q 4. Which two bars are the closest in value?

A four minute time limit was placed on each graph. The participants were instructed to answer the questions as quickly as possible but keep a high accuracy. The time taken by the participants to answer all four questions was noted. Before the experiment started, four practice graphs were given to the participants for training purposes. In the experiment, the sequence of which condition came first was randomized in order to counter-balance the effect of learning. After the experiment, each participant was asked to fill in a questionnaire regarding the workload they experienced in the experiment. The screen activity, which shows the users' cursor movement, was recorded on video tapes.

### 3.2 Experiment Results

Three major measurements were taken in the evaluation process: number of correct answers, time taken to answer all questions, and the NASA task load index. The task load index is designed to assess the workload placed on a user based on his/her subjective responses on the contributing factors (Harts and Wicken, 1990). There are six factors including mental demand, physical demand, temporal demand, effort, performance and frustration.

Figure 5a shows the results of the average number of correct answers per question across all 6 bar charts. The results of the first three questions in both conditions are very similar. Participants were able to obtain a large percentage of correct answers. A significant drop in accuracy occurs in Question 4. Only 29.17% correct answers was obtained in Condition 1. Situation in Condition 2 is better with 40.28% accuracy. The t-test revealed a significant difference between them ( $T_{11}=2.345$ ,  $p=0.0388$ ). The overall accuracy in conditions 1 and 2 is 77.78% and 80.21% respectively. However, no significance has been found ( $T_{11}=1.865$ ,  $p=0.0891$ ).



**Figure 5.** (a) Average correct answers, (b) Average workload index (Condition 1 – bars separated, Condition 2 – bars close together).

The task completion time recorded varies between participants. The average time taken by the participants in condition 1 is 2 minutes and 7 seconds which was longer than the time required in condition 2, 1 minute and 49 seconds. The difference between them was not significant as the t-test only shows  $T_{11}=1.985$  and  $p=0.0727$ .

A similar situation happens to the workload index. The participants perceived lower workload in condition 2 and thus gave a workload index 44.9%. The workload index of condition 1 was rated 50.7%. Again no significant difference was found ( $T_{11}=0.975$ ,  $p=0.351$ ). The average distributions of participants' ratings on the six contribution factors are plotted in Figure 5b.

### 3.3 Discussion

According to participants' performance, the first three questions are easier to answer than the last question. The accuracy of the answers to the last question dropped significantly which means that it is difficult to compare the height of different bars. The two different arrangements of the bars had very little effect on participants' answers to the first three questions due to the ceiling effect. High accuracy is maintained in both conditions. The difference between two conditions shows in the answer to the last question (the comparison). Condition 2 received slightly more correct answers (40.28%) than condition 1 (29.17%). This implies



that arranging bars close together helps participants to compare different bar heights more than separating bars by a gap. However, the overall correct answers did not show a significant difference between two conditions.

Task completion time varies between individuals, some participants needed more time, the other did not. The average task completion time in condition 2 is shorter than the one in condition 1. However, a t-test did not reveal a significant difference between them.

In the workload index assessment, the overall workload index in condition 2 is again less than the one in condition 1. However, t-test did not produce a significant result. On the other hand, the variance of the participants' ratings showed a significant difference between two conditions ( $F_{11} > 2.82$ ,  $p = 0.024$ ). The variances in condition 1 and 2 are 18.07 and 5.13 respectively. This indicates that the workload perceived by participants in condition 2 is more consistent.

Mental demand and effort received much higher ratings from participants than other factors. Obviously, to perceive information through haptic and audio channels is much harder than using vision. The narrower bandwidth of these two channels has placed a heavy burden on the participants when they were trying to assemble pieces of information picked up through the interface into a complete picture. Moreover, all participants are sighted people, it is very unusual for them to ignore vision and use senses of touch and hearing alone.

Participants' ways of exploring the graphs during the experiment were observed. The common technique adopted by most participants is to move the mouse cursor across the bars. Based on this pattern of cursor movement, it can be seen that most of the participants were using the haptic feedback as a navigation guide rather than a tool to extract information about the bars, i.e. bar height. The haptic feedback informs the participants when they were entering the graph or moving into a new bar. Participants relied on the audio feedback on the bars to determine their value. Force feedback was rarely used for this purpose, or as a backup to the audio interpretations. This may explain why participants' answers to the last question were so inaccurate. The nature of the graph representation might affect users' way of using haptics. As detecting the height of the bars by haptics is an indirect and time-consuming process in comparison with the audio representation. Users would not bother to move the mouse up and down on the bars to find the answers, instead they would use the audio provided by the interface. On the other hand, a MIDI note stored in a Wave file is an abstract form of information. It is very difficult to make comparisons based on the MIDI notes alone. Moreover, very few people have perfect pitch. More precise forms of data representation are required for comparison. At the moment, while this kind of information cue is not available at the interface, users should be encouraged to use haptic features on the graph to assist their exploration. Hopefully, the performance of people who can successfully use both haptic and audio features can be improved.

Based on the overall number of correct answers, task completion time and workload index, both experiment conditions provided participants with similar performance. Participants' comments were also noted. In condition 1, a user commented that the graph was "tough to navigate" in comparison to condition 2. Participants often seemed to get slightly more disorientated on condition 1 graphs,

especially where there were dramatic changes in bar height. The fact that condition 2 bars led onto each other was preferred by some of the participants. To them, the feeling of empty space between the bars with no haptic or audio feedback was disconcerting and perhaps putting a slightly greater demand. This has coincided with the design guidelines for tactile display proposed by Challis and Edwards (2001). Space between the bars can be confusing and cause disorientation to users.

On the other hand, some participants preferred having a gap between bars. With bars close together, some users felt there was no definite distinction between the bars, with the possibility arising of mistaking one bar for another. In summary, although the bars arranged close together has little effect on the participants' performance except in the data comparison, majority of participants would prefer this layout.

## 4. Conclusions

Web-based multimodal graphs have been developed. Three types of graphs have been implemented including bar charts, pie charts and line graphs. They were designed to address problems faced by visually impaired people in accessing graphical contents on Web pages. The Web graphs are presented to users in visual, haptic and audio forms. At the current stage, audio feedback is available on the bar charts. Haptic features can be explored by using a Logitech WingMan Force Feedback mouse.

An evaluation, which was conducted to test the effectiveness of using haptic and audio effects in conveying information the users, has shown that participants are able to use the features after a very short training. A high accuracy of participants' answers has revealed the successful use of the Web pages. It has been observed that most participants use haptic features for navigation. They were guided to explore the bar charts and able to discover their position on the graph. Audio features were used to access the information presented on the bars. This works well with locating general and approximate information from the graph. A problem arises when precise information is required for data comparison. It shows that non-speech audio alone cannot give participants a very good indication of the data value therefore a low accuracy is obtained in the last question. The audio used in this implementation is hence regarded as useful in trend detection but not in detailed information representation. However, when synthesized speech is introduced, the amount of information that can be presented will be increased.

The experiment only gives inconclusive results on the difference between two types of bar charts. There are not significant differences with respect to the overall number of correct answers, task completion time and overall workload index. In general, very few differences exist whether the bars are arranged close together or having a gap in between. However, a difference exists in the accuracy of the answers to the last question. Participants did better in finding out two similar bars in height in condition 2 and gave more consistent ratings to the workload index which is slightly lower than in condition 1. Moreover, based on the feedback from participants, most of them prefer to have the bars closely arranged together to

minimize the chance of disorientation in the empty space. Therefore, to render an effective haptic bar chart, it is recommended to have bars placed next to each other closely.

## 5. Future Work

Constructing the Web-based graphs is a very tedious process. The implementation of the graphics and audio clips, as well as the positioning of these components is very laborious. The next stage of development will be to produce a system by which online haptic graphs can be produced automatically from raw data. We are currently looking at the possibility of using Java Applets in conjunction with JavaScript to make the graph production easier and more flexible. Moreover, we will conduct case studies on the new system with blind people.

## 6. Acknowledgements

Multivis project is funded by EPSRC Grant GR/M44866, ONCE (Spain) and Virtual Presence Ltd.

## 7. References

- Beggs, J., Thede, D., (2001) Designing Web Audio- RealAudio, MP3, Flash and Beatnik. O'Reilly.
- Challis B. P., Edwards A. D. N., (2001) Design Principles for Tactile Interaction. Haptic Human-Computer Interaction, 2001. LNCS 2058: p. 17-24, Springer.
- Department of Health, Hospital Activity Statistics, (2001) <http://www.doh.gov.uk/hospitalactivity/index.htm>. Accessed July 2001
- Fritz J. P., Barner .K.E., (1999) Design of a Haptic Data Visualisation System for People with Visual Impairments. IEEE Transactions on Rehabilitation Engineering, 7(3): p. 372-384.
- Hardwick A., Furner S., Rush J., (1998) Tactile display of virtual reality from the World Wide Web - a potential access method for blind people. Displays, (18): p. 153-161.
- Hart S. G., Wickens C., (1990) Workload assessment and predication, in MANPRINT, an approach to systems integration, Van Nostrand Reinhold: New York.
- Mynatt E. D., Weber G., (1994) Nonvisual Presentation of Graphical User Interfaces: Contrasting Two Approaches. ACM, CHI 94: p. 166-172, ACM Press Addison-Wesley.
- Nguyen K., (1996) The Accessible Web: Web Access Through Adaptive Technology. <http://www.utoronto.ca/atrc/rd/library/papers/WebAccess.html> ATRC Library, Accessed July 2001.
- Ramloll, R., Yu, W., Riedel, B. and Brewster, S.A., (2001) Using Non-speech Sounds to Improve Access to 2D Tabular Numerical Information for Visually Impaired Users. In Proceedings of BCS IHM-HCI 2001 (Lille, France), Springer, pp 515-530.
- Ramstein C., Martial O., Dufresne A., Carignan M., Chasse P., Mabillean P., (1996) Touching and Hearing GUI's: Design Issues for the PC-Access System. ACM ASSETS 96, p. 2-9, ACM Press.

- Royal National Institute for the Blind (RNIB), (2001) Accessible Web Design. <http://www.rnib.org.uk/digital/hints.htm>, Accessed July 2001
- Web Content Accessibility Guidelines (W3C Recommendation 5-May-1999), <http://www.w3.org/TR/1999/WAI-WEBCONTENT-19990505>, Accessed July 2001.
- Web Content Accessibility Guidelines 1.0, (W3C). (2001) <http://www.w3.org/TR/WCAG10/>, Accessed July 2001.
- Wies E. F., Gardner J. A., Sile O'Modhrain M., Hasser C. J., Bulatov V. L., (2001) Web-based Touch Display for Accessible Science Education. *Haptic Human-Computer Interaction*, LNCS 2058: p. 52-60 Springer.
- World Wide Web Consortium (W3C), (2001) Style Sheet Activity Statement. <http://www.w3.org/Style/Activity>, Accessed July 2001.
- Yu W., Ramloll R., Brewster S., Ridel B., (2001) Exploring Computer-Generated Line Graphs through Virtual Touch. *Proceedings of The Sixth International Symposium on Signal Processing and Its Applications*, 1: p. 72-75, IEEE.
- Yu W., Brewster S., (2002) Comparing Two Haptic Interfaces for Multimodal Graph Rendering, Accepted by IEEE VR2002, 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (Haptics 2002), in Orlando Florida, March 24-28, 2002, IEEE.

**Author list**

Wai Yu  
Department of Computing Science  
University of Glasgow  
Glasgow  
G12 8RZ, U. K.  
rayu@dcs.gla.ac.uk  
<http://www.multivis.org>

Douglas Reid  
Department of Computing Science  
University of Glasgow  
Glasgow  
G12 8RZ, U. K.  
reiddh@dcs.gla.ac.uk

Stephen Brewster  
Department of Computing Science  
University of Glasgow  
Glasgow  
G12 8RZ, U. K.  
stephen@dcs.gla.ac.uk