Chapter 5. Usability Implications of Perception and Memory - Introduction

Table of Contents

Context

Unit 4 introduced the area of cognitive psychology and focused on the important cognitive process of visual perception. Although vision is, at present, the key sense used in interaction, there are and will be roles for our other #input channels#. This unit looks at the user#s capabilities and limitations in terms of the auditory (hearing) and haptic (touch) senses and also, more briefly, at the possibilities of using taste and smell in interaction. Technologies and applications relevant to the different channels will be discussed. The unit continues with an exploration of two other cognitive elements: the process of attention and human memory. You will learn about these elements and, importantly, how this knowledge can be applied in the design process.

Objectives

At the end of this unit you will be able to:

- Discuss the use of audio at the interface.
- Illustrate how current audio use might be extended both verbally (speech synthesis and recognition) and non-verbally (e.g. auditory icons and earcons).
- Discuss the current limited role of the haptic sense and possible extensions (e.g., Braille output, use in virtual reality etc).
- Give example potential applications and technologies for smell and taste output / input (e.g. e-commerce).
- Define the cognitive processes of attention.
- State why designers should accommodate a user’s attention mechanism.
- Show how designs can be improved in terms of the way they focus a user’s attention.
- Distinguish between the 3 forms of human memory.
- Show how knowledge of the 3 forms of memory can be used in design.
- Distinguish between the process of recall and recognition and discuss why and how recognition can be promoted.

Introduction

The previous unit discussed the properties of the visual medium and the implications of what we know about human vision for the design of interactive systems. Of course, vision is not the only way we are able to perceive the world around us, and in this section we shall discuss the properties of some of our other senses, especially hearing, and the implications that these properties have for user interface design.

In addition to considering how senses other than vision may be used in human-computer interaction, we will also look at the mechanisms by which sensory inputs are processed by the human cognitive system. In particular, we shall look at how we are able to pay attention to certain sensory stimuli and mental activities, while ignoring others, and what implications this has for interface design. Finally in this unit, we will outline how human memory works, and once again take a look at the relevance this has for interface designers.

Sound in the Interface

The vast majority of computer-based user interfaces that we encounter rely almost totally on the visual medium. However, another medium that is frequently used in interface design is sound. Before discussing how interface designers do and may make use of sound, we will review some of the properties of the audio medium consisting of sound together with the human sense of hearing.

Properties of sound and hearing

A number of properties of the audio medium are pertinent to the use of sound in the interface.

Sounds can vary in a number of dimensions: pitch (or frequency), timbre (or musical tone), and intensity (or loudness). Not all sounds or variations is sound are audible to humans. The ear is capable of hearing frequencies ranging from about 20Hz up to about 15KHz, and differences of around 1.5 Hz can be discerned (though this is less accurate at high frequencies).

The capability of computer sound output devices to produce variations along each of the dimensions of pitch, timbre and intensity means that sound output is potentially a rich and sophisticated medium for conveying information to users.

The audio medium, like any other, has a number of inherent properties that constrain the way humans process and make sense of the sounds they hear. Therefore, understanding these constraints will be crucial to the successful use of sound in interactive systems.

In contrast to vision, sound is a "volatile" medium in the sense that sounds do not persist in the same way that visual images do. Or to put it another way, the visual field can be regarded as a parallel array of information elements, each of which may remain constant or may vary over time. Sound, on the other hand, can be seen as a single element (described by its pitch, timbre and intensity) that may vary over time, and the rate at which it varies or carries information, or is perceived, is not under the control of the listener. Its potential as a means of conveying information, and the amount and type of information that can be carried, is therefore rather different from that of the visual channel. Consequently, the visual channel can be regarded as frequently having a much faster access time. For example, a large amount of information may be made simultaneously available in the visual channel, whereas presenting the information in the audio channel may take longer as the information must be "serialised". A further implication of this is that an audio information stream may place greater demands on the user’s memory: while listening to part of a message, the user must remember what has gone before. When reading a visual display, the parts previously read remain visible.

Another relevant property of the audio medium is the fact that, unlike vision, hearing is non-directional. While binaural hearing does grant us a limited ability to determine the direction of a sound source, there is no sense in which we can listen in a particular direction. Similarly, while we can very easily control our visual sense (e.g., by looking in a particular direction), it is much harder to be selective about what one listens to.

It is well known that people are rather good at noticing changes in sound – such as the onset or cessation of a noise or tone. However, we are rather less good at extracting information from a stream that remains relatively constant. In fact, if a background sound remains relatively constant, over a period of time we will tend to become less aware of it, and eventually will filter it completely and cease to notice it at all. See the later section on ATTENTION.

A further property of sound and hearing that designers should be aware of is that we are relatively poor at separating out a single sound or sequence of sounds from a background of similar sounds. Imagine trying to hold a conversation in a noisy environment where many other people are talking. Or trying to follow two conversations at once.

For most of us, in our everyday lives, sound plays a very important part. It is often said that the majority of information we receive about the world comes to us through our visual sense. While this is true, it is also the case that sound plays a central role in communicating with others (through speech and other sounds), receiving information and entertainment (through radio broadcasts, musical performances, and so on), and allowing us to be aware of events – some of which may be outside our visual field (police sirens, ringing telephones, etc). Sounds that are apparently in the ‘background’ often give us vital clues about the status of ongoing processes (e.g., the sound made by a car engine as we are driving, the noise made by machinery in a factory).

Despite the apparent limitations described above, sound is a remarkably important channel for conveying information. Next we will look at some of the ways sound is used in current user interfaces, and how it could be used in the future, and will identify some guidelines that can help designers to make the use of sound more successful.

Activity 1 - Sound and vision

You can carry out this activity on your own if you like, but it might be easier if you can work in pairs. Compare the time taken to read a written passage, with that taken to listen to the same text being spoken. For the latter part, you can simply read the text out loud, or use a computer speech synthesiser.

Activity 2 - Menu Structure

It is often said structured menus should be made wide (many choices in a menu) and shallow (few levels of menu), rather than deep (many levels). Studies suggest that doing this will improve allow users to carry out tasks faster, while making fewer errors. See (Shneiderman, 1998, pages 249-250) or (Norman, 1990, chapter 5) for discussions.

Suppose that instead of menus on a visual display, we are designing menus in an audio interface for telephone-based services. Do the same guidelines about wide/shallow or deep/narrow structures apply? How would you investigate the efficiency of different audio menu schemes?

A Discussion on this activity can be found at the end of the chapter.

Use of sound in current and future interfaces

Sound is currently used in many computer interfaces and other interactive devices, sometimes as an important source of information, and at other times simply as a means of making the interface seem more impressive. Sound can be used for both input and output functions in an interface, and the kinds of sounds that are used include abstract tones and bleeps, naturally occurring sounds, music, and, of course, speech.

The audio channel is, for most of us, a rich and important source of information in interactions with other people or with the environment. However, sound has played a more limited role in human-computer interaction than in other aspects of life. Part of the reason for this may be that it is not always clear what kinds of functions sound is appropriate for, and how to make effective use of sounds to support those functions.

Audio Alerts

While sound might be ineffectively exploited, that is not to say it is unused. Almost all computer systems "bleep" when an error of some sort occurs. What is the purpose of such a bleep, and what can the user infer from it? Clearly the sound indicates to the user that something has happened, but it is typically left up to the user to determine by other means the nature of the event (an erroneous input, or a system generated warning, or simply the arrival of email?), its source (one of several applications, the operating system, networking software?), and what should be done about it. The following, rather extreme, example highlights the kind of problematic situation that can arise.

Note

In 1979, at the Three Mile Island nuclear power station in the USA, one of the most serious nuclear accidents in history took place. As the incident was unfolding, operators in the plant control room were faced with a bewildering array of instruments and displays giving information about the state of the power station. As problems arose, auditory alarms sounded to alert the operators to what was going on. The following commentary on the incident (drawing on the report of the subsequent investigation by a Presidential Commission) explains one aspect of the operators' predicament:

Frederick and Faust [two of the operators] were in the control room when the first alarm sounded, followed by a cascade of alarms that numbered 100 within minutes. … Later Faust would recall for the Commission his reaction to the incessant alarms: "I would have liked to thrown away the alarm panel. It wasn't giving us any useful information."

Partly as a result of the confusing alarms, the operators failed to diagnose one of the serious problems with the reactor for two hours, during which time extensive damage and leakage of radioactive material occurred.

The point is that although alarms and alerts can be successful for indicating that some event has occurred, they often carry little useful information about the nature or location of the event or what should in response. And if several notable events occur together, then providing an auditory indication of each is simply going to confuse users. The world of power stations and control rooms may seem very far removed from everyday design, but the same issues are relevant in the interfaces of desktop systems and the design of web pages.

Sound output, therefore must be used with care. Sometimes it is appropriate to indicate a change of status or a particular event with a sound, but we must be aware that beeps and similar alerting sounds often provide the user with too little information about either the nature of the event or what action will need to be taken as a result. Used with care, however, sound can enhance an interface and provide users with an important source of information.

Providing Information

So far, we have been discussing sounds without saying much about what kinds of sounds are typically – or in the future could be – useful features of an interface. The simplest and most basic sounds are simple bleeps that indicate an event or change of state. However, other possibilities exist.

An idea that several researchers have experimented with is to add sounds to many of the familiar features of existing visual user interfaces. So, in addition to the visual cue provided by an icon, an auditory cue would be provided as a further memory aid. The sounds used in some of these experiments were natural ones, chosen to match the kind of interface objects with which they are associated. For example, folders on a desktop might have the sound of paper crumpling (presumably because in the real world, we put paper into folders), and dropping an item in the wastebasket might produce a sound of breaking (because things sometimes break when we throw them in the bin).

While this might seem an attractive way of providing the user with additional feedback and extra ways of remembering how the interface works, it has some problems. One is that the same sound might mean different things to different people – so the sounds are not as "natural" as they first seem. Another problem is that while auditory equivalents may readily be found for some interface elements, there are a great many computer-based objects and operations for which no real-world audio counterpart exists.

One researcher in this area, Stephen Brewster, has taken up the idea of augmenting existing, visually based interfaces with sounds to assist the user, but has used abstract sounds rather than naturally occurring ones. These sounds – known as earcons (icons for the ear!) – are made up of short musical phrases that vary in the sequence of notes, overall pitch, tempo, and so on. Earcons have been added to icons, menu items, and so on, of conventional computer interfaces.

Speech Output

In communicating with other people, we most commonly make use of speech, and our interactions with computers can be similarly speech-based. Speech synthesis has been possible for quite some time, but it is only relatively recently that using synthetic speech has become a reality in everyday interfaces. However, speech is becoming increasingly popular ad an addition to more conventional user interfaces.

Note

Example: Voice as an alert

Recent versions of Apple's MacOS operating system include a speech synthesiser that will speak the contents of pop-up alert windows. If a pop-up window appears with a text message, this will be accompanied by a voice speaking the text.

While this will undoubtedly be a valuable feature for some users, others simply find it unnecessary, and turn it off (as we already discussed, it can take rather longer to listen to the verbalisation than it does to read the corresponding text). In fact, in Apple's implementation of this feature, there is a short delay between the pop-up box appearing and the start of the speech, and if the pop-up is closed or cancelled during this delay, then the speech is never started. Some users are quite surprised when, the first time they are slow to close the dialogue box, the computer starts speaking to them!

Note

Example: Voice can reduce demands on the visual sense

On modern aircraft flight decks, synthesised speech is used for many applications. For example, in the final stages of an approach to landing, the aircraft's height above the ground is critical. On some aircraft, the computer system monitors the height and provides a verbal read-back of the height in intervals of 10 feet as the aircraft approaches the ground.

This is one instance where speech has the potential to add real value to an interface. On approach to landing, a pilot's visual sense is usually in high demand – looking out of the window at the runway below, as well as monitoring other instruments. If the critical height information were available only on a visual display, then the pilot's vision would be stretched even further, by the requirement to look at yet another display, and by the need to continually re-focus on different things. On the other hand, if some of the information can be "off loaded" and presented acoustically instead of visually, then the demand on the pilot is lessened, reducing the potential for error and catastrophe.

Sound Input

Sound can also be used as an input device. Almost always in human-computer interaction, this means using voice recognition – a relatively new technology that is beginning to achieve the level of performance needed to be effective. One often hears a view, presumably inspired by science fiction movies, that if only we could just speak to our computers, then all our usability problems would be solved. Indeed, we would no longer need user interfaces! It must be emphasised, though, that voice recognition is yet another kind of user interface technology, and just like the mouse and keyboard. It therefore is appropriate for some things and not for others, and using voice input has just as much (or perhaps more) potential to create usability problems as any other technology.

A number of products exist that allow standard desktop computer systems to take speech as input. For instance, recent versions of Apple's MacOS operating systems allow the user to speak standard commands (such as Save, Print, Open, and so on), which the computer interprets and executes. Typically, though, users spend only a small part of their time entering commands into a computer, and much more time entering and editing the content of their documents, so a facility for speaking commands is only of limited use.

A number of products are available that allow users to speak text to the computer, which is then used as the input for word processors and other software.

While voice input might be a nice addition to existing desktop computer systems, it is likely to play a much more central role in interactions where using familiar visual output and mouse and keyboard input is difficult or infeasible. One such situation that we have mentioned before is the provision of services that will be accessed by telephone using only the auditory channel. Several products exist that allow designers interactive telephone-based services to incorporate voice recognition technology. This allows users to provide information, commands and so on. For example, the Vocalis group markets a range of products and has made demonstrations available over the telephone.

Implications for user interface design

In the following sections we will present some suggestions and guidelines for when and where sound might most successfully be employed in the user interface for both the output and input channels.

Sound output

Audio output may be the appropriate channel to use for:

- Giving immediate feedback that an action has taken place (buttons on ATMs, telephones, etc., that "beep" or "click" as they are pressed);

- Presenting different kinds of information to that made available using the visual channel (e.g., non-static information such as alerts);
- Augmenting visual interfaces by providing additional information and cues
- Supporting users for whom the visual interfaces are not an option (e.g., those with visual impairments)
- Supporting users whose visual senses are already heavily used for other parts of their task (e.g., aircraft pilots)
- In interfaces where visual information cannot be presented (e.g. mobile or handheld devices with small or no screens).

Audio output may not be particularly good for:

- Constantly changing status information
- Use in shared offices or workplaces, where privacy is important or where the output of many users' computers would lead to confusion or disturbance
- Noisy environments where sound may be difficult to hear (e.g., the user interface of machinery on a building site)
- Quiet environments where sounds could cause distraction (e.g., a library or recording studio)

Voice input

Sound can also be used as an input device. Almost always, this means using voice recognition – a relatively new technology that is beginning to achieve the level of performance needed to be effective.

Voice input may be appropriate for:

- Users who are unable to use more familiar input devices like the mouse or keyboard (e.g. those with severe motor impairments)
- Users who are busy doing other things (e.g., an aircraft pilot, whose visual channel is occupied monitoring the approaching runway during landing)
- Interactions that may be cumbersome using the other interaction devices that are available (e.g., menus used in mobile phone services).
- Small and pre-determined range of spoken commands.

Voice input may not be particularly good for:

- Noisy environments
- Use by many different users with different voices and accents
- Wide range of words or a specialised technical vocabulary. Note that some commercially available voice recognition systems have versions adapted to technical specialisms such as medicine or law.
- Things that are not easily verbalised (e.g., diagrams)

Other Senses, Other Devices

So far, we have looked at how sound and vision can form parts of the interface between user and computer. We now ask the question of whether some of our other senses and physical abilities can be used in order to facilitate human-machine communication. The use of other senses is usually proposed for one of two reasons. The first is to replace a sensory channel that is more often employed in user interfaces. This is typically so as to make a system accessible to users with sensory impairments. For instance, if touch can be used to do some or all of the work of vision, then a system will be more accessible to blind users. The second reason is to add to or augment the experience created by senses already used. For example, if the web site about cookery were able to allow users to smell the results of cooking, the effect might be a richer and more enjoyable experience for the user.

Touch

In everyday life, we use touch, and related senses (that give us information about temperature, texture, pressure, body position and orientation, and so on) as yet another way of gaining information and feedback about objects and events around us – especially those with which we are interacting most directly.

Currently, little use is made by interface designers of users' sense of touch. However, a number of experimental prototypes and products meeting specialised requirements serve as excellent examples of what is possible and what may become more significant for interface design in the future.

Several devices have been designed that allow the computer to produce output in a form that can be sensed by touch and related senses. Such devices seem to be most successful is in providing output that is tactile (i.e., that is amenable to the sense of touch – especially by having distinctive texture or shape) and providing force feedback (where the user experiences the computer output as a force on their body). In some cases, these devices are intimately connected to the input devices through which the user issues commands to the computer, and are therefore capable of providing instant and very direct feedback to the user.

Tactile Output

Output devices exist that convert the output of the computer into tactile form. For example, the text normally displayed on a computer screen can be rendered in a tactile language, such as Braille, making it available to users with vision impairments.

For example there is product called "Braille ‘n' Speak" that allows blind or partially sighted users to take notes using a specialised personal organiser. Notes can be "read back" via either a synthesised voice or a "refreshable Braille cell" that turns text into tactile form.

Tactile output devices of this type essentially use touch as an alternative output medium. Other devices rely on an observation about the way most human interaction with the physical world takes place. The separation between input and output that is so clear-cut in many conventional computer interfaces is not so sharp in other contexts. As we take an action in the real world (such as taking hold of a cup and lifting it up) we get immediate feedback (about the texture, temperature, and weight of the cup), through the sensation of pressure in our muscles, joints, and fingers. Thus the actions we take and the feedback we get as a result are very closely connected.

A simple but effective way of combining input and output for user interfaces has been to allow a pointing device such as a mouse or joystick to provide physical feedback my making it resist the user's actions in a context sensitive way, and provide tactile feedback as the mouse pointer is moved over screen objects.

For example, the Moose is a mouse-like device that gives various kinds of physical feedback. If the user is dragging an object, the mouse will feel heavier than normal, and the user will be given tactile cues as the pointer moves over different kinds of screen object (for instance, moving across the edge of a window could produce a "click" sensation).

Such force-feedback mice are a relatively simple way of combining user input and feedback, but more sophisticated products exist. For example, the Phantom allows the user to move a pen-like stylus through space in order to interact with a computer using gestures. Feedback is provided by mechanisms attached to the stylus that make it easier or harder to move, or simulate the effects of different objects and textures that it comes into contact with.

A similar input/output technique that has been developed for use with virtual reality systems involves the use of a "data glove", which senses the position, orientation, and movement of the user's hand, allowing gestures to be used as inputs, and objects in a virtual reality system to be grasped. Several products exist that add force, pressure or vibration feedback to a data glove, allowing the user to "feel" virtual objects they touch or grasp.

Other Senses

So far, we have discussed how the senses of vision, hearing and touch might be employed by user interface designers to provide a richer and more compelling experience for the user. But that still leaves several senses that have not been utilised in human-computer interaction. The senses of smell and taste are unlikely to be very effective as ways to convey large amounts of structured information, but both these senses are highly evocative and profoundly shape our experience of real-life situations.

Surely, though, smell is not a feasible medium for computers to produce their output? At least one company is developing a product that allows computers to generate olfactory output: you can read about the "iSmell" device in an article in Wired magazine [http://www.wired.com/wired/archive/7.11/digiscent_pr.html].

Activity 3 – Smell, taste and touch in the interface

It is clear how designers can make use of user's auditory and visual senses in their designs. Assuming the necessary hardware was available, how could the senses of touch, smell and taste be used as well as vision and audition for e-commerce applications over the Internet?

A Discussion on this activity can be found at the end of the chapter.

Multiple modalities

We have now discussed a number of modalities – or human sensory channels – that interface designers routinely require users to make use of (and a couple of modalities that aren't yet used in human-computer interaction, but may one day be). Each of these has properties that make its use more or less suitable for particular kinds of function in particular kinds of situation. A design challenge that we haven't mentioned yet, however, arises when more than one sense or means of providing input to the computer is employed in an interface in a way that requires the user to use them in co-ordination.

Several systems have been developed that allow the user to mix commands expressed in various ways. One example is MATIS, a system for finding information about air travel. The user may specify their constraints (such as origin and destination, date of travel) using a mixture of typed information, pointing with the mouse, and speech. The design challenge for such a system is to be able to correctly interpret a command expressed in a simultaneous mixture of these forms. So the user may specify travel by speaking "from here to there", while pointing at origin and destination cities with the mouse, and the system must match the occurrence of words like "here" with the corresponding mouse actions.

Not all users are the same!

As a final note in our discussion of perception, it should be emphasised that not all of our users have the same perceptual capabilities.

Graphical user interfaces are – not surprisingly, given the name – highly visual in the demands they make on users. Much of the user interface technology that is being designed for personal computers and used on the world wide web makes strong assumptions about the perceptual capabilities of users, and designers are therefore often excluding sections of the population whose vision, hearing, or other senses and abilities are impaired. Aside from any moral considerations about making technology accessible to all, there are good business reasons for considering access to our designs: the more people who are unable to use out designs, the more potential customers we have lost. This is especially true on the web, where a large number of potential customers will look at a site, and decide whether to continue to use it on the basis of whether its design helps them to carry out their tasks easily.

Sometimes, designing for inclusiveness and accessibility isn't only an ethical consideration or a matter of commercial good sense. In some cases it is also a legal requirement. For example, in the United States, the Americans with Disabilities Act (ADA) requires that certain organisations and bodies (particularly governmental bodies) provide "effective communication" when they use the Internet to make information and services available. Specifically, this ruling notes that web sites must be accessible to those using screen readers and receiving web information through auditory rather than (or in addition to) their visual sense.

Review Question 1

What are the properties of the audio channel that make it different from vision? What are the implications for the use of sound in the user interface of ordinary desktop computers?

Answer to this question can be found at the end of the chapter.

Review Question 2

Most telephones in the UK provide a facility for users to find the number of the last person who called them. By dialling 1471 a computer synthesised voice informs the user of the phone number and the time at which the call took place, for example, "Telephone number 07934 363 001 called today at 0930 hours.." Mobile phones, on the other hand, can show the numbers of the last several callers on the display. Compare the two presentations of roughly the same information and identify some of the advantages and disadvantages of using the auditory and visual senses for presenting this kind of information.

Answer to this question can be found at the end of the chapter.

Review Question 3

Voice synthesisers provide an obvious way for blind and partially sighted users to browse the web. What might some of the problems be with this way of making the web more accessible?

Answer to this question can be found at the end of the chapter.

Attention

The previous section and the previous unit have discussed some of the ways that humans are able to experience the world they inhabit using their visual, auditory, and other senses, and the importance of understanding the properties of these senses when designing user interfaces. It is clear that the world presents an extremely rich and complex set of stimuli to the senses. It is only by filtering out some of this, and ignoring the rest that we are able to make sense of the world around us, instead of being hopelessly confused by it. The processes by which humans are able to focus on some things and filter out others, to concentrate on certain mental or perceptual events and devote less mental resource to the rest, are known collectively as attention.

Attention is..

Every one knows what attention is. It is the taking possession by the mind in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are of its essence. It implies a withdrawal from some things in order to deal with others.

This oft-quoted passage from one of the founders of modern psychology captures some of the essence of what we mean by attention. At the heart of the concept is the idea that there are limits to our cognitive and perceptual capabilities, and that multiple demands are almost always made on these limited resources. Some form of "filter" or selection process is required so that we can focus on only a subset of the possible mental and perceptual activities; this process of focusing is known as attention.

We are able to choose a single stimulus from the many of offer on which to focus on using what is often referred to as selective attention. However, we can also engage in a number of activities and focus on a number of things simultaneously, in which case out attention is said to be divided. The selection of what to attend to may be the result of conscious deliberation, or deciding to concentrate on something specific, or may arise from unconscious or involuntary responses to things that "grab our attention".

Much of the psychological research on attention has focused on the ways that people are able to direct their attention to particular perceptual events in their auditory experience. However, many of the points that are made and theories that have been developed can be generalised to other senses and other mental processes.

The consequences of inattention

We can fail to pay attention to the appropriate thing for a number of reasons, and the consequence is likely to be that we fail to complete a task as we intended, or substitute an incorrect task instead.

Does the following story, as told by psychologist James Reason, sound familiar?

Note

Imagine you have a visitor who has requested a cup of tea, while you only drink coffee. You go to the kitchen intending to prepare both tea and coffee, but return with two cups of coffee. The reason for this slip is clear; you failed to make an attentional check on your plan at the point where the initial common pathway, boiling a kettle, branches into its separate tea- and coffee-making components. As a result, you proceed along the habitual coffee route.

In general, a number of possible consequences can result from inattention. The kind of unintentional error or "slip" exemplified above, where through a failure to pay attention, a person carries out a more familiar task instead of a correct, but less familiar one, is known as a capture error. Related to this is a memory lapse, where something is simply forgotten because we aren't paying attention. Common examples are using a photocopier, but failing to remove our original document at the end, or using a vending machine, but failing to take the change.

A third class of attention-related problem occurs when we are simply distracted from one task by some other stimulus or task. An example when we stop working at the computer to answer the telephone. Sometimes the distraction is unwanted and unnecessary (such as the distraction caused by blinking text on a web page) and sometimes, even if the distraction is desirable (such as an important telephone call) it creates problems when we attempt to resume the original task.

A further kind of problem arises if our attention is not focused on something important. For example, we can miss the arrival of new e-mail unless we happen to be attending to the icon that announces its arrival. Of course, the effect of failing to attend to the right thing is not always so trivial, as the following example illustrates.

Note

A serious air accident resulted when the pilots of an aircraft incorrectly located the source of an engine problem their aircraft was having as being in the right engine (it was actually in the left engine). The pilots shut down the healthy engine and the aircraft subsequently crashed as it was attempting to land.

An instrument on board the flight deck could possibly have revealed the source of the problem, had the pilots noticed it. It is, of course, all to easy to blame the pilots for making an incorrect decision, but it is much more worthwhile to look at how attention getting characteristics designed into the user interface had a part to play.

The relevant instrument (lower right, underlined in white, in the picture below) was part of a panel containing several other dials, itself part of a mass of controls, dials, instruments, and computer screens on the flight deck. The dial in question has no alert or other mechanism to attract the pilots' attention to the fact that it was displaying an out-of-normal-range value. In fact the normal and abnormal ranges are not even shown on the dial at all (e.g., with a "red zone", as on many of the instruments) Therefore, there was nothing to encourage the pilots to focus on this item in preference to all the other possibilities.

Note that the contribution of the dial design was only one component in a complicated series of occurrences that lead eventually to the accident, which is described from a different perspective in Unit 2.

Factors that affect attentional focus

As has already been said, at the core of the concept of attention is the idea that we have only a limited capacity for perceiving the world and carrying out mental and physical tasks. We therefore need choose from the large number of competing possibilities, a small number of things to concentrate on. A number of factors connected with the stimuli we receive, the way our mental apparatus processes them and our state of mind at the time affect the way we are able to achieve this focusing.

These include the following properties of the stimulus

- Movement and change – the human eye is able to discern great detail, but only in the central part of the visual field. In other parts of the visual field (the "peripheral vision"), we are not particularly good at seeing detail and colour, but are very good at spotting movement and change towards which we can direct our attention. This is why it is possible to notice that something is happening "out of the corner of ones eye", without seeing exactly what.

- Colour, size, intensity– certain colours tend to be particularly good at attracting our attention. Similarly if one object is physically larger, brighter, bolder or more visually intense than surrounding ones, then it is more likely to attract attention
- Number of competing stimuli – the greater the number of possible things to attend to, the less attentional resource there is to devote to each.

In addition to these properties of stimuli, two factors concerning the human observer have an impact on how and where attention is directed.

- Meaningfulness – we are more likely to focus on things that are meaningful or understandable or make sense to us. One aspect of this is a phenomenon sometimes referred to as the "cocktail party effect" Imagine you are at a party or in some other situation where a crowd of people are talking. There is a constant level of background sound, yet it is possible to pick out words and threads of other conversations across the room. We are especially good at doing this if the content of the conversation is meaningful to us – for instance, if we hear our name mentioned in someone else's conversation, we can "tune in" and listen to what is being said.
- Emotional state – in a particularly stressful or pressured situation, a person's attentional capacity tends to be diminished, meaning that they are able to focus on fewer things at once. This fact may not be very significant for the design of web sites and desktop systems. However it is highly relevant for the design of systems such as aircraft flight decks. If designers are not careful, there is a tendency to place the greatest attentional burden (e.g., lots of visual tasks to be carried out) on the user at precisely the most stressful times (e.g., landing the aircraft).

Activity 4 – Structuring information and focusing attention

The following exercise is taken from Preece et al. (1994), pages 102-103.

You can carry out this exercise by yourself, but it will be easier and more reliable if you are able to get someone else to help by carrying out the timing. The two figures below are alternative presentations of the same kind of data. With the first display find out how long it takes you to find (1) the phone number of Howard Johnsons in Columbia and (2) the name of a motel with a double room for $46. With the second display find out how long it takes to find (1) the telephone number of a Holiday House and (2) the name of a hotel with a double room for $27.

Which display takes the longest? Why do you think this is?

City	Motel/Hotel	Area Code	Phone	Rates Single	Rates Double
Charleston	Best Western	883	747-8961	$26	$38
Charleston	Days Inn	883	881-1888	$18	$24
Charleston	Holiday Inn N	883	744-1621	$36	$46
Charleston	Holiday Inn SW	883	556-7188	$33	$47
Charleston	Howard Johnsons	883	524-4140	$31	$36
Charleston	Ramada Inn	883	774-8281	$33	$48
Charleston	Sheraton Inn	883	744-2401	$34	$42
Columbia	Best Western	883	796-9400	$29	$34
Columbia	Carolina Inn	883	799-8200	$42	$48
Columbia	Days Inn	883	736-0828	$23	$27
Columbia	Holiday Inn NW	883	794-9448	$32	$39
Columbia	Howard Johnsons	883	772-7288	$25	$27
Columbia	Quality Inn	883	772-8278	$34	$41
Columbia	Ramada Inn	883	796-2700	$36	$44
Columbia	Vagabond Inn	883	796-6240	$27	$38

Pennsylvania
Bedford Motel/Hotel: Crinoline Courts
(814) 623-9511 S: $18 D: $28
Bedford Motel/Hotel: Holiday Inn
(814) 623-9006 S: $29 D: $36
Bedford Motel/Hotel: Midway
(814) 623-8107 S: $21 D: $26
Bedford Motel/Hotel: Penn Manor
(814) 623-8177 S: $18 D: $25
Bedford Motel/Hotel: Quality Inn
(814) 623-5188 S: $23 D: $28
Bedford Motel /Hotel: Terrace
(814) 623-5111 S: $22 D: $24
Bradley Motel/Hotel: De Soto
(814) 326-3567 S: $28 D: $24
Bradley Motel/Hotel: Holiday House
(814) 362-4511 S: $22 D: $25
Bradley Motel/Hotel: Holiday Inn
(814) 362-4581 S: $32 D: $40
Breezewood Motel/Hotel: Best Western Plaza
(814) 735-4352 S: $28 D: $27
Breezewood Motel/Hotel: Motel 78
(814) 735-4385 S: $16 D: $18

A Discussion on this activity can be found at the end of the chapter.

Implications for user interface design

In order for a user to be able to successfully use a computer system, it is imperative that they be able to direct their attention to the right thing at the right time. In the remainder of this section, we outline some of the measures that designers can take to help users to focus their attention on the parts of the user interface relevant to their task, and to remain focused and not be unnecessarily distracted.

Design features affecting attention

Let us take as an example the design of web pages. The kind of question a designer or usability analyst should be asking is: what does the user need to focus on in order to carry out their task? is the user's attention likely to be focused on these things? Or are there features of the design that tend to encourage them to focus elsewhere? If the user does get distracted and focus their attention elsewhere (sometimes this is unavoidable, and results from events outside the user interface – such as the telephone ringing or a colleague walking into the office), can they pick up the task again later and continue by re-focusing their attention appropriately?

As designers, there are a number of techniques we can use to try and ensure that users' attention will be directed towards the appropriate things for the task at hand. Although these are exemplified by pages and sites on the World Wide Web, the points made apply equally to other interactive systems.

- Flashing, motion, change – these techniques may be a powerful way to attract attention, but they can often simply be a distracting annoyance to users – see the section below on attention grabbing versus helping the user to decide.
- Difference – differences in colour, size, shape, intensity, and so on, can be used to make some items of a display more attention getting. For example, colour can be used to segment the page into regions that help users to focus, and to highlight important items.
- Meaningfulness and importance. Users will tend to find it easier to focus on items that are easy to understand and important to their task.
- Visual structure and grouping of items meaningful units can help the user to focus on the items that are relevant to their concerns. See for example BBC News website [http://news.bbc.co.uk/], where a large amount of information is presented in a structured way.

- Simplify the choices presented to the user. Look, for example, at UCT's Department of Computer Science [http://www.cs.uct.ac.za]
- Alerting mechanisms such as animations, and pop-up screens should be used with care, and only to provide warnings or supplementary information.
- Spatial cues, indicating, for instance, where in a fill-in form the user currently is, can help the user focus on the parts of the a task they are currently working on.
- Temporal cues, indicating where in a process or sequence of sub-tasks the user currently is. See, for example, the sequence of steps involved in purchasing an aeroplane ticket at Mango Airline [http://www.flymango.co.za]. This way of depicting the steps in a process is becoming increasingly common in e-commerce sites.

Activity 5 - Attention focusing in Web design

Look at the following collection of web sites, and for each, consider how the page designer has built in features that can help or hinder the user in their ability to focus their attention on what's important, and to remain focused on the task at hand.

- http://www.bbc.co.uk/
- http://www.amazon.com
- http://www.telkom.co.za

Attention grabbing versus helping the user to decide

It is all very well to design user interfaces with features that attract the user's attention in one way or another, but does this really contribute to the usability of the interface?

A typical situation, which might be familiar to users of the web, is that the designer includes features that are likely to attract the users attention, presumably because the designer felt that certain parts of the interface, or certain pieces of information were especially important. If you are a designer, this might seem quite easy to do: simply decide what you what the user to look at, and make it larger than surrounding elements of the interface. Or make it more brightly coloured, or make it move or flash. It you really feel that the item is very important (e.g., if its an advert for a product the user hadn't yet realised they wanted to buy), then use all of these techniques together, and include a large, attractively coloured animated graphic in the interface. Surely, the user will be virtually unable take their eyes off the graphic and will soon forget what it was that they originally came to your site for!

Actually, if this technique does work at all, and the user's attention is captured, they are likely to find it distracting and annoying. Users are likely to stop using the site and choose a different one that allows them to carry out their task without being bothered animations and so forth.

A better alternative might be to accept that the user might be best placed to decide what is important and what is not for purposes of the task that they are trying to accomplish. The design should therefore try to support users in identifying what possibilities the interface is offering them, assessing their importance and relevance, selecting an appropriate alternative, and remaining focused on their task without unnecessary distractions.

In any case, recent research suggests that crude attempts to attract web users' attention are less than successful. A study in which eye-tracking technology was used to find out where web site users direct their gaze showed that users tend to prefer looking at text, and avoid looking at images and animations. Even when the graphics contained useful information! A finding that suggests that users have become habituated to web designers tricks and have trained themselves to overcome basic physiological and psychological responses, so as to ignore web features they imagine to be superfluous. This tends to indicate that banner advertising and other attempts to encourage the user attend to things that may not be relevant to their interests and intentions may be far less effective than designers would like to believe. It should be emphasised that this research finding is specific to the users of particular kinds of web site, and cannot necessarily be generalised to other settings and other kinds on user interface.

Of course, if the user interface emphasises nothing in particular (or makes everything look equally eye-catching) then it is failing to help the user to decide what to do. There are sometimes situations where it is useful or even vital to distract the user form what they are currently doing and encourage them to think about something else instead. These include situations where safety is at stake (such as the aircraft example described earlier), as well as more everyday systems.

In deciding what to emphasise and make more available, web page designers can often find out (from server access logs) which parts of a site are the most popular. The site can then be re-designed so as to make it easy for the user to focus on the most popular information and links.

Activity 6 - Getting the user to look in the right place

Study the description given by Bruce Tognazzini at his web site [http://www.asktog.com/columns/000maxscrns.html] of how progress was made on a particular, and simple sounding design problem. At each stage, consider where users were focusing their attention, and how what the designers did had an effect on this.

Review Question 4

What is attention? What are some of the consequences of failing to focus attention on the appropriate things in a user interface?

Answer to this question can be found at the end of the chapter.

Review Question 5

What user interface design features affect users attention focusing? In what situations might the use of each be appropriate for in a user interface?

Answer to this question can be found at the end of the chapter.

Memory

A further aspect of human psychology that is highly pertinent to the design of user interfaces is the study of memory: that is how we are able to store and remember information that is relevant to the computer systems we use and the tasks we use them for.

Every activity we engage in involves the use of our memories in one way or another. Not only is memory used in processing every input we receive from our senses, but also memory is essential for the fleeting, temporary or intermediate thoughts that are part of just about every task we carry out. At the opposite end of the spectrum, we continually need to remember facts we have learned previously about the world we live in – and the computer systems we use – in order to make decisions and carry out actions. Even when we are not aware that we are remembering things, we probably are! Skills that we acquire are often performed without conscious thought, but still require us to recall remembered abilities. Memory is just as central for an expert cyclist or a proficient typist, who may not be conscious of remembering the skills they have acquired, as it is for someone trying to recall facts during an examination, who may be painfully aware of their attempts to remember (or failures to do so).

Failing to remember can result in a number of possible errors in carrying out tasks. The kind of "lapse" or error of forgetting mentioned in the attention section, which can lead to us omitting part of a task or carrying out the wrong task, is one example of this. An HCI-related example here would be failing to save a document before quitting a word processor, or failing to bookmark a web site before quitting the browser.

Another effect of failing to remember is when we simply cannot recall how to carry out a task at the time when we need to, or fail to recall some important fact about the workings of the computer system. An example is if we are unable to recall the meanings of visual elements of a user interface. See if you can you recall the meanings of each of the icons on this toolbar from a popular word processing program:

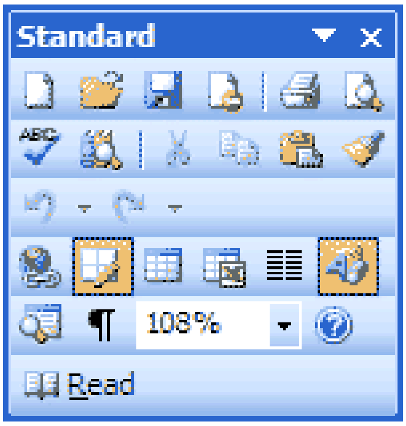

Some of the meanings are easy – the icon that looks like a printer invokes the print command – but others are more difficult – what is the meaning of the icon to the right of the printer?

A potentially more damaging instance of much the same phenomenon is if we mis-remember facts about the computer system or tasks we wish to perform with it. For example, if we incorrectly recall the meanings of the icons on the title bars of windows, we can end up quitting an application and losing work rather than simply minimising a window.

In the remainder of this section, we will look at how our memories work, and the three different kinds of memory store that studies have identified. We will then look at how what is known about memory can be used to improve user interface design.

The psychology of human memory

Psychological studies have revealed several types of memory system have a role to play in human cognition, each with its own specialised function and its own properties. Central properties that distinguish these types of memory are the capacity or size of the store, the decay time or the time for which items will persist before being forgotten, and the time it takes to recall or access items.

Three classes of memory that psychologists have studied extensively are sensory memory, short-term (or working) memory and long-term memory. The following diagram (extracted from the Model Human Processor of Card, Moran and Newell, 1983) summarises the characteristics of these memories and the relationships between them. Although the figure only shows sensory stores for the visual and acoustic senses, specialised memories exist for the other senses as well.

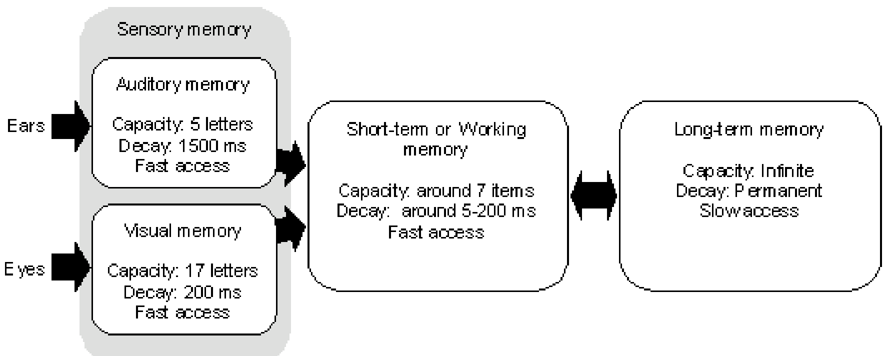

Sensory Memory

The first type of memory is responsible for recording information as it comes from our senses, prior to processing by other cognitive mechanisms. This so-called sensory memory acts as a kind of buffer store between the sensory organs and other forms of memory. In fact, the sensory memory is not a single entity, but a series of specialised stores, one for each of our senses. Thus images perceived by our eyes are retained briefly in the visual (or iconic) memory, and sounds perceived by our ears are stored in the acoustic (or echoic) memory. These memories store "images" of what is perceived for a short period of time, and can be thought of as being continually overwritten as new sense data arrives. In fact there are several different sensory memories: one for each of the senses, each having a different capacity and decay time, and storing images specific to a single sense.

Short-term or working memory

The second form of memory is referred to as either short term memory or working memory, and can be thought of as a kind of scratch pad in which information that is the current focus of our conscious attention is stored. The short-term memory typically stores sensory data from the sensory memories, and other short-term information that is needed as mental tasks are carried out.

Typically, our senses are bombarded by a vast amount of information about the external environment, and therefore a large volume of sensory stimuli enters the sensory memories. However, we choose to pay attention to only a small amount of this, and this attentional focus allows only a part of our sensory experience to be passed on to working memory.

In addition to processing sensory experiences, working memory also serves as a kind of "scratchpad" in which thoughts are stored while they are our focus of attention. For example, if you wish to multiply 39 by 8, it is necessary to remember these two numbers, multiply 9 by 8, to produce an intermediate result, remember the 30 and 8 and multiply them together, producing another intermediate result. The two intermediate results must then be remembered and added together, a process that involves further remembering of small amounts of information, which are stored for short periods of time, and then discarded.

Working memory has only a very limited capacity – studies have shown that it typically holds between 5 and 9 items (though what counts as an item is by no means fixed). Items persist in working memory for a relatively short time – from a matter of a few milliseconds up to tens of seconds). This means that working memory is a very limited resource, and if tasks make excessive demands on it, then errors are likely to result.

Activity 7 - Working memory

"Modes" are states of a system or user interface where the controls and inputs have the same effect. A user interface may have several modes, and the user's inputs will do something different in each, is which case the user must remember which mode they are in and what effect their actions will have. (For instance, Microsoft Word has Normal and Overwrite modes. In the former, characters that are typed are inserted in to the document, in the latter, existing text is overwritten by typed characters.)

Look at the user interface of your computer (if you want, select a particular application to look at), and find places where different modes exist, and where the user must remember the current mode. What can be done to reduce the errors that are likely to result from the presence of modes and the limits of working memory?

A Discussion on this activity can be found at the end of the chapter.

Long-term Memory

The third kind of memory is referred to as long-term memory, and is an effectively permanent and infinitely large store for semantically linked information.

A number of distinctions between different forms of long term memory have proved useful in the study of memory by psychologists. The subtleties of these different classes of long term memory, and the theories about how they function are not really relevant to this module, but it is worth noting that a variety of different kinds of information are stored in long-term memory. These include:

- facts about the world, computers and other devices we use, and tasks we carry out (sometimes referred to as semantic memory)
- things that have happened in the past (autobiographical memory)
- things to do in the future (prospective memory).

The processes by which we are able to remember items that are stored in long-term memory are complex and are the subject of much psychological research. However, a factor that greatly affects our ability to remember something is its familiarity, both in terms of its meaningfulness, and the frequency with which items have been remembered in the past.

This is why the use of graphical icons in a user interface can aid usability: the more meaningful a graphical element or a name appears to be, the more likely its meaning is to be remembered. For this reason, the meaning of menu items like "Print" and icons like

are more likely to be remembered than "Control-X" or "F4". In addition, in a given situation we are more likely to successfully recall a fact or piece of information that we have recalled often in the recent past, than something that is only occasionally required.

Note

Example: The Unix operating system.

In order to carry out any command, a Unix user must remember a great deal about the command, such as its name, how to supply parameters, the form of the output, and so on. Even to be able to use the online help system, the user needs to remember the command name. In some cases, the names bear a resemblance to the command they denote (the move command is "mv", the copy command is "cp", the print command is "lpr"). For some commands, the name is equally suggestive of the function, but the function is less frequently used and therefore harder to recall. In other cases, the connection between the name and its function is bizarre. Take, for instance the "biff" command, whose purpose is to notify the user when email arrives:

Did You Know?

"Biff was Heidi Stettner's dog, back when Heidi (and I and Bill Joy) were all grad students at U.C.Berkley and the early versions of BSD were being developed. Biff was popular among the residents of Evans Hall, and was known for barking at the mailman, hence the name of the command."

It is, of course, possible that if an amusing story is attached to a command, then a chain of associations between the command and the means of invoking it is produced, making it more memorable. But this is hardly a workable principle for the design of user interfaces!

Activity 8 - Long Term Memory and Interface Design

Investigate how much expert users know and can recall about a computer system they use routinely (obvious examples would be popular word processors, email programs, or web browsers). In particular, find out how much they know about the commands that are available, how the commands are named, and how they are accessed through the toolbars, menus and icons on the screen. A good way to do this would be to plan a structured interview or a questionnaire, to probe the user's knowledge of a selection of frequently used and less popular functions that the product supports.

A Discussion on this activity can be found at the end of the chapter.

Knowledge in the head and knowledge in the world

Memory research has shown that humans are far better at recognising material presented to them than they are at recalling material from memory.

One of the difficulties with text-based command-line user interfaces, as the Unix example above shows, is that the user interface affords very few opportunities for recognition, and relies almost totally on the user being able to recall the relevant commands from memory. Graphical user interfaces, on the other hand, make use of menus, icons, and similar devices to provide the user with on-screen representations of the commands that are available. The mental act of recalling how to carry out a particular function is therefore transformed into the rather less demanding task of recognising the visual representation of a command. In general, we are able to combine information we recognise in the environment with that recalled from memory in deciding how to act.

Experimental research in HCI suggests this is precisely the way people make use of menus and similar visual interface features. Even experienced users of a system have a very poor recollection of what menu items appear where in an interface, how they're labelled, and so on. What this means is that in order to be proficient at using a computer program, it is not necessary to remember precisely how the user interface works. It appears that instead, users decide what action to take largely on the basis of matching options and functions displayed on the screen (e.g., as menu items and headings) with their current goals and intentions. For example, if the user wishes to save the current document, they may well look for items that are relevant to this goal. Of the options that are available, the "File" menu tab will be recognised as the most relevant, leading the user to pull down that menu. Once the menu appears, a number of options appear, allowing the user a further opportunity for recognising the one that is most relevant: in this case "Print".

Donald Norman (see Norman, 1988/1990) has discussed the difference between recall and recognition in terms of the location of knowledge as being "in the head" or "in the world". What this means is that our actions are often governed by bringing together information that we recall from memory and information that is present in the external world. In everyday life, we frequently arrange our environment so as to place important information "in the world" where it is less susceptible to being forgotten or mis-remembered. Familiar examples of this are when we write notes to ourselves or lists of things to do at a later time, thus "off-loading" the responsibility for storing them.

As designers of interactive systems, it is possible to achieve similar effects, by placing important information about how a system works in the user interface, rather than relying on the users remembering it.

Implications for user interface design

From the discussion above, it should be clear that an understanding of how human memory works is very useful in designing human-computer interfaces. Some key points to bear in mind when designing user interfaces are:

- Reduce working memory load – try to minimise situations where the user must remember temporary information. For example, if the user needs to makes use of selections made earlier in a dialogue, then the choices made should be displayed so the user doesn't have to remember them.
- Recognition of interface elements presented to the user is far more effective that requiring the user to recall how to use a system
- Therefore, where possible put the knowledge about how a system should be used into the user interface.
- Graphical user interfaces typically require far less mental effort that text-based ones as they put more knowledge and information in the user interface. This allows the user to recognise what is in the interface, rather than having to recall facts from their memory.

- The use of meaningful and easy to understand icons, text labels, and menu items is a way to do this and therefore reduces the burden on the user's memory.
- User interfaces in which commands work in a consistent way reduce memory demands on the user. For instance, in Microsoft Word, the TAB key inserts a tab character – most of the time. When the cursor is inside a table, the TAB key moves the cursor on to the next box in the table! The designers have therefore forced the users to remember and recall an extra item from memory: the rule that governs how TAB behaves inside a table.

Review Question 6

Working memory is said to be limited in capacity to 7 +/- 2 items. What implications does this have for the design of menus in standard computer systems? Does it have implications for any kind of menu?

Answer to this question can be found at the end of the chapter.

Review Question 7

What is the difference between knowledge in the head and knowledge in the world? Why is it important for user interface design?

Answer to this question can be found at the end of the chapter.

Discussion Topics

People are good at adapting their environment in order to remember things. This is specially true for things they have to do in the future (prospective memory) and thing they have done in the past (autobiographical memory), but applies to other kinds of knowledge too. Examples are writing various kinds of notes (e.g., a list of things to do) and placing objects so that they are in view (e.g. placing letters to be posted near the door so that they're not forgotten). Some techniques are specific to certain jobs (such as aircraft pilots removing a glove as a reminder that a task is unfinished) or to certain cultures (such as tying a knot in a handkerchief as a reminder that something needs to be done).

How do you adapt your environment to help you remember things? How might the kinds of "external memories" you create provide inspiration for designing memory aids in user interfaces?

Answers and Discussions

Answer to Review Question 1

Sound is not persistent, carries relatively less information, conveys a single "linear" stream in information, tends to perceivable at a greater distance than visual output, is more or less non-directional, and it is difficult for us to choose to focus on only a part of our "audio field".

This means that sound is not necessarily appropriate conveying all types of information, and not good for conveying large volumes of information. But it is particularly good for providing certain kinds of alerts, giving short term feedback that actions have been carried out, augmenting interfaces to provide additional cues to the user (e.g., Earcons).

The non-directionality and range of sound may also have implications for privacy and use in situations where sounds would create a disturbance.

Answer to Review Question 2

Main advantage of voice is that it doesn't require a special phone.

The voice message may take longer to play than the visual presentation takes to read.

If the time at which the call was made is the important thing, the need to wait until the earlier part of the audio message has played might be frustrating.

Audio interfaces tend to place a greater burden on the user's memory.

Recording contains unnecessary words and could therefore be made more efficient.

Answer to Review Question 3

Time taken may be much greater than for sighted users.

Many web sites have substantial non-text content, such as graphics, animations and so forth, which a voice synthesiser would be unable to cope with. (The problem can be alleviated, but not eliminated if web page designers use ALT tags when images are included in web pages.)

Web pages are often designed in an inherently two-dimensional way, making use of layout and organisation in significant ways. A voice synthesiser, though, reduces all pages to a one-dimensional stream of words, losing much of the structure that is otherwise visually presented.

Answer to Review Question 4

Attention is the selection of one of a number of possible perceptual and mental activities to focus on. The idea is that we have limited cognitive capabilities, and therefore have to be selective in what we perceive or think about.

Inattention can result in failing to carry out tasks (forgetting them – a "lapse") or carrying out a similar, more familiar, task in place of the desired one (a "capture error"). Focusing attention of the wrong elements of a user interface can lead to distractions, making the user's task harder, and to the user missing important stimuli (such as notification of an emergency).

Answer to Review Question 5

Visual and audio alerts – to draw the user's attention to specific and important events and conditions (e.g., emergencies, or system generated warnings).

Motion, blinking, annoyance, etc – can be a distraction and an annoyance. Use only occasionally and with care!

Colour, intensity, size, etc. – to make specific pieces of information stand out from the background.

Spatial and temporal cues – to help the user to focus on where they are currently working in a large space (e.g., a form) or what stage they are at in a process with several stages (e.g., online purchasing).

Structure and layout – structuring can help users decide what to focus attention on when a large amount of information is presented, and to stay focused.

Answer to Review Question 6

Size limits of working memory have no implications for visual menu design - the user is not expected to retain all menu items in memory, that is what the menu labels are for!

However, working memory size and decay time is highly relevant for the design of audio menus, where the user may have to remember the menu items until after the menu has completely finished playing before making a selection.

Answer to Review Question7

Knowledge or information that is in the world (i.e., in the environment or the user interface) must be perceived and recognised by the user. Knowledge in the user's head must be recalled before it is of use.

People are generally better at recognising things than recalling material from memory. Therefore, if we can build more of the information about how to use a system into its user interface, we have reduced the demands on the user's ability to recall knowledge. Therefore the mental effort of using the system should be reduced.

Discussion on Activity 2

Wide / shallow schemes are unlikely to be as effective for audio menus. The benefit of having many options in a menu doesn't necessarily apply, since the user must wait for all of the options to be spoken.

However, one problem with deep audio menu structures occurs if user gets lost, makes an error, or has to backtrack.

One way to investigate this without the need to implement a voice menu system is to carry out a "Wizard of Oz" experiment. An experimenter plays the role of the computer, and speaks the menu items in response to the user's actions. See, for example, (Dix et al., 1998) or (Preece et al., 1994) for more details about this technique. The results, in terms of time taken, errors, and so on can be compared with results from using the menu structure of an application program.

Discussion on Activity 3

Applications are many and include selling anything where being able to smell, touch, or taste a sample of the product would make the user more confident that it is what they're looking for. For instance, users could smell food products, perfumes or cosmetics before they buy.

Touch could be similarly used to allow users to feel the texture of products before buying. Some sites marketing clothing attempt to give an idea of the texture of fabric used by providing pictures. This could be greatly enhanced if the user were able to feel "for real".

Discussion on Activity 4

You probably found that it takes longer to find information in the second display than the first. A study cited in Preece et al (1994) found that on average people took 3.2 seconds to find information in the first screen and 5.5 seconds in the second. Although the two screens contain roughly the same amount on information, the structuring of information in the first makes it easier for users to focus their attention on the particular type of information they are searching for.

Discussion on Activity 7

Many examples or modes exist. Already mentioned are is the TAB key which works differently in different places. Similarly, In Word, Control-D sometimes opens the Font dialogue box, and sometimes duplicates drawing objects; a title-bar icon with an X in it sometimes closes a document and sometimes an application plus documents; and so on.

Measures to reduce the problems include: reducing the number of modes and mode-dependent commands, making modes work in a more consistent way, making it very clear what the current mode is (so that the current mode can be perceived rather than having to be remembered). The best solution is to eliminate modes entirely and make commands work the same way in all circumstances, thus reducing demands on both memory and attention.

Discussion on Activity 8

Typically, expert users have a surprising poor recall of the functioning of the systems they use. What they do tend to have a better knowledge of is how to interpret information that is available on the screen, and well as information that is not available on the screen, See (Preece et al. 1994, Chapter 5) for a fuller discussion.